The Heterogeneous Consequences of Reduced Labor Costs on Firm Productivity*

Francesco Del Prato† Paolo Zacchia‡

July 18, 2023

Abstract

We explore the effect of a reduction in overall labor costs, indirectly induced by an Italian reform that weakened employment protection legislation, on the productivity distribution of manufacturing firms. Due to the unique institutional features of the Italian collective bargaining system, in the manufacturing sector the reform led to a clean reduction in average worker compensation, without altering the average structure of employment relationships. This decrease in labor cost resulted in a reduction in average total factor productivity (TFP) among less productive firms, and an increase at the upper end of the distribution. We pair these findings with increased entry and exit dynamics among low-productivity firms, suggesting the presence of an adverse selection mechanism at the bottom of the TFP distribution, enhanced by the reform. We formalize this concept via a general equilibrium model that links productivity to frictions in the markets for inputs.

JEL classification: D21, D22, D24, E24, J08, O14.

Keywords: Productivity, TFP, labor flexibility, EPL, labor cost.

*We express our gratitude to Raffaele Saggio for generously sharing his data on the renewal of collective agreements. We are also thankful to Edoardo M. Acabbi, Bruno Cassiman, Luca Citino, Edoardo Di Porto, Kenan Huremović, Nikolas Mittag, Mattia Nardotto, Matteo Paradisi, Santiago Pereda Fernández, Francesco Serti, and Cristina Tealdi for their valuable comments and suggestions. We extend our thanks to the participants of the seminars at VisitINPS, KU Leuven, the SIE's 63rd RSA, the XX Brucchi Luchino workshop, the 2022 European Winter Meeting of the Econometric Society, and SOLE 2023. We would like to acknowledge the support of the "VisitINPS Scholars" program, which made this project possible. We are particularly grateful to Edoardo Di Porto and Paolo Naticchioni for their assistance with the data and helpful feedback. We also thank the entire INPS' *Direzione Studi e Ricerche* for their logistical support, as well as the INPS personnel at large, particularly Roberto Bruno. Any errors are entirely ours.

†Paris School of Economics. Contact: francesco.delprato@psemail.eu.

‡CERGE-EI (Charles University and the Czech Academy of Sciences). Contact: paolo.zacchia@cerge-ei.cz.

1 Introduction

The relationship between labor market institutions and firms' productivity is a complex one. Employment protection legislation (EPL) plays a crucial role. By allowing for easier workforce adjustments, lower EPLs can minimize distortions and enhance efficiency (Autor et al., 2007). Conversely, higher protection can suppress job creation and destruction rates, diminishing aggregate productivity (Lazear, 1990; Lagos, 2006). Institutional mechanisms such as minimum wages (Dustmann et al., 2021), unionization (Haucap and Wey, 2004), and the structure of collective bargaining (Jäger et al., 2022) can also influence productivity, both directly and by interplaying with changes in EPLs. Yet, implicitly, changes affecting EPLs or labor institutions are often thought to be associated with alterations in labor costs. Thus, to understand how shaping the labor markets influence productivity, it is crucial to comprehend the repercussions of labor cost shifts.

This paper leverages the distinctive nature of Italian collective bargaining institutions to isolate the effect of an indirect reduction in workers' salaries on the total factor productivity (TFP) in the manufacturing sector, documenting substantial variation across the productivity distribution. We demonstrate that increased flexibility—achieved through a reform that reduced firms' labor costs by decreasing workers' bargaining power—adversely affects total factor productivity (TFP) among less productive firms. This negative effect diminishes as we move up the pre-intervention productivity distribution and eventually reverses, becoming positive for the most productive firms. To interpret these results, we construct a comprehensive general equilibrium model that links the equilibrium TFP distribution across sectors with labor and capital frictions. We use this model to theoretically decompose the cost reduction effect along the productivity distribution, attributing it to a combination of a selection mechanism specific to the left tail and an incentive for productivity-enhancing investments due to downward pressure on labor costs. Highlighting these heterogeneous effects underscores the need for nuanced policy approaches, such as addressing capital misallocation based on the model's insights, to effectively address the diverse productivity levels exhibited by firms.

To empirically investigate the heterogeneous effects of flexibilization on productivity, we leverage a 2001 Italian reform designed to lower the barriers for firms to initiate new temporary contracts. We conduct a series of event studies to examine the staggered adoption of the intervention across collective bargaining agreements (Contratti Collettivi Nazionali del Lavoro, CCNL). As in Daruich et al. (2023),oour identification strategy relies on the dynamic comparison of firms that implemented the reform early—based on the prevalent contract they used prior to the reform—with those that adopted it later.

Contrary to Daruich et al. (2023), our empirical analysis primarily focuses on the firm level, examining the effect of the reform on firms' TFP. In the manufacturing sector, this intervention led to a decrease in labor costs triggered by a reduction in workers' bargaining power, indirectly caused by the reform's original mechanism—a change in the use of temporary jobs—which we document did not occur within our sample. This setting enables us to more accurately isolate a causal relationship between pure labor cost reduction and productivity.

Interestingly, the adoption of temporary contracts has remained consistently low in manufacturing over time despite the reform—a contrast to what occurs in other industries, as documented by Daruich et al. (2023). However, we still observe a significant reduction in labor costs. This paradox is explained by the loss in workers' bargaining power that followed the reform, which propagated through collective contracts across industries. We present three key facts to support this claim. First, the vast majority of Italian firms use a single collective contract for nearly their entire workforce. Second, collective agreements are employed across sectors, regardless of the specific industry they were initially designed for. Third, there exists a relationship between workforce composition and salaries across macro-sectors. In other words, changes in the share of temporary workers used in the service industries are associated with changes in salaries in manufacturing.

Our empirical analysis utilizes a comprehensive, matched employer-employee administrative dataset from the Italian Social Security Institute (Istituto Nazionale di Previdenza Sociale, INPS), further enriched by incorporating firm-level financial data from balance sheets.[1] We provide evidence of a significant heterogeneous response across the ex-ante industry-specific TFP distribution for firms within the Italian manufacturing sector. The reform reduced average TFP among already unproductive firms, and increased it for the most productive ones. We further complement these linear specifications with a quantile treatment effect approach, enabling a more nuanced examination of heterogeneity in productivity outcomes, and we provide evidence of reduced exit rates and increased entry rates for the least productive firms. Overall, our empirical results support a dual mechanism where a reduction in labor cost aids the survival and new entrance of unproductive firms, thereby impeding allocative efficiency and overall productivity, while productive firms benefit from enhanced investment incentives and potential TFP growth due to a more flexible labor market.

In more detail, we first document that the reform led to a substantial decrease in labor costs. On average, firms experienced a reduction of up to 6% in per-worker earnings within two years of adopting the new framework. Importantly, this outcome is not attributed to a shift in firms' reliance on flexible work arrangements in our sample. Indeed, the adoption of temporary contracts has remained unaltered in manufacturing following the reform. As we have argued, the combination of this evidence, the distinctive patterns in the utilization of CCNLs both across and within firms, and the lack of a response to the reform in terms of temporary work arrangements usage, allow us to interpret the reform itself as a pure labor cost shifter.

Turning to productivity, we estimate the effect of the reform on a variety of alternative measures of TFP. Plain staggered difference-in-differences around the firm-specific year of the

[1] This data allows us to construct three distinct TFP measures using various estimation techniques—namely, those proposed by Levinsohn and Petrin (2003), Ackerberg et al. (2015), and Gandhi et al. (2020)—and to incorporate worker, firm, and province-by-sector level outcomes into our analysis.

reform adoption reveal a marginally negative average impact on firm-level TFP, which remains consistent across the three measures employed. However, this overall effect conceals significant heterogeneity, as a large portion of the impact is attributable to firms in the lowest quartile of the ex-ante productivity distribution. Post-reform, the average effect within this bottom quartile is negative, resulting in a productivity loss of up to 14% three years after adoption. Simultaneously, we observe a positive and nearly symmetrical average effect within the highest quartile of the distribution, albeit estimated with slightly less precision than the impact on the lower end. Notably, these findings remain qualitatively consistent across all three productivity measures. The evidence we present stems from linear specifications that measure the average effect of the EPL reduction within given quartiles of the ex-ante TFP distribution. To refine our heterogeneity analysis, we employ a non-parametric specification for quantile treatment effects, comparing post-reform TFP distribution to a counterfactual distribution, had the reform never been implemented. Our findings reveal a monotonic influence on the TFP distribution, with negative effects on the lowest deciles and positive effects on the most productive firms.

Additionally, we demonstrate that the labor cost reduction had heterogeneous impacts on firms' turnover. While there is no effect at all on the top quartile of the distribution regarding entry and exit rates at the province-by-subsector level, there is strong evidence of a substantial reduction (approximately 7% after two years) in exit events and an increase (up to 10% after two years) in entries for less productive firms. In essence, firms seemed to exploit the reduced labor costs to bolster their chances of survival, compared to a scenario without the reform. Simultaneously, this decrease in labor costs attracted more unproductive firms to enter the market.

We propose a new interpretation of these heterogeneous results on productivity, suggesting the presence of a mixed mechanism at work, which can be understood through a straightforward intuition. A reduction in EPL leads to a change in workers' outside options that in turn decreases both the absolute and relative prices of labor faced by firms. Consequently, this change offers survival opportunities to unproductive firms that would have had a higher probability of exiting the market if the reform had not occurred. As a result, an adverse selection effect arises on the left side of the TFP distribution, where firms that should not have continued production manage to survive, entry barriers are lowered, allowing low-productivity firms to enter the market, and the reduced labor cost discourages capital deepening and investments. This combination, in turn, impairs allocative efficiency and suppresses overall productivity. In contrast, firms on the right side of the distribution, which are already productive, face increased incentives to continue investing due to the relative price change of production factors. In this case, no negative selection occurs, allowing these firms to potentially experience TFP growth as a result of efficiency gains in labor force adjustments driven by a more flexible labor market.

To lend greater rigor to the arguments supporting our findings, we develop a model that incorporates the proposed mechanisms and builds upon the foundation of general equilibrium models featuring monopolistic competition (Dixit and Stiglitz, 1977) and heterogeneous firms (Melitz, 2003). Our analysis integrates two key elements previously introduced in the literature, with novel approaches in each case: financial frictions (see, for example, Manova, 2013) and endogenous productivity (Bustos, 2011; Zhelobodko et al., 2012). We address financial frictions as an asymmetric information issue: firms require financial intermediaries (FIs) to provide credit for market entry, while FIs only have access to a noisy signal of the firms' true productivity. We model endogenous productivity in a similar manner to Bustos (2011), but treat the cost of productivity-enhancing investments (PEIs) as a continuous variable, rather than binary. Our model predicts that stronger EPLs result in reduced entry of low-productivity firms and hinder PEIs, particularly on the right tail. Our empirical findings offer robust evidence supporting the former mechanism and mixed evidence for the latter. Excluding the consideration of EPL's utility value for workers, the net welfare effect of these two mechanisms remains ambiguous, dependent on the relative impact at the tails.

Related literature Our paper intersects with several strands of literature.

First, it contributes to the literature on labor market institutions and collective bargaining. A few recent studies have reviewed the various models of industrial relations, predominantly in Germany (Jäger et al., 2021, 2022) and in Italy (Boeri et al., 2021), underscoring significant cross-country differences. We harness specific characteristics of the Italian collective bargaining institutions to isolate labor cost variations following a reform aimed at diminishing barriers to temporary contracts. Other recent research focused on characteristics of the Italian collective bargaining. Devicienti et al. (2019) focuses on the influence of collective bargaining evolution on wage inequality. Garnero (2018) analyzes compliance levels with sectoral minimum wages, revealing relatively high wage floors alongside non-negligible rates of non-compliance. Fanfani (2022) show that an increase in contractual wage levels results in the reduction of actual pay levels and has negative effects on employment. In a related study, Devicienti and Fanfani (2021) examine how firms adapt to the increases in CCNL-specific growth in minimum wages after renewals. Using the same dataset of our paper, they highlight a strong heterogeneity in the impacts on different outcomes along the productivity distribution, which aligns with our results. Yet, our paper is the first to document some unique features of the Italian collective bargaining institutions and to employ them to isolate an indirect change in labor cost.

Second, our paper advances the literature examining the micro-level determinants of the disappointing productivity of Southern European economies, and Italy in particular: a topic that has attracted both scholarly and policy interests. Different scholars have proposed different mechanisms at play, for example: relatively lower social trust levels with implications for firm internal organization (Bloom et al., 2012); financial frictions—a key feature of our model—leading to enhanced factor misallocation (Gamberoni et al., 2016); additional determinants of misallocation (Calligaris et al., 2018); the relatively lower ability of managers to reap the gains of the IT revolution (Schivardi and Schmitz, 2020). Our paper proposes a new mechanism: the interplay between labor costs and financial frictions at enhancing misallocation *in the left tail* of the productivity distribution, as detailed in our analytical framework.

Third, our paper contributes to the vast literature linking EPL and productivity. In particular, three papers lie closer to some of our setting. Autor et al. (2007) examined US plant-level data to study the effects of state courts adopting wrongful-discharge protection provisions, finding a decrease in job flows, entries, and TFP. Cappellari et al. (2012) used the same Italian reform as our study to demonstrate productivity losses resulting from the substitution of temporary employees for external staff and a decrease in capital intensity. Dolado et al. (2016) associated the cost gap between permanent and temporary jobs with firms' TFP, arguing that a larger gap lowers the temp-to-perm conversion rate, which in turn reduces worker effort and firm-level paid-for training—thus decreasing productivity.[2] While our results are consistent with many of the findings from these three works, we diverge from them in several aspects. First, we do not directly use the variation in labor force composition, rather we focus on the unique change in labor cost. Second, we emphasize the heterogeneous effects of increased labor flexibility on TFP based on ex-ante productivity. Third, we introduce a new mechanism to link the change in labor cost that followed the reform to the heterogeneous impact on TFP by combining a selection effect with altered investment incentives. Gnocato et al. (2020) examine a different channel of the heterogeneous effect of easing temporary contracts by considering the size-productivity covariance as a measure of allocative efficiency, following the approach of Hsieh and Klenow (2009). In their study, heterogeneity is driven by geographical differences in the length of labor court disputes. Our results remain consistent with their proposed mechanism of heterogeneous gains in labor productivity, which suggests that more productive firms tend to gain market shares due to longer tenures at the workplace for fixed-term workers.

Outline of the paper The paper is organized as follows. Section 2 presents the reform we study, discusses some peculiar facts about collective bargaining institutions in Italy, and illustrates the data and the sample. Section 3 outlines the empirical approach; then, it presents and discusses our findings. Section 4 illustrates the model. Section 5 concludes.

2 Institutional setting and data

This section is composed of three parts. Initially, we present and analyze the reform in question, emphasizing how our identification strategy capitalizes on its staggered implementation across different Italian national collective contracts. Subsequently, we explore the distinctive characteristics of Italian collective bargaining institutions and their implications for our empirical analyses. Lastly, we delineate the data sources utilized in our study and provide information regarding our selected sample.

[2]in comparison to Cappellari et al. (2012), we access an exceptionally rich administrative dataset, allowing us to observe the entire worker-firm match universe and the specific collective bargaining adopted by each worker, as seen in Daruich et al. (2023) and Acabbi and Alati (2021).

2.1 The Italian 368/2001 decree

Labor legislation in Italy distinguishes between regulations for permanent and temporary employment contracts. Permanent contracts, lacking a predefined termination date, necessitate substantial severance packages if an employer decides to terminate an employee. These costs depend on variables such as the size of the company and the employee's length of service. In contrast, temporary contracts, agreements with a set termination date, allow employers to dismiss employees post-contract without additional costs. Prior to the enactment of Decree 368, compliant with the EU directive 1999/70/CE on September 6, 2001, Italian companies could only employ temporary contracts under certain conditions, which needed explicit reporting to the Italian social security institute (INPS). The reform removed numerous restrictions associated with temporary contracts, leaving the permanent ones unaffected. Consequently, this led to a easier use of fixed-term employment.[3]

Though the reform was officially enacted on a specific date, it only took effect in different occupations upon the renewal of the corresponding Collective Bargaining Agreements (CCNLs). Each Italian union or union group negotiates its own CCNL, hence expiration dates differ and are known well in advance. This resulted in a staggered implementation of the new regulations on temporary contracts across CCNLs, without disrupting their usual renewal timelines.

Given this framework, we utilize the staggered renewal of 181 Italian CCNLs, exploiting the staggered adoption of more liberal temporary employment policies across these contracts. The timing of this setup allows us to leverage a potentially exogenous shift in the reform's application, yielding a quasi-experimental variation in labor market flexibility across collective agreements, as seen in Acabbi and Alati (2021) and Daruich et al. (2023). This approach enables us to identify the causal impact of the reform on temporary employment on productivity and other crucial firm-level outcomes. However, due to the unique characteristics of Italy's collective contract negotiation system, we refrain from attributing the changes prompted by the reform solely to the modified conditions for using temporary labor arrangements. The following section delves into the reasons for this stance.

2.2 The collective bargaining agreements in Italy

Institutional Overview Collective bargaining in Italy is a well-structured process, characterized by the existence of hundreds of national sector-wide collective national labor contracts. These contracts (*Contratti Collettivi Nazionali del Lavoro*, CCNLs), negotiated by trade unions and employers' associations, primarily aim to establish minimum pay levels at the national, industry-wide level within the private sector. These compensation floors, known as *contractual wages*, are set for each job title, typically encompassing between five and ten occupations. They are immune to reductions at the local level and apply to all employees

[3]It is worth noting that the reform did not amend the existing employment protection measures for ongoing and permanent contracts, increasing the difference in worker protection levels across contract types. Moreover, even post-reform, there were restrictions on the duration a firm could employ a worker under temporary contracts.

Table 1: Share of firms applying a single CCNL to 80%+ of workforce

	All panel		Within year	
	All industries	Manufacturing	All industries	Manufacturing
All	0.83	0.80	0.95	0.96
>15 empl.	0.81	0.83	0.94	0.96
>50 empl.	0.82	0.86	0.92	0.96

Note. The table presents the proportion of companies in the sample that apply a single CCNL code to at least 80% of their workforce, broken down by size. The left panel reports the share for the entire 1996-2016 panel, while the right panel displays the within-year averages. Source: Istituto Nazionale della Previdenza Sociale (INPS).

within the contract, regardless of their union membership. Contractual wages are not just seen as a wage floor, but also as a fixed component of the wage (Fanfani, 2022).

The number and nature of collective agreements within an industry are not uniform, due to both historical and organizational factors, and because the activities defined and regulated by each collective agreement do not map to a standard sector classification. As a result, as we will document in this section, it is common to observe multiple collective contracts coexisting within a single sector, with multi-sector contracts also being a frequent occurrence. The activities governed by these agreements are delineated by the bargaining parties and explicitly articulated within each contract. Employers are formally obligated to apply the contract that is most relevant to the activities performed by each employee, and this contract must bear the signatures of the most representative unions and employers' associations at the national level.

We now document three main facts regarding the use of collective bargaining in Italy. This evidence is noteworthy not only *per se*, as it contributes to a better understanding of how firms utilize CCNLs, but it will also be useful for later discussions of the mechanisms behind the effect of the reform on firm-level productivity.

Fact 1. *Most firms apply a single collective contract to the vast majority of their workforce.*

The evidence of the presence of a dominant contract within a firm is presented in Table 1. This table shows the proportion of companies in the sample that apply a single CCNL code to at least 80% of their workforce, broken down by size and sector. The vast majority of Italian companies apply a single collective contract to almost all of their workforce, regardless of their size or sector. Furthermore, such a contract remains stable over time.

Fact 2. *Collective contract types are not segregated by macro-sector (manufacturing vs. services). Instead, they are typically used across these sectors.*

Despite their names, which often refer to industry-specific occupations, the same CCNLs are widely used across sectors. Evidence of this fact is presented in Figure 1, which shows, for each collective contract, the proportion of that contract's use in manufacturing against the count ratio of observations in services vs manufacturing. The count is carried out both at the worker level, i.e., by counting the single contracts, and at the firm level, i.e., by as-

FIGURE 1: CCNLs' overlap across macro-industries

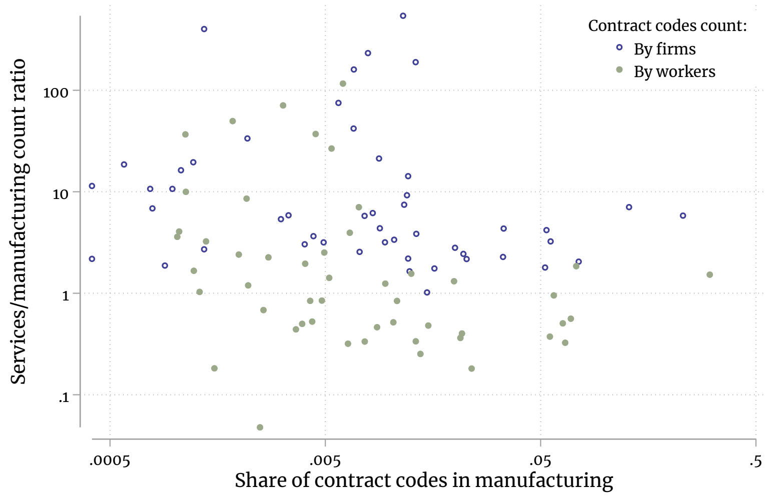

Note. This figure illustrates the overlap in the use of CCNLs across manufacturing and service industries. It plots the proportion of a collective contract's use in manufacturing against the count ratio of observations in services vs manufacturing. The count is carried out both at the worker level (blue empty circles) and at the firm level (grey filled circles), i.e., by assigning a modal collective contract to each firm. Axes are in log-scale. Source: Istituto Nazionale della Previdenza Sociale (INPS).

signing a modal collective contract to each firm. The high dispersion of the points indicates an overlap of contracts across industries: if this were not the case, we would observe points clustered rather than spread across the area. This dispersion remains consistent regardless of the count, a result that aligns with the previously established evidence of within-firm contracts' homogeneity. Fact 2 can be explained as a by-product of low enforcement (Garnero, 2018) and legal ambiguity regarding the sectoral specificity of the contracts.[4]

Fact 3. *Wages in one macro-sector display significant residual correlation with the share of temporary workers in the other macro-sector.*

The evidence supporting this last fact is presented in Table 2. It reports the estimated coefficient of the OLS regression

$$E^{I}_{ct} = \beta_1 \text{TempSh}^{M}_{ct} + \beta_2 \text{TempSh}^{S}_{ct} + \beta_3 \text{SouthSh}^{M}_{ct} + \beta_4 \text{SouthSh}^{S}_{ct} + \alpha_c + \tau_t + \gamma_c \text{Year} + \varepsilon_{ct}$$

where the dependent variable is the CCNL-year average earnings for macro-industry $I \in \{S, M\}$, regressed on the shares of workers under a temporary workers and workers in the South of Italy (as a control) for both the manufacturing and services macro-sectors. We include CCNL and year fixed effects (α_c and τ_t, respectively), and we control for CCNL-specific linear time trends (γ_c).

[4]The debate on whether a firm can choose to apply a Collective Bargaining Agreement (CCNL) that differs from its sector's has been settled by the Italian Supreme Court, favoring the view that supports a company's freedom to choose its CCNL. This principle hinges on union freedom and the contract's effectiveness between the parties involved. Despite this, choosing a different CCNL may indirectly affect aspects like contractual wage, fiscal benefits of social burdens, and other legal facilitations. Labor inspectors, though unable to compel a firm to change its CCNL, can address wage differences, recalculate minimum contributions, and reclaim economic benefits derived during the period the firm used a non-corresponding CCNL.

TABLE 2: Within-CCNL temporary-to-earnings transmission

Dep. variable	*Earnings*		*Log-earnings*	
	Manufacturing	Services	Manufacturing	Services
Temporary share				
in services	-1.37**	-6.04*	-0.040**	-0.31**
	(0.54)	(3.20)	(0.020)	(0.14)
in manufacturing	-8.13**	-0.45	-0.34**	-0.029
	(3.81)	(0.54)	(0.13)	(0.023)
South share				
in services	-0.98	-5.84***	-0.051*	-0.23**
	(0.65)	(2.06)	(0.030)	(0.11)
in manufacturing	-4.03***	0.066	-0.20***	0.014
	(1.38)	(0.48)	(0.057)	(0.022)
Observations	5,245	5,245	5,245	5,245
Adj, R-squared	0.996	0.987	0.995	0.988
CCNL + Year FE	✓	✓	✓	✓
CCNL-specific time trends	✓	✓	✓	✓

Note. The table presents the estimates of an OLS regression of earnings on the share of temporary workers and on the share of workers in the South of Italy as a control, broken down by industry. Regressions are weighted for workers numerosity in each macro-industry. Standard errors are provided in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$. Source: Istituto Nazionale della Previdenza Sociale (INPS) and CNEL.

We are mainly interested in estimating β_2 when $I = M$, to assess the correlation between the share of temporary workers in the service macro-industry and the wages in manufacturing. This coefficient is negative and statistically significant in both the levels and logs specifications. A 10 percentage points increase in the number of temporary workers in the service macro-industry is estimated to be associated with a 0.4% decrease in earnings in manufacturing, on average. Such an association remains qualitatively consistent (although statistically not significant) when considering the opposite relationship, i.e., the one between the share of temporary contracts in manufacturing and earnings in services. We interpret these results as evidence of interdependence between macro-industries within a CCNL. An increase in the use of temporary contracts can affect salaries even outside the sector where this shift occurs due to the reduction in bargaining power that follows the change in the industrial relationship. The presence and relevance of such a transmission channel are confirmed by Facts 1 and 2.

2.3 Data, sample, and summary statistics

Our paper relies on a diverse set of administrative data sources to create a comprehensive panel that connects workers and firms. This panel is enhanced by extensive details on *i)* national labor contracts (CCNLs), and their renewal dates and *ii)* data from firms' balance sheets, which we leverage to construct various total factor productivity measures. In the following section, we delve deeper into these data sources, elaborate on the TFP measures we use, and present some descriptive statistics of our sample.

2.3.1 Data sources

The empirical analysis in this paper is based on three distinct data sources:

1. The firm-level balance sheet panel data for incorporated firms in Italy from 1996 to 2016 (*Cerved* dataset).
2. The matched employer-employee panel data that covers the entirety of employment relationships in the Italian non-agricultural private sector over the same time period (*Uniemens* database).
3. Information on national collective bargaining agreements (CCNLs) detailing their renewal dates.

The first two datasets are sourced from the Italian Social Security Institute (*Istituto Nazionale di Previdenza Sociale*, INPS), and were exclusively accessible to us through the VisitINPS Scholars program.

Firms' Financial Data (Cerved) We use proprietary firm-level data on balance sheets from the Cerved database to construct three distinct TFP measures, which are discussed in detail in Section 2.3.2. The sample encompasses a period of twenty years from 1996 to 2016, incorporating standard account variables such as revenues, value-added, labor costs, tangible and intangible assets, and the cost of materials. These measures are deflated using three indices: monetary value, industry prices, and industry costs at the three-digit sector level. The deflation primarily targets the manufacturing industry, as the sample is limited to sectors for which reliable deflators are available from the Italian National Institute of Statistics (*Istituto Nazionale di Statistica*, ISTAT). The procedure for deflation and data cleaning is further detailed in Appendix B.

MEE Data (INPS' Uniemens) The National Institute for Social Security (INPS) provided us with detailed matched employer-employee records for all non-agricultural firms in the Italian private sector employing at least one worker. This unique panel comprises monthly employment histories of workers, yielding comprehensive employee-level information on demographic characteristics, labor earnings, contract type (temporary, permanent, apprenticeship), and working time arrangement (part-time or full-time). It also includes data on the collective contract applied to each worker. On the firm side, we observe company demographics—including establishment and cessation, suspension periods—the industry they operate in, and their workforce composition, size, and the total labor cost. To clean this dataset, we first select the primary employment relationship for each worker-year pair based on duration and, secondarily, earnings. We limit our sample to establishments employing at least five workers for a minimum of one year within our sample period to exclude very small firms for which a reliable TFP measure is unattainable. To match this information with firm-level data, we further restrict our *Uniemens* sample to firms present in the *Cerved* database.

Appendix B provides additional details on our data cleaning choices for this dataset.

CCNL Data Our panel is supplemented with data on the renewal dates of each CCNL, provided by the National Center for Economy and Labor (*Centro Nazionale dell'Economia e del Lavoro*, CNEL).[5] This data allows us to capitalize on the staggered implementation of the employment legislation reform, forming the cornerstone of our reduced-form analysis, as elaborated in Section 3.1.

2.3.2 TFP measures

Throughout the analysis, we focus on total factor productivity as the main measure of a firm's efficiency in production. TFP allows us to assess the effect of changes induced by the reform on a margin resulting from mechanisms that differ from (or complement) the simple adjustment of production factors that follow a change in their relative prices.

We leverage the panel structure of the data to derive a parametric estimate of the residual, ω_{it}, from a firm-specific Cobb-Douglas production function, represented as:

$$Y_{it} = K_{it}^{\beta_K} L_{it}^{\beta_L} M_{it}^{\beta_M} \exp \omega_{it}$$

where Y_{it} are deflated sales, K_{it} is capital (assets), L_{it} is the labor force, and M_{it} is the deflated cost of materials for firm i at time t.

We have employed three distinct methodologies to measure TFP, each applied as firm-specific estimations within the industry:

1. The semi-parametric estimation using the control function approach proposed by Levinsohn and Petrin (2003)—referred to as LP.

2. The output-based proxy variable approach, which originated from Olley and Pakes (1996) and further refined by Ackerberg et al. (2015)—referred to as ACF.[6]

3. The nonparametric approach introduced by Gandhi et al. (2020)—referred to as GNR.

Incorporating these three measures is designed to reduce the dependence of our results on a specific model, given the potential for TFP variation across the numerous estimation methods outlined in the literature. Moreover, our use of output-based measures permits a larger sample size in our analysis. This contrasts with value-added measures that restrict the sample to only those firm-year pairs where the value remains non-negative after logarithmic transformation.

[5] We express our gratitude to Raffaele Saggio and his co-authors for sharing this dataset, which they initially gathered and used in their study Daruich et al. (2023).

[6] Ackerberg et al. (2015) discuss that a purely Cobb-Douglas production function, similar to the one presented above, cannot be identified via a control function approach if one of its inputs is used as a proxy variable, *unless* there are frictions that impact firms' input timing decisions. EPLs, the focus of our paper, serve as a perfect example of such frictions influencing labor input. Consequently, we find the ACF method appropriate for estimating the entire production function in this setting.

FIGURE 2: Permanent vs. temporary contracts by macro-industry

(A) Share of temporary contracts (B) Temp / perm earnings

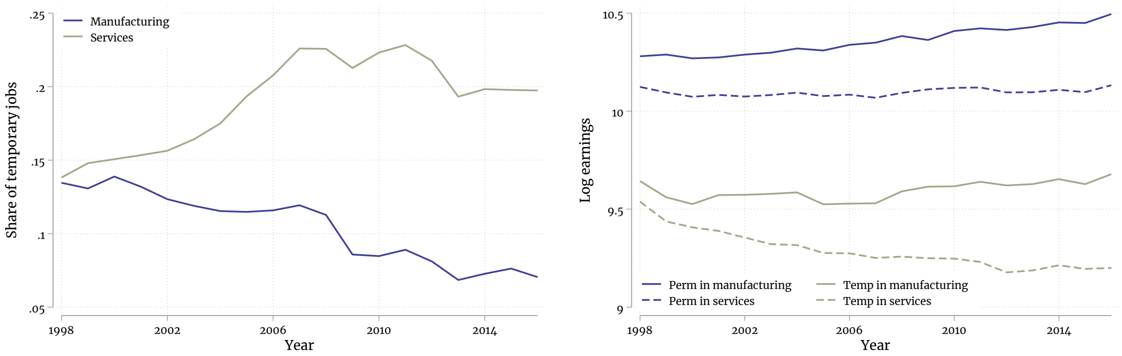

Note.
Source: Istituto Nazionale della Previdenza Sociale (INPS).

2.4 Descriptive Statistics of the Sample

Following the data cleaning process detailed in Appendix B, we are left with a selected sample comprising 50,000 to 70,000 firms per year across approximately 65 three-digit sectors in manufacturing. The variability in the sample (Figure A.1) is primarily due to changes in the coverage of the Cerved database, which included an increasing number of firms over the years. Moreover, the sample size is influenced by two main factors: firstly, we retain firms for which we can observe balance sheet data from Cerved; secondly, we limit our analysis to sectors for which we can use industry-specific prices and cost indexes as meaningful deflators.

Figure 2 presents two notable descriptive statistics. First, the proportion in the use of temporary contracts has demonstrated different patterns across the manufacturing and service sectors. In the manufacturing sector, the share of temporary contracts remained largely stable over time, even showing a slight decrease. Conversely, in the service sector, this proportion rose from about 13% in 1998 to 23% in 2007, as depicted in Panel A. Concurrently, Panel B reveals that the real earnings of workers in the manufacturing sector remained largely constant over time, regardless of contract type, with a minor increase for permanent contracts. In contrast, temporary contracts in the service sector experienced a steady decline in real earnings, negatively mirroring the increase in their usage.[7]

Figure A.2 presents the time trend of selected statistics for the sample: mean, variance, in-

[7]The slight decrease in temporary contract usage within the manufacturing industry may be attributed to two factors. First, the manufacturing macro-sector is characterized by unique technological and institutional attributes. Technologically, the less volatile nature of manufacturing production limits the necessity for rapid workforce adjustments in response to unexpected demand fluctuations. Institutionally, manufacturing firms frequently use *cassa integrazione*—an Italian wage support scheme during downturns—maintaining workforce flexibility when needed. Second, manufacturing firms might prefer to outsource labor through temporary employment agencies. This approach potentially minimizes administrative and bureaucratic burdens and mitigates litigation risk, which is highly relevant in the union-dense manufacturing sector. However, the dataset that we are currently using does not allow us to verify this hypothesis or to delve into the reasons behind the differential use of outsourced labor between the service and manufacturing sectors. As such, a comprehensive quantitative assessment of these issues is left for future research.

terquartile range (p75-p25), and inter-extreme range (p90-p10) for the three considered TFP measures. The TFP growth trend is notably negative across all three measures. The overall variance shows an increase in the early years of the sample and a subsequent reduction that continued until 2016, the last year we consider. This trend remains consistent across the TFP measures. The same observation applies to the interquartile and the inter-extreme ranges, which display a much more stable evolution over time.

Figure A.3 depicts the dynamics of firm entry and exit in our sample, taken at the province-by-3-digit-sector level each year.[8] The number of exit events, considered as potentially temporary suspensions, has increased over time. The same trend is observed when looking at the rates. This may reflect the stagnation in growth that the country has experienced over the last 30 years. Consistently, the number of entry events, defined as possible reactivations, began to decline after the 2008 financial crisis when viewed in absolute numbers, and even earlier when considering the rates. However, both panels show that new firms continue to be established and enter the market at a higher rate than the rate at which they cease their activity. In addition to these macro-trends, this paper will examine how changes in EPLs have impacted firm dynamics and through which channels.

3 Empirical strategy and results

This section outlines the empirical approach utilized in our reduced-form analysis for assessing the impact of the labor cost shifts, induced by the reform, on firm productivity and other outcomes. Moreover, here we present and discuss our findings.

We first establish that the reform led to a decrease in average labor cost without altering the use of temporary contracts in the manufacturing industry. Next, we provide robust causal evidence showing this decrease negatively impacted average firm productivity. Specifically, we demonstrate the reform's uneven distributional effects, which negatively impacted firms on the lower end of the ex-ante TFP distribution and moderately benefited those at the upper end. This finding indicates the average negative effect conceals substantial heterogeneity, mainly driven by numerous already underperforming firms. Lastly, we illustrate how the reform impacted entry and exit rates variably, enabling previously unproductive companies to endure.

3.1 Overview of the empirical strategy

Our empirical strategy, following Daruich et al. (2023) and Acabbi and Alati (2021), leverages the quasi-experimental variation introduced by the 368/2001 decree, as detailed in Section 2.1.

The reform allowed firms to apply a collective contract with eased use of temporary contracts right after its predetermined renewal date. To leverage this staggered implementation, we initially assign each firm a contract-specific renewal date by computing the

[8] For details on the construction of the two groups of entry and exit measures, please refer to Appendix B.

modal collective contract code applied by a firm in 2001—the last year before the reform. In essence, we regard a firm as treated in the year its 2001 modal collective contract was renewed. This assignment procedure is reinforced by the evidence of a single contract's high representativeness within a firm, as indicated in Table 1 and discussed in Section 2.2. As a result, we can examine within-firm changes in productivity or other pertinent firm-level outcomes before and after the reform. This is achieved by dynamically comparing firms primarily operating under the new rules with a control group of firms that, within the same year, are still adhering to pre-reform requirements. Therefore, causal identification arises from comparing as-good-as-randomly early-treated firms to later-treated ones, given the covariates.

Our empirical analysis begins by using this identification strategy to establish two facts: the reform had no uptake within our manufacturing industry sample, and it led to an average, reduction in labor cost. We then estimate the impact on firm-level TFP. Initially, we run our specification on the whole sample to gauge the average effect of the reform-induced changes on overall productivity. We then broaden our analysis to evaluate the heterogeneity of effects based on firms' positions along the sector-specific TFP distribution *before* the reform. Considering the heterogeneity in pre-reform TFP allows us to determine if labor market changes, induced by the intervention, impacted the productivity of already less productive firms differently than more productive ones. We perform this heterogeneity analysis in two ways: first, we assign a time-invariant pre-reform TFP quartile to each firm in our sample, and then run separate specifications for the top and bottom quartile of pre-reform TFP. Second, we shift from comparing within-quartile average effect estimates to the distributional one by conducting a quantile treatment effect analysis. This exercise compares the observed TFP distribution after the reform to an imputed counterfactual distribution if the reform had not occurred.

Lastly, we shift our focus to other outcomes that we deem significant for explaining the mechanisms we believe underpin the productivity results. Specifically, we transition from a firm-level specification to a cell (province-by-industry) one to evaluate the reform's effect on firms' entry and exit events., again dividing the sample by ex-ante (for exits) and ex-post (for entries) TFP quartiles.

3.2 Specifications and results

In what follows, we discuss the empirical specification we run for each different outcome of interest, and we present the associated results.

3.2.1 Baseline event study (average effects)

We start our empirical analysis with a baseline event study to assess the average effects of the reform on different firm level outcomes.

Specification We quantify the effect of the reform's adoption on the use of temporary contracts estimating the average treatment on the treated firms following Callaway and

Sant'Anna (2021).[9] More in detail, we estimate the following specification:

$$\text{ATT}(g, t; X) = \mathbb{E}\left[Y_{f,t} - Y_{f,g1} | G_f = g\right] - \mathbb{E}\left[Y_{f,t} - Y_{f,g-1} | G_f \in \mathcal{G}\right] \quad (1)$$

where $\text{ATT}(g, t; X)$ gives the average treatment effect at time t for the cohort of firms treated in time g: for example, $\text{ATT}(2003, 2005)$ measures the effect of the reform in 2005 on the group of firms that adopted the reform in 2003. For each firm f, we have that $g = c(f, 2001)$ where c is a function assigning the modal collective contract employed by firm f in 2001. Thus, g is the treatment year for all firms that in 2001 used to employ a modal CCNL that was renewed in year g. As discussed, identification comes from comparing the expected change in the outcome of interest for cohort g between periods $g-1$ and t to that of a control group of not-yet treated firms in t. This set of dynamic controls is represented by $\mathcal{G}$, i.e., $\mathcal{G} \equiv \{g' : g' > t\}$. $Y_{f,t}$ is the outcome of interest for firm f at time t—in this section, the ratio of temporary contracts use by the firm in each year, the firm's per-worker labor cost and the TFP measures. More specifically, we run specification 1 on the residuals obtained from the intermediate specification

$$\tilde{Y}_{ft} = \psi_f + \lambda_{p(f),t} + \eta_{s(f),t} + \varepsilon_{ft}$$

where ψ_f is a firm fixed effect, $\lambda_{p(f),t}$ and $\eta_{s(f),t}$ are province- and sector-by-time fixed effects, respectively, and ε_{ft} is an error term.

To visualize the estimated effect in time deviation from the reform, we are interested in an *event study* parameter that represents the weighted average of the treatment effect k periods away from the adoption across the cohorts:

$$\text{ATT}_k = \sum_g w_g \text{ATT}(g, g + k) \quad (2)$$

where weights w weight cohorts for their relative frequencies in the treated population. Specifically, we consider the integers $k \in [-3, 3]$, thereby concentrating on a symmetric three-year window around the modal CCNL's renewal. Unless otherwise stated, we cluster the standard errors at the firm level.

Results on the use of temporary contracts and labor cost On average, the reform did not result in an increased use of temporary contracts in the manufacturing industry. As shown in Figure 3, which illustrates the event study coefficients computed as in (2), the estimated effect is precisely null, with a perfectly sharp trend (Panel A). Concurrently, labor cost underwent a significant decrease. Panel B depicts the estimated coefficients on per-worker labor earnings, demonstrating a reduction of up to 7% after two years.

[9]The methodology proposed by Callaway and Sant'Anna (2021) addresses two main issues associated with standard dynamic TWFE specifications. Firstly, it offers accurate estimations even when dealing with variable treatment effects, as in our case. Specifically, this method not only avoids negative weighting but also provides control over how effects across cohorts are weighted. Secondly, this approach explicitly defines the units used as a control group to infer unobserved potential outcomes. This is in contrast to traditional TWFE models, which can result in perplexing comparisons during staggered implementations.

Figure 3: Avg. effect of the reform on use of temporary contracts and labor cost

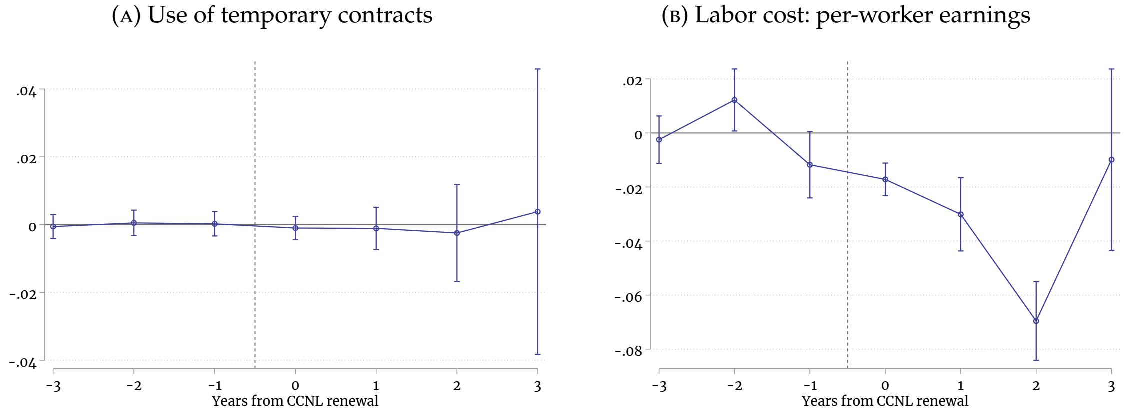

Note. This figure presents the event study coefficients calculated in accordance with equation (2) for the share of temporary contracts in the workforce (Panel A) and the labor cost as per-worker earnings (Panel B). Confidence intervals at the 95 percent level are obtained from firm-level cluster-robust standard errors. Source: Istituto Nazionale della Previdenza Sociale (INPS) and Cerved.

These results can be explained by focusing on the specificities of national collective bargaining institutions discussed in Section 2.2. Firstly, other studies that exploit the same reform across samples including non-manufacturing firms demonstrate that the intervention led to an average increase in the use of temporary contracts (Daruich et al., 2023). This suggests that the effect on labor cost we observe may be attributed to a transmission mechanism across sectors through collective bargaining. Specifically, we suggest that the increased use of temporary contracts and lower conversion rate from temporary to permanent in the service sector documented by Daruich et al. (2023) decreased the overall workers' bargaining power *within* a collective bargaining agreement. As these agreements are extensively shared *across* industries (Figure 1) and exhibit minimal within-firm heterogeneity (Table 1), the change in the workforce composition induced in the service sector by the intervention deploys effects in the manufacturing sector, even without direct uptake of temporary work arrangements. This mechanism is particularly intriguing as it enables us to interpret the effect within manufacturing as a plausibly exogenous shift in labor cost—a change not endogenous to alterations in firms' choices regarding work arrangements.

Results on the average TFP The average overall impact on TFP due to the labor cost shift induced by the reform becomes noticeable after two years, and it's negative when statistically different from zero. Figure 4 shows that between two and three years after the modal 2001 CCNL has been renewed, a firm experiences, on average, a TFP reduction between 2 (two years) and 6 (three years) percent. These results remain qualitatively consistent across the three TFP estimation methods we use. Moreover, it's worth noting that our negative average result on productivity aligns with the findings of Cappellari et al. (2012), who also leverage the same reform but use a much smaller CCNL sample and survey data on firms' sectors.

FIGURE 4: Avg. effect of the reform on TFP

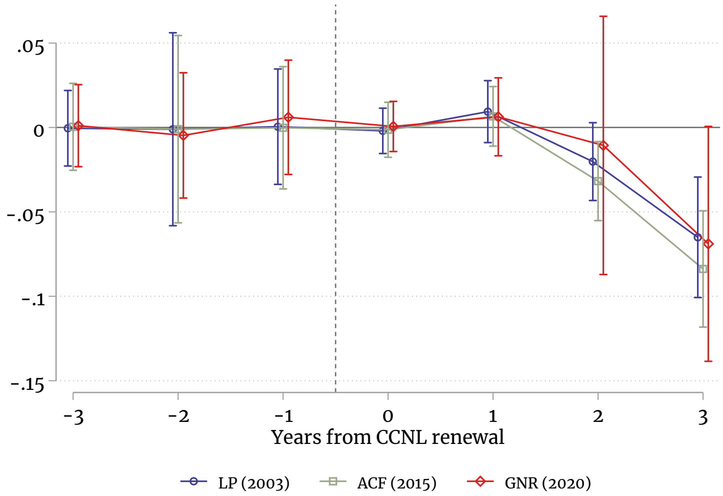

Note. This figure presents the event study coefficients calculated in accordance with equation (2) for the three different measures of TFP detailed in Section 2.3.2. Confidence intervals at the 95 percent level are obtained from three digit sector cluster-robust standard errors. Source: Istituto Nazionale della Previdenza Sociale (INPS) and Cerved.

Results on the heterogeneous ex-ante TFP quartiles A natural question arises within our conceptual framework concerning the driver of this average effect. Specifically, we aim to understand whether the reform's impact varies across the heterogeneity margin defined by ex-ante productivity. In other words, we seek to answer: did the reform differentially impact firms that were initially more or less productive? To this end, we assign each firm to a time-invariant quartile of TFP, following a process similar to the one proposed by Devicienti and Fanfani (2021). More precisely, we first calculate the firm's position in the TFP distribution within a given sector-year pair. Subsequently, we assign each firm the modal quartile in which it was classified in the five years preceding the reform's enactment.

Figure 5 showcases the outcomes of this exercise, focusing on the top and bottom quartiles of the ex-ante TFP distribution. The labor cost shift caused by the reform had symmetrical effects on these quartiles. For firms that were already significantly unproductive, the impact is detrimental, with TFP declining between 4 and 8% after two years, depending on the estimation method applied. Conversely, among the most productive firms prior to the reform, the average effect is positive, though slightly less precisely estimated, with a 10% increase in TFP two years post-reform when productivity is calculated using the ACF method (Panel B). Despite minor discrepancies in the statistical significance of the estimates, the symmetrical nature of the effect remains qualitatively consistent across all three methods. Furthermore, all three panels reveal no differences between treated and yet-to-be-treated firms in the pre-reform period, suggesting that these firms followed parallel trends prior to the reform. These findings substantiate considerable heterogeneity behind the results presented in Figure 4. Our estimates indicate that the reform negatively impacted the lower end of the distribution, with these underperforming firms primarily driving the overall negative effect.

FIGURE 5: Avg. effect of the reform on TFP, by pre-reform TFP quartiles

(A) Levinsohn and Petrin (2003)

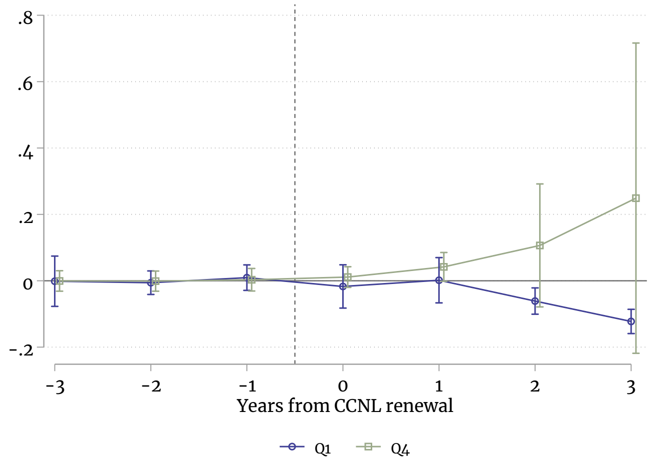

(B) Ackerberg et al. (2015)

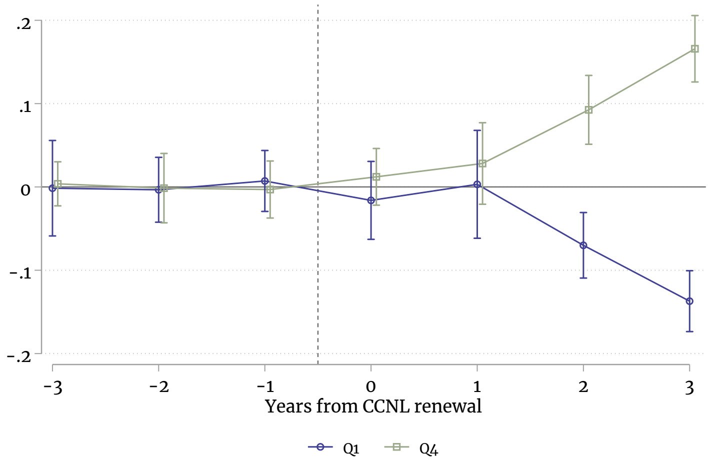

(C) Gandhi et al. (2020)

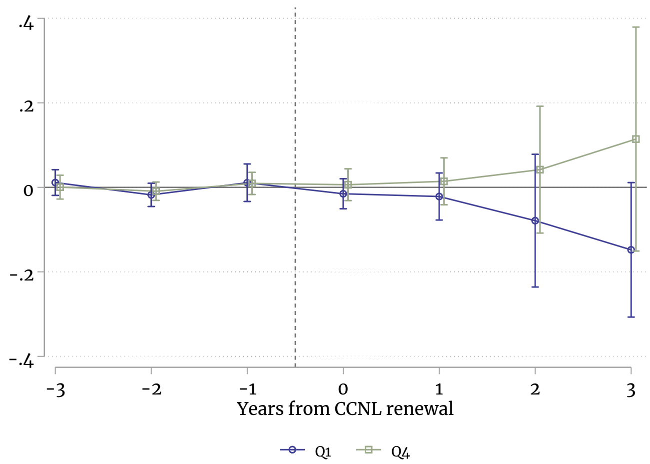

Note. This figure presents the event study coefficients calculated in accordance with equation (2) for firms within the bottom and the top time-invariant quartile of the pre-reform TFP distribution. Each panel shows the results using a different measure of TFP as the dependent variable: Levinsohn and Petrin (2003) (A); Ackerberg et al. (2015) (B); Gandhi et al. (2020) (C). Confidence intervals at 95 percent are obtained from three digit sector level cluster-robust standard errors. Source: Istituto Nazionale della Previdenza Sociale (INPS) and Cerved.

3.2.2 Quantile Treatment Effects (heterogeneous effects on the TFP distribution)

In this section, we examine how the change in labor market flexibility induced by the reform has differently impacted total factor productivity based on firms' pre-reform productivity levels. So far, we've explored the average effects of the intervention *within* the quartiles of the ex-ante TFP distribution. However, we have yet to assess the direct distributional impact of the policy. We aim to compute the Quantile Treatment Effects (QTE) of the reform on the TFP distribution for firms. Specifically, we want to evaluate how the changes in labor market conditions has differentially affected the TFP distribution itself. Essentially, our aim is to estimate the TFP distribution across firms following the reform compared to the distribution had the liberalization not occurred.

Specification To estimate the Quantile Treatment Effect on the Treated (QTT), we use the following equation:

$$QTT(\tau) = F^{-1}_{TFP_{1,t}\,|\,D=1}(\tau) - F^{-1}_{TFP_{0,t}\,|\,D=1}(\tau) \tag{3}$$

Here, τ is a quantile of the TFP distribution, and $F_{TFP_{1,t}\,|\,D=1}$ and $F_{TFP_{0,t}\,|\,D=1}$ represent, respectively, the distribution of a firm's potential productivity $TFP_{1,t}$ and $TFP_{0,t}$, conditional on running under the new rules. To accurately estimate the QTT, we need to identify the marginal distributions of potential productivity, which requires us to make two empirical assumptions.

Empirical Assumption 1 (Distributional Parallel Trends)**.** *Define* $\Delta TFP_{0,t} = TFP_{0,t} - TFP_{0,t-1}$. *Then,*

$$\Delta TFP_{0,t} \perp\!\!\!\perp D$$

In words, the distribution of the change in the untreated potential TFP must not depend on the treatment status. This is a generalization of the standard difference-in-differences parallel trends assumption applied to a non-linear context. Essentially, conditioned on covariates, the TFP trajectory observed after the reform should have been the same had the temporary contracts not been liberalized. Despite the usefulness of the Distributional Parallel Trends assumption, it isn't sufficient to fully identify the counterfactual distribution of the outcome on its own, as demonstrated by Fan and Yu (2012). To point identify the counterfactual distribution, an additional assumption, as suggested by Callaway and Li (2019), is needed.

Empirical Assumption 2 (Copula Stability)**.** *Let* $C\left(\Delta TFP_{0,t}, TFP_{0,t-1}\,|\,X, D = d\right)$ *be the copula between the change in untreated potential TFP and its starting level, conditional on covariates* X *and being treated. Then,*

$$C\left(\Delta TFP_{0,t}, TFP_{0,t-1}\,|\,X, D = 1\right) = C\left(\Delta TFP_{0,t-1}, TFP_{0,t-2}\,|\,X, D = 1\right)$$

The Copula Stability Assumption presumes that the copula, which describes the statistical dependence between the change in untreated potential TFP and its baseline level, remains

FIGURE 6: Quantile Treatment Effect for selected cohorts

(A) Reform adopted in 2002 (B) Reform adopted in 2003

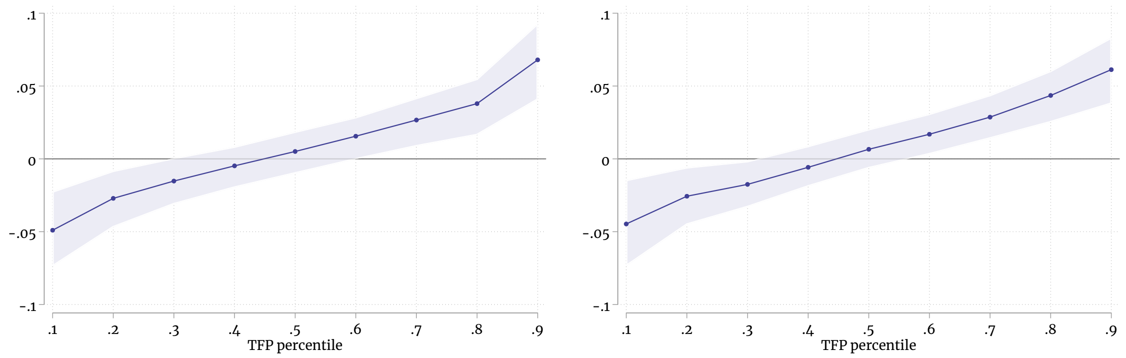

Note. This figure presents the estimate of the QTT(τ) as specified in equation (3) for $\tau = (.1, \ldots, .9)$ on the residualized TFP. The results reveal considerable heterogeneity across the TFP distribution: the effect of the reform strictly increases with the quantiles, leading to a sign reversal in the coefficients. Confidence intervals at the 95 percent level are derived through bootstrapping with 1000 iterations. Source: Istituto Nazionale della Previdenza Sociale (INPS) and Cerved.

constant over time for treated firms. [10] This is to say, if firms with higher TFP have historically exhibited greater increases in TFP, this trend will continue in the present, assuming no treatment is applied.

Given these two assumptions, we can identify the counterfactual marginal distribution in (3), thereby estimating the QTT.[11] Conceptually, the Copula Stability Assumption aids in identifying the joint distribution of $(\Delta TFP_{0,t}, TFP_{0,t-1} \mid D = 1)$, from which one can derive $F_{TFP0,t \mid D=1}$.[12]

Results Figure 6 presents the estimated distributional effects of the reform on firm-level TFP for ten deciles.[13] These results represent the short-term impact of the reform on the productivity distribution. Specifically, the specification (3) does not account for a dynamic effect k periods post-event; instead, it offers a straightforward pre-vs-post comparison. Therefore, we separately estimate the QTT for two specific cohorts, which comprise approximately 90% of the treated firms in our sample (Figure A.4): 2002 (Panel A) and 2003 (Panel B). Each panel thus illustrates the impact of the reform along the TFP distribution for the two firm cohorts, two years subsequent to the intervention.

Our findings not only corroborate the significant heterogeneity observed in the event-study

[10] It is important to note that this assumption does not necessitate a specific parametric copula or a particular form of dependence, provided that a form of dependence exists and is consistent over time.

[11] To estimate the first term of equation (3), we simply need to invert the observed empirical distribution of the TFP for firms that adopt the reform within a specific year.

[12] It is worth noting that while the marginal distributions of $\Delta TFP_{0,t}$ and $TFP_{0,t-1}$ are identified by the Distributional Parallel Trend Assumption and data respectively, this does not inherently allow for the identification of the joint distribution. As proposed by Callaway and Li (2019), we utilize the observed dependence (the past copula) to construct the necessary information to identify $F_{\Delta TFP_{0,t}, TFP_{0,t-1} \mid D=1}$, leveraging the connection between the joint distribution and the copula function as established by Sklar's Theorem (Sklar, 1959).

[13] In this case, we employ the measure based on Ackerberg et al. (2015). The results remain consistent across other TFP specifications.

estimates, but also reveal a monotonic effect across the TFP distribution for both treated cohorts. In both instances, the impact of the reform varies from a 5% decrease for the bottom decile to a positive effect of the same magnitude on the top decile, relative to a counterfactual scenario where the reform did not occur. The negative effect at the lower end of the distribution diminishes as we ascend the distribution, eventually reversing. This result imparts additional insights beyond the heterogeneity analysis conducted through linear models in preceding sections. Through the event studies, we have quantified the average marginal effect of belonging to a specific segment of the pre-reform TFP distribution on the TFP. The QTT enables us to examine the direct effect of the reform *on* the quantiles of the TFP distribution, offering evidence of the heterogeneous distribution shifts that ensued from the increase in labor flexibility.

3.2.3 Cell-level event studies (entry and exit dynamics)

This section evaluates the impact of the reform on entry and exit dynamics at the cell level. A cell is defined as a specific combination of province and industry (classified to the 3-digit level). We apply the same specification detailed in equation (1) to this revised unit of observation. The appearance of a firm within the panel dataset is considered an entry, while its disappearance indicates an exit.[14] For this analysis, the dependent variable is the ratio of the number of entries or exits to the total number of firms in the cell in 2001. This ratio is residualized for year and industry fixed effects, in addition to industry-specific linear time trends. Each cell is allocated an event year, during which it is considered treated based on the modal CCNL implemented in that particular cell in 2001. Similar to the firm-level event studies, we execute quartile-specific estimations by assigning each firm to a time-invariant TFP quartile prior to the renewal of the CCNL (for exits) or following it (for entries).[15] Again, our aim is to determine whether the reform had differential effects on firm dynamics across the top and bottom of the productivity distribution within each cell.

Results The estimated ATT computed following (2) is reported in Figure 7. On average, the reform caused an increase in entry and a derease in exit, up to 10% and 8%, respectively, within two years from the CCNL renewal. This happened only within the bottom quartile of the ex-ante TFP distribution. On the contrary, the reform led completely unaltered the entry/exit dynamics for firm the where already productive.

We interpret these results as a sharp evidence of the use of the decrease in labor cost low-productivity firms made to survive—or to enter the market. Thus, by lowering labor cost, the reform gave the opportunity to firms that would have probably left the market otherwise to continue operating, and allowed low-productivity firms that wouldn't have make it to the market, to enter it. The next Section discusses the intuition behind this mechanism and its links with results on productivity, thus introducing the model presented in Section 4.

[14]Firms persist in the MEE records as long as they maintain at least one employee in the given period.

[15]For this analysis, we utilize the TFP measure based on Ackerberg et al. (2015), but the results remain qualitatively consistent across the other measures.

FIGURE 7: Reform's effect on entry and exit, by TFP quartile

(A) Firms' exits (B) Firms' entry

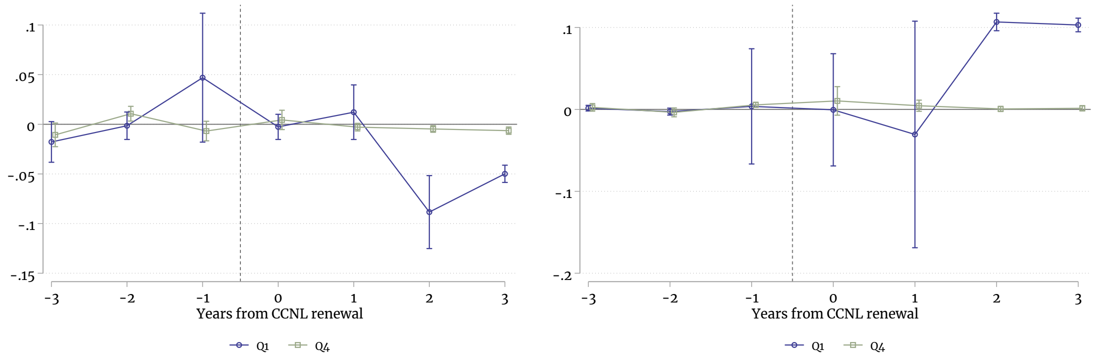

Note. This figure presents the event study coefficients calculated in accordance with equation (2) for entry (Panel A) and exit (Panel B) dynamics at the province-sector level, for the two extreme quartiles of ex-ante TFP. Confidence intervals at 95 percent are obtained from three-digit-sector-by-province level cluster-robust standard errors. Source: Istituto Nazionale della Previdenza Sociale (INPS) and Cerved.

3.3 A brief discussion on the mechanism at play

Here, we briefly discuss the intuition behind our empirical findings thus far. This discussion wraps up our empirical observations, links them to specific characteristics of the Italian labor market, and introduces some key ideas for the model presented in the next section.

First, we demonstrated that for manufacturing industries, the reform did not lead to increased use of temporary contracts. Yet, labor costs consistently decreased despite the absence of change in the composition of work arrangements. We explain the coexistence of these two phenomena with the unique characteristics of Italian labor bargaining, as discussed in Section 2.2. Briefly, the observed surge in the use of temporary jobs in the service sector, both in raw descriptive statistics (as shown in Figure 2) and following the reform (as shown in Daruich et al., 2023), weakened the bargaining power of workers *within* a collective contract. Since collective contracts are highly homogeneous within a firm and the same contract is applied across sectors, such propagation explains the observed decrease in labor costs. Thus, we consider the shift induced by the reform as a shift in firm-level labor costs, and interpret our results based on this assumption.

Next, we showed that the reform adversely affected the productivity of already unproductive firms, while improving it for already productive ones. Notably, this relationship holds not only within each quartile of ex-ante productivity, but also along the entire distribution in a monotonic way. Furthermore, the reform resulted in less market exit and more entries for unproductive firms.

We hypothesize that, thanks to the decrease in labor costs, a double negative selection mechanism was triggered at the bottom of the productivity distribution. This mechanism allowed unproductive firms to survive or enter the market. In equilibrium, the TFP composition was altered, demonstrating the productivity effects we documented. We formalize this intuition in the following section through a general equilibrium model that links labor wedges to the productivity distribution in the market.

4 Model

In this section, we present a comprehensive theoretical framework to help interpret our empirical findings. This framework connects the equilibrium productivity distributions across economic sectors with frictions in both labor and capital markets. Besides offering predictions that largely align with our primary empirical results, the model underscores the potential for labor wedges to have heterogeneous effects and an ambiguous net impact, as they may mitigate misallocation effects stemming from other types of distortions.

Our model builds on the well-established closed-economy monopolistic competition framework featuring heterogeneous firms, as developed by Melitz (2003). This approach allows us to separate various aspects of firm behavior—such as entry, exit, and investment—from other elements of the economy like labor supply. We incorporate financial frictions (FFs) due to asymmetric information and post-entry productivity-enhancing investments (PEIs) into our model. While both FFs and PEIs have been explored in prior work, our combined treatment is both innovative and tractable, as we will elaborate. Throughout our discussion, we maintain notation consistent with the original Melitz (2003) model.

We examine a standard Melitz (2003) economy where preferences for individual goods are described by a Constant Elasticity of Substitution (CES) with $\sigma > 1$. A representative consumer has a utility function $U^{\frac{\sigma-1}{\sigma}} = \int_{\omega \in \Omega} q(\omega)^{\frac{\sigma-1}{\sigma}} d\omega$, where Ω represents the set of varieties available in equilibrium, and $q(\omega)$ denotes the quantity of product $\omega \in \Omega$ consumed. Firms, which provide these varieties, are heterogeneous in productivity $\varphi(\omega) > 0$ and are characterized by a linear cost function with increasing returns due to fixed costs $f > 0$ incurred during each period of operation. The labor demand function, also linear in quantity, is expressed as $l(q) = f + q/\varphi$. In this economy, labor is inelastically supplied by a mass of workers L. Each unit of labor receives a wage w, which we normalize to unity ($w = 1$).

This economy inherits all the standard properties of the monopolistic competition model by Dixit and Stiglitz (1977) as developed later by Melitz (2003). In particular, for each firm optimal quantity is a power function of productivity with exponent σ, whereas both revenue and profit scale with exponent $\sigma - 1$. Hence, for any two firms with productivity φ_1 and φ_2, the ratio of their equilibrium revenues is $r(\varphi_1)/r(\varphi_2) = (\varphi_1/\varphi_2)^{\sigma-1}$. As in Melitz (2003), the probability distribution of productivity in this model is endogenous, represented by a density function $\mu(\varphi)$. However, the determination of this distribution in equilibrium differs in our model. Alongside a modified entry decision that firms undertake, they also need to obtain setup financing from financial intermediaries. In what follows, we refer to financial intermediaries more simply as "banks."[16]

[16]This is fitting in the Italian setting, where commercial banks dominate capital markets.

The set of varieties Ω with associated productivity $\varphi(\omega)$ is determined through the interaction between *entrepreneurs* and *banks*. We define an entrepreneur as a pair (φ, θ), where $\theta > 0$ is an individual signal about the productivity φ. These two random variables are drawn from a joint probability distribution $G(\varphi, \theta)$, which is common knowledge; however, they are initially unobserved by all agents involved. Banks are instead described as a mass B of risk-neutral workers endowed with the ability to convert any unit of labor into a unit of "capital," a unique good used solely to set up firms.[17] The process of creating firms goes as follows:

1. A given mass of entrepreneurs decides whether to attempt setting up a firm. To do so, they must incur a one-time, sunk experimentation cost f_n. This provides information about the signal θ, which both entrepreneurs and banks observe.

2. Next, firms must secure capital financing equal to f_b units of labor, which only banks can provide. The true productivity φ is revealed only after both f_n and f_b are paid. In return for paying f_b, banks demand a permanent claim over a share $b(\omega) \in (0, 1]$ of *all* future profits $\pi(\omega)$ of a firm supplying variety ω. The capital market is perfectly competitive: entrepreneurs can purchase capital from any bank without frictions.

3. Lastly, all extant firms set their prices and quantities and may even choose to exit and supply zero output if the optimal profits conditional on producing are negative (due to fixed costs). Firms then operate in the economy until an event occurring with exogenous probability δ forces them to exit.

This enhanced entry-stage model incorporates financial frictions (FFs) due to informational asymmetries. At the financing stage, banks are unable to see or verify entrepreneurs' true productivity (irrespective of whether the entrepreneurs themselves can). Existing versions of the Melitz model that incorporate FFs (see Manova, 2013 and Chaney, 2016) typically introduce liquidity constraints that firms face only when they encounter costs for entering foreign markets. In our model, FFs also impact entry into the domestic market, as our primary objective is to elucidate potential sources of misallocation while abstracting from trade considerations.[18]

Our choice to model FFs through informational frictions enhances the tractability of the model. It helps isolate one specific channel: the firm selection on the extensive margin in the left tail of the productivity distribution.[19] For the sake of clarity, we will momentarily abstract from post-entry investments (PEIs). Following our analysis of the (closed) economy with financial frictions at the entry stage, we will discuss the implications of post-entry

[17] This is a simplification and normalization, as a more elaborate production function for the capital good would not significantly alter the analysis.

[18] Our distinction between the two fixed costs f_n and f_b can be viewed as a form of liquidity constraint: of the full Melitz entry cost f_e, entrepreneurs are only able to pay $f_n < f_e$ upfront, with $f_b < f_e - f_n$ to be financed by banks.

[19] In a recent contribution, Unger (2021) introduced an augmented Melitz model where firms face financial frictions in the post-entry stage, as they need to anticipate part of both variable and fixed production costs in every period before realizing revenues. Contrary to our model, his framework (which features moral hazard) predicts that financial frictions lead to a more intense selection effect, as the least productive firms face tighter access to credit.

choices about PEIs made by firms.

4.1 Analysis

Once the set of firms that have paid both entry fixed costs f_n and f_b is determined, firm behavior proceeds as in the Melitz model. To understand how financial frictions affect firms' selection, we solve the entry stage recursively, starting from the financing stage. To clarify the trade-offs that banks face, we introduce an innocuous assumption.

Assumption 1. *Signal informativeness*: if $\theta_1 > \theta_2$ are two different realizations of the signal θ, then $G(\varphi|\theta_1) \leqslant G(\varphi|\theta_2)$ for any $\varphi > 0$.

This assumption states that signals are ordered in a way that higher values lead to conditional distributions of productivity that first-order stochastically dominate those from lower values.[20]

There are two key implications of Assumptions 1: first, lower signals imply a higher risk for banks; second, as θ is the only information that banks receive about firms, set shares $b(\theta)$, that is the fraction of total equity they demand to entrepreneurs in exchange for f_b, which is only a function of the signal. Therefore, when financing an entrepreneur with signal θ, a bank's expected profit is $\widetilde{\pi}(\theta)\, b(\theta)/\delta - f_b$, where $\widetilde{\pi}(\theta)$ is the *unconditional* per-period profit (which incorporates the probability that a firm exits after observing φ) that one can expect from setting up a firm under signal θ.

Perfect competition in capital markets leads to an equilibrium where banks make zero profits in expectation. The reason is straightforward: there cannot be an equilibrium where $\widetilde{\pi}(\theta)\, b(\theta)/\delta > f_b$ for any value of $\theta > 0$, or else any subsets of banks with mass $B' < B$ would find it profitable to set a strictly lower share $b'(\theta) < b(\theta)$ and capture all the profits from firms generated by signal θ. Banks would not make negative profits either, as they would simply deny financing to all entrepreneurs with signal values such that $\widetilde{\pi}(\theta)\, b(\theta)/\delta < f_b$. If such a strict inequality is theoretically possible in the support of θ for some fixed primitives of the model, Assumption 1 implies the existence of a threshold signal that makes banks indifferent towards financing an entrepreneur under the assumption that they would capture all the profits of the resulting firm, i.e., the smallest positive number θ^* such that

$$\frac{\widetilde{\pi}(\theta^*)}{\delta} - f_b = 0. \tag{4}$$

We guess that a suitable value of θ^* exists; we verify *ex post* whether this is true.

This analysis implies that in equilibrium, only those firms with signal $\theta \geqslant \theta^*$ receive financing, $b(\theta^*) = 1$, and for any two signals $\theta_1 \geqslant \theta^*$ and $\theta_2 \geqslant \theta^*$, banks set shares that yield zero profits in expectation with the property that $b(\theta_1)/b(\theta_2) = \widetilde{\pi}(\theta_2)/\widetilde{\pi}(\theta_1)$, and $\widetilde{\pi}(\theta)\, b(\theta) = \widetilde{\pi}(\theta^*)$ for any $\theta \geqslant \theta^*$. [21] Since (4) completely summarizes the trade-off faced

[20]This comes without any loss of generality: as signals are abstract, they can always be transformed in such a way that Assumption 1 holds by construction.

[21]This result can be formulated formally as a Bayes-Nash equilibrium.

by banks and the equilibrium in the capital markets, we call it (with some abuse of terminology) the Arbitrage Condition (AC), as it subsumes the fact that banks demand higher shares in exchange for riskier signals.

The initial entry decision by entrepreneurs is conceptually simpler. The expected value of generating a business idea is $v_n = \delta^{-1} \int_{\theta^*}^{\infty} \widetilde{\pi}(\theta) [1 - b(\theta)] \, dC(\theta)$, where $C(\theta)$ is the marginal cumulative distribution of the signal θ. Since entrepreneurs are free to attempt entering the economy and generate new signals, they would only refrain from doing so if the value of entry v_n falls shorter of the experimentation cost f_n. Thus, incorporating the equilibrium in the subsequent financing subgame and the value of the bank share $b(\theta)$ implies the following Free Entry (FE) condition in the economy:

$$\int_{\theta^*}^{\infty} \frac{\widetilde{\pi}(\theta)}{\delta} dC(\theta) - [1 - C(\theta^*)] \, f_b - f_n = 0. \tag{5}$$

Together with the Arbitrage Condition (4), this equation characterizes the economy's equilibrium. As (5) shows, entrepreneurs anticipate the probability of bearing the financing cost f_b, which they only bear if they receive a signal $\theta \geqslant \theta^*$.

To complete the analysis, it is necessary to characterize the function $\widetilde{\pi}(\theta)$. Following the analysis of the post-entry phase of the Melitz model, given a value of θ one has:

$$\widetilde{\pi}(\theta) = \mathbb{E}_{\varphi|\theta} [\pi(\varphi)| \, \theta] = f \left\{ \int_{\varphi^*}^{\infty} \left(\frac{\varphi}{\varphi^*} \right)^{\sigma - 1} g(\varphi | \, \theta) \, d\varphi - [1 - G(\varphi^* | \, \theta)] \right\}, \tag{6}$$

where $g(\varphi | \theta)$ is a conditional density function derived from $G(\varphi | \theta)$, and φ^* is the threshold value of productivity below which, in equilibrium, firms find production unprofitable and exit. Note that (6) implicitly embeds a "Zero Profit Condition" *à la* Melitz, which is specific to θ. A pair of thresholds (θ^*, φ^*), one for the signal and one for productivity, completely determines the equilibrium—if one exists.

Here, we demonstrate the existence and uniqueness of equilibrium in a specific case outlined in the following Assumption.

Assumption 2. *Log-normality*: $G(\varphi, \theta)$ is a cumulative bivariate (joint) log-normal distribution with standard log-normals as its marginals. Let $\rho \equiv \mathbb{C}\text{orr}(\log \theta, \log \varphi)$.

This assumption of normality simplifies the analysis without compromising the model's realism, given that a (truncated) normal distribution is a well-known approximation for log-productivity distributions. The assumption that the marginals are standard is a normalization that does not limit the generality of the model. It should be noted that Assumptions 1 and 2 imply that $\rho \geqslant 0$: the signal and the true productivity are non-negatively correlated.

Proposition 1. An equilibrium pair (θ^*, φ^*) always exists, is unique, and is identified by the intersection of the curve of points satisfying the AC, given by $\varphi^* = A(\theta^*)^{\rho}$ for a suitable constant $A > 0$, and a globally concave curve tracing the points that satisfy the FE condition. The intersection always occurs at the global maximum of the implicit function of φ^* for θ^*, as traced out by the FE curve.

Proof. The proof is provided in Appendix C.1. □

Figure 8 illustrates the equilibrium as the intersection between the two solid lines. The AC curve is consistently increasing because higher threshold values set by banks result in higher average productivity due to the selection of better firms, and vice-versa. The FE curve, on the other hand, is concave due to the interaction of two mechanisms. As in the Melitz model, the higher the productivity threshold, the higher the profits required to motivate entry. The same mechanism applies to the signal threshold; hence for low values of θ^*, the latter increases alongside φ^* on the FE curve. However, a higher signal threshold implies a lower probability that firms repay the financing cost f_b, thereby increasing the relative entry value. This latter effect dominates at high values of θ^* and causes the FE curve to decrease in that section. The equilibrium is located at the maximum of the FE curve due to perfect competition among banks: they lend the financing cost f_b as long as the benefits outweigh the costs.

Proposition 2. Introducing a wedge $\tau > 0$ to firms' labor costs (but not to either entry cost f_n or f_b) such that the effective wage increases from $w = 1$ to $w_{(\tau)} = 1 + \tau$, results in an equilibrium $\left(\theta^*_{(\tau)}, \varphi^*_{(\tau)}\right) \gg (\theta^*, \varphi^*)$ where both thresholds are higher with the wedge.

Proof. The proof is provided in Appendix C.2. □

The intuition behind this is straightforward: higher labor costs make it more challenging for firms to repay their fixed costs and remain in the market, leading to a more stringent selection. The Melitz model translates this effect in a downward rotation of the Zero Profit Condition curve. In our model, the wedge induces a leftward rotation of the AC curve and a rightward shift of the FE curve. As the two curves must still intersect at the maximum of the implicit function for θ^* as traced out by the FE curve, both equilibrium thresholds inevitably increase. This is graphically represented by the two dashed curves in Figure 8.

4.2 Welfare Implications

This model presents intriguing and non-trivial implications for welfare. There are two primary distinctions from the Melitz model. Firstly, there are three types of labor to be compensated: entrepreneurial (L_n), banking (L_b), and production (L_p) labor, with the total labor force being $L = L_n + L_b + L_p$. Secondly, the presence of FFs leads to the optimality result by Dhingra and Morrow (2019), which posits that the Melitz economy provides the optimal product diversity and firm size distribution, to collapse. This is summarized in the following proposition.

Proposition 3. Social welfare increases with ρ, and reaches its maximum at $\rho = 1$.

Proof. The proof is provided in Appendix C.3. □

Figure 8: Equilibrium of the model and comparative statics

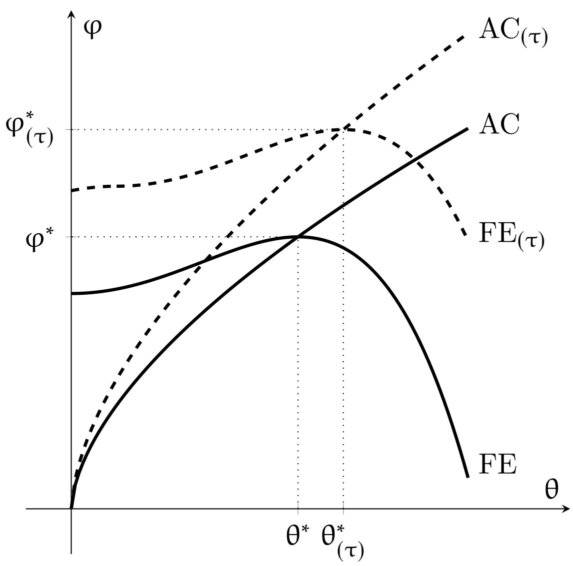

Note. This figure illustrates the equilibrium pair (θ^*, ϕ^*) as the intersection between the solid lines representing the AC condition (4) and the FE condition (5). The AC curve slopes upward because a higher signal threshold leads to higher productivity due to selection, and vice versa. For $\theta \leqslant \theta^*$, increasing productivity and signal thresholds necessitate higher profits to incentivize entry, hence they increase alongside the two. Conversely, if $\theta \geqslant \theta^*$, a higher signal threshold reduces the probability of repaying the financing cost f_b, thereby increasing the relative entry value. In this region, the FE curve slopes downward. The dashed line illustrates the shift in the FE curve and the leftward rotation of the AC curve resulting from the introduction of a wedge $\tau > 0$ in the labor cost.

The intuition behind this is straightforward: more informative signals result in a more efficient allocation of the financing cost f_b. This is best exemplified by two extreme, degenerate cases, also discussed in the Appendix, where $\rho = 0$ and $\rho = 1$ respectively, which can be viewed as two Melitz economies with different primitives. If $\rho = 0$, all entrepreneurs are financed in equilibrium, even if they exit the economy after the revelation of their true productivity. Conversely, if $\rho = 1$, no financial resources are wasted as all financed entrepreneurs remain in the economy.

A significant implication is that price distortions may lead to second-best outcomes.

Proposition 4. Introducing a wedge $\tau > 0$ to firms' labor costs (but not to either entry cost f_n or f_b), which does not factor into workers' utility, increases average productivity in equilibrium and has ambiguous effects on social welfare.

Proof. The proof is provided in Appendix C.4. □

This concept is illustrated in the Appendix, also referring to the two extreme cases mentioned above. The intuition is as follows: in the presence of FFs, higher labor costs make entry less profitable (thus reducing social welfare), but they also raise the signal threshold θ^* as per Proposition 2. This in turn leads to higher equilibrium productivity (due to a pure *selection* effect) and less "waste" on the financing cost f_b. Both of these mechanisms enhance welfare, counteracting the negative impact on welfare from reduced entry. This holds under the assumption that τ does not affect workers' compensation and social welfare *per se*. If τ is due to EPLs, it is likely that workers derive utility from it.

5 Conclusions

This paper introduces new evidence concerning the relationship between labor market institutions and productivity within the manufacturing sector. We analyze an Italian labor reform intended to ease restrictions on the use of temporary contracts to evaluate its impact on total factor productivity (TFP), following an indirect reduction in labor cost. Fundamental characteristics of Italian collective bargaining institutions are documented, enabling us to leverage the within-collective contract relationship between changes in contract composition and labor cost across macro-industries.

Indeed, we demonstrate that while the reform did not affect the use of temporary contracts within the manufacturing sector, it did reduce labor cost. We then illustrate how the reform decreased the TFP within the lowest quartile of the ex-ante productivity distribution, while inversely affecting the top of the distribution. We also note that these distributional effects grow progressively within the distribution, suggesting the more productive the firm, the greater the reform's impact. Combining this evidence with the observations of reduced exits and increased entries solely among the lowest quartile of ex-ante productivity, we propose that the reform induced a negative selection mechanism at the bottom of the TFP distribution.

We rationalize our findings through a steady-state general equilibrium model featuring monopolistic competition and heterogeneous firms. Financial frictions necessitate an up-front investment for firms to enter the market, supplied by financial intermediaries. An asymmetric information problem arises, as financial intermediaries receive only a noisy signal about firms' productivity. Our model indicates that increased labor costs contribute to selection at the lower end of the productivity distribution, leading to fewer entries. This mechanism aligns with our empirical evidence, as lower labor costs and decreased exits among unproductive firms jointly explain the observed heterogeneous effect on TFP. Our model also accounts for productivity gains on the right tail due to incentives to invest arising from labor cost savings.

Overall, this work shows that large-scale policy interventions aimed at improving labor market flexibility can have ambiguous interpretations. A variety of labor market mechanisms contribute to determining observable outcomes, with some them warranting further investigation. A full-fledged welfare analysis of our model's implications is left for future refinement.

References

Acabbi, Edoardo Maria and Andrea Alati (2021) "Defusing Leverage: Liquidity Management and Labor Contracts," *SSRN Electronic Journal*, 10.2139/ssrn.3768825. (Cited on pages 6, 7, and 14).

Ackerberg, Daniel A., Kevin Caves, and Garth Frazer (2015) "Identification Properties of Recent Production Function Estimators," *Econometrica*, 83 (6), 2411–2451, 10.3982/ECTA13408. (Cited on pages 3, 12, 19, 21, and 22).

Autor, David H., William R. Kerr, and Adriana D. Kugler (2007) "Does Employment Protection Reduce Productivity? Evidence from US States," *The Economic Journal*, 117 (521), F189–F217, 10.1111/j.1468-0297.2007.02055.x. (Cited on pages 2 and 6).

Bloom, Nicholas, Raffaella Sadun, and John Van Reenen (2012) "The Organization of Firms Across Countries*," *The Quarterly Journal of Economics*, 127 (4), 1663–1705, 10.1093/qje/qje029. (Cited on page 5).

Boeri, Tito, Andrea Ichino, Enrico Moretti, and Johanna Posch (2021) "Wage Equalization and Regional Misallocation: Evidence from Italian and German Provinces," *Journal of the European Economic Association*, 19 (6), 3249–3292, 10.1093/jeea/jvab019. (Cited on page 5).

Bustos, Paula (2011) "Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms," *The American Economic Review*, 101 (1), 304–340. (Cited on page 5).

Callaway, Brantly and Tong Li (2019) "Quantile Treatment Effects in Difference in Differences Models with Panel Data," *Quantitative Economics*, 10 (4), 1579–1618, 10.3982/QE935. (Cited on pages 20 and 21).

Callaway, Brantly and Pedro H. C. Sant'Anna (2021) "Difference-in-Differences with Multiple Time Periods," *Journal of Econometrics*, 225 (2), 200–230, 10.1016/j.jeconom.2020.12.001. (Cited on pages 15 and 16).

Calligaris, Sara, Massimo Del Gatto, Fadi Hassan, Gianmarco I P Ottaviano, and Fabiano Schivardi (2018) "The Productivity Puzzle and Misallocation: An Italian Perspective," *Economic Policy*, 33 (96), 635–684, 10.1093/epolic/eiy014. (Cited on page 5).

Cappellari, Lorenzo, Carlo Dell'Aringa, and Marco Leonardi (2012) "Temporary Employment, Job Flows and Productivity: A Tale of Two Reforms*," *The Economic Journal*, 122 (562), F188–F215, 10.1111/j.1468-0297.2012.02535.x. (Cited on pages 6 and 17).

Chaney, Thomas (2016) "Liquidity Constrained Exporters," *Journal of Economic Dynamics and Control*, 72, 141–154, 10.1016/j.jedc.2016.03.010. (Cited on page 25).

Daruich, Diego, Sabrina Di Addario, and Raffaele Saggio (2023) "The Effects of Partial Employment Protection Reforms: Evidence from Italy," *The Review of Economic Studies*, rdad012, 10.1093/restud/rdad012. (Cited on pages 2, 3, 6, 7, 12, 14, 17, and 23).

Devicienti, Francesco and Bernardo Fanfani (2021) "Firms' Margins of Adjustment to Wage Growth: The Case of Italian Collective Bargaining," *SSRN Electronic Journal*, 10.2139/ssrn.3888551. (Cited on pages 5 and 18).

Devicienti, Francesco, Bernardo Fanfani, and Agata Maida (2019) "Collective Bargaining and the Evolution of Wage Inequality in Italy," *British Journal of Industrial Relations*, 57 (2), 377–407, 10.1111/bjir.12444. (Cited on page 5).

Dhingra, Swati and John Morrow (2019) "Monopolistic Competition and Optimum Product Diversity under Firm Heterogeneity," *Journal of Political Economy*, 127 (1), 196–232, 10.1086/700732. (Cited on page 28).

Dixit, Avinash K. and Joseph E. Stiglitz (1977) "Monopolistic Competition and Optimum Product Diversity," *The American Economic Review*, 67 (3), 297–308. (Cited on pages 4 and 24).

Dolado, Juan J., Salvador Ortigueira, and Rodolfo Stucchi (2016) "Does Dual Employment Protection Affect TFP? Evidence from Spanish Manufacturing Firms," *SERIEs*, 7 (4), 421–459, 10.1007/s13209-016-0150-9. (Cited on page 6).

Dustmann, Christian, Attila Lindner, Uta Schönberg, Matthias Umkehrer, and Philipp Vom Berge (2021) "Reallocation Effects of the Minimum Wage," *The Quarterly Journal of Economics*, 137 (1), 267–328, 10.1093/qje/qjab028. (Cited on page 2).

Fan, Yanqin and Zhengfei Yu (2012) "Partial Identification of Distributional and Quantile Treatment Effects in Difference-in-Differences Models," *Economics Letters*, 115 (3), 511–515, 10.1016/j.econlet.2012.01.001. (Cited on page 20).

Fanfani, Bernardo (2022) "The Employment Effects of Collective Bargaining," *LABORatorio R. Revelli Working Paper no. 184*. (Cited on pages 5 and 8).

Gamberoni, Elisa, Claire Giordano, and Paloma Lopez-Garcia (2016) "Capital and Labour (Mis)Allocation in the Euro Area: Some Stylized Facts and Determinants," *Bank of Italy, Questioni di Economia e Finanza*, 349. (Cited on page 5).

Gandhi, Amit, Salvador Navarro, and David A. Rivers (2020) "On the Identification of Gross Output Production Functions," *Journal of Political Economy*, 128 (8), 2973–3016, 10.1086/707736. (Cited on pages 3, 12, and 19).

Garnero, Andrea (2018) "The Dog That Barks Doesn't Bite: Coverage and Compliance of Sectoral Minimum Wages in Italy," *IZA Journal of Labor Policy*, 7 (1), 3, 10.1186/s40173-018-0096-6. (Cited on pages 5 and 9).

Gnocato, Nicolò, Francesca Modena, and Chiara Tomasi (2020) "Labor Market Reforms and Allocative Efficiency in Italy," *Labour Economics*, 67, 101938, 10.1016/j.labeco.2020.101938. (Cited on page 6).

Haucap, Justus and Christian Wey (2004) "Unionisation Structures and Innovation Incentives," *The Economic Journal*, 114 (494), C149–C165. (Cited on page 2).

Hsieh, Chang-Tai and Peter J. Klenow (2009) "Misallocation and Manufacturing TFP in China and India," *The Quarterly Journal of Economics*, 124 (4), 1403–1448. (Cited on page 6).

Jäger, Simon, Shakked Noy, and Benjamin Schoefer (2022) "The German Model of Industrial Relations: Balancing Flexibility and Collective Action," *The Journal of Economic Perspectives*, 36 (4), 53–80. (Cited on pages 2 and 5).

Jäger, Simon, Benjamin Schoefer, and Jörg Heining (2021) "Labor in the Boardroom*," *The Quarterly Journal of Economics*, 136 (2), 669–725, 10.1093/qje/qjaa038. (Cited on page 5).

Lagos, Ricardo (2006) "A Model of TFP," *The Review of Economic Studies*, 73 (4), 983–1007. (Cited on page 2).

Lazear, Edward P. (1990) "Job Security Provisions and Employment," *The Quarterly Journal of Economics*, 105 (3), 699–726, 10.2307/2937895. (Cited on page 2).

Levinsohn, James and Amil Petrin (2003) "Estimating Production Functions Using Inputs to Control for Unobservables," *The Review of Economic Studies*, 70 (2), 317–341. (Cited on pages 3, 12, and 19).

Manova, Kalina (2013) "Credit Constraints, Heterogeneous Firms, and International Trade," *The Review of Economic Studies*, 80 (2), 711–744, 10.1093/restud/rds036. (Cited on pages 5 and 25).

Melitz, Marc J. (2003) "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71 (6), 1695–1725. (Cited on pages 5 and 24).

Schivardi, Fabiano and Tom Schmitz (2020) "The IT Revolution and Southern Europe's Two Lost Decades," *Journal of the European Economic Association*, 18 (5), 2441–2486, 10.1093/jeea/jvz048. (Cited on page 5).

Sklar, M. (1959) *Fonctions de Répartition À N Dimensions Et Leurs Marges*: Université Paris 8. (Cited on page 21).

Unger, Florian (2021) "Credit Frictions, Selection into External Finance and Gains from Trade," *Canadian Journal of Economics/Revue canadienne d'économique*, 54 (3), 1206–1251, 10.1111/caje.12529. (Cited on page 25).

Zhelobodko, Evgeny, Sergey Kokovin, Mathieu Parenti, and Jacques-François Thisse (2012) "Monopolistic Competition: Beyond the Constant Elasticity of Substitution," *Econometrica*, 80 (6), 2765–2784, 10.3982/ECTA9986. (Cited on page 5).

Appendix A Additional figures and tables

Figure A.1: Sample size evolution

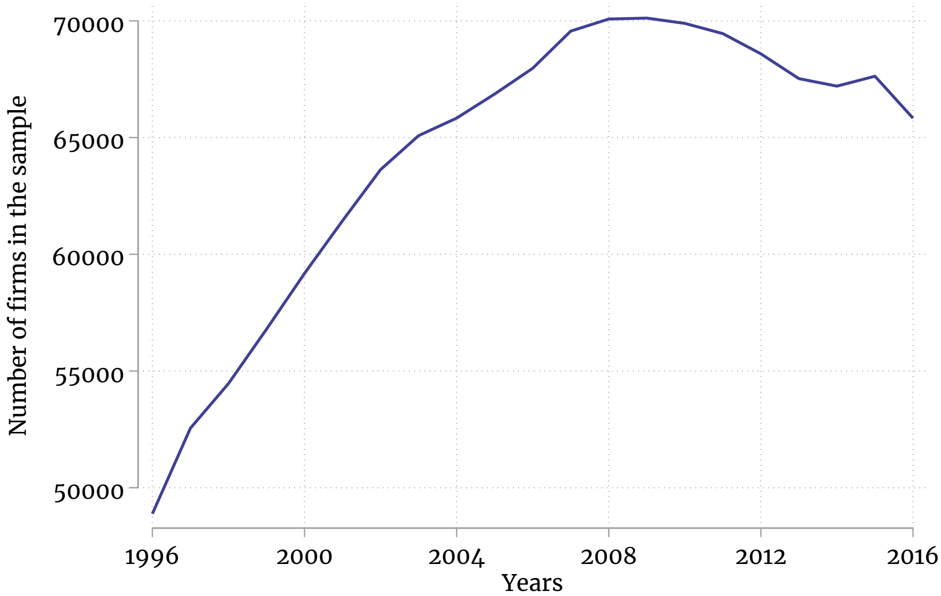

Note. This figure illustrates the evolution of the sample size over the sample period from 2006 to 2016. The observed change is primarily due to the expansion of firms' balance sheets recorded in the Cerved database. Source: Cerved.

Figure A.2: Descriptive statistics of the sample

(A) Mean

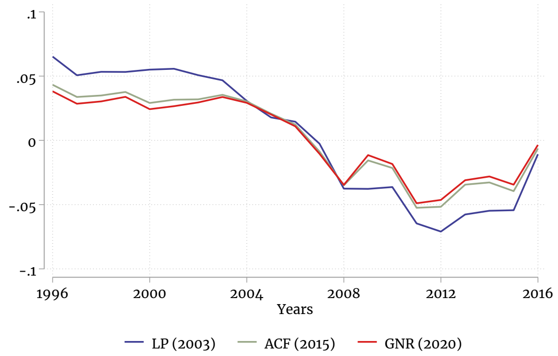

(B) Variance

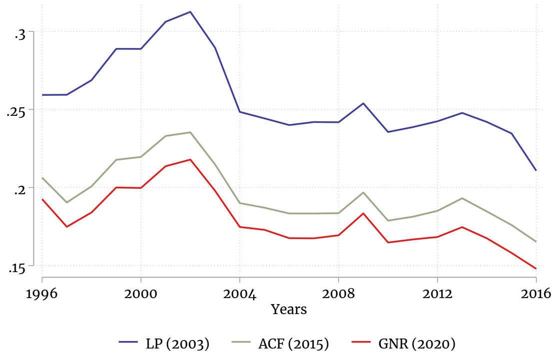

(C) Inter-quartile range (p75-p25)

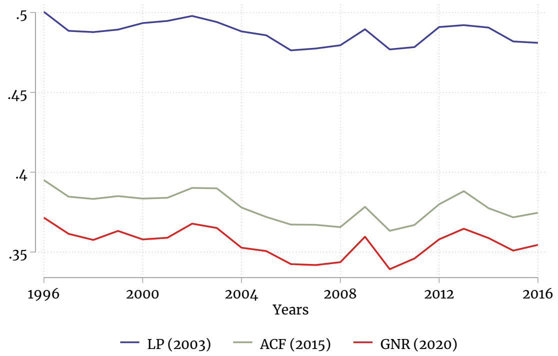

(D) Inter-extreme range (p90-p10)

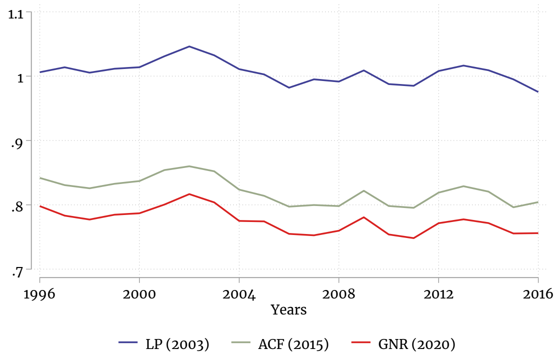

Note. This figure reports time series of different descriptive statistics of the sample—mean (A), variance (B), inter-quartile range (C), inter-extreme range (D)—for the three TFP measures employed in the paper. Source: Istituto Nazionale della Previdenza Sociale (INPS) and Cerved.

Figure A.3: Firm dynamics in the sample

(A) Exit / entry events

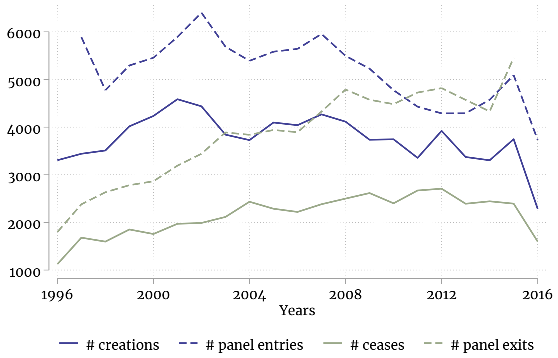

(B) Exit / entry rates

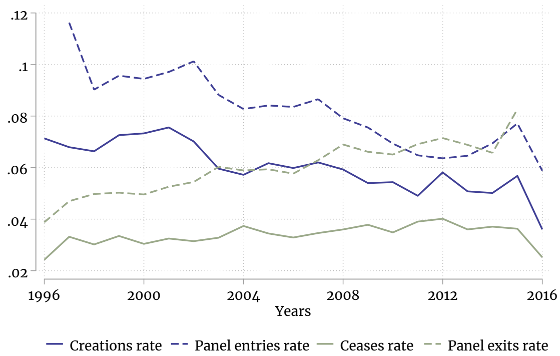

Note. This figure reports the entry and exit events (Panel A) and the entry and exit rates (defined as the ratio between the events out of the firms population in each year; Panel B) for the sample. Source: Istituto Nazionale della Previdenza Sociale (INPS).

Figure A.4: Event years distribution

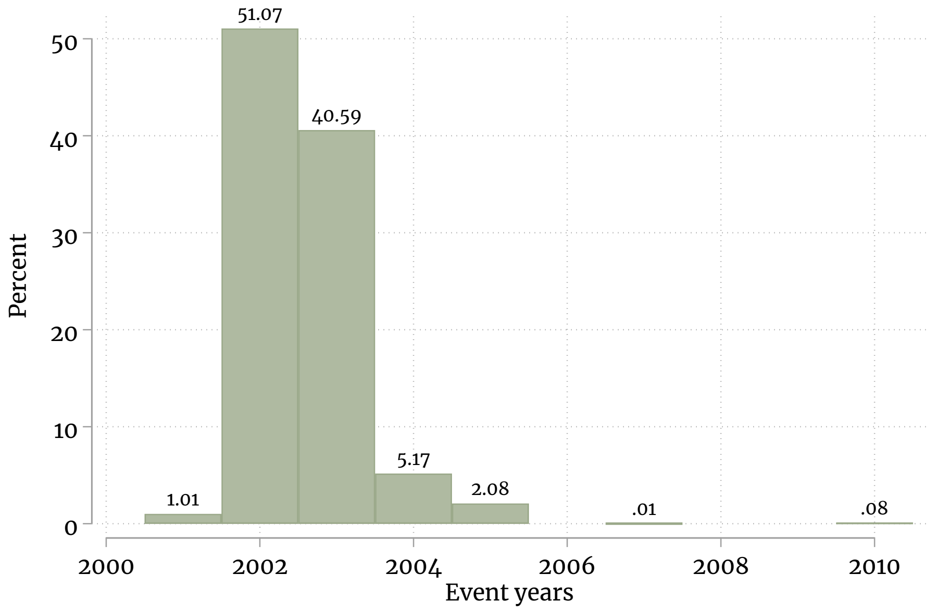

Note. This figure presents the relative percentages of event years at the firm level. An event year is defined as the year of renewal for the most commonly used National Collective Labor Contract (CCNL) among a firm's workforce in 2001. Source: Istituto Nazionale della Previdenza Sociale (INPS) and Consiglio Nazionale dell'Economia del Lavoro (CNEL).

Appendix B Additional Information on Data Cleaning and Deflation

B.1 Deflation of Cerved Measures

We adjust the firm-level balance sheet measures for inflation using three different indexes from ISTAT, the Italian National Institute of Statistics. Firstly, we adjust capital measures—fixed assets and liquidity—using a purchasing power index. Secondly, we adjust revenues using an industry-specific price index, matched at the finest available digit each year. Lastly, we use an imputed cost index at the industry level to adjust production inputs—net purchases and labor costs, matching these indexes at the best possible digit.

In more detail, we construct this latter measure through the following steps.

1. we normalize the input-output table so that each matrix element represents the relative weight an input has in the output costs in a given year;
2. for each output sector-year pair, we construct a cost index through a weighted sum of the cost indexes of the input sectors;
3. each cost index is assigned to the best available industry-specific price index in ISTAT.

We set the base year of all three indexes in 2015. When the industry price indexes are unavailable from 1996 to 1999, we extrapolate the available points of the series to predict these observations. Specifically, we use an ARIMA$(0, 1, 0)$ with a subset of external predictors, primarily the series of the lagged salary index.

B.2 Details on Data Cleaning

Here, we detail all the data-cleaning operations performed on the different datasets.

We construct a firm-year panel dataset starting from the Uniemens database. In this dataset, we assign a unique province and industry for each observation—since a single firm might operate in more than one sector or geographical region with some branch—keeping the observation with the highest number of employees.

For the matched employer-employee dataset, we first discard the contracts that lasted less than nine weeks in a year. Then, we assign to each worker-year firm only one establishment. Firstly, we resolve multiple spells *within* the same employer in a year by keeping the one that pays more. Then, we resolve multiple spells in different employers within the same year by adopting a double criterion: we keep the one that pays more and, subordinately, the one that involves more worked months. Finally, we discard contracts with no wage, and we winsorize the wage outliers on the right at the 99.7 percentile.

We clean the Cerved firm-level panel dataset by winsorizing all the relevant balance sheet variables at the 0.1 and 99.9 percentiles to remove outliers and by replacing negative values among variables that should not contain one (costs, revenues, purchases, and assets) as missing values: this way, we can still use the observation in its valid information. Besides the labor cost measure that this dataset contains, we also add the one obtained by the matched employer-employees, where we can collapse individual wages at the firm-year level. We then remove industries with less than forty firms in the entire period and industry-province-year cells that do not contain at least three establishments. Finally, after each TFP estimation, we discard the estimates that report at least one negative coefficient, and we further restrict the sample to industries for which we have non-missing estimates of all three productivity measures.

B.3 Details on Exit and Entry Events

We use the record of a registered cessation of a firm in a specific year from the INPS dataset to identify permanent exit events, i.e., those cases in which a firm permanently exits the markets, communicating this fact to the Social Security Institute. Similarly, we use the registered firms' creations to identify entries of newborn employers.

Moreover, since firms appear in the INPS panel as long as they employ at least one worker (either full- or part-time; either with a temporary or permanent contract), we interpret disappearances from the dataset in specific years as a signal of firms' inactivity in those periods. Specifically, we flag as a *disappearance* event a period of at least two consecutive years in which the firm does not appear in the panel. Conversely, we consider a reappearance in the panel as an indication of a firm's reactivation—with the same reasoning applied for the disappearances from the panel. It is important to note that in both cases, the second definition of an exit or entry event encompasses the first one.

Appendix C Additional analysis of the model

C.1 Analysis of Proposition 1

It is useful to establish some auxiliary notation first. Let:

$$\begin{aligned} t &= \log\theta \\ p &= \log\varphi \\ u &= -\log\theta \\ u' &= -\log\theta + \rho(\sigma-1) \\ z &= \frac{\log\varphi - \rho\log\theta}{\sqrt{1-\rho^2}} \end{aligned}$$

and use asterisks to denote the values of these transformations evaluated at the corresponding threshold value of their argument(s): thus, $t^* = \log\theta^*$, $p = \log\varphi^*$, *et cetera* (but $z^* = (p^* - \rho t)/\sqrt{1-\rho^2}$ is a function of $t = \log\theta$). In addition, let $\phi(x)$ be the probability density function of the standard normal distribution and $\Phi(x)$ the corresponding cumulative distribution – both evaluated at a given point x—and $\Phi_\rho(x, y)$ be the cumulative bivariate normal distribution with standard normal marginals and correlation parameter ρ—evaluated at point (x, y).

We start by elaborating expression (6) under the model's assumptions:

$$
\begin{aligned}
\frac{\widetilde{\pi}\left(e^{t}\right)}{f} &= \int_{p^*}^{\infty} \frac{e^{(\sigma-1)(p-p^*)}}{\sqrt{1-\rho^2}} \phi\left(\frac{p-\rho t}{\sqrt{1-\rho^2}}\right) dp - \left[1-\Phi\left(\frac{p^*-\rho t}{\sqrt{1-\rho^2}}\right)\right] \\
&= \int_{z^*}^{\infty} e^{(\sigma-1)\left(\sqrt{1-\rho^2}z+\rho t-p^*\right)} \phi(z)\, dz - \left[1-\Phi(z^*)\right] \\
&= e^{(\sigma-1)(\rho t-p^*)+\frac{1}{2}(\sigma-1)^2\left(1-\rho^2\right)} \int_{z^*}^{\infty} \phi\left(z-(\sigma-1)\sqrt{1-\rho^2}\right) dz - \Phi(-z^*) \\
&= e^{(\sigma-1)(\rho t-p^*)+\frac{1}{2}(\sigma-1)^2\left(1-\rho^2\right)} \Phi\left((\sigma-1)\sqrt{1-\rho^2}-z^*\right) - \Phi(-z^*) \\
&= e^{(\sigma-1)(\rho t-p^*)+\frac{1}{2}(\sigma-1)^2\left(1-\rho^2\right)} \Phi\left(\frac{\rho t-p^*+(\sigma-1)\left(1-\rho^2\right)}{\sqrt{1-\rho^2}}\right) - \Phi\left(\frac{\rho t-p^*}{\sqrt{1-\rho^2}}\right).
\end{aligned}
$$

Therefore, the Arbitrage Condition (4) reads:

$$
e^{(\sigma-1)(\rho t^*-p^*)+\frac{1}{2}(\sigma-1)^2\left(1-\rho^2\right)} \Phi\left(\frac{\rho t^*-p^*+(\sigma-1)\left(1-\rho^2\right)}{\sqrt{1-\rho^2}}\right) - \Phi\left(\frac{\rho t^*-p^*}{\sqrt{1-\rho^2}}\right) - \frac{\delta f_b}{f} = 0,
$$

with an associated implicit function $p^* = \rho t^* + a$ where $a = \log A$—as one can verify by setting the total differential at zero. It is also possible to verify that plugging this implicit function back into the right-hand side of the above AC delivers a decreasing monotone function of a which cuts the x-axis so long as $\delta f_b / f > 0$. Therefore, a (and hence A) is unique, and it is both decreasing in f_b and increasing in f.

To analyze the Free Entry Condition, it is helpful to define $\widetilde{v} \equiv \int_{\theta^*}^{\infty} \widetilde{\pi}(\theta)\, dC(\theta)$ as the expected joint profits that accrue to both the entrepreneur and the bank following the experimentation stage. This quantity can be expressed as the following function of the two threshold values (t^*, p^*):

$$
\begin{aligned}
\widetilde{v}(t^*, p^*) &= f \int_{t^*}^{\infty} \pi\left(e^{t}\right) \phi(t)\, dt \\
&= f \int_{t^*}^{\infty} e^{(\sigma-1)(\rho t-p^*)+\frac{1}{2}(\sigma-1)^2\left(1-\rho^2\right)} \Phi\left(\frac{\rho t-p^*+(\sigma-1)\left(1-\rho^2\right)}{\sqrt{1-\rho^2}}\right) \phi(t)\, dt \\
&\quad - f \int_{t^*}^{\infty} \Phi\left(\frac{\rho t-p^*}{\sqrt{1-\rho^2}}\right) \phi(t)\, dt \\
&= f e^{\frac{1}{2}(\sigma-1)^2-(\sigma-1)p^*} \int_{-\infty}^{-t^*+\rho(\sigma-1)} \Phi\left(\frac{-\rho u'-p^*+(\sigma-1)}{\sqrt{1-\rho^2}}\right) \phi(u')\, du' \\
&\quad - f \int_{-\infty}^{-t^*} \Phi\left(\frac{-\rho u-p^*}{\sqrt{1-\rho^2}}\right) \phi(u)\, du \\
&= f\left[e^{\frac{1}{2}(\sigma-1)^2-(\sigma-1)p^*} \Phi_\rho\left(-p^*+\sigma-1, -t^*+\rho(\sigma-1)\right) - \Phi_\rho\left(-p^*, -t^*\right)\right],
\end{aligned}
$$

where the last line follows from the analysis of the moments of the standard normal cumulative distribution as in **?**. Write the Free Entry condition as follows:

$$\mathcal{H}\left(p^{*},t^{*}\right)=e^{\frac{1}{2}(\sigma-1)^{2}-(\sigma-1)p^{*}}\Phi_{\rho}\left(-p^{*}+\sigma-1,-t^{*}+\rho\left(\sigma-1\right)\right)-\\ -\Phi_{\rho}\left(-p^{*},-t^{*}\right)-\frac{\delta f_{b}}{f}\Phi\left(-t^{*}\right)-\frac{\delta f_{n}}{f}=0.$$

The derivative of the above with the respect to the log-productivity threshold p^{*} is, following some manipulation, shown to be always negative:

$$\frac{\partial\mathcal{H}\left(p^{*},t^{*}\right)}{\partial p^{*}}=-\left(\sigma-1\right)e^{\frac{1}{2}(\sigma-1)^{2}-(\sigma-1)p^{*}}\Phi_{\rho}\left(-p^{*}+\sigma-1,-t^{*}+\rho\left(\sigma-1\right)\right)<0.$$

Instead, the derivative with respect to the log-signal threshold t^{*} is shown to be:

$$\frac{\partial\mathcal{H}\left(p^{*},t^{*}\right)}{\partial t^{*}}=-\left[e^{(\sigma-1)(\rho t^{*}-p^{*})+\frac{1}{2}(\sigma-1)^{2}\left(1-\rho^{2}\right)}\Phi\left(\frac{\rho t^{*}-p^{*}+\left(\sigma-1\right)\left(1-\rho^{2}\right)}{\sqrt{1-\rho^{2}}}\right)-\\ -\Phi\left(\frac{\rho t^{*}-p^{*}}{\sqrt{1-\rho^{2}}}\right)-\frac{\delta f_{b}}{f}\right]\phi\left(t^{*}\right)$$

which is not a monotone function of t^{*}. However, an analysis of this derivative shows that, for a fixed p^{*}, it is $\lim_{t^{*}\to-\infty}\partial\mathcal{H}\left(p^{*},t^{*}\right)/\partial t^{*}=\lim_{t^{*}\to\infty}\partial\mathcal{H}\left(p^{*},t^{*}\right)/\partial t^{*}=0$; that the derivative equals exactly 0 whenever $t^{*}=\left(p^{*}-a\right)/\rho$ (observe that the expression in brackets matches the Arbitrage Condition); and that to the left of this value, the derivative is positive, while on the right, it is negative. These results give rise to the pattern depicted in Figure 8, with the interpretation given in the text. Also, observe that the line $p^{*}=\rho t^{*}+a$ can only intersect the implicit function of p^{*} with respect to t^{*} based on the Free Entry condition at a stationary point of the implicit function because a is unique. Since there is only one such stationary point, there is only one intersection point and, therefore, only one equilibrium of the model.

C.2 Analysis of Proposition 2

This is straightforward: as already mentioned a (and thus A) is increasing in f, while $\partial\mathcal{H}\left(p^{*},t^{*};f\right)/\partial f=\delta\left[f_{b}\Phi\left(-t^{*}\right)+f_{n}\right]f^{-2}>0$. Hence, the AC and FE curves shift, following an increase of the fixed cost of production from f to $f\left(1+\tau\right)$—with f_{n} and f_{b} staying unchanged—according to the pattern depicted in Figure 8. Since the two curves must always meet at the maximum of the implicit function of p^{*} over t^{*}, both threshold values are higher in the new equilibrium.

C.3 Analysis of Proposition 3

This is a particular instance where informational frictions lead to a deadweight welfare loss, which is larger the more marked frictions are. As in the original Melitz model, we analyze the welfare implications of the model's steady state. First, define:

$$\mathcal{P}_\theta^* \equiv \mathbb{P}\mathrm{r}(\theta \geqslant \theta^*)$$
$$\mathcal{P}_\varphi^* \equiv \mathbb{P}\mathrm{r}(\varphi \geqslant \varphi^*)$$

as the two unconditional probabilities that in equilibrium, before the draw of (θ, φ) pair, a firm-entrepreneur passes either threshold. Also define:

$$\widetilde{\pi} \equiv \mathbb{E}[\pi | \theta \geqslant \theta^*]$$

that is the expected market profits (including the share to be paid out to banks) that firms expect in equilibrium conditional upon passing the signal threshold. Thus, in steady state the mass of entering firms M_e and that of active firms M must comply to $\delta M = \mathcal{P}_\varphi^* M_e$; the total remuneration of entrepreneurial labor is $L_n = M_e f_n$; and bank labor amounts in equilibrium to $L_b = M_e \mathcal{P}_\theta^* f_b$. Moreover, free entry implies the following relationship in steady state:

$$\mathcal{P}_\theta^* (\widetilde{\pi} - \delta f_b) - \delta f_n = 0.$$

Lastly, recall that $\bar{r}$ and $\bar{\pi}$, in the original Melitz model, indicate average equilibrium revenues and profits conditional on successful entry, respectively.

Combining everything, it is:

$$\begin{aligned} L = L_p + L_b + L_n &= M(\bar{r} - \bar{\pi}) + M_e \left(\mathcal{P}_\theta^* f_b + f_n\right) \\ &= M\left[\bar{r} - \bar{\pi} + \frac{\delta}{\mathcal{P}_\varphi^*}\left(\mathcal{P}_\theta^* f_b + f_n\right)\right] \\ &= M\left(\bar{r} - \bar{\pi} + \widetilde{\pi}\frac{\mathcal{P}_\theta^*}{\mathcal{P}_\varphi^*}\right) \end{aligned}$$

where the first line exploits $L_p = M(\bar{r} - \bar{\pi})$ as in Melitz; the second line leverages stationarity, and the third lines makes use of the Free Entry condition in steady state. Therefore, welfare per worker $\mathcal{W}$ equals the inverse of the price level, that is:

$$\mathcal{W} = \frac{\sigma - 1}{\sigma} L^{\frac{1}{\sigma-1}} \left(\bar{r} - \bar{\pi} + \widetilde{\pi}\frac{\mathcal{P}_\theta^*}{\mathcal{P}_\varphi^*}\right)^{-\frac{1}{\sigma-1}} \widetilde{\varphi}$$

where $\widetilde{\varphi}$, using the same notation as in Melitz, is the productivity of the representative firm. Note that $\widetilde{\varphi}$ is increasing in ρ: an argument akin to that of Proposition 2 would show that a higher ρ leads to higher equilibrium thresholds (θ^*, φ^*), and hence to a higher average productivity (thanks to a sharper selection by banks).

Further observe (although this is tedious to show), that for $\rho \geqslant 0$ it is:

$$\frac{\mathcal{P}_\theta^*}{\mathcal{P}_\varphi^*} \geqslant \frac{\bar{\pi}}{\widetilde{\pi}} \geqslant 1,$$

with both relationships becoming equalities if and only if $\rho = 1$. In addition, the two in-equalities widen the closer ρ gets to zero. To conclude, social welfare is maximized under the perfect information case $\rho = 1$ when the economy reduces to Melitz's, and hence the optimality result by **?** is restored. A deviation of ρ from the optimal benchmark leads to two sources of inefficiency: first, representative productivity $\widetilde{\varphi}$ falls due to a selection effect; second, the number of available varieties decreases by a factor $\bar{r}/\left(\bar{r} - \bar{\pi} + \widetilde{\pi}\mathcal{P}^*_\theta/\mathcal{P}^*_\varphi\right) < 1$, as some resources in the economy are wasted to finance entrepreneurs-firms that pass the signal threshold θ^* but fail to meet the productivity threshold φ^*.

C.4 Analysis of Proposition 4

Adding a wedge τ to firms labor costs, holding everything else equal, raises the two equilibrium threshold are raised (per Proposition 2), hence $\widetilde{\varphi}$ increases while the gap between $\bar{\pi}$ and $\widetilde{\pi}\mathcal{P}^*_\theta/\mathcal{P}^*_\varphi$ also narrows. At the same time, fewer firms can repay production costs and survive in the economy, leading to higher average revenues $\bar{r}$ and fewer product varieties. This makes the overall welfare effects of the wedge ambiguous and dependent on the specific parametrization of the model.

C.5 Extensions

We next sketch a version of the model that features PEIs. In the analysis developed so far, the equilibrium productivity distribution of the model obtains as a truncated version of the distribution firms draw their productivity from, as in the Melitz model: $\mu(\varphi) = [1 - G_0(\varphi^*)]^{-1} g_0(\varphi)$, where $g_0(\varphi)$ and $G_0(\varphi)$ are the marginal p.d.f. and c.d.f. for φ, respectively. We now allow firms to adjust their productivity after entry. Specifically, we add a further, final stage of the model where firms are allowed to *set* a productivity level $\breve{\varphi}$ subject to a decreasing cost in their original draw φ.

We specify the firm optimization problem as the difference between the additional profits obtained by raising productivity from φ to $\breve{\varphi}$ and the cost of the raise:

$$\max_{\breve{\varphi}} B\left(\breve{\varphi}^{\sigma-1} - \varphi^{\sigma-1}\right) - \kappa\left(\frac{\breve{\varphi}}{\varphi}\right)^{\alpha}, \tag{7}$$

where B is a constant that comes from the Dixit-Stiglitz analysis of monopolistic competition, while κ and α are two technological constants. We assume $\alpha > \sigma - 1$ to ensure that the cost of the raise scales faster than the benefit, thereby making the problem salient. The problem is globally concave, and the solution is straightforward:

$$\breve{\varphi} = \left[\frac{B(\sigma-1)\varphi^{\alpha}}{\alpha\kappa}\right]^{\frac{1}{\alpha+1-\sigma}}. \tag{8}$$

This delivers a monotone increasing mapping $\breve{\varphi}(\varphi)$ and an equilibrium productivity distribution expressed by $\mu(\varphi) = [1 - G_0(\breve{\varphi}^*)]^{-1} g_0\left(\breve{\varphi}^{-1}(\varphi)\right) \frac{d}{d\varphi}\breve{\varphi}^{-1}(\varphi)$, where $\breve{\varphi}^*$ is the new productivity threshold that obtains in the new equilibrium where firms have enhanced their productivity.

The implications of adding PEIs to our analysis of labor market distortions differ slightly depending on the interpretation one gives to the cost side of (7). On the one hand, wedges to labor costs definitely reduce the benefit side of (7), as they depress equilibrium profits. On the other hand, they may also raise the cost of PEIs, if the latter depends, at least in part, on human labor. We summarize these considerations with the following statement.

Proposition 5. When firms can perform PEIs as in (7), adding a wedge $\tau > 0$ to firms' labor costs (but not to either entry cost f_n or f_b), has ambiguous effects on average productivity: the positive effect due to a higher threshold (per Proposition 4) is mitigated by a negative effect due to lower PEIs. This negative effect is larger if the wedge also leads to a multiplicative increase in the cost side of PEIs.

Proof. The proof can be found in Appendix C.6. □

Adding PEIs helps make sense of our empirical results at the distribution level. Under our specification of PEIs (7), firms that are *ex ante* highly productive (on the right tail of the distribution) benefit from labor market reforms that decrease effective labor costs. Conversely, on the left tail, the selection effect dominates, which contributes to depressing average productivity. Note that our analysis is silent on the overall welfare effects of the reform: even if the net impact on average productivity is lower, consumers may still benefit from lower product varieties. In future work, we plan to provide structural estimates of the model that would let us make preliminary conclusions about the overall welfare effects.

The analysis so far was confined to a closed economy and neglected considerations about trade, as this is not the key concern of this paper. In this regard, we plan to develop a suitable extension in future work, which is natural for an extension of the Melitz framework like ours. We expect to formalize the intuition according to which adding (removing) labor market distortions harms (helps) those firms in the right tail of the productivity distribution that are more likely to engage in foreign markets.

C.6 Analysis of Proposition 5

Use the (τ) subscript to denote the values of the constants featured in (7) following the addition of a wedge τ to labor costs. From the Dixit-Stiglitz analysis of monopolistic competition one has:

$$B_{(\tau)} = \frac{B}{(1+\tau)^{\sigma-1}},$$

as both revenues and profits decrease because of higher labor costs. Let the wedge τ also cause the cost side of (7) to increase, say because part of the cost of enhancing productivity involves human resources, as follows:

$$\kappa_{(\tau)} = \kappa\,(1+\tau)^{\zeta},$$

for some $\zeta \geqslant 0$. Therefore, by (8) the wedge leads to a multiplicative transformation of the equilibrium productivity distribution, which is expressed as follows:

$$\breve{\varphi}_{(\tau)} = (1+\tau)^{\frac{\sigma-1+\zeta}{\sigma-1-\alpha}}\,\breve{\varphi} < \breve{\varphi},$$

where $\breve{\varphi}$ is as in (8), while $\breve{\varphi}_{(\tau)}$ is the updated value of post-investment productivity following the addition of the wedge.

C.7 Analysis of the two extreme cases

The critical properties of the model are perhaps best appreciated by looking at two "extreme" cases about the statistical relationship between the signal θ and productivity φ. In one case, signals are not informative at all, and the two random variables are fully independent. In the other case, the signal is fully informative, and the two random variables are perfectly correlated. The analysis of these two cases can be conducted without maintaining either Assumptions 1 or 2. Under these assumptions, however, the cases in question correspond to those where $\rho = 0$ and $\rho = 1$, respectively.

Signals are not informative. If signals deliver no information about productivity, the two random variables are independent: $G(\varphi|\,\theta) = G_0(\varphi)$ for all pairs (φ, θ), and $\widetilde{\pi}(\theta) = (1 - G_0(\varphi^*))\,\bar{\pi}$, where $\bar{\pi}$ are the expected profits conditional upon successful entry as in Melitz, for all values of θ. At stage 2. of the model, banks set their share uniformly for all firms: hence $\theta^* = 0$ and $b(\theta) = \delta f_b/(1 - G_0(\varphi^*))\,\bar{\pi}$. Back in the firm entry stage (stage 1.) it is $C(\theta^*) = 0$ and free entry reduces to:

$$\frac{(1 - G_0(\varphi^*))}{\delta}\bar{\pi} - f_b - f_n = \frac{f}{\delta}(1 - G_0(\varphi^*))\,k(\widetilde{\varphi}(\varphi^*)) - f_e = 0$$

for $f_e = f_b + f_n$ and where $\bar{\pi} = fk(\widetilde{\varphi}(\varphi^*))$ is the Zero Profit Condition (ZPC) as in Melitz. This is precisely the equilibrium condition of the original Melitz model.

Signals are fully informative. If signals predict productivity with probability one, θ^* and φ^* are jointly determined, hence one can safely focus on productivity φ only while disregarding the signal θ. To solve the model, observe that at stage 2. banks will finance only firms that are able to repay f_b in present value. This translates, relative to the Melitz benchmark, into an actual *per-period* post-entry fixed cost of $f + \delta f_b$: therefore the Zero Profit Condition becomes as $\bar{\pi} = (f + \delta f_b)\,k(\widetilde{\varphi}(\varphi^*))$. In the firm entry stage (stage 1.) entrepreneurs only need to bear their own entry cost f_n, and Free Entry implies $\bar{\pi} = \delta f_n/(1 - G_0(\varphi^*))$. Combining everything, the equilibrium solution is given as follows, and it shown to be unique by Appendix B in Melitz.

$$\frac{(1 - G_0(\varphi^*))}{\delta}\bar{\pi} - f_n = \frac{f + \delta f_b}{\delta}(1 - G_0(\varphi^*))\,k(\widetilde{\varphi}(\varphi^*)) - f_n = 0$$

One can show analytically that the second scenario leads to a higher threshold productivity value φ^*. An easier way to appreciate this is by comparing Melitz's ZCP and FE curves between the two cases: when moving from the first scenario (no information) to the second (full information), both curves shift outward in such a way that leads to a higher value of φ^*, as it is shown in Figure C.5. As in Proposition 4, the second scenario is more efficient for two reasons: entering firms are, on average, more productive, and no intermediary-specific fixed entry cost f_b is wasted on firms that eventually fail to pass the final threshold and produce.

FIGURE C.5: Analysis of the two extreme cases

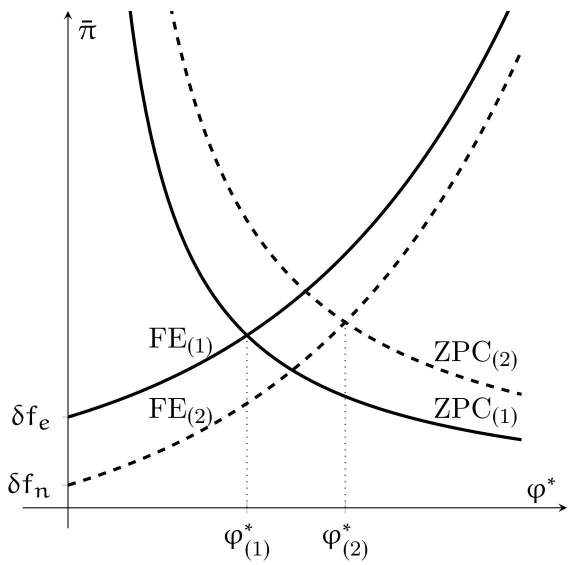

Note. Numbers apposed to curves or variables denote one of the two "extreme" scenarios, as described above.

It is interesting to analyze the effect of adding a wedge τ to labor costs in the first of the two extreme scenarios, where signals provide no information. Conditional on expected post-entry profits $\bar{\pi}$ staying constant, the wedge does not affect the cost side of firm entry decisions. However, it obviously affects the benefits side, in a way that is summarized by the ZPC, which becomes $\bar{\pi} = (1+\tau) fk(\widetilde{\varphi}(\varphi^*))$. Hence, graphically the ZPC curve shifts outward thus leading to a higher productivity threshold (from φ^* to $\varphi^*_{(\tau)}$ in the representation given by Figure C.6). Thus, the wedge τ can in principle be tailored to make the resulting productivity threshold equal to that of the "full information," efficient outcome shown in Figure C.5. Observe that this would not, however, restore the full efficiency properties of the model! The intuition is that by introducing the wedge, only the ZCP curve shifts, but the FE curve does not. In equilibrium, the lower expected profits dissuade some firms from entering, thus decreasing the number of varieties and increasing average profits. Therefore, the overall welfare effect is ambiguous, as expressed by Proposition 4 for the general case.

FIGURE C.6: Introduction of labor frictions in the Melitz model

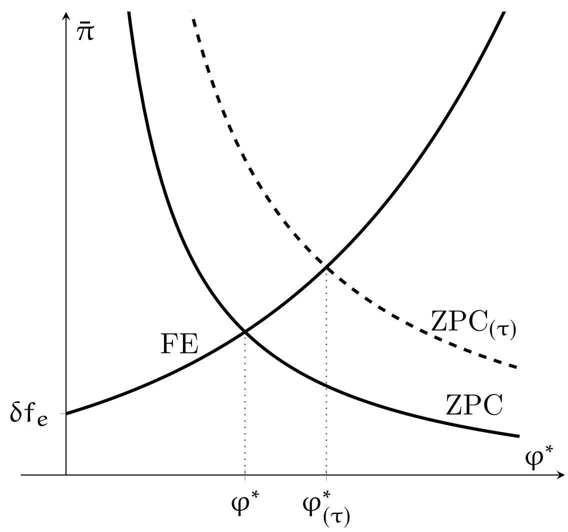

Note. τ included in a curve's or variable's subscript represents the implications of introducing labor price wedges equal to τ on it.