Regulating Flexible Labor Contracts: Substitution and Allocative Efficiency*

Alejandro Micco
University of Chile†

Francisca Pérez
Adolfo Ibañez University‡

February 3, 2026

Abstract

Regulations governing flexible labor contracts shape firms' adjustment margins and the allocation of labor across heterogeneous producers. We study the aggregate consequences of such regulation using Chile's 2006 reform of temporary agency work. Combining plant-level data with a dynamic industry equilibrium model featuring non-convex firing costs and imperfect substitution between regular and agency workers, we show that the reform shut down a key adjustment margin by raising both fixed and variable costs of agency labor. Aggregate productivity declines, even though average firm-level productivity remains unchanged, reflecting increased dynamic misallocation rather than within-firm efficiency losses. Output prices rise as higher labor costs and impaired reallocation contribute in roughly equal measure, with their relative importance governed by the elasticity of substitution across worker types. Our results highlight that regulating flexible employment arrangements can generate sizable productivity losses by weakening firms' dynamic adjustment margins, effects that are not captured by static or reduced-form approaches.

Keywords: labor adjustment costs, temporary agency workers, rent sharing, misallocation, firm dynamics

JEL Codes: J23; J31; J38; D22; E24

*We wish to thank Francois Gourio, Alisdair McKay, and Daniel Hamermesh for helpful discussions and comments during the early stages of this project, as well as audiences at the 2022 SOLE Annual Meeting, 2021 SED Annual Meeting, 5th IZA Junior/Senior Symposium, 2020 EALE/SOLE/AASLE World Conference, 2020 ASSA Annual Meeting, 2019 Chilean Economic Society Annual Meeting, and 2nd IZA/World Bank/NJD Conference on Jobs and Development, Universidad de Chile, Universidad Católica, UAI. This research was partially supported by the supercomputing infrastructure of the NLHPC (ECM-02). This paper previously circulated as "Costs and Benefits of Firms' Use of Outsourced Labor: A Structural Estimation".

†amicco@fen.uchile.cl

‡Corresponding author. Email: mfrancisca.perez@uai.cl

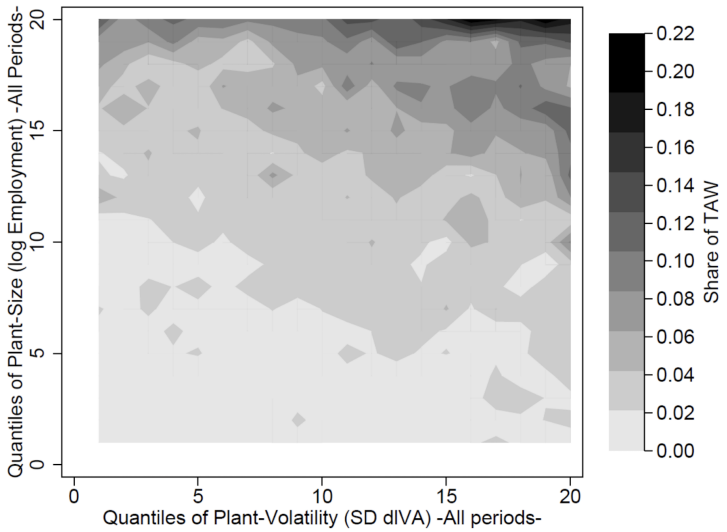
1 Introduction

Regulations governing flexible labor contracts are a central feature of modern labor markets, particularly in economies where employment protection and collective bargaining shape firms' adjustment margins (Lazear, 1990; Botero et al., 2004; Besley and Burgess, 2004). Such policies are often evaluated through their effects on wages, employment, or prices (Micco and Muñoz, 2024; Daruich et al., 2023; Dolado et al., 2021; Bentolila et al., 2012). Yet by altering the relative costs of different types of labor, regulation affects not only firms' labor expenses but also how labor is allocated across heterogeneous producers. These two channels—higher within-firm labor costs and changes in the allocation of labor across firms—have distinct implications for aggregate productivity and prices (Hsieh and Klenow, 2009; Restuccia and Rogerson, 2008; Asker et al., 2014), but they are difficult to disentangle using reduced-form evidence or models that abstract from firm-level labor choices. Assessing the aggregate consequences of regulating flexible labor arrangements therefore requires an explicit framework that models firms' labor choices, substitution across labor types, and the resulting patterns of misallocation in equilibrium.

Flexible labor arrangements such as outsourcing and temporary agency work (henceforth, TAW) provide firms with an important margin of adjustment in environments characterized by employment protection and firing costs (Bentolila and Saint-Paul, 1992; Cahuc and Postel-Vinay, 2002; Houseman, 2001). When adjustment costs for regular employment are non-convex—due to hiring frictions or severance pay—firms cannot smoothly re-optimize their workforce in response to shocks, and flexibility acquires economic value (Hopenhayn and Rogerson, 1993; Caballero, 1999; Cooper and Haltiwanger, 2006; Bloom, 2009; Bentolila and Bertola, 1990). In such settings, temporary agency workers allow firms to absorb demand and productivity fluctuations without triggering costly lumpy adjustments (Caggese and Cuñat, 2008; Blanchard and Landier, 2002; Autor, 2003). Exploiting rich plant-level data from Chile, we find evidence consistent with this role: the use of temporary agency workers is systematically higher among plants facing greater volatility, even conditional on size. Figure 1 illustrates this pattern for Chilean manufacturing plants, showing that the use of agency workers increases with both plant size and value-added volatility. At the same time, flexible labor arrangements often come with lower wages, weaker bargaining power, or distinct regulatory treatment, raising concerns about job quality and worker protection (Goldschmidt and Schmieder, 2017; Drenik et al., 2023; Weil, 2014; Autor, 2001). Policies that regulate TAW therefore embody a central trade-off: improving labor standards while potentially increasing firms' labor costs and reshaping how labor is allocated across heterogeneous firms.

We study these issues in the context of Chile's 2006 reform regulating temporary agency

Figure 1: Distribution of TAW Usage by Size and Volatility



Note: the figure shows the share of TAW across quantiles of plant size and value-added volatility. Plant size is measured as the log of total employment, and volatility is proxied by the five-year standard deviation of the log change in value added. Each cell represents the average TAW employment share for plants in a given size–volatility bin, with darker shades indicating higher reliance on TAW employment.

Source: Own elaboration using data from the Chilean National Annual Industrial Survey (ENIA).

work, which substantially increased the obligations and liabilities faced by user firms. Similar in spirit to reforms implemented across Europe—such as the Directive 2008/104/EC on Temporary Agency Work—the Chilean reform aimed to improve working conditions for TAW employment by raising compliance requirements and strengthening their legal protections. Crucially for our purposes, the reform generated a sharp and plausibly exogenous increase in the cost of using flexible labor contracts, while leaving the broader institutional environment unchanged. This setting provides a clean regulatory shock to examine how firms adjust their labor choices along both the extensive and intensive margins, how regulation reshapes substitution across labor types, and how these firm-level responses translate into aggregate outcomes. While recent reduced-form evidence documents that firms more reliant on temporary agency work prior to the reform experienced larger contractions in employment and output (Micco and Muñoz, 2024), the extent to which such regulation distorts the allocation of labor across heterogeneous firms and affects aggregate productivity remains an open question that we address in this paper.

To analyze these mechanisms and quantify their aggregate consequences, we develop a dynamic model of heterogeneous firms in the tradition of Hopenhayn and Rogerson (1993) that choose between regular and agency workers in the presence of non-convex adjustment costs. Worker types differ in both costs and productivity, with wages determined through

a rent-sharing bargaining process following Card et al. (2018). The model features a rich labor cost structure, with wages and administrative and fixed outsourcing costs, allowing firms to rely on agency labor as a flexible adjustment margin. This dynamic structure enables us to disentangle the effects of regulation operating through direct within-firm labor costs from those arising through the dynamic misallocation channel. By explicitly modeling the intertemporal trade-offs induced by firing costs and hiring frictions, the framework captures the regions of inaction that characterize firm-level labor choices, providing a structural link between micro-level lumpiness and aggregate productivity losses.

We estimate the model using plant-level data from the Chilean National Annual Industrial Survey (ENIA, for its initials in Spanish), which uniquely records firms' employment of both regular and agency workers at the user-firm level. This feature allows us to directly observe firms' reliance on agency labor along both the extensive and intensive margins, overcoming a key limitation of most administrative datasets that link outsourced workers only to contracting agencies rather than to production sites.[1] We estimate the model to match a set of moments that capture firms' size distribution, employment dynamics, exit behavior, and the cross-sectional patterns of temporary agency work usage. By ensuring that the model reproduces these micro-level features, we obtain a quantitatively disciplined framework in which to evaluate the aggregate implications of changes in the cost of flexible labor contracts for productivity, prices, and resource allocation.

The model accounts well for the joint distribution of firm size, employment, exit, and the use of agency workers. It closely replicates the highly skewed firm-size distribution, the concentration of employment in large establishments, and the sharp decline in exit rates with firm size. A key strength of the framework is its ability to match observed patterns of agency-worker use along both the extensive and intensive margins. Fixed administrative costs of outsourcing and heterogeneity in adoption incentives reflecting managerial preferences and workplace frictions, rationalize why many firms never use agency workers, while adoption rises steeply with firm size. Conditional on adoption, the model reproduces the modest increase in agency-worker employment shares with firm size. This pattern reflects proportional outsourcing costs that constrain intensive-margin reliance on agency labor, along with an increasing regular–agency wage gap as firms grow larger. While the model slightly overstates the time-series volatility of agency-worker employment, this reflects a trade-off in matching cross-sectional heterogeneity in agency-worker adoption. Overall, the framework captures the main aggregate and cross-sectional patterns in the data, providing a disciplined quantitative

[1]For instance, Katz and Krueger (2019) rely on survey data to link contingent workers to firms, while Goldschmidt and Schmieder (2017) infer on-site outsourcing from worker flows across establishments. Exceptions include Drenik et al. (2023), who use Argentine administrative data, and Micco and Muñoz (2024), who employ ENIA to study the Chilean reform.

environment to study labor adjustment and outsourcing in regulated labor markets.

We model the 2006 reform as an increase in the cost of using temporary agency work along both extensive and intensive margins. In the model, the reform raises the variable cost of agency labor through higher wages and mandated benefits, increases fixed costs associated with administering agency contracts, and strengthens workers' outside options in wage bargaining. Together, these changes make temporary agency work a more expensive adjustment margin for firms, while leaving firing costs and other forms of employment protection unchanged. Rather than imposing a mechanical equalization of costs across worker types, this approach captures how regulation reshapes firms' effective labor cost structure and adjustment margins following the reform.

Our estimates imply that the 2006 reform raised output prices by about 0.49 percent and reduced aggregate productivity by 0.24 percent under our benchmark elasticity of substitution between regular and agency workers ($\phi = 3.74$). Importantly, this productivity decline is driven entirely by increased misallocation rather than by changes in average firm-level productivity. As access to agency workers becomes more costly, firms lose an important margin to buffer idiosyncratic shocks, weakening the reallocation of output toward more productive producers. By contrast, when we impose a lower elasticity of substitution closer to Cobb–Douglas ($\phi = 1.5$), the misallocation channel is substantially muted and the productivity loss falls to 0.10 percent, while the direct cost effect becomes more prominent. These results highlight that the aggregate consequences of regulating flexible labor contracts depend critically on the degree of substitutability between labor types, and that standard Cobb–Douglas assumptions can substantially understate the role of dynamic misallocation.

Literature review. Our paper connects three strands of literature. First, we contribute to the macroeconomic literature on misallocation. While seminal frameworks (Hopenhayn and Rogerson, 1993) lay the groundwork for understanding firm dynamics and selection, the subsequent misallocation literature (e.g., Restuccia and Rogerson, 2008; Hsieh and Klenow, 2009) has often focused on static distortions to explain productivity gaps. These static approaches typically abstract from the role of labor adjustment margins under non-convex costs. We bridge this gap by connecting our work to the growing literature on dynamic misallocation (Asker et al., 2014; David and Venkateswaran, 2019; Kehrig and Vincent, 2021). We show that flexible labor contracts act as an intertemporal buffer against idiosyncratic shocks, mitigating the dynamic misallocation generated by firing costs rather than simply lowering marginal labor costs.

Second, we build on the empirical literature on labor market reforms of temporary contracts (Daruich et al., 2023; Micco and Muñoz, 2024; Dolado et al., 2021; Bentolila et al., 2012). Unlike Micco and Muñoz (2024) or Daruich et al. (2023), who provide essential evidence on

the average impact of reforms, our structural approach allows us to capture regions of inaction and to quantify the option value of TAW. We show that the loss of this flexibility—the "wait-and-see" channel described by Bloom (2009)—is a primary driver of the productivity drop. Our framework allows us to decompose the aggregate effects of regulation into a direct cost effect and a dynamic misallocation channel.

Finally, we contribute to the literature on labor outsourcing and firm boundaries. Our work is concurrent with Bostanci (2025), who also uses an industry dynamics model to study outsourcing. However, the two papers differ fundamentally in the nature of the friction and the role of outsourcing. While that paper focuses on information-based frictions and high-skill services, we model temporary agency work as a task-equivalent instrument for managing labor risk. The use of plant-level data allows us to identify a non-convex cost structure—including fixed costs of adoption—and to replicate the extensive-margin response to regulation. This allows us to show that the 2006 reform does more than just increase marginal costs; it effectively shuts down a key adjustment margin, locking firms into inefficient employment states. Related work by Bilal and Lhuillier (2025) studies the effects of domestic outsourcing, where firms bypass internal wage ladders to reduce labor costs. In that framework, efficiency gains arise primarily from static cost-saving and reallocation effects. In contrast, our analysis emphasizes how regulating flexible labor affects firms' dynamic adjustment margins, generating productivity losses through intertemporal misallocation.

This mechanism is highly sensitive to the elasticity of substitution between worker types. We show that the standard Cobb–Douglas benchmark—common in the misallocation literature (e.g., Asker et al., 2014; David and Venkateswaran, 2019)—substantially understates the dynamic misallocation channel, as it fails to capture how firms reallocate labor away from rigid margins when flexibility is available. Our findings align with recent frontier work by Chan (2023), who finds high elasticities of substitution (ranging from 1.5 to 4) when accounting for fixed costs and endogenous selection in firms' outsourcing decisions. By empirically disciplining these substitution patterns rather than assuming them, we avoid the pitfalls of functional-form restrictions that can lead to a significant misinterpretation of the economic costs of labor regulation as emphasized by (Gandhi et al., 2020).

The paper proceeds as follows. Section 2 reviews the regulatory environment for regular workers and TAW in Chile and provides details about the 2006 reform. Section 3 presents the theoretical model. Section 4 describes the data and the estimation method. Section 5 presents the SMM estimation results and the goodness of fit of the benchmark model. Section 6 quantifies the impact of the 2006 TAW reform, assessing its effects on productivity and misallocation. Finally, Section 7 offers concluding remarks.

2 Institutional Context

This section describes the regulations governing labor contracts in Chile and the specific features of the 2006 reform. We first characterize the differences between regular and agency workers, focusing on their contractual nature, their degree of substitutability in production, and the enforcement mechanisms that generate compliance costs. We then detail the 2006 reform, which shifted the legal liability toward user firms, effectively increasing both the fixed and variable costs of agency labor. This institutional setting provides the empirical basis for the firm-level trade-offs that we formalize in our theoretical model.

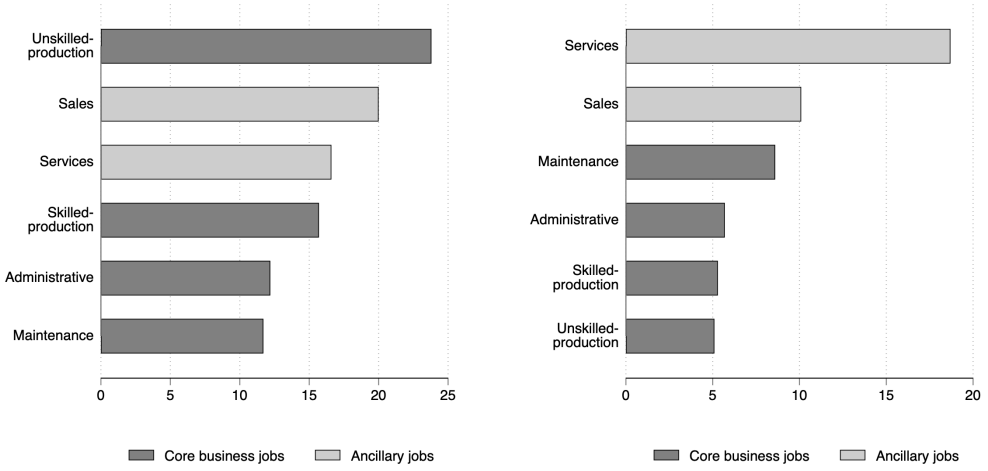
Workers' contractual structure. In Chile, manufacturing plants hire workers either directly as regular workers or indirectly through agencies or contracting firms providing TAW. The relationship between user firms and agency workers is triangular, mediated by the agency, with no direct contractual tie between user firms and the workers supplied. Prior to the 2006 reform, user firms bore *subsidiary* responsibility for TAW obligations; after the reform, they became *jointly and severally liable* alongside agencies. As a result, user firms must actively monitor agency compliance and remain legally accountable for violations related to safety, health, and social security, even though agencies are the employers of record. In day-to-day work, user firms manage tasks, while agencies, as employers of record, assume these formal liabilities.

Workers' substitutability. Agency workers are employed both in core production occupations (e.g., skilled and unskilled production, administration, and machine maintenance) as well as in ancillary services (e.g., security, cleaning, janitorial, and secretarial duties) within the Chilean manufacturing industry.[2] Conceptually, TAWs differ from service outsourcing or subcontracting; in the latter, the subcontracting firm provides a specific service and typically oversees both the execution of the work and the management of the workforce. In contrast, TAWs are integrated into the user firm's daily operations. Within plants, TAWs are most prevalent in unskilled production, sales, and services jobs (Figure 2, Panel (a)). However, their shares within core occupations are relatively uniform, ranging from 5% to 9% (Figure 2, Panel (b)). This intra-occupational overlap indicates that TAWs are not specialized inputs but operate alongside regular workers within the same production frontiers. This motivates our modeling approach, in which regular workers and TAWs are treated as imperfect substitutes within occupations.

[2]For core functions, TAWs typically work under user firms' supervision and resources, while agencies manage work execution primarily in ancillary jobs (Direccion del Trabajo, 2011).

From a cost perspective, TAWs allow firms to avoid firing procedures and non-wage benefits, reflecting their weaker bargaining power (Houseman et al., 2003; Drenik et al., 2023; Goldschmidt and Schmieder, 2017). As a result, they provide a flexible margin to address short-term labor shortages or demand spikes. However, TAWs cannot fully replace regular workers; certain tasks require continuity, firm-specific knowledge, or loyalty that is better incentivized through long-term regular employment (Bostanci, 2025). In our framework, the joint use of both worker types is an equilibrium decision, where firms balance the lower cost and flexibility of TAWs against the stability and specific productivity of regular workers. For estimation, we restrict the sample to core production occupations—where the presence of both worker types is most balanced—and exclude service workers. This ensures that our parameters are identified from the margin where regular and agency workers most clearly compete within the same production tasks.

Figure 2: TAW use by plant, 2005-6



(a) Distribution of TAWs across occupations

(b) Proportion of TAWs within occupations

Note: Figure (a) presents the percentage of TAWs by occupation over total TAWs per plant, $S_{ij}/\sum_j S_{ij}$ for plant i in occupation j. Panel (b) presents the percentage of TAWs over total employment by occupation, $S_{ij}/(S_{ij}+N_{ij})$ for plant i in occupation j. Averaged across plants, 2005–6. Core business jobs include unskilled and skilled production, administrative, and machine maintenance. Ancillary jobs include security, cleaning, janitorial, secretarial duties (under services jobs), and sales jobs.

Employment protection for regular workers. The Chilean Labor Code governs contracts for both regular and agency workers, providing strong employment protection. Termination of regular workers requires one month's notice (or equivalent compensation) and severance

pay equal to one monthly wage per year of service, capped at eleven months. Wrongful or unjustified dismissals trigger penalties of 30–100% of severance pay, implying firing costs that often exceed statutory severance.[3] These substantial non-convex firing costs create significant zones of inaction—where firms find it optimal to maintain their current labor force despite productivity shocks—thus reinforcing the role of agency workers as the primary margin for labor adjustment..

Labor inspections and compliance costs. Compliance with labor regulations is enforced by the Chilean Labor Directorate through two primary mechanisms: programmed and complaint-driven inspections. Programmed inspections target firms based on idiosyncratic characteristics, such as size and industry, while complaint-driven inspections are triggered by worker grievances regarding labor, social security, or safety standards. The use of agency labor is a key determinant in this enforcement process; TAW is frequently associated with a less supportive work environment, increased job insecurity, and higher anxiety among regular workers, which collectively raise the likelihood of formal complaints and inspection requests (Bryson, 2013; Weil, 2014).

These inspections disrupt plant operations even in the absence of sanctions, as they require management to pause production to review documentation, interview staff, and conduct facility tours. In cases of non-compliance, firms face further penalties, including fines, operational suspensions, or temporary closures.[4] Consequently, the institutional environment in Chile implies that employing agency workers entails fixed compliance and monitoring costs that arise whenever TAW is utilized, regardless of the scale of its use.

2.1 2006 TAW reform

The 2006 reform of TAW regulation in Chile was prompted by growing public and policy concerns over the precarious working conditions of temporary agency workers. Compared to regular workers performing similar tasks within the same firms, agency workers faced heightened job insecurity and lacked adequate legal protections, prompting calls for reform (Direccion del Trabajo, 2011). The Chilean case is particularly compelling, as the reform shares key features with contemporaneous policy efforts in Europe. For example, in 2008, the European Commission adopted the Directive on Temporary Agency Work (Directive 2008/104/EC), aimed at ensuring that TAWs receive the same basic working conditions as regular workers, thereby

[3]Chile ranks among OECD countries with the highest firing costs; see Figure A.1 in Appendix A, based on Botero et al. (2004).

[4]For further details on Chile's inspection procedures, see the *Manual de Procedimiento de Fiscalización de la Dirección del Trabajo*.

balancing labor market flexibility with worker protections.[5]

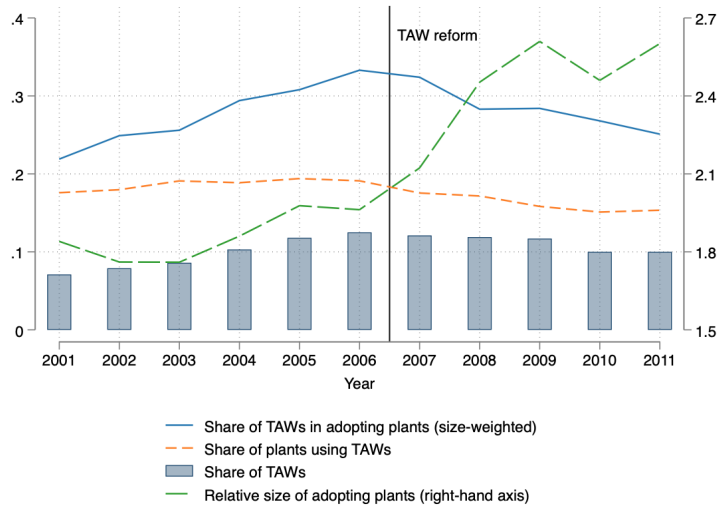
Legal change: joint and several liability. The central change introduced by the 2006 reform was the establishment of *joint and several liability* for user firms and agencies in meeting TAWs' labor obligations, including wages, severance pay, unemployment insurance, and social security contributions. In effect, user firms became responsible for ensuring compliance with labor regulations governing agency workers' safety and health, regardless of contractual arrangements. In the event of noncompliance or direct violations of the Labor Code, TAWs were granted the legal right to pursue action against either the agency or the user firm. To mitigate this liability, user firms now incur additional monitoring costs, including exercising their legal *Right to Information* and invoking their *Withholding Right* to suspend payments if obligations are unmet. Agencies, in turn, are required to establish financial guarantees to cover potential defaults. Prior to the reform, user firms were only secondarily liable, and the scope of these obligations was often ambiguous.

Evidence of cost changes induced by the reform. The 2006 reform aimed to equalize working conditions between TAW and regular employees, fundamentally altering the cost structure for user firms. On the one hand, variable costs rose due to stricter requirements regarding workplace safety and social security obligations—such as the mandatory provision of protective equipment—alongside a shift in bargaining power that narrowed the wage gap relative to regular workers. On the other hand, the expansion of joint and several liability introduced a more complex contractual relationship, raising fixed overhead costs associated with continuous monitoring and legal compliance efforts invariant to the number of agency workers employed.

Figure 3 illustrates the impact of these cost changes for Chilean manufacturing plants through three salient patterns: (i) a decline in the share of plants employing TAWs (orange dashed line), (ii) an increase in the relative size of plants that continue to use them (green long-dashed line), and (iii) a reduction in the share of TAWs within adopting plants (solid blue line). Furthermore, Figure 4 shows that TAW adoption fell across all plant-size groups, with the most pronounced declines among smaller plants. Taken together, these patterns are consistent with an increase in both fixed and variable costs, operating through extensive- and intensive-margin adjustments. This combination of institutional shifts and firm-level responses motivates our use of the 2006 reform as a setting to study how the regulation of flexible labor contracts reshapes firm behavior and aggregate resource allocation.

[5]Figure A.1 in Section A of the Supplemental Appendix shows that the rapid growth in the share of temporary agency workers across EU countries up to the mid-2000s largely levels off thereafter, coinciding with the adoption of a stricter regulatory framework (Directive 2008/104/EC).

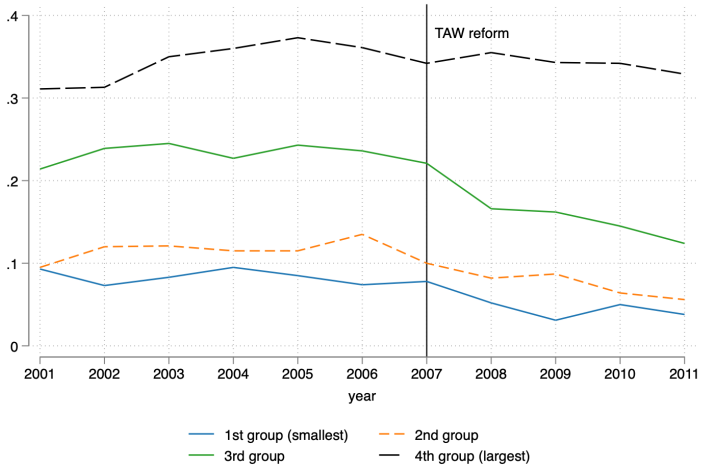
Figure 3: Firm-Level TAW Usage: Pre- and Post-Reform



Note: The figure shows plant-level trends in the use of TAWs between 2001 and 2011. Bars represent the industry-wide share of TAWs. The dashed orange line shows the share of plants using TAWs, the solid blue line shows the average share of TAWs in adopting plants weighted by plant size, and the long-dashed green line shows the relative size of adopting plants, measured by employment, compared to the average size of all plants (whether or not they use TAWs).

Source: Own elaboration using data from the Chilean Annual Industrial Survey (ENIA).

Figure 4: Share of Plants using TAWs by Plant Size: Pre- and Post-Reform



Note: The figure shows the share of plants using TAWs by plant-size. Plant-year observations are divided into four groups based on the fourth lag of total employment.

Source: Own elaboration using data from the Chilean National Annual Industrial Survey (ENIA).

3 Model

In this section, we present a model of firm dynamics in the tradition of Hopenhayn and Rogerson (1993), extended to allow firms to hire both regular and agency workers. The model features imperfect substitution between worker types, non-convex adjustment costs for regular employment, and a wage-setting structure that links worker compensation to firm productivity. We solve the firm's dynamic labor demand problem, taking output and input prices as given, and use the model to quantify how changes in the regulation of TAW affect firms' adjustment behavior and the allocation of labor across heterogeneous producers.

3.1 Firms' Technology

There is an industry composed of a continuum of firms that produce a homogeneous good.[6] Firms behave competitively, taking output and materials prices as given. Throughout the paper, we use the terms "firms" and "plants" interchangeably. Each firm operates a decreasing returns-to-scale production technology using labor and intermediate inputs. Materials are assumed to be fully flexible, as is standard in the literature (e.g., Garcia-Marin and Voigtländer, 2019). We abstract from capital and focus on labor adjustment as the primary margin of short-run input reallocation.[7] In a Cobb–Douglas production structure, changes in labor costs affect substitution between the labor bundle and capital but do not alter the composition of labor inputs, which is the object of interest in this paper.

Production is given by the following Cobb–Douglas technology:

$$Y_t = Z_t \left[V_t^{\gamma} M_t^{1-\gamma} \right]^{\alpha}, \tag{1}$$

where $\alpha \in (0,1)$ governs returns to scale, γ is the labor share, and Z_t is firm-specific productivity. Productivity follows a first-order Markov process over a finite support and is i.i.d. across firms, implying no aggregate uncertainty. Hence, all firms face the same stochastic productivity process ex ante, while heterogeneity in outcomes arises endogenously from firms' dynamic responses to idiosyncratic shocks.

Labor input V_t is a composite of regular workers, N_t, and TAW, S_t, combined through a

[6] Consistent with our plant-level data, we treat each production site as an independent decision-maker, assuming that multi-plant firms manage employment at the establishment level (see, e.g., Cooper et al., 2015).

[7] We abstract from capital adjustment as it typically faces higher non-convex costs and lower short-run elasticity than labor, making it an unlikely substitute for the flexible buffer role of TAW (Cooper and Haltiwanger, 2006). Under capital-labor complementarities, labor rigidities likely discourage investment, meaning our productivity loss estimates are likely a lower bound (Bloom, 2009).

Constant Elasticity of Substitution (CES) aggregator with elasticity of substitution $\phi > 1$:

$$V_t = \left[(a_N)^{\frac{1}{\phi}} (N_t)^{\frac{\phi-1}{\phi}} + (a_S)^{\frac{1}{\phi}} (S_t)^{\frac{\phi-1}{\phi}} \right]^{\frac{\phi}{\phi-1}}. \tag{2}$$

where a_N and a_S capture the relative productivity of regular workers and TAW, respectively.

To capture labor substitution, we model regular and temporary agency workers as imperfect substitutes within a CES labor aggregate. This formulation allows both worker types to be used within the same occupations while accounting for differences in relative productivity. As documented in Section 2, agency workers in Chile are typically integrated into core production tasks alongside regular staff rather than being confined to peripheral activities. This evidence motivates a framework where substitution operates within occupations, reflecting the operational reality of the user firms.

3.2 Workers' total compensation

We depart from the standard competitive labor market framework and allow wages to depend on firm productivity. A large empirical literature documents substantial heterogeneity in firm productivity and shows that these differences translate into systematic wage dispersion across firms (Syverson, 2011; Dunne et al., 2004; Faggio et al., 2010; Barth et al., 2016). Using matched employer–employee data, Card et al. (2018) show that firm-specific components account for a large share of wage inequality.

To capture this empirical regularity, we model wages through a rent-sharing mechanism, whereby workers receive a share of firm-specific quasi-rents (Abowd and Lemieux, 1993; Card et al., 2014; Lamadon et al., 2022). Under this structure, productivity shocks are only partially passed through to wages, reflecting bargaining frictions and limited worker mobility.

While reduced-form regressions provide suggestive evidence of such wage–rent relationships, they face inherent identification challenges in this setting, as wages mechanically enter firm-level quasi-rents. The structural framework allows us to separately identify rent-sharing elasticities by exploiting exogenous productivity variation and firms' dynamic labor demand responses.

We assume that wages for regular workers follow a wage curve of the form:

$$\log W_N(Z_t) = \log \overline{W}_N + \psi_N \log \left(\frac{Z_t}{\overline{Z}} \right), \tag{3}$$

where $W_N(Z_t)$ denotes the wage paid to regular workers in a firm with productivity Z_t, $\overline{W}_N$ is the wage at a reference productivity level $\overline{Z}$, and ψ_N captures the elasticity of wages with

respect to firm productivity.

Similarly, the cost of employing TAWs is given by:

$$\log W_S(Z_t) = \log \overline{W}_S + \psi_S \log \left(\frac{Z_t}{\overline{Z}} \right), \tag{4}$$

where ψ_S governs the responsiveness of TAW costs to firm productivity.

We interpret $W_S(Z_t)$ as the per-worker price of outsourced labor services paid by the user firm, rather than as a pure wage, consistent with models of domestic outsourcing in which firms purchase labor services rather than hire workers directly (Bilal and Lhuillier, 2025). This price reflects the compensation of the worker as well as the agency markup embedded in the service contract, but excludes fixed regulatory or administrative costs, which are modeled separately.

Although agency workers do not bargain directly with the user firm, they bargain with the agency over their compensation, and we assume that these negotiated labor costs are passed through one-for-one to the user firm through the service fee charged by the agency. Under this assumption, bargaining between agency workers and the agency is economically equivalent to direct bargaining between agency workers and the user firm: the agency acts as a passive intermediary in terms of pass-through rather than as a buffer that absorbs rents. As a result, firm-level quasi-rents can affect the total cost of agency labor. Under this interpretation, ψ_S captures the extent to which productivity-driven quasi-rents at the user firm are reflected in the price of outsourced labor services, consistent with the weaker bargaining position of agency workers relative to regular employees.

As motivating evidence, simple reduced-form estimates using Chilean manufacturing data indicate that wages of TAWs respond significantly less to firm quasi-rents than wages of regular workers. These patterns are consistent with weaker bargaining power among agency workers and motivate our structural estimation, which exploits exogenous productivity variation to identify the underlying rent-sharing elasticities. Supporting reduced-form evidence documenting lower rent-sharing for agency workers relative to regular workers is reported in the Supplemental Appendix, Section B.

3.3 TAWs fixed costs

Consistent with the empirical evidence that TAW use is concentrated among larger firms, we model temporary agency work as requiring the payment of fixed costs that grant access to labor flexibility. These costs capture the idea that outsourcing is not a frictionless margin, but rather an operational technology that firms adopt only when the option value of flexibility is sufficiently high.

We consider three types of fixed costs associated with employing agency workers.

Regulation compliance costs. Firms must devote internal resources to monitoring regulatory compliance and managing legal exposure associated with the use of temporary agency workers under joint liability, beyond the standard administrative costs of regular employment. Although user firms do not manage individual employment contracts for agency workers, they remain legally responsible for compliance with labor and safety regulations. We model this administrative burden as the cost of assigning c_s regular workers, paid wage W_N, to oversee TAW-related compliance. This cost enters the profit function as $c_s W_N$ and is incurred whenever a firm employs agency workers, regardless of the scale of agency employment.

Inspection disruption costs. As discussed in Section 2, the use of TAWs increases firms' exposure to labor inspections, which disrupt production regardless of whether violations are detected. We model the expected cost of these disruptions as a fraction δ of firm revenues, following Cooper et al. (2015). Firms employing TAWs face a higher inspection probability due to increased regulatory scrutiny and a greater likelihood of worker complaints, making these disruption costs an additional fixed burden of outsourcing that is independent of the scale of agency employment.

Together, the compliance cost $c_s W_N$ and the disruption cost δ constitute the fixed regulatory entry price to accessing labor flexibility. By paying these costs, firms acquire the ability to adjust employment without incurring the non-convex firing costs associated with regular contracts. Regulations that increase either component therefore act as a tax on firm agility, disproportionately affecting producers that rely on TAWs to hedge against idiosyncratic productivity shocks.

Nonpecuniary adoption costs. Finally, beyond these regulatory costs, we allow for heterogeneity in firms' willingness to use TAWs by introducing a nonpecuniary adoption cost ν, drawn from a $N(0, \sigma_\nu)$ distribution. This subjective cost captures managerial preferences and workplace frictions that may deter TAW adoption even when expected profits are positive. Importantly, ν is realized after the firm has chosen its regular workforce N_t and affects only the subsequent decision to employ TAWs. Following Gomes (2018), this component prevents implausibly sharp changes in the fraction of firms employing agency workers.

To remain faithful to the institutional setting, we assume that agency workers can perform the same tasks as regular workers within the firm and that terminating agency contracts does not involve dismissal costs for the user firm.

3.4 Regular workers' adjustment costs

Employment protection regulation imposes substantial firing costs on firms when terminating regular workers, including advance notice and severance pay (Section 2). Severance payments increase with tenure but are not guaranteed in all separations; for instance, workers who quit voluntarily are not entitled to them. This dependence on tenure and separation type complicates the modeling of firing costs, as it would require tracking workers' individual histories. To balance tractability and realism, we distinguish between regular workers with tenure and newly hired workers who have not yet acquired it.

This institutional structure gives rise to fundamentally non-convex labor adjustment costs. Unlike quadratic adjustment cost specifications, commonly used for tractability, firing costs tied to tenure and dismissal generate regions of inaction in which firms optimally tolerate deviations from their desired employment level rather than adjust continuously. As a result, employment dynamics are inherently intertemporal, and firms' current labor decisions depend on expectations about future shocks and adjustment costs.

Let L_{t-1} denote the stock of regular workers with tenure inherited from the previous period. In each period, firms choose the total number of regular workers N_t, which may exceed or fall short of L_{t-1}. Regular workers with tenure quit exogenously with probability λ_1, corresponding to voluntary separations that do not trigger severance payments. Newly hired regular workers gain tenure with probability $1 - \lambda_2$ and become eligible for severance in subsequent periods. As a result, the stock of tenured workers evolves stochastically over time and constitutes a state variable in the firm's dynamic problem.

In addition, firms incur a fixed turnover cost F whenever they adjust their regular workforce, capturing hiring, screening, onboarding, and dismissal-related expenses. The total labor adjustment cost faced by a firm that begins period t with L_{t-1} tenured workers and chooses to operate with N_t regular workers is given by:

$$g(N_t, L_{t-1}) = \begin{cases} F, & \text{if } N_t > L_{t-1}, \\ F + \tau W_N \times (L_{t-1} - N_t), & \text{if } N_t < L_{t-1}, \\ 0, & \text{if } N_t = L_{t-1}, \end{cases} \tag{5}$$

where τ denotes the severance obligation measured in units of per-period wages for each dismissed tenured worker. In Chile, severance pay is typically equivalent to one monthly wage per year of tenure, though payments are not guaranteed in all separations and may depend on judicial rulings. The parameter τ therefore captures the average effective cost of dismissal implied by the institutional environment, while the fixed cost F captures turnover-related expenses associated with hiring and dismissal.

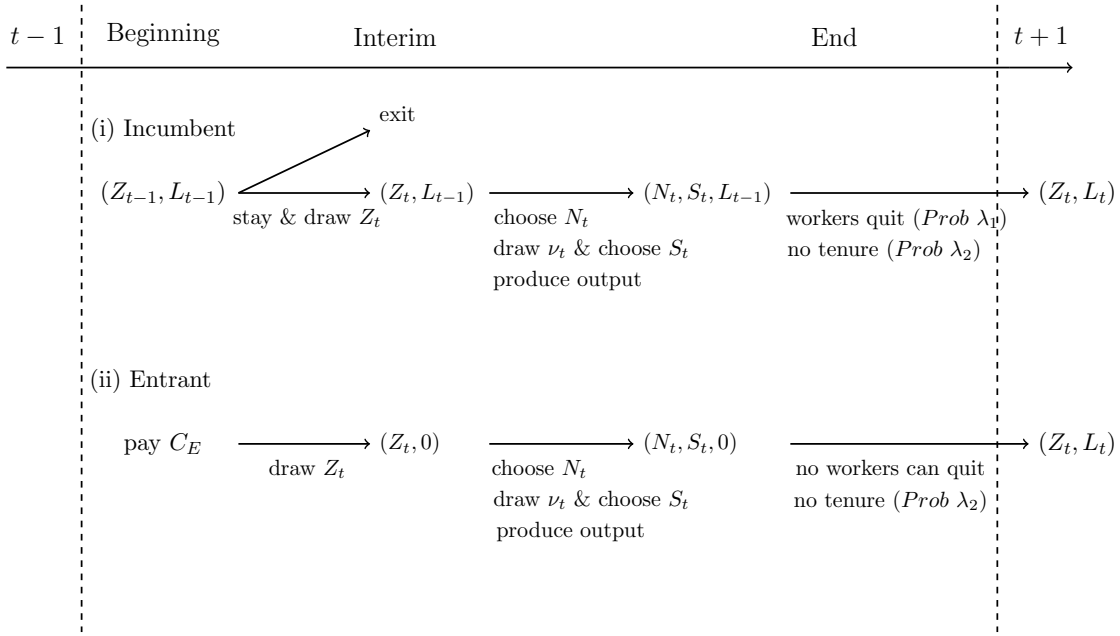
This adjustment cost structure implies non-convex labor adjustment: the marginal cost of changing employment does not vanish as the size of the adjustment approaches zero. Firms therefore optimally follow threshold-based employment policies, characterized by periods of inaction punctuated by discrete adjustments, consistent with the lumpy adjustment patterns documented in micro data (Cooper and Willis, 2009; Bloom, 2009). In environments with high severance pay and dismissal protections, such non-convexities are first-order and generate substantial frictions in firms' responses to idiosyncratic productivity shocks, amplifying the economic relevance of short-run volatility.

Temporary agency work emerges in this setting as a key adjustment margin that allows firms to respond to transitory shocks without triggering the discrete firing costs associated with regular employment. By increasing the fixed and variable costs of accessing this margin, the 2006 reform tightened firms' adjustment constraints, effectively expanding the regions of inaction induced by employment protection legislation. Capturing these mechanisms requires a dynamic model with non-convex labor adjustment costs, which cannot be replicated by frameworks relying on quadratic or static adjustment costs.

3.5 Timing of events for firms

Figure 5 summarizes the within-period timing of decisions and shocks for incumbent firms and potential entrants.

Figure 5: Within-period timing of events for incumbents and entrants



Incumbent firms. Incumbent firms enter period t with two predetermined state variables inherited from the previous period: productivity Z_{t-1} and the stock of tenured regular workers L_{t-1}. The within-period sequence is as follows:

1. *Exit or continuation decision.* The firm decides whether to exit or remain active.
   - If it exits, the firm fires all remaining tenured workers and pays the associated separation cost $g(0, L_{t-1}) = F + \tau W_N L_{t-1}$, after which it exits permanently and earns zero profits.[8]
   - If it stays, the firm pays the fixed operating cost C_F, observes its current productivity draw Z_t, and proceeds to choose labor inputs and produce.
2. *Regular labor decision.* Given (Z_t, L_{t-1}), the firm chooses its desired number of regular workers N_t. Hiring or firing regular workers entails a fixed adjustment cost F and, in the case of separations, severance payments for dismissed tenured workers.
3. *Temporary agency worker decision.* After choosing N_t, the firm draws the nonpecuniary cost ν_t associated with employing TAWs and chooses the number of agency workers $S_t \geq 0$, conditional on N_t. Production then takes place.

At the end of period t, tenured regular workers quit with probability λ_1, while newly hired regular workers fail to acquire tenure with probability λ_2. These flows determine the next-period stock of tenured workers L_t.

Potential entrants. Potential entrants face the following sequence:

1. *Entry decision.* The firm decides whether to enter the industry.
   - If it enters, it pays a one-time entry cost C_E (denominated in units of output) and draws its initial productivity Z_t from the stationary distribution $\nu(Z)$.
2. *Operations.* Conditional on Z_t, the entrant follows the same within-period sequence as an incumbent. Since entrants begin with no tenured workers ($L_{t-1} = 0$), their initial regular labor choice is not affected by firing costs.

[8] Firms exit to avoid paying the fixed operating cost rather than remaining inactive while waiting for a better productivity realization. Exiting firms may re-enter later as potential entrants.

3.6 Firms' current-period profit function

Firms' current-period profits are defined as expected operating profits net of labor adjustment costs and fixed operating costs:

$$\begin{aligned}\Pi(Z_t, N_t, \nu_t) = \mathbb{E}_{\nu_t}\Big[OP(Z_t, N_t, S(Z_t, N_t, \nu_t)) - C_S(S(Z_t, N_t, \nu_t))\Big] \\ - g(N_t, L_{t-1}) - C_F(Z_t),\end{aligned} \tag{6}$$

where operating profits are given by

$$OP(Z_t, N_t, S_t) = P_t Y_t(Z_t, N_t, S_t) - W_N(Z_t)N_t - W_S(Z_t)S_t - P_M M_t(Z_t, N_t), \tag{7}$$

with output Y_t defined in equation (1).

The function $g(N_t, L_{t-1})$ captures hiring and firing costs for regular workers, as defined in equation (5), and depends on the inherited stock of tenured workers L_{t-1}. Fixed operating costs are productivity-dependent and parameterized as $C_F(Z_t) = P\, c_f\,(Z_t/\overline{Z})$, where $\overline{Z}$ represents a reference productivity level, c_f is the cost per unit of output at that reference point and P is output price. This specification allows more productive firms to incur proportionally higher overhead costs.

TAW adoption decision. Consistent with the institutional setting, firms first choose the size of their regular workforce N_t. Conditional on this choice, they decide whether to employ temporary agency workers. Employing TAWs requires paying fixed costs that grant access to labor flexibility, reflecting the idea that outsourcing is an operational technology rather than a frictionless margin.

After choosing N_t, the firm draws a nonpecuniary subjective cost $\nu_t \sim N(0, \sigma_\nu)$ associated with the use of TAWs. This cost captures heterogeneity in managerial preferences and workplace frictions that may deter TAW adoption even when expected profits are positive, following Gomes (2018).

The expected operating profits net of TAW-related fixed costs can be written as

$$\begin{aligned}\mathbb{E}_{\nu_t}[OP(Z_t, N_t, S_t) - C_S(S_t)] = \chi(Z_t, N_t)\Big(OP(Z_t, N_t; S_t > 0) \\ - \delta P_t Y_t(Z_t, N_t; S_t > 0) - c_s W_N\Big) \\ + (1 - \chi(Z_t, N_t))OP(Z_t, N_t; S_t = 0),\end{aligned} \tag{8}$$

where $c_s W_N$ represents the administrative cost of compliance under joint liability, and $\delta P_t Y_t$ captures expected revenue losses from inspection-related disruptions. Firms employing TAWs face higher expected disruption costs due to increased regulatory scrutiny and a higher likeli-

hood of worker complaints.

The firm hires TAWs whenever the realized subjective cost ν_t is below a threshold ν^* defined as

$$\nu^* = \frac{OP(Z_t, N_t, S_t > 0) - \delta P_t Y_t(Z_t, N_t; S_t > 0) - c_s W_N - OP(Z_t, N_t; S_t = 0)}{P_t Y_t(Z_t, N_t; S_t > 0) - P_t Y_t(Z_t, N_t; S_t = 0)}. \tag{9}$$

The denominator rescales the net gains from TAW use by firm size, ensuring that the subjective cost has comparable implications across firms. The probability of employing TAWs is therefore

$$\chi(Z_t, N_t) = \Phi(\nu^*; 0, \sigma_\nu), \tag{10}$$

where $\Phi(\cdot)$ denotes the standard normal cumulative distribution function. The firm's policy rule for TAW employment is

$$S(Z_t, N_t, \nu_t) = \begin{cases} S(Z_t, N_t) > 0, & \text{if } \nu_t \leq \nu^*(Z_t, N_t), \\ 0, & \text{otherwise.} \end{cases} \tag{11}$$

Why do firms use TAWs? In this framework, firms use TAWs both to enhance productive efficiency through substitution within the labor aggregate and to mitigate exposure to non-convex firing costs associated with regular employment. Fixed costs of TAW use generate an extensive margin: only sufficiently large or volatile firms find it optimal to adopt agency labor. Conditional on adoption, the intensity of TAW use depends on relative labor costs and the degree of substitutability between worker types. Differences in rent-sharing elasticities imply that regular workers' wages rise more steeply with productivity, making TAWs relatively more attractive in larger firms, while task-based constraints prevent full displacement of regular employment.

3.7 Dynamic Optimization

We study a stationary equilibrium with constant prices, $\vec{p}_t = [P, P_M, W_N, W_S]$. Firms solve a dynamic labor demand problem under uncertainty. In this equilibrium, firms may grow, contract, exit, or enter, while the distribution of firms across states remains constant.

3.7.1 Incumbent Firms

This section describes the value of a firm that operates in the industry in period t. The firm enters the period with L_{t-1} tenured regular workers and, after paying the fixed operating cost,

observes its current productivity draw Z_t. Its value function, denoted $V(Z_t, L_{t-1}; \vec{p})$, is given by the solution to the following dynamic programming problem:

$$V(Z_t, L_{t-1}; \vec{p}) = \max_{N_t \geq 0} \Bigg\{ \Pi(Z_t, N_t, S(Z_t, N_t, \nu_t)) + \beta \mathbb{E}_{L_t} \Bigg(\max \Big\{ \mathbb{E}_{Z_{t+1}} V(Z_{t+1}, L_t; \vec{p}), -g(0, L_t) \Big\} \mid N_t, L_{t-1} \Bigg) \Bigg\}. \tag{12}$$

The firm chooses the number of regular workers N_t to maximize current-period profits and its expected continuation value. This continuation value reflects the firm's forward-looking assessment of whether to remain in the industry or exit in the following period.

At the time of choosing N_t, the firm faces uncertainty along two dimensions. First, the future stock of tenured workers L_t depends on current employment decisions and stochastic quits and tenure acquisition. Second, conditional on L_t, the firm faces uncertainty about next-period productivity Z_{t+1}.

The exit decision is evaluated after L_t is realized but before drawing Z_{t+1}. In particular, the firm compares the expected value of continuing operations, $\mathbb{E}_{Z_{t+1}} V(Z_{t+1}, L_t)$, to the cost of exit, which entails dismissing all tenured workers and paying the associated separation cost $g(0, L_t)$. Since both the continuation value and the exit cost depend on L_t, expectations over its possible realizations are central to the firm's problem.

The evolution of L_t links current employment decisions to future adjustment costs. If the firm expands ($N_t > L_{t-1}$), newly hired workers may acquire tenure with probability $1 - \lambda_2$; if it contracts ($N_t < L_{t-1}$), some tenured workers are dismissed. In addition, tenured workers quit exogenously with probability λ_1. As a result, today's choice of N_t affects not only current production and labor costs, but also the firm's future state, its flexibility, and its potential exit costs.

In each period, incumbent firms therefore make two interrelated decisions within a single dynamic optimization problem. After observing current productivity Z_t and inherited tenure L_{t-1}, the firm chooses its regular workforce N_t and, conditional on that choice, its use of temporary agency workers. Subsequently, given the realized stock of tenured workers L_t, the firm decides whether to continue operating or exit. Although these decisions occur at different stages of the period, they are jointly determined through the firm's value function.

3.7.2 Entrants

Potential entrants compare the expected value of operating a new plant to a fixed entry cost C_E. Free entry implies:

$$V^e(\vec{p}) = \int V(Z_t, 0) dG(Z_t) \leq C_E = Pc_e, \tag{13}$$

Entrants draw initial productivity from the stationary distribution and begin operations without tenured workers.

3.8 Stationary Equilibrium

We focus on a stationary equilibrium in which the distribution of firms over productivity and employment states is time-invariant. Let $\mu^*(Z, L)$ denote the invariant distribution of firms over productivity Z and the stock of tenured regular workers L. This distribution is induced by firms' optimal labor, exit, and entry decisions, together with the stochastic evolution of productivity and tenure.

In equilibrium, incumbents that remain in operation transition across states according to their optimal policies and the exogenous processes governing productivity and worker flows, while entrants draw initial productivity from the invariant distribution and start without tenured workers. Entry occurs until the expected value of entering equals the entry cost.

Aggregate outcomes are obtained by integrating firm-level decisions over the stationary distribution μ^* and the mass of entrants. Prices are constant in equilibrium.

A stationary equilibrium is therefore defined by firm value and policy functions, prices, an invariant distribution of firms, and a mass of entrants such that: (i) firms solve their dynamic optimization problem, (ii) markets clear, (iii) the distribution of firms is stationary, and (iv) the free-entry condition holds. The law of motion for μ^* and the numerical procedure used to compute the equilibrium are described in Appendix B.

4 Estimation method

This section describes the data and the estimation strategy. We follow a two-step approach. In the first step, a subset of parameters is directly estimated or calibrated using micro-level data (Section 4.2). In the second step, the remaining structural parameters are estimated via a Simulated Method of Moments (SMM) procedure, targeting moments that capture plant-level employment dynamics, the firm size distribution, and the use of temporary agency workers in Chilean manufacturing (Section 4.3).

4.1 Data

We use panel data from the Chilean Annual Industrial Survey (ENIA) for the period 2001–2011, collected by Chile's *Instituto Nacional de Estadísticas* (INE).[9] The survey covers all manufacturing plants with more than ten regular workers and is updated annually to include new entrants and continuing plants. For simplicity, we treat each plant as a firm, noting that multi-plant firms are rare in the data (97% of firms operate a single plant).

We restrict the estimation sample to low-technology manufacturing industries following Levinsohn and Petrin (2003), classified according to the OECD ranking based on direct and embodied R&D intensity (Hatzichronoglou, 1997). This restriction improves comparability in production structures and ensures sufficient data coverage. The low-technology group includes food, textiles, leather, wood, printing, and furniture (ISIC Rev. 3 codes 15, 17–20, 22, 36).[10]

Each year, ENIA reports total employment at the plant level, distinguishing between regular workers—directly hired under open-ended contracts—and TAWs, who are hired through third parties.[11] The dataset also records the occupational composition of both groups, including production (skilled and unskilled workers, maintenance), administrative (e.g., accountants, lawyers), and service roles (e.g., cleaning, security, janitorial, secretarial). We exclude plant-year observations with zero or missing regular employment. Consistent with the institutional discussion (Section 2), we exclude service occupations from the analysis and focus on production workers, for whom temporary agency labor operates within the same production frontiers as regular employment.

The dataset also includes gross revenues (our measure of output), intermediate input expenditures, and total labor compensation for both regular workers and TAWs. ENIA does not report wages at the worker level, but rather total payments made by the user firm to regular workers or to the agency supplying TAWs. All nominal variables are deflated to 2013 Chilean pesos using the annual GDP deflator from the Central Bank of Chile. Summary statistics are reported in Appendix A.

To quantify the costs of the 2006 TAW reform, we estimate the benchmark model using data from 2005–2006, with 2004 serving as the base year to compute employment growth rates.[12] To assess the impact of the reform, we re-estimate the model using post-reform data from 2010–2011, a period chosen to avoid confounding effects from the 2008–2009 Great

[9] This dataset has been widely used in previous studies; see, for example, Levinsohn and Petrin (2003); Asker et al. (2014); Gandhi et al. (2020).

[10] High- and medium-technology industries include, among others, chemicals, machinery, electrical equipment, and motor vehicles (ISIC Rev. 3 codes 21, 24–35).

[11] The ENIA questionnaire defines TAWs as "all personnel who do not have a direct contractual relationship with the firm but perform activities for it."

[12] These years are characterized by relatively stable industrial growth, in contrast to the slower expansion observed in the early 2000s.

Recession.

4.2 Predefined Parameters

We partition the full set of model parameters into two subsets, $\theta = \{\theta_p, \theta_e\}$, and estimate them in two steps. This section describes the calibration or direct estimation of a subset of predefined parameters:

$$\theta_p = \{\beta, \lambda_1, \lambda_2, \alpha, \phi, \gamma, a_N, a_S, \rho\}.$$

We set the discount factor β to 0.85, consistent with the cost of capital for medium-sized firms in the food processing sector.[13]

The annual quit rate of tenured regular workers is set to $\lambda_1 = 0.05$, and the probability that newly hired regular workers fail to transition into tenure is set to $\lambda_2 = 0.08$. Both parameters are calibrated using monthly administrative data from Chile's Unemployment Insurance system, exploiting observed separation and reemployment patterns.

The remaining parameters in θ_p, which govern production technology and productivity dynamics, are estimated using plant-level data and described in the next subsection. We estimate the elasticity of substitution between regular and agency workers outside the full dynamic model, exploiting reduced-form variation induced by the 2006 reform. Estimating this parameter jointly with the remaining structural parameters would weaken identification, as multiple mechanisms affect firms' employment responses along the intensive margin in the dynamic model.

4.2.1 Identification of Production Technology Parameters

Our identification strategy for production technology parameters follows a similar spirit to Micco and Muñoz (2024), in that we discipline key parameters using relationships implied by firms' input choices. However, the economic rationale underlying this strategy differs in our framework. While Micco and Muñoz (2024) rely on quadratic adjustment costs, which ensure that static first-order conditions hold pointwise at the firm–period level, our model features non-convex adjustment costs for regular workers that generate inaction regions and lumpy employment adjustments. As a result, static first-order conditions need not hold in every period. Nevertheless, in a stationary equilibrium without aggregate productivity trends, these conditions continue to hold on average, implying that long-run expenditure shares and

[13] Based on the cost of capital for traded firms in the food processing sector in emerging markets (Damodaran Online: Home Page for Aswath Damodaran), plus a 5 percentage point premium to account for differences between large traded firms and the median-sized firms in our sample.

relative factor intensities remain informative about the underlying production technology, even though short-run employment choices are distorted by adjustment frictions.

Returns to scale (α). We discipline the curvature of the profit function using long-run expenditure shares on variable inputs. In a stationary equilibrium with no trend in productivity, the sum of expenditures on labor and intermediate inputs equals α times revenue on average, even in the presence of adjustment costs (Bertola and Caballero, 1990). Appendix C reports the static first-order conditions underlying this relationship. In the data, we proxy this long-run relationship using the median ratio of total labor compensation and materials expenditures to nominal output, which provides a robust estimate that is less sensitive to outliers. This yields $\alpha = 0.71$, consistent with previous estimates for Chilean manufacturing plants (Asker et al., 2014).

Labor share parameter (γ). Given $\hat{\alpha}$, we recover the labor share parameter γ directly from observed materials expenditure shares. Since materials are not subject to adjustment costs, the first-order condition with respect to materials holds at the firm–period level. Using the condition $\alpha(1-\gamma) = P_M M/(PY)$, we obtain a median value of $\hat{\gamma} = 0.29$ for plants in 2005–2006.

Elasticity of substitution between regular workers and TAWs (ϕ). To discipline the elasticity of substitution between regular workers and temporary agency workers, we exploit the long-run relationship between relative input use and relative costs implied by firms' static cost minimization problem. Although non-convex adjustment costs prevent this condition from holding pointwise, it continues to hold on average in a stationary environment. In addition, variation in relative labor costs induced by the 2006 TAW reform provides an informative source of exogenous shifts along this margin, which we exploit in alternative specifications.

Combining the first-order conditions for labor inputs yields a log-linear relationship between a cost-based measure of TAW intensity and the relative employment of TAWs and regular workers. Estimating this relationship allows us to recover the elasticity of substitution ϕ and the relative productivity parameters of the two labor types. Details of the estimating equation, robustness checks, and alternative specifications are provided in Appendix C. Our preferred estimate implies an elasticity of substitution of $\hat{\phi} = 3.74$, indicating that regular workers and TAWs are imperfect but highly substitutable inputs.

Our estimate is closely aligned with recent frontier work. In particular, Chan (2023) estimates elasticities between labor and purchased inputs in the range of 1.5 to 4 using a task-level

model with explicit make-or-buy decisions, fixed costs, and endogenous selection. Despite the richer structure and heavier data requirements of his framework, our preferred estimate of $\hat{\phi} = 3.74$ lies well within this range. This similarity is reassuring, especially given that our approach relies on a more parsimonious aggregation and exploits equilibrium relationships implied by firms' input choices rather than task-level data. While factor shares in our model need not map one-to-one into technological parameters, the resulting elasticity captures the economically relevant intensive-margin substitution governing firms' labor adjustments. Variation induced by the 2006 TAW reform further supports this interpretation by generating plausibly exogenous shifts in relative labor costs that move firms along this substitution margin without directly affecting relative productivity.

Using observable variables for the median firm in our sample together with the estimated elasticity $\hat{\phi}$, we recover the relative productivity terms implied by the CES labor aggregator. Specifically, we obtain $(a_N)^{1/\hat{\phi}} = 0.98$ for regular workers and $(a_S)^{1/\hat{\phi}} = 0.47$ for temporary agency workers. These values indicate that, conditional on input quantities, regular workers are more productive than TAWs, consistent with observed task allocation within firms.

While routine or short-term tasks can be performed by either worker type, more sensitive or strategic activities—those requiring continuity, loyalty, or firm-specific knowledge—are more frequently assigned to regular workers, raising their effective productivity (Bostanci, 2025). Additional mechanisms reinforce this productivity gap. First, workers employed under temporary or agency contracts accumulate less firm-specific human capital and receive less firm-sponsored on-the-job training, reflecting both shorter expected job duration and weaker incentives for firms to invest in their skills (Garcia-Louzao et al., 2023; Bratti et al., 2021). Second, higher turnover and weaker attachment to the firm may reduce effort when the perceived probability of conversion to a permanent contract is low, further lowering effective productivity (Dolado and Stucchi, 2008; Bryson, 2013).

Productivity persistence (ρ). Finally, we estimate the persistence of firm-level productivity by constructing a measure of total factor productivity using the estimated production parameters and fitting an AR(1) process. We obtain a persistence parameter of $\rho = 0.94$, in line with existing evidence for Chilean manufacturing firms and comparable datasets (Asker et al., 2014; Figal et al., 2020; Eberhardt and Helmers, 2019).

4.3 Fitted Parameters

We take as given the calibrated and previously estimated parameters from the first step and estimate the remaining structural parameters using a Simulated Method of Moments (SMM)

procedure. The set of parameters to be estimated is:

$$\theta_e = \{c_f, \delta, \sigma_\nu, F, \tau, \overline{W}_N, \overline{W}_S, \mu, \sigma, \psi_N, \psi_S, c_s, c_e\}.$$

To implement SMM, we select a vector of empirical moments $\hat{\Psi}^A$ computed from the data and, for each candidate parameter vector θ_e, simulate the model and compute the corresponding simulated moments $\Psi^S(\theta_e; \theta_p)$. We then match the data and simulated moments using the average across simulations. Appendix B.3 provides implementation details (simulation size, burn-in, and numerical settings).

The SMM estimator $\hat{\theta}_e$ minimizes the distance between empirical and simulated moments:

$$\hat{\theta}_e = \arg\min_{\theta_e \in \Theta_e} \left[\hat{\Psi}^A - \bar{\Psi}^S(\theta_e; \theta_p)\right]' W \left[\hat{\Psi}^A - \bar{\Psi}^S(\theta_e; \theta_p)\right], \tag{14}$$

where $\bar{\Psi}^S(\theta_e; \theta_p) \equiv \frac{1}{S}\sum_{s=1}^{S} \Psi_s^S(\theta_e; \theta_p)$ denotes the average simulated moments. In the baseline estimation we set $W = I$, giving equal weight to all moments.

4.3.1 Selection of Moments and Identification

The selection of moments is guided by their informativeness for identifying the key structural parameters of the model. We target moments capturing four dimensions of firm behavior: the cross-sectional distribution of firm size and exit, firms' use of temporary agency workers, employment dynamics over time, and wage schedules by firm size. Table 1 summarizes the mapping between moment groups and the parameters they primarily discipline. Appendix E provides a formal sensitivity analysis of the identification of structural parameters using the methodology proposed by Andrews et al. (2017).

The first group of moments targets the cross-sectional distribution of firm size and exit. We match the distribution of plants and employment across standard firm-size categories (5–19, 20–49, 50–199, and 200+ employees), as well as overall and size-specific exit rates. These moments discipline the fixed operating cost c_f, the entry cost c_e, and the parameters governing the productivity distribution, (μ, σ), which jointly shape market selection and the size distribution of firms. Higher fixed or entry costs raise the productivity threshold for survival and entry, while the mean and dispersion of productivity determine the strength of selection. For instance, stronger selection—driven by higher average productivity or lower dispersion—leads to a higher average plant size and a greater concentration of employment in large, productive firms. Conversely, a lower mean productivity or higher volatility increases exit rates and shifts the firm-size distribution toward smaller establishments.

The second group of moments targets firms' use of agency workers and is informative

Table 1: Identification of Structural Parameters by Moment Groups

Moment group	Targeted moments	Parameters
Firm size and exit	Distribution of firms across size bins Distribution of employment by firm size Overall and size-specific exit rates	c_f, c_e, μ, σ
Use of TAWs	Incidence of TAW use by firm size Conditional TAW employment shares Unconditional TAW employment shares	c_s, σ_ν, δ
Employment dynamics	Volatility of regular employment growth Volatility of TAW employment growth Volatility of total employment growth	τ, F, σ
Wages by firm size	Average regular-worker wages by size Wage distribution across size bins	$\overline{W}_N, \overline{W}_S, \psi_N, \psi_S$

Note: The table summarizes the primary mapping between moment groups and model parameters. Moments may discipline parameters both directly and indirectly through equilibrium interactions; see text for details.

about both the fixed and variable costs associated with this margin. We match the overall incidence of agency-worker use and its sharp increase with firm size, a pattern that disciplines the fixed administrative cost of outsourcing labor, c_s. This cost scales with the regular wage bill and captures administrative, legal, and monitoring costs that arise whenever firms rely on agency labor, regardless of the number of agency workers employed. Empirically, the share of firms using agency workers rises from 7.0% among small plants to over 40% among the largest establishments, while many firms never adopt them at all. This extensive-margin pattern is informative about both the level of c_s and the dispersion of the nonpecuniary adoption taste-shock, captured by σ_ν. The variance of this taste shock governs heterogeneity in agency-worker adoption among observationally similar firms and is disciplined by the coexistence of adopters and non-adopters within narrow firm-size bins.

We also target the average employment share of TAWs conditional on use, and its more modest increase with firm size. These moments are informative about the intensive margin of TAW adjustment and discipline the proportional cost parameter δ, which scales with revenues or profits. This cost limits the optimal reliance on TAWs among firms that already outsource labor and prevents the model from predicting excessive use of agency workers. The relatively smooth increase in the conditional TAW share with firm size reflects both the presence of this proportional cost and differences in wage-setting across worker types. Because the rent-sharing elasticity of wages is larger for regular workers than for TAWs, regular workers' wages rise more rapidly with productivity, making TAWs relatively more cost-effective for larger, more productive firms. Together, these moments identify the fixed cost of adopting TAWs,

the variable cost governing their intensity of use, and the degree of rent sharing for each labor type.

The third group of moments focuses on employment growth volatility for regular workers, agency workers, and total employment. To capture both discrete adoption decisions and intensive-margin adjustments, we measure employment growth using the centered growth-rate measure introduced by Davis and Haltiwanger (1992). This measure is well suited to settings with entry, exit, and zero employment levels, as on/off transitions map to the bounds of the measure, $[-2, 2]$.

In the data, fluctuations in regular employment are smaller and less frequent than those in agency-worker employment, while total employment volatility closely tracks that of regular workers. This pattern is consistent with the presence of adjustment frictions affecting regular workers and with firms relying on agency workers as a more flexible margin of adjustment. Higher productivity volatility raises the value of flexibility, incme greasing firms' reliance on agency labor, while higher firing costs and fixed adjustment costs dampen adjustments along the regular employment margin and shift adjustment toward agency workers. Accordingly, these moments discipline the volatility of idiosyncratic productivity shocks, σ, as well as firing costs for regular workers, τ, and fixed adjustment costs, F. At the same time, volatility in agency-worker employment reflects switching behavior in agency use and is therefore related to the dispersion of the agency-adoption taste shock, σ_ν, although this parameter is primarily disciplined by cross-sectional patterns of adoption across firms.

The final group of moments matches wage levels and wage gradients by firm size for regular workers. These moments identify the baseline wage of regular workers, $\overline{W}_N$, and their rent-sharing elasticity, ψ_N, capturing how wages respond to productivity differences across firms. Because firm size is positively correlated with productivity in the data, variation in regular-worker wages across size bins is informative about the slope of the wage–productivity relationship and, hence, the degree of rent sharing. While wages of temporary agency workers are not directly targeted in the estimation, their level and gradient ($\overline{W}_S$, ψ_S) are disciplined indirectly through the joint fit of regular-worker wages and firms' TAW employment choices by size. Given the model's wage-setting structure, differences in bargaining power across worker types affect relative labor costs and, in turn, firms' optimal reliance on agency labor. As a result, the interaction between regular-worker wage gradients and TAW usage patterns by firm size identifies both the rent-sharing elasticity of TAWs and the relative cost advantage of agency labor as productivity increases.

5 Estimation Results

This section presents the estimation results for the benchmark model, with a focus on the parameters fitted using the Simulated Method of Moments, holding first-stage values fixed. For completeness, we report the full parameter vector, although the predefined parameters and their sources were discussed in detail in Section 4.2. Our emphasis here is on the magnitude and plausibility of the fitted parameters and on how they map into economically meaningful objects such as fixed costs, adjustment frictions, and wage-setting behavior.

Table 2 reports the full parameter vector for the benchmark calibration. Unless otherwise noted, cost parameters are expressed in units of output and evaluated at a plant with productivity close to the mean of the size distribution. Output prices are normalized to one ($P = 1$), so that differences in prices and productivity are absorbed into the productivity term. The model period is annual, but wages are reported in monthly terms to match the institutional setting and measurement in the data.

The fixed operating cost parameter c_f is estimated at 34.30 units of output, while the entry cost c_e is approximately twice as large.[14] Fixed operating costs scale with productivity according to $C_F(Z_t) = Pc_f(Z_t/\overline{Z})$, where $\overline{Z}$ is normalized to a productivity level close to the mean of the plant-size distribution. As a result, the reported values of c_f and c_e correspond to plants near the center of the size distribution, with fixed costs rising (falling) proportionally for more (less) productive firms.[15]

The inspection disruption cost associated with the use of agency labor, captured by the parameter δ and measured as a fraction of firm revenues, is estimated at 4.2%. While this figure is not directly comparable to other structural estimates of similar costs, it falls within a plausible range. For instance, Cooper et al. (2015) estimate that disruption costs to production and sales processes during employment adjustments range from 2.4% to 8.1% of revenues. The fixed cost of managing agency contracts and ensuring compliance with legal obligations, c_s, is estimated at 1.096. This value is expressed in units of a regular worker's annual wage bill and implies that the fixed administrative burden is roughly equivalent to the cost of employing one regular worker for a full year. Because regular workers' wages rise with productivity, this fixed cost of outsourcing labor also scales with plant size.[16]

[14] Because output prices and productivity enter multiplicatively in the firm's objective function, it is not possible to separately identify low (high) prices from low (high) productivity levels. To address this identification issue, we follow Hopenhayn and Rogerson (1993) and normalize output prices by setting $P = 1$. The entry cost c_e is then chosen so that the free-entry condition (equation 13) holds under this normalization.

[15] To provide a sense of magnitude, consider a plant with approximately 10 employees, a size representative of the least productive firms in the data. For such a plant, the implied fixed operating cost is comparable to the annual cost of renting basic production space plus six regular workers for a full year. This comparison is intended solely as an order-of-magnitude benchmark.

[16] Although c_s scales with productivity through the regular wage bill, it remains independent of the firm's

Table 2: Estimated Parameters in the Benchmark Model

Parameter	Description	Value	S.E.
Technology and Preferences:			
β	Discount factor	0.85	–
α	Returns to scale	0.71	–
γ	Labor share in production	0.29	–
ϕ	Elasticity of substitution (labor types)	3.74	–
$(a_N)^{1/\phi}$	Efficiency of regular workers	0.98	–
$(a_S)^{1/\phi}$	Efficiency of TAW	0.47	–
Productivity Process:			
μ	Mean productivity	0.527	0.005
σ	Std. dev. productivity shocks	0.248	0.010
ρ	Productivity persistence	0.94	–
Labor Market Frictions:			
λ_1	Quit rate, tenured workers	0.05	–
λ_2	Separation before tenure	0.08	–
τ	Firing cost (fraction of annual wage)	0.879	0.103
F	Fixed adjustment cost (annual wage units)	1.780	0.886
TAW Costs:			
c_s	TAW administrative cost (annual wage units)	1.096	0.433
δ	TAW disruption cost (share of revenues)	0.042	0.009
σ_ν	Dispersion of TAW adoption taste-shock	0.590	0.143
Wages and Rent Sharing:			
$\overline{W}_N$	Baseline wage, regular workers (monthly)	5.535	0.090
$\overline{W}_S$	Baseline wage, TAW (monthly)	3.689	0.175
ψ_N	Rent-sharing elasticity, regular workers	0.239	0.032
ψ_S	Rent-sharing elasticity, TAW	0.138	0.051
Entry and Operating Costs:			
c_f	Fixed operating cost (output units)	30.98	1.767
c_e	Entry cost (output units)	63.67	–

Note: The table reports the parameter values of the benchmark model. Parameters reported without standard errors in column S.E. are calibrated, estimated outside the SMM procedure, or pinned down by equilibrium conditions, as described in Section 4.2. Parameters with reported standard errors are estimated using the Simulated Method of Moments in Section 4.3. Wages are reported in monthly units; all other parameters are annual. Output prices are normalized to one ($P = 1$).

Penalty costs for firing a regular worker, captured by the parameter τ, are estimated at approximately 88% of the regular-worker wage in the average plant, corresponding to about ten months of the annual wage bill. While substantial, this magnitude is consistent with Chile's stringent employment protection regime. Severance payments increase with job tenure for workers with at least one year of continuous service, subject to a cap of eleven years. Importantly, this estimate captures not only statutory severance payments but also other variable, worker-level costs that scale with the number of layoffs, such as legal expenses, bargaining frictions, and disruptions associated with separating individual workers from the firm. Although these estimates are not directly comparable across studies—given differences in methods, institutional contexts, and data—the implied magnitudes are broadly in line with the existing literature. For example, Aguirregabiria and Alonso-Borrego (2014) estimate hiring and firing costs of approximately 15% and 50% of a worker's annual salary, respectively, in Spain. Roys (2016), using a similar approach, estimates considerably lower adjustment costs in France; around two months of the annual wage bill per worker. Their framework abstracts from the use of temporary contracts as an adjustment margin, which may account for these lower estimates.

The other fixed cost component of adjusting regular workers, F, is estimated at 1.78 annual wage bills of a regular worker, or roughly 21 months of pay. This magnitude is comparable to fixed adjustment costs estimated in other settings, such as Rota (2004), who report median fixed costs on the order of fifteen months of labor costs. In contrast to τ, F captures firm-level costs that arise whenever regular employment is reorganized, regardless of the number of workers adjusted. These include costs associated with restructuring tasks, reallocating responsibilities within the firm, managerial coordination, and temporary disruptions to production during reorganization, rather than marginal firing or severance costs.

The estimated baseline wage of agency workers is about two-thirds that of regular workers. At a mean-size plant, we estimate a monthly wage of 553,500 Chilean pesos for regular workers (5.535 in the table) and 368,900 pesos for agency workers (3.689 in the table), equivalent to approximately USD 13,500 and USD 8,900 annually, respectively. The estimated rent-sharing elasticity for regular workers is 0.239, while for agency workers it is substantially lower, at 0.138. These estimates imply that regular workers' wages respond more strongly to firm-level productivity than those of agency workers, consistent with differences in bargaining power across employment types.[17] Taken together, these estimates imply that agency workers face

current input choices or output level in any given period, preserving its interpretation as a fixed cost.

[17] Our estimated rent-sharing elasticities are somewhat higher than standard benchmarks in the literature. As documented in Card et al. (2018), wage–productivity elasticities typically lie in the range of 0.05–0.15. The higher values we obtain likely reflect the use of plant-level variation and the fact that wages in the model respond to persistent productivity differences rather than transitory shocks. What matters for our analysis is the differential rent sharing across worker types, which shapes relative labor costs and firms' reliance on

both a baseline wage penalty and weaker rent sharing. This pattern is consistent with existing evidence on outsourcing wage gaps and heterogeneous rent-sharing within firms (Dube and Kaplan, 2010; Goldschmidt and Schmieder, 2017; Daruich et al., 2023; Cho and Krueger, 2022). Using comparable administrative data from Argentina, where labor market institutions are broadly similar, Drenik et al. (2023) find that agency workers capture roughly half of the workplace-specific pay premia earned by regular workers, closely aligning with our findings.

5.1 Goodness of Fit: Benchmark Model

Table 3 reports the thirty-seven moments from ENIA data, along with their standard errors, the corresponding simulated moments from the benchmark model, and the squared weighted relative deviations,

$$SQRD_i = (\Psi_i^A - \bar{\Psi}_i^S(\theta_e; \theta_p))^2 \times W_i.$$

The model performs well in replicating the aggregate structure of the industry. It closely matches average firm size and reproduces both the distribution of plants and the distribution of employment across firm-size categories observed in Chilean manufacturing. In particular, it captures well the concentration of employment in large establishments alongside the predominance of small plants, a key feature of the data. The model also matches the overall exit rate and reproduces its strong decline with firm size, although it tends to underpredict exit among larger firms. A key modeling choice underlying this dimension is the treatment of exit costs. We assume that exiting firms incur only a fraction of the total adjustment cost—specifically, 30% of severance obligations—reflecting the idea that firm closure may allow firms to partially avoid or renegotiate separation costs, based on administrative and legal evidence on firm bankruptcies in Chile (Ministerio de Economía, Fomento y Turismo, 2015). While this assumption improves the fit of exit rates, it may still understate exit among larger firms, for whom closure is likely to involve additional costs not explicitly modeled, such as contractual obligations, reputational concerns, or coordination costs associated with downsizing large workforces.

One of the main contributions and strengths of the model lies in its ability to match observed patterns of agency-worker use, both along the extensive and intensive margins. In particular, the model closely replicates the overall incidence of agency-worker adoption and its sharp increase with firm size, a pattern that is central to the identification of the fixed administrative cost of outsourcing labor, c_s, and the dispersion of nonpecuniary adoption incentives, σ_ν. The strong fit along this margin indicates that the model successfully captures the substantial heterogeneity in agency-worker adoption observed among firms of similar size.

agency labor.

The model also performs well in matching outcomes along the intensive margin. It accurately predicts the mean (size-weighted) employment share of agency workers among user firms, as well as the mildly increasing relationship between firm size and the conditional agency-worker share reported in the data. This pattern reflects the interaction between proportional agency-worker costs, captured by δ, and differences in wage-setting across worker types. Because the rent-sharing elasticity is larger for regular workers, their wages rise more rapidly with productivity, making agency labor relatively more attractive for larger, more productive firms. Quantitatively, the model slightly overpredicts the conditional agency-worker share among small plants (20–49 employees), estimating a value of 0.284 compared to 0.251 in the data, and underpredicts the share among medium-sized plants (50–199 employees), estimating 0.312 versus 0.390. These discrepancies suggest that, in the data, the cost advantage of agency labor may increase with firm size at a somewhat faster rate than implied by the model, although the overall fit of both extensive and intensive margins remains strong.

The model also replicates well the volatility of employment growth for regular workers and for total employment. However, it overstates the volatility of agency-worker employment. In the simulations, the variance of agency-worker growth is 3.528, compared to 1.793 in the data. This gap suggests that, in practice, adjustments involving agency workers are either less frequent or less abrupt than implied by the model. Because the centered growth-rate measure assigns extreme values to on/off transitions, agency employment volatility is particularly sensitive to adoption switching and, hence, to the dispersion of the adoption taste shock. In our framework, matching cross-sectional heterogeneity in agency-worker adoption can therefore come at the cost of overstating time-series volatility of agency employment. Sharpening the identification of the adoption taste-shock dispersion with additional cross-sectional moments capturing within-size-bin overlap in adoption is a natural avenue for future work.

Finally, the model closely matches the overall (size-weighted) average wage of regular workers, estimated at roughly 533,000 Chilean pesos per month (about USD 12,900 annually).[18] It also captures the upward wage gradient by firm size, although it does not match wages equally well across all size bins: it overpredicts wages for the smallest plants and underpredicts wages for medium-sized plants. One plausible reason for the remaining misfit is that firm size is an imperfect proxy for productivity in the presence of adjustment frictions, so wages by size bins are sensitive to composition differences in productivity and workforce histories within bins. Notably, the misfit is concentrated in size categories where selection and discrete employment adjustments are most salient, suggesting that composition effects within size bins play an important role.

[18] Wages are reported in monthly Chilean pesos, the customary unit in the data. All wages are expressed in 2013 Chilean pesos, as they are deflated using the annual GDP deflator from the Central Bank of Chile. For international comparability, we use the average 2013 exchange rate of 1 USD = 495 Chilean pesos.

Table 3: Model Fit to Targeted Moments: Benchmark Model

Panel A: Cross-Sectional Firm-Level Outcomes

	Data (1)	S.E. (2)	Model (3)	$SQRD_i$ (4)
Average firm size	79.9	3.944	80.8	0.762
Fraction of plants in each bin:				
5-19 employees	0.421	0.010	0.417	0.000
20-49 employees	0.284	0.009	0.293	0.000
50-199 employees	0.206	0.008	0.215	0.000
200+ employees	0.088	0.006	0.074	0.000
Share of employment in each bin:				
5-19 employees	0.060	0.004	0.062	0.000
20-49 employees	0.108	0.007	0.112	0.000
50-199 employees	0.252	0.016	0.253	0.000
200+ employees	0.580	0.023	0.573	0.000
Exit rate (overall)	0.112	0.005	0.084	0.001
5-19 employees	0.141	0.008	0.176	0.001
20-49 employees	0.110	0.009	0.071	0.002
50-199 employees	0.078	0.009	0.019	0.004
200+ employees	0.044	0.011	0.006	0.001
Distance ($\sum_i SQRD_i$)				0.771

Panel B: Firms' Use of TAWs

	Data (1)	S.E. (2)	Model (3)	$SQRD_i$ (4)
TAW share among users	0.377	0.028	0.381	0.000
5-19 employees	0.295	0.029	0.284	0.000
20-49 employees	0.251	0.022	0.284	0.001
50-199 employees	0.390	0.033	0.312	0.006
200+ employees	0.384	0.037	0.406	0.001
Share of firms with TAWs	0.168	0.008	0.167	0.000
5-19 employees	0.070	0.008	0.066	0.000
20-49 employees	0.158	0.014	0.165	0.000
50-199 employees	0.268	0.019	0.278	0.000
200+ employees	0.436	0.034	0.418	0.000
TAWs as a share of total emp.	0.136	0.014	0.144	0.000
5-19 employees	0.024	0.003	0.021	0.000
20-49 employees	0.043	0.005	0.049	0.000
50-199 employees	0.109	0.013	0.089	0.000
200+ employees	0.177	0.023	0.200	0.001
Distance ($\sum_i SQRD_i$)				0.009

Panel C: Employment Dynamics Over Time

	Data (1)	S.E. (2)	Model (3)	$SQRD_i$ (4)
Volatility of emp. growth:				
Regular workers	0.774	0.023	0.744	0.001
TAW	1.793	0.068	3.528	3.011
Total	0.777	0.023	0.782	0.000
Distance ($\sum_i SQRD_i$)				3.012

Panel D: Regular Workers' Wages by Firm Size

	Data (1)	S.E. (2)	Model (3)	$SQRD_i$ (4)
Regular workers' wage (mean)	5.331	0.173	5.181	0.022
5-19 employees	3.227	0.067	3.971	0.554
20-49 employees	4.314	0.103	4.307	0.000
50-199 employees	5.448	0.175	4.789	0.433
200+ employees	5.752	0.301	5.740	0.000
Distance ($\sum_i SQRD_i$)				1.009

Note: This table reports the fit of the benchmark model to the empirical moments targeted in estimation. Column (1) reports the data moments computed from Chilean manufacturing plants in low-technology industries during 2005–06, with standard errors in column (2) obtained from 1,000 bootstrap replications. Column (3) shows the corresponding simulated moments from the model, averaged over $S = 300$ simulated datasets of size (N, T) with $N = 10{,}000$ firms and $T = 50$ periods. Column (4) reports the squared weighted relative deviation (SQRD), which summarizes each moment's contribution to the SMM objective function. Panel A reports moments on firm size and exit, Panel B on agency-worker adoption and intensity, Panel C on employment growth volatility, and Panel D on regular-worker wages by firm size.

6 Impact of the 2006 TAW Reform

This section analyzes the aggregate consequences of the 2006 reform. First, we quantify how the reform reshaped firms' effective costs of employing temporary agency workers along both the extensive and intensive margins. Second, we assess how these cost changes affected aggregate productivity through changes in allocative efficiency across firms. In our framework, misallocation arises when higher adjustment costs weaken the reallocation of production toward more productive plants. We measure misallocation using the Olley–Pakes covariance between plant-level productivity and output shares, which summarizes the extent to which resources are concentrated in more productive firms.

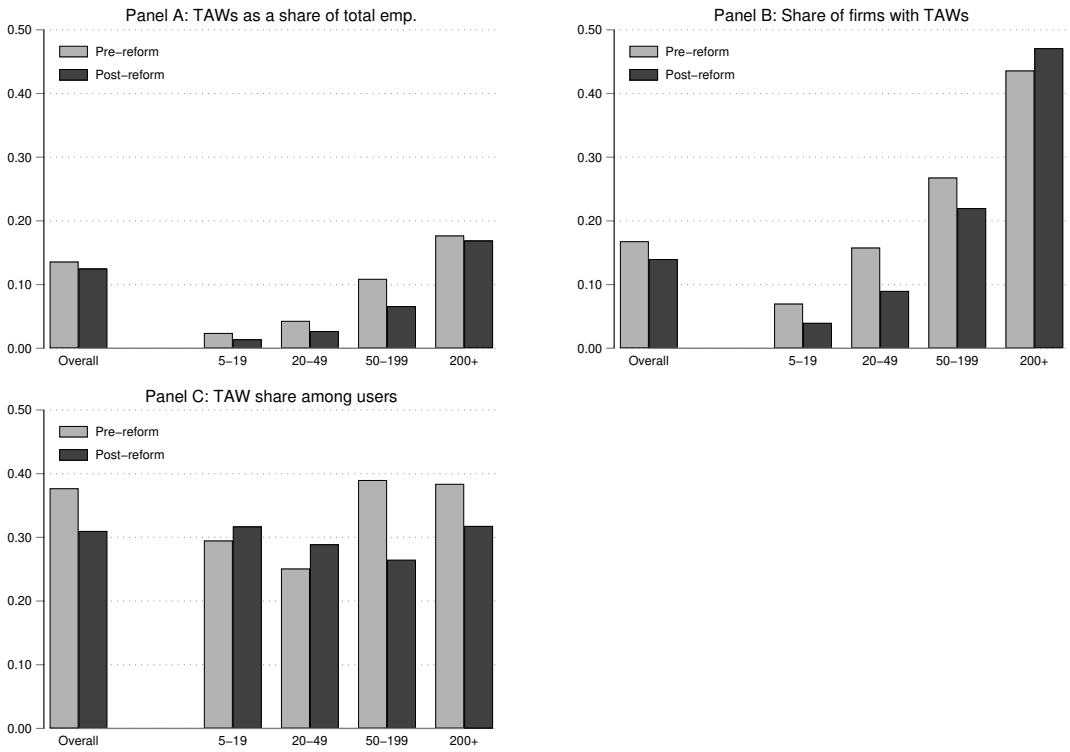
In 2006, Chile introduced a reform to TAW legislation that increased both the fixed and variable costs of employing agency labor. On the one hand, the reform raised variable costs by requiring user plants to provide equal working conditions to TAWs (e.g., food, clothing, and safety equipment), thereby strengthening their bargaining position and increasing per-worker costs. On the other hand, it increased fixed costs by establishing joint and several liability for TAWs' labor obligations (e.g., severance payments and unemployment insurance), which required firms to monitor agencies and ensure regulatory compliance.

Figure 6 illustrates changes in TAW usage before and after the reform. Between 2005–06 (pre-reform) and 2010–11 (post-reform), Panel A shows that the overall share of TAWs in total employment declined from 13.6% to 12.5%, with a relatively smaller decline among large firms (200+ employees). Panel B reports a fall in the fraction of firms employing TAWs, from 16.8% to 14%, driven primarily by small (5–49 employees) and mid-sized plants (50–199 employees), consistent with an increase in fixed adoption costs and a rise in the relative size of firms that continue to rely on TAWs. Panel C shows that, conditional on employing TAWs, the average TAW share per plant declined from 37.7% to 31%, with the largest reductions among mid-sized and large firms. This pattern suggests a reduction in the relative cost advantage of TAWs for larger, more productive firms following the reform.

These changes are consistent with a narrowing of the gap in wage-setting and rent-sharing between regular and agency workers induced by the reform. Because larger firms typically provide more generous benefits to regular workers, the requirement that agency workers receive equivalent labor and safety conditions implies a relatively larger increase in agency-worker costs for bigger firms, precisely along the margins captured by the wage and rent-sharing parameters in the model.

To quantify the change in TAW costs implied by the reform, the next subsection re-estimates the model parameters governing the fixed and variable costs of employing TAWs using post-reform data (2010–11). Specifically, we match a subset of moments related to TAW

Figure 6: TAW usage by firm size Pre- and Post-Reform



Note: The figure reports moments characterizing firms' use of agency workers before and after the 2006 reform, distinguishing between 2005–06 (pre-reform) and 2010–11 (post-reform). Panel A shows the share of agency workers in total employment. Panel B reports the fraction of firms employing at least one agency worker. Panel C shows the share of agency workers among firms that use them.

usage and compare the resulting parameter estimates to their benchmark (pre-reform) values. The differences between these estimates identify the increases in TAW costs attributable to the reform. In the following subsection, we assess how these higher costs affect aggregate productivity and misallocation.

6.1 Estimated Increase in TAW Variable and Fixed Costs

To quantify the increase in the cost of employing agency workers induced by the reform, we re-estimate the model using post-reform data and the same Simulated Method of Moments (SMM) approach. We select a subset of fifteen moments from the benchmark estimation, all related to firms' use of agency workers, now computed using ENIA data for the post-reform period (2010–11). The post-reform parameter vector, $\hat{\theta}_e^{post}$, is obtained by minimizing the weighted distance between these empirical moments and their model-implied counterparts.

We restrict the set of parameters allowed to change after the reform to those directly governing the cost of agency labor, namely $\{W_S, \psi_S, c_s, \delta\}$, while holding all other parameters fixed at their pre-reform benchmark values. This restriction reflects the institutional features of the reform, which primarily altered firms' costs of employing agency labor and operated through short-run adjustment margins, rather than through changes in production technology or long-run profitability affecting entry and exit.

Table 4: Estimated TAW Cost Parameters: Pre- and Post-Reform

Parameter	Description	Estimates	
		2005–06 (Pre-reform) (1)	2010–11 (Post-reform) (2)
δ	TAW disruption cost (share of revenues)	0.042	0.042
		(0.009)	(0.002)
$\overline{W}_S$	TAW baseline wage (monthly)	3.689	3.852
		(0.175)	(0.141)
ψ_S	TAW rent-sharing elasticity	0.138	0.168
		(0.051)	(0.044)
c_s	TAW administrative cost (annual wage units)	1.096	2.161
		(0.433)	(0.420)

Note: The table reports estimated parameters governing the cost of agency labor before and after the 2006 reform. Column (1) reports benchmark parameter estimates and standard errors (in parentheses) based on pre-reform data (2005–06), reproduced from Table 2. Column (2) reports the corresponding estimates using post-reform data (2010–11). All parameters not listed are held fixed at their pre-reform benchmark values.

Table 4 reports the resulting parameter estimates and associated standard errors. The estimates suggest that the reform eroded the cost advantage of employing agency workers along

several dimensions. Three key cost components increased after the reform: the baseline wage paid to agency workers ($\overline{W}_S$), their rent-sharing elasticity (ψ_S), and the fixed administrative cost of managing agency contracts and ensuring regulatory compliance (c_s). Prior to the reform, the fixed administrative cost of agency use was equivalent to slightly more than one full-time worker; after the reform, it exceeded the equivalent of two full-time workers.

Stricter enforcement of labor and safety standards also raised the fees charged by agencies, increasing the effective per-worker cost of agency labor, as reflected in the higher estimated value of $\overline{W}_S$. In addition, the point estimates indicate an increase in the rent-sharing elasticity of agency workers, from $\psi_S = 0.138$ to $\psi_S = 0.168$, consistent with a strengthening of their bargaining position within user firms following the reform. While the increase in the fixed administrative cost of agency use is quantitatively large and precisely estimated, the changes in agency wages and in the rent-sharing elasticity of agency workers are smaller in magnitude and should therefore be interpreted with caution. More broadly, this result is consistent with earlier analyses arguing that partial employment protection reforms can affect the distribution of bargaining power across worker groups, with implications for insider–outsider dynamics within firms (Bentolila and Dolado, 1994; Daruich et al., 2023; Boeri and Jimeno, 2016).

By contrast, the estimated disruption cost associated with labor inspections, δ, remains essentially unchanged. This stability likely reflects offsetting forces: while sanctions associated with non-compliance became more severe after the reform, the effective intensity of inspections may not have increased—and may even have declined—if enforcement relied more on deterrence through higher penalties rather than on inspection frequency. This interpretation is consistent with standard models of regulatory enforcement in which inspection intensity and sanction severity act as substitutes (Polinsky and Shavell, 2000; Kanbur and Ronconi, 2018).

Overall, the direction of these estimated changes is consistent with the institutional features of the reform, which required user firms to monitor agencies' compliance with labor and safety regulations and to cover outstanding obligations—including wages, severance payments, unemployment insurance, and social security contributions—in the event of agency noncompliance. Taken together, the post-reform estimates imply a higher expected cost of agency labor, both through higher fixed compliance costs and through a narrowing of the wage gap between regular and agency workers.

6.1.1 Goodness of Fit: Post-Reform Economy

Table 5 reports the fifteen post-reform moments targeted in estimation, their associated standard errors, the corresponding simulated moments from the model using the post-reform parameter estimates, and the squared weighted relative deviations, $SQRD = (\Psi_i^A - \bar{\Psi}_i^S(\theta_e;\theta_p))^2 \times$

W_i.

Table 5: Model Fit to Agency-Worker Use: Post-Reform Economy

	Data (1)	S.E. (2)	Model (3)	$SQRD_i$ (4)
TAW share among users	0.310	0.027	0.336	0.001
5-19 employees	0.317	0.047	0.273	0.002
20-49 employees	0.289	0.040	0.268	0.000
50-199 employees	0.265	0.037	0.284	0.000
200+ employees	0.318	0.032	0.354	0.001
Share of firms with TAWs	0.140	0.008	0.126	0.000
5-19 employees	0.040	0.007	0.034	0.000
20-49 employees	0.090	0.011	0.118	0.001
50-199 employees	0.220	0.020	0.230	0.000
200+ employees	0.471	0.035	0.355	0.013
TAWs as a share of total emp.	0.125	0.016	0.108	0.000
5-19 employees	0.014	0.003	0.011	0.000
20-49 employees	0.027	0.005	0.035	0.000
50-199 employees	0.066	0.011	0.070	0.000
200+ employees	0.169	0.023	0.150	0.000
Distance ($\sum_i SQRD_i$)				0.020

Note: This table reports the fit of the model to the empirical moments targeted in estimation. Column (1) reports post-reform data moments from Chilean manufacturing plants in the low-technology-intensity group during 2010–11. Column (2) reports bootstrap standard errors based on 1,000 replications. Column (3) shows the corresponding simulated moments from the model, averaged over $S = 300$ simulated datasets of size (N, T) with $N = 10{,}000$ firms and $T = 50$ periods. Column (4) reports the squared weighted relative deviation (SQRD), which summarizes each moment's contribution to the SMM objective function.

The model reproduces well the aggregate share of agency workers in total employment, the fraction of firms employing agency workers, and the average conditional share of agency workers among user firms in the 2010–11 data. Relative to the benchmark estimation, the post-reform model continues to fit aggregate outcomes well, but shows some loss of precision in jointly matching adoption and intensity patterns across firm sizes under the policy-consistent parameter restrictions.

Among large firms, the model generates a higher intensity of agency employment than observed in the data, while at the same time underpredicting the fraction of firms that employ agency workers. This pattern reflects a tension between the extensive and intensive margins: relative to the data, the post-reform estimates imply adoption that is too costly, pushing down extensive-margin use, while conditional intensity is not sufficiently penalized, leading to excessive reliance on agency labor among adopters.

In terms of parameters, this discrepancy is consistent with fixed adoption costs being

relatively high (e.g., a larger c_s) combined with proportional costs of agency labor that are not strong enough to curb intensive use among large firms (e.g., a lower δ or a cost schedule that rises too slowly with firm size through W_S and ψ_S). Allowing the dispersion of the nonpecuniary adoption shock to change post-reform, or introducing additional heterogeneity in compliance costs, could provide additional flexibility to jointly match extensive and intensive margins across firm sizes, and constitutes a natural avenue for future work.

For medium-sized firms (20–199 employees), by contrast, the model predicts a higher overall share of agency workers in total employment than in the data, driven primarily by an overprediction of adoption along the extensive margin. This suggests that, for this segment, the post-reform parameterization may make adoption too easy relative to the data, even though the model matches the intensity of agency use among adopters reasonably well.

6.2 Impact on productivity and misallocation

To quantify the impact of the reform on productivity and misallocation, we follow Olley and Pakes (1996) and decompose aggregate productivity into two sources of variation: the unweighted average of plant-level productivity and the sample covariance between productivity and output shares (which captures misallocation). Consider industry productivity at time t, denoted Z_t, defined as the output-share–weighted average of plant-level productivity:

$$Z_t = \sum_{i=1}^{N_t} \omega_{i,t} Z_{i,t}, \tag{15}$$

were $\omega_{i,t} = \frac{Y_{i,t}}{\sum_j Y_{j,t}}$ is plant i's share of industry output at time t, $Z_{i,t}$ is its productivity level and N_t the number of plants at time t.

We can decompose Z_t into two terms as follows:

$$\begin{aligned} Z_t &= \sum_{i=1}^{N_t} \left(\bar{Z}_t + \Delta Z_{i,t}\right)\left(\bar{\omega}_t + \Delta\omega_{i,t}\right) \\ &= N_t \bar{Z}_t \bar{\phi}_t + \sum_{i}^{N_t} \Delta\omega_{i,t}\Delta Z_{i,t} \\ &= \bar{Z}_t + \sum_{i}^{N_t} \Delta\omega_{i,t}\Delta Z_{i,t}, \end{aligned} \tag{16}$$

where $\bar{Z}_t$ and $\bar{\omega}_t$ denote the unweighted means of productivity and output shares at time t, $\Delta Z_{i,t} = Z_{i,t} - \bar{Z}_t$ and $\Delta\omega_{i,t} = \omega_{i,t} - \bar{\omega}_t$. By construction, these deviations sum to zero, so the

cross terms vanish: $\bar{\omega}_t \sum_i \Delta Z_{i,t} = 0$ and $\bar{Z}_t \sum_i \Delta\omega_{i,t} = 0$. Since output shares sum to one, the average output share is $\bar{\omega}_t = 1/N$, which implies $N_t\bar{\omega}_t = 1$.

This decomposition implies that increases in aggregate productivity Z_t can arise from two channels: (i) growth in average plant-level productivity ($\bar{Z}_t$), or (ii) more efficient reallocation of resources toward productive firms, measured by the covariance term $\sum_i^{N_t} \Delta\omega_{i,t}\Delta Z_{i,t}$. The larger this covariance, the greater the share of output concentrated in more productive firms, and the lower the degree of misallocation.

Aggregate effects of the reform (benchmark elasticity). To measure the reform's impact on aggregate productivity using the benchmark elasticity of substitution, $\phi = 3.74$, we simulate the stationary industry equilibrium under two scenarios: the benchmark pre-reform economy and a counterfactual economy in which the parameters governing the cost of temporary agency work are adjusted to reflect the increased regulatory burden introduced by the 2006 reform, as estimated in Table 4. Comparing these two equilibria allows us to isolate the reform's aggregate effects and to decompose them into changes in average plant-level productivity and changes in allocative efficiency across firms.

Table 6, column (1), reports the results using the benchmark elasticity of substitution ($\phi = 3.74$). When the cost of employing temporary agency workers rises by the magnitude implied by the 2006 reform, output prices increase by 0.49%. For demand elasticities in the range $\eta = 2$ to $\eta = 3$, this price increase implies a contraction in output of approximately 1–1.5%, indicating that the reform had economically meaningful aggregate effects.[19]

The increase in TAW costs raises both fixed and marginal production costs, with the burden falling disproportionately on firms that initially relied more heavily on temporary agency workers.[20] The productivity decomposition reveals that aggregate productivity declines by 0.24%, entirely due to an increase in misallocation, as captured by the covariance term. Average plant-level productivity remains essentially unchanged, consistent with the reform not being large enough to materially affect firms' entry and exit decisions. Intuitively, as temporary agency work becomes more costly, firms lose an important margin to buffer idiosyncratic productivity shocks, spending longer periods operating away from their optimal scale. This loss of flexibility weakens the reallocation of output toward more productive firms and drives the observed decline in aggregate productivity.

Interpreting output prices as being determined by marginal costs, the difference between

[19]The demand function does not affect the aggregate price, which is pinned down by firms' marginal costs; it only influences the number of firms operating in equilibrium.

[20]Importantly, the reform does not reduce the inherent flexibility of TAW contracts by imposing firing costs. Instead, it raises both the fixed and variable costs of accessing this flexible margin of adjustment. As a result, even holding TAW employment constant, higher TAW costs distort firms' input choices and worsen allocative efficiency.

the increase in prices and the decline in aggregate productivity reflects the direct cost effect of the reform. In the benchmark case, higher wages, benefits, and administrative costs associated with temporary agency work raise within-firm labor costs by about 0.25 percent. Roughly half of the total increase in output prices therefore operates through higher marginal costs at the firm level, while the remaining half reflects increased misallocation arising from the erosion of firms' adjustment margins. Because tighter adjustment constraints hinder the reallocation of production toward more productive plants, they raise aggregate marginal costs and, in turn, output prices.

Aggregate effects under lower substitutability. We next examine how the estimated effects would differ under an elasticity of substitution closer to Cobb–Douglas, a common assumption in the literature. The exact Cobb–Douglas case ($\phi = 1$), frequently adopted in related work (see, e.g., Asker et al., 2014; Cahuc et al., 2016), is not feasible in our setting, as it implies that firms must employ at least one temporary agency worker to produce positive output, while many firms in our data operate without TAWs. We therefore consider $\phi = 1.5$, which represents the lowest elasticity consistent with the observed prevalence of zero-TAW firms, while remaining close to a Cobb–Douglas technology.

To implement this exercise, we re-estimate the model under $\phi = 1.5$ for both the pre- and post-reform economies. For the pre-reform period (2005–06), all remaining parameters are estimated using SMM and the same set of moments as in the benchmark case. For the post-reform economy, we follow the same identification strategy as in the benchmark analysis and allow only the parameters directly affected by the reform—those governing the cost of employing TAWs ($\overline{W}_S$, ψ_S, c_s, and δ)—to change, while holding all other parameters fixed. Supplemental Appendix C reports the full set of results.

Table 6, column (2), reports the results for $\phi = 1.5$. Under this lower elasticity of substitution, the misallocation channel is substantially weaker: aggregate productivity declines by only 0.10% (compared to 0.24% in the benchmark), reflecting the fact that firms rely less on TAWs ex ante and therefore lose less flexibility when their costs increase. By contrast, the direct cost channel becomes more important. With limited substitutability, higher TAW costs translate more directly into higher marginal costs, raising prices by 0.37% through increased labor costs alone, compared to 0.25% in the benchmark case.

Taken together, these results underscore that the aggregate consequences of regulating flexible labor contracts depend critically on the elasticity of substitution between labor types. From an economic perspective, the elasticity of substitution between regular and agency workers, ϕ, captures how easily firms can reallocate tasks across contractual forms in response to shocks. A higher elasticity reflects production environments in which tasks are more modular,

Table 6: Aggregate Effects of the 2006 TAW Reform: Productivity and Misallocation

	Percent change (post-reform vs. pre-reform)	
	$\phi = 3.74$ (1)	$\phi = 1.5$ (2)
Output price	0.49	0.47
Aggregate productivity: Z_t	-0.24	-0.10
Average productivity: $\bar{Z}_t$	0.02	0.00
Olley–Pakes covariance (misallocation): $\sum_i \Delta\omega_i \Delta Z_i$	-0.26	-0.10

Note: The table reports percent changes (post-reform vs. pre-reform) implied by the estimated increase in the regulatory burden associated with the 2006 TAW reform. Aggregate productivity Z_t is decomposed into average productivity $\bar{Z}_t$ and the Olley–Pakes covariance term $\sum_i \Delta\omega_i \Delta Z_i$, where ω_i denotes firm output shares, $\Delta Z_i = Z_i - \bar{Z}_t$, and $\Delta\omega_i = \omega_i - \bar{\omega}_t$. Column (1) uses the benchmark elasticity of substitution, $\phi = 3.74$; column (2) presents results for $\phi = 1.5$. All results are from model simulations.

monitoring is easier, and flexible labor can substitute for regular employment with limited losses in productive efficiency.

Assuming a Cobb–Douglas–like technology, as much of the literature does, would therefore understate the role of misallocation while overstating the importance of direct cost increases. By contrast, our benchmark estimate with a higher elasticity reveals that restricting access to flexible labor arrangements can generate sizable productivity losses through worsened allocative efficiency, even when average firm productivity remains unchanged.

While higher substitutability mechanically amplifies misallocation when adjustment margins are restricted, this mechanism is not transparent in static or reduced-form frameworks, nor in models that impose Cobb–Douglas technologies. In such settings, changes in ϕ primarily affect aggregate levels but do not alter the decomposition of productivity into average productivity and reallocation components.

These findings speak to a broader trade-off in labor market policy. While regulating flexible contracts can improve worker protection, doing so in the presence of substantial employment protection for permanent workers may generate unintended efficiency costs by distorting firms' adjustment margins. Although our analysis abstracts from labor supply responses and other general equilibrium channels, the dynamic misallocation mechanism we highlight is likely to operate in a wide range of institutional settings.

7 Concluding Remarks

This paper studies the aggregate consequences of regulating flexible labor contracts in environments characterized by strong employment protection for regular workers. We focus on Chile's 2006 reform regulating TAWs, which substantially increased both the fixed and variable costs of employing agency labor while leaving the broader institutional environment unchanged. To analyze firms' responses to this regulatory shock, we develop and estimate a dynamic model of heterogeneous firms that choose between regular workers and agency workers in the presence of non-convex labor adjustment costs. The model is disciplined using plant-level data that uniquely link agency workers to user firms, allowing us to quantify how changes in the cost of flexible labor arrangements affect firms' adjustment margins and aggregate outcomes.

Our results show that the reform led to a meaningful increase in output prices and a decline in aggregate productivity. Importantly, this productivity loss is not driven by changes in average firm-level productivity, but rather by increased misallocation across firms. As access to temporary agency work becomes more costly, firms lose a key margin to buffer idiosyncratic shocks without triggering costly firing decisions for regular workers. This erosion of adjustment flexibility weakens the reallocation of production toward more productive plants, generating aggregate productivity losses even in the absence of technological change.

The core mechanism highlighted in the paper operates through the interaction between non-convex employment protection and flexible labor arrangements. In our framework, firing costs for regular workers generate regions of inaction and intertemporal trade-offs in firms' labor choices. Temporary agency work serves as an adjustment margin that allows firms to respond to transitory shocks without crossing these adjustment thresholds. By raising the cost of this margin, the reform effectively tightens firms' adjustment constraints, amplifying misallocation in equilibrium. Capturing this mechanism requires a dynamic framework with non-convex adjustment costs; it cannot be recovered from static or reduced-form approaches that abstract from intertemporal adjustment and inaction regions.

A key determinant of the magnitude of these effects is the elasticity of substitution between regular and agency workers. Economically, this elasticity captures how easily firms can reallocate tasks across contractual forms in response to shocks. When substitutability is high, flexible labor arrangements closely substitute for regular employment, making the loss of adjustment margins particularly costly in terms of allocative efficiency. Under a Cobb–Douglas specification, as commonly assumed in the literature, the misallocation channel is substantially muted and the productivity effects of regulation appear modest. By contrast, our benchmark estimates reveal that higher substitutability can generate sizable productivity losses through worsened allocation, even when average firm productivity remains unchanged.

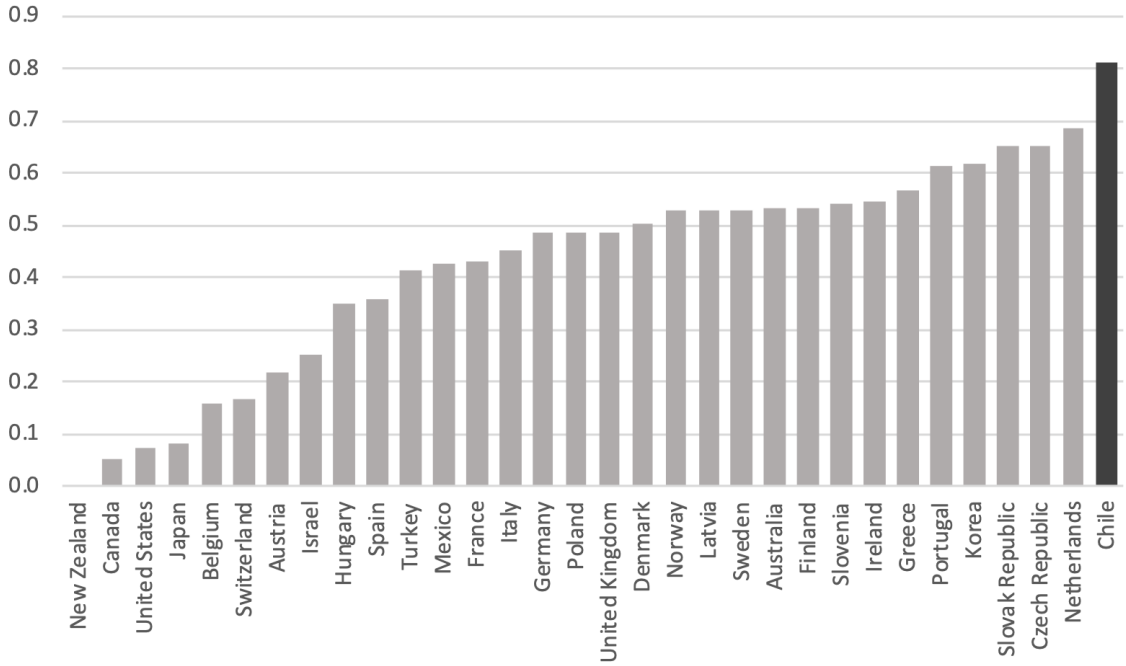
Our analysis abstracts from several potentially important channels. In particular, we do not quantify the welfare benefits of regulation associated with improved worker protection, changes in labor supply, or reductions in monopsony power. We also abstract from capital adjustment as an additional margin of response. While capital may play a role in longer-run firm dynamics, the 2006 reform primarily affected short-run labor adjustment margins by increasing the cost of temporary agency work. Given that capital typically faces its own non-convex adjustment costs and is characterized by lower short-run elasticity than labor, it is unlikely to serve as a seamless substitute for the buffer role provided by TAWs in response to transitory shocks. If anything, capital-labor complementarities suggest that restricting labor flexibility could further discourage capital deepening, potentially amplifying our estimated productivity losses. Consequently, our findings should be interpreted as a focused quantification of the efficiency costs of regulating flexible labor contracts through the lens of firms' adjustment margins, rather than as a full welfare evaluation of the reform, providing a necessary building block for a comprehensive assessment.

More broadly, our results suggest that labor market reforms cannot be evaluated in isolation and highlight a fundamental trade-off in labor market policy design. Regulating flexible labor contracts can improve labor standards, but when combined with stringent employment protection for permanent workers, such regulation may generate unintended efficiency costs by distorting firms' adjustment margins and worsening allocative efficiency. This dynamic misallocation mechanism is likely relevant in any institutional setting where high firing costs coexist with flexible labor arrangements. Understanding how different regulatory instruments interact is therefore essential for evaluating the aggregate consequences of labor market reforms.

Appendix

A Additional Tables and Figures

Figure A.1: Cost of Firing Workers in OECD Countries



Note: the cost of firing workers measures the cost of firing 20% of the firm's workers (10% are fired for redundancy and 10% without cause). The cost of firing a worker is calculated as the sum of the notice period, severance pay, and any mandatory penalties established by law or mandatory collective agreements for a worker with three years of tenure with the firm. If dismissal is illegal, the cost of firing is set equal to the annual wage. The new wage bill incorporates the normal wage of the remaining workers and the cost of firing workers. The cost of firing workers is computed as the ratio of the new wage bill to the old one.
*S*ource: Botero et al. (2004).

B Model Details and Numerical Solution

B.1 Stationary Distribution

In this model, the state of an individual firm is fully characterized by its productivity and employment status, (Z_t, L_t). The state of the industry is then described by the distribution of firms across these state variables. Let $\mu_t(Z_t, L_t)$ denote the measure of incumbent firms at the beginning of period t, after they have made their exit/stay decisions and observed their productivity draws Z_t. Let E_t denote the mass of entering firms.

Table A.1: Descriptive statistics

Variables	Obs.	Mean	Std. Dev.	Min	Max
Panel A: All plants					
Total emp.	6,261	79.91	186.30	1.00	3489
Regular emp.	6,261	68.73	159.69	0.00	2785
TAW	6,261	11.18	69.19	0.00	2359
Dummy var. TAWs	6,261	0.17	0.37	0.00	1.00
Share of TAWs (simple average)	6,261	0.06	0.17	0.00	0.99
Nominal Output (log)	6,261	13.40	1.89	8.86	19.26
Labor Share	6,261	0.24	0.16	0.01	0.91
Inputs Share	6,261	0.47	0.20	0.01	0.89
Panel B: Plants with TAW					
Total emp.	1,044	173.03	304.45	4.00	3489
Regular emp.	1,044	108.07	214.64	0.00	2653
TAWs	1,044	50.41	158.78	1.00	2359
Share of TAWs (simple average)	1,044	0.33	0.30	0.00	1
Nominal Output (log)	1,044	14.55	1.86	9.01	19
Labor Share	1,044	0.22	0.15	0.01	0.91
Inputs Share	1,044	0.47	0.19	0.01	0.89

Note: Sample period 2005–6. Includes only plants in sectors in the low-technology intensity group.

The law of motion for the distribution of firms is given by

$$\mu_{t+1}(Z_{t+1}, L_{t+1}) = \int [1 - X(Z_t, L_t; \vec{p})]\, F(Z_{t+1} \mid Z_t) \cdot \Pr(L_{t+1} \mid N(Z_{t+1}, L_t; \vec{p}), L_t)\, d\mu_t(Z_t, L_t) + \int E_t \cdot \Pr(L_{t+1} \mid N(Z_{t+1}, 0; \vec{p}), 0)\, dG(Z_{t+1}). \tag{B.1}$$

In this expression the first terms captures the transition of incumbent firms that choose to remain in the industry. These firms observe their new productivity Z_{t+1}, choose employment $N_t(Z_{t+1}, L_t)$ and stochastically evolve into states (Z_{t+1}, L_{t+1}). The evolution of tenured workers L_t reflects both endogenous hiring decisions and exogenous transitions, as newly hired workers may gain tenure with probability $1 - \lambda_2$, and tenured workers may quit with probability λ_1.

The second term accounts for entrants, who begin with no tenured workers and draw productivity from the invariant distribution $G(Z_{t+1})$. Like incumbents, they choose employment $N_t(Z_{t+1}, L_t = 0)$ and stochastically transition into future states (Z_{t+1}, L_{t+1}), based on the same tenure dynamics.

In a stationary equilibrium, the distribution of firms over states remains constant over time, i.e., $\mu_{t+1} = \mu_t = \mu^*$. This stationary distribution over the employment–productivity state space is shaped by the stochastic process governing productivity, the endogenous evolution of tenured employment (driven by hiring decisions and worker flows), firm exit behavior

(selection), the productivity distribution of entrants, and the mass of entry. Once μ^* is known, all aggregate outcomes in the industry can be computed from it.

Aggregate supply in the industry is given by the following expression:[21]

$$\begin{aligned} Q^s(\mu^*, E^*; \vec{p}) = &\int f(Z, N(Z, L; \vec{p}), S(N, Z; \vec{p}), M(N, Z; \vec{p}))\, d\mu^*(Z, L) \\ &+ E^* \int f(Z, N(Z, 0; \vec{p}), S(N, Z; \vec{p}), M(N, Z; \vec{p}))\, dG(Z). \end{aligned} \tag{B.2}$$

Aggregate demand for the industry is represented by a standard downward-sloping demand curve, $Q^d = D(p)$.[22]

B.2 Definition of Equilibrium

A ***stationary industry equilibrium*** with positive entry and exit is a set of value functions and decision rules, an output price $P^* \geq 0$, wages W_N^*, W_S^*, and a materials price P_M^*, a stationary distribution of firms μ^*, and a mass of entrants $E^* > 0$ such that the following conditions hold:

i. **Firm optimization:** Given prices $(P^*, W_N^*, W_S^*, P_M^*)$, the value and policy functions solve the firm's dynamic problem.

ii. **Market clearing:** Product markets clear:

$$P^* = D(Q^*) \quad \text{and} \quad Q^* = Q^s(\mu^*, E^*; \vec{p}^*).$$

iii. **Stationary distribution:** The distribution of firms over states is invariant:

$$\mu^* = T(\mu^*, E^*; \vec{p}^*).$$

iv. **Free-entry condition:** The expected value of entry equals the entry cost:

$$V^e(\vec{p}^*) = P^* c_e.$$

[21] Since we focus on the stationary equilibrium, where the distribution of firms and the entry mass have converged to their steady-state values ($\mu_t = \mu^*$ and $E_t = E^*$), and prices are constant ($\vec{p}_t = \vec{p}$), we omit time subscripts in the expression for aggregate supply.

[22] Note that while we include aggregate supply and demand to close the model, our primary interest is in firms' employment decisions, rather than in the overall size of the industry or the determination of the equilibrium price.

B.3 Solution Method

The model has no closed-form solution and is therefore solved numerically using the following iterative procedure:

Step 1: Solve for the equilibrium output price P^* that satisfies the free-entry condition

(1) Guess initial price P_0 and solve incumbent's Bellman equation for the value function $V(Z', L; P_0)$

(2) Check that price P_0 satisfies the free entry condition:

$$EC(P_0) \equiv \int V(Z', 0; P_0) dG(Z)/P_0 - c_e = 0.$$

(3) If yes, set $P^* = P_0$ and proceed to Step 2. If no, update the guess to P_1 and return to (1). For P_1, if $EC(p_0) > 0$, then set $P_1 < P_0$, otherwise set $P_1 > P_0$, and repeat.

Step 2: Solve for the stationary distribution

$$\mu^*(Z, L)$$

Given P^* that satisfies the free-entry condition and the associated value and policy functions of incumbent firms, solve for the stationary distribution $\mu^*(Z, L)$, assuming an arbitrary initial mass of entrants $E_0 = 1$.

Step 3: Find the equilibrium mass of entrants E^* that clears the goods market

(1) Given $\mu^*(Z, L)$ for Step 2, check if with the initial measure $E_0 = 1$ the market clearing condition is satisfied:

$$EQ(\mu, E_0) \equiv Q^d - Q^s(\mu(E_0), E_0; P^*) = 0.$$

(2) If no, return to (1) with a new guess E_1. For E_1, if $EQ(\mu, E_0) > 0$, then set $E_1 > E_0$, otherwise set $E_1 < E_0$. When $EQ(\mu, E_i) \approx 0$ then $(\mu, E_i) = (\mu^*, E^*)$.

The model period is one year. We assume firm-level productivity evolves according to an AR(1) process in logs:

$$\log Z_t = \mu + \rho \log Z_{t-1} + \sigma \varepsilon_t$$

where μ is a constant, ρ is the persistence of productivity, and ε_t is a standard normal distribution random shock. To approximate the distribution of idiosyncratic shocks, we use the quadrature-based method of Rouwenhorst (1995), which is well-suited for highly persistent processes. We set the number of grid points for productivity to $g_z = 97$. The upper bound of the productivity grid corresponds to firm sizes between 2,400 and 5,000 workers and is not binding.

The initial distribution $G(Z_0)$, is set to the stationary distribution implied by the AR(1) process. This distribution closely matches the empirical size distribution of firms aged zero to one year in the data.

To discretize the space for the choice variable N_t, we construct a grid with $g_N = 686$. The grid for N_t includes 100 equally spaced points from 0 to 99, 67 equally spaced points from 102 to 300, and 519 equally spaced points from 309 to 4,971. Since regular employment is an endogenous variable, we verify that the number of grid points does not affect the results. Sensitivity analysis confirms that the chosen discretization is adequate.

C Parameters of the production function

This section provides the formal derivations and additional details underlying the identification of the production technology parameters discussed in Section 4 of the main text.

Static First-Order Conditions. The static first-order conditions (FOCs) with respect to materials M_t and temporary agency workers (TAWs) S_t are given by:

$$\frac{P_M M_t}{PY_t} = \alpha(1-\gamma) \quad \text{and} \quad \frac{W_S S_t}{PY_t} = \alpha\gamma\Omega_t \tag{C.3}$$

where

$$\Omega_t = \frac{a_S^{1/\phi} S_t^{(\phi-1)/\phi}}{a_N^{1/\phi} N_t^{(\phi-1)/\phi} + a_S^{1/\phi} S_t^{(\phi-1)/\phi}}$$

denotes the share of effective labor services provided by agency workers.

In the absence of adjustment costs, the FOC with respect to regular workers is:

$$\frac{W_N N_t}{PY_t} = \alpha\gamma(1-\Omega_t). \tag{C.4}$$

Together, equations (C.3) and (C.4) imply that the sum of expenditures on variable inputs

as a share of revenue equals the returns-to-scale parameter:

$$\alpha = \frac{W_S S_t + W_N N_t + P_M M_t}{P Y_t}. \tag{C.5}$$

In the presence of non-convex adjustment costs for regular workers, this condition need not hold pointwise. However, as shown by Bertola and Caballero (1990), it continues to hold on average in a stationary equilibrium without aggregate productivity trends. With adjustment costs, firms deviate from static optimal employment in individual periods, but these deviations average out over time in a steady state with stationary productivity.

Accordingly, we discipline the curvature of the profit function by setting $\alpha = 0.71$, equal to the median ratio of labor compensation and materials expenditures to nominal output in the data. We use the median rather than the mean to reduce sensitivity to outliers. This value closely matches the estimate reported by Asker et al. (2014) for Chilean manufacturing plants.[23]

Elasticity of substitution between regular workers and TAW. Combining the FOCs in equation (C.3) yields a relationship linking the firm's choice between TAWs and regular workers to relative costs and relative input use:

$$\left(\frac{W_S S_{it}}{\frac{\gamma}{1-\gamma} P_M M_{it} - W_S S_{it}} \right) = \left(\frac{a_S}{a_N} \right)^{\frac{1}{\phi}} \left(\frac{S_{it}}{N_{it}} \right)^{\frac{\phi-1}{\phi}}. \tag{C.6}$$

Because non-convex adjustment costs prevent static labor FOCs from holding pointwise, equation (C.6) should be interpreted as a long-run equilibrium relationship rather than as a per-period optimality condition. Taking logs on both sides yields the estimating equation:

$$\log \left(\frac{W_S S_{it}}{\frac{\gamma}{1-\gamma} P_M M_{it} - W_S S_{it}} \right) = \alpha_0 + \beta \log \left(\frac{S_{it}}{N_{it}} \right). \tag{C.7}$$

where $\alpha_0 = \frac{1}{\phi} \log\left(\frac{a_S}{a_N}\right)$ and $\beta = \frac{\phi-1}{\phi}$. The coefficient β identifies the elasticity of substitution as $\phi = 1/(1-\hat{\beta})$, while the intercept recovers the relative productivity ratio a_S/a_N, given $\hat{\phi}$ and the normalization $a_N + a_S = 1$.

Estimation results. Table C.2 reports estimates of equation (C.7). Columns (1) and (2) present OLS estimates using plant-year observations in which both regular workers and TAWs

[23] Asker et al. (2014) employ a monopolistic competition framework, whereas we assume perfect competition with decreasing returns to scale. The two formulations are equivalent when $\alpha = 1/\mu$, where μ denotes the gross markup.

are employed. The left-hand side is constructed using $\frac{\hat{\gamma}}{1-\hat{\gamma}} P_M M_{it}$, while column (2) uses the theoretically equivalent expression $\hat{\gamma} PY_{it}$ to assess robustness. Column (3) addresses potential endogeneity in $\log(S_{it}/N_{it})$ using an instrumental variables strategy based on lagged relative employment and changes in total employment. Across these specifications, the estimated slope coefficient is stable and close to -0.7, implying an elasticity of substitution around 3.3.

We further estimate the relationship in first differences using data from the years surrounding the 2006 reform, which generated exogenous shifts in relative labor costs. Columns (4) and (5) report OLS and quantile regression estimates for 2006–2007. Using the quantile regression specification in column (5), our preferred estimate is $\hat{\phi} = 3.74$. Columns (6)–(9) confirm that this estimate is robust to perturbations of $\hat{\gamma}$ by $\pm 5\%$.

Table C.2: Estimation Results of the Elasticity of Substitution Between TAWs and Regular Workers from Equation (C.7)

Model:	OLS		IV	OLS	QR	OLS	QR	OLS	QR
			$\hat{\gamma} = 0.29$			$\hat{\gamma} = 0.29 - 5\%$		$\hat{\gamma} = 0.29 + 5\%$	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
$log(S_{it}/N_{it})$	0.707***	0.694***	0.715***						
	(0.0341)	(0.0319)	(0.0996)						
$\Delta log(S_{it}/N_{it})$				0.759***	0.731***	0.753***	0.710***	0.735***	0.714***
				(0.0627)	(0.0545)	(0.0698)	(0.0628)	(0.0563)	(0.0545)
Year FE	Yes	Yes	Yes	No	No	No	No	No	No
Plant FE	Yes	Yes	Yes	No	No	No	No	No	No
Observations	2,355	2,475	1,540	259	259	256	263	256	263
$F - Stat$			51.48						
p-value J test			0.466						
Sample years	2001–6	2001–6	2001–6	2007	2007	2007	2007	2007	2007

Note: All columns estimate equation (C.7). The left-hand side (LHS) is computed using $\frac{\hat{\gamma}}{1-\hat{\gamma}} P_M M_{it}$, except in column (2), which uses the theoretically equivalent expression $\hat{\gamma} PY_{it}$ to assess robustness to using different data inputs. Columns (1) and (2) report OLS estimates using plant-year observations that employ both TAWs and regular workers. Column (3) uses an instrumental variables (IV) approach, where $\log(S_{it-1}/N_{it-1})$ and $\Delta \log(S_{it} + N_{it})$ are instruments for $\log(S_{it}/N_{it})$. Columns (4) and (5) estimate the model in first differences using OLS and Quantile Regression, respectively, with data from 2006–2007. Columns (6) to (9) replicate the specifications in columns (4) and (5), varying $\hat{\gamma}$ by $\pm 5\%$ to assess robustness. All regressions are restricted to sectors classified as low-technology intensity. Robust standard errors are reported in parentheses. Significance at *10 percent, **5 percent and ***1 percent.

Relative productivity. Using observable variables for the median firm together with the estimated elasticity $\hat{\phi}$, we recover relative productivity terms implied by the CES labor aggregator: $(a_N)^{1/\hat{\phi}} = 0.98$ for regular workers and $(a_S)^{1/\hat{\phi}} = 0.47$ for TAWs. These values imply higher effective productivity for regular workers, consistent with observed task alloca-

tion within firms (Bostanci, 2025). Additional mechanisms reinforce this gap, including lower firm-sponsored training and higher turnover among agency workers (Garcia-Louzao et al., 2023; Bratti et al., 2021; Bryson, 2013; Dolado and Stucchi, 2008).

D Simulated Method of Moments Estimation

This section describes the implementation details of the Simulated Method of Moments (SMM) procedure used to estimate the remaining structural parameters of the model discussed in Section 4 of the main text.

We take as given the parameters calibrated or estimated in the first step and estimate the remaining parameter vector

$$\theta_e = \left\{ c_f, \delta, \sigma_\nu, F, \tau, W_N, W_S, \mu, \sigma, \psi_N, \psi_S, c_s, c_e \right\}.$$

Simulation procedure. For any candidate parameter vector θ_e, we solve the model and generate simulated firm-level data to compute the model-implied moments. Each simulation consists of $T = 50$ periods with $N = 10{,}000$ active firms per period. To ensure convergence to the stationary distribution, we discard an initial burn-in period of 20 periods. For each parameter vector, we generate $S = 300$ independent simulated datasets and compute the average of the simulated moments,

$$\bar{\Psi}^S(\theta_e; \theta_p) \equiv \frac{1}{S} \sum_{s=1}^{S} \Psi_s^S(\theta_e; \theta_p).$$

To facilitate comparability across simulations, we fix the random seed for the productivity shocks.

Objective function. The SMM estimator minimizes the distance between the empirical moments $\hat{\Psi}^A$ and their simulated counterparts:

$$\hat{\theta}_e = \arg\min_{\theta_e \in \Theta_e} \left[\hat{\Psi}^A - \bar{\Psi}^S(\theta_e; \theta_p) \right]' W \left[\hat{\Psi}^A - \bar{\Psi}^S(\theta_e; \theta_p) \right],$$

where W is a positive-definite weighting matrix. In the baseline estimation, we set $W = I$, assigning equal weight to all moments. The minimization problem is solved using a Nelder–Mead simplex algorithm. To mitigate concerns about local minima, the optimization is initialized from 1,000 distinct starting values drawn from the parameter space.

Inference. Standard errors are computed using the asymptotic variance formula for SMM estimators (Gouriéroux and Monfort, 1996):

$$V_S(W) = (J'WJ)^{-1} J'WSWJ (J'WJ)^{-1},$$

where S denotes the variance–covariance matrix of the empirical moments and $J = E\left(\frac{\partial \Psi^S(\theta_e)}{\partial \theta_e}\right)$ is the Jacobian of simulated moments with respect to parameters.

We estimate S using 1,000 bootstrap replications with replacement on the original data, resampling entire firm-level time series to preserve the dependence structure. Because the value functions are non-differentiable due to adjustment costs, the Jacobian J is computed numerically using finite differences. For each parameter, we evaluate derivatives at perturbations of $\pm 1\%$, $\pm 2\%$, and $\pm 3\%$ around the midpoint of the parameter space and take the median derivative, following Bloom (2009).

E Identification and Sensitivity of Fitted Parameters

To assess whether the selected moments are informative for identifying the structural parameters of the model, we compute the local sensitivity of the estimated parameters to moments following Andrews et al. (2017). Specifically, we evaluate how marginal perturbations in each empirical moment affect the estimated parameter vector. We report this sensitivity measure in elasticity form, as in Kogan et al. (2020).

Formally, the elasticity of parameter θ_i with respect to moment Ψ_j is given by

$$\hat{\lambda}_{ij} = \Lambda_{ij} \frac{\Psi_j^S(\hat{\theta})}{\theta_i},$$

where Λ_{ij} denotes element (i, j) of the sensitivity matrix

$$\Lambda = (J'WJ)^{-1} J'W,$$

$J = E\big(\partial \Psi^S(\theta_e)/\partial \theta_e\big)$ is the Jacobian of simulated moments with respect to parameters, and W is the weighting matrix used in the SMM objective function. The reported elasticities therefore measure the percentage change in parameter θ_i induced by a one-percent change in moment Ψ_j.

Given the non-convexities and discontinuities of firms' policy functions, numerical derivatives are computed using symmetric finite differences:

$$f'(x) = \frac{f(x+\varepsilon) - f(x)}{\varepsilon},$$

for ε equal to $\pm 3\%$, $\pm 2\%$, and $\pm 1\%$ of the midpoint of the parameter space. We report the median derivative across these perturbations, following Bloom (2009). All derivatives are evaluated at the benchmark parameter vector.

Heuristically, a parameter is well identified when it exhibits sizable sensitivity to multiple moments that are themselves precisely estimated. Large values of $|\hat{\lambda}_{ij}|$ indicate that moment j plays an important role in disciplining parameter i; meaning both that the moment is well estimated and that the parameter varies when the moment does. Figures E.2, E.3, and E.4 plot the sensitivity elasticities for the key structural parameters. Each figure is indexed by a parameter, and each bar represents the elasticity of that parameter with respect to a particular moment.

The sensitivity patterns are fully consistent with the identification strategy described in subsection 4.3.1. Fixed operating costs c_f are primarily identified by moments related to average firm size, the firm-size distribution, and exit rates. An increase in c_f raises the productivity threshold for survival, leading to larger average plant size and lower exit rates, particularly among smaller firms.

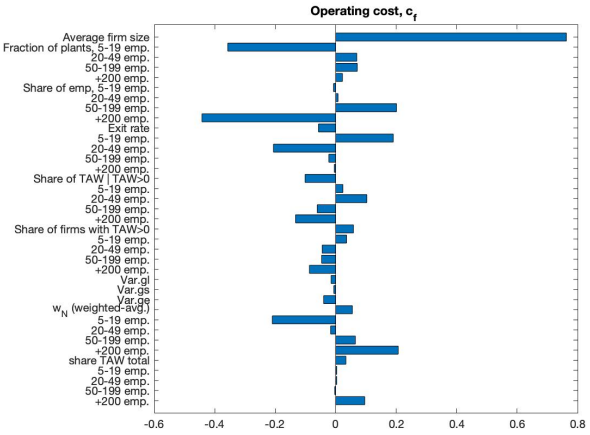
Parameters governing the use of temporary agency workers—including the proportional cost δ, the fixed administrative cost c_s, and the dispersion of nonpecuniary adoption costs σ_ν—are primarily disciplined by moments capturing agency-worker adoption along the extensive and intensive margins. In particular, the sharp increase in TAW adoption with firm size and the relatively flat conditional TAW share across size bins jointly identify these parameters. Larger values of δ or c_s reduce both the probability of using TAWs and the employment share of agency workers among adopters, with stronger effects for larger firms.

Firing costs for regular workers, τ, are identified mainly through moments capturing employment-growth volatility. Higher firing costs dampen regular-worker adjustments, reduce volatility of regular employment, and shift adjustment toward agency workers, while also increasing exit rates by limiting firms' ability to respond to negative productivity shocks.

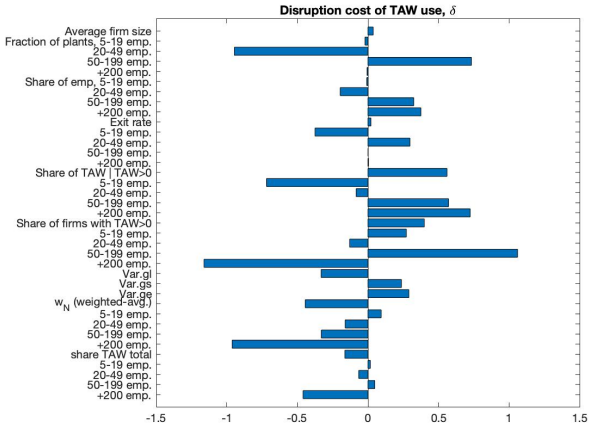
The baseline wage levels and rent-sharing elasticities, $(\overline{W}_N, \psi_N)$ and $(\overline{W}_S, \psi_S)$, are identified through wage moments by firm size and their interaction with firms' agency-worker use. Differences in wage gradients across worker types affect relative labor costs and, in equilibrium, firms' optimal reliance on agency labor.

Finally, the parameters governing the productivity process, (μ, σ), are primarily disciplined by moments related to average firm size, the firm-size distribution, and employment volatility. While both parameters affect average firm size, they do so with opposite signs, allowing the estimation procedure to disentangle mean productivity from idiosyncratic volatility.

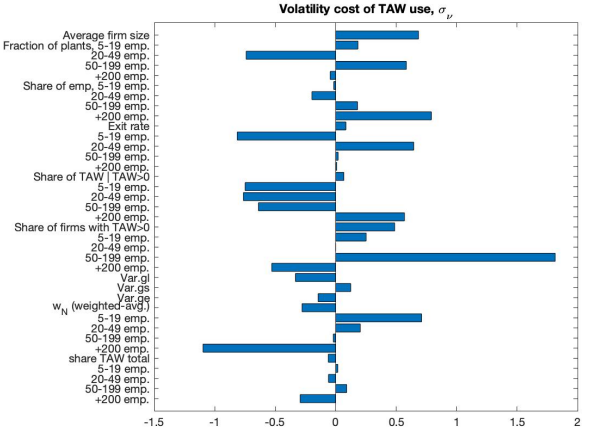
Figure E.2: Sensitivity of parameters to moments



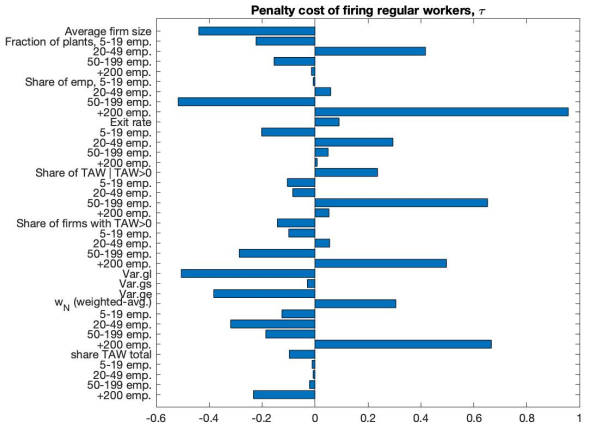
(a) Fixed operating cost, c_f



(b) TAW disruption cost, δ



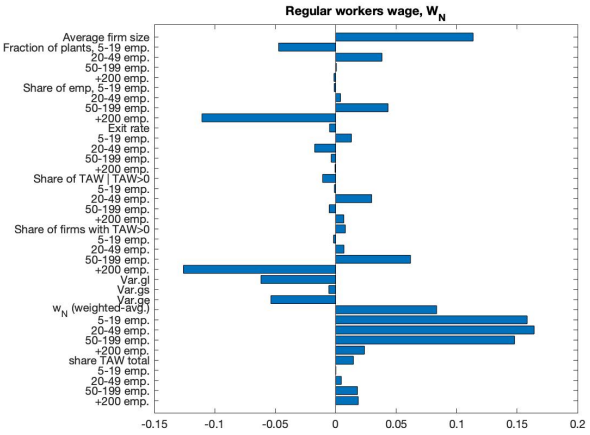
(c) Dispersion of TAW adoption taste-shock, σ_ν



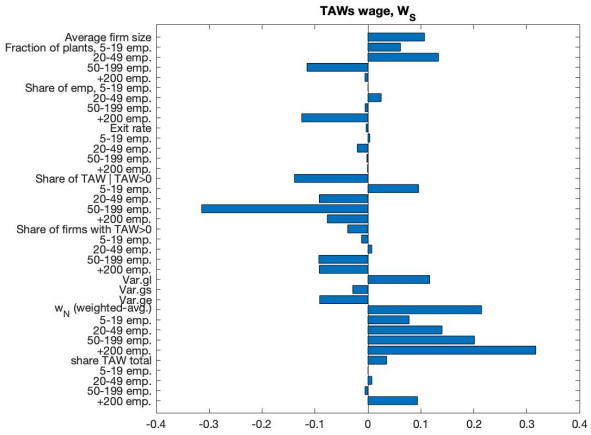
(d) Firing cost, regular workers, τ

Note: This figure reports the Andrews et al. (2017) sensitivity measure of the estimated parameters to moments, expressed in elasticity form. Specifically, $\hat{\lambda}_{ij} = \Lambda_{ij} \frac{\Psi_j^S(\hat{\theta})}{\theta_i}$, where $\Lambda = (J'WJ)^{-1}J'W$. Each bar can be interpreted as the elasticity of parameter i with respect to moment j, i.e. $\partial \log \hat{\theta}_i / \partial \log \Psi_j$.

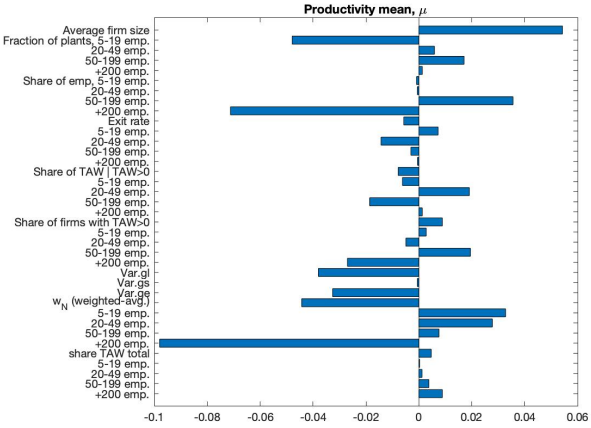
Figure E.3: Sensitivity of parameters to moments



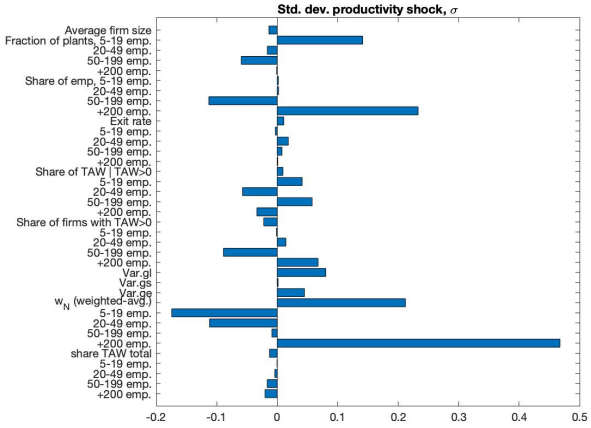
(a) Baseline wage, regular workers, $\overline{W}_N$



(b) Baseline wage, TAW, $\overline{W}_S$



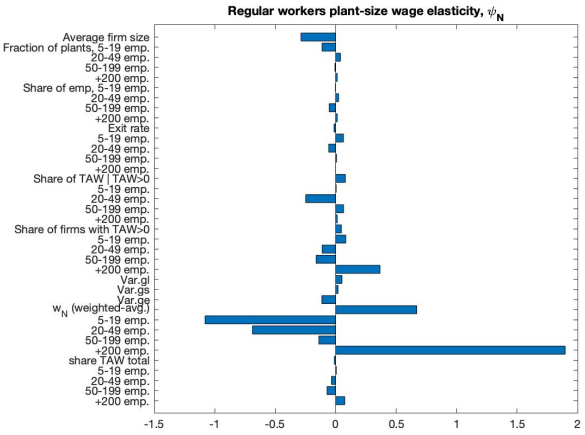
(c) Mean productivity , μ



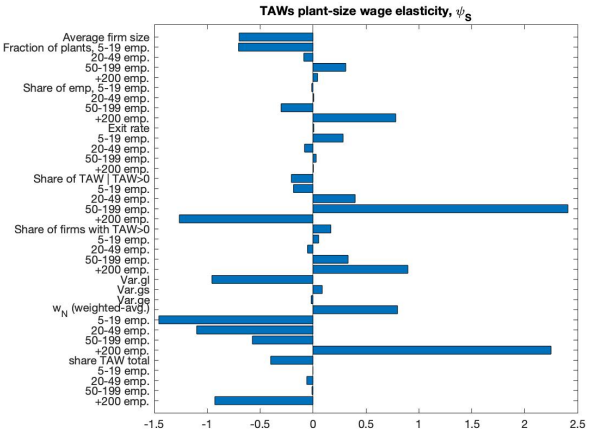
(d) Std. dev. productivity shocks, σ

Note: This figure reports the Andrews et al. (2017) sensitivity measure of the estimated parameters to moments, expressed in elasticity form. Specifically, $\hat{\lambda}_{ij} = \Lambda_{ij} \frac{\Psi_j^S(\hat{\theta})}{\theta_i}$, where $\Lambda = (J'WJ)^{-1}J'W$. Each bar can be interpreted as the elasticity of parameter i with respect to moment j, i.e. $\partial \log \hat{\theta}_i / \partial \log \Psi_j$.

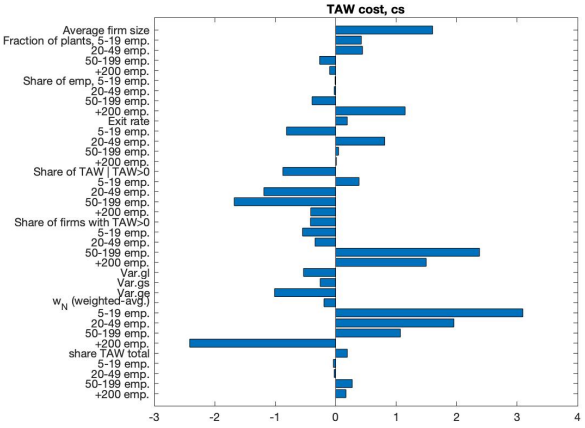
Figure E.4: Sensitivity of parameters to moments



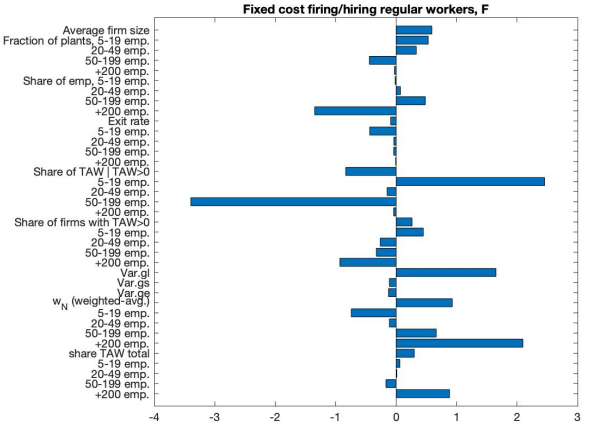
(a) Rent-sharing elasticity, regular workers, ψ_N



(b) Rent-sharing elasticity, TAW, ψ_S



(c) TAW administrative cost, c_s



(d) Fixed adjustment cost, regular workers, F

Note: This figure reports the Andrews et al. (2017) sensitivity measure of the estimated parameters to moments, expressed in elasticity form. Specifically, $\hat{\lambda}_{ij} = \Lambda_{ij} \frac{\Psi_j^S(\hat{\theta})}{\theta_i}$, where $\Lambda = (J'WJ)^{-1}J'W$. Each bar can be interpreted as the elasticity of parameter i with respect to moment j, i.e. $\partial \log \hat{\theta}_i / \partial \log \Psi_j$.

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. *The Quarterly Journal of Economics*, 108(4):983–1014.

Aguirregabiria, V. and Alonso-Borrego, C. (2014). Labor Contracts and Flexibility: Evidence from a Labor Market Reform in Spain. *Economic Inquiry*, 52(2):930–957.

Andrews, I., Gentzkow, M., and Shapiro, J. M. (2017). Measuring the Sensitivity of Parameter Estimates to Estimation Moments. *The Quarterly Journal of Economics*, 132(4):1553–1592.

Asker, J., Collard-Wexler, A., and De Loecker, J. (2014). Dynamic inputs and resource (mis)allocation. *Journal of Political Economy*, 122(5):1013–1063.

Autor, D. H. (2001). Why Do Temporary Help Firms Provide Free General Skills Training? *The Quarterly Journal of Economics*, 116(4):1409–1448.

Autor, D. H. (2003). Outsourcing at Will: The Contribution of Unjust Dismissal Doctrine to the Growth of Employment Outsourcing. *Journal of Labor Economics*, 21(1):1–42.

Barth, E., Bryson, A., Davis, J. C., and Freeman, R. (2016). It's Where You Work: Increases in the Dispersion of Earnings across Establishments and Individuals in the United States. *Journal of Labor Economics*, 34(S2):S67–S97.

Bentolila, S. and Bertola, G. (1990). Firing costs and labour demand: how bad is eurosclerosis? *The Review of Economic Studies*, 57(3):381–402.

Bentolila, S., Cahuc, P., Dolado, J. J., and Le Barbanchon, T. (2012). Two-tier labour markets in the great recession: France versus spain. *The economic journal*, 122(562):F155–F187.

Bentolila, S. and Dolado, J. J. (1994). Labour flexibility and wages: lessons from spain. *Economic policy*, 9(18):53–99.

Bentolila, S. and Saint-Paul, G. (1992). The macroeconomic impact of flexible labor contracts, with an application to spain. *European Economic Review*, 36(5):1013–1047.

Bertola, G. and Caballero, R. J. (1990). Kinked adjustment costs and aggregate dynamics. *NBER macroeconomics annual*, 5:237–288.

Besley, T. and Burgess, R. (2004). Can labor regulation hinder economic performance? evidence from india. *The Quarterly journal of economics*, 119(1):91–134.

Bilal, A. and Lhuillier, H. (2025). Outsourcing, inequality and aggregate output,. Manuscript.

Blanchard, O. and Landier, A. (2002). The perverse effects of partial labour market reform: fixed-term contracts in france. *The Economic Journal*, 112(480):F214–F244.

Bloom, N. (2009). The impact of uncertainty shocks. *Econometrica*, 77(3):623–685.

Boeri, T. and Jimeno, J. F. (2016). Learning from the great divergence in unemployment in europe during the crisis. *Labour Economics*, 41:32–46.

Bostanci, G. (2025). Productivity Gains from Labor Outsourcing: The Role of Trade Secrets. Manuscript.

Botero, J. C., Djankov, S., La Porta, R., Lopez-de Silanes, F., and Shleifer, A. (2004). The Regulation of Labor. *The Quarterly Journal of Economics*, 119(4):1339–1382.

Bratti, M., Conti, M., and Sulis, G. (2021). Employment protection and firm-provided training in dual labour markets. *Labour Economics*, 69:101972.

Bryson, A. (2013). Do temporary agency workers affect workplace performance? *Journal of Productivity Analysis*, 39:131–138.

Caballero, R. J. (1999). Aggregate investment. *Handbook of macroeconomics*, 1:813–862.

Caggese, A. and Cuñat, V. (2008). Financing constraints and fixed-term employment contracts. *The Economic Journal*, 118(533):2013–2046.

Cahuc, P., Charlot, O., and Malherbet, F. (2016). Explaining the spread of temporary jobs and its impact on labor turnover. *International Economic Review*, 57(2):533–572.

Cahuc, P. and Postel-Vinay, F. (2002). Temporary jobs, employment protection and labor market performance. *Labour economics*, 9(1):63–91.

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.

Card, D., Devicienti, F., and Maida, A. (2014). Rent-sharing, holdup, and wages: Evidence from matched panel data. *Review of Economic Studies*, 81(1):84–111.

Chan, M. (2023). How substitutable are labor and intermediates? Manuscript.

Cho, D. and Krueger, A. B. (2022). Rent sharing within firms. *Journal of Labor Economics*, 40(S1):S17–S38.

Cooper, R., Haltiwanger, J., and Willis, J. L. (2015). Dynamics of labor demand: Evidence from plant-level observations and aggregate implications. *Research in Economics*, 69(1):37–50.

Cooper, R. and Willis, J. L. (2009). The cost of labor adjustment: Inferences from the gap. *Review of Economic dynamics*, 12(4):632–647.

Cooper, R. W. and Haltiwanger, J. C. (2006). On the nature of capital adjustment costs. *The Review of Economic Studies*, 73(3):611–633.

Daruich, D., Di Addario, S., and Saggio, R. (2023). The effects of partial employment protection reforms: Evidence from italy. *Review of Economic Studies*, page rdad012.

David, J. M. and Venkateswaran, V. (2019). The sources of capital misallocation. *American Economic Review*, 109(7):2531–2567.

Davis, S. J. and Haltiwanger, J. (1992). Gross job creation, gross job destruction, and employment reallocation. *Quarterly Journal of Economics*, 107(3):819–863.

Direccion del Trabajo (2011). Encla 2011. Informe de Resultados. Séptima Encuesta Laboral.

Dolado, J. J., Lalé, E., and Siassi, N. (2021). From dual to unified employment protection: Transition and steady state. *Quantitative Economics*, 12(2):547–585.

Dolado, J. J. and Stucchi, R. (2008). Do Temporary Contracts Affect TFP? Evidence from Spanish Manufacturing Firms. *IZA Discussion Paper*, 3832.

Drenik, A., Jäger, S., Plotkin, P., and Schoefer, B. (2023). Paying outsourced labor: Direct evidence from linked temp agency-worker-client data. *Review of Economics and Statistics*, 105(1):206–216.

Dube, A. and Kaplan, E. (2010). Does Outsourcing Reduce Wages in the Low-Wage Service Occupations? Evidence from Janitors and Guards. *ILR Review*, 63(2):287–306.

Dunne, T., Foster, L., Haltiwanger, J., and Troske, K. R. (2004). Wage and productivity dispersion in United States manufacturing: The role of computer investment. *Journal of Labor Economics*, 22(2):397–429.

Eberhardt, M. and Helmers, C. (2019). Untested Assumptions and Data Slicing: A Critical Review of Firm-Level Production Function Estimators. Technical Report 513, University of Oxford, Department of Economics.

Faggio, G., Salvanes, K. G., and van Reenen, J. (2010). The evolution of inequality in productivity and wages: panel data evidence. *Industrial and Corporate Change*, 19(6):1919–1951.

Figal, L., López, P., Maffioli, A., and Ruzzier, C. (2020). Firm-level productivity in Latin America and the Caribbean. *Research in Economics*, 74(2):186–192.

Gandhi, A., Navarro, S., and Rivers, D. A. (2020). On the Identification of Gross Output Production Functions. *Journal of Political Economy*, 128(8):2973–3016.

Garcia-Louzao, J., Hospido, L., and Ruggieri, A. (2023). Dual returns to experience. *Labour Economics*, 80:102290.

Garcia-Marin, A. and Voigtländer, N. (2019). Exporting and Plant-Level Efficiency Gains: It's in the Measure. *Journal of Political Economy*, 127(4):1777–1825.

Goldschmidt, D. and Schmieder, J. F. (2017). The rise of domestic outsourcing and the evolution of the German wage structure. *The Quarterly Journal of Economics*, 132(3):1165–1217.

Gomes, P. (2018). Heterogeneity And The Public Sector Wage Policy. *International Economic Review*, 59(3):1469–1489.

Gouriéroux, C. and Monfort, A. (1996). *Simulation-Based Econometric Methods*. Oxford University Press.

Hatzichronoglou, T. (1997). Revision of the High-Technology Sector and Product Classification. Technical Report 1997/02, OECD Science, Technology and Industry.

Hopenhayn, H. and Rogerson, R. (1993). Job turnover and policy evaluation: A general

equilibrium analysis. *Journal of Political Economy*, 101(5):915–938.

Houseman, S. N. (2001). Why Employers Use Flexible Staffing Arrangements: Evidence From an Establishment Survey. *ILR Review*, 55(1):149–170.

Houseman, S. N., Kalleberg, A. L., and Erickcek, G. A. (2003). The Role of Temporary Agency Employment in Tight Labor Markets. *ILR Review*, 57(1):105–127.

Hsieh, C.-T. and Klenow, P. J. (2009). Misallocation and manufacturing tfp in china and india. *The Quarterly journal of economics*, 124(4):1403–1448.

Kanbur, R. and Ronconi, L. (2018). Enforcement matters: The effective regulation of labour. *International Labour Review*, 157(3):331–356.

Katz, L. F. and Krueger, A. B. (2019). The Rise and Nature of Alternative Work Arrangements in the United States, 1995–2015. *ILR Review*, 72(2):382–416.

Kehrig, M. and Vincent, N. (2021). The micro-level anatomy of the labor share decline. *The Quarterly Journal of Economics*, 136(2):1031–1087.

Kogan, L., Papanikolaou, D., and Stoffman, N. (2020). Left Behind: Creative Destruction, Inequality, and the Stock Market. *Journal of Political Economy*, 128(3):855–906.

Lamadon, T., Mogstad, M., and Setzler, B. (2022). Imperfect Competition, Compensating Differentials, and Rent Sharing in the US Labor Market. *American Economic Review*, 112(1):169–212.

Lazear, E. P. (1990). Job security provisions and employment. *Quarterly Journal of Economics*, 105(3):699–726.

Levinsohn, J. and Petrin, A. (2003). Estimating production functions using inputs to control for unobservables. *The Review of Economic Studies*, 70(2):317–341.

Micco, A. and Muñoz, P. (2024). The impact of extending employment protection to agency workers on firms. *American Economic Journal: Macroeconomics*, 16(1):66–101.

Ministerio de Economía, Fomento y Turismo (2015). Insolvencia y Quiebra en Chile. Principales Estadisticas desde 1982 a la Fecha.

Olley, G. S. and Pakes, A. (1996). The Dynamics of Productivity in the Telecommunications Equipment Industry. *Econometrica*, 64(6):1263–1297.

Polinsky, A. M. and Shavell, S. (2000). The economic theory of public enforcement of law. *Journal of economic literature*, 38(1):45–76.

Restuccia, D. and Rogerson, R. (2008). Policy Distortions and Aggregate Productivity with Heterogeneous Establishments. *Review of Economic Dynamics*, 11(4):707–720.

Rota, P. (2004). Estimating labor demand with fixed costs. *International Economic Review*, 45(1):25–48.

Rouwenhorst, K. G. (1995). Asset pricing implications of equilibrium business cycle models. In Cooley, T. F., editor, *Frontiers of Business Cycle Research*. Princeton University Press.

Roys, N. (2016). Persistence of shocks and the reallocation of labor. *Review of Economic Dynamics*, 22:109–130.

Syverson, C. (2011). What determines productivity? *Journal of Economic literature*, 49(2):326–365.

Weil, D. (2014). *The Fissured Workplace*. Harvard University Press.