Collective Bargaining Networks, Rent-sharing, and the Propagation of Shocks

Santiago Hermo*

Job Market Paper
(Click here for latest version)

March 20, 2024

Abstract

I study the role of collective bargaining, a prevalent wage-setting institution in many countries, in determining the effects of economic shocks on workers and firms. To do so, I leverage novel administrative data which allows me to construct the network linking employers to collective bargaining agreements in Argentina. I exploit changes in world import demand between 2009 and 2013 to construct exogenous shocks to firms' product demand and study how these shocks propagate through the collective bargaining network. My findings indicate that a shock equivalent to a 10% rise in firm revenue leads to a 1.2% increase in wages, and a shock equivalent to a 10% rise in average revenue of firms under the same agreement leads to a 4.2% increase in wages. The wage floor set by the agreements increases as well, which means that economic shocks propagate to firms and workers who are not directly affected. I develop a general equilibrium model of the labor market with collective bargaining and firm heterogeneity to study how the collective bargaining network affects the propagation of shocks. The estimated model indicates that the degree of shock propagation is decreasing in the average dispersion of productivity within agreements. As a result, the extent of propagation is hump-shaped in the degree of centralization of bargaining, as more centralized networks tend to result in agreements that cover more heterogeneous firms.

Keywords: collective bargaining, networks, rent-sharing, wage floor, trade shocks.

*I am indebted to Jesse M. Shapiro, Lorenzo Lagos, Peter Hull, and Neil Thakral for continuous feedback and support throughout the development of this project. I am thankful to Victoria Castillo, David Trajtemberg, and the staff of the *Dirección General de Estudios y Estadísticas Laborales* at the Ministry of Labor of Argentina. They not only provided access to the administrative data but also offered invaluable guidance in using it. I am thankful to Constanza Abuin, Pedro Dal Bo, Joao Garcia, Andrew Garin, Brian Kovak, Diego Gentile Passaro, Ana Moreno-Maldonado, Matthew Pecenco, Pascuel Plotkin, Santiago Pérez, Sara Spaziani, Darío Tortarolo, Diego Verdugo, and seminar participants at Brown University, the 17th North American Meeting of the Urban Economics Association, and the PhD-EVS online seminar for helpful comments and exchanges. I gratefully acknowledge support from the Orlando Bravo Center for Economic Research at Brown University. All errors are my own.
E-mail: santiago_hermo@brown.edu.

1 Introduction

Collective bargaining (CB, for short) establishes common provisions for workers, such as wage floors, across all firms covered in a bargaining unit. CB can therefore influence how wages respond to changes in the economic conditions of employers. In Europe, for example, differential responses of national labor markets to the Great Recession and the rise of import competition from China have been linked to differences in bargaining institutions (e.g., Ronchi and di Mauro 2017; Barth et al. 2023). Understanding these responses is, in turn, important for the many public policies that seek to mitigate the effect of shocks on workers.[1] In the US, discussions about bargaining institutions have gained traction in recent years, and new bargaining policies have already been implemented.[2] These endeavors require an understanding of the trade-offs involved in different bargaining structures. While many papers examine how economic shocks, such as changes in product demand to firms, affect workers' wages (e.g., Van Reenen 1996; Autor et al. 2013) and propagate across firms and regions (e.g., Giroud et al. 2021; Adão et al. 2022), there is little direct evidence on the role of CB in mediating these effects.

CB typically defines a network of firms, with two firms connected in the network if they fall under the same CB unit and therefore share the same commonly bargained provisions. The structure of the CB network differs widely across countries, with some characterized by large sectoral CB units, and others by employer-level bargaining (Bhuller et al. 2022). When economic shocks affect firms in a particular region or sector, these different forms of CB network seem likely to respond differently.[3] For instance, if firms in an affected region are connected to firms in another region via CB, then updated negotiations in the CB unit may result in wage changes in the unaffected region as well. This could be a cause for concern if it results in low-wage regions covered by high wage floors that are not aligned with local conditions (Boeri et al. 2021; Adamopoulou et al. 2022). However, it could also be desirable if it dissipates the effects of concentrated shocks, effectively sharing the risk of shocks across workers and firms.

Studying the role of the CB network in determining the responsiveness of wages to economic shocks is challenging for several reasons. First, it requires detailed information on the CB network, which is not usually available in administrative datasets. Second, it requires measurable and exogenously determined shocks to employers. Finally, studying the effects of shocks under counterfactual CB networks requires an economic model where firms exposed to shocks are subject to common provisions negotiated by a CB network.

In this paper, I study how CB shapes the impact of economic shocks to firms. First, I leverage novel administrative data from Argentina to construct the network linking firms to CB units which, coupled with rich labor market data, allows me to study how shocks propagate through

[1] See, for example, the US Trade Adjustment Assistance program or the EU Globalisation Adjustment Fund.

[2] See, for example, the recent California law granting bargaining power to fast-food workers. Washington Post (2022) notes that the "California law is widely seen as a step toward sectoral bargaining."

[3] As noted by Katz and Autor (1999, pp. 1540), "the same labor market shocks [...] may have different impacts on the wage structure depending on how unions and government regulations affect wage setting."

the CB network. Second, I construct plausibly exogenous product-demand shocks by exploiting changes in international demand for exported products, and compare the evolution of outcomes across firms that are shocked directly or via their CB unit. Finally, I develop and estimate a structural model where firms in different regions and economic sectors are connected by a CB network, and use it to study the incidence of shocks on wages under counterfactual networks.

The case of Argentina in the aftermath of the Great Recession provides an excellent case study for several reasons. First, CB is widespread in the country, with 93% of workers covered by some agreement in 2014 (Ministerio de Trabajo, Empleo y Seguridad Social 2023). As in several other countries, agreements establish minimum working conditions for covered workers and firms.[4] Second, an unexpected shift in international demand for Argentina's exports in the aftermath of the Great Recession makes for a compelling natural experiment to study the role of CB in the propagation of shocks. Finally, due to the country's legal structure, CB units are not entirely determined by region or economic sector, resulting in an idiosyncratic network that allows me to compare otherwise similar firms that are nevertheless part of different CB units.

To construct product-demand shocks I leverage variation in world import demand arising from granular country-product pairs exported by Argentine firms in 2009–2013. Following Hummels et al. (2014), firms are exposed to changes in world import demand via the share of their value exported to each country-product, as reflected in Argentine Customs data. CB units are also exposed to product-demand changes via their employment shares in firms that export to each country-product. I define shocks to firms and CB units as the weighted average of the changes in world import demand, weighting by the appropriate exposure shares to each country-product.

I use administrative data on firms' wages and employment and a difference-in-differences strategy to estimate the effects of firm and CB shocks. To identify the effect of *firm shocks* I compare the evolution of outcomes for firms in the same province and economic sector that are subject to a similar CB shock but a different firm shock. To identify the effect of *CB shocks* I proceed analogously, but comparing firms with a similar firm shock and a different CB shock. The key assumption is that, within province and sector cells, growth in world import demand (WID) affected firms because of changes in their product demand and in the outcomes of bargaining, and not because growing firms sorted into markets with trending WID or were affected by changes in WID through channels other than their CB unit. Following Garin and Silvério (2023), I see this as a plausible assumption given the nature of the changes in WID, which were driven by unexpected recessions and recoveries of varying magnitudes across countries following the Great Recession. In an appendix, I show that these assumptions can be cast in terms of conditional quasi-random assignment of product-demand shocks, as in recent work by Borusyak et al. (2022b).

I start discussing the effects of shocks on wages and wage floors. A 10% increase in world import demand at the firm level increases wages by 0.13% (SE=0.07%), while a 10% increase in average world import demand at the CB unit level increases wages by 0.46% (SE=0.20%). Combining these results with evidence from a survey of businesses, I estimate that a shock equivalent

[4] See Bhuller et al. (2022) for a recent review of collective bargaining institutions in OECD countries.

to a 10% increase in average CB revenue would increase wages by 4.2% among covered firms. These effects are stronger for larger CB units. I use the administrative data to construct a novel dataset of wage floors, the main outcome of negotiations, by detecting bunching in the distribution of wages within CB units and occupation categories. I then estimate the models using the wage floor as outcome and find that wage floors increase by a similar magnitude as paid wages following a CB shock. This implies a pass-through of wage floors to wages close to one, larger than found by Card and Cardoso (2022), suggesting strong "wage norms" in the Argentine economy. I consistently find that wage floors respond to CB shocks but not to firm or sectoral shocks.

I then delve into the effects of shocks on employment. Importantly, theory suggests that a wage floor hike should result in heterogeneous employment responses across firms with different productivity (Ahlfeldt et al. 2022; Berger et al. 2022). Low productivity firms are demand constrained by the wage floor, and so they should respond by decreasing employment as they move along their labor demand curve. Medium productivity firms are supply constrained implying that they should respond by increasing employment. High productivity firms should not be affected. As a proxy for productivity, I study heterogeneity of effects of the CB shock by the level of pre-period wages. I find that low wage firms in CB units that experience growth in world import demand decrease employment, consistent with this theory. I also find evidence of positive employment effects for medium wage firms, and null effects for high wage firms.

A potential threat to the view that CB matters is that shocks actually propagate through other networks that are correlated with CB units, such as geographic proximity or input-output linkages. Several additional exercises, however, suggest that the causality runs through CB units. First, I find robust estimates when varying controls for connections via local labor markets, suggesting that the effects are not driven by spillovers through proximity. Second, I find that CB shocks also affect non-exporting firms in the same CB unit, suggesting that exporter-specific unobservables are not responsible for the results. Finally, I exploit firms with workers under multiple CB units to construct a falsification test. Unlike the main analysis, where I assign the most common CB unit to each firm, in this test I compare workers bound by different CB units *within the same firm*. If a shock to a firm's main CB unit reflects spillovers through other networks, then we should not observe CB unit-specific effects within the firm. I find that workers' wages respond to shocks to their own CB unit, suggesting that the effects are driven by CB.

In the final part of the paper, I develop a general equilibrium model to investigate how the structure of the CB network affects the incidence of shocks on wages. The model assumes the existence of heterogeneous local labor markets, defined by region and economic sector, which are affected by wage floors. CB units, defined as partitions of local labor markets, bargain over a common wage floor anticipating the response of covered local markets. I use a "Nash-in-Nash" solution concept, first introduced by Horn and Wolinsky (1988), to model the equilibrium of the bargaining game between unions and employer associations. The key parameters that govern heterogeneity in the model are estimated by inversion of equilibrium conditions. In particular, I use the share of firms bunching at the wage floor to obtain an unobserved productivity parameter in each local labor market. Similarly, I invert first order conditions of the Nash bargaining problem to obtain bargaining power parameters of unions. I simulate the export shocks in the model, and verify that it is able to replicate the qualitative empirical patterns in the data.

The model allows me to assess the role of the CB network in the propagation of shocks. I define counterfactual networks that vary in the mean size of CB units, a proxy for the degree of centralization of bargaining, and simulate the effects of the export shocks. Countries with decentralized bargaining include the US and the UK, while Sweden or Norway are examples of more centralized bargaining systems (Bhuller et al. 2022). I find that the degree of shock propagation is hump-shaped in the mean size of CB units, meaning that shock propagation is larger in CB networks of intermediate size. This result can be explained by the extent of connections created by the CB network and the dispersion of productivity within CB units. Decentralized bargaining results in a high "bite" of wage floors, but trivially does not propagate shocks across local labor markets. Increasing the size of CB units leads to more connections, so shocks propagate more. It also increases the dispersion of productivity within CB units, which results in lower wage floors as unions want to prevent employment costs in low productivity local markets. Eventually, the second effect dominates, and as the wage floor "bite" declines changes in the wage floor as response to shocks become less impactful on wages. This result suggests that the degree of risk sharing of shocks between firms is higher in networks where CB units are of moderate size.

This article contributes to several strands of literature. First, to my knowledge this is the first paper to show direct evidence that the wages paid by a given firm are affected by shocks to other firms in the same CB unit. Rose (1987) and Abowd and Lemieux (1993) find wage responses to shocks to CB units in different contexts.[5] Gürtzgen (2009b) and Rusinek and Rycx (2013) find that wages are more responsive to firm shocks when bargaining is decentralized. Card and Cardoso (2022) study the association between changes in average value added in the CB unit and changes in wage floors in Portugal. None of these papers distinguish between the effects of shocks to a firm and to other firms in the same CB unit, nor do they explore the response of firms according to their productivity. Garin and Silvério (2023) study the effects of "common" and "idiosyncratic" export shocks in Portugal, but do not explore the role of CB.

My result of stronger wage responses to CB shocks than to firm shocks suggest that the different magnitudes of rent-sharing elasticities in the literature may be explained by the level at which shocks are aggregated. For example, findings in Abowd and Lemieux (1993) using contract wages from CB agreements suggest an elasticity of 0.2, twice as high as the average elasticities across worker-level specifications reported in reviews by Card et al. (2018) and Jäger et al. (2020).

Second, the paper shows that the effects of trade shocks are mediated by labor market institutions, pointing towards a novel channel for the regional contagion of shocks. This finding contributes to our understanding of the spatial effects of economic shocks, such as those arising from international trade (e.g., Topalova 2010; Autor et al. 2013; Dix-Carneiro and Kovak 2017).

[5] Rose (1987) studies the response of wages in a deregulation episode in the US trucking industry. Abowd and Lemieux (1993) study the response of contract wages to changes in international prices in Canadian manufacturing.

Adão et al. (2022) study the role of spatial linkages in the propagation of shocks, and Felix (2022) discusses the role of local labor market concentration. Giroud and Mueller (2019) and García-Lembergman (2022) find that shocks propagate spatially through firms' internal networks. Borusyak et al. (2022a) note that the effects of regional shocks on migration depend on a region's shock and on the shocks of other regions as well. None of these papers discuss the role of CB, nor do they study how different CB networks affect the propagation of shocks.

Additionally, this is the first paper to incorporate bargaining and endogenous wage floors in a spatial economic model, allowing for varying degrees of centralization in bargaining. While prior work exists modelling bargaining institutions in different settings, it usually does not feature firm heterogeneity or spatial linkages.[6] In principle, the model can be used to study other questions in a spatial setting for which bargaining may be important, such as the connection between collective bargaining and regional inequality.

Finally, the paper contributes to a broad literature studying unions and collective bargaining (e.g., Freeman and Medoff 1984; Card 1990, 1996; Moene et al. 1993). In particular, several papers have studied the properties of different bargaining regimes both theoretically and empirically (e.g., Calmfors and Driffill 1988; Holden 1988; Cardoso and Portugal 2005; Plasman et al. 2007; Boeri et al. 2021). Bhuller et al. (2022) calls for "closing the gap between how economists tend to model wage setting and how wages are actually set." This paper takes a step in that direction by modeling the institutional structure of collective bargaining explicitly, and using an estimated model to evaluate the effects of different bargaining structures in a principled way.

The paper is organized as follows. Section 2 presents a stylized theoretical framework. Section 3 describes the context and data, Section 4 discusses the empirical strategy used to estimate the effects of shocks, and Section 5 dives into the results. I present a structural model of the labor market with collective bargaining in Section 6, discuss its estimation and validation in Section 7, and use it to conduct counterfactual exercises in Section 8. Section 9 concludes.

2 Theoretical Framework

This section presents a motivating theoretical framework with heterogeneous firms covered by a single CB unit. The model shows that shocks propagate to wages and employment through the wage floor, and that employment responses are heterogeneous across firms, yielding testable implications for the empirical analysis discussed in Section 4. While this section presents a stylized discussion of partial equilibrium effects under a single CB unit, the model of Section 6 will allow for multiple CB units and general equilibrium effects.

[6] Several articles present theoretical models to study the role of bargaining. For example, Calmfors and Driffill (1988) study the relationship between centralized bargaining and the level of unemployment, and Corneo (1995) and Naylor (1998) study how the bargaining structure affects the response of national labor markets to trade integration.

2.1 Firm heterogeneity and wage floors

Firms j face an upward-sloping labor supply curve $\ell_j(w)$ at wage $w = w_j$. In Section 6, ℓ_j will also depend on a measure of aggregate wages, capturing general equilibrium effects. Firms are heterogeneous in productivity φ_j, and the production function is $y_j = \varphi_j f(\ell_j(w_j))$ for a concave $f(\cdot)$. For simplicity, the output price is normalized to 1. A wage floor $\underline{w}$ sets a minimum to the wage that firms can pay. The firm's static decision problem is thus given by

$$\max_{w_j} \varphi_j f(\ell_j(w_j)) - w_j \ell_j(w_j) \quad \text{s.t.} \quad w_j \geq \underline{w}.$$

The solution (w_j^*, ℓ_j^*) is characterized by two thresholds that determine firms' behavior given their productivity (Ahlfeldt et al. 2022), which I denote by $\underline{\varphi}$ and $\overline{\varphi}$. Firms with $\varphi_j > \overline{\varphi}$ are *unconstrained* by the wage floor. They optimally choose wages above the floor, so we can write $w_j^* = w_j(\varphi_j)$ and $\ell_j^* = \ell_j(w_j^*)$. The unconstrained wage exhibits a wage markdown over the marginal product of labor, as in monopsonistic labor markets (Manning 2011; Card et al. 2018). The model implies that these firms will not be affected by marginal changes in the wage floor.

Firms with $\varphi_j \leq \overline{\varphi}$ are *constrained* by the wage floor, so that $w_j^* = \underline{w}$. These firms come in two varieties, depending on their productivity (Ahlfeldt et al. 2022; Berger et al. 2022). When $\varphi_j \leq \underline{\varphi}$ firms are *demand constrained* and their optimal choice of labor is such that $d\ell_j(\underline{w})/d\underline{w} < 0$. These firms may also close down if not profitable, resulting in employment losses as well. When $\varphi_j \in (\underline{\varphi}, \overline{\varphi}]$ firms are *supply constrained* and their optimal choice of labor is such that $d\ell_j(\underline{w})/d\underline{w} > 0$.[7]

2.2 Nash bargaining problem

I incorporate a CB unit that determines wage floors and study the partial equilibrium effects of shocks to firms. Aggregate revenue and the wage bill are given by $R(\underline{w}) = \sum_{j \in \mathcal{J}} \varphi_j f(\ell_j(w_j))$ and $WB(\underline{w}) = \sum_{j \in \mathcal{J}} w_j \ell_j(w_j)$, respectively, where $\mathcal{J}$ is the set of firms in the CB unit. The union objective is to maximize the wage bill $U(\underline{w}) = WB(\underline{w})$, and the employer objective to maximize aggregate profits $\Pi(\underline{w}) = R(\underline{w}) - WB(\underline{w})$. In this formulation, the union is concerned with maximizing total wages instead of worker welfare. In Argentina, unions receive a fee proportional to each payslip, so this is equivalent to assuming that unions try to maximize their income.[8]

Unions and employers within a CB unit Nash bargain over a wage floor $\underline{w}$. I assume that, in case of a floor $\underline{w}'$ that is not binding, aggregate profits are positive ($R(\underline{w}') > WB(\underline{w}')$). Additionally, the presence of firm heterogeneity implies that the wage bill is hump-shaped in the wage floor (Ahlfeldt et al. 2022), so there exists a $\underline{w}$ that maximizes the wage bill. Letting $\beta \in [0, 1]$ be

[7] The response of constrained firms depends on the level at which their marginal revenue product of labor (MRPL) and marginal cost of labor (MCL) intersect in the equilibrium without a wage floor. For demand constrained firms the MRPL equates the MCL at a level lower than $\underline{w}$. Thus, as the MCL goes up to $\underline{w}$, these firms increase the MRPL to equate it to $\underline{w}$ by reducing employment. For supply constrained firms the MRPL equates the MCL at a level higher than $\underline{w}$. A wage floor induces these firms to lower the MRPL by increasing employment.

[8] Gürtzgen (2009a) theoretically explores the effect of different bargaining structures in an oligopoly model and similarly assumes that unions maximize the wage bill.

the bargaining power of the union, the Nash bargaining problem is

$$\max_{\underline{w}} U\left(\underline{w}\right)^{\beta} \Pi\left(\underline{w}\right)^{1-\beta}. \tag{1}$$

If $\beta = 1$ the problem amounts to selecting the wage bill-maximizing $\underline{w}$, whereas if $\beta = 0$ the problem amounts to selecting the profit-maximizing $\underline{w}$ (which will be non-binding). Now consider the case $\beta \in (0, 1)$. In an internal solution the optimal wage floor $\underline{w}^*$ is implicitly given by

$$WB(\underline{w}^*) = \omega(\underline{w}^*)R(\underline{w}^*) \tag{2}$$

where $\omega \in (0, 1)$ is a weight given by

$$\omega(\underline{w}^*) = \frac{\beta}{\beta + (1-\beta)\left(-\frac{d\Pi}{d\underline{w}}(\underline{w}^*)/\frac{dU}{d\underline{w}}(\underline{w}^*)\right)}. \tag{3}$$

The solution sets the wage floor so that the wage bill equals a fraction of aggregate revenue. If we divide (2) by total employment, and allow for an outside option for workers in the objective function U, then we would get the familiar result that average wages are a weighted average of revenue per worker and the outside option. The solution is thus analogous to the classical Nash bargaining problem in which unions set wages instead of wage floors. The main difference is the weight ω, which may differ from β. For instance, if the wage floor is more "painful" for firms relative to how "beneficial" it is for unions ($-d\Pi/d\underline{w} > dU/d\underline{w}$), then ω will be smaller.

2.3 Response to shocks

I now consider the response of wages and employment to shocks to firms, and how this response is mediated by the wage floor. I refer to productivity shocks, but the results apply to any shock that affects the revenue of firms, such as product demand shocks. A key feature that determines the response of wages and employment is firms' baseline productivity, which determines whether they are constrained or not. The following proposition shows how shocks affect the wage floor.

Proposition 1 (Response of wage floor to shocks)**.** *Assume a change in the productivity profile* $\{d\ln\varphi_j\}$*. Holding* ω *fixed, the resulting change in the wage floor is*

$$d\ln\underline{w} = \frac{WB}{\tilde{WB}^{\text{co}}}\left[\sum_{j\in\mathcal{J}} s_j^R d\ln\varphi_j - (1-\omega)\sum_{j\in\mathcal{J}} \iota_j s_j^{WB} d\ln\varphi_j\right] \tag{4}$$

where $s_j^R = R_j/R_c$ *is the share of revenue of firm* j *in the CB unit and* $s_j^{WB} = WB_j/WB_c$ *is the analogous wage bill share,* $\tilde{WB}_c^{\text{co}}$ *is the wage bill of constrained firms adjusted for its response to the wage floor, and* ι_j *is the elasticity of the wage bill to productivity* φ_j*. Allowing* ω *to change results in a third term that depends on the responses of the ratio* $(-d\Pi/d\underline{w})/(dU/d\underline{w})$ *to the changes in productivity and the wage floor.*

Proofs are available in Appendix A. Proposition 1 shows that $\underline{w}$ responds to a weighted average of changes in log productivity minus a weighted average of wage bill elasticities times changes in log productivity. This latter term reflects the fact that firms respond to the shock changing wages in proportion to revenue, yet revenue is down-weighted in the Nash split, so the wage bill increases "in excess." We also note that the response of $\underline{w}$ is decreasing in the adjusted constrained wage bill $\tilde{WB}_c^{\text{co}}$. As $\tilde{WB}_c^{\text{co}}$ is decreasing in β, stronger unions will result in more responsive wage floors.

The following proposition shows how wages and employment respond to shocks.

Proposition 2 (Response of wages and employment to shocks)**.** *Assume a change in the productivity profile $\{d\ln\varphi_j\}$ that leads to the wage floor change given by (4).*

1. For unconstrained firms with $\varphi > \overline{\varphi}$ we have $d\ln w_j = (dw_j(\varphi_j)/d\varphi_j)\,\varphi_j d\ln\varphi_j$ and $d\ln\ell_j = (d\ell_j(w_j)/dw_j)\,w_j d\ln w_j$. Wages and employment respond positively to shocks.

2. For demand constrained firms with $\varphi < \underline{\varphi}$ we have $d\ln w_j = d\ln\underline{w}$ and $d\ln\ell_j = (\partial\ell_j(\underline{w})/\partial w)$ $\underline{w}d\ln\underline{w}$. Wages respond in the same direction as $\underline{w}$ and employment in the opposite direction.

3. For supply constrained firms $\varphi \in [\underline{\varphi}, \overline{\varphi}]$ the expressions are the same as for demand constrained firms, but the employment response is in the same direction as the wage floor.

Proposition 2 establishes the response of wages and employment to the entire profile of shocks. The model implies that unconstrained firms should only be affected by their own shock.[9] For constrained firms, Proposition 2 implies that shocks that increase average CB revenue will unequivocally lead to higher wages. However, the effect of shocks on average employment is heterogeneous: the lowest productivity firms will respond to shocks reducing employment, those with intermediate productivity will increase employment, and the highest productivity firms will not be affected.

In this section I showed that shocks propagate across firms via changes in the wage floor. The responses of firms to shocks via CB depend on their productivity, with employment responding in opposite directions for low and medium productivity firms. The empirical analysis that follows will test these predictions to establish whether the effects can be explained by the wage floor.

3 Context and Data

In this section I describe the labor market institutions in Argentina, the economic environment of the period under study, and the data. The goal is to provide context for the comparisons that I make in the empirical analysis and discuss the role of wage floors. I offer a succinct overview here, while Appendix B provides more details.

3.1 The socio-economic context

This subsection describes the socio-economic context surrounding the years 2009–2013, which is the period of interest for the construction of the export shocks. Figure 1 visualizes the evolution

[9] Unconstrained firms may be affected by the wage floor if spillover effects are present. For example, Cengiz et al. (2019) in the US and Milgrom and Verdugo (2022) in Chile find evidence of spillovers from minimum wage policies.

of the economy in 2004–2019. Panel (a) shows the real GDP and Panel (b) the inflation rate.

Exports experienced an increase in the aftermath of the Great Recession, partly due to a surge in international demand for Argentine products. Panel (c) shows that exports of manufactured products increased between 2009 and 2013, a period that—as I discuss in the next section—coincides with an increase in international demand. Exports declined afterwards due to deteriorating macroeconomic conditions such as the appreciation of the real exchange rate.

The increase in inflation prompted an increase in the frequency of negotiations to revise wages. Panel (d) shows that CBA alterations increased almost threefold between 2005 and 2010, and remained high afterwards. CBA alterations are agreements that modify some provisions in comprehensive "master CBAs," and tend to update wage floors. (I discuss the different types of CBAs in the next subsection.) The fact that nominal wage revisions are frequent turns out to be a feature of the empirical context, as *real* wage floors can potentially change significantly in response to changes in the economic conditions faced by firms within CB units.

3.2 Labor market institutions in Argentina

The Argentine law sets general regulatory standards for workers, and various CB agreements (or CBAs) establish additional regulations for specific worker subgroups. The law stipulates that individual employers cannot modify the terms of the CBAs to the detriment of workers, but can improve upon them. Consequently, CBAs establish minimum standards for covered workers.

The government grants bargaining privileges to a single union per "area of representation," defining the scope of CBAs. These areas may be defined by industry, occupation, geography, or a single employer. Only unions with privileges can sign either one of two types of agreements. First, a "master CBA" is a comprehensive contract that sets standards for all workers, and defines the CB units.[10] Second, "CBA alterations" act as amendments to master CBAs and are signed more frequently. Usually, national master CBAs allow for a proportional difference in wage floors across regions (most typically for southern provinces) that is determined in the master CBA and thus is constant over time. Additionally, the law allows a single employer to hire employees under multiple CBAs, a scenario that applies to around 16% of firms.

CBAs are extended to all workers and firms in a given area of representation, resulting in nearly universal CB coverage. Managerial positions and specific occupations are usually excluded. In 2014, of a total 6.4 million formal workers, 80% were covered by CBAs under the main private sector bargaining regime, 7% were not covered by any CBA, and the remaining workers were under special regimes (Ministerio de Trabajo, Empleo y Seguridad Social 2023).

Heterogeneity in CB coverage. The legal structure results in an idiosyncratic CB network in which firms in the same economic sector and province usually operate under different CB units. One reason is that the definition of coverage of master CBAs, which follows from the union's

[10] A master CBA may be modified by a new master CBA that supersedes it, or a CBA alteration that updates some provisions in it. Expired master CBAs remain in force until renegotiated due to a clause known as "ultra-activity."

area of representation, is based on inexact verbal descriptions. As an example, Appendix Table 1 shows the description of three textile CBAs. It is hard to tell from the descriptions what CBA would apply to a given sector within textile manufacturing, such as "weaving of textiles" or "manufacture of wearing apparel." For instance, CBA 0500/07 covers "workers in the textile industry" and 0501/07 covers "workers ... from clothing industrial companies and related fields."

Figure 2 illustrates the heterogeneity in CB coverage, using the primary CBA in the firm (as defined later in this section). The figure shows the concentration of firms across different CB units within 4-digit sector and province cells using an HHI index. An exact alignment between CB units and sector-province cells would result in an HHI of 1 for all cells. However, we observe a large dispersion of HHIs. 40% of cells have an HHI lower than 0.5, equivalent to 2 equal-sized CB units in the cell. As a concrete example, Appendix Figure 1 shows the proportion of firms in different CB units in three textile-related economic sectors. The CB units described in Appendix Table 1 appear in all sectors.

3.3 Data

I obtain information on labor market outcomes from two administrative registries. The first one is Argentina's matched employer-employee dataset, covering 2007–2020. The key variable of interest is total monthly compensation. The data also include the worker's gender and age, the hiring modality for the job, and the firm's fiscal location (postal code) and 6-digit economic sector (built from ISIC codes).[11] Importantly, the data do not include hours or full-time status, though I use the hiring modality to proxy for the latter. Second, I use data from *Simplificación Registral*, a national system introduced in 2008. The dataset contains information at worker hiring and termination dates, and includes the (master) CBA code, the category within the CBA, and an occupation code. I join these datasets using firm and worker identifiers.

Upon cleaning the CBA code variable, I assign a "primary CBA" to each firm, which defines its CB unit.[12] I do so by selecting the most common CBA code in the firm. Sometimes all workers in a firm have a missing CBA code. In this case I assign the primary CBA based on the firm's location and economic sector. All firms in the economy are assigned to a CB unit.

I use two additional data sources for the empirical analysis. First, to construct export shocks I rely on international trade flows data from BACI-CEPII (Gaulier and Zignago 2010) for 2007–2020, and data from Argentine Customs for 2011–2020. I harmonize product and country codes to match the datasets. Second, I use a survey of businesses conducted jointly by the Ministry of Labor and the Ministry of Science to study the effect of shocks on firm revenue and expenditures.

[11] A firm's fiscal postal code may not coincide with its production site. However, fiscal provinces are likely to coincide, which is why I use this geography in my empirical analysis.

[12] I take several steps to impute missing values for the CBA code, increasing coverage from 70% to 85% after 2011. The most important step involves imputing the CBA to workers with missing values in a firm that only declares a single CBA. I also harmonized master CBA codes that change over time. Appendix B.2.3 provides details.

Estimating wage floors. While the administrative data provides a CBA code and a within-CBA occupation code, it does not directly include wage floors.[13] I estimate wage floors by detecting bunching in the distribution of total wages within CBAs, categories, CBA-regions, and month cells, starting in 2011 when the CBA codes are more reliable. Then, I smooth the wage floors so that the percent difference across categories is constant and all categories experience the same growth over time, reflecting the usual structure of CBAs. I obtain $243,400$ monthly wage floors between 2011 and 2017, corresponding to 450 CBAs and $2,619$ CBA-occupations.[14]

The estimated wage floors align well with a sample of manually collected floors. The correlation between the levels of the estimated and manually collected wage floors is 0.87. However, the wage floors collected manually exhibit on average somewhat larger first differences. While measurement error is likely present, there is no reason to suspect it would be correlated with the export shocks. Appendix C discusses the details of the estimation and validation of the wage floors.

Wage floors are strongly binding in the labor market. Appendix Figure 2 shows the distribution of the wage to wage floor ratio for workers in the main hiring modality in 2012. The average ratio is 1.28 and 47.1% of workers are below the value of 1.1, though these numbers vary by CBA. Cardoso and Portugal (2005) pioneer the usage of the "wage cushion," defined as log ratio of the wage and the wage floor, to measure wage flexibility under CB. Card and Cardoso (2022) find an average wage cushion 0.20 in Portugal, and Bhuller et al. (2022) of 0.15 in Norway.

A yearly panel of firms. For my main analysis, I create a panel dataset at the firm-year level for the period 2007–2017, and for 2011–2017 for wage floors. I compute the average monthly real wage and wage floor for each firm-year, deflating monthly wages by the consumer price index. I also tally the number of workers that appear in a firm in the year and compute the share of workers in the main hiring modality (which proxies for the share of full-time workers). I detail the construction of the export shocks and the baseline analysis sample in the next section.

4 Empirical Strategy

In this section, I present the empirical strategy used to study the propagation of shocks through collective bargaining units. The strategy leverages fluctuations in world import demand for granular products to construct economic shocks at the firm and CB unit levels. Then, I use these shocks in a difference-in-differences strategy. Additionally, I discuss the construction of the baseline analysis sample and the identification assumptions underlying the analysis.

4.1 Overview

This paper aims to uncover the causal effects of a shock to the firm's product demand and to the product demand of a firm's peers in the CB unit. An ideal experiment to identify these

[13] A dataset containing wage floors is not readily available. Acquiring this information typically requires reviewing the text of CBAs, usually available in scanned PDF format, making manual collection of wage floors impractical.

[14] I exclude June and December, as these months correspond to the 13th-month salary payments.

effects would consist of randomly changing the demand for different granular products to generate exogenous variation in the total product demand to each firm. The effect of a product demand change at the firm could be estimated by comparing firms with different changes, conditional on the changes to the firm's peers. Section 2 suggests that bargaining outcomes are determined by average conditions across firms in the CB unit. As such, the experiment should generate variation in the average of the firm-level changes across CB units. Then, conditioning on the product demand change of the firm, the effect of the average product demand change at the CB unit could be estimated comparing firms across CB units.

To approximate this experiment I construct shocks to firm and CB units using variation arising from changes in world import demand for granular products exported by Argentine firms. I then rely on a difference-in-differences strategy to estimate the effects of these shocks on firm outcomes. I detail this strategy in the remainder of this section.

4.2 Identifying trade shocks

Building on Hummels et al. (2014) and recent shift-share literature (Borusyak et al. 2022b), I use variation in international demand for country-products exported by Argentine firms interacted with exposure shares to construct shocks at the firm and CB unit levels. Let WID_{pt} denote the world import demand of country-product $p \in \mathcal{P}$ from the world (excluding Argentina) in year t, where $\mathcal{P}$ is the set of country-products. The country-product level shock is

$$f_p = \frac{1}{2}\sum_{t=2012}^{2013} \ln WID_{pt} - \frac{1}{2}\sum_{t=2009}^{2010} \ln WID_{pt}.$$

I then define a shock to each firm j as

$$z_j = \sum_{p \in \mathcal{P}} s_{jp} f_p,$$

where $s_{jp} = EX_{pj} / \left(\sum_{p' \in \mathcal{P}} EX_{p'j}\right)$ and EX_{pj} is the sum of the value exported to country-product p in 2011 and 2012. Similarly, I define a shock to each CB unit c as

$$z_c = \sum_{j \in \mathcal{J}} s_{cj} z_j = \sum_{j \in \mathcal{J}} s_{cj} \left(\sum_{p \in \mathcal{P}} s_{jp} f_p\right) = \sum_{p \in \mathcal{P}} s_{cp} f_p.$$

where $s_{cj} = L^{EX}_{cj} / \left(\sum_{j' \in \mathcal{J}} L^{EX}_{cj'}\right)$ is the share of workers in exporting firm j in CB unit c, and $s_{cp} = \sum_j s_{cj} s_{jp}$ denotes the contribution of p to c's shock.[15] The CB shock z_c is the average product demand change across firms in the CB unit. As such, one would expect it to shift average revenue in the CB unit, which according to Section 2 should determine changes in wage floors.

[15]Shocks are not defined for CB units without exporting firms, which are excluded from the analysis.

The data suggest an exogenous shift in international demand for different products following the Great Recession that translated into heterogeneous shocks to firms and CB units. To see this I compute time-varying versions of the shocks using changes in WID relative to 2009. Figure 3 visualizes the evolution of the average shock for different levels of z_j and z_c. Panel (a) presents the firm shocks and Panel (b) the CB shocks. We observe stable trends that start to diverge around 2010. I discuss evidence supporting shock exogeneity in the next subsection.

Baseline analysis sample. To ensure that outliers do not drive the results, I define a "baseline sample" of exporting firms that incorporates several restrictions. First, I keep firms that were operational in 2007 through 2009, had an average of 1 to 500 workers in the same period, and their value exported in 2011–2012 was within the 1st and 99th percentiles. I drop firms that experienced a change in average world import demand between 2007–2008 and 2009–2010 in the bottom 1% or top 1% of the distribution as well.[16] Second, I exclude firms in CB units with less than 5 exporting firms in 2011–2012, and CB units with shocks falling within the bottom or top 1% of the distribution. I discuss the robustness of the results to these restrictions.

Appendix Table 2 shows cross-sectional statistics of the baseline sample. The sample contains $7,972$ firms, spanning 222 4-digit and 467 6-digit sectors. In the pre-period the average firm has 46.0 employees, and 23.5% of firms have less than 10 workers. Appendix Table 3 shows cross-sectional statistics of the 174 CB units that cover these firms. Lastly, Appendix Table 4 shows statistics of the main panel of firms used for estimation, spanning from 2007 to 2017.

Appendix Figure 3 shows the distribution of both z_j and z_c in the baseline sample. The distribution of the firm shock is bell-shaped around a positive mean. The variation in CB shocks is somewhat larger, suggesting some within-CB unit correlation of firm shocks.

4.3 A difference-in-differences strategy

Let $I\{\cdot\}$ be an indicator function. The static difference-in-differences (DiD) model is

$$y_{jt} = \theta z_{c(j)} I\{t \geq 2012\} + \lambda z_j I\{t \geq 2012\} + \alpha_j + X_j' \psi_t + \delta_{\ell(j)t} + \varepsilon_{jt}. \tag{5}$$

where y_{jt} is firm j's outcome in year t, $c(j)$ indicates the CB unit of firm j, and z_c and z_j are the CB and firm shocks, respectively. The parameters of interest, λ and θ, can be interpreted as the effect of an increase in the average product demand at each level on the evolution of y_{jt}. I study the effect of the firm shock on revenue in a sample of firms observed in a survey of businesses to interpret the magnitude of these effects in terms of revenue as well.

The model controls for several potential confounders. First, I include firm effects α_j to control for time-invariant firm characteristics. Second, I include baseline characteristics X_j interacted

[16] As a result of matching imperfections between Argentine customs data and international trade flows data, 0.81 percent of exporting firms in 2011–2012 have less than 99 percent of their exporting value matched to a country-product (see Appendix B.2.2 for details). In the terminology of Borusyak et al. (2022b), these firms have "incomplete shares". I drop these firms from the analysis.

with time to control for time-varying factors that may affect different firms differently, such as the increase in the federal minimum wage in 2007–2011. In particular, X_j consists of pre-period firm size and wage level dummies interacted with economic sector dummies.[17] Third, I include local labor market by year fixed effects $\delta_{\ell(j)t}$, where $\ell(j)$ indicates j's local market, to control for common trends in the local labor market. As baseline, I define local labor markets as the interaction between provinces and 4-digit ISIC sectors, though I vary this definition in robustness checks. I further show that the results are robust to different sets of controls.

To study treatment effect dynamics I use a specification analogous to (5) that includes interactions of year dummies with the shocks. Formally,

$$y_{jt} = \sum_{s \in \mathcal{S}} \theta_s z_{c(j)} I\{t = s\} + \sum_{s \in \mathcal{S}} \lambda_s z_j I\{t = s\} + \alpha_j + \delta_{\ell(j)t} + X_j' \psi_t + \varepsilon_{jt},$$

where $\mathcal{S}$ is the set of years from 2007 to 2017. Since Figure 4 suggests that effects start in 2011, I exclude 2010 from $\mathcal{S}$. I exclude 2011 for the wage floor as the wage floor data starts in that year. Throughout the paper, I cluster standard errors at the CB unit level.

Estimation of these models requires independent variation in each shock. Appendix Figure 4 plots the firm shock versus the CB shock for each firm. Panel (a) focuses on the raw data, and Panel (b) shows the same plot after residualizing on local labor market fixed effects. There is a small positive correlation in the raw data, which is somewhat smaller after controlling for local labor market effects.[18] Overall, there is significant independent variation in each of the shocks.

Identification. Borusyak et al. (2022b) demonstrate that a shift-share model can be cast in terms of the shifting variable and present assumptions on the assignment process of the "shifts" (in my case, world import demand shocks for each country-product pair) that are sufficient for identification. Appendix D develops the identification argument from Borusyak et al. (2022b) in my setting, extending it to a scenario with two shift-share variables.

The key assumption is that changes in world import demand for each country-product p are quasi-randomly assigned with respect to average residuals of firms that export in that country-product, both when weighted by exposure via firms and via CB units.[19] This assumption would be violated if, for example, firms that are growing select into products in which world import demand is growing. The assumption that shocks are quasi-randomly assigned implies that shocks should not be correlated with pre-period observables.

Is the assumption of quasi-random assignment of shocks plausible? As argued by Garin and Silvério (2023), the 2009 international crisis generated unexpected changes in world import demand for different products that should be uncorrelated with pre-determined economic conditions.

[17] More precisely, I define two categorical variables using 2007–2009 averaged data: one for firm size (1–19, 20–124, 125–500) and one for the wage level (above or below the median wage).

[18] The raw correlation is slightly stronger when I focus on larger firms. This is to be expected given that shocks to CB units are defined as employment-weighted averages of firms shocks.

[19] A condition for this interpretation to hold is that the shares in the definition of the shift-share variables sum to 1 (Borusyak et al. 2022b). By construction, this holds for both z_j and z_c.

To explore this idea, Appendix Figure 5 shows estimates of firm-level regressions of a pre-period outcome on the firm shock (from years 2007–2009). Using the baseline sample and the firm controls I find that the firm shock does not predict the pre-period firm shock, the firm size, the mean wage among workers, the share of workers in the main hiring modality, and an indicator for the retail CB unit 0130/75. These results support the quasi-randomness assumption. Results are similar when using all firms and no controls, except for the pre-period firm shock. Panel (a) of Appendix Figure 6 shows that the apparent negative auto-correlation in the firm shock is mainly driven by mean reversion at the tails of the distribution, in particular the left tail.[20] Panel (b) shows that this correlation disappears once I drop firms with extreme pre-period shocks in the baseline sample.

The quasi-randomness assumption implies that CB-level shocks should not be correlated with pre-determined CB-level outcomes. Appendix Figure 7 shows that the CB shocks are not correlated with the pre-period CB shock, the number of firms in the CB unit, and the share of employment in exporting firms.

Panel event-study. To estimate the effects of firm shocks on revenue I rely on the survey of businesses, which has fewer observations only in the years 2010–2012 and 2014–2016, and a panel event-study design in the spirit of Freyaldenhoven et al. (forthcoming). This estimation strategy has two advantages. First, it uses variation in world import demand beyond the change between 2009 and 2013, leveraging more observations in the smaller survey sample of firms. Second, it estimates pre-period coefficients asking whether future shocks are correlated with current outcomes, meaning that outcome data from the pre-period are not needed to estimate pre-period coefficients.

5 Empirical Results

This section presents the main empirical results of the paper. First, I explore the effects of both the CB shock and the firm shock on wages and wage floors. Second, I explore the effect of shocks on employment, and use heterogeneity analysis to study whether responses to CB shocks can be rationalized by a wage floor mechanism. Third, I compare my estimates with the literature. Finally, I present a set of falsification tests and robustness checks.

5.1 Effects on wages and wage floors

The raw data suggest stronger effects of CB shocks on wages than of firm shocks. Figure 4 visualizes the average evolution of wages by level of the CB shock and firm shock, relative to 2009. CB units with larger shocks consistently experience larger wage increases, whereas there seems to be a small increase in wages, if any, following a firm shock. The figure also shows small

[20] Using similar data for Portugal, Garin and Silvério (2023, Figure A.4) also finds evidence of mean reversion at the left of the distribution of export shocks.

divergences between firms before the Great Recession, suggestive of parallel trends before the shocks. Next, I ask whether these patterns are robust to controls using the DiD model.

The estimates of the dynamic DiD model indicate that both mean wages and mean wage floors increase in response to a CB shock, and that mean wages increase in response to a firm shock. Figure 5 shows the estimates. Panel (a) shows a strong and stable increase in wages following a CB shock, and a smaller and marginally significant increase as response to a firm shock. We cannot reject that pre-trends are zero in anticipation of either shock, consistent with the parallel trends assumption. Panel (b) shows estimates of the same model but using the log mean wage floor as outcome. As discussed in Section 3.3, the wage floor is only available since 2011, so this model uses 2011 as omitted year. We observe a strong increase in wage floors following a CB shock, of a similar magnitude as the increase in wages, and a precisely estimated null effect of the firm shock. This is reassuring as firm shocks are not supposed to affect wage floors.

I summarize the magnitude of the estimates using a static DiD model. Column (1) of Table 1 shows that an estimate of the elasticity of wages to the average world import demand in the CB unit of 0.0458, and to the average world import demand in the firm of 0.0134. While the coefficient on the CB shock is strongly significant ($t = 2.27$), the coefficient on the firm shock is only marginally so ($t = 1.92$). Column (2) of Table 1 shows an effect of the CB shock on log mean wage floors of 0.0541 ($t = 2.54$), using only observations with a valid wage floor value. Column (3) of Table 1 shows the effect of the CB shock on the "wage cushion" (the log ratio between the wage and wage floor), which is indistinguishable from zero.

The fact that wages and wage floors respond similarly to a CB shock suggests a pass-through of wage floors to wages close to 100%. Card and Cardoso (2022) estimate the effect of changes in wage floors on wages in Portugal, finding a smaller pass-through of about 50%. The authors do not use exogenous variation in wage floors. My estimates suggest that, in Argentina, wage negotiations in a CB unit serve as a reference point for wage-setting even from firms that pay wages higher than the floor. Another possibility is that other components of pay are affected by the CB shock, such as productivity bonuses, leading to the impression of a large pass-through.

The magnitude of effects. To better interpret the magnitude of these effects, I scale them by the effect of the firm shock on revenue. Appendix E discusses estimates of the firm shock on revenue and expenditures using a panel event-study estimation strategy. I obtain that a 10% increase in average world import demand raises product market sales by 1.13%. The results show null pre-trends, consistent with the identification assumption. I also show estimates of the effect of the firm shock on labor expenditures, which also increase but by a smaller magnitude.

Following these results, assume that a 10% change in z_j translates into a 1.1% increase in firm revenue. Then, a 10% CB shock can be interpreted as a 1.1% increase in (weighted) average firm revenue in the CB unit. Using the results in Table 1, this implies that a 10% increase in revenue of all firms in the CB unit increases wages by $10 \times (0.0458/0.11) \approx 4.2\%$. Similarly, a 10% increase in firm revenue increases wages by $10 \times (0.0134/0.11) \approx 1.2\%$.

Heterogeneity by CB unit size. A natural question is whether the effects differ across CB units. Estimates in Table 3 interact the CB shock with an indicator for a large CB unit, defined as having more firms than the median number of firms in the 4-digit sector by province cell. These estimates exclude firms in cells with less than one CB unit. Columns (1) through (4) show that the wage effects of CB shocks are driven by larger CB units, suggesting that they are more effective at raising wages in response to shocks.

5.2 Effects on employment

Trends in the raw data are not as clear for employment as they are for wages. Appendix Figure 8 shows the evolution of employment by level of the CB shock and firm shock, relative to 2009. We observe lower employment growth in firms that experienced a low CB shock, and stronger growth in firms with a high firm shock.

Figure 6 shows the dynamic DiD estimates. Panel (a) shows a clear increase in employment following a firm shock, while Panel (b) reveals a noisy, statistically indistinguishable from zero, response of employment to the CB shock. The theory discussed in Section 2 suggests that this is to be expected, as firms with different productivity levels will systematically respond differently to a wage floor. Table 1 shows estimates of the static DiD model. Column (4) shows the response of employment, which is positive and marginally significant for the firm shock ($t = 1.42$), and indistinguishable from zero for the CB shock ($t = -0.66$).[21] Columns (5) and (6) show that neither shock seems to affect the share of workers in the main hiring modality nor firm exit.

Heterogeneity by CB unit size. Columns (5) and (6) of Table 3 show that the effect of CB shocks on employment is not statistically different from zero for either small or large CB units.

5.3 Responses to CB shocks by firm productivity

A threat to the interpretation that CB negotiations are responsible for observed CB effects is that these are actually driven by unobserved shocks to product demand that are correlated with CB shocks, such as a concurrent increases in local demand. If this is the case, then a positive CB shock should increase both wages and employment across different types of firms. However, as discussed in Section 2, a positive CB shock that affects wage floors should cause low-productivity firms, which are constrained by the wage floor, to decrease employment.

To differentiate these hypotheses I explore the response of firms with varying productivity levels to CB shocks. I employ a static DiD model as in (5), but I interact the treatment variables with a proxy for firm productivity. In particular, to account for regional differences in cost of living I construct quartiles of the distribution of average wages in 2007–2009 within each province, excluding provinces with less than 6 firms. Appendix Figure 9 shows that the estimates are similar

[21] The non-significant effect of the firm shock on employment is driven by the fact that coefficients start to increase in 2010. The years 2010–2011, which are in the pre-period in the static model, bias the estimate towards zero.

when using a different definition of quartiles. I interact the local labor market effects with quartile indicators, as a result I compare firms with similar pre-period level of wages.

The evidence of the wage effects of CB shocks is consistent with the hypothesis that they operate mainly through wage floors. Table 2 shows the results. Because I drop firms in local labor markets with less than 6 firms and I increase the number of fixed effects, the estimates are noisier. Column (1) shows that the effects of the CB shock on wages are declining in the pre-period level of wages, and column (2) shows that this is not the case for the wage floor. This is consistent with a wage floor that is more binding for low-wage firms. Column (3) shows suggestive evidence that the wage cushion increases in low-wage firms and decreases in high-wage firms, although the coefficients are not statistically different from zero.

The evidence on employment effects is consistent with the presence of wage floors as well. Column (4) shows that the effect of the CB shock on employment is negative and statistically significant for firms in the lowest quartile of pre-period wages, and positive and marginally significant for firms in the third quartile. This is consistent with low-wage firms moving along the labor demand curve and with reductions in monopsony power for firms in the third quartile. The point estimate for the second quartile is positive and for the fourth quartile is negative, but neither is statistically significant. The hypothesis of joint significance of coefficients can be confidently rejected ($p = 0.0003$). Column (5) shows positive and significant effects of the CB shock on the share of workers hired in the main hiring modality for firms in the second and third quartiles.

The estimates suggest that employment elasticity to the firm wage differs across the productivity distribution. Taking the point estimates at face value, for the first quartile I obtain $-0.1857/0.1509 \approx -1.23$ and for the third quartile $0.0926/0.0871 \approx 1.06$. For comparison, Cengiz et al. (2019) estimates an employment elasticity to own wage of 0.41. However, the elasticity estimates are quite noisy, making it hard to draw strong conclusions about their magnitude.

5.4 Comparison with the rent-sharing literature

The estimates of the effects of firm shocks are comparable to those in the rent-sharing literature. The most closely related paper is Garin and Silvério (2023), who use a similar estimation strategy in Portugal. The authors find that an "idiosyncratic" export shock to a firm increases log sales by 0.143 (Panel A of Table 4) and log monthly wages by 0.022 (Panel A of Table 6). The implied elasticity is 0.15, comparable to my estimate of 0.12. The authors also study the effects of "common" export shocks to firms, but do not study CB shocks. More generally, Jäger et al. (2020, Figure II) review the rent-sharing literature and find that elasticity estimates using worker-level micro-data are on average 0.099, which is comparable to my estimates as well. Card et al. (2018, Table 1) find an average rent-sharing elasticity of studies using worker-level specifications of 0.08.

The larger estimates of effects of CB shocks on wages suggest that the magnitude of rent-sharing elasticities might depend on the level of aggregation of the shocks. Jäger et al. (2020, Figure II) also find that industry-level specifications, and calibrations of bargaining power in macro models, tend to find larger elasticity estimates. Similarly, industry-level specifications in Card et al. (2018, Table 1) also tend to find larger elasticities. For example, Abowd and Lemieux (1993) use contract wages from CB agreements and shocks to international prices in Canadian manufacturing. Their estimates suggest a rent-sharing elasticity of around 0.2, twice as high as the average worker-level estimates. I find that the response to a change in average product demand at the CB unit is larger than the response to a comparable change in demand at the firm, implying stronger wage responsiveness to CB shocks. Card and Cardoso (2022) find that the mean change in value added in the CB unit is associated with the change in wage floors in Portugal, but do not compare the effects of shocks to firms and CB units.[22]

5.5 Falsification tests and robustness checks

In this subsection I present additional empirical exercises to assess the robustness and validity of the empirical results.

Firms with multiple collective bargaining agreements. A minority of firms declare workers under multiple CB units, presenting an interesting case study to determine whether the relevant factor is the CB unit covering the worker or other unobserved channels that may propagate the shocks. If the shocks to firms in the same CB unit propagate through other channels, we should not observe CB unit-specific effects within the firm. However, if the shock propagates through CB, we should observe an effect on wages of workers declared under that CB unit.

I conduct a worker-level analysis and find that wages respond to shocks to their own CB unit, even when controlling for firm by year fixed effects. To abstract from the extensive margin, I focus on a sample of workers that are observed in an exporting firm before and after 2011.[23] Figure 7 shows the estimates. Panel (a) replicates the baseline results: a positive shock to the primary CB unit in the firm increases wages. Panel (b) estimates a model that uses the CB shock of the worker's CB unit, instead of the primary CB unit in the firm, as treatment variable. Importantly, as it controls for firm by year fixed effects, this specification allows for arbitrary time-varying unobservables at the firm level that should capture any spillover effects of the CB shock through other channels. The figure shows a consistent positive effect of the CB shock on wages.

Appendix Table 5 provides detailed regression results using static DiD models. Columns (1) through (3) show consistent positive effects of CB shocks to the primary CB unit of the firm when varying local labor market controls, with column (3) corresponding to the specification in Panel (a) of Figure 7. Column (4) shows that the effect of the primary CB shock is driven by workers that are actually declared under the primary CB unit in the firm. The effect on workers declared under a different CB unit is smaller in magnitude and not statistically significant. Finally, Column (5) replicates the within-firm model in Panel (b) of Figure 7.

[22] Card and Cardoso (2022) do not rely on quasi-random variation, and find an implied elasticity of around 0.1. These estimates are hard to compare to my results as I cannot compute the effect of CB shocks on value added.

[23] Specifically, I focus on workers employed in the same exporting firm in 2008, 2011, and 2014, and with a non-missing CB agreement code.

Controls for local-labor-market effects. A key concern is that the results may be driven by spillovers between firms through proximity instead of the CB network. I explore this possibility in Appendix Table 6, where I show estimates that use three alternative specifications to control for local labor market trends: the baseline specification that interacts 4-digit economic sector with province and year fixed effects; the same specification but using 6-digit economic sector instead; and a third specification that uses 2-digit sector and additionally controls for an export shock computed similarly to the CB shock but at the 6-digit sector level. The effects of firm and CB shocks are similar across specifications, suggesting that spillovers through proximity are not driving the results. Columns (3), (6), and (9) indicate that 6-digit sector shocks affect wages and employment, but do not affect wage floors, consistent with spillovers through proximity.

Effects on non-exporters. I have so far focused on exporting firms, but CB shocks should affect non-exporting firms as well. To test for this, in Appendix Figure 10 I estimate the effect of CB shocks on wages and the wage floor using only firms that do not export. I find a similarly-sized effect of CB shocks on mean wages, implying that non-exporting firms are also affected by the CB shock. Given that the effect on wages seems to start in 2011, the effect on wage floors computed relative to 2011 appears somewhat smaller in this sample.

Placebo CB units. To assess the validity of the design I estimate the DiD model using several placebo CB shocks. To do so, I construct placebo CB networks by randomly shuffling the CB unit code across all firms, across firms within 1-digit sector and province cells, and across firms within 2-digit sector and province cells. Appendix Figure 11 shows estimates using the placebo CB shocks, and reveals no evidence of effects of the placebo shocks on wages or wage floors.

Other robustness checks. A final concern is that regression controls or sample restrictions may be driving the results. Appendix Table 7 varies the set of controls included in the regression, and shows very similar results when dropping the firm-level controls, controlling for the pre-period firm shock directly, and excluding the control for the pre-period CB shock. Appendix Figure 12 replicates the dynamic estimates on wages and employment keeping firms with extreme values of the pre-period firm shock, relaxing one of the sample restrictions. The effects of the CB shock and those of the firm shock on employment are very similar to the baseline estimates. However, for the effect of the firm shock on wages there is a significant pre-period coefficient in 2007, and no effect of the firm shock after 2011. While I find the exclusion of these firms to be justified, it is important to note how the results change when they are included.

Appendix Table 8 varies the set of CB units included in the sample. Column (2) shows that keeping CB units with less than 5 exporting firms in the sample does not affect the results. Column (3) shows that dropping CB units with less than 30 exporting firms seems to increase the effect of the CB shock. Consistent with Table 3, larger CB units seem to respond more strongly to CB shocks. Column (4) shows results of the static model when excluding the retail CB unit. Columns (5) through (8) replicate the finding of non-significant effects of the CB shock on employment.

6 Structural Model

While the empirical estimates indicate that shocks propagate through the existing CB network in Argentina, they cannot inform us about the role of the network in shaping the effects of shocks. To study this question I develop a spatial economic model of the labor market with collective bargaining. Unlike in Section 2, the model in this section allows for multiple local labor markets, explicitly incorporates workers' labor supply decisions, adds formal structure to the firm's problem to obtain tractability, and introduces multiple CB units. The model is estimated and used to analyze how shocks propagate under counterfactual CB networks in the upcoming sections.

6.1 Set-up

There is a fixed population of N_r in each region $r \in \mathcal{R}$, each with a separate labor market. Regions are divided into local labor markets $g \in \mathcal{G}$, each characterized by a single economic sector $k \in \mathcal{K}$. I denote by $\mathcal{K}1$ the broadest grouping of economic sectors. A given sector and region cell may be partitioned into multiple local markets. Each local market contains a continuum of firms, indexed by j, and the measure of firms is given by M_g. The collective bargaining network is a partition of local labor markets denoted by $\mathcal{C}$. Within each CB unit $c \in \mathcal{C}$ a union and an employer association bargain over a single wage floor $\underline{w}_c$ binding for all local markets in the unit.[24]

Different actors take decisions sequentially. First, CB units play a simultaneous-move game to determine wage floors $\{\underline{w}_c\}_{c\in\mathcal{C}}$. Second, firms j in each g draw productivities $\varphi(j)$ from a known distribution $F_g(\varphi)$. Third, workers in each region r decide, first, whether to enter the formal labor market or not, and second, their supply of labor to each firm. Simultaneously, firms decide employment and wages under monopsonistic competition. Wage indexes adjust to clear regional labor markets. Finally, production takes place and workers earn the wage $w(j)$ in their firm j.

To focus on the role of the CB network, I abstract from the goods market. I will model economic shocks as changes in the productivity distributions of local labor markets.

6.2 Solving the model

I solve the model backwards starting with the decisions of workers and firms, and then moving into the bargaining game between CB units.

6.2.1 Labor supply

Conditional on formal labor market entry, a worker i in region r has an indirect utility of working for firm j given by $V_{ri}(j) = A_{k1(j)}w(j)\xi_i(j)$. $A_{k1(j)}$ is an amenity value specific to j's broad sector, $k1(j)$, and $w(j)$ is j's wage. I assume that the idiosyncratic component $\xi_i(j)$ follows a Fréchet (or type-2 extreme value) distribution with scale parameter η, as is standard in the literature of

[24] As discussed in Section 3.2, there may be constant regional differences in wage floors within a CB unit. These differences do not respond to shocks, os CB units effectively bargain over a single wage floor.

discrete choice models. Then, it can be shown that labor supply to a firm located in r is

$$\ell(j) = \left[\frac{A_{k1(j)}w(j)}{W_r}\right]^\eta \tag{6}$$

where W_r is a region-specific aggregate wage index. Derivations are available in Appendix F.1. Note that η can be interpreted as the elasticity of labor supply to the firm.

Since $\xi_i(j)$ follows a Fréchet distribution, the expected utility of formal employment in region r, V_r, is proportional to the regional wage index. Let $\Gamma(\cdot)$ denote the gamma function, then

$$V_r = \Gamma\left(\frac{\eta-1}{\eta}\right)W_r. \tag{7}$$

Before choosing a firm, workers decide whether to work formally or not. Workers' preferences for formal employment in r depend on the value of formal employment V_r, the value of non-formal employment b_r, and an idiosyncratic shock that follows a Gumbel (or type-1 extreme value) distribution with shape parameter ζ. The resulting formal employment share is given by

$$\mu_r = \frac{V_r^\zeta}{V_r^\zeta + b_r}. \tag{8}$$

A larger value of ζ implies a larger elasticity of the employment share to changes in V_r.

6.2.2 Labor demand

Each firm j draws a productivity $\varphi(j) \geq 0$ from a local labor market-specific distribution $F_g(\varphi)$ defined over $[\varphi_{g0}, \infty)$. Firms in local labor market g maximize profits with a linear technology:[25]

$$(\ell(j)^*, w(j)^*) = \arg\max_{\ell,w}\left\{\varphi(j)\ell - w\ell \;\middle|\; \ell = \left(\frac{A_{k1(j)}w}{W_{r(j)}}\right)^\eta,\ w \geq \underline{w}_c\right\}. \tag{9}$$

The wage floor may be binding or not. If not, we say that the firm is *unconstrained*, in which case the solution to the firm's problem for a firm with productivity $\varphi(j)$ is $w(\varphi(j)) = \mu\varphi(j)$, where $\mu = (\eta/(\eta+1))$ is the markdown factor. If the wage floor is binding and the firm is *constrained* we have that $w(\varphi(j)) = \underline{w}_c$. The quantity of labor can be obtained by replacing the wage in the labor supply curve (6). Appendix F.1 shows exact expressions for employment and profits.

If the wage floor is high enough, the threshold that determines whether firms are constrained in g is given implicitly by $w(\overline{\varphi}_g) = \underline{w}_c$, which results in

$$\overline{\varphi}_g = \mu^{-1}\underline{w}_c. \tag{10}$$

Firms with $\varphi \leq \overline{\varphi}_g$ will pay exactly the wage floor. Similarly, firms leave the market if ex-post

[25] I assume a linear technology for analytical convenience. The main results will be qualitatively similar under alternative specifications.

they experience negative profits $\pi(\underline{\varphi}_g) < 0$, implying thresholds

$$\underline{\varphi}_g = \underline{w}_c. \tag{11}$$

Consistent with Section 2, employment may decline following wage floor increases when firms with $\varphi < \underline{\varphi}_g$ exit. Unlike in Section 2, the model does not feature demand-constrained firms, so the only reason for employment declines is firm exit. If $\underline{w}_c < \mu\varphi_{g0}$, then no firm will be constrained.

The distance between the thresholds depends on η. A higher η implies lower heterogeneity in idiosyncratic worker-level preferences $\xi_i(j)$, and so more responsiveness to wages. In the extreme case of $\eta \to \infty$, the markdown vanishes and the wage floor is only binding for the lowest productivity firm. A lower η implies more heterogeneity in $\xi_i(j)$, and so lower responsiveness to wages. If $\eta \to 0$, the markdown goes to 0 and so the bargained wage floor becomes binding for all firms.

Parametrization and the share of firms bunching. I assume that productivities in local labor market g are drawn from a Pareto distribution with shape $\alpha > 1$ and cdf given by

$$F_g(\varphi) = 1 - \left(\frac{\varphi_{g0}}{\varphi}\right)^{\alpha} \tag{12}$$

for $\varphi \geq \varphi_{g0}$ and zero otherwise. The conditional cdf of productivity given a minimum value of $x > \varphi_{g0}$ is given by $F_g(\varphi|x) = 1 - (x/\varphi)^{\alpha}$. The fraction of "possible" productivities that will be observed is given by $1 - F_g(\varphi^{\min})$, where $\varphi^{\min} = \max\{\varphi_{g0}, \underline{\varphi}_g\}$. The total measure of firms in the market is then $M_g\left(1 - F_g(\varphi^{\min})\right)$. Note that this quantity is non-increasing in the wage floor (if $\underline{w}_c > \varphi_{g0}$, then the measure of firms is decreasing in the floor).

The share of observed firms paying the wage floor, or "bunching," can be computed using the conditional cdf $F_g(\varphi|\varphi^{\min})$. If $\underline{w}_c < \mu\varphi_{g0}$ the share is zero. If $\underline{w}_c > \varphi_{g0}$ then $\varphi^{\min} = \underline{w}_c$ and the share of firms bunching takes its maximum value of $1 - \mu^{\alpha}$. Otherwise, the share varies between 0 and this maximum value. Appendix Figure 13 illustrates the productivity distribution with a binding wage floor, for a situation where the share of firms bunching is in the intermediate case. Panel (a) shows the pdf. Panel (b) illustrates the wage paid at each productivity level.

Computing aggregate quantities. To compute aggregate quantities I integrate over the distribution of firms in g. To obtain average wages across active firms, or the share of firms bunching, I integrate over "observed" productivities. The minimum value of this integration might be φ_{g0}, $\underline{\varphi}_g$, or $\overline{\varphi}_g$ depending on the value of $\underline{w}_{c(g)}$. For example, with $\underline{\varphi}_g < \varphi_{g0} < \overline{\varphi}_g$ the average wage is given by $w_g = \int_{\varphi_{g0}}^{\overline{\varphi}_g} \underline{w}_c f_g(\varphi)d\varphi + \int_{\overline{\varphi}_g}^{\infty} w(\varphi)f_g(\varphi)d\varphi$, where $f_g(\varphi)$ is the pdf of the Pareto distribution and $w(\varphi) = \mu\varphi$. To compute aggregate labor demand I integrate over the distribution of "possible" productivities in g, so the minimum value will always be φ_{g0}. Formally, $L_g = M_g \int_{\varphi_{g0}}^{\infty} \ell(\varphi)f_g(\varphi)d\varphi$. I proceed analogously to determine revenue and the wage bill. These are the relevant quantity for Nash bargaining, as the measure of firms in the market will then respond to the wage floor. Appendix F.1 shows closed form expressions for these quantities.

6.2.3 Nash bargaining

Each CB unit negotiates over wage floors $\{\underline{w}_c\}_{c\in\mathcal{C}}$. I assume that both unions and employer associations are risk neutral, and that both parties have rational expectations in the sense that they know the distributions $\{F_g(z)\}_{g\in\mathcal{G}}$ and correctly anticipate the outcomes following their choice of $\underline{w}_c$. As a result, I do not use expectations below.

Following the framework in Section 2, I assume that the objective function of the union is given by $U_c\left(\underline{w}_c, \underline{\mathbf{w}}_{-c}\right) = \sum_g WB_g$, where WB_g is the wage bill in g, and the objective of the employer is $\Pi_c\left(\underline{w}_c, \underline{\mathbf{w}}_{-c}\right) = \sum_g \left(R_g - WB_g\right)$, where R_g is revenue in g. Preferences depend on the wage floor $\underline{w}_c$ by altering the share of firms bunching and the cost of labor. They also depend on the wage floor of other CB units $\underline{\mathbf{w}}_{-c}$, which affects the equilibrium wage index in each region.

Solving the Nash bargaining problem. Letting β_c denote the bargaining power of the union, the Nash bargaining problem can then be written as in (1). Denoting $WB_c = \sum_g WB_g$ and $R_c = \sum_g R_g$, the solution to a single Nash problem is given by the split rule $WB_c\left(\underline{w}_c, \underline{\mathbf{w}}_{-c}\right) = \omega_c R_c\left(\underline{w}_c, \underline{\mathbf{w}}_{-c}\right)$, where $\omega_c \in (0,1)$ is a weight that depends on the wage floors, given in (3). Importantly, the weight ω_c is affected by the derivatives of U_c and Π_c with respect to $\underline{w}_c$, which incorporate the effect of the wage floor on the aggregate wage index.

Nash-in-Nash solution. Horn and Wolinsky (1988) introduced this solution concept in the analysis of bilateral monopolies.[26] In a Nash-in-Nash solution, each individual bargaining problem results in a Nash equilibrium given that the wage floors of other CBAs are in equilibrium as well. The solution assumes that players in a given CB unit do not take into account the effect of their decision on the choice of other CB units.[27] In the setting of this model, the wage floor profile $\{\underline{w}_c^*\}_{c\in\mathcal{C}}$ is a Nash-in-Nash solution if

$$\underline{w}_c^* = \arg\max_{\underline{w}} U_c\left(\underline{w}, \underline{\mathbf{w}}_{-c}^*\right)^{\beta_c} \Pi_c\left(\underline{w}, \underline{\mathbf{w}}_{-c}^*\right)^{1-\beta_c} \tag{13}$$

for all $c \in \mathcal{C}$.

6.3 Equilibrium

Given a collective bargaining network $\mathcal{C}$ and a set of parameters, an equilibrium is defined as a set of wage floors $\{\underline{w}_c^*\}_{c\in\mathcal{C}}$, regional wage indexes $\{W_r^*\}_{r\in\mathcal{R}}$, and employment shares $\{\mu_r^*\}_{r\in\mathcal{R}}$ such that: (1) the Nash-in-Nash bargaining game is solved, (2) labor markets clear in each region. Appendix F.2.1 formally defines the equilibrium. In general, there is no closed form solution for the vector of equilibrium wage floors and regional wage indexes. Appendix F.2.2 discusses the algorithm I use to compute the equilibrium for a given set of parameters.

[26] Davidson (1988) studies a two-union bargaining game and uses a similar solution concept.

[27] This assumption is debatable if a single union takes part in multiple CB units. While I assume separate unions across CB units, Collard-Wexler et al. (2019) show that the Nash-in-nash solution can be micro-founded in a fully non-cooperative environment where players internalize the interdependence of their potentially multiple bargains.

7 Estimation and Validation of Structural Model

In this section I estimate and validate the model with aggregate data. To do so, I first construct a dataset at the local labor market level and use it to estimate the model parameters. Second, I show that the model is able to replicate several features of the data. I postpone the counterfactual exercises, which rely on the estimated model, to the next section.

7.1 Local labor markets data

To keep the number of local labor markets and CB units manageable, I simplify the economic sectors and the CB network. I specify sectors $k \in \mathcal{K}$ as a coarsening of the 4-digit sectors, and regions $r \in \mathcal{R}$ as a grouping of provinces in Centro (center, including Buenos Aires), Cuyo (west), Norte (north), and Patagonia (south).[28] I further divide the economic sector by region cells using the exporting status of firms, the CB unit codes, and (if the cell is large enough) provinces. To make sure that all local markets have a wage floor, I create "local" CB units by grouping CB units that have less than 200 employees or fewer than 1% of workers with a valid wage floor. Importantly, the definition of local CB units does not use the exporting status, so local CB units may include both an exporter and a non-exporter local labor market.

Once the spatial economy is defined, I aggregate the firm-level data to the local labor market level. To estimate the share of firms constrained by the wage floor I use the workers that were assigned a wage floor. To account for part-time work and measurement error, I define a worker as a "buncher" (i.e., with a deviation of zero) if the wage is between 90 and 105% of the floor, and between 40 and 60% if the worker is classified as part-time. Then, the share of firms constrained is simply the share of firms where all workers are bunchers. To estimate wage floors for each CB unit, I take the mean wage floor across local labor markets weighting by the share of workers with an assigned wage floor in each local labor market. I allow for regional wage floors within a CB unit only if the region's average wage floor is sufficiently different from the rest. I also compute the average wage and total employment using the firm-level data, adjusting for (estimated) part-time work. Appendix F.3 provides details on the construction of the local labor markets and the adjustments made to the aggregate data to correct for possible measurement error.

Appendix Table 11 provides summary statistics. There are 894 CB units, of which 456 are local. 63.5% of local labor markets are in the region "Centro," 13.6% are covered by a local CB unit, and 14.9% are covered by the retail CB unit. There are a few small local labor markets where the mean wage is lower than the wage floor. This is caused by two things: (1) merging different CB units together, and (2) small local labor markets relying on part-time work. Using the part-time adjusted mean wage solves the problem, but introduces some missing values. The measure of firms in each local labor market M_g is normalized so that $\frac{1}{|\mathcal{G}|}\sum_g M_g = 1$. The table also shows model-estimated minimum productivities, which are discussed in the next subsection.

[28] I drop several broad economic sectors, including agriculture, education and government, in which bargaining is regulated by different regimes. Additionally, the agricultural sector represents a small share of employment.

7.2 Calibration and Estimation

I use several strategies to estimate the model parameters, which are summarized in Table 4 and detailed in this section. Appendix F.4 provides details.

Worker problem. To estimate the preference heterogeneity parameter η and the amenity values $\{A_{k1}\}_{k1\in\mathcal{K}1}$, I leverage the relationship between firm size and firm wages implied by the labor supply to the firm (6). In particular, I estimate a regression of log mean wage on log firm size, controlling for region and broad sector fixed effects.[29] I find a value of $\eta \approx 4.10$, which is in line with the literature. This value implies a markdown factor of $\mu \approx 0.80$. I calibrate the parameter indicating heterogeneous preferences for formal employment using evidence on the extensive-margin labor supply elasticity from Chetty et al. (2011) and data on the share of formal employment in the labor market. This results in $\zeta \approx 0.28$.

Productivity distributions. I calibrate the common shape parameter α to 5.50. This value, which is in line with estimates calibrated in the literature, corresponds to the model admitting all shares below the 98.8th percentile of the distribution of observed shares of firms bunching at the wage floor. Shares above this threshold are set to the maximum feasible value in the model. To estimate the minimum productivities $\{\varphi_{g0}\}_{g\in\mathcal{G}}$, I invert the expression for the share of firms bunching at the wage floor implied by the model. I use the wage floor in each local labor market and the values of η and α from above. Appendix Table 11 shows summary statistics of the estimated $\{\varphi_{g0}\}_{g\in\mathcal{G}}$, in ARS. We observe that the φ's are on average larger than the minimum wage, which per equation (11) is the lower threshold that defines whether firms exit the market.

Outside options. Using estimates from above, the share of formal firms in each region from household survey data, and the market-clearing condition in each region, I compute $\{W_r\}_{r\in\mathcal{R}}$ and $\{V_r\}_{r\in\mathcal{R}}$. Then, I invert equation (8) to obtain the outside option parameters $\{b_r\}_{r\in\mathcal{R}}$.

Bargaining power. I invert equation (2) to obtain β_c for each c. This condition will hold in any interior equilibrium. The key variation used to pin down these parameters comes from the ratio of the derivatives of U and Π with respect to the wage floor, evaluated at the equilibrium. If the ratio is low, the union is relatively closer to its optimal wage floor (where $dU/d\underline{w}_c = 0$), suggesting stronger bargaining power. The distribution of the ratio of derivatives is shown in Panel (a) of Appendix Figure 14. After estimating the parameters I check whether they correspond to an interior equilibrium or not, and find that this is not the case for the retail CB unit. This unit is so large that it has significant general equilibrium effects. As a result, the objective function is not strictly concave, and the bargaining power actually corresponds to a local maximum. To deal with this, I hold the wage floor of this CB unit fixed when computing counterfactual equilibria.

[29] While the model implies an exact fit of this regression, it is straightforward to incorporate hours to allow for an error term, as discussed in the appendix.

Panel (b) of Appendix Figure 14 shows the estimated bargaining power parameters. There is a great deal of heterogeneity that cannot be accounted for by characteristics of CB units.

7.3 Validation

I validate the model comparing data moments with model-based predictions. To do so, I use the parameters estimated before to compute the model-implied mean wages in each local labor market in 2011–2012. I start by comparing the correlation between mean wages and wage floors in the data and the model. Panel (a) of Figure 8 shows that the model does a good job at replicating the correlation between mean wages and wage floors, where log mean wages are roughly a constant shift of log wage floors. Unsurprisingly, given that the production technology is very simple in the model, the data show more variation in average wages than the model. As discussed before, the data construction results in a minority of local labor markets where the mean wage is lower than the wage floor. Appendix Figure 15 shows that the model does a good job at replicating the correlation between the share of firms bunching at the wage floor and the mean wage.

The second validation exercise is to compare the distribution of wages in the data and the model. Importantly, mean wages in the data were not used in the estimation of the model parameters. Additionally, mean wages are weakly correlated with the share of firms bunching at the wage floor (as shown in Appendix Figure 15), which is the key input used to estimate the minimum productivities. Panel (b) of Figure 8 shows that the observations cluster around the 45-degree line, indicating that the model does a good job at replicating the mean wage in each local labor market. Once again, mean wages in the data are more volatile than in the model.

7.4 Replicating the effects of CB shocks

To construct shocks in the model I rely on the estimated effects of shocks on local labor market revenue. To do so, I first estimate the effect of trade shocks on aggregate revenue using data from the survey of businesses. Appendix Table 12 shows that aggregate revenue increases by about 22% in exporting local labor markets, and this conclusion is similar when excluding the largest CB units in the data. There is also a non-significant effect on non-exporting local labor markets of about 6%. Then, I take the estimated model in 2011–2012 and simulate shocks to minimum productivity that would result in the same changes in aggregate revenue as suggested by the survey data. I re-compute the model equilibrium using the new minimum productivities, and estimate the effect of the shocks on wages using a shift-share DiD strategy.

The model is able to replicate the effects of CB shocks on wages observed in the data. Table 5 shows the estimates using the aggregate data and the model-generated data. CB shocks affect wages in the aggregate data, just like in the firm-level estimates, although the results are noisier. The model is able to replicate the effects of CB shocks on wages. However, the magnitudes of the effects are smaller in the model, which understates the degree of spillovers. The main reason is that the model does not account for spillover effects of the wage floor on unconstrained firms, which the heterogeneity analysis of Table 2 suggests can be quite important. There is no reason to think that this bias would be different for counterfactual CB networks, suggesting that the model is useful to explore how shocks propagate under different bargaining networks even if the magnitudes are not exactly correct.

Overall, these results suggest that the estimated model does a good job at capturing the main features of the data. In the next section I use the model to explore the propagation of shocks under different bargaining networks.

8 Model-Based Counterfactuals

In this section, I use the estimated model to explore how shocks propagate through the labor market under different CB networks. First, I illustrate how shocks propagate through the labor market using a simple simulated economy to visualize the mechanisms at play. Second, I define the counterfactual networks used in the Argentine data. Third, I explore how the different networks affect the propagation of shocks. Finally, I discuss the implications of the findings and their relationship to the literature.

8.1 Illustration of model mechanisms

I simulate a simple economy to illustrate how shocks propagate in the model. I define 3 local labor market types (L, M, H), 3 regions (L, C, R), and 15 economic sectors within each region and type cell. The elasticities η, ζ, and α are the same as in the main estimates, and the bargaining power parameters are defined to $\beta_c = 0.35$ for all CB units. I allow the baseline minimum productivity to differ by local market type and region, and I shock a selection of local markets in regions L and C.[30] Then, I observe how the shock propagates to wages in different CB networks.

The main takeaway from this exercise is that the response of CB units to shocks depends on the distribution of productivity of covered local markets, suggesting that productivity dispersion within CB units may be an important determinant of the response of wages to shocks. Figure 9 shows the results, focusing only on 4 local labor markets in each type by region cell. Panels (a) and (b) show baseline productivities and the changes in productivity.[31] Panel (c) shows how shocks affect wage floors for a network in which CB units connect local labor markets with the same sector and type across regions. Consider the CB units labelled H0, M0, and L2. While they received a similar shock, they differ in the dispersion of productivity. Wage floors do not change in H0, which has the highest dispersion, and they change relatively more in L2, which has the lowest dispersion. Wage floor responses in M0 are in between, as is the dispersion of productivity. Panel (e) shows the response of average wages to shocks.

[30] A few more parameters need to be defined for the simulation. I assume $N_r = 100$ and $b_r = 1$ for all regions, and $A_{k1} = 1$ for all local labor markets. I assume $\varphi_{g0} = 1$ for local markets of type L in region L and all local markets in region C, $\varphi_{g0} = 1.025$ for those of type M in region L and those of type R, and $\varphi_{g0} = 1.05$ for those of type H in region L. I let $M_g = 3$ for local markets of type L, $M_g = 2$ for type M, and $M_g = 1$ for type H.

[31] Local labor markets omitted from the figure are not shocked directly.

Figure 9 illustrates a couple more interesting features of the model. First, the exercise illustrates that CB units that cross regions will result in cross-regional propagation of shocks. Comparing across networks in Panels (c) and (d), we note that the cross-type network does not lead to regional contagion of shocks. Second, the response of wage floors will depend on whether the low- or high- productivity local labor markets are shocked. Consider the second network, depicted in Panels (d) and (f), and the CB units labelled L1 and L2, both of which have two local markets shocked. In L1 the high-productivity local market is shocked, while in L2 the low-productivity one is shocked. Wage floors respond more strongly in L2 relative to L1. Finally, there may be wage floor responses in CB units that are not directly affected by the shocks. For instance, Panel (d) shows that CB unit L3 increases its wage floor even though none of its local labor markets are directly affected by the shock.

8.2 Counterfactual CB networks

I manipulate the CB network observed in the data to create counterfactual networks. The first counterfactual, which I call Baseline Simple network, is defined by assigning some local labor markets with a local CB unit to the most common non-local CB unit in the region by economic sector cell.[32] I define a Baseline Region network by splitting the CB units in the Baseline Simple network by region, and a 2-d Sector and 3-d Sector networks by splitting it by 2-digit and 3-digit economic sector, respectively. Finally, I define a Local network, which consists of a separate CB unit for each local labor market, and a non-bargaining economy, for which I simply compute the market clearing regional wage indexes with non-binding wage floors. In order to hold the average economy-wide level of bargaining power constant, I assign bargaining power parameters to the new CB units by averaging the parameters of the baseline CB network across local labor markets, weighting by the measure of local labor market size M_g.[33]

These networks differ in the degree of centralization of bargaining or, in terms of the classification in Bhuller et al. (2022), of "vertical coordination" of bargaining. Examples of countries with mostly sectoral bargaining are Italy, Portugal, or France, while examples of countries with mostly local bargaining are the US, the UK, or Japan. I classify the counterfactual networks based on the average number of local labor markets per CB unit. The Baseline and Baseline Region networks have an average of about 4 local markets per CB unit. The Baseline Simple network is more centralized, with an average of about 6. The 2-d and 3-d Sector networks are more decentralized, with an average of about 2 to 3. Finally, the Local network is the most decentralized.

These networks also differ in the degree of productivity dispersion within CB units. To capture this dimension I compute, for each network, the average log ratio of the 90th to the 10th percentile of minimum productivity across local labor markets within CB units. The ratios oscillate between

[32] Specifically, I replace local CB units in a region by province by coarse 4-digit sector cell if the most common CB unit in the cell is non-local.

[33] A second consideration is the regional differences in wage floors within CB units. I preserve these differences in the new CB units and calculate them similarly to the bargaining power parameters to ensure their average remains constant. As stated before, these regional differences do not respond to shocks.

about 0.02 and 0.06. By definition, the local network gets a value of 0, as each local labor market is its own CB unit and there is no dispersion.

8.3 Shock propagation under counterfactual CB networks

I use the counterfactual CB networks and the simulated shocks to compute the equilibrium of the model pre- and post-shock. As discussed in Section 7.4, the simulated shocks consist of a change in revenue that mimics the effect of the export shocks on aggregate revenue. Then, for each network I compute the wage floors and mean wages pre- and post-shock, and their change (in logs) as a result of the simulated shocks. I also simulate the effect of shocks in a non-bargaining economy, where wage floors are not binding in any local labor market.

I compare wage changes as response to shocks to the non-bargaining economy to summarize the extent of shock propagation in a given CB network. Specifically, I compute the share of local labor markets for which the absolute difference in the change in wages in a given network and the non-bargaining economy is more than 0.25%. If this share is large, then the economy is resulting in a large degree of shock propagation, so wage changes are mostly determined by factors outside the local labor market. This means that the effects of shocks are being shared across connected local labor markets instead of being fully absorbed by the directly shocked local labor markets.

The counterfactual estimates reveal a hump-shaped relationship between the degree of shock propagation and the degree of centralization of bargaining. This result is shown in Panel (a) of Figure 10. Both highly decentralized and highly centralized bargaining result in low degree of shock propagation. On the other hand, bargaining with medium levels of centralization results in relatively higher levels of shock propagation. Overall, there is significant variability in the degree of shock propagation across networks.

This pattern can be explained by extent to which wage floors are binding in the economy and the extent of the connections generated by CB. Panel (b) of Figure 10 shows that more centralized networks result in a lower share of firms bunching at the wage floor on average. In decentralized networks the bite of the wage floor is high, but since there are not many connections across local labor markets, the shock does not propagate. As the network becomes more centralized there are more connections across local markets, leading to more propagation. However, the wage floor has less bite, so shocks propagates less. At some point the bite becomes too low, and even though the network connects many local markets, the degree of shock propagation declines.

The key driver of wage floor bite is the average dispersion of productivity within CB units. Panel (c) of Figure 10 shows that the bite of the wage floor is decreasing in the dispersion of productivity within CB units. More centralized bargaining networks tend to result in CB units that are very heterogeneous. In order to avoid negative employment consequences in low-productivity local markets, unions in these networks set wage floors that are not as high. As a result, the bite of the wage floor is lower in these networks, and changes in wage floors following shocks are less impactful on wages.

The degree of bargaining centralization has consequences for the responses of employment to shocks as well. Appendix Figure 16 shows a similar hump-shaped relationship between the degree to which employment responses to shocks differ from the non-bargaining economy and bargaining centralization. Because the share bunching is far from the maximum, employment responses to the wage floor are generally positive. As a result, the model implies that CB networks that lead to more propagation of shocks also relocate employment towards the positively shocked CB units.

8.4 Discussion

This paper shows that the degree of propagation of shocks across CB units is hump-shaped in the degree of centralization of bargaining. This suggests that labor markets in countries with moderate degrees of centralization may be more resilient to shocks, as their effects are shared across firms connected by the CB network. The key mechanism at play is the endogenous level of wage floors across different networks: moderate levels of centralization result in wage floors that have significant bite and also connect many local labor markets.

This paper contributes to a recent literature studying the effects of bargaining centralization on the labor market adjustment to shocks. Barth et al. (2023) find that the effect of import competition from China on employment was stronger in European countries with uncoordinated wage bargaining. The authors point out that wage coordination may work as an insurance device against the risk of shocks. My model does not allow for cross-union wage coordination, but it does incorporate a risk-sharing mechanism as local shocks are not fully transmitted to wages. Countries with coordinated bargaining also tend to have more centralized bargaining, so it is possible that the mechanisms proposed in this paper explain part of Barth et al.'s (2023) result.[34] Ronchi and di Mauro (2017) find that firms in countries with more centralized bargaining faced stronger wage rigidity during the Great Recession, resulting in employment losses. This finding seems to contradict Barth et al. (2023), who find better employment performance in countries with more centralized bargaining. Studying the effects of these shocks in a unified framework, such as the one proposed in this paper, may help reconcile these results.

This paper also contributes to literature studying the effects of bargaining centralization on unemployment and the spatial allocation of labor. Calmfors and Driffill (1988) argue that extreme degrees of centralization of bargaining perform best in terms of unemployment rates.[35] The evidence for this theoretical prediction has been mixed (see Moene et al. 1993 for a review). Boeri et al. (2021) compare wage-setting institutions in Italy and Germany, and find that the more centralized Italian system results in higher unemployment in low-productivity regions. In light of these findings, the results in this paper highlight the potential beneficial effects of centralization in dissipating the effects of shocks across regions, leading to risk sharing among firms connected

[34]The authors use a measure of wage coordination from Visser (2019). This dataset also provides a measure of the degree of centralization of bargaining that is positively correlated with wage coordination.

[35]Their key insight is that medium degrees of decentralization will generate negative externalities in other sectors that are ignored by unions, whereas these negative effects will be internalized by the competitive forces of decentralized bargaining or by the bargains of a monopolistic union.

by the CB network. However, if the level of centralization is too high the extent of risk sharing will be lower, as the bite of wage floors is endogenously lower.

9 Conclusions

This paper studies the role of collective bargaining (CB) in mediating the effects of shocks. I find that firms respond to product-demand shocks to other firms in the same CB unit, indicating that shocks propagate through CB. The results are consistent with collective rent-sharing: when average economic conditions among covered firms improve, unions negotiate higher wage floors that affect wages and employment among all firms in the CB unit. Additionally, I find that wages respond more strongly to CB shocks than to comparable firm shocks. Although this result is robust across several specifications, it may not readily generalize to other settings. In particular, the strength of unions in Argentina and the frequent nominal wage adjustments during the study period could be factors that contribute to the strong response of wages to CB shocks.

The empirical findings suggest that CB may play an important role in mediating the effects of shocks. However, they are uninformative about the effects of shocks under alternative CB networks. To study this question, I develop and estimate a structural model of the labor market with CB. The model illustrates how the CB network can affect the propagation of shocks, leading to a sharing of risk across firms. Interestingly, I find a hump-shaped relationship between the degree of shock propagation and the degree of centralization of bargaining, suggesting that wage responsiveness to local shocks is lower with intermediate levels of centralization. This can be explained by the extent to which firms are constrained by the wage floor across CB networks.

This article opens up several avenues for future research. First, the article does not speak to the value for workers of decreasing the risk of shocks. This might be important to quantify the welfare contribution of the insurance provided by sectoral CB. Second, the analysis abstracts from several features of labor markets that may be important to determine the role of CB in wage inequality, such as worker heterogeneity. Third, the analysis ignores potential effects of CB on informality, which may be particularly important in developing countries. Addressing these questions can provide a more comprehensive understanding of the impacts of different CB structures on the labor market.

References

Abowd, J. A. and Lemieux, T. (1993). The effects of product market competition on collective bargaining agreements: The case of foreign competition in Canada. *Quarterly Journal of Economics*, 108(4):983–1014.

Adamopoulou, E., Díez-Catalán, L., and Villanueva, E. (2022). Staggered contracts and unemployment during recessions. IZA Discussion Papers 15801, IZA Institute of Labor Economics.

Adão, R., Arkolakis, C., and Esposito, F. (2022). General equilibrium effects in space: Theory and measurement. *Working Paper*.

Ahlfeldt, G. M., Roth, D., and Seidel, T. (2022). Optimal minimum wages. Technical Report 16913, Centre for Economic Policy Research.

Autor, D., Dorn, D., and Hanson, G. (2013). The China syndrome: Local labor market effects of import competition in the United States. *American Economic Review*, 103(6):2121–68.

Barth, E., Finseraas, H., Kjelsrud, A., and Moene, K. (2023). Hit by the silk road: How wage coordination in Europe mitigates the China shock. *Scandinavian Journal of Economics*, 125(1):32–72.

Berger, D. W., Herkenhoff, K. F., and Mongey, S. (2022). Minimum wages, efficiency and welfare. Working Paper 29662, National Bureau of Economic Research.

Bhuller, M., Moene, K. O., Mogstad, M., and Vestad, O. L. (2022). Facts and fantasies about wage setting and collective bargaining. *Journal of Economic Perspectives*, 36(4):29–52.

Boeri, T., Ichino, A., Moretti, E., and Posch, J. (2021). Wage equalization and regional misallocation: evidence from Italian and German provinces. *Journal of the European Economic Association*, 19(6):3249–3292.

Borusyak, K., Dix-Carneiro, R., and Kovak, B. (2022a). Understanding migration responses to local shocks. *Available at SSRN 4086847*.

Borusyak, K., Hull, P., and Jaravel, X. (2022b). Quasi-experimental shift-share research designs. *Review of Economic Studies*, 89(1):181–213.

Calmfors, L. and Driffill, J. (1988). Bargaining structure, corporatism and macroeconomic performance. *Economic policy*, 3(6):13–61.

Card, D. (1990). Unexpected inflation, real wages, and employment determination in union contracts. *American Economic Review*, 80(4):669–688.

Card, D. (1996). The effect of unions on the structure of wages: A longitudinal analysis. *Econometrica*, pages 957–979.

Card, D. and Cardoso, A. R. (2022). Wage flexibility under sectoral bargaining. *Journal of the European Economic Association*, 20(5):2013–2061.

Card, D., Cardoso, A. R., Heining, J., and Kline, P. (2018). Firms and labor market inequality: Evidence and some theory. *Journal of Labor Economics*, 36(S1):S13–S70.

Cardoso, A. R. and Portugal, P. (2005). Contractual wages and the wage cushion under different bargaining settings. *Journal of Labor Economics*, 23(4):875–902.

Cengiz, D., Dube, A., Lindner, A., and Zipperer, B. (2019). The effect of minimum wages on low-wage jobs. *Quarterly Journal of Economics*, 134(3):1405–1454.

Chetty, R., Guren, A., Manoli, D., and Weber, A. (2011). Are micro and macro labor supply elasticities consistent? A review of evidence on the intensive and extensive margins. *American Economic Review*, 101(3):471–475.

Collard-Wexler, A., Gowrisankaran, G., and Lee, R. S. (2019). "Nash-in-Nash" bargaining: A microfoundation for applied work. *Journal of Political Economy*, 127(1):163–195.

Corneo, G. (1995). National wage bargaining in an internationally integrated product market. *European Journal of Political Economy*, 11(3):503–520.

Datta, N. (2023). The measure of monopsony: The labour supply elasticity to the firm and its constituents. Discussion Paper 1930, Centre for Economic Performance, London School of Economics and Political Science.

Davidson, C. (1988). Multiunit bargaining in oligopolistic industries. *Journal of Labor Economics*, 6(3):397–422.

Dix-Carneiro, R. and Kovak, B. K. (2017). Trade liberalization and regional dynamics. *American Economic Review*, 107(10):2908–46.

Felix, M. (2022). Trade, labor market concentration, and wages. *Unpublished Manuscript*.

Freeman, R. B. and Medoff, J. L. (1984). *What do unions do?* Basic Books.

Freyaldenhoven, S., Hansen, C., Pérez, J. P., and Shapiro, J. M. (forthcoming). Visualization, identification, and estimation in the linear panel event-study design. *Advances in Economics and Econometrics - Twelfth World Congress*.

Galichon, A. (2022). 'math+econ+code' masterclass on equilibrium transport and matching models in economics.

García-Lembergman, E. (2022). Multi-establishment firms, pricing and the propagation of local shocks: Evidence from us retail. *Unpublished manuscript*.

Garin, A. and Silvério, F. (2023). How Responsive Are Wages to Firm-Specific Changes in Labour Demand? Evidence from Idiosyncratic Export Demand Shocks. *Review of Economic Studies*, page rdad069.

Gaulier, G. and Zignago, S. (2010). BACI: International trade database at the product-level. the 1994-2007 version. Working Papers 2010-23, CEPII.

Giroud, X., Lenzu, S., Maingi, Q., and Mueller, H. (2021). Propagation and amplification of local productivity spillovers. Technical Report 29084, National Bureau of Economic Research.

Giroud, X. and Mueller, H. M. (2019). Firms' internal networks and local economic shocks. *American Economic Review*, 109(10):3617–49.

Gürtzgen, N. (2009a). Firm heterogeneity and wages under different bargaining regimes: does a centralised union care for low-productivity firms? *Jahrbücher für Nationalökonomie und Statistik*, 229(2-3):239–253.

Gürtzgen, N. (2009b). Rent-sharing and collective bargaining coverage: Evidence from linked employer–employee data. *Scandinavian Journal of Economics*, 111(2):323–349.

Holden, S. (1988). Local and central wage bargaining. *Scandinavian Journal of Economics*, 90(1):93–99.

Horn, H. and Wolinsky, A. (1988). Bilateral monopolies and incentives for merger. *The RAND Journal of Economics*, 19(3):408–419.

Hummels, D., Jørgensen, R., Munch, J., and Xiang, C. (2014). The wage effects of offshoring: Evidence from danish matched worker-firm data. *American Economic Review*, 104(6):1597–1629.

Jäger, S., Schoefer, B., Young, S., and Zweimüller, J. (2020). Wages and the value of nonemployment. *Quarterly Journal of Economics*, 135(4):1905–1963.

Katz, L. F. and Autor, D. (1999). Changes in the wage structure and earnings inequality. In Ashenfelter, O. and Card, D., editors, *Handbook of Labor Economics*, volume 3, pages 1463–1555. Elsevier.

Liao, S., Kim, I. S., Miyano, S., and Zhu, F. (2020). concordance: Product concordance. R package version 2.0.0.

Manning, A. (2011). Imperfect competition in the labor market. In Ashenfelter, O. and Card, D., editors, *Handbook of Labor Economics*, volume 4, pages 973–1041. Elsevier.

Milgrom, M. and Verdugo, D. (2022). The drivers of minimum wage spillovers: Evidence from Chile. *Unpublished manuscript*.

Ministerio de Trabajo, Empleo y Seguridad Social (2022a). Simplificación Registral. Data processed by the Dirección de Estudios y Estadísticas Laborales. Last accessed: June 2023.

Ministerio de Trabajo, Empleo y Seguridad Social (2022b). Sistema Integrado Previsional Argentino. Data processed by the Dirección de Estudios y Estadísticas Laborales. Last accessed: June 2023.

Ministerio de Trabajo, Empleo y Seguridad Social (2023). Cobertura e incidencia de los convenios colectivos de trabajo. Data produced by the Dirección de Estudios y Estadísticas de Relaciones de Trabajo. Date accessed: June 2023.

Moene, K. O., Wallerstein, M., and Hoel, M. (1993). Bargaining structure and economic performance. In Flanagan, R., Moene, K. O., and Wallerstein, M., editors, *Trade Union Behavior, Pay Bargaining and Economic Performance*, pages 63–132. Oxford University Press, Oxford.

Monte, F., Redding, S. J., and Rossi-Hansberg, E. (2018). Commuting, migration, and local employment elasticities. *American Economic Review*, 108(12):3855–3890.

Naylor, R. (1998). International trade and economic integration when labour markets are generally unionised. *European Economic Review*, 42(7):1251–1267.

Palomino, H. and Trajtemberg, D. (2006). Una nueva dinámica de las relaciones laborales y la negociación colectiva en la argentina. *Revista de trabajo*, 2(3).

Parente, R. (2022). Minimum wages, inequality, and the informal sector. *Unpublished manuscript.*

Plasman, R., Rusinek, M., and Rycx, F. (2007). Wages and the bargaining regime under multi-level bargaining: Belgium, Denmark and Spain. *European Journal of Industrial Relations*, 13(2):161–180.

Pontoni, G. and Trajtemberg, D. (2017). Estructura, dinámica y vigencia de los convenios colectivos de trabajo sectoriales del ámbito privado (1975-2014). *Estudios del trabajo*, (54):5–26.

Ronchi, M. and di Mauro, F. (2017). Wage bargaining regimes and firms' adjustments to the Great Recession. Working paper no. 2051, European Central Bank.

Rose, N. L. (1987). Labor rent sharing and regulation: Evidence from the trucking industry. *Journal of Political Economy*, 95(6):1146–1178.

Rusinek, M. and Rycx, F. (2013). Rent-sharing under different bargaining regimes: Evidence from linked employer–employee data. *British Journal of Industrial Relations*, 51(1):28–58.

Topalova, P. (2010). Factor immobility and regional impacts of trade liberalization: Evidence on poverty from India. *American Economic Journal: Applied Economics*, 2(4):1–41.

Van Reenen, J. (1996). The creation and capture of economic rents: Wages and innovation in UK manufacturing plants. *Quarterly Journal of Economics*, 111:195–226.

Visser, J. (2019). ICTWSS database. version 6.1.

Washington Post (2022). Newsom signs law with sweeping protections for California fast-food workers. Accessed: 2023-10-11.

Figure 1: Performance of Argentine economy, 2004–2019

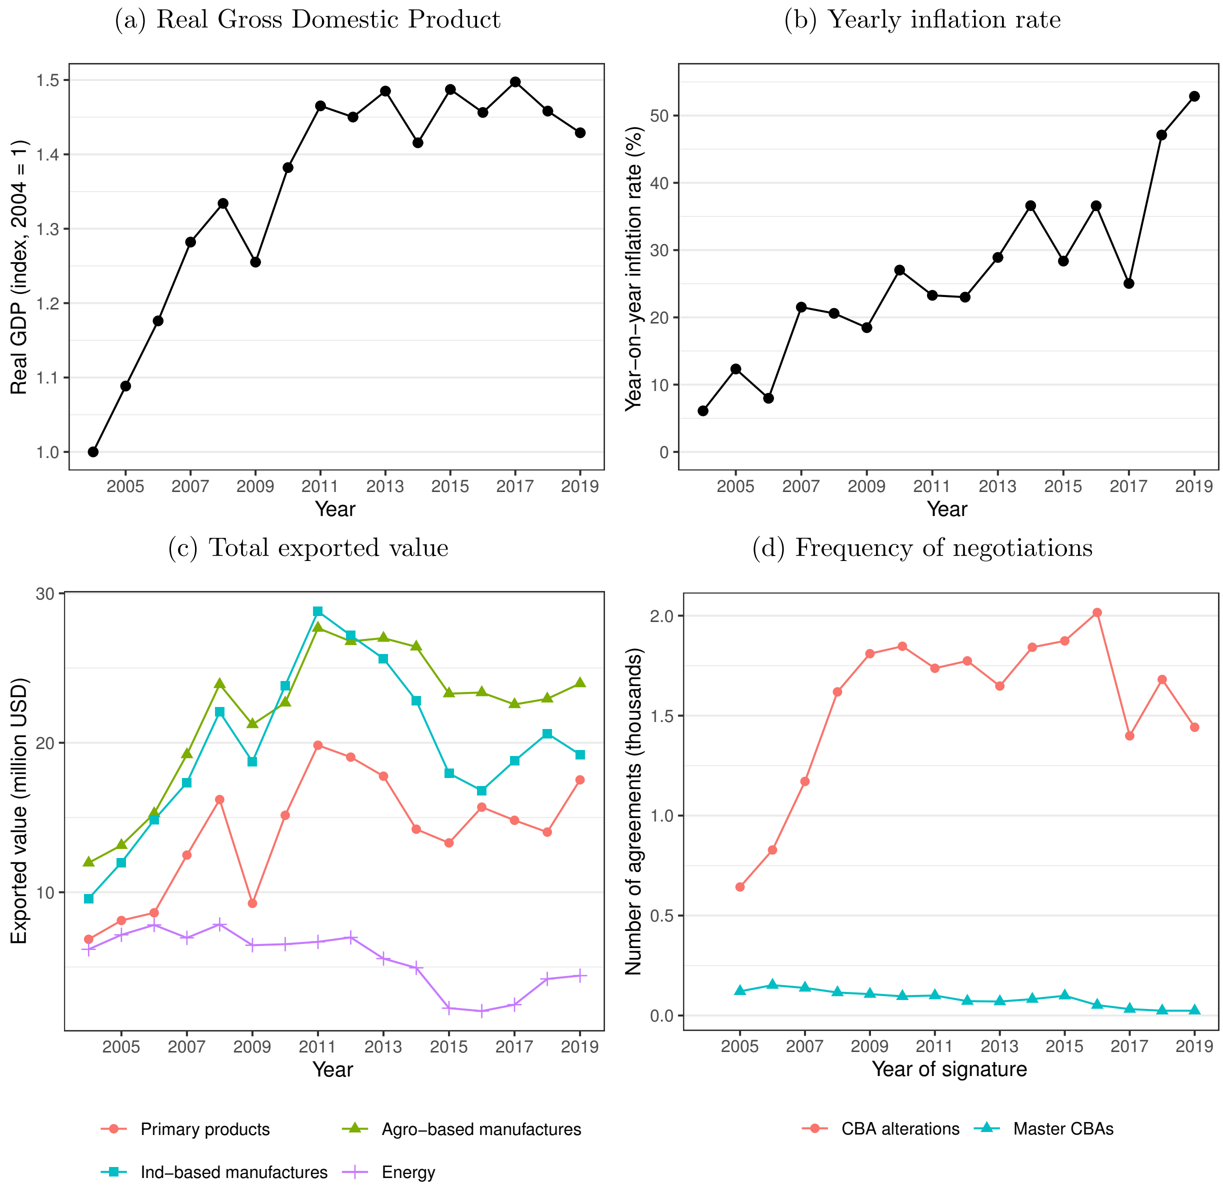

Notes: The figure shows the evolution of key macroeconomic variables in Argentina for the period 2004–2019. Panel (a) shows the evolution of the real gross domestic product (GDP), using data from the National Institute of Statistics and Censuses (INDEC). The real GDP is measured in constant 2004 Argentine pesos and normalized to 1 in 2004. Panel (b) shows the yearly inflation rate constructed from INDEC, and using data from regional statistics offices for the period 2007–2015. The inflation rate is measured as the yearly percentage change in the consumer price index as of December of each year. Panel (c) shows the total exported value, using data from INDEC. The total exported value is measured in millions of current US dollars. Panel (d) shows the number of collective bargaining agreements (CBAs) signed each year, using data obtained from the public archive of Collective Bargaining Agreements.

Figure 2: Concentration of firms in different CB units within 4-digit sector and province cells

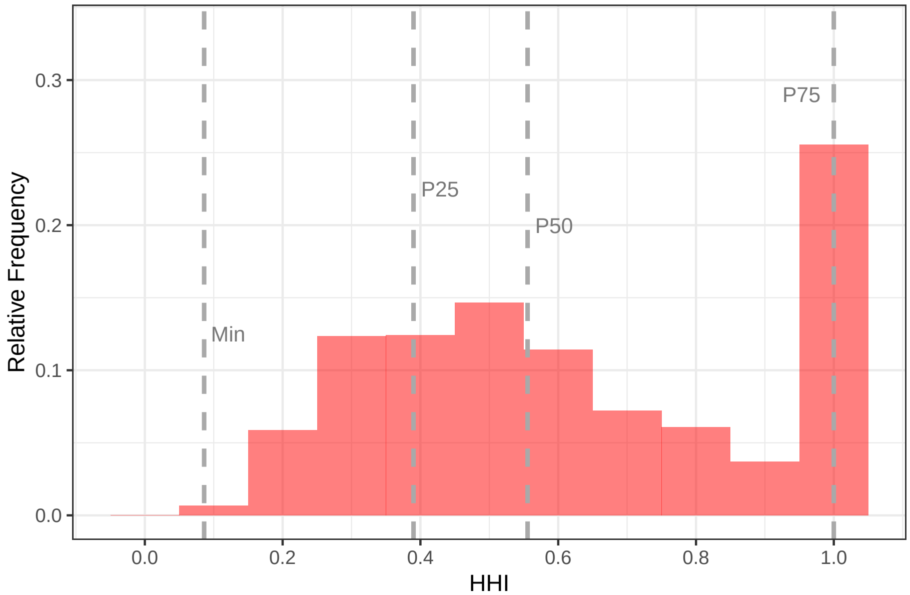

Notes: Data include all firms that had positive employment in 2012 and declared an economic sector in "Manufacturing" or "Wholesale and Retail Trade". The figure shows the distribution of the Herfindahl-Hirschman Index (HHI) of the number of firms across different collective bargaining units within 4-digit sector and province cells. The index is computed as $HHI_\ell = \sum_c h_{c\ell}^2$, where $h_{c\ell}$ is the share of firms in CB unit c within the given sector by province cell ℓ. In case of a firm with multiple CB units I assign it to the one with the highest number of employees.

Figure 3: Evolution of export shocks to firms and collective bargaining units

(a) Shocks to firms

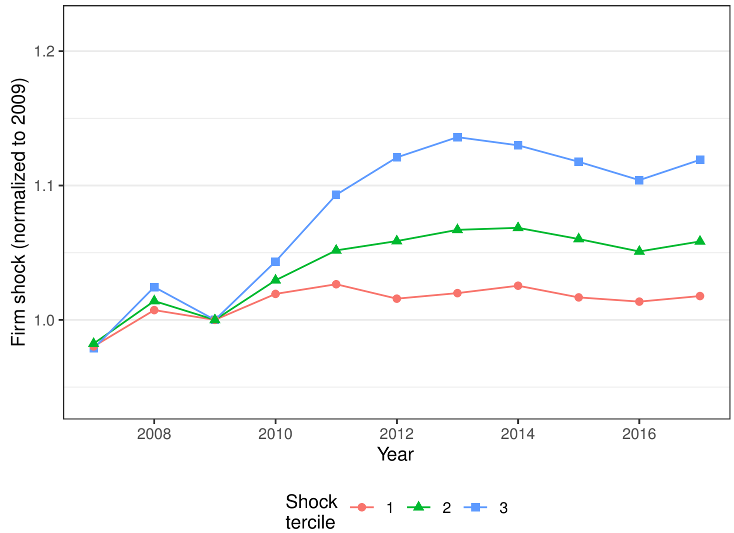

(b) Shocks to CB units

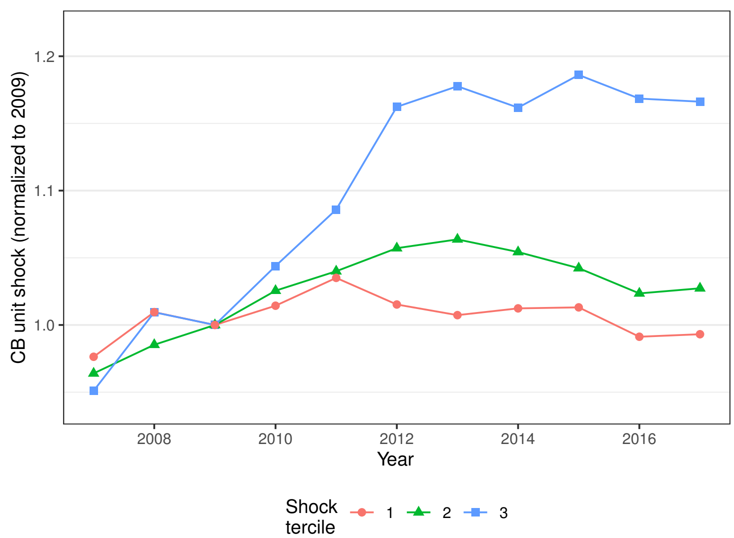

Notes: Data are from the baseline sample of exporting firms. The figure illustrates the evolution of time-varying firm and CB shocks, for different levels of the static firm and CB shocks. Each line depicts the average of the time-varying shock in a given tercile of the distribution of the static shock. The time-varying versions of the shocks are constructed as averages in world import demand in a given year, weighting by appropriate exposure shares, and then normalized by the 2009 value. The static shocks are constructed from the average change in world import demand for a given country-product between 2009–2010 and 2012–2013, also weighted by exposure shares. Firm exposure shares are equal to the share of a firm's exports in a given country-product. CB exposure shares are defined as the sum across firms of the employment share of a firm in the CB unit times the firm's value exported share. The averages in Panel (b) are weighted by the number of firms in a given CB unit.

Figure 4: Evolution of wages by level of CB unit and firm shock, baseline sample

(a) Shocks to firms

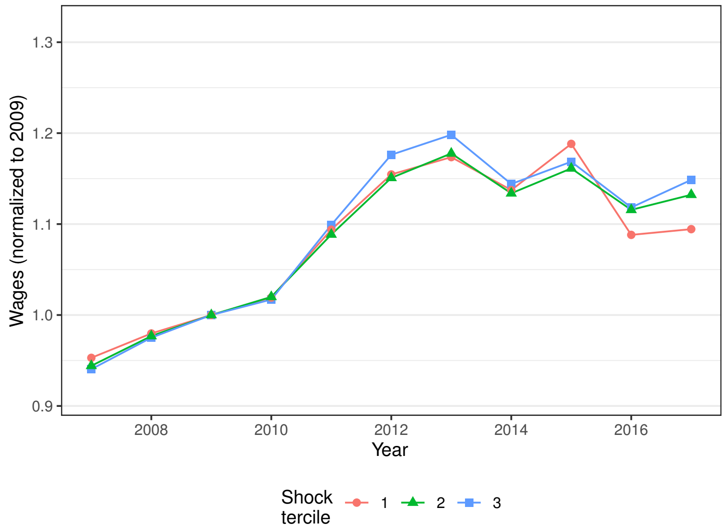

(b) Shocks to CB units

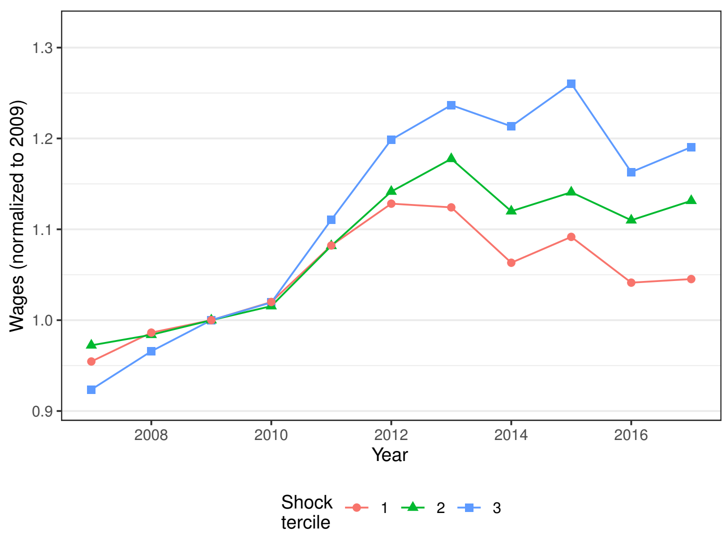

Notes: Data are from the baseline sample of exporting firms. The figure shows the average evolution of mean wages for firms in different terciles of the distribution of the firm shock (Panel a) and the CB shock (Panel b), relative to 2009. The firm and CB shocks are defined as the average changes in world import demand between 2009–2010 and 2012–2013, weighting by appropriate exposure shares.

Figure 5: Effect of export shocks on mean wages and mean wage floors

(a) Log mean wages

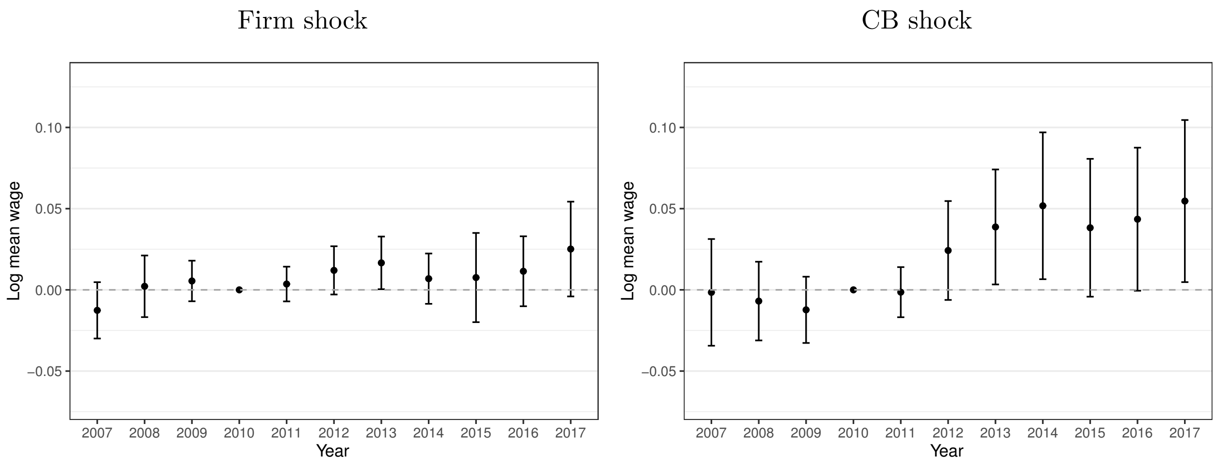

(b) Log mean wage floors

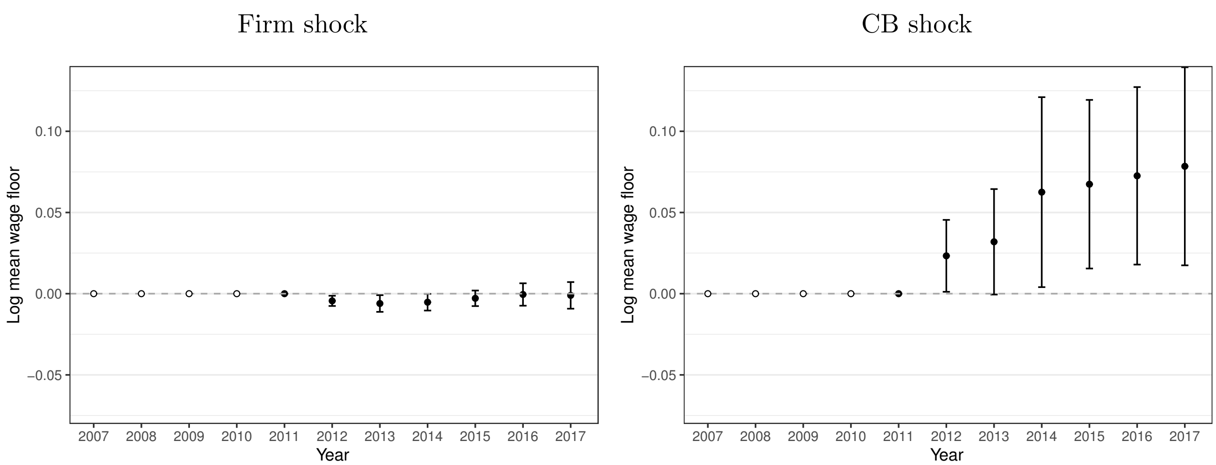

Notes: Data are from the baseline sample of exporting firms. The figure shows the dynamic effects of firm and CB shocks on log mean wages and log mean wage floors, interacting the shocks with year dummies. The regression omits the year 2010 for the wage variable, and the year 2011 for the wage floor variable. The regression includes firm fixed effects, 4-digit economic sector by province by year fixed effects, time-varying firm controls, and a similar CB shock for the pre-period interacted with year dummies. Firm controls consist of a firm size indicator (categories 1-9, 10-24, 25-99, and 100-500) interacted with 2-digit sector and year and the pre-period share of workers in the main hiring modality interacted with year. The firm and CB shocks are defined as the average changes in world import demand between 2009–2010 and 2012–2013, weighting by appropriate exposure shares. Standard errors are clustered at the CB unit level.

Figure 6: Effect of export shocks on employment

(a) Firm shock

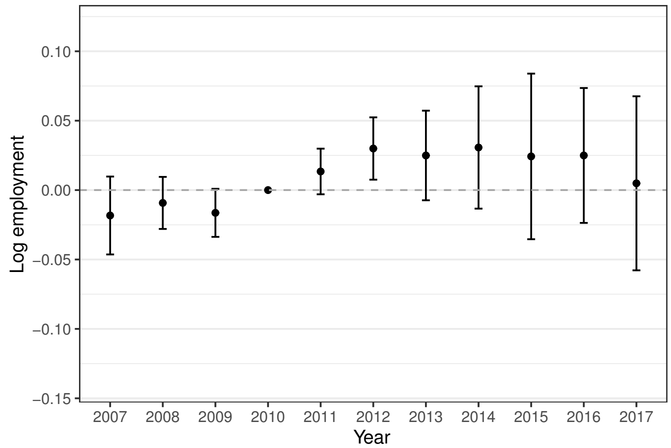

(b) CB shock

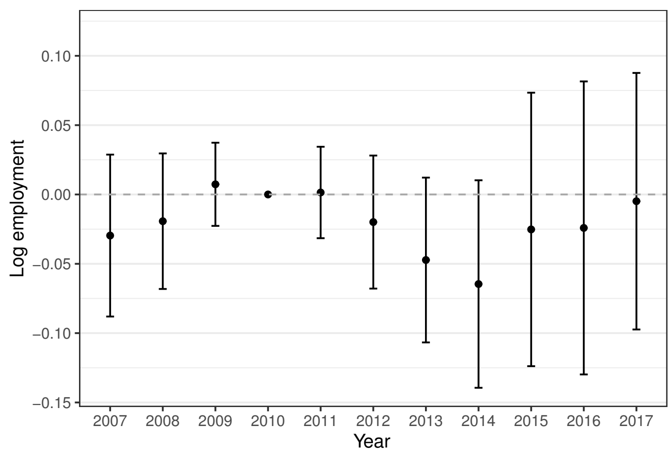

Notes: Data are from the baseline sample of exporting firms, including only firm-years for which wage floors are available. The figure shows the dynamic effects of firm and CB shocks on log employment, interacting the shocks with year dummies and omitting the year 2010. The regression includes firm fixed effects, 4-digit economic sector by province by year fixed effects, time-varying firm controls, and a similar CB shock for the pre-period interacted with year dummies. Firm controls consist of a firm size indicator (categories 1-9, 10-24, 25-99, and 100-500) interacted with 2-digit sector and year and the pre-period share of workers in the main hiring modality interacted with year. The firm and CB shocks are defined as the average changes in world import demand between 2009–2010 and 2012–2013, weighting by appropriate exposure shares. Standard errors are clustered at the CB unit level.

Figure 7: Effect of export shocks to CB units, worker-level estimates

(a) Shock to firm's primary CB

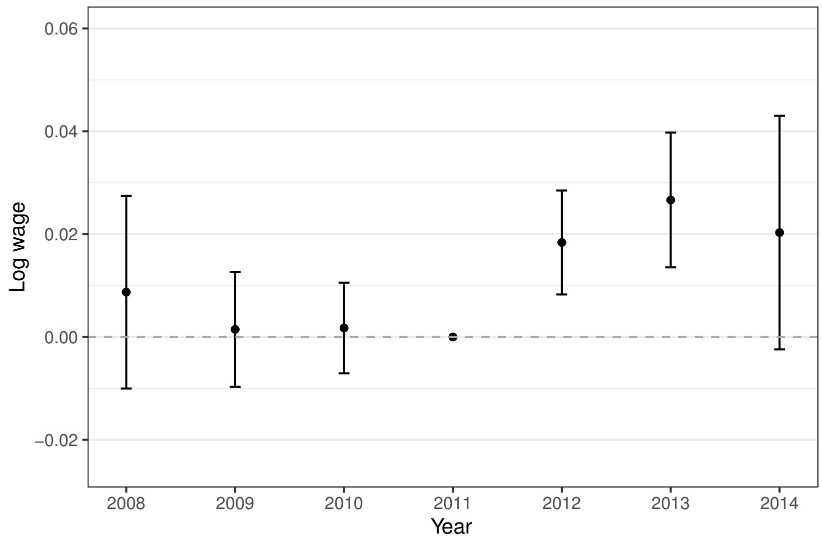

(b) Shock to worker's CB unit, controlling for firm-year effects

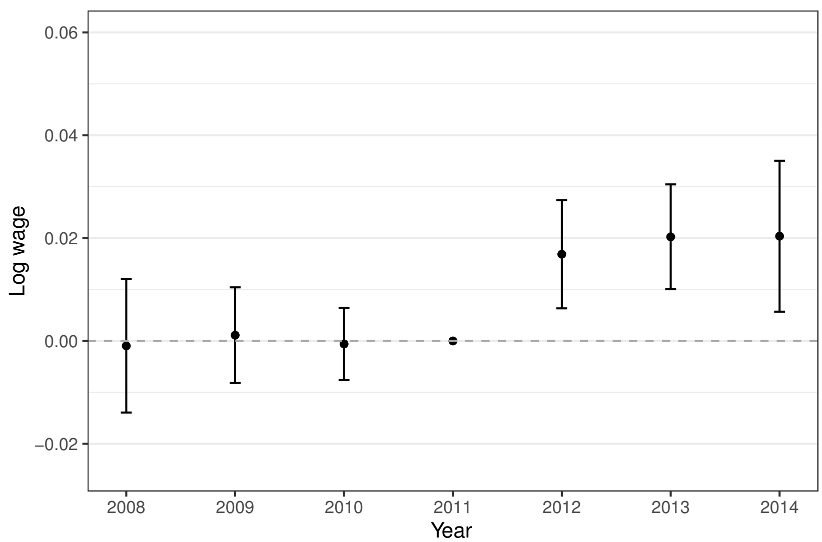

Notes: Data are from a panel of workers that worked in 2008, 2011, and 2014 in a firm in the baseline sample. The figures show estimates of the effect of CB shocks on mean monthly wage. Panel (a) estimates a difference-in-differences model using the primary CBA of the firm to define the treatment. It includes controls for the firm shock, worker by firm ("match") fixed effects, 6-digit economic sector by province by year fixed effects, hiring modality by year fixed effects, and 2-digit economic sector by an indicator for whether the worker's CB is the primary CB unit in the firm by year fixed effects. Panel (b) estimates a difference-in-differences model as well, but instead uses the CB shock that is specific to each worker and controls for worker fixed effects and firm by year fixed effects. Standard errors are clustered at the CB unit level.

Figure 8: Summary of model fit to the data

(a) Wage floor versus average wage

In the data

In the model

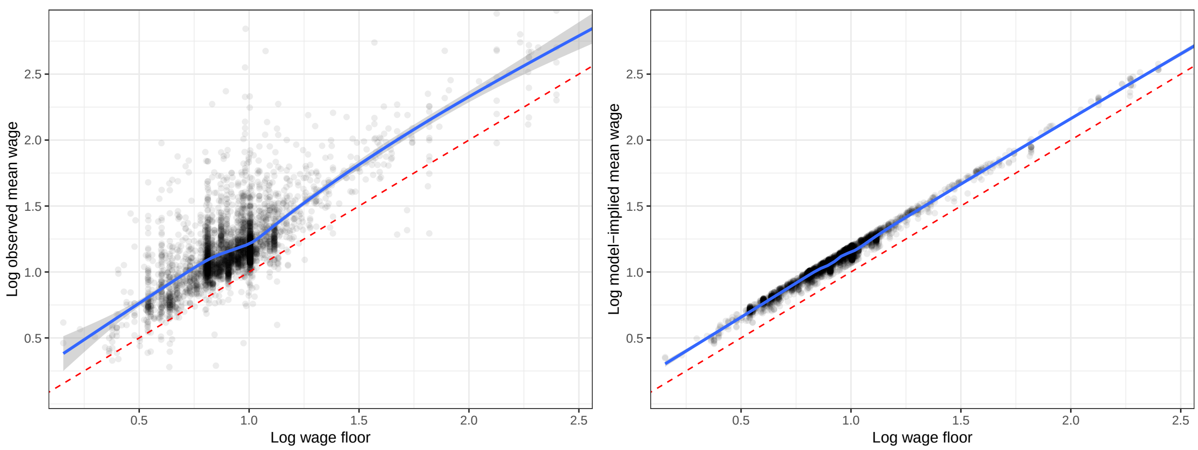

(b) Wages in the data versus in the model

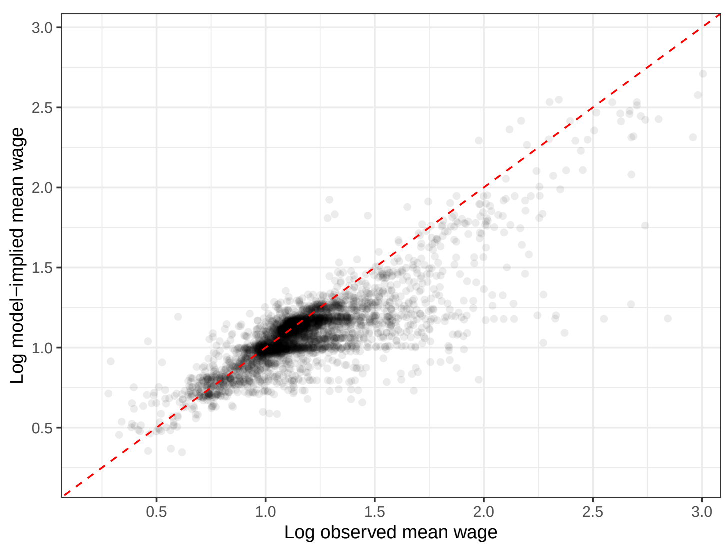

Notes: This figure illustrates the model fit to the local labor market data. Panel (a) shows the correlation between the log wage floor and the log mean wage. The left column shows the data, the right column shows the model. Panel (b) shows a scatter plot of observed mean wages versus model-implied mean wages in each local labor market. The measure of mean wages used in these figures is not adjusted by part-time employment.

Figure 9: Illustration of collective bargaining networks and shock propagation, simulated data

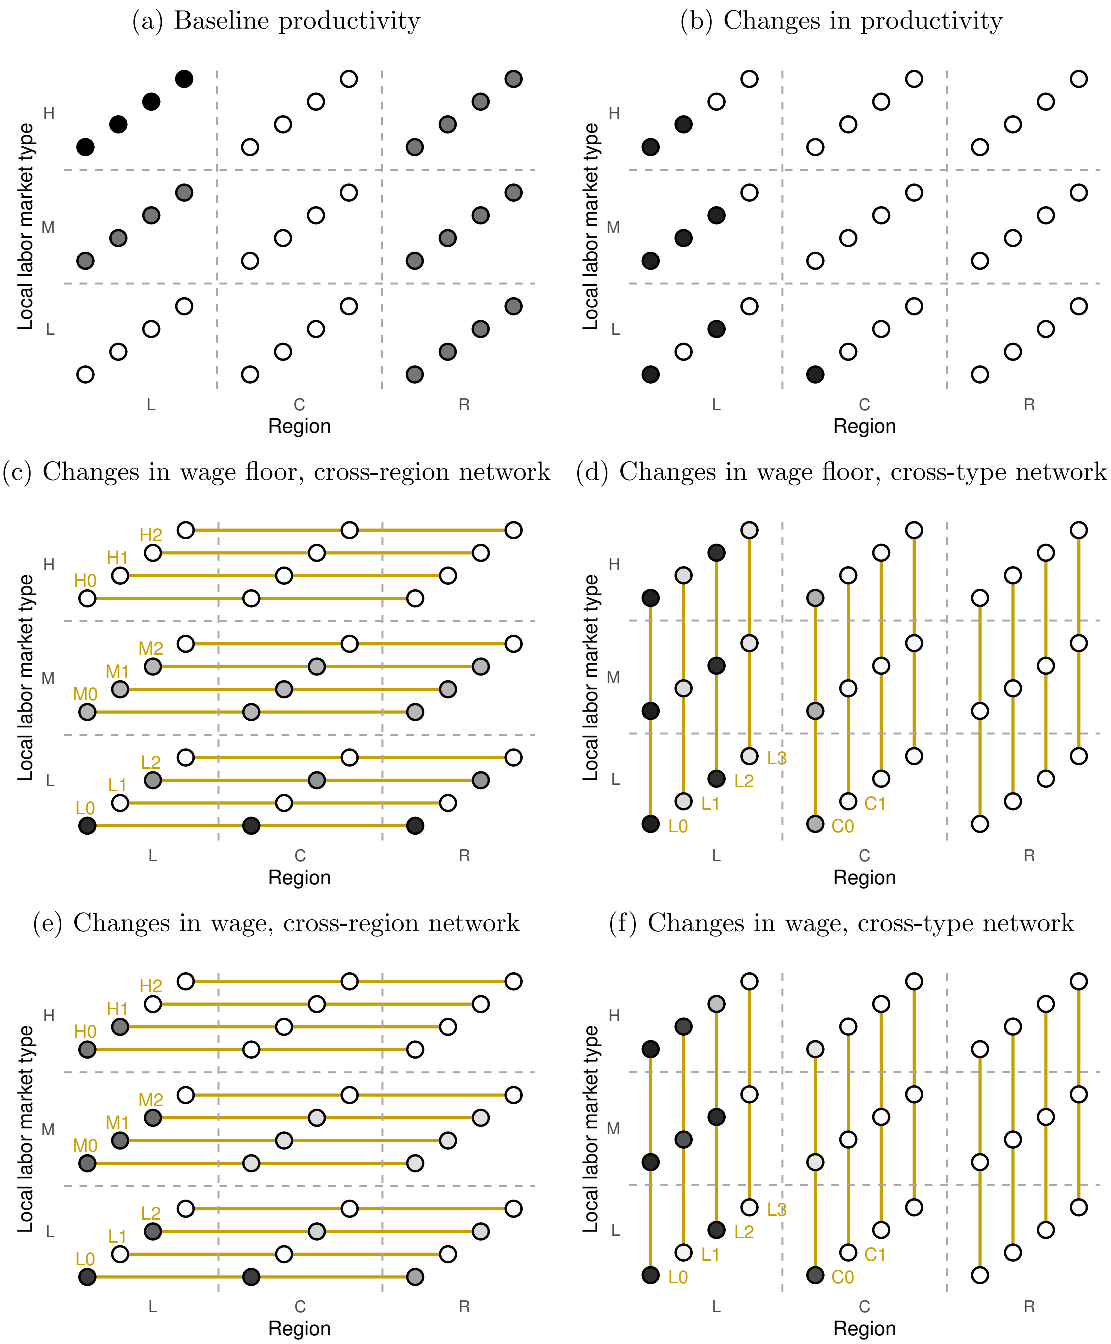

Notes: The figure illustrates the propagation of productivity shocks to wages across different collective bargaining (CB) networks using simulated data. Each dot represents a local labor market, and its color shows the level of a variable in the pre-shock equilibrium or the percent change between the pre-shock and post-shock equilibria. Color palettes are defined within each outcome, and darker colors indicate larger values. Line segments connect local labor markets in the CB unit. Panel (a) shows the baseline productivity level, and Panel (b) shows the change in productivity. Panel (c) and Panel (d) show the change in the wage floor in each local labor market for two different CB networks, respectively. Panel (e) and Panel (f) show the change in the average wage for the same two CB networks, respectively. The model is simulated with 15 local labor markets in each of the nine Productivity by Region cells, of which only 4 are shown. There are no productivity changes in the omitted local labor markets.

Figure 10: Centralization of bargaining and shock propagation across CB networks

(a) Centralization and shock propagation

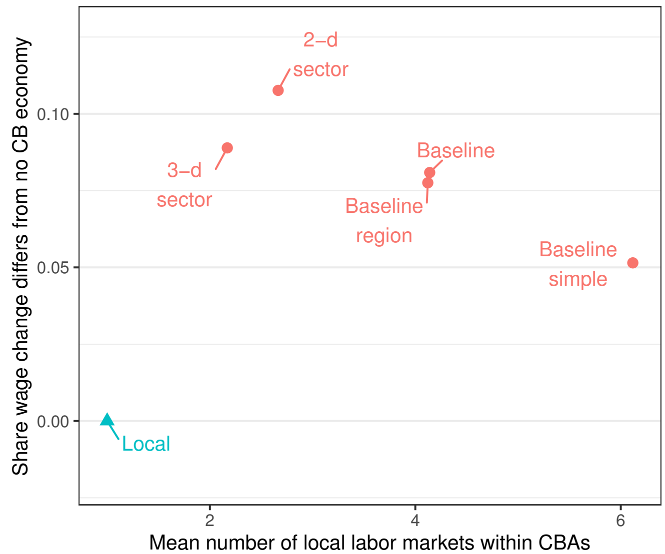

(b) Centralization and wage floor "bite"

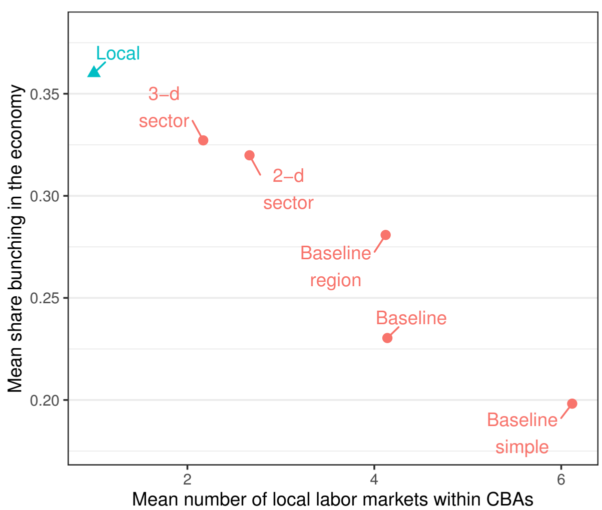

(c) Productivity dispersion and wage floor "bite"

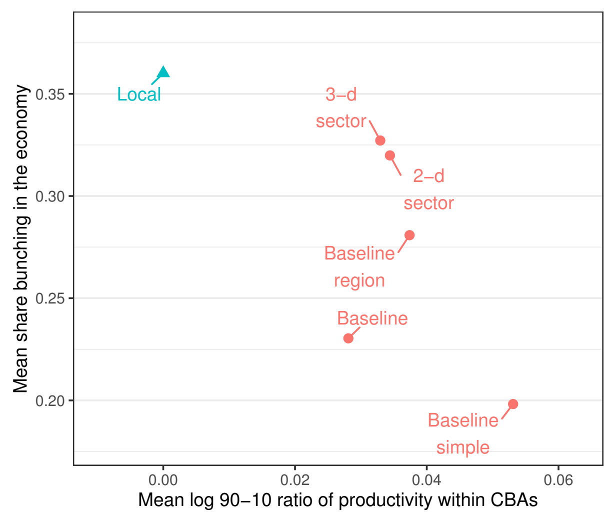

Notes: Data are from model simulations pre- and post-export shocks under different CB networks. Panel (a) shows the degree of shock propagation against the level of bargaining centralization. Panel (b) shows the average share of firms bunching (wage floor "bite") against bargaining centralization. Panel (c) shows the average share of firms bunching against the level of productivity dispersion. The degree of shock propagation is measured as the share of local labor markets with an absolute average wage change following the shock more than 0.25% different from the counterfactual wage change in an economy without CB. This computation excludes local labor markets that correspond to the retail CBA (0130/75) at baseline and to CBAs with less than 5% of employment in exporting firms. Bargaining centralization is measured as the average number of local labor markets per CBA. The productivity dispersion is measured as the average ratio of the 90th to the 10th percentile of the productivity distribution within CBAs. The average share of firms bunching is measured as the simple mean across local labor markets.

Table 1: Static difference-in-differences estimates

	Log mean wage	Log mean wage floor	Log wage cushion	Log employment	Sh. main modality	Firm exit
	(1)	(2)	(3)	(4)	(5)	(6)
CB shock × Post	0.0458	0.0541	-0.0001	-0.0233	0.0162	-0.0050
	(0.0202)	(0.0213)	(0.0171)	(0.0351)	(0.0118)	(0.0077)
Firm shock × Post	0.0134	-0.0036	0.0100	0.0297	0.0027	-0.0008
	(0.0070)	(0.0021)	(0.0079)	(0.0210)	(0.0071)	(0.0040)
Firm FE	Y	Y	Y	Y	Y	Y
Local market-year FE	Y	Y	Y	Y	Y	Y
Firm controls	Y	Y	Y	Y	Y	Y
Pre-period CB shock	Y	Y	Y	Y	Y	Y
Num. firms	7,972	7,654	7,654	7,972	7,972	7,972
Num. fixed effects	27,976	19,860	19,860	27,976	28,031	28,234
Num. observations	85,777	50,703	50,699	85,777	86,238	87,692
Adjusted R^2	0.8480	0.9253	0.8266	0.8965	0.5851	0.3212

Notes: Data are from the baseline sample of exporting firms. The table show regression coefficients on the firm and CB shocks variables interacted with an indicator for year greater than or equal to 2012. The firm and CB shocks are defined as the average changes in world import demand between 2009–2010 and 2012–2013, weighting by appropriate exposure shares. The regression includes firm fixed effects, 4-digit economic sector by province by year fixed effects, firm controls, and a similar CB shock for the pre-period interacted with year dummies. Firm controls consist of the following fixed effects interacted: firm size categories (categories 1-19, 20-124, 125-500) by an indicator for above median wages in 2007–2009 by a 4-digit economic sector by year. Standard errors are clustered at the CB unit level.

Table 2: Static difference-in-differences estimates, heterogeneity by pre-period mean wage

	Log mean wage	Log mean wage floor	Log wage cushion	Log employment	Sh. main modality
	(1)	(2)	(3)	(4)	(5)
CB shock × Pre wage 1	0.1509	0.0884	0.0624	-0.1857	-0.0134
	(0.0400)	(0.0215)	(0.0414)	(0.0944)	(0.0370)
CB shock × Pre wage 2	0.1433	0.1112	0.0018	0.0516	0.0616
	(0.0375)	(0.0260)	(0.0306)	(0.1037)	(0.0302)
CB shock × Pre wage 3	0.0871	0.0729	0.0068	0.0926	0.0659
	(0.0253)	(0.0272)	(0.0330)	(0.0591)	(0.0265)
CB shock × Pre wage 4	0.0483	0.1167	-0.0531	-0.0719	-0.0199
	(0.0352)	(0.0564)	(0.0644)	(0.0667)	(0.0448)
p-value joint significance	0.0000	0.0000	0.6086	0.0003	0.0049
p-value equality $1 = 2$	0.8414	0.4579	0.2777	0.0832	0.0421
p-value equality $1 = 3$	0.1704	0.5468	0.2408	0.0033	0.0271
p-value equality $1 = 4$	0.0849	0.6655	0.1752	0.3577	0.8736
Firm shock	Y	Y	Y	Y	Y
Local market-Pre wage-Size-year FE	Y	Y	Y	Y	Y
Pre-period CB shock	Y	Y	Y	Y	Y
Num. firms	7,145	6,937	6,937	7,145	7,145
Num. fixed effects	24,894	17,853	17,853	24,894	24,947
Num. observations	76,911	46,104	46,100	76,911	77,308
Adjusted R^2	0.8591	0.9246	0.8373	0.9000	0.5946

Notes: Data are from the baseline panels of exporting firms excluding CB unit by province cells with fewer than 6 firms. The figure shows regression coefficients on the CB shocks interacted with an indicator for year greater than or equal to 2012. The CB shock variable is further interacted with an indicator for whether the average firm wage in 2007–2009 is in the first ("Pre wage 1"), second ("Pre wage 2"), third ("Pre wage 3") or fourth ("Pre wage 4") quartile of the distribution of firm wages within the province. The first p-value tests the joint significance of the four coefficients presented in the table, and the remaining p-values test the null hypothesis that the coefficient on the first quartile is equal to the coefficient on the second, third, and fourth quartiles, respectively. The regression models controls for the firm shock, local labor market by year by pre wage fixed effects, and a similar CB shock for the pre-period interacted with year dummies. Standard errors, included for the hypothesis testing, are clustered at the CB unit level.

Table 3: Static difference-in-differences estimates, by size of CB unit

	Log mean wage		Log mean wage floor		Log employment	
	(1)	(2)	(3)	(4)	(5)	(6)
CB shock	0.0477		0.0385		0.0072	
	(0.0222)		(0.0241)		(0.0358)	
CB shock × Small CB		-0.0011		0.0023		-0.0266
		(0.0206)		(0.0256)		(0.0419)
CB shock × Large CB		0.1063		0.0911		0.0702
		(0.0265)		(0.0290)		(0.0521)
p-value equality		0.0021		0.0235		0.1671
Firm FE	Y	Y	Y	Y	Y	Y
Firm shock	Y	Y	Y	Y	Y	Y
Local market-CB size-year FE	Y	Y	Y	Y	Y	Y
Pre-period CB shock	Y	Y	Y	Y	Y	Y
Num. firms	7,013	7,013	6,741	6,741	7,013	7,013
Num. fixed effects	23,772	16,125	17,047	12,316	23,772	16,125
Num. observations	75,529	75,529	44,709	44,709	75,529	75,529
Adjusted R^2	0.8533	0.8519	0.9322	0.9351	0.8979	0.8979

Notes: Data are from the baseline sample of exporting firms, excluding 4-digit sector by province cells with less one CB unit. The table show regression coefficients on the CB shocks variable interacted with an indicator for year greater than or equal to 2012. Columns (1), (3), and (5) replicate the baseline results. Columns (2), (4), and (6) further interact the CB shock variable with an indicator for CB unit size. A "small" CB unit is defined as one with fewer than the median number of firms in the 4-digit sector by province cell. A "large" CB unit is defined as one with more than the median number of firms in the 4-digit sector by province cell. The regression controls for the firm shock, and includes firm fixed effects, CB size by 4-digit economic sector by province by year fixed effects, and a similar CB shock for the pre-period interacted with year dummies. Standard errors are clustered at the CB unit level.

Table 4: Estimation strategy for different model parameters

Parameter	Description	Source
ζ	Elasticity of formal employment	Literature
η	Elasticity of labor supply to the firm	OLS estimation
$\{A_{k1}\}$	Amenity values of 1-digit sectors	OLS estimation
α	Curvature of productivity processes	Calibration
$\{\varphi_{g0}\}$	Minimum productivity of Pareto distributions	Model inversion
$\{b_r\}$	Outside option of workers in each r	Model inversion
$\{\beta_c\}$	Bargaining power parameters	Model inversion

Notes: The table shows the parameters of the structure model and the approach used to estimate them.

Table 5: Effects of CB shocks on wages and wage floors, aggregate data vs model estimates

	Data, 2011 vs 2014			Model	
	Log wage (adj. part-time)	Log wage	Log wage floor	Log wage	Log wage floor
CB shock	0.0619	0.0369	0.0411	0.0203	0.0086
	(0.0313)	(0.0493)	(0.0374)	(0.0016)	(0.0073)
Share exporting empl. CB unit	Y	Y	Y	Y	Y
Region by 4d by exporter FE	Y	Y	Y	N	N
Region by exporter FE	N	N	N	N	N
Observations	1,390	1,492	1,404	1,079	1,079

Notes: The table shows the effect of CB shocks on log wages and log wage floors, in the aggregate data and the model-generated data. The model is estimated using 2011–2012 data. I simulate shocks in the model so that changes in minimum productivities mimic the effects of exporting shocks on revenue at the local labor market level, and then re-compute the model equilibrium using the new minimum productivities. Then, I regress the change in an outcome on the average local labor market shock at the CB level, using pre-period employment shares to weight the shocks. Both the aggregate data and the model-based data exclude local labor markets covered by the retail CB unit, and those in CB units with less than 5 percent of employment in exporting firms. Standard errors are clustered at the CB level.

Appendix

A Proofs for Theoretical Framework

Proof of Proposition 1. The Nash split equation is given by

$$\sum_{j\in\mathcal{J}} w_j \ell_j = \omega \sum_{j\in\mathcal{J}} \varphi_j f\left(\ell_j\right),$$

where $w_j = \underline{w}$ and $\ell_j = \ell_j(\underline{w}, \varphi_j)$ for constrained firms and $w_j = w(\varphi_j)$ and $\ell_j = \ell_j(w_j)$ for unconstrained ones.

Let $\mathcal{J}^{\text{co}}$ be the non-empty set of constrained firms, and $\mathcal{J}^{\text{uco}} = \mathcal{J}\backslash\mathcal{J}^{\text{co}}$ the set of unconstrained firms. Assume that ω is fixed. Differentiating with respect to φ_j and $\underline{w}$, and reordering terms yields

$$\begin{aligned}\sum_{j\in\mathcal{J}^{\text{co}}} \left(\ell_j + \underline{w}\frac{d\ell}{dw_j} - \omega\varphi_j f_\ell \frac{d\ell_j}{dw}\right) d\underline{w} &= \omega \sum_{j\in\mathcal{J}} f(\ell_j) d\varphi_j \\ &+ \sum_{j\in\mathcal{J}^{\text{co}}} \left(\omega\varphi_j f_\ell - \underline{w}\right)\frac{d\ell_j}{d\varphi} d\varphi_j + \sum_{j\in\mathcal{J}^{\text{uco}}} \left(\varphi_j f_\ell \frac{d\ell_j}{dw} - \ell_j + w_j \frac{d\ell_j}{dw_j}\right)\frac{dw_j}{d\varphi_j} d\varphi_j.\end{aligned} \tag{A.1}$$

where $L = \sum_{j'\in\mathcal{J}} \ell_{j'}$ is aggregate employment in the CBA.

From the first order conditions of the firm problem we know that $\ell_j + w_j \frac{d\ell_j}{dw_j} = \varphi_j f_\ell \frac{d\ell_j}{dw_j}$ for all $j \in \mathcal{J}^{\text{uco}}$, and that $\underline{w} = \varphi_j f_\ell$ for all $j \in \mathcal{J}^{\text{co}}$. I will add and subtract terms to drop the terms involving f_ℓ. I also substitute $d\underline{w} = \underline{w} d\ln\underline{w}$ and $d\varphi_j = \varphi_j d\ln\varphi_j$, and construct elasticities by multiplying and diving by appropriate terms. This yields

$$\begin{aligned}\sum_{j\in\mathcal{J}^{\text{co}}} \underline{w}\ell_j \left(1 + (1-\omega)\eta_j\right) d\ln\underline{w} &= \omega \sum_{j\in\mathcal{J}} \varphi_j f(\ell_j) d\ln\varphi_j \\ &- (1-\omega) \sum_{j\in\mathcal{J}^{\text{co}}} \underline{w}\ell_j \rho_j^\ell d\ln\varphi_j - (1-\omega) \sum_{j\in\mathcal{J}^{\text{uco}}} w_j\ell_j \left(1+\eta_j\right)\rho_j^w d\varphi_j,\end{aligned}$$

where $\eta_j = \frac{d\ell_j}{dw}\frac{w}{\ell_j}$, $\rho_j^\ell = \frac{d\ell_j}{d\varphi}\frac{\varphi_j}{\ell_j}$, and $\rho_j^w = \frac{dw_j}{d\varphi_j}\frac{\varphi_j}{w_j}$ are elasticities. Let us now define the adjusted wage bill

$$\tilde{WB}^{\text{co}} = \sum_{j\in\mathcal{J}^{\text{co}}} \underline{w}\ell_j \left(1 + (1-\omega)\eta_j\right). \tag{A.2}$$

Using this definition, recalling that $R_j = \varphi_j f(\ell_j)$ and $WB_j = w_j\ell_j$, and dividing and multiplying appropriately to obtain shares, we can write

$$d\ln\underline{w} = \frac{WB}{\tilde{WB}^{\text{co}}} \sum_{j\in\mathcal{J}} s_j^R d\ln\varphi_j - (1-\omega_c)\frac{WB}{\tilde{WB}^{\text{co}}} \left[\sum_{j\in\mathcal{J}^{\text{co}}} s_j^{WB} \rho_j^\ell d\ln\varphi_j - \sum_{j\in\mathcal{J}^{\text{uco}}} s_j^{WB}\left(1+\eta_j\right)\rho_j^w d\varphi_j\right].$$

Defining the elasticity of the j's wage bill to the productivity shock φ_j as $\iota_j = \rho_j^\ell$ for constrained firms and $\iota_j = (1+\eta_j)\,\rho_j^w$ for unconstrained firms, we can use the previous expression to obtain equation (4).

Now, if ω is not fixed, then (A.1) will have an extra term on the right hand side, namely

$$R\left(\frac{d\omega}{d\underline{w}_c} + \sum_j \frac{d\omega}{d\varphi_j}\right) = R\omega(1-\beta)\left(\frac{d\left[\left(-\frac{d\Pi}{d\underline{w}}/\frac{dU}{d\underline{w}}\right)\right]}{d\underline{w}} + \sum_j \frac{d\left[\left(-\frac{d\Pi}{d\underline{w}}/\frac{dU}{d\underline{w}}\right)\right]}{d\varphi_j}\right), \tag{A.3}$$

where $R = \sum_j \varphi_j f(\ell_j)$ is the aggregate revenue. After the corresponding algebraic manipulations, the final expression will resemble equation (4) but with an additional term that corresponds to (A.3) divided by $\tilde{W}B^{\text{co}}$.

□

Proof of Proposition 2. The results follow directly from the effects of the wage floor on firms discussed in Section 2.1. □

B Details on Context and Data

B.1 Labor market institutions in Argentina

The Law of Labor Contracts (N° 20.744) sets the general standards for all labor relations. Above this base, a set of CB agreements (CBAs) establishes standards that are binding for subgroups of workers in different industries, occupations, and firms. Private-sector CBAs are governed by the regime in Law N° 14.250, first sanctioned in 1953. Different regimes regulate CBAs for government employees and educators.

The CBAs are negotiated between unions and employer associations, and sometimes they are adhered to by other unions that did not participate directly in the negotiations. The government serves as a mediator and legal validator of these agreements. The terms established in CBAs set minimum standards for workers, which individual firms cannot alter to the workers' detriment.

Types of unions. The law enables any group of workers to form a union. Unions exist in 3 legal forms: basic unions (*sindicatos*), which directly represent workers, are the most common; federations (*federaciones*), which are groups of unions; and confederations (*confederaciones*), which agglomerate federations and basic unions.

Despite this freedom of association, not all unions are legally allowed to negotiate CBAs. Only one union per "area of representation" is endowed with "bargaining privileges." The government grants privileges to the union that meets certain requirements, such as being the one with more affiliates among the workers it aims to represent.[1] It is not uncommon to find unions with bargaining privileges that simply adhere to existing CBAs. An example is the retail sector (*comercio*): its

[1]Unions that are registered but do not have bargaining privileges are known as *sindicatos simplemente inscriptos* (roughly, "unions simply registered"), while those with privileges are *sindicatos con personería gremial* ("unions with legal recognition"). The criteria for assigning *personería gremial* to unions are outlined in Law N° 23.551.

most important CBA (0130/75) is adhered by many regional basic unions. Furthermore, a single union can participate in multiple CBAs and a single CBA can have multiple adhering unions. In the paper I focus on the role of CB units, abstracting from these complexities.

Areas of representation and coverage. An area of representation can be determined by industry, occupation, geographical location, or even a single employer, and is formally defined when the government grants bargaining privileges to a union. Areas of representation effectively delimit the scope of CB units. The government has the authority to change these areas by granting new bargaining privileges or revoking existing ones. However, areas of representation have been stable in the recent past, especially in mature sectors exposed to international trade. CBAs signed within these areas of representation are binding for all workers and firms within them. This is so because the law establishes "universal coverage," meaning that CBAs are automatically extended to workers not affiliated with the union, and "automatic extension," meaning that CBAs are automatically extended to all firms in the area of representation.

CBAs, CBA alterations, and the negotiation process. A union with bargaining privileges and an employer association will typically negotiate a comprehensive CBA that outlines labor regulations applicable to the workers they represent. I refer to these as "master CBAs," or simply CBAs when the context is clear. Procedural rules for the negotiation process, established by law, define protocols for unions to formally request meetings with employers, facilitate information exchange (e.g., employers providing details about labor costs and organizational structures), among other considerations.[2] The use of strikes, as well as the procedure in case of employer crisis, is regulated by law as well. If an agreement is not reached, the government can issue an arbitral award to determine the regulations for the labor contracts of the involved workers.

New agreements are legally validated by the Ministry of Labor in a process termed *homologación*. The government archives these master CBAs under unique codes, which align with the CBA codes I observe in the data. A master CBA may be modified by either a new master CBA that supersedes it, or a "CBA alteration" that simply updates some provisions within it.[3] While master CBAs have an expiration date, even if they are not renegotiated they remain in force until a new master CBA replaces them due to a clause known as "ultra-activity." CBA alterations act as amendments to the master CBA and are negotiated more frequently. They typically relate to updates in wage scales, although they may also entail modifications in other provisions.

The dynamics of negotiations. Panel (d) of Figure 1 shows the dynamics of collective negotiations in 2005–2019. After a period of low activity in the 90s, the number of negotiations reignited in the early 2000s.[4] The recovery from the 2001-2002 crisis triggered government interventions

[2]The procedural rules for collective bargaining are established in Law N° 23.546.

[3]If the CBA is completely revised the code of the CBA in the data changes as well. I reviewed these cases so that a constant code appears in my data for the same CB unit.

[4]The 90s negotiations took place in a context of pro-market reforms that weakened traditional unions. See Palomino and Trajtemberg (2006) for a discussion of the dynamics of negotiations in 1991–2006.

affecting private sector wages, such as minimum wage increases and even wage supplements by decree. These developments revitalized the negotiations, which at first incorporated these provisions into existing CBAs. The new ruling party, elected in 2003, introduced legislation that further encouraged negotiations, galvanizing unions into signing new master CBAs. These factors account for the peak of 150 new master CBAs signed in 2006. By 2014, 52% of active master CBAs had been signed in 2003 or later (Pontoni and Trajtemberg 2017). However, soaring inflation since 2007 prompted unions and employers to meet nearly annually to revise wages via CBA alterations, resulting in a steep increase in the number of CBA alterations.[5]

B.2 Data

B.2.1 Main labor market data

The primary source of information on the formal labor market is Argentina's matched employer-employee dataset. The data are collected by the tax authority for social security purposes under a system known as *Sistema Integrado Previsional Argentino* (SIPA). I have gained access to a version of this data maintained by the Ministry of Labor (Ministerio de Trabajo, Empleo y Seguridad Social 2022b), covering 2007–2020. The data contain worker identifiers, worker monthly compensation, and worker characteristics such as age and gender. I also observe firm identifiers, fiscal province and postal code, and their 6-digit economic sector which corresponds to a custom version of the ISIC classification system, version 4.[6] While the dataset does not contain information on hours or full-time status, it does contain a "hiring modality." This variable contains tens of categories, but the most common one (number 8) usually corresponds to full-time workers under a permanent contract. However, as suggested by Appendix Figure 2, part-time workers are oftentimes declared under the main hiring modality as well.

I use a second administrative dataset to obtain additional information on labor relations (Ministerio de Trabajo, Empleo y Seguridad Social 2022a). This dataset, collected also by the tax authority under a system known as *Simplificación Registral*, is constructed from employer's online declarations during hiring or termination of workers. The system was introduced progressively since 2008. Appendix Table B.1 lists the size thresholds that determined whether firms were required to or had the option of entering the system at different times. The idea was to accommodate larger firms who might need more time to adjust to the new system.

The goal of the new system was to simplify the process of registering workers and to collect information that could be used to, for example, determine eligibility for government programs or family allowances. The dataset contains the CBA code, the within-CBA category, and an occupation code, among other variables that I did not get access to for the paper.

[5]The 2017 dip in the number of CBA alterations can also be attributed to inflation-related developments. Encouraged by the government, some CB units introduced a "trigger clause" that would automatically update wages in case of unanticipated inflation, thus reducing the need to negotiate.

[6]From this sector code I compute 2- through 4-digit codes as needed. The 1-digit code I use in the paper corresponds to the "letter" of the ISIC codes, and is the broadest category available.

Importantly, both datasets contain the same worker and firm identifiers. As a result, I can use *Simplificación Registral* to add information on the CBA code to the matched employer-employee dataset. I describe the processing of the CBA code variable later in this appendix.

Appendix Table B.1: Mandatory usage of *Simplificación Registral* by firm employment

Period	**Threshold (employment)**
January 2008 to July 2012	10
August 2012 to March 2013	25
April 2013 to March 2014	100
April 2014 to March 2015	200
April 2015 to April 2016	300
May 2016 to July 2017	400
August 2017 to November 2017	600
December 2018 to July 2018	2000
August 2018 onwards	Any

Source: Resolución general AFIP 4265/2018.

B.2.2 Other data sources

First, I collected data from the publicly available BACI-CEPII dataset (Gaulier and Zignago 2010), which contains yearly trade flows between any pair of countries in each 6-digit product from the Harmonized System (HS) of product classification. In particular, I use the data coded with the 2007 version of the HS system, which covers 2007–2020. Second, I obtained data from Argentine customs (*Dirección General de Aduanas*) which details the value exported to each country and product for every Argentine firm. As a member of Mercosur, Argentina's product classification system is based on the *Nomenclatura Común del Mercosur* (NCM), which is an 8-digit code that is compatible with the HS. Using concordance tables from Liao et al. (2020), I convert NCM codes into 6-digit 2007 HS codes.[7] I use names to match country codes between the Customs data and BACI-CEPII. Then, I join the datasets using country and HS codes. Because this matching is imperfect 0.81 percent of exporting firms in 2011–2012 have less than 99 percent of their exporting value matched to a country-product. I drop these firms from the analysis.

To study how a firm's accounting outcomes respond to economic shocks, I use survey data from a national business survey. This survey, known as *Encuesta Nacional de Dinámica de Empleo e Innovación* (ENDEI), was conducted jointly by the Ministry of Labor and the Ministry of Science and Technology in two waves, in 2012 and 2016. The survey asked about the firm's situation in regard to innovation in the 3 previous years. Fortunately, the survey also asked about revenue and expenditures. Some firms were sampled twice, and are thus observed for 6 years. I identify 1,800

[7] To convert codes I proceed as follows. First, I convert NCM 8-digit to NCM 6-digit by keeping the first 6 digits. These 6 digits directly correspond to the HS, although not necessarily the 2007 version. I then convert codes in different years to HS 6-digit version 2007 using the appropriate concordance table. I impute a handful of codes that are not present in the concordance manually to minimize missing values.

firms that participated in the survey at least once and exported in 2011–2012, which accounts for 22.6% of the firms in my primary estimation sample (as seen in Appendix Table 2).

B.2.3 Imputation of CBA codes

Cleaning CBA codes at worker-firm level. Due to the progressive introduction of *Simplificación Registral* many workers in SIPA do not have a CBA code. Additionally, sometimes the CBA code is outdated because a new master CBA was signed that supersedes the old one.

To increase the number of workers with a CBA code and update the codes so that they reflect a constant CB unit, I proceed in three steps. First, in the early years of the system most workers are observed by their termination date only. I use this information to fill the declared CBA code backwards, which increases coverage quite a bit for 2008 through 2010. Second, I impute the CBA code to workers with a missing code in a firm-year cell if a single code is observed. If a single CBA code is observed in increasingly large cells (such as firm-occupation, 6-digit ISIC-postal code, and occupation-postal code), I impute that code to workers with a missing code in that cell. Finally, I update the CBA codes forward to take into account updates to master CBAs that result in a new code. To do so, I scraped data from an online search engine of CBAs constructed by the Ministry of Labor. The search engine tells the user whether a master CBA updates a previous one. Results were manually reviewed to ensure that the information was correct.

Appendix Figure B.1 shows the share of workers in the employer-employee dataset with a non-missing CBA code in the raw data and after the imputation described below. The imputation increases coverage significantly. The most important steps in the imputation are the backwards filling, which raises coverage in the yearly years, and the imputation using firm-year cells, which increases coverage by around 13% after 2010.

Defining CB units. For firms employing workers across multiple CBAs, I assign as primary CBA code the modal CBA code. These primary CBA codes correspond to CB units. If a firm does not have any workers associated with a CBA code, I assign the most frequent CBA code in the postal code and 4-digit economic sector. A few postal codes and 4-digit sector cells do not have any worker with a declared CBA code. In such cases I use wider cells defined by province and 4-digit economic sector. About 25% of codes across the economy are imputed. Among firms with a non-imputed CBA code, the average share of workers with the primary CBA code is 97%.

C Computation of Wage Floors

The data lacks information on the wage floors set by the CBAs. I therefore use the distribution of wages within a CBA, worker category, CBA-region, and month to infer the wage floors.[8] I use workers declared in the main hiring modality and exclude the first and last month of a spell as

[8] Card and Cardoso (2022) observe wage floors but not within-CBA categories. They rely on a lengthy matching process of within-CBA categories to occupations in the administrative data to assign wage floors to workers.

Appendix Figure B.1: Share of workers with non-missing CBA code, raw data

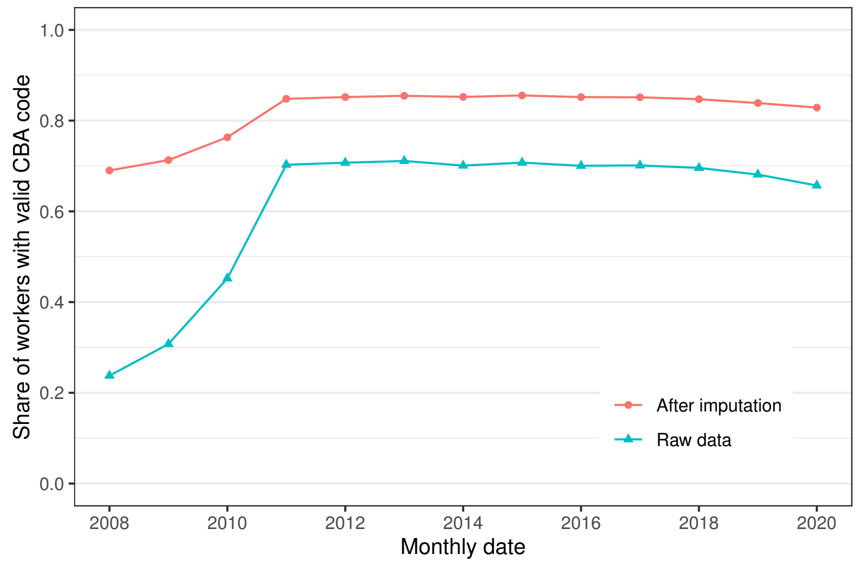

Notes: The figure shows the share of workers in the employer-employee dataset that can be matched to a valid CB agreement code.

well as the months of June and December, as these months correspond to the 13th-month salary payments. Then, I smooth the resulting time series of wage floors to reduce noise.

Figure C.2 shows the distribution of wages for workers in the main hiring modality within two categories of a given CBA, CBA-region, and month. We observe a clear bunching of wages at the wage floor. There is a significant mass of workers earning below the wage floor, which is to be expected as we do not control for hours worked and simply observe the total monthly wage received by the worker.[9] We also observe a smaller bunching point at around half the wage floor, which corresponds to workers that actually work part-time.

The goal is to identify the bunching point. One possible way is to follow the approach suggested by Cardoso and Portugal (2005) and pick the mode of the distribution. However, as shown by the left panel of Appendix Figure C.2, the mode may be higher than the wage floor. A second option is to pick the point at which the distribution increases more rapidly. I implement this approach by selecting the bin at which the difference between a 20-bin average and the subsequent 5-bin average is maximized. This is the preferred approach as it seems to capture the actual wage floor in a variety of cases in which the mode is not a good approximation.

Now, after identifying the wage floors, I end up with a monthly time series of wage floors for each CBA, category, and CBA region cell. I start by dropping cells that appear for less than 3.5 years. I also drop categories that show an "implausible" behavior, i.e., for which the wage floor decreases for a period of time and then increases again. These restrictions attempt to drop the noisy estimated wage floors from CBAs with few workers. Finally, I smooth the time series of wage

[9]Some reasons people may earn less than the wage floor are part-time workers declared in the main hiring modality or workers on vacation not receiving the usual pay supplement for attendance.

Appendix Figure C.2: Distribution of wages within a CBA and CBA-region, Jan. 2012

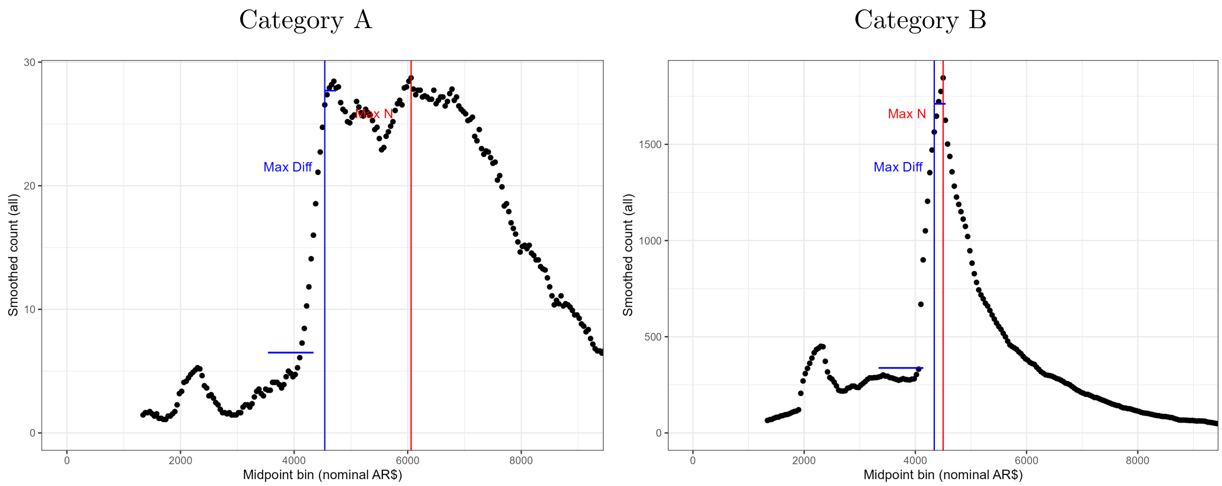

Notes: The figures show the distribution of wages within a CBA, category, CBA region, and month. The bins are equal to 20 pesos in 2012, but change with the inflation rate in subsequent years. The bins were smoothed using a moving average with a window of 5 bins at each side. The blue line shows the "maximum difference" between a 20-bin average with respect to the subsequent 5-bin average. The red line shows the mode.

floors using a fixed effects model that imposes that log wage floors are a linear function of CBA by category by CBA region fixed effects and CBA-region by month fixed effects. The structure of most contracts is such that the relative difference between categories is constant over time and all wage floors increase at the same rate. The fixed effects model captures this structure, reducing the noise in the data. I conclude by smoothing each series with a 1-month moving average.

Appendix Figure C.3 shows a comparison between manually-collected wage floors and the wage floors inferred from the distribution of wages. The manual collection is challenging as it requires reading the actual agreements, which are usually in PDF format and do not follow a consistent structure over time. Furthermore, it is common for agreements to include non-compensatory payments, or one-time payments, that are hard to identify. Panel (a) shows that the levels of the wage floors are similar, with the data-inferred wage floors being slightly higher (suggesting that the manual collection may be missing some mandatory non-compensatory payments). The comparison is similar when using the average wage floor in each semester, as shown in Panel (b). Finally, Panel (c) compares the changes in the wage floors across semesters. The manually-collected wage floors exhibit a higher volatility than the data-inferred wage floors. This is likely due to the fact that the manual collection is missing some form of compensation. For example, if the manual collection missed the update of the wage scales in a given period, then it would appear to jump by a larger amount in the next wage update.

These data sources are in approximate agreement, suggesting that the data-inferred wage floors serve as a good approximation to the actual floors set by CBAs. I acknowledge that they may be noisy. However, this noise is unlikely to be correlated with the CB shocks, and therefore should not affect the empirical results of the paper.

Appendix Figure C.3: Validation of wage floors with manually-collected wage floor data

(a) Monthly comparison

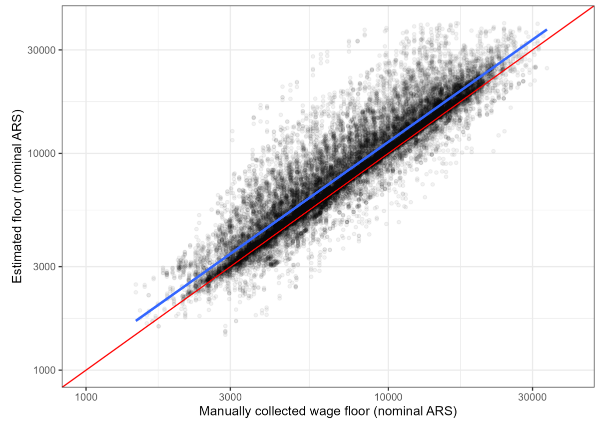

(b) Semester average comparison

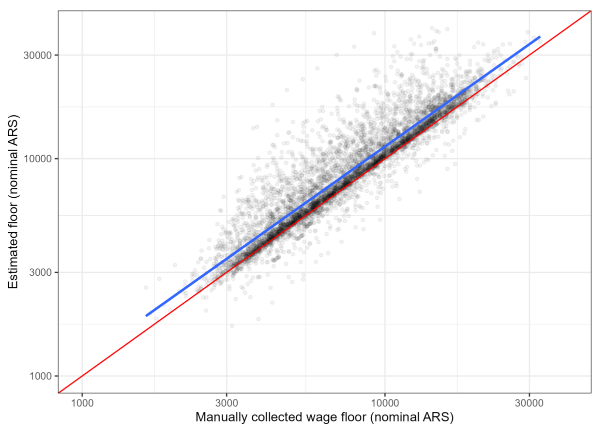

(c) Changes in semester average

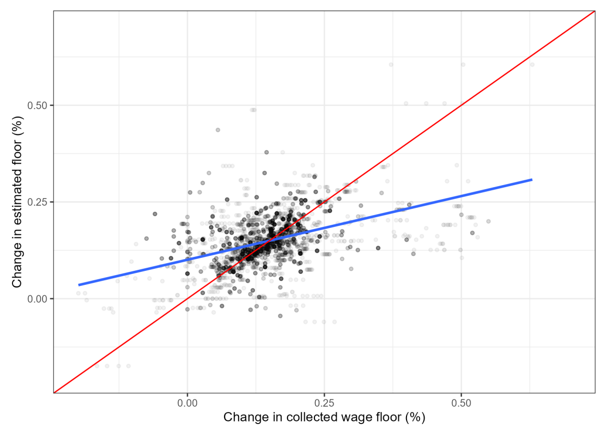

Notes: The figure compares the manually-collected wage floors and the data-inferred wage floors. Panel (a) shows a monthly comparison of the wage floors. Panel (b) shows a similar comparison but using the average wage floor in each semester. Panel (c) compares the changes in the semester average wage floor.

D Shift-share Identification with Two Treatments

In this section I will show that identification of equation (5) can be cast in terms of country-product shocks, following Borusyak et al. (2022b). For simplicity, I consider two time periods and assume a single intercept for all firms. The reasoning would go through with more time periods and local labor market intercepts, but the notation would be more cumbersome.

Consider the following causal model for the change in outcome y_j

$$\Delta y_j = \beta_1 z_{j1} + \beta_2 z_{j2} + \varepsilon_j, \tag{D.1}$$

where Δy_j is the demeaned change in y for firm j and z_{j1} and z_{j2} are defined as

$$z_{jn} = \sum_{p \in \mathcal{P}} s_{jpn} f_p$$

for $n \in \{1,2\}$. s_{jpn} are exposure shares (with $\sum_{p\in\mathcal{P}} s_{jpn} = 1$) and f_p are common shifts. This setting aligns with the main paper since we have two shift-share variables (the firm and CB shocks) that differently weight the same set of country-product demand shocks. Given this non-*iid* setting, I consider identification of β_1 and β_2 by the full-data moment conditions

$$\mathbb{E}\left[\frac{1}{|\mathcal{J}|}\sum_{j\in\mathcal{J}} z_{j1}\varepsilon_j\right] = 0 \quad \text{and} \quad \mathbb{E}\left[\frac{1}{|\mathcal{J}|}\sum_{j\in\mathcal{J}} z_{j2}\varepsilon_j\right] = 0 \tag{D.2}$$

where ε_j is the residual from (D.1) and $\mathcal{J}$ is the set of firms. When these conditions hold the parameters are identified (provided the shift-share variables are not perfectly collinear).

To show that an assumption on shocks f_p allow identification of the parameters β_1 and β_2, rewrite the moment conditions as

$$\begin{aligned} 0 &= \mathbb{E}\left[\frac{1}{|\mathcal{J}|}\sum_{j\in\mathcal{J}} z_{jn}\varepsilon_j\right] = \mathbb{E}\left[\frac{1}{|\mathcal{J}|}\sum_{j\in\mathcal{J}}\left(\sum_{p\in\mathcal{P}} s_{jpn} f_p\right)\varepsilon_j\right] = \mathbb{E}\left[\sum_{p\in\mathcal{P}} f_p \frac{1}{|\mathcal{J}|}\sum_{j\in\mathcal{J}} s_{jpn}\varepsilon_j\right] \\ &= \mathbb{E}\left[\sum_{p\in\mathcal{P}} s_{pn} f_p r_{pn}\right], \end{aligned}$$

where $n \in \{1,2\}$, $s_{pn} = (1/|\mathcal{J}|)\sum_{j\in\mathcal{J}} s_{jpn}$ are shock-level weights (where $\sum_p s_{pn} = 1$ as well), and $r_{pn} = \frac{(1/|\mathcal{J}|)\sum_{j\in\mathcal{J}} s_{jpn}\varepsilon_j}{(1/|\mathcal{J}|)\sum_{j\in\mathcal{J}} s_{jpn}}$ are average residuals for each category p weighted by appropriate shares. These conditions define a p-level GMM problem.

The key sufficient assumption for identification is $\mathbb{E}\left[f_p|s,r\right] = \mu$, which amounts to a quasi-randomness assumption on the shocks f_p with respect to the shares $s = \{s_{pn}\}_{p,n}$ and shocks $r = \{r_{pn}\}_{p,n}$, and is analogous to Assumption 1 in Borusyak et al. (2022b).[10] This assumption guarantees that the previous conditions hold since

$$\mathbb{E}\left[\sum_{p\in\mathcal{P}} s_{pn} f_p r_{pn}\right] = \mathbb{E}\left[\sum_{p\in\mathcal{P}} s_{pn}\,\mathbb{E}\left[f_p|s,r\right] r_{pn}\right] = \mu\,\mathbb{E}\left[\sum_{p\in\mathcal{P}} s_{pn} r_{pn}\right] = 0.$$

The first equality follows from the law of iterated expectations, and the second from the fact that $\mathbb{E}\left[\sum_{p\in\mathcal{P}} s_{pn} r_{pn}\right] = 0$.[11] Given that the moment conditions of the p-level problem are equivalent to the moment conditions (D.2), the parameters β_1 and β_2 in (D.1) are identified.

The interpretation of the quasi-randomness assumption is that the shifts f_p have the same mean μ regardless of the realization of the unobservables r_{pn} (and shares s_{pn}). As discussed by

[10] A second assumption, in the spirit of Assumption 2 in Borusyak et al. (2022b), is required for consistency. Namely, that $\mathbb{E}[s_{pn}^2] \to 0$ for $n \in \{1,2\}$ and $\mathrm{Cov}\left(f_p, f_{p'}|s,r\right) = 0$ for all (p,p') with $p \neq p'$. This assumption requires a large effective number of shocks, and that shocks are uncorrelated given the unobservables and shares. The authors discuss ways in which these assumptions can be relaxed.

[11] To show this result recall that $\sum_{p\in\mathcal{P}} s_{jpn} = 1$ and $\mathbb{E}\left[(1/|\mathcal{J}|)\sum_{j\in\mathcal{J}}\varepsilon_j\right] = 0$ by construction. Then, $\mathbb{E}\left[\sum_{p\in\mathcal{P}} s_{pn} r_{pn}\right] = \mathbb{E}\left[\sum_{p\in\mathcal{P}}(1/|\mathcal{J}|)\sum_{j\in\mathcal{J}} s_{jpn}\varepsilon_j\right] = \mathbb{E}\left[(1/|\mathcal{J}|)\sum_{j\in\mathcal{J}}\varepsilon_j\left(\sum_{p\in\mathcal{P}} s_{jpn}\right)\right] = 0.$

Borusyak et al. (2022b), this assumption can be relaxed by the inclusion of controls. First, adding controls in (D.1) allows one to remove variation in the error term that may be correlated with the shift-share variables. For instance, if one thinks that firms in some industries tend to both grow faster and also experience increases in world demand, then controlling for industry fixed effects will mean that shocks f_p must be quasi-random with respect to the within-industry residuals. Second, controlling for average exposure to clusters of p's allows the average of f_p to differ by cluster, relaxing the quasi-randomness assumption as well. I take these two points as a motivation for the inclusion of controls in the main paper. We see that the logic of Borusyak et al. (2022b) applies to the case of two shift-share variables that use the same set of shocks but weighted by different shares.

E The effects on shocks on revenue in an event-study design

This section describes the estimates of the effects of firm shocks on firm revenue and employment. This section uses data from a survey of businesses (ENDEI) conducted jointly by the Ministry of Labor and the Ministry of Science and Technology in the years 2010–2016 (excluding 2013).

E.1 Empirical Strategy

Consider the following panel model:

$$y_{jt} = \theta z_{c(j)t} + \beta z_{jt} + \alpha_j + \delta_{\ell(j)t} + \varepsilon_{jt}$$

where variables are defined as in equation (5). This model imposes that only the period-t shocks affect the outcome. However, it is possible that the effect of the shock takes time to fully materialize. To account for this, I follow Freyaldenhoven et al. (forthcoming) and add a specific transformation of the leads and lags of the shock variables into the model:

$$\begin{aligned} y_{jt} = {} & \beta_{-\underline{R}-1}\left(-z_{j,t+\underline{R}}\right) + \sum_{r\in\mathcal{R}} \beta_k \Delta z_{j,t-r} + \beta_{\overline{R}}\left(z_{j,t-\overline{R}}\right) \\ & + \theta_{-\underline{R}-1}\left(-z_{c(j),t+\underline{R}}\right) + \sum_{r\in\mathcal{R}} \theta_k \Delta z_{c(j),t-r} + \theta_{\overline{R}}\left(z_{c(j),t-\overline{R}}\right) \\ & + \alpha_j + \delta_{\ell(j)t} + \varepsilon_{jt}. \end{aligned} \tag{E.1}$$

$\mathcal{R}$ is a set of "relative years" running from $-\underline{R}$ to $\overline{R}-1$. Given that the end-points and the first-differences of each z sum to 0, a normalization is required to avoid perfect multicollinearity.

The parameters of this model are identified by different comparisons relative to equation (5). The outcome at each year t is affected by a set of event-study coefficients of each treatment, which are identified by the conditional covariance of the period-t change in outcome to the period-t change in the treatment, relative to the normalized period. Future values of the treatment (i.e., forward values up to $\overline{R}$ years) pin down the pre-period coefficients. If there is no correlation between future shocks and current outcomes, then these coefficients should be zero.[12] Similarly, past values of the z's (i.e., lagged values up to $\underline{R}$ years) identify the post-period coefficients.

I set $\underline{R} = 4$ and $\overline{R} = 3$. Since the shock variables are observed in all years 2007–2020, this allows me to compute an entire set of event-study coefficients for any firm in the survey data in 2010–2012 or 2014–2016. For those observed in 2010, I compute post-coefficients using data starting from $2010 - \overline{R} = 2007$, and for those observed in 2016, I compute pre-coefficients using data up to $2016 + \underline{R} = 2020$. A second decision concerns the period to normalize to 0. Considering that the treatment variables exhibit some AR(1) autocorrelation, I normalize the relative year -2 for both sets of event-study variables.

Appendix Figure E.4: Effect of firm shocks on firm's accounting outcomes, panel event-study design

(a) Log wage expenditures (b) Log product market sales

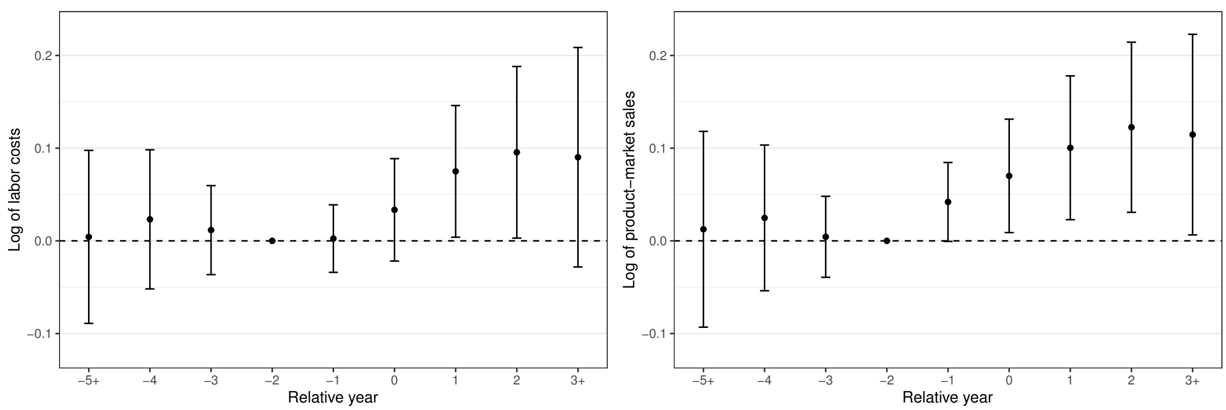

Notes: Data are from a sample of firms that exported in 2011–2012 that were surveyed in the first or second waves of the *Encuesta Nacional de Dinámica del Empleo y la Innovación* (ENDEI). The figure shows coefficients on the effect of firms shocks on a firm accounting outcomes, using a panel event-study design as described in Appendix E. The time-varying firm shock z_{jt} is defined as the average world import demand across the country-products exported by the firm in 2011–2012, weighting by the share of the value exported to each country-product. The regressions include controls for firm fixed effects and 4-digit economic sector by province by year fixed effects. Standard errors are clustered at the firm level.

E.2 Results

Appendix Figure E.4 shows effects of the firm shock on product-market sales and labor expenditures labor using data from the survey of businesses. We observe an increase in both variables following the shock. A 10% increase in average world import demand rises product-market sales by 1.13%, and labor costs by 0.87%, on average in relative years 1–3. Both panels show stable pre-trends, supporting the underlying parallel-trends assumption of the model.

[12] Contrast this with the DiD model where zero pre-trends for those treated at t result from no correlation between *current* shocks and *past* outcomes.

F Details on Structural Model

F.1 Derivations

F.1.1 Labor supply: supply to the firm decision

The cdf and pdf of a Fréchet distribution with shape η, scale equal to one and location equal to zero are given by

$$F(\xi) = e^{-\xi^{-\eta}}, \quad f(\xi) = \eta\xi^{-\eta-1}e^{-\xi^{-\eta}}.$$

Let Ω_r be the set of firms operating in region r. The share of formal-sector workers in region r who optimally choose firm j is given by

$$\begin{aligned}
\ell(j) &= \int_0^1 \Pr\big(V_{ir}(j) \geq V_{ir}(j') \forall j' \neq j\big)\, di \\
&= \int_0^1 \int_0^\infty f(\xi(j)) \prod_{j' \in \Omega_r \setminus \{j\}} F\left(\frac{w(j)A_{k1(j)}\xi(j)}{w(j')A_{k1(j')}}\right) d\xi(j)di \\
&= \int_0^1 \int_0^\infty f(\xi(j)) \prod_{j' \in \Omega_r \setminus \{j\}} F\left(\frac{w(j)A_{k1(j)}\xi(j)}{w(j')A_{k1(j')}}\right) d\xi(j)di \\
&= \int_0^1 \int_0^\infty \eta\xi(j)^{-\eta-1}e^{-\xi(j)^{-\eta}} \prod_{j' \in \Omega_r \setminus \{j\}} e^{-\left(\frac{w(j)A_{k1(j)}\xi(j)}{w(j')A_{k1(j')}}\right)^{-\eta}} d\xi(j)di \\
&= \int_0^1 \int_0^\infty \eta\xi(j)^{-\eta-1}e^{\int_{j' \in \Omega_r} -\left(\frac{w(j)A_{k1(j)}\xi(j)}{w(j')A_{k1(j')}}\right)^{-\eta} dj'} d\xi(j)di \\
&= \int_0^1 \int_0^\infty \eta\xi(j)^{-\eta-1}e^{-\left(w(j)A_{k1(j)}\xi(j)\right)^{-\eta}} e^{\int_{j' \in \Omega_r}\left(w(j')A_{k1(j')}\right)^{\eta} dj'} d\xi(j)di.
\end{aligned}$$

Defining $s = \frac{\left[\int_{j' \in \Omega_r}\left(w(j')A_{k1(j')}\right)^{\eta} dj'\right]^{1/\eta}}{w(j)A_{k1(j)}}$ we can rewrite as

$$\ell(j) = \int_0^1 \int_0^\infty \eta\xi(j)^{-\eta-1}e^{-\left(\frac{\xi(j)}{s}\right)^{-\eta}} d\xi(j)di,$$

and further manipulations yield

$$\ell(j) = s^{-\eta} \int_0^1 \int_0^\infty \frac{\eta}{s}\left(\frac{\xi(j)}{s}\right)^{-\eta-1} e^{-\left(\frac{\xi(j)}{s}\right)^{-\eta}} d\xi(j)di,$$

Note that, as workers are homogeneous, the integral over i is irrelevant. Also, the expression inside the integrals is the pdf of a Fréchet distribution with shape η and scale s. Thus, the double integral above equals 1, and we have

$$\ell(j) = \left(\frac{w(j)A_{k1(j)}}{\int_{j' \in \Omega_r}\left(w(j')A_{k1(j')}\right)^{\eta} dj'}\right)^{\eta}.$$

The expression for expected utility can be obtained following similar steps. See Appendix D in Parente (2022).

F.1.2 Labor supply: Formal employment decision

Derivations are available in Appendix D.2.1 of Ahlfeldt et al. (2022).

F.1.3 Labor demand: solution to firm problem

For *unconstrained* firms the solution to the firm's problem is

$$w(\varphi) = \left(\frac{\eta}{\eta+1}\right)\varphi, \quad \ell(\varphi) = W_r^{-\eta}\left(A_{k1}\left(\frac{\eta}{\eta+1}\right)\varphi\right)^{\eta}, \quad \pi(\varphi) = \frac{\eta^{\eta}}{(\eta+1)^{\eta+1}}\varphi^{\eta+1}A_{k1}^{\eta}W_r^{-\eta}.$$

For *constrained* firms we have

$$w(\varphi) = \underline{w}_c, \quad \ell(\varphi) = W_r^{-\eta}\left(A_{k1}\underline{w}_c\right)^{\eta}, \quad \pi(\varphi) = \left(\varphi - \underline{w}_c\right)\underline{w}_c^{\eta}A_{k1}^{\eta}W_r^{-\eta}.$$

F.1.4 Local labor markets: derivations

Recall that the cdf of productivity is $F_g(\varphi) = 1-(\varphi/\varphi_{g0})^{\alpha}$ if $\varphi > \varphi_{g0}$ and zero otherwise. The conditional cdf on some value x can be found by simply using x instead of φ_{g0} in the expression above. Also recall the value of the thresholds is $\overline{\varphi}_g = \underline{w}_c(\eta+1)/\eta$ and $\underline{\varphi}_g = \underline{w}_c$.

Share of firms bunching at the wage floor First, if $\varphi_{g0} > \overline{\varphi}_g > \underline{\varphi}_g$ all firms pay above the wage floor, thus $S_g = 0$. Second, if $\overline{\varphi}_g > \varphi_{g0} > \underline{\varphi}_g$, then the minimum productivity is φ_{g0}. The share of firms bunching at the wage floor is

$$\begin{aligned} S_g &= F_g\left(\overline{\varphi}(\underline{w}_{c(g)})\right) - F_g\left(\varphi_{g0}\right) = \left[\left(\frac{\varphi_{g0}}{\varphi_{g0}}\right)^{\alpha}\right] - \left[\left(\frac{\varphi_{g0}}{\overline{\varphi}(\underline{w}_{c(g)})}\right)^{\alpha}\right], \\ &= 1 - \left(\frac{\varphi_{g0}}{\underline{w}_{c(g)}}\right)^{\alpha}\left(\frac{\eta}{\eta+1}\right)^{\alpha}. \end{aligned} \tag{F.1}$$

Finally, if $\overline{\varphi}_g > \underline{\varphi}_g > \varphi_{g0}$, then the minimum productivity is $\underline{\varphi}_g = \underline{w}_c$. Thus,

$$\begin{aligned} S_g &= F_g\left(\overline{\varphi}(\underline{w}_{c(g)})\right) - F_g\left(\underline{\varphi}(\underline{w}_{c(g)})\right) = \left[\left(\frac{\underline{w}_c}{\underline{\varphi}(\underline{w}_{c(g)})}\right)^{\alpha}\right] - \left[\left(\frac{\underline{w}_c}{\overline{\varphi}(\underline{w}_{c(g)})}\right)^{\alpha}\right] \\ &= 1 - \left(\frac{\eta}{\eta+1}\right)^{\alpha}. \end{aligned}$$

This is the maximum share of observed firms bunching that is compatible with the model.

Average wage For any threshold x, the average wage of *unconstrained* firms is given by

$$\begin{aligned}\bar{w}_g^u &= \int_x^\infty w(z) f_g(z|x) dz \\ &= \int_x^\infty \left(\frac{\eta}{\eta+1}\right) z \alpha x^\alpha z^{-\alpha-1} dz \\ &= \left(\frac{\eta}{\eta+1}\right)\left(\frac{\alpha}{\alpha-1}\right) x\end{aligned}$$

Then, if $\varphi_{g0} \geq \overline{\varphi}_g$ there are no constrained firms, so that

$$\bar{w}_g^u = \left(\frac{\eta}{\eta+1}\right)\left(\frac{\alpha}{\alpha-1}\right)\varphi_{g0}.$$

Now, if $\overline{\varphi}_g > \varphi_{g0} > \underline{\varphi}_g$, the minimum productivity of unconstrained firms is $\overline{\varphi}_g$, implying

$$\bar{w}_g^u = \left(\frac{\alpha}{\alpha-1}\right)\underline{w}_c.$$

This holds whether $\varphi_{g0} > \underline{\varphi}_g$ or not.

The average wage of *constrained* firms is given by $\underline{w}_c$ if there are any constrained firms. If $\varphi_{g0} \geq \overline{\varphi}_g$ no firm is constrained so this quantity is not defined.

Then, the average wage in a local labor market is given by a weighted average of the unconstrained and constrained wages, where the weights are given by the share of firms bunching. Going over the computations for each case, we get

$$\bar{w}_g = \begin{cases} \left(\frac{\eta}{\eta+1}\right)\left(\frac{\alpha}{\alpha-1}\right)\varphi_{g0} & \text{if } \varphi_{g0} > \overline{\varphi}, \\ \left(1 + \left(\frac{\varphi_{g0}}{\underline{w}_c}\right)^\alpha \left(\frac{1}{\alpha-1}\right)\left(\frac{\eta}{\eta+1}\right)^\alpha\right)\underline{w}_c & \text{if } \overline{\varphi} > \varphi_{g0} > \underline{\varphi}_g, \\ \left(1 + \left(\frac{1}{\alpha-1}\right)\left(\frac{\eta}{\eta+1}\right)^\alpha\right)\underline{w}_c & \text{if } \overline{\varphi} > \underline{\varphi}_g > \varphi_{g0}. \end{cases}.$$

Aggregate quantities To compute aggregate labor demand in g we need to solve $L_g = M_g \int_{\varphi_{g0}}^\infty \ell(\varphi) f_g(\varphi) d\varphi$. Solving the integral for each case we get:

$$L_g = M_g W_r^{-\eta} A_{k1}^\eta \begin{cases} \left(\frac{\eta}{\eta+1}\right)^\eta \left(\frac{\alpha}{\alpha-\eta}\right)\varphi_{g0}^\eta & \text{if } \varphi_{g0} > \overline{\varphi}_g, \\ \underline{w}_c^\eta \left\{1 + \left(\frac{\eta}{\alpha-\eta}\right)\left(\frac{\eta}{\eta+1}\right)^\alpha \left(\frac{\varphi_{g0}}{\underline{w}_c}\right)^\alpha\right\} & \text{if } \overline{\varphi}_g > \varphi_{g0} > \underline{\varphi}_g, \\ \underline{w}_c^\eta \left(\frac{\varphi_{g0}}{\underline{w}_c}\right)^\alpha \left\{1 + \left(\frac{\eta}{\alpha-\eta}\right)\left(\frac{\eta}{\eta+1}\right)^\alpha\right\} & \text{if } \overline{\varphi}_g > \underline{\varphi}_g > \varphi_{g0}, \end{cases}$$

and the partial derivative with respect to the wage floor, holding constant the wage index, is given by

$$\frac{\partial L_g}{\partial \underline{w}_c} = M_g W_r^{-\eta} A_{k1}^{\eta} \begin{cases} 0 & \text{if } \varphi_{g0} > \overline{\varphi}_g, \\ \underline{w}_c^{\eta-1} \eta \left\{ 1 - \left(\frac{\eta}{\eta+1}\right)^{\alpha} \left(\frac{\varphi_{g0}}{\underline{w}_c}\right)^{\alpha} \right\} & \text{if } \overline{\varphi}_g > \varphi_{g0} > \underline{\varphi}_g, \\ \underline{w}_c^{\eta-1} \left(\frac{\varphi_{g0}}{\underline{w}_c}\right)^{\alpha} (\eta - \alpha) \left\{ 1 + \left(\frac{\eta}{\alpha-\eta}\right) \left(\frac{\eta}{\eta+1}\right)^{\alpha} \right\} & \text{if } \overline{\varphi}_g > \underline{\varphi}_g > \varphi_{g0}. \end{cases}$$

These equations imply that the effect of the wage floor on employment is hump-shaped, a result first derived by Ahlfeldt et al. (2022) in a more general model where there are also firms that can ration employment.

To compute the aggregate wage bill in g we need to solve $WB_g = M_g \int_{z^{\min}}^{\infty} w(z)\ell(z) f_g(z) dz$. The result is given by

$$WB_g = M_g W_r^{-\eta} A_{k1}^{\eta} \begin{cases} \varphi_{g0}^{\eta+1} \left(\frac{\eta}{\eta+1}\right)^{\eta+1} \left(\frac{\alpha}{\alpha-\eta-1}\right) & \text{if } \varphi_{g0} \geq \overline{\varphi}_g, \\ \underline{w}_c^{\eta+1} \left\{ 1 + \left(\frac{\eta+1}{\alpha-\eta-1}\right) \left(\frac{\eta}{\eta+1}\right)^{\alpha} \left(\frac{\varphi_{g0}}{\underline{w}_c}\right)^{\alpha} \right\} & \text{if } \overline{\varphi}_g > \varphi_{g0} \geq \underline{\varphi}_g, \\ \underline{w}_c^{\eta+1} \left(\frac{\varphi_{g0}}{\underline{w}_c}\right)^{\alpha} \left\{ 1 + \left(\frac{\eta+1}{\alpha-\eta-1}\right) \left(\frac{\eta}{\eta+1}\right)^{\alpha} \right\} & \text{if } \varphi_{g0} < \underline{\varphi}_g, \end{cases}$$

and the derivative of WB_g with respect to the wage floor, holding constant the wage index, is

$$\frac{\partial WB_g}{\partial \underline{w}_c} = M_g W_r^{-\eta} A_{k1}^{\eta} \begin{cases} 0 & \text{if } \varphi_{g0} \geq \overline{\varphi}_g, \\ \underline{w}_c^{\eta} (\eta+1) \left[1 - \left(\frac{\eta}{\eta+1}\right)^{\alpha} \left(\frac{\varphi_{g0}}{\underline{w}_c}\right)^{\alpha} \right] & \text{if } \overline{\varphi}_g > \varphi_{g0} \geq \underline{\varphi}_g, \\ \underline{w}_c^{\eta} \left(\frac{\varphi_{g0}}{\underline{w}_c}\right)^{\alpha} \alpha \left[\frac{\eta+1}{\alpha} \left(1 - \left(\frac{\eta}{\eta+1}\right)^{\alpha} \right) - 1 \right] & \text{if } \overline{\varphi}_g > \underline{\varphi}_g > \varphi_{g0}. \end{cases}$$

Finally, aggregate revenue can be obtained from $R_g = M_g \int_{z^{\min}}^{\infty} z\ell(z) f_g(z) dz$. The solution is given by

$$R_g = M_g W_r^{-\eta} A_{k1}^{\eta} \begin{cases} \varphi_{g0}^{\eta+1} \left(\frac{\eta}{\eta+1}\right)^{\eta} \left(\frac{\alpha}{\alpha-\eta-1}\right) & \text{if } \varphi_{g0} \geq \overline{\varphi}_g, \\ \underline{w}_c^{\eta} \varphi_{g0} \left(\frac{\alpha}{\alpha-1}\right) \left\{ 1 + \left(\frac{\eta}{\alpha-\eta-1}\right) \left(\frac{\eta}{\eta+1}\right)^{\alpha-1} \left(\frac{\varphi_{g0}}{\underline{w}_c}\right)^{\alpha-1} \right\} & \text{if } \overline{\varphi}_g > \varphi_{g0} \geq \underline{\varphi}_g, \\ \underline{w}_c^{\eta+1} \left(\frac{\varphi_{g0}}{\underline{w}_c}\right)^{\alpha} \left(\frac{\alpha}{\alpha-1}\right) \left\{ 1 + \left(\frac{\eta}{\alpha-\eta-1}\right) \left(\frac{\eta}{\eta+1}\right)^{\alpha-1} \right\} & \text{if } \overline{\varphi}_g > \underline{\varphi}_g > \varphi_{g0}. \end{cases} \tag{F.2}$$

Once again, the partial derivative with respect to the wage floor is

$$\frac{\partial R_g}{\partial \underline{w}_c} = M_g W_r^{-\eta} A_{k1}^{\eta} \begin{cases} 0 & \text{if } \varphi_{g0} \geq \overline{\varphi}_g, \\ \underline{w}_c^{\eta-1} \varphi_{g0} \eta \left(\frac{\alpha}{\alpha-1}\right) \left[1 - \left(\frac{\eta}{\eta+1}\right)^{\alpha-1} \left(\frac{\varphi_{g0}}{\underline{w}_c}\right)^{\alpha-1} \right] & \text{if } \overline{\varphi}_g > \varphi_{g0} \geq \underline{\varphi}_g, \\ \underline{w}_c^{\eta} \left(\frac{\varphi_{g0}}{\underline{w}_c}\right)^{\alpha} \alpha \left[\frac{\eta}{\alpha-1} \left(1 - \left(\frac{\eta}{\eta+1}\right)^{\alpha-1} \right) - 1 \right] & \text{if } \overline{\varphi}_g > \underline{\varphi}_g > \varphi_{g0}. \end{cases} \tag{F.3}$$

To solve the Nash bargaining problem, it is useful to compute the total derivative with respect to the wage floor. Let $X \in \{L, WB, R\}$ be any of the aggregate quantities defined above. Then, the total derivative is given by

$$\frac{dX_g}{d\underline{w}_c} = \frac{\partial X_g}{\partial \underline{w}_c} - \eta \frac{X_g}{W_r}\frac{dW_r}{d\underline{w}_c},$$

where $dW_r/d\underline{w}_c$ is the derivative of the wage index with respect to the wage floor.

F.2 Equilibrium of structural model

F.2.1 Definition of equilibrium

Given a CB network $\mathcal{C}$, values for $\{\{N_r\}_{r\in\mathcal{R}}, \{M_g\}_{g\in\mathcal{G}}\}$, parameters for the worker problem $\{\zeta, \{b_r\}_{r\in\mathcal{R}}, \eta, \{A_{k1}\}_{k1\in\mathcal{K}1}\}$, bargaining power parameters $\{\beta_c\}_{c\in\mathcal{C}}$, and productivity processes parameters $\{\alpha, \{\varphi_{g0}\}_{g\in\mathcal{G}}\}$, an equilibrium is defined as a vector of employment shares $\{\mu_r\}_{r\in\mathcal{R}}$, a set of wage floors $\{w_c^*\}_{c\in\mathcal{C}}$, and a set of wage indices $\{W_r^*\}_{r\in\mathcal{R}}$ such that:

1. Employment shares satisfy equation (8), with expected utility given by (7).

2. Firms at every φ choose wages given labor supply according to (9).

3. Wage floors simultaneously solve the Nash-in-Nash problem given by (13) for all $c \in \mathcal{C}$.

4. Wage indices clear the labor market in each region:

$$\sum_{g:r(g)=r} M_g \int_{\varphi_{g0}}^{\infty} \ell(\varphi) f_g(\varphi) d\varphi = \mu_r N_r. \tag{F.4}$$

Existence and uniqueness of equilibrium. Let $\hat{\Gamma} = \Gamma(\frac{\eta-1}{\eta})^{\zeta}$ and write $\hat{L}_r = \sum_{g:r(g)=r} \hat{L}_g$. Then, the market clearing condition can be written as

$$W_r^{-\eta}\hat{L}_r = \frac{\hat{\Gamma}W_r^{\zeta}}{\hat{\Gamma}W_r^{\zeta} + b_r} N_r. \tag{F.5}$$

This is a continuous function of W_r, and applying the intermediate value theorem we can conclude that a unique solution exists for any vector of wage floors.[13]

Then, it remains to discuss the equilibrium of the Nash-in-Nash problem. If all objective functions were concave, then the problem would be a concave maximization problem over a convex set, and thus a unique solution would exist. However, the objective functions may actually be decreasing at first, increasing later on, and then decreasing again. These "convex regions" would occur if, starting from a non-binding wage floor, an increase in the floor decreases the wage bill and thus the problem's objective function. This could be due to either general equilibrium effects or because some low productivity local labor markets are fully constrained by the wage floor and

[13]Define $h(W_r) = W_r^{-\eta}\hat{L}_r - \frac{\hat{\Gamma}W_r^{\zeta}}{\hat{\Gamma}W_r^{\zeta}+b_r}N_r$. Then, $W_r \to 0$ implies that $h(0) \to \infty$ and $W_r \to \infty$ implies that $h(\infty) \to -N_r$. Thus, by the intermediate value theorem, there exists a W_r^* such that $h(W_r^*) = 0$.

experience negative wage floor effects. These convex regions are problematic if they result in multiple maxima, which occurs in a knife-edge case where the level of utility with a non-binding wage floor is equal to the level of utility of a binding wage floor that is also a local maximum.

In summary, the Nash-in-Nash problem has a unique solution if all objective functions in the Nash-in-Nash problem are concave. In the case some objective function has a convex region, the problem has multiple equilibria only in a knife-edge case that is unlikely to occur in practice.

Effect of wage floor on regional wage index. A closed form solution to (F.5) does not exist. However, we can obtain an expression for the derivative of the wage index with respect to the wage floor using the implicit function theorem:

$$\frac{dW_r}{d\underline{w}_c} = \frac{\left(\hat{\Gamma} W_r^{\zeta} + b_r\right)}{\hat{\Gamma} W_r^{\zeta-1}\left[(\zeta+\eta) W_r^{\eta} N_r - \zeta \hat{L}_r\right]} \sum_{g:c(g)=c} \frac{\partial \hat{L}_g}{\partial \underline{w}_c}. \tag{F.6}$$

This expression will be positive as long as the summation is positive. I.e., a wage floor increase that increases employment will have a positive effect on the wage index.

F.2.2 An algorithm to compute the equilibrium

Given a set of parameters, I solve for equilibrium wage indices $\{W_r\}_{r\in\mathcal{R}}$ and wage floors $\{w_c\}_{c\in\mathcal{C}}$ using an algorithm that mixes fixed-point iteration and Gauss-Seidel's coordinate update algorithm (See Galichon 2022). The algorithm proceeds as follows:

1. Take an initial guess $\{\{W_r^0\}, \{\underline{w}_c^0\}\}$
2. Use fixed-point iteration to find equilibrium wage indices $\{W_r^1\}$ using the previous wage floors vector.
3. Iteratively find equilibrium wage floors $\{\underline{w}_c^1\}$ numerically solving the Nash bargaining split for each c, using the equilibrium wage indices $\{W_r^1\}$ and the previously found wage floors.
4. Take $\{\{W_r^1\}, \{\underline{w}_c^1\}\}$ as new starting point and repeat.
5. Iterate until convergence.

I set as starting point $W_r = .5$ for every r, and $\underline{w}_c = 1.05 \times \min_{c:g(c)=c}\{z_{g0}\}$ for every c. I set as criteria that the mean absolute difference between consecutive iterations is less than 10^{-8}, both for the fixed-point iteration and for the entire loop.

Finding wage indices. Using the four labor market-clearing equations it is straightforward to write $W_r = T(W_r)$, where $T(\cdot)$ is a contraction mapping. In particular, using (F.5), we can write

$$W_r = \left(\frac{\hat{\Gamma} W_r^{\zeta} + b_r}{\hat{\Gamma} W_r^{\zeta}} \frac{\hat{L}_r}{N_r}\right)^{1/\eta}.$$

Thus, to find wage indices I use a fixed-point iteration using the most up-to-date wage floors to compute aggregate labor demand.

Finding wage floors. In this step I use the Nash bargaining FOCs, defined in equation (2). However, as the FOC may not always hold in a maximum, I select several candidate points for maximizers and then evaluate the objective function to select the global maximum.

The first issue with the FOCs is that the derivatives of the wage bill and revenue are discontinuous when $\varphi_{g0} = \underline{w}_c$. Fortunately, this happens only in a small share sectors, usually belonging to the largest CB units. As a result, derivatives are generally well-defined in the feasible range where the wage floor will be set. However, for different counterfactuals this might not be the case. As a result, I add all points where $\varphi_{g0} = \underline{w}_c$ for some g as candidate points for maximizers.

Second, when there are multiple local labor markets in a CB unit, the objective functions are sums of terms, some of which do not depend on the wage floor (the unconstrained ones). Because derivatives are discontinuous, when a local labor market enters this sum the FOC may change discontinuously, generating the possibility of local optima. To deal with this, I search for all wage floors at which equation (2) holds and add them to the set of candidate maximizers.

Finally, I evaluate candidate points and pick the one that yields the global maximum. When evaluating the objective function I adjust wage indexes to account for general equilibrium effects.

F.3 Defining local labor markets and aggregating data

I construct the aggregate data from the universe of firms observed in 2011–2012. I construct local labor markets, and then compute aggregate data used for model estimation.

Defining local labor markets and CB units

The goal is to avoid small local labor markets as much as possible, since these add computation time to the model without significantly affecting the conclusions of the analysis. I start by dropping broad sectors that negotiate under a different bargaining regime, such as agriculture.[14] Then, I define exporter firms as those with any exports in 2011–2012, and I group provinces in 4 regions: Centro (center), Cuyo (west), Norte (north), and Patagonia (south).[15]

The next step is to create a coarse 4-digit sector variable. First, if an exporter, region, 4-digit sector and CB unit cell has less than 10 firms, then I assign the most common 4-digit sector in the CB unit to the cell.[16] I do so to keep firms in the correct CB unit, while simplifying the large heterogeneity observed within CB units. Second, I group some 4-digit sectors to broader categories to reduce the number of sectors. I try to keep the 4-digit resolution in Manufacturing,

[14] Specifically, I drop the broad categories, or 1-digit sectors, A (Agriculture, forestry and fishing), B (Mining and quarrying), P (Education), R (Arts, entertainment and recreation), T (Activities of households as employers), and U (Activities of extraterritorial organizations and bodies). I also drop the 4-digit sectors that correspond to waste management and employment agencies.

[15] The regions are (1) "Centro," which includes Ciudad de Buenos Aires, Buenos Aires, Córdoba, Entre Ríos, La Pampa, Mendoza, and Santa Fe; (2) "Cuyo," which includes La Rioja, Mendoza, San Juan, and San Luis; (3) "Norte," including Catamarca, Chaco, Corrientes, Formosa, Jujuy, Misiones, Salta, and Tucumán; and (4) "Patagonia," grouping Chubut, La Pampa, Neuquén, Río Negro, Santa Cruz, and Tierra del Fuego.

[16] Additionally, I define as non-exporter all firms in cells that have less than 2 employees, regardless of their actual exporting status.

since this is where most of the export shocks take place.[17] After these changes I still observe 10.6% of firms in coarse 4-digit sector by region cells with less than 5 firms. For these cells, I change the 4-digit sector code so that they are assigned to a related group.[18] This process results in a coarse 4-digit sector that has at least 5 firms in every category by region cell.

The final step before grouping the data is to define local labor markets and CB units. The baseline cell is the coarse 4-digit sector by region grouping created above. I further split each cell using provinces if the coarse 4-digit sector by region by province cell has more than 75 firms.[19] I also use a modified version of the observed CB code to split the cells. The reason to modify the CB code is, once again, to reduce the number of local labor markets and CB units in the economy. The key modification is to define "local" CB units if a given CB unit has low employment or a small share of firms observed with a wage floor.[20] Then, my local labor markets are defined by region, province, exporter status, coarse 4-digit sector, and CB unit. The CB network is simply defined by the modified CB unit codes.

Aggregating data

I group the firm-level data to the local labor market level and compute aggregate statistics. I compute the mean **wage** and **wage floor** in each local labor market by taking the mean of the firm-level average wage in each CB unit. I adjust the firm-level average wages by part-time work. To do so, I define a worker to be part-time if she earns less than 90% of the wage floor, and double their wage when computing the local labor market average. To obtain the **share of firms bunching**, I compute the average of a "firm buncher" indicator. There are some concerns of measurement error in wage floors, which complicate the definition of this indicator. First, wage floors are estimated from data, which means they are noisy.[21] Second, workers earning less than the wage floor are common. I define a worker to be a "buncher" (that is, a worker with a deviation

[17]To be precise, I take the following steps. First, if a firm has a 4-digit sector that corresponds to any broad category *excluding* "Manufacturing" (broad category C), "Wholesale and retail trade" (G), "Transportation and Storage" (H), and "Accommodation and food service activities" (I), I group the 4-digit sector code to the broad category level. Second, for any 4-digit sector *excluding* those in "Manufacturing" I group the 4-digit sector code to the 3-digit level. Third, I group to the 2-digit level all 4-digit sectors included in 2-digit sectors that have less than 50 firms and broad categories that have more than 99 firms. Similarly, I group to the 3-digit level all 4-digit sectors included in 3-digit categories with less than 50 firms, 2-digit categories with more than 49 firms, and broad categories with more than 99 firms.

[18]More specifically, I change the 4-digit code for cells that were not grouped so far and have less than 5 firms to the most common 4-digit sector in the corresponding 3-digit grouping. There still remain 5.9% of firms in cells with less than 5 firms. I group all these 4-digit sector cells to the broad category level.

[19]Only 1% of these cells have more than 75 firms.

[20]More precisely, I take the following steps to simplify the CB unit code. First, if a CB unit has less than 200 workers I assign it to the *local* category. The second step concerns CB units with more than 200 workers. If a local labor market is non-exporter, has less than 25 firms and 1000 employees, the share of employment with an assigned wage floor is less than 10%, and it corresponds to less than 10% of the employment in the CB unit, then I assign it to the *modal* CB unit in the region by coarse 4-digit sector cell. This introduces a small error in the CB network, but makes it more likely that these small local markets will get a valid wage floor estimate. Finally, if a CB unit has less than 2000 workers and the share of workers with a valid wage floor is less than 1%, then I assign the corresponding local markets to the "local" CB unit.

[21]Small CBAs and categories within CBAs with a few workers are both lost in the estimation step.

of 0 from the wage floor) if she is full-time and her wage is between 90 and 105% of the floor, or if she is part-time and her wage is between 40 and 60% of the floor. Then, I define a "firm buncher" as a firm with all workers being bunchers.[22] To compute the **number of firms** in each local labor market, I simply count the number of distinct firm IDs. To compute **employment** I sum all workers in each firm.

The mean wage floors computed above are noisy and vary across local labor markets within a CB unit. I re-define wage floors at the CB unit level, which is the level at which they are determined. I start by computing the average of the local market mean wage floors in each CB unit by region cell, weighting for the share of workers with a valid wage floor in the local labor market. If all regional wage floors are within 5% of each other, I take the average and use it to define the CB unit's wage floor. I allow regional wage floors within a CB unit only if the region's average wage floor is sufficiently different from the rest.[23] I consider regional differences in wage floors as constant when estimating the effects of shocks, and allow a single wage floor per CB unit to be determined in the negotiations. This aligns with usual practice in Argentina, where regional differences within CBAs are usually constant.

The estimated shares of firms bunching are also noisy, the reason being that many workers do not have a valid wage floor assigned.[24] As a result, I adjust the estimated share of firms bunching in several ways. First, I notice that exporting local labor markets systematically report a lower share of firms bunching, even after controlling for the mean wage floor and the mean wage. This is problematic because it would lead to an over-estimate of the productivity of these local labor markets, and since these local markets are the ones hit by shocks these could have consequences for the model simulations. To deal with this, I simply add a random number between 0 and 0.03 to the share. Second, I compute a GLM regression of the share bunching on local market covariates and use it to predict the share of firms bunching for "problematic" local labor markets.[25] It is important to note that the model used to assign shares for a given region is estimated excluding that region, so all assigned shares are out-of-sample predictions. The model is used to impute 15.6% of local labor markets that have less than 10% of workers with a valid wage floor or an implausible share given the ratio between the mean wage and mean wage floor. Finally, I shrink the estimated share of firms bunching using the James-Stein estimator towards the region by coarse 4-digit sector average.

[22] There are some workers earning less than 40% of the estimated wage floor. The average firm-level share of workers that earn less than 40% of the wage floor is approximately 0.13. I drop firms with more than 50% of workers earning less than 40% of the wage floor from the computation of the aggregate share of firms bunching.

[23] Specifically, if the difference between regions is greater than 5% and the maximum is in Patagonia, I assign the maximum to Patagonia and average the rest. I then implement a handful of manual adjustments.

[24] They may not get a valid wage floor if they are in small CB units or if their occupation code within the CB unit is not observed.

[25] The covariates of the model are the mean wage floor, the mean wage, a dummy for exporter status, the share of workers with a valid wage floor, the share of part time workers, and coarse 4-digit sector fixed effects. The model is estimated only on local labor markets with more than 10% of workers with a valid wage floor and with a share of firms bunching between .5% and 99.5%.

F.4 Estimating model parameters

Preference heterogeneity parameter η. To estimate η, I exploit the relationship between firm size ℓ and wages w implied by (6). To do so, I regress log wages on log employment at the firm level controlling for region and 1-digit sector fixed effects. However, the structure of labor supply implies an exact fit of this regression.

To allow for a non-perfect fit the model can be extended to incorporate hours. In particular, assume that production takes place before the firm draws a random value h such that realized employment is $\ell^* = h\ell$. In that case, the labor supply to the firm equation can be written as $\ell^*(j) = hW_r^{-\eta} A_{k1}^{\eta} w(j)^{\eta}$, and so by taking logs and rearranging terms I can write observed log wages as a function of observed log employment and $\ln h$ plays the role of the error term.

Specifically, I use the firm-level data and estimate the regression

$$\ln w_{jt} = \bar{\eta} \ln \ell_{jt} + b_{k1(j)} + \delta_{r(j)t} + \nu_{jt}.$$

where $k1(j)$ is the 1-digit sector of firm j and $\delta_{r(j)t}$ corresponds to year by region fixed effects. Under the assumption that unobservables (hours) are uncorrelated with firm size within regions and time periods, the coefficient $\bar{\eta}$ identifies the inverse of η.

Appendix Table 9 shows the estimates, which are stable when varying the set of fixed effects included in the regression. My preferred specification is column (3), which yields $\eta = 4.0995 \approx 1/0.2439$. The literature provides other values for this parameter. Monte et al. (2018) in the US use county-level data and estimate a value of 3.3. Ahlfeldt et al. (2022) for Germany use a firm-level dataset and find a value of 5.2. Parente (2022) uses a calibration approach in Brazil and finds values of 4.52 and 4.22 for 1996 and 2012, respectively. Datta (2023) estimates a labor supply elasticity of 4.6 in the UK using HR data from a multi-establishment firm.

I also rely on column (3) of Appendix Table 9 to obtain my estimates of the amenity values $\{A_{k1}\}_{k1\in\mathcal{K}1}$. The omitted category is chosen so that all b_{k1} are positive. Then, using the model structure I compute $A_{k1} = \exp(-b_{k1})$, which results in values ranging from 1 to 1.81. The amenity values allow firms to have different sizes conditional on the wage across 1-digit sectors.

Preference for formal employment ζ. From equation (8) the extensive-margin labor supply elasticity can be computed to be

$$\frac{d\mu_r}{dV_r}\frac{V_r}{\mu_r} = \zeta(1-\mu_r).$$

The elasticity is estimated to be around 0.2 in the literature (Chetty et al. 2011). Using the average share of formal employment from Appendix Table 10, I set $\zeta = 0.2813$. Differences in the shares of formal employment across regions will load on outside options.

Shape of productivity distributions α. I calibrate the shape of the productivity distributions α to 5.50. As discussed in Section 6.2.2, the maximum share of firms constrained by the wage floor that is consistent with the model is given by

$$S_g^{\max} = 1 - \left(\frac{\eta}{\eta+1}\right)^{\alpha}.$$

Inverting this expression for α and plugging-in the 98.8th percentile of the share of firms constrained by the wage floor for $S_g^{\max}$ and the value of η estimated above results in $\alpha \approx 5.50$.[26] For comparison, Parente (2022) uses a Pareto-LogNormal distribution where the shape of the Pareto is calibrated to 6.02 and 6.33 in two different periods.

Minimum productivities. I invert the share of firms constrained by the wage floor to obtain the minimum productivities $\{\varphi_{g0}\}_{g\in\mathcal{G}}$. First, if the share of firms constrained by the wage floor is zero, then I set $\varphi_{g0} = (\eta/(\eta+1))\underline{w}_c$. This assumes that the productivity is the minimum value that allows the local labor market to be unconstrained. Second, if the share of firms constrained is equal to the maximum possible share, then I set $\varphi_{g0} = \underline{w}_c$. This assumes that the productivity is the minimum value that makes the local labor market to be fully constrained. These two cases are rare, and happen only in a few small local labor markets. Finally, for all other observed values $S_g = s$, I invert equation (F.1) to get

$$\varphi_{g0} = \underline{w}_c \left(1-s\right)^{1/\alpha} \left(\frac{\eta}{\eta+1}\right).$$

For a given $\underline{w}_c$, a larger observed share s implies that the minimum productivity must be smaller.

Outside options b_r. To obtain outside options I use equation (8), which requires computing V_r. Furthermore, V_r requires knowledge of regional wage indexes W_r. So, I proceed as follows. First, noting that $L_g = W_r^{\eta}\hat{L}_g$, I compute wage indexes using the labor market clearing condition:

$$W_r = \left(\frac{\sum_g \hat{L}_g}{\mu_r N_r}\right)^{1/\eta},$$

where μ_r and N_r are obtained from household survey data. In particular, N_r is the number of private sector workers (formal or informal), and μ_r the share of formal private sector workers.[27] Then, equation (7) is used to obtain V_r. Finally, I compute b_r inverting equation (8). Appendix Table 10 shows the results.

Bargaining power parameters. I invert the closed form expression for the Nash bargaining solution of each CB unit to obtain the bargaining power parameters, which is given by (2). This condition will hold in any equilibrium with $\underline{w}_c < \min_{g:c(g)=c}\{\varphi_g\}$, i.e., when no market is maximally constrained by the wage floor. Importantly, the inversion relies on computing the equilibrium value of the derivative of the objective functions of the union and the employer

[26] Matching higher percentiles requires increasingly larger values of α, with $\alpha \to \infty$ if we want to match $S_g^{\max} = 1$.
[27] Formal workers are those that declare to contribute to the social pension system.

association with respect to the wage floor, both of which enter the bargaining weight. This condition may not hold if a CB unit's optimum is not found at an interior solution.

Let $\gamma_c = WB_c/R_c$ be the share of the wage bill of the union in local labor market, computed using equilibrium quantities. Then, the bargaining power parameters are computed as

$$\beta_c = \frac{\gamma_c}{\gamma_c + (1-\gamma_c)\left(-\frac{dU}{d\underline{w}_c}\Big/\frac{d\Pi}{d\underline{w}_c}\right)}.$$

Note that these derivatives are evaluated at the equilibrium wage floors and wage indexes, and include the general equilibrium term $dW_r/d\underline{w}_c$ from equation (F.6). Since γ_c is bound by construction in the model, the key variation used to obtain the bargaining parameters comes from the ratio of the derivatives of U and Π with respect to the wage floor. The distribution of the ratio of derivatives is shown in Panel (a) of Appendix Figure 14. Panel (b) of Appendix Figure 14 shows the estimated bargaining power parameters.

I check whether the inverted bargaining power parameters actually result in maximizers of the Nash bargaining objective functions. To do so I test whether the wage floor can be recovered from the model equilibrium using the bargaining power parameters. While I do recover the correct wage floor in most cases, for one important CB unit I do not: the retail CB unit 0130/75. For this CB unit the condition that delivers the bargaining power parameter is a local maximum, but for this bargaining power parameter the global maximum is a non-binding wage floor. To deal with this issue I keep fixed the wage floor of this CB unit in counterfactual exercises.

G Additional Tables and Figures

Appendix Table 1: Description of coverage of selected collective bargaining agreements in the textile industry

CBA	Spanish	English
0500/07	Obreros de la industria textil, son los ocupados en establecimientos cuya actividad principal comprenda procesos destinados a la confección de colchones, bolsas, tejer, lavar, clasificar, peinar, cardar, hilar, urdir, tramar, retrocer, estrusar, devanar, desfibrar, teñir, aprestar, texturizar, bordar, cortar, coser, atar, anudar, bobinar, planchar, estampar, terminar, o similares y que se lleve a cabo sobre cualquier tipo de fibras, sean naturales o manufacturadas, ya sea manualmente o mediante la utilización de maquinarias subordinadas al proceso industrial textil.	Workers in the textile industry are those engaged in establishments whose main activity involves processes intended for making mattresses, bags, knitting, washing, sorting, combing, carding, spinning, warping, weaving, twisting, extruding, winding, defibring, dyeing, finishing, texturing, embroidering, cutting, sewing, tying, knotting, spooling, ironing, stamping, completing, or similar activities, carried out on any type of fibers, whether natural or manufactured, either manually or through the use of machinery subordinate to the textile industrial process.
0501/07	Trabajadores, empleados, supervisores, encargados, mecánicos, personal auxiliar de ambos sexos, de administración, de comercialización y de fábrica únicamente de las empresas industriales de indumentaria y afines.	Workers, employees, supervisors, managers, mechanics, auxiliary staff of both genders, from administration, sales, and exclusively from clothing industrial companies and related fields.
0746/17	Todos los trabajadores de la industria de la confección de indumentaria y afines según se especifican en los respectivos capítulos de la misma, comprende también a las empresas que fabrican toldos en general y sus respectivos accesorios, en artículos con tela de lona, plástica, sintéticas y/o similares, empresas que confeccionan y arman colchones en general y con sus respectivos accesorios, con todo tipo de materiales. Están comprendidos también los lavaderos industriales de los procesos de producción, tanto internos como externos. También incluye las empresas que producen avíos y accesorios internos para todo tipo de prendas de vestir en general, cualquiera fuere el material empleado en su producción y/o elaboración.	All workers in the apparel manufacturing industry and related fields as specified in the respective [CBA] chapters. This also includes companies that manufacture awnings in general and their respective accessories, in articles made of canvas fabric, plastic, synthetic and/or similar materials. It also covers companies that make and assemble mattresses in general and their accessories, with all types of materials. Also included are industrial laundries involved in the production process, both internal and external. It also includes companies that produce fittings and internal accessories for all types of clothing in general, regardless of the material used in their production and/or manufacturing.

Notes: The figure shows the description of the areas of representation for three selected collective bargaining agreements (CBAs) in the textile industry. The three CBAs specify the entire country as their regional scope, and were signed by different unions. CBA 0500/07 was signed by *Asociación Obrera Textil*, CBA 0501/07 was signed by *Sindicato de Empleados Textiles de la Industria y Afines*, and CBA 0746/17 was signed by *Federación Obrera Nacional de la Industria del Vestido y Afines*. The description in Spanish is verbatim from the CBA, and the description in English is a translation obtained using the large language model GPT-4.

Appendix Table 2: Summary statistics of main estimating sample, firm cross-section

	N	Mean	Std. Dev	Min	Max
Unique 4d sector	222				
Unique 6d sector	467				
Firm shock (2013-12 vs. 2010-09)	7,972	0.4325	0.3661	-3.6234	3.2670
Pre firm shock (2010-09 vs. 2008-07)	7,972	0.1363	0.2951	-0.8630	1.3006
Average employment 2007-09	7,972	46.04	55.66	1.00	396.33
Indicator employment 2007-09 $\leq$10	7,972	0.2346	0.4238	0.0000	1.0000
Average monthly wage 2007-09 (2010 ARS)	7,972	3,593.46	2,503.61	428.75	57,295.78
Log mean value exported 2011-2012	7,972	11.2357	2.1860	6.0638	16.9391
Observed in survey data	7,972	0.2258	0.4181	0.0000	1.0000

Notes: Data are from the baseline sample of firms that exported in 2011–2012. The tables show summary statistics of the data used in the main difference-in-differences estimates.

Appendix Table 3: Summary statistics of main estimating sample, CB units cross-section

	N	Mean	Std. Dev	Min	Max
Shock 2013-2012 minus 2010-09	174	0.5305	0.7440	-1.5745	4.2882
Shock 2010-2009 minus 2007-08	174	0.2912	1.0960	-8.1718	4.2070
Num. firms 2011-12	174	2,975.14	20,682.06	6	263,629
Num. firms in baseline sample 2011-12	174	45.82	246.99	1	2,310
Share employment exporting firms 2011-12	174	0.4065	0.2834	0.0162	0.9837

Notes: Data are from a panel of firms that exported in 2011–2012. The tables show summary statistics of the data used in the main difference-in-differences estimates.

Appendix Table 4: Summary statistics of main estimating sample, firm panel

	N	Mean	Std. Dev	Min	Max
Year	86,238	2,011.94	3.15	2,007	2,017
CB shock 2013-2012 minus 2010-09	86,238	0.4739	0.3780	-1.5745	4.2882
Firm shock 2013-2012 minus 2010-09	86,238	0.4328	0.3660	-3.6234	3.2670
Log average monthly wage	50,703	7.8383	0.1942	6.5794	9.6323
Log average monthly wage floor	85,777	3.2267	1.1976	0.0000	8.5142
Log employment	85,564	3.1432	1.1880	0.0000	8.4933
Share main hiring modality	86,238	0.6871	0.2839	0.0000	1.0000
Indicator active firm	87,692	0.9834	0.1277	0.0000	1.0000

Notes: Data are from a panel of firms that exported in 2011–2012. The tables show summary statistics of the data used in the main difference-in-differences estimates.

Appendix Table 5: Static difference-in-differences estimates, worker-level estimates

	Log mean wage				
	(1)	(2)	(3)	(4)	(5)
CB shock × Post	0.0165	0.0108	0.0189		0.0192
	(0.0058)	(0.0067)	(0.0101)		(0.0066)
CB shock × Post × Main CB				0.0209	
				(0.0108)	
CB shock × Post × Secondary CB				0.0060	
				(0.0132)	
Firm shock	Y	Y	Y	Y	N
Worker-firm FE	Y	Y	Y	Y	N
Firm-year FE	N	N	N	N	Y
Worker FE	N	N	N	N	Y
2d sector-province-year FE	Y	N	N	N	N
4d sector-province-year FE	N	Y	N	N	N
6d sector-province-year FE	N	N	Y	Y	N
2d sector-secondary CB-year FE	Y	Y	Y	Y	N
Hiring modality-year FE	Y	Y	Y	Y	N
Num. fixed effects	124,661	128,979	132,666	132,666	186,814
Num. observations	793,751	793,751	793,751	793,751	850,164
Adjusted R^2	0.8727	0.8795	0.8846	0.8846	0.9102

Notes: Data are a panel of workers that worked in exporting firms in 2011–2012 in 2008, 2011, and 2014 The table shows estimates of the effect of CB shocks on the log mean real monthly wage. Columns (1) to (3) show estimates a difference-in-differences model using the main CB shock variable in the firm as treatment. Column (4) replicates column (3), but interacts the primary CB shock with an indicator for whether the worker's CB unit is the primary CB unit in the firm. Column (5) shows estimates of a difference-in-differences model using the CB unit of the worker to define the treatment. The "2d sector-secondary CB-year" fixed effects include interactions between 2-digit sector, an indicator for whether the worker's CB is the primary CB in the firm, and year indicators. The "Hiring modality-year" fixed effects include interactions between all possible hiring modality indicators as of 2008 with year indicators. Standard errors are clustered at the CB unit level.

Appendix Table 6: Static difference-in-differences estimates, robustness to local labor market controls

	Log mean wage			Log mean wage floor			Log employment		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
CB shock × Post	0.0458	0.0338	0.0466	0.0541	0.0560	0.0526	-0.0233	-0.0174	0.0022
	(0.0202)	(0.0236)	(0.0172)	(0.0213)	(0.0239)	(0.0199)	(0.0351)	(0.0372)	(0.0305)
Firm shock × Post	0.0134	0.0106	0.0119	-0.0036	-0.0035	-0.0031	0.0297	0.0211	0.0257
	(0.0070)	(0.0063)	(0.0064)	(0.0021)	(0.0025)	(0.0018)	(0.0210)	(0.0175)	(0.0194)
6d sector shock × Post			0.0144			-0.0013			-0.0229
			(0.0064)			(0.0033)			(0.0138)
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y	Y
Firm controls	Y	Y	Y	Y	Y	Y	Y	Y	Y
2d sector by province by year FE	N	N	Y	N	N	Y	N	N	Y
4d sector by province by year FE	Y	N	N	Y	N	N	Y	N	N
6d sector by province by year FE	N	Y	N	N	Y	N	N	Y	N
Num. firms	7,972	7,972	7,953	7,654	7,654	7,636	7,972	7,972	7,953
Num. fixed effects	27,976	33,671	21,289	19,860	23,252	15,808	27,976	33,671	21,289
Num. observations	85,777	85,777	85,634	50,703	50,703	50,644	85,777	85,777	85,634
Adjusted R^2	0.8480	0.8499	0.8540	0.9253	0.9267	0.9282	0.8965	0.8992	0.9002

Notes: Data are from the baseline sample of exporting firms. The figure show regression coefficients on the firm, CB, and 6-digit sector shocks variables interacted with an indicator for year greater than or equal to 2012 ("Post"). The regressions are analogous to the ones in Table 1, but changing the approach to control for local labor market effects. The economic sector categories are constructed from a granular 6-digit economic sector variable, which was created by the tax authority of Argentina and is based on ISIC, rev 4. Standard errors are clustered at the CB unit level.

Appendix Table 7: Static difference-in-differences estimates, robustness to model specification

	Log mean wage				Log employment			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
CB shock × Post	0.0458	0.0366	0.0458	0.0484	-0.0233	-0.0119	-0.0243	-0.0232
	(0.0202)	(0.0164)	(0.0202)	(0.0221)	(0.0351)	(0.0272)	(0.0354)	(0.0351)
Firm shock × Post	0.0134	0.0141	0.0134	0.0128	0.0297	0.0309	0.0298	0.0297
	(0.0070)	(0.0084)	(0.0070)	(0.0070)	(0.0210)	(0.0165)	(0.0210)	(0.0208)
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y
Firm controls	Y	N	Y	Y	Y	N	Y	Y
Pre-period firm shock	N	N	Y	N	N	N	Y	N
Pre-period CB shock	Y	Y	Y	N	Y	Y	Y	N
Local labor market by year FE	Y	Y	Y	Y	Y	Y	Y	Y
Num. firms	7,972	7,972	7,972	7,972	7,972	7,972	7,972	7,972
Num. fixed effects	27,976	18,530	27,976	27,965	27,976	18,530	27,976	27,965
Num. observations	85,777	85,777	85,777	85,777	85,777	85,777	85,777	85,777
Adjusted R^2	0.8480	0.8491	0.8479	0.8478	0.8965	0.8984	0.8965	0.8965

Notes: Data are from the baseline sample of exporting firms. The table show regression coefficients on the firm and CB shocks variables interacted with an indicator for year greater than or equal to 2012 ("Post"), varying the set of included controls. All outcomes are computed using the full set of workers in the firm-year. The regressions are analogous to the ones in Table 1, but changing the set of included controls. Standard errors are clustered at the CB level for the CB shock variable, and at the firm level for the firm shock variable.

Appendix Table 8: Static difference-in-differences estimates, robustness to inclusion of CB units

	Log mean wage				Log employment			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
CB shock × Post	0.0458	0.0463	0.0687	0.0265	-0.0233	-0.0215	-0.0490	-0.0316
	(0.0202)	(0.0202)	(0.0226)	(0.0174)	(0.0351)	(0.0349)	(0.0414)	(0.0428)
Firm shock × Post	0.0134	0.0129	0.0131	0.0170	0.0297	0.0279	0.0283	0.0098
	(0.0070)	(0.0069)	(0.0070)	(0.0092)	(0.0210)	(0.0214)	(0.0214)	(0.0268)
Excluded CB units (n. exp. firms)	< 5	None	< 30	< 5	< 5	None	< 30	< 5
Include 0130/75 (retail CB unit)	Y	Y	Y	N	Y	Y	Y	N
Firm FE	Y	Y	Y	Y	Y	Y	Y	Y
Firm controls	Y	Y	Y	Y	Y	Y	Y	Y
Pre-period CB shock	Y	Y	Y	Y	Y	Y	Y	Y
Local labor market by year FE	Y	Y	Y	Y	Y	Y	Y	Y
Num. firms	7,972	8,019	7,754	5,733	7,972	8,019	7,754	5,733
Num. fixed effects	27,976	28,083	27,412	22,571	27,976	28,083	27,412	22,571
Num. observations	85,777	86,273	83,441	61,741	85,777	86,273	83,441	61,741
Adjusted R^2	0.8480	0.8485	0.8488	0.8504	0.8965	0.8966	0.8962	0.8969

Notes: Data are from the baseline sample of exporting firms. The table show regression coefficients on the firm and CB shocks variables interacted with an indicator for year greater than or equal to 2012 ("Post"). The regressions are analogous to the ones in Table 1, but changing the sample of CB units included in the regression. "n. exp. firms" refers to the number of exporting firms in CB unit. Standard errors are clustered at the CB level for the CB shock variable, and at the firm level for the firm shock variable.

Appendix Table 9: Estimates of preference heterogeneity

	Log average monthly pay				
	(1)	(2)	(3)	(4)	(5)
Log employment	0.2372	0.2368	0.2439	0.2307	0.2351
	(0.0108)	(0.0109)	(0.0065)	(0.0061)	(0.0065)
Year	Y	N	N	N	N
Year-Region FE	N	Y	Y	N	N
1d sector FE	N	N	Y	N	N
Year-Region-CBA FE	N	N	N	Y	N
Year-Region-CBA-6d sector FE	N	N	N	N	Y
Observations	1,243,640	1,243,640	1,240,627	1,240,627	1,240,627

Notes: Data are from labor market administrative records from Argentina. The sample is a panel of firm-years between 2012–2017 with 99% or more of their workers with a declared CBA code, and no workers with a declared wage below 40% of the wage floor. The table shows estimates of preference heterogeneity $1/\eta$ in the theoretical model. The dependent variable in all models is the log average monthly pay. The independent variable is the log of employment, computed in "full-time equivalent" workers by weighting part-time workers by 1/2. Columns show estimates changing the value of the fixed effects. Standard errors are clustered at the province by CBA level.

Appendix Table 10: Region-level data and 2011–2012 model estimates

Region	Centro (center)	Cuyo (west)	Norte (north)	Patagonia (south)
Private-sector workforce	5,099,128	399,917	837,599	204,902
Formal workforce	2,588,641	177,622	280,195	123,478
Share formal	0.5077	0.4441	0.3345	0.6026
W_r	0.8690	0.8530	0.8548	1.2763
V_r	1.0580	1.0384	1.0406	1.5538
b_r	0.9853	1.2648	2.0118	0.7465

Notes: Data are from the national household survey from INDEC and estimates from the structural model.

Appendix Table 11: Summary statistics of local labor markets

Variable	N	Mean	Std. Dev.	Min	Max
Unique regions	4				
Unique sectors	681				
Unique non-local CB units	438				
Unique local CB units	456				
Indicator region "Centro"	3,701	0.6347	0.4816	0.0000	1.0000
Indicator local CB unit	3,701	0.1359	0.3427	0.0000	1.0000
Indicator retail CB 0130/75	3,701	0.1489	0.3560	0.0000	1.0000
M_g	3,701	1.0000	4.7276	0.0070	171.2625
Mean wage (2010 ARS, 000s)	3,701	2.8210	1.6467	0.1842	28.2819
Mean wage adj. part-time (2010 ARS, 000s)	3,593	3.4498	1.5810	1.3225	28.8316
Mean wage floor (2010 ARS, 000s)	3,701	2.6216	0.8856	1.1674	13.3039
Share of firms bunching	3,701	0.2258	0.1696	0.0007	0.9962
Estimated minimum productivity	3,701	3.1008	1.0562	1.4388	15.0234

Notes: Data are from local labor market aggregates constructed using the administrative labor market data of Argentina. The table shows summary statistics of local labor markets in 2011–2012.

Appendix Table 12: Effects of export shocks on aggregate revenue, survey sample

	Change aggregate revenue			
	(1)	(2)	(3)	(4)
Sectoral shock by Exporter	0.2238	0.1436	0.2475	0.2223
	(0.1192)	(0.0697)	(0.1381)	(0.1188)
Sectoral shock by Non-exporter	0.0849	0.0527	0.0689	0.0713
	(0.0564)	(0.0508)	(0.0577)	(0.0533)
Excluded CBA	None	Metal	Retail	Plastic
Exporter-specific intercept	Yes	Yes	Yes	Yes
R2	0.0510	0.0391	0.0524	0.0547
Observations	154	110	147	149
Number of firms	4,840	3,268	4,405	4,513

Notes: Data are from a sample of firms surveyed in the first wave of the *Encuesta Nacional de Dinámica del Empleo y la Innovación* (ENDEI). The table show estimates of export shocks on aggregate revenue at the 4-digit economic sector level on exporters and non-exporters. The baseline sample includes economic sectors with at least 4 exporting firms and 6 total firms. The sectoral shock is computed as follows. First, I compute the change in average log world import demand at the firm level between 2010 and 2012. Second, I define the sectoral shock as the average firm-level shock, weighting by employment in 2010, for all exporting firms. The change in aggregate revenue is computed as the log difference in the sum of revenue declared by surveyed firms in each exporter status by 4-digit economic sector cell. The Metal CBA has code 0260/75, the Retail CBA has code 0130/75, and the Plastic CBA has code 0419/05. The row "Number of firms" reports the number of firms used to compute aggregate revenue. Standard errors are clustered at the 4-digit sector level.

Appendix Figure 1: Illustration of heterogeneity in coverage of CB units, textile industries

(a) Weaving of textiles

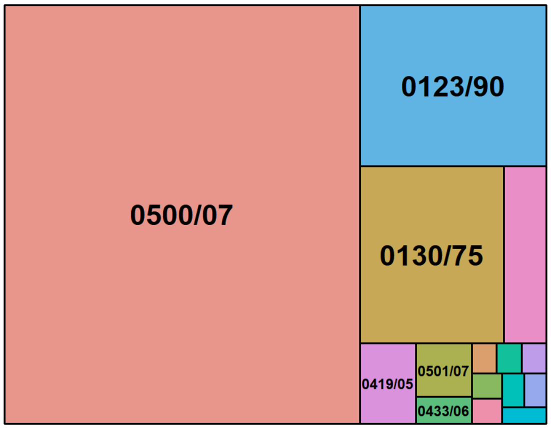

(b) Manuf. of textiles (ex. apparel)

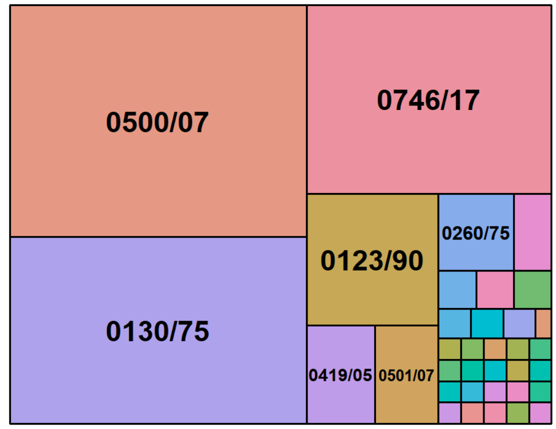

(c) Manuf. of wearing apparel (ex. fur)

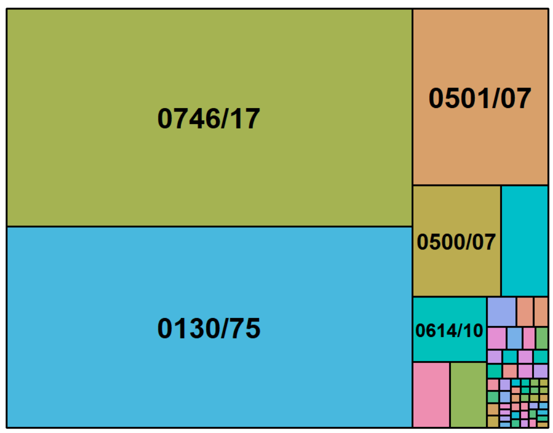

Notes: Data include all firms that had positive employment in 2012. Each box represents the sum of firms in the given economic sector, and the rectangles within each box represent the number of firms in each CB unit. The figure is based on the CB unit code of each firm obtained by the procedure described in Appendix B.2.3.

Appendix Figure 2: Distribution of wage to wage floor ratio, worker-level data in 2012

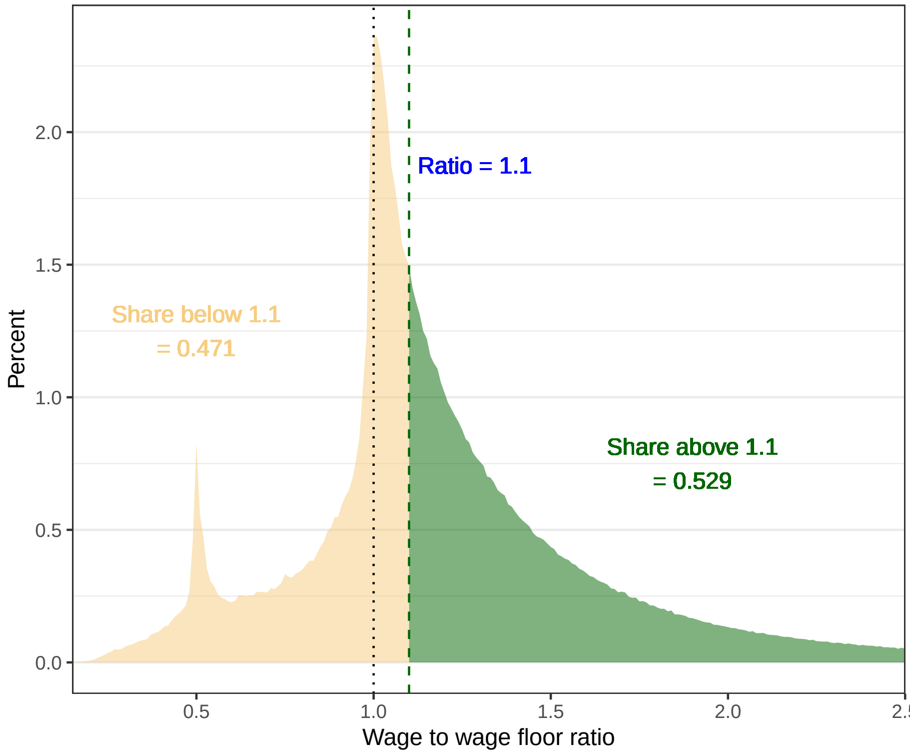

Notes: Data include all workers with a valid wage floor in 2012 declared in the main hiring modality. The figure shows a histogram of the ratio of the worker's wage to the wage floor that applies to the respective worker (excluding tenure-based compensation). The color fill varies for workers earning more or less than 1.1 times the wage floor. The dotted black line marks the ratio of 1. The green dashed line marks the 1.1 threshold that determines the color of the distribution. The blue dashed line marks the average wage cushion.

Appendix Figure 3: Distribution of export shocks

(a) Shocks to firms

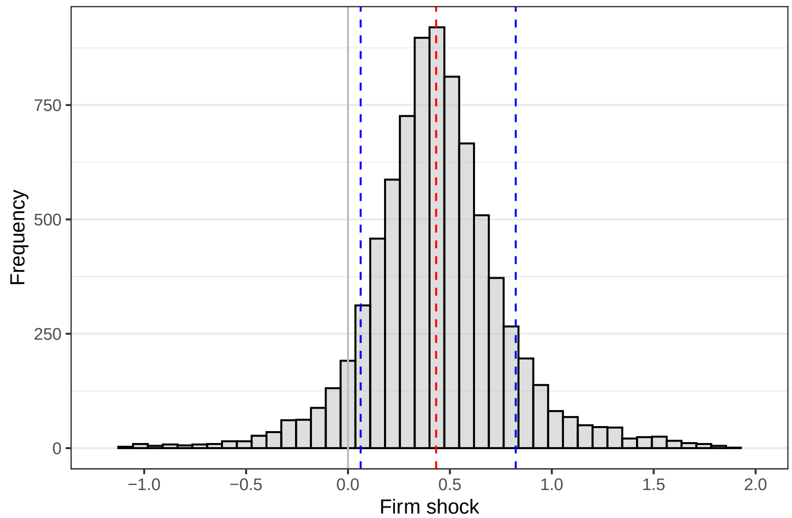

(b) Shocks to CB units

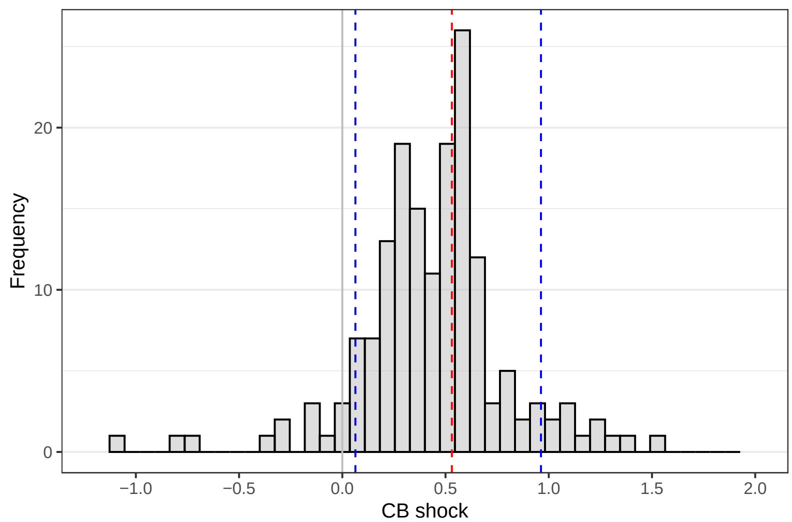

Notes: Data are constructed from a panel of firms that exported in 2011–2012. The figure illustrates the histogram of the firm and CB unit shocks. The figure excludes values lower than -1 and larger than 2 to increase visual clarity. The red dotted line in the middle shows the average shock, whereas the blue dotted lines on the sides show the 10th and 90th percentiles, respectively.

Appendix Figure 4: Correlation of export shocks to firms with export shocks to CB units

(a) Raw data

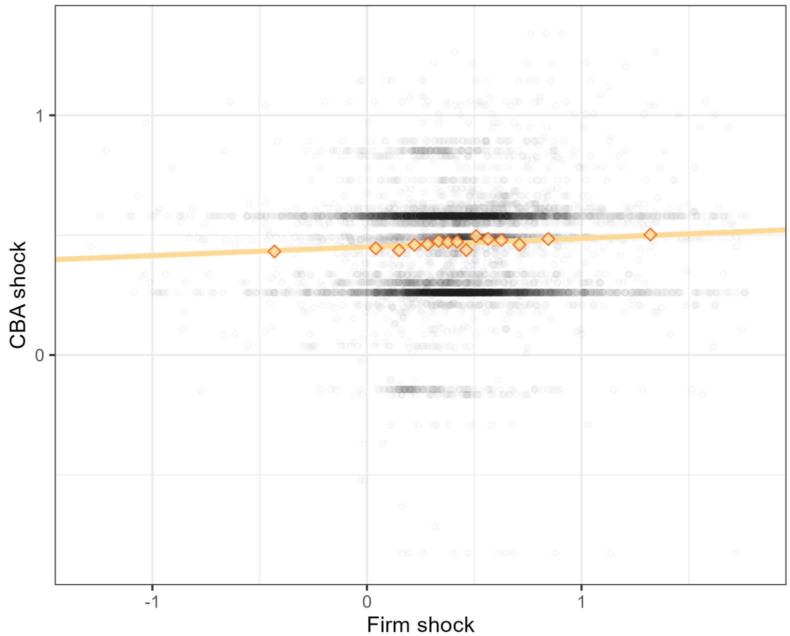

(b) Residualized on local labor market effects

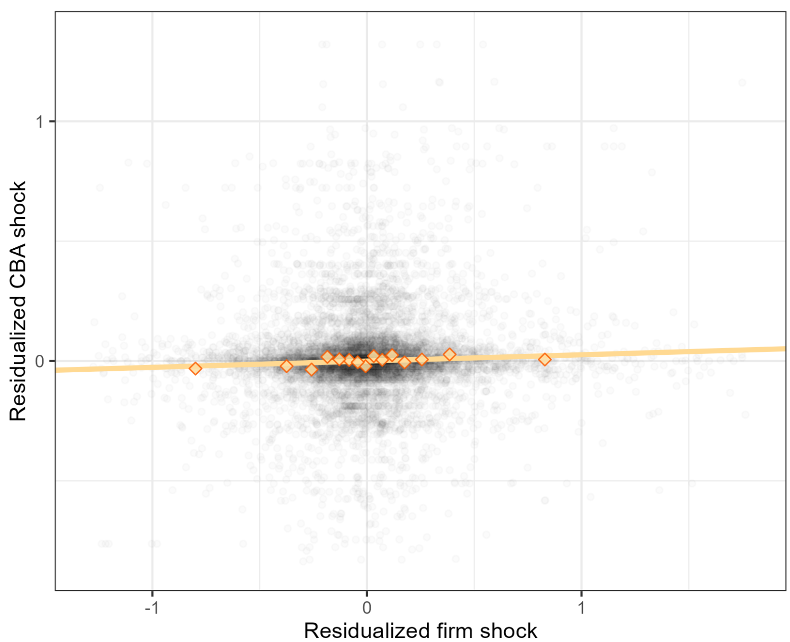

Notes: Data are constructed from a panel of firms that exported in 2011–2012. Panel (a) shows the correlation of firm shocks with CB shocks. Panel (b) shows the same correlation after residualizing for 4-digit sector by province fixed effects. The firm shock is defined as the difference between the 2012–2013 average to the 2009–2010 average in the value-weighted average world import demand for the firm, The CB shock is defined as the difference between the 2012–2013 average to the 2009–2010 average in the employment-weighted firm shock. To increase visual clarity Panel (a) excludes values lower than -1 and larger than 2 of both shocks, and Panel (b) excludes values lower than -1.5 and larger than 1.5 of both shocks. The blue line in both plots shows a non-parametric best fit.

Appendix Figure 5: Conditional correlation of firm shocks with baseline outcomes, firm-level cross-section

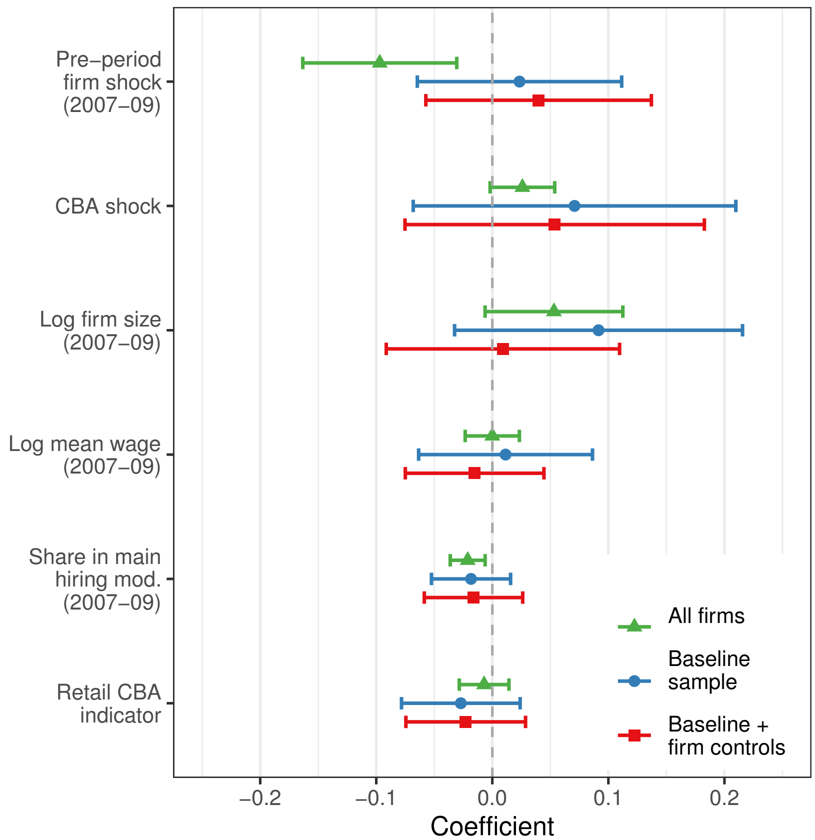

Notes: Data are constructed from a cross-section of firms that exported in 2011–2012. The figure shows the estimated coefficient on the firm shock on a regression of the outcome variable on the firm shock, controlling for 4-digit sector by providence fixed effects. Panel (a) shows the correlation of firm shocks, defined as the difference between the 2012–2013 average to the 2009–2010 average in the value-weighted average world import demand for the firm, with firm characteristics. Panel (b) shows the correlation of CB shocks, defined as the difference between the 2012–2013 average to the 2009–2010 average in the employment-weighted firm shock, with CB characteristics. The figure excludes values lower than -1 and larger than 2 of the CB shock to increase visual clarity. Standard errors are clustered at the CB level for the CB shock variables.

Appendix Figure 6: Auto correlation of export shocks to firms

(a) All firms

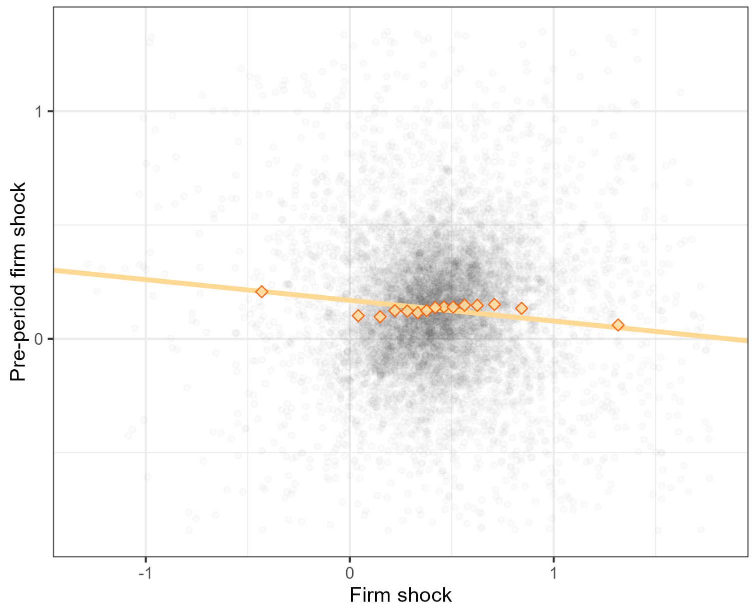

(b) Excluding top and bottom 1%

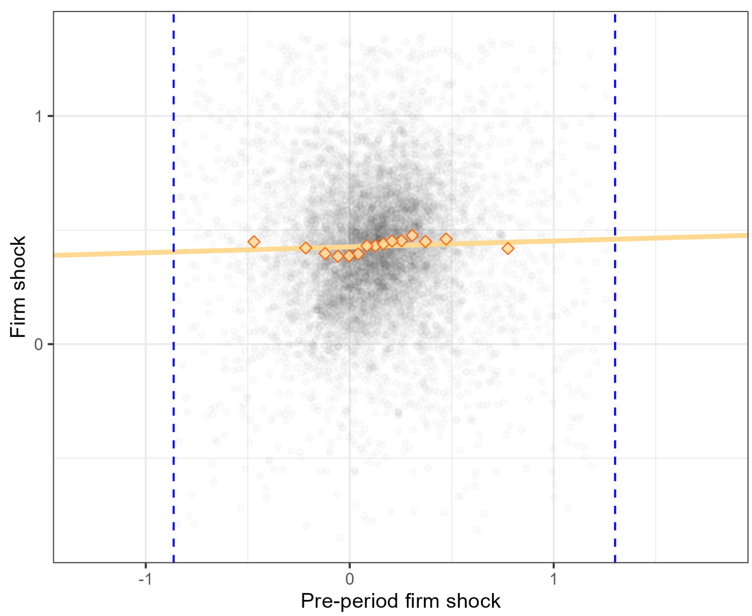

Notes: Data are constructed from a panel of firms that exported in 2011–2012. The figure shows the correlation of firm shocks, defined as the difference between the 2012–2013 average to the 2009–2010 average in the value-weighted average world import demand for the firm, with an analogous firm shock computed as the difference between 2009–2010 and 2007–2008. Panel (a) uses all exporting firms, while Panel (b) excludes firms that are in the top or bottom 1% of the pre-period distribution of firm shocks. The orange diamonds show the average firm shock within 15 bins of the pre-period firm shock. The orange line shows a linear fit to the data. The blue dashed lines indicate the 1st and 99th percentiles of the pre-period firm shock.

Appendix Figure 7: Correlates of export shocks to CB units with observables

(a) Pre-period CBA shock

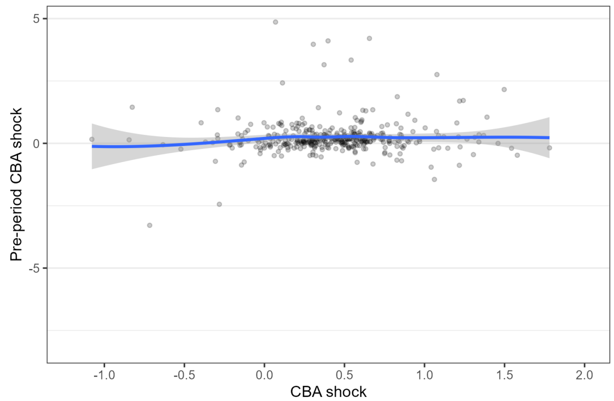

(b) Log number of firms

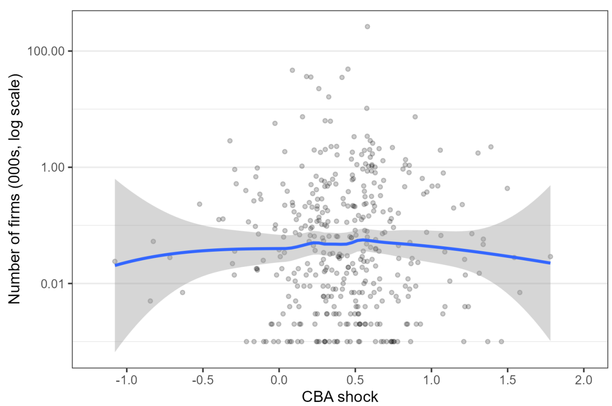

(c) Share employment in exporting firms

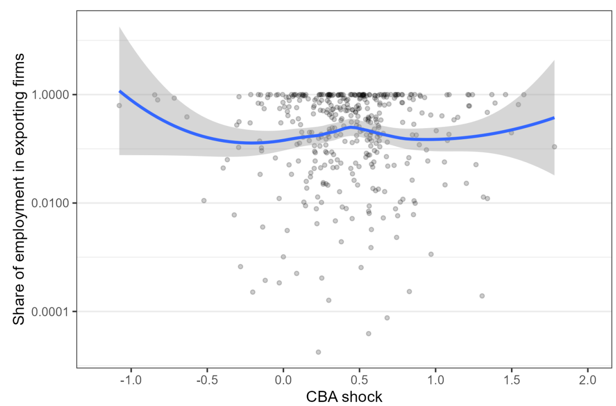

Notes: Data are constructed from a panel of firms that exported in 2011–2012. The figure shows a scatter plot of the CB unit shock with CB unit-level observables. Panel (a) uses the pre-period CB shock, constructed by differencing the proxy z_{jt} between 2008–2009 and 2008–2007. Panel (b) uses the number of firms. Panel (c) uses the share of employment in exporting firms. The line represents a locally estimated scatterplot smoothing curve fitted to the data.

Appendix Figure 8: Evolution of employment by level of CB unit and firm shock, baseline sample

(a) Shocks to firms

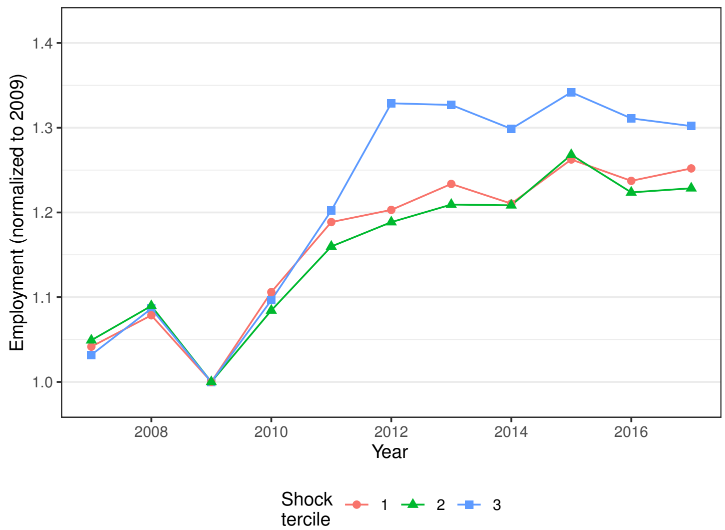

(b) Shocks to CB units

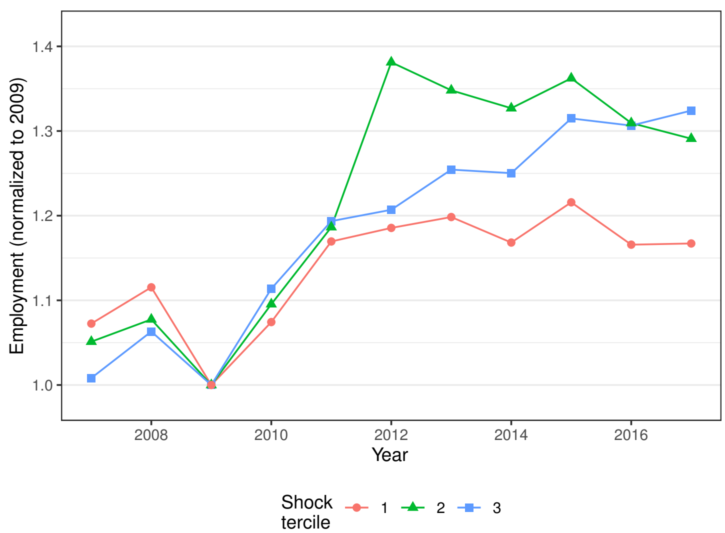

Notes: Data are from the baseline sample of exporting firms. The figure shows the average evolution of employment for firms in different terciles of the distribution of the firm shock (Panel a) and the CB shock (Panel b), relative to 2009. The firm and CB shocks are defined as the average changes in world import demand between 2009–2010 and 2012–2013, weighting by appropriate exposure shares.

Appendix Figure 9: Effect of export shocks to CB units, heterogeneity analysis

(a) Wage outcomes

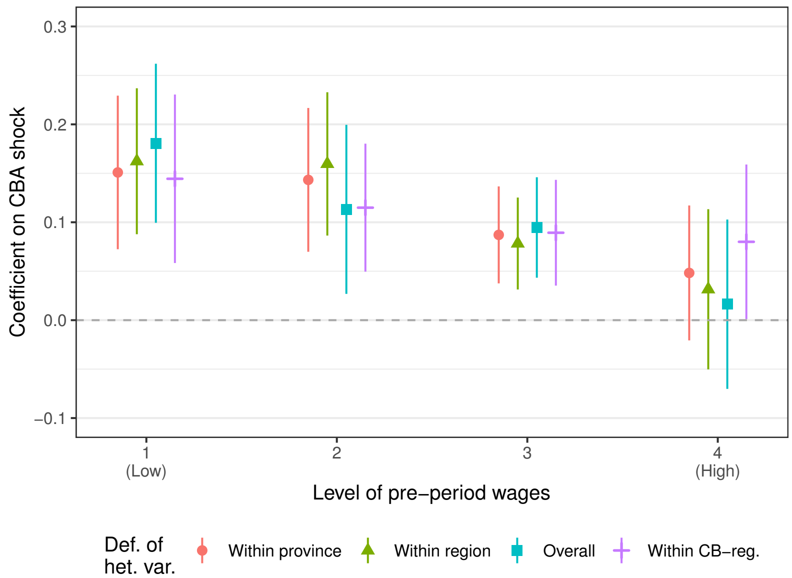

(b) Employment outcomes

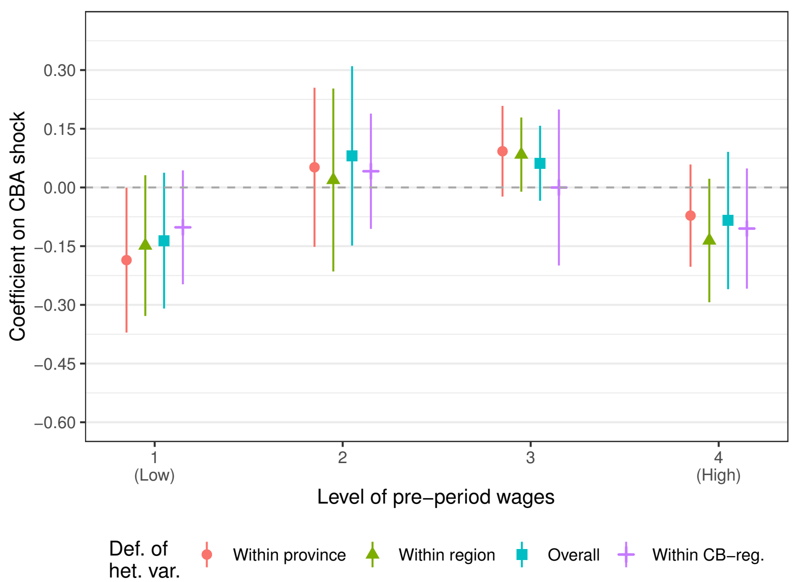

Notes: Data are from the baseline sample of exporting firms, excluding CB unit by province cells with less than 6 firms. The figure shows estimates of the effects of CBA shocks on a given outcome interacted with dummies for different quartiles of the pre-period level of wages (2007–2009), defined in within province, within region, within CB unit by region, and overall. The top figure shows the effect of the place CBA shock on log mean wage and log mean wage floor, and the bottom figure shows the effect on log employment and the share of workers hired in the main hiring modality in the firm. The regression model follows 2. Standard errors are clustered at the CB unit level.

Appendix Figure 10: Effect of export shocks to CB units on non-exporting firms

(a) Log mean wage

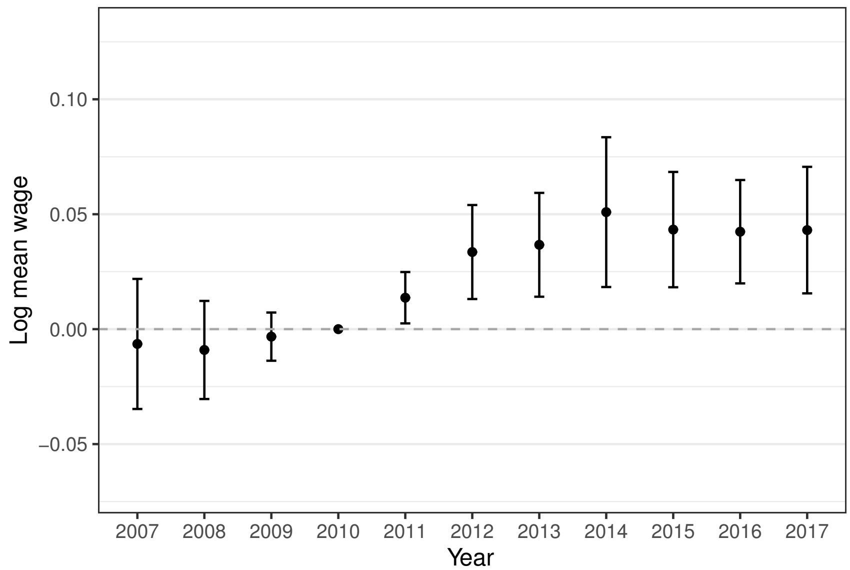

(b) Log mean wage floor

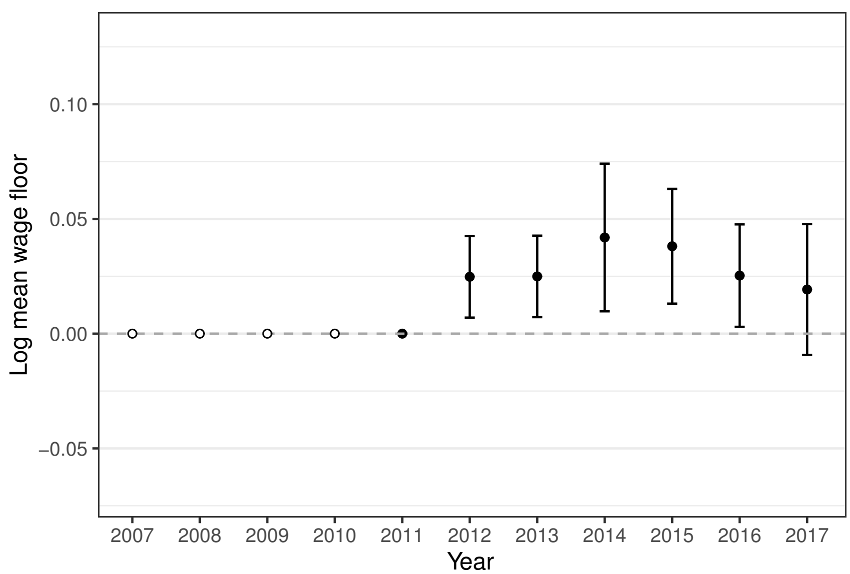

Notes: Data are from a panel of non-exporting firms that are covered by an exporting CB unit in 2011–2012. The figure shows the dynamic the effects of CB shocks on mean wages and mean wage floors. The regression omits the year 2010 for the wage variable, and the year 2011 for the wage floor variable. The sample includes firms smaller than 500 employees at baseline, firms in CB units with more than 30 firms in 2011, and firms that were active in 2007 and 2009. The regression includes firm fixed effects, 4-digit economic sector by exporter status by province by year fixed effects, time-varying firm controls, and a similar CB shock for the pre-period interacted with year dummies. Firm controls consist of a firm size indicator (categories 1-9, 10-24, 25-99, and 100-500) interacted with 2-digit sector and year and the pre-period share of workers in the main hiring modality interacted with year. The firm and CB shocks are defined as the average changes in world import demand between 2009–2010 and 2012–2013, weighting by appropriate exposure shares. Standard errors are clustered at the CB unit level.

Appendix Figure 11: Effect of export shocks to CB units, placebo network exercise

(a) Log mean wage

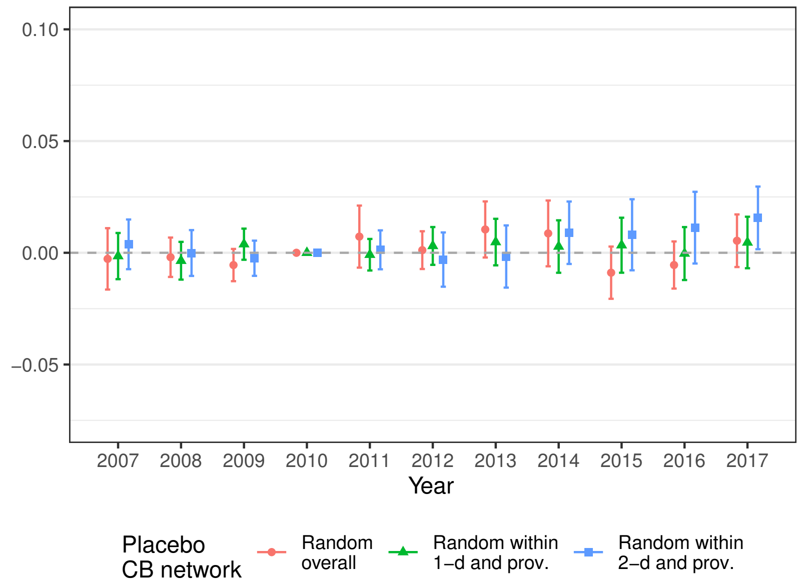

(b) Log mean wage floor

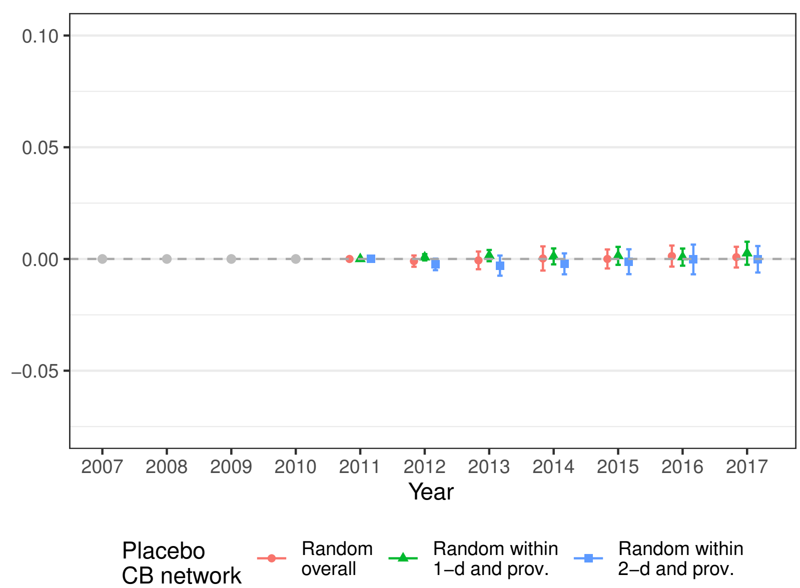

Notes: Data are from the baseline sample of exporting firms. The figure shows estimates of placebo shocks to CB units. The placebo network was constructed by randomly shifting the CB agreement code across all firms, across firms within 1-digit sector and province, and within 2-digit sector and province. The top figure shows the effect of the place CB shock on log mean wages, and the bottom figure shows the effect on the log mean wage floor. The regression includes the firm shock, firm fixed effects, 4-digit economic sector by province by year fixed effects, time-varying firm controls, and a similar CB shock for the pre-period interacted with year dummies. Firm controls consist of a firm size indicator (categories 1-9, 10-24, 25-99, and 100-500) interacted with 2-digit sector and year and the pre-period share of workers in the main hiring modality interacted with year. Standard errors are clustered at the CB unit level.

Appendix Figure 12: Effect of export shocks on wages, not excluding firms with extreme values of the pre-period shock

(a) Log mean wages

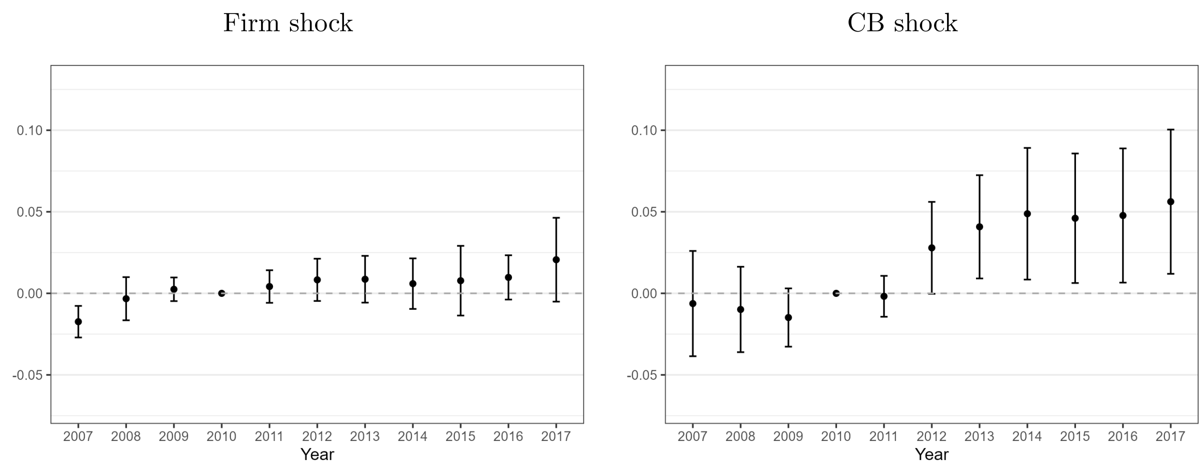

(b) Log employment

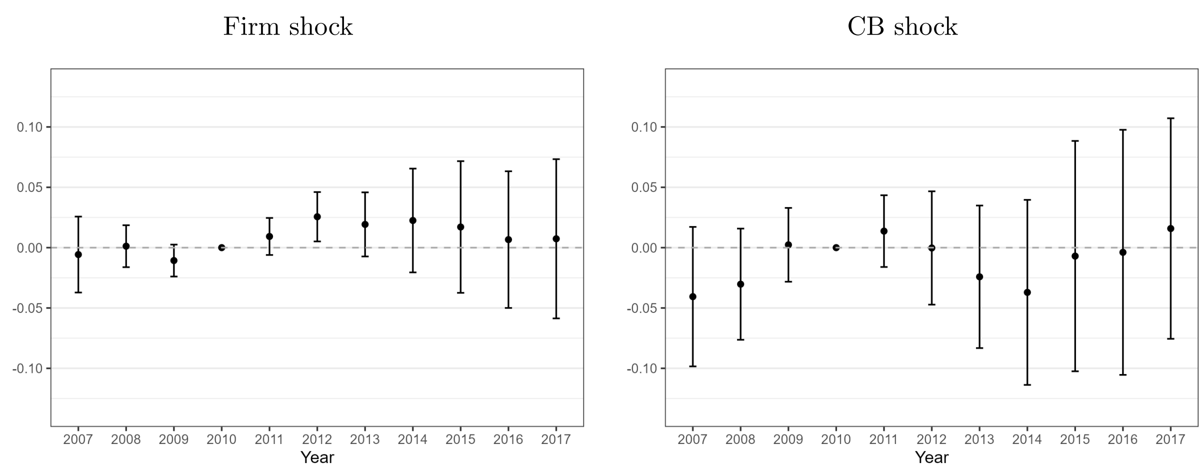

Notes: Data are from a panel of firms that are covered by an exporting CB in 2011–2012. The figure shows estimates of the effects of firm and CB shocks on log mean wages in Panel (a) and log employment in Panel (b). Estimation is done using a difference-in-differences strategy as in Figure 5 but keeping firms with extreme values of the pre-period shock. Standard errors are clustered at the CB unit level.

Appendix Figure 13: Illustration of firms productivity distribution

(a) Share of firms bunching

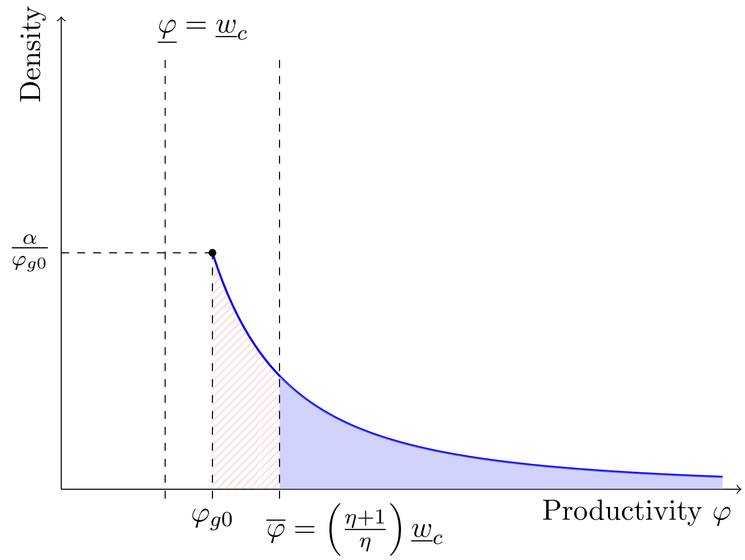

(b) Wages

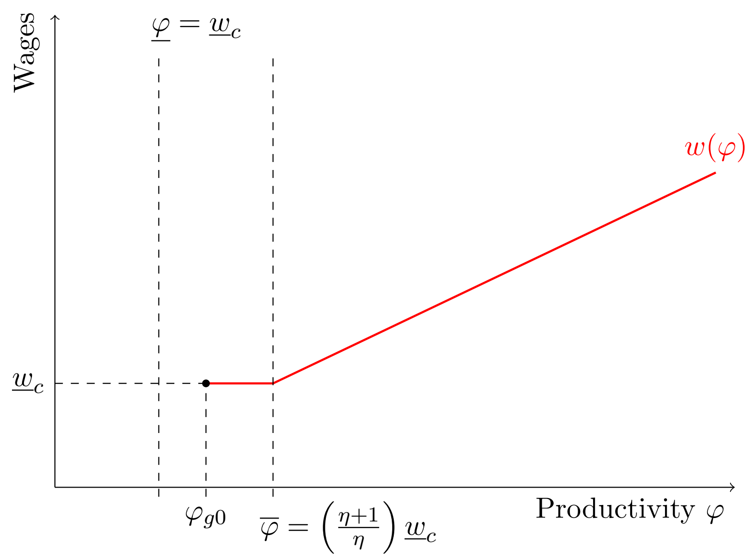

Notes: The figure illustrates the productivity distribution of firms in local labor market g with binding wage floor $\underline{w}_c$, shape productivity parameter α, and minimum value φ_{g0}.

Appendix Figure 14: Estimates of bargaining power parameters

(a) Ratio of derivatives of objective functions

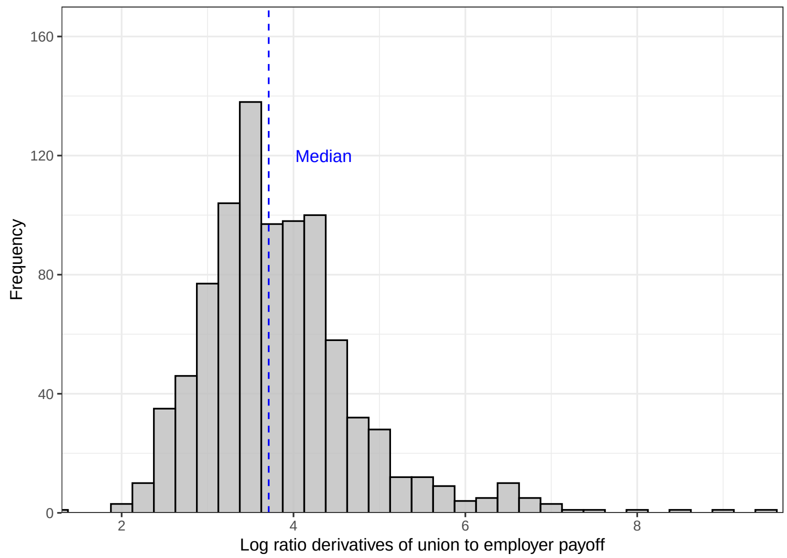

(b) Estimates of bargaining parameters

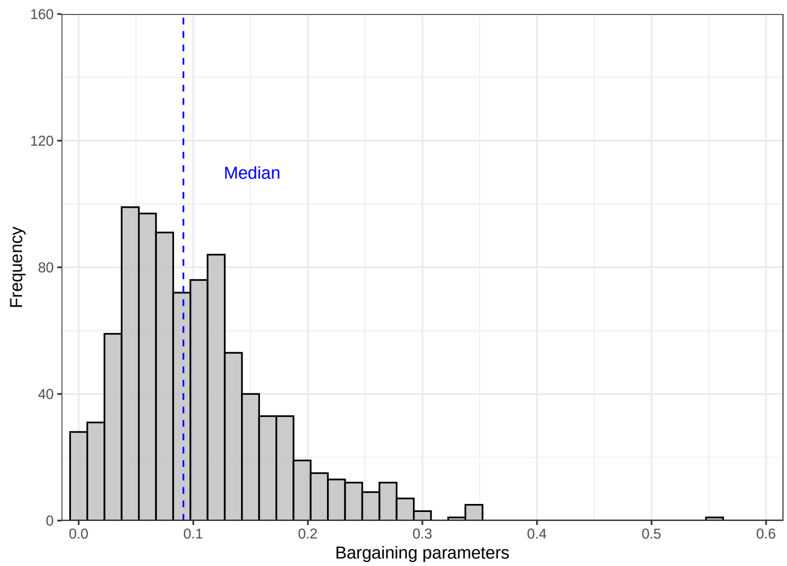

Notes: The figure shows model estimates. Panel (a) shows the estimated ratio of derivatives of the union objective function and the employer association objective function, both with respect to the wage floor. Panel (b) shows the estimated bargaining parameters. I exclude the estimates for the retail CBA 0130/75.

Appendix Figure 15: Share of firms bunching and mean wage in the data and in the model, for 2011–2012

(a) Data

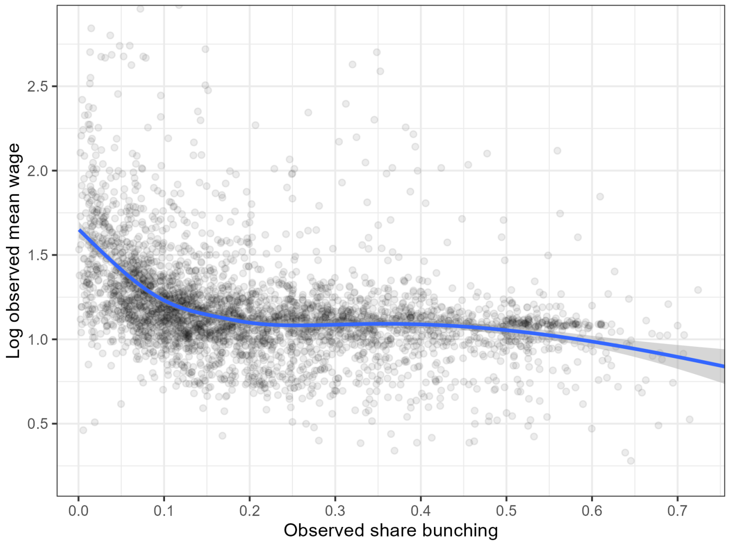

(b) Model

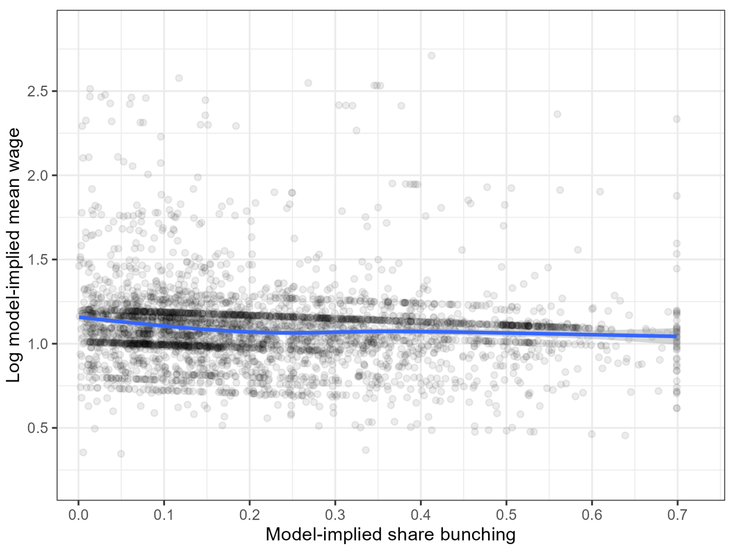

Notes: The figure shows a correlation between the share of firms bunching and the mean wage in each local labor market. Panel (a) shows the data, and panel (b) shows the model-based data. The line represents a locally estimated scatterplot smoothing curve fitted to the data.

Appendix Figure 16: Centralization of bargaining and employment responses to shocks across CB networks

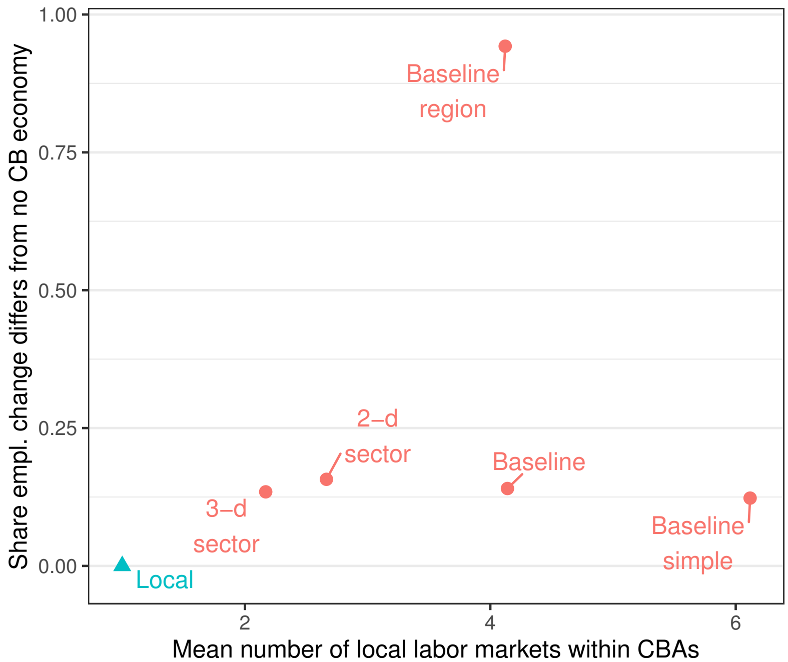

Notes: Data are from model simulations pre- and post-export shocks under different CB networks. The figure shows the share of local labor markets with an employment response to shocks different from a non-bargaining economy against the level of bargaining centralization. Specifically, the y-axis measures the share of local labor markets for which the absolute employment change following the shocks is more than 0.25% different from the counterfactual employment change in an economy without CB. This computation excludes local labor markets that correspond to the retail CBA (0130/75) at baseline and to CBAs with less than 5% of employment in exporting firms. Bargaining centralization is measured as the average number of local labor markets per CBA.