

The Heterogeneous Consequences of Reduced Labor Costs on Firm Productivity*

FRANCESCO DEL PRATO[†]

PAOLO ZACCHIA[‡]

February 2024

Abstract

We document how a reduction in labor costs led to heterogeneous effects on manufacturing firms' total factor productivity (TFP). Leveraging an Italian labor legislation reform and unique institutional features of the local collective bargaining system, we show that such effects vary along the TFP distribution. Relative to the counterfactual, TFP markedly declines on the left tail, which we explain via selection mechanisms; on the right, TFP mildly increases as firms are able to expand and reallocate their workforce. We develop a general equilibrium model featuring firm selection and frictions in input markets to guide the evaluation of welfare implications.

JEL classification: D22, D24, J08, O14.

Keywords: labor costs, productivity, collective bargaining, quantile effects

*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, Štěpán Jurajda, 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, SOLE 2023, the Spring 2023 UNCE workshop, EALE 2023, and the 2023 VisitINPS conference. 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, especially Roberto Bruno. Any omissions and mistakes are entirely our own.

[†]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

Economic theory and empirical evidence alike suggest that secular rises in living standards are driven by increases in the total factor productivity (TFP) of firms in the economy. While much scholarly effort went to explain TFP via technological or institutional determinants, little is known about the interplay between TFP and key input markets, particularly the labor market. This paper seeks to explore how variations in labor costs, possibly induced by policy or institutional changes, can affect the competitive environment where businesses operate and thus, endogenously, their TFP.¹ More specifically, we address three related questions. First, how does a reduction in labor costs affect firm TFP at different positions along the productivity distribution? Second, what mechanisms can account for these effects? Third, what are the associated welfare implications? The answers to these inquiries can inform economic policy in a variety of settings: from advanced economies with extensive labor legislation and industrial relations centered around collective bargaining, yet heterogeneous in their institutional details (Bhuller et al., 2022), to countries that have experienced a more recent, rapid industrialization, but whose input markets may be subject to distortions, like China (see e.g. Song et al., 2011; Cooper et al., 2018).

To address these questions, in this study we leverage an Italian reform (Decree 368/2001) originally designed to introduce new temporary labor contracts and weaken employment protection legislation (EPL), alongside some unique institutional features of the Italian system of industrial relations that we document for the first time. We show empirically that the intended effects of the Decree were muted among manufacturing firms, and yet these underwent a marked decline in their labor costs following the reform. This puzzle is most easily explained via specific characteristics of the Italian system of collective bargaining. In brief, manufacturing workers shared collective agreements with workers in services, where the reform did lead to increased diffusion of temporary contracts and, consequently, to lower bargaining power of workers (Daruich et al., 2023). At the same time, the new temporary contracts were not designed to be suitable for manufacturing firms. As a result, manufacturing firms were subject to a “clean” exogenous shock to their labor costs with no direct implications for other dimensions of their business decisions or internal organization. To the best of our knowledge, this is the first study that exploits a variation of this sort to assess consequences on firm-level outcomes, and in particular, TFP. Our key finding is that the effects are heterogeneous, as they vary along the *ex ante* productivity distribution: they are negative on the left tail and positive on the right tail.

Using comprehensive administrative data about the universe of Italian workers and firms

¹We foresee two types of mechanisms by which labor costs may affect the TFP distribution: competitive selection effects, as in conventional models of competitive equilibrium with entry and exit (Hopenhayn, 1992; Melitz, 2003); and micro-level mechanisms, where TFP changes as a consequence of firm-level choices that reflect changes in the business environment (for example, incentives for technology adoption, such as tax breaks, that are conditional on the education and qualifications of a firm’s workforce).

from the private sector, we evaluate the causal effects of the reform on manufacturing firms. We achieve this goal by exploiting its staggered implementation over time: firms whose workers were subject to different collective bargaining agreements were typically first exposed to the reform on different dates. This empirical strategy is reminiscent of those adopted in the same setting (to answer related questions) by Daruich et al. (2023), Acabbi and Alati (2021), as well as earlier by Cappellari et al. (2012). Our analysis, however, innovates in three key respects. First, restricting the analysis to manufacturing allows us, for the aforementioned reasons (which we support through a number of novel stylized facts), to evaluate the reform as a pure labor cost shifter, whereas previous studies concentrated on the EPL-reduction angle. Second, we focus on estimating heterogeneous effects on TFP at different points of the TFP distribution. We attain this in two complementary ways: by estimating quantile treatment effects and by estimating distinct average treatment effects on subsamples of firms distinguished by their *ex ante* TFP. Third, we track other firm-level outcomes, such as employment, labor productivity, capital intensity, and more, on the same subsamples. The aim of this exercise is to search for mechanisms that drive the main effects on TFP at either end of the TFP distribution.

Our main empirical results are cast in an event-study framework. In particular, causal effects are evaluated by comparing firms subject to the reform (the “treated”) against firms yet to be covered by the new rules (“not-yet-treated”). We first assess the originally intended effects of the reform: in manufacturing, there is no evidence of increased use of temporary contracts. Yet, manufacturing firms experienced a decline in total labor costs by about 5 percentage points, on average. With these results at hand, which support our proposed inter-sector wage spillover mechanism induced via collective bargaining, we proceed to estimate the main effects on the TFP of manufacturing firms. We find that, regardless of which particular method is used to measure TFP, the latter declined by 5-10 percentage points, on average, two years after the reform. However, average effects in this case hide a major source of heterogeneity. Upon restricting the sample to those firms which, before the reform, most often ranked either in the bottom or in the top quartile of the TFP distribution *of their industry*,² we find that in the former group, the decline in question amounts to 20 per cent or more. In the latter group, by contrast, the effects are reversed, as TFP increased by about 5 per cent. These results are corroborated by estimates of quantile treatment effects (based on Callaway and Li, 2019), which turn from negative to positive as one moves up along the *ex ante* unconditional distribution of manufacturing TFP.

²We adapt this approach from a related study (Devicienti and Fanfani, 2021) about the effect of collective bargaining renegotiation in Italy on a number of firm-level outcomes, including TFP. Beyond details about measurement and data selection, there are two notable differences between our contributions. First, we focus exclusively on a single reform that delivers quasi-experimental variation in labor costs. Second, their results are based on a semiparametric regression model, as opposed to our non-parametric event study tailored for the evaluation of causal effects. Whereas their results about heterogeneous effects on TFP differ qualitatively from ours, they are based on rather different measures, data, and approach, and are thus hardly comparable.

Intrigued by these findings, we search for additional evidence that would highlight what mechanisms drive our main results at both ends of the productivity distribution. We first adapt our event studies to evaluate the causal effect of the reform on the entry and exit dynamics of firms that share similar characteristics: such as industry, geographic location, and TFP level. We find that the reform decreased the exit rate of low-productivity firms, while it encouraged entry of high-productivity ones. As one would expect, lower labor costs lead to an improved business environment, making it easier for firms, especially low-productivity ones, to survive and thrive. Other results also support looser competitive selection as the mechanism driving the effects on the left tail. In particular, low-productivity firms appear to downsize following the reform, while their labor productivity also decreases. We do not rule out concurrent explanations, such as increased managerial slack induced by lower operational costs. We find it more challenging to explain what drives the positive effects on the right tail. We observe that among firms with already high TFP, lower labor costs led to faster growth (measured in terms of total employment), lower capital intensity, and lower labor productivity. All these effects match conventional predictions of economic theory. To reconcile them with the observed rise in TFP, we conjecture that firms already close to the technological frontier may draw additional efficiency gains by expanding their workforce and reallocating it across production tasks.

How to evaluate the welfare consequences of a policy that would lead to lower labor costs and, in expectation, to effects on TFP and firm dynamics similar to the ones we observe? The answer to this question is not straightforward, as these effects have ambiguous welfare implications. For example, allowing low-productivity firms to linger in the economy may lead to misallocation of capital (which could be better used to finance higher-productivity firms), but also to a more diverse set of goods available to final consumers. To guide the analysis, we propose an extension of the closed-economy version of the model by Melitz (2003), which features monopolistic competition, heterogeneous firms, and an endogenous productivity distribution. To make the model's predictions non-trivial, we introduce two types of frictions: an *ad valorem* tax on labor, which the policymaker can manipulate, and financial frictions in the firm entry stage, which enhance capital misallocation.³ As a result of decreased labor frictions, the model, despite its parsimony, reproduces many of our empirical findings, especially about firms' capital intensity, size, and productivity in the left tail of the distribution.⁴ The net welfare effects, however, depend on the quantitative

³Previous contributions in international trade extended the model by Melitz (2003) by introducing financial frictions that affect the export decisions (Manova, 2013; Chaney, 2016). By contrast, the financial frictions we introduce, modeled as a problem of asymmetric information, lead to adverse selection in the domestic market.

⁴In its current version, the model does not predict the positive TFP effects on the right tail of the distribution. In a previous version, these effects were reproduced through "productivity-enhancing investments" available to high-productivity firms, and inspired by previous contributions by Bustos (2011) and Zhelobodko et al. (2012). Such a mechanism, however, does not come with trade-offs associated with a manipulation of labor costs; in addition, this mechanism does not match the empirical evidence, as we observe decreased capital intensity in the right tail. As a result, for parsimony's sake, we removed it from the current version.

extent of financial frictions. If these are large, lower labor costs might exacerbate the costs of capital misallocation. Conversely, if they are small, the positive effects derived from the quantity and variety of goods available in the economy dominate. This suggests that the evaluation of major labor market policies should take into account their general equilibrium implications and their interplay with distortions in other input markets.

Related literature. This paper relates to several strands of literature in the fields of labor and personnel economics, and beyond.

First, we contribute to the literature on the consequences of labor market institutions centered around collective bargaining. Studies in this tradition have examined the implications of different models of industrial relations (Boeri et al., 2021; Jäger et al., 2021, 2022), the response of firms to changes in labor costs induced by changes in collective bargaining agreements (Devicienti and Fanfani, 2021; Bustos, 2023),⁵ the influence of unions on employment and wages (Farber et al., 2021), the strategies that employers adopt during collective bargaining negotiations (Prager and Schmitt, 2021), and more. Thanks to institutional features unique to Italy, this paper develops novel insights about collective bargaining, as it shows how a confined negative shock to workers' bargaining power can spill over to other sectors of the economy, leading to a more widespread fall in labor costs. In addition, it shows that labor market policies enacted within a collective bargaining framework can have far-reaching consequences on firms' productivity distribution.

Second, we extend the literature on the relationship between EPL and productivity. Previous research (for example: Autor et al., 2007; Cappellari et al., 2012; Dolado et al., 2016), explored the consequences of EPL on outcomes such as job flows, temporary employment, and firm productivity, revealing a variety of effects including productivity losses and changes in labor force composition. Our study provides a novel perspective on EPL, for two reasons. First, we show that, as a consequence of institutional features of the labor market, EPL can affect firm labor costs indirectly, that is, even for firms that are not directly exposed to the legislation. In addition, we show that the response of firm-level outcomes is likely to depend on a firm's position in the productivity distribution, and we explore the mechanisms that lead to such heterogeneous response. In a paper related to ours, Gnocato et al. (2020) show that temporary contracts improve allocative efficiency, as they reinforce the statistical link between firm size and productivity.⁶ By interpreting some of our empirical results via our conceptual framework, we show how loosened EPL, if associated with a reduction in labor

⁵The findings in the recent paper by Bustos (2023), in particular, partly echo ours: using Swedish data, he finds that labor cost *increases* following collective bargaining lead to faster firm growth, both in employment and sales. Unlike ours, that paper does not look specifically at TFP, does not examine heterogeneous effects, and only links firms to specific collective agreements via an indirect approach based on firms' skill composition.

⁶Gnocato et al. (2020) study labor market reforms enacted in Italy since 2006 (later than the one examined in this paper) which further eased the use of both fixed-term and apprenticeship contracts. In an ancillary analysis of the heterogeneous effects of these contracts on firm TFP, they obtain results aligned with ours.

costs, may lead to an excessive allocation of resources among low-productivity firms.⁷

Finally, our paper advances the literature examining the micro-level determinants of the disappointing productivity performance of Southern European economies (especially Italy) in recent decades: a topic that has attracted both scholarly and policy interests. The literature has proposed various mechanisms to explain this phenomenon: Bloom et al. (2012) suggest that lower levels of social trust affect firm internal organization; Gopinath et al. (2017) point to financial frictions as the catalyst of capital misallocation; Calligaris et al. (2018) assess the role of general resource misallocation, which they quantify for Italy following Hsieh and Klenow (2009);⁸ while Schivardi and Schmitz (2020) argue that managers struggled to capitalize on the IT revolution because of relatively lower capabilities. The empirical findings of this paper support a complementary hypothesis, originally proposed by Dew-Becker and Gordon (2012), which partly attributes the productivity slowdown to labor market reforms. Furthermore, this paper offers a framework to evaluate how different sources of allocative frictions considered in this literature: in both the labor and the capital markets, interact to jointly determine both productivity and welfare.

Paper outline. The remainder of this paper is organized as follows. Section 2 illustrates the institutional background, the reform we examine, and the stylized facts that support our empirical strategy. Section 3, alongside brief outlines of our empirical methodologies, presents our results. Section 4 describes the conceptual framework we build to draw welfare evaluations. Lastly, Section 5 concludes. A set of appendices provides additional details about data construction, summary statistics, empirical results, and model analysis.

2 Setting and data

This section comprises four parts. In the first part, we summarize some institutional features of the Italian system of industrial relations. In the second part, we discuss the labor market reform that we leverage for identification. In the third part, we describe and summarize the data. In the fourth part, we illustrate some stylized facts that are central to motivating our empirical strategy.

2.1 Collective bargaining

Italian industrial relations are regulated via collective bargaining agreements, not unlike other European countries (for example, France or Belgium). These agreements, negotiated at

⁷More indirectly, our paper also relates to studies that, theoretically and empirically, examine the effect of labor market policies other than EPL (like, for example, unemployment benefits) on firm TFP. Examples include Marimon and Zilibotti (1999), Lagos (2006) and Ortego-Martí (2020). These studies typically highlight channels that lead to better, productivity-enhancing matches between the two sides of the labor market.

⁸Gamberoni et al. (2016) conduct a similar analysis extended to the whole Euro-area.

the national level between the trade unions and the employers' associations, yield hundreds of sector-specific national collective labor contracts called *Contratti Collettivi Nazionali del Lavoro* (CCNLs). CCNLs establish minimum pay levels known as contractual floors (*minimi contrattuali*) that apply nationwide and that are set for each job title or occupation specific to an industry (typically between five and ten). Contractual floors cannot be suppressed and apply to all employees subject to the contract, regardless of union membership. As shown by Faia and Pezone (2023), to whom we defer for a more extended description of CCNLs, the passthrough between changes in the floors and wages is substantial, since it typically also affects workers whose *ex ante* wage is higher than the newly bargained floor. In practice, contractual floors function as a fixed "base" component of the wage (Fanfani, 2022). CCNLs are re-bargained according to a regular schedule (typically every three years) on contract-specific dates that are known in advance, irregularly distributed, and ultimately dependent upon the idiosyncratic history of industrial relations in a sector. The process leads contractual floors to *de facto* increase along the inflation rate (Faia and Pezone, 2023).

Italian labor law traditionally prescribes that firms apply the CCLN specific to the firm's industry and that is most relevant to the activities performed by each employee. However, the evolution of the economy over time led to an imperfect overlap between CCNLs and standard sector classifications. Because of resulting issues about the enforcement of the traditional prescription, one can typically observe multiple collective contracts coexisting within a single sector and at the same time, individual CCNLs applying to firms from very different industries (see again Fanfani, 2022, for additional discussion). This fact, which we document more extensively in section 2.4, is central to our empirical strategy.

2.2 The reform

Labor legislation in Italy distinguishes between permanent and temporary employment contracts. Permanent contracts lack a predefined termination date and require substantial severance packages (as a function of, e.g., company size and employee tenure) if an employer resolves to dismiss an employee. Temporary contracts, which are also covered by collective bargaining agreements, instead feature an explicit termination date, allowing employers to dismiss employees thereafter 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 required explicit reporting to the Italian Social Security Institute (INPS). The reform removed numerous restrictions associated with temporary contracts, leaving permanent ones unaffected and leading to the expansion of fixed-term employment.⁹

⁹The reform did not amend the existing employment protection measures for ongoing and permanent contracts: thus, differences in worker protection levels across contract types were amplified. Even post-reform, there were restrictions on the maximum duration of employment relationships based on temporary contracts.

The reform took effect across occupations upon the renewal of their corresponding CCNLs, thus resulting in staggered implementation. This allows us to leverage quasi-experimental variation in the degree of labor market flexibility (as set by the reform) across CCNLs, as done recently by other scholars. Whereas previous contributions exploited this setting to evaluate the causal effect of increased labor market flexibility on outcomes such as wages, profits, and employment (Daruich et al., 2023) as well as financial outcomes (Acabbi and Alati, 2021), in this paper we use it to estimate the causal effects a fall in labor costs on the *productivity distribution* of manufacturing firms, a decrease that is *independent of the use of temporary contracts* in a firm. This is enabled by some unique characteristics of the Italian labor market, which we detail in section 2.4, that allow us to isolate changes in labor costs *indirectly* due to the reform, that come with no direct effects (in terms of increased use of temporary contracts) among manufacturing firms.

2.3 Data

Our analysis is based on extensive administrative data about workers and firms from the Italian private sector. The two key datasets we combine, which we describe more elaborately below, are *INPS-Uniemens* (employment spells) and *Cerved* (firm-level financial data). Both are sourced from the Italian Social Security Institute (*Istituto Nazionale di Previdenza Sociale*, INPS) and accessed via the VisitINPS Scholars program. The resulting matched dataset is supplemented with information about the renewal dates of each CCNL, which are central to our empirical strategy.¹⁰

MEE Data (INPS-Uniemens) INPS data provide detailed matched employer-employee records (with monthly cadence) for all non-agricultural firms in the Italian private sector employing at least one worker. This yields comprehensive employee-level information on demographic characteristics, labor earnings, contract type (temporary, permanent, apprenticeship), and working time arrangement (part-time or full-time). Importantly, the data inform us about the CCNL applied to a worker at any point in time. On the firm side, we observe information about the establishment and cessation of a business, periods of suspended activity, industry, province of operation,¹¹ and total labor cost. We restrict 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 reliable TFP measures are unattainable. We further restrict our MEE data to only those firms also featured in the *Cerved* database; see Appendix A for details on data construction and selection.

¹⁰We express our gratitude to Raffaele Saggio and his co-authors (Daruich et al., 2023), for sharing this auxiliary dataset, which they originally elaborated from information provided by the Italian National Center for Economy and Labor (*Centro Nazionale dell'Economia e del Lavoro*, CNEL).

¹¹Provinces are intermediate administrative divisions of Italy, smaller than regions, encompassing multiple municipalities. The number of provinces has been fluctuating over time around 90-100.

Firms' Financial Data (Cerved) To construct firm-level TFP measures we use the *Cerved* proprietary firm-level data about the universe of registered balance sheets of private sector firms. The database covers the 1996-2016 period and features balance-sheet information such as revenues, value-added, labor costs, tangible and intangible assets, and the cost of materials. Where applicable, these measures are deflated using indices about industry prices and costs down to the three-digit sector level obtained from the Italian National Institute of Statistics (*Istituto Nazionale di Statistica*, ISTAT). We only deflate revenues and costs of manufacturing firms (the ones entering our analysis), as their sectors are the only ones for which ISTAT provides reliable cost indices. Procedures about data cleaning and deflation, as well as about TFP estimation, are further detailed in Appendix A.

Summary. Following data cleaning and sample restriction, we are left with a selected sample comprising 50,000 to 70,000 firms per year, which are distributed approximately over 65 three-digit manufacturing sectors. Appendix B displays and discusses relevant descriptive statistics calculated on our final dataset. Notably, the number of covered firms, while initially increasing, plateaus since 2008, as firm exit appears to increase over time. Consistently with the stagnation of the Italian economy observed over the last decades, all three TFP measures we construct display a negative trend between 2003 and 2015.

2.4 Stylized facts

We document next four key facts about the practice of collective bargaining in Italy and the adoption of different labor contracts. Jointly, these facts motivate our strategy of using the 2001 reform to evaluate the effect of a decrease in labor costs on the productivity distribution of (manufacturing) firms.

TABLE 1: Share of firms applying a single CCNL to 80% or more of their workforce

	Whole 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 businesses in the sample that apply a single CCNL to at least 80% of their workforce, broken down by size. In the left panel, the calculation is performed by pooling all yearly contracts associated with a firm between 1996 and 2016; in the right panel, calculations are conducted separately for each firm-year. empl.: employees. Source: *Istituto Nazionale della Previdenza Sociale* (INPS).

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

Supporting evidence is presented in Table 1, which shows the proportion of companies in the sample that apply a single CCNL code to at least 80 per cent of their workforce (whether

across all industries or in manufacturing only), broken down by size category and sector. Most Italian companies apply a single collective contract to almost all of their workforce, regardless of size or sector. Furthermore, such a contract remains stable over time. This is consistent with the legal prescription that firms must apply the CCNL that most closely corresponds with their production activity.¹² This fact comforts us when, in our empirical analysis, we assign to firms a *treatment*, which indicates whether Decree 368 applies to (most of) their workforce.

FIGURE 1: Diffusion of the largest CCNLs across macro-industries



Note. This Figure illustrates the extent to which the most prevalent CCNLs are used across manufacturing and service industries. Specifically, for each of the 100 most common CCNLs (ranked in terms of total number of contracts observed at the worker level over the entire panel), the Figure plots the proportion of that contract's use in manufacturing against the count ratio of observations in services vs manufacturing. The latter count is carried out both at the worker level (golden full circles) and at the firm level (blue hollow circles); in the latter case, each firm is assigned the modal CCNL observed among its workers over time. Both axes are cast on logarithmic scales. Source: *Istituto Nazionale della Previdenza Sociale (INPS)*.

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

Notwithstanding their original design and current nomenclature (which typically refers to industry-specific occupations), individual CCNLs are widely used *between*, rather than just within, sectors. Evidence of this is presented in Figure 1, which shows for each of the 100 most common CCNLs (ranked in terms of total contract count in our MEE data) 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 individual contracts, and at the firm level, i.e., by assigning a modal collective contract to

¹²There may be other reasons why firms may want to apply the same CCNL to most of its workers: for example, to decrease costs related to influencing negotiations at the national level, or to ensure predictable and equitable career progression paths among the workforce, with potential efficiency implications.

each firm (which is easily motivated via Fact 1). The wide dispersion of the points in the figure, for both counts, reflects the high degree to which CCNLs overlap across industries (if contracts were more segregated, one would expect the points to be distributed around a monotonically decreasing curve on the XY-plane). Fact 2 can be explained as a by-product of low enforcement (Garnero, 2018) and legal ambiguity regarding the matching between industries and CCNLs. In short, whereas firms appear compelled to apply a single CCNL to most of their workers, they also display discretion in the choice of *which* particular CCNL to adopt, as adherence between a CCNL and a firm’s actual production activity is hard to verify and enforce. Anecdotes about CCNLs specific to metal workers being applied to workers in professional consulting services are all but unheard of.¹³

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

Dep. variable	<i>Earnings</i>		<i>Log-earnings</i>	
	Manufacturing	Services	Manufacturing	Services
<i>Temporary share</i>				
in services	-1.37** (0.54)	-6.04* (3.20)	-0.04** (0.02)	-0.31** (0.14)
in manufacturing	-8.13** (3.81)	-0.45 (0.54)	-0.34** (0.13)	-0.03 (0.02)
<i>South share</i>				
in services	-0.98 (0.65)	-5.84*** (2.06)	-0.05* (0.03)	-0.23** (0.11)
in manufacturing	-4.03*** (1.38)	0.07 (0.48)	-0.20*** (0.06)	0.01 (0.02)
Observations	5,245	5,245	5,245	5,245
CCNL & Year fixed effects	✓	✓	✓	✓
CCNL-specific linear 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, by macro-sector. Regressions are weighted for workers’ numerosity in each macro-sector. Heteroscedasticity-consistent standard errors are in parentheses. The asterisk series *, **, and *** denote statistical significance at the 10, 5, and 1 per cent levels, respectively. Source: *Istituto Nazionale della Previdenza Sociale (INPS)* and *Centro Nazionale dell’Economia e del Lavoro (CNEL)*.

Fact 3. *Wages in one macro-sector (services or manufacturing) display a significant partial correlation with the share of temporary workers in the other macro-sector.*

The evidence supporting this last fact is presented in Table 2, which reports the estimated

¹³Whereas in the last two decades, the Italian legislation and governmental practice have moved towards increased monitoring and enforcement of the “correct” application of CCNLs, legal scholarship has debated the extent to which companies are free to choose which CCNLs to apply. Some key judicial decisions (e.g. *Corte di Appello di Cagliari* on July 13th, 2011; *TAR Lombardia* on September 4th, 2023), established that firms can choose any CCNL specific to an industry that is “not too different” from that where they actually operate.

coefficients for OLS regressions of the following kind:

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

Here, c indexes CCNLs and t time (i.e. years). The dependent variable E_{ct}^I is the average earnings, which we measure either in levels or in logs, in each CCNL, year, and macro-sector $I \in \{S(\text{ervices}), M(\text{anufacturing})\}$. Our key variables of interest are the shares of workers under temporary contracts in a macrosector-CCNL-year: TempSh_{ct}^I , for $I \in \{S, M\}$. Instead, the shares of workers located in the South of Italy in a macrosector-CCNL-year: SouthSh_{ct}^I for $I \in \{S, M\}$, control for regional differences. We additionally include CCNL and year fixed effects (α_c and ϑ_t , respectively), and CCNL-specific linear time trends (γ_c). Lastly, ε_{ct} is an error term.

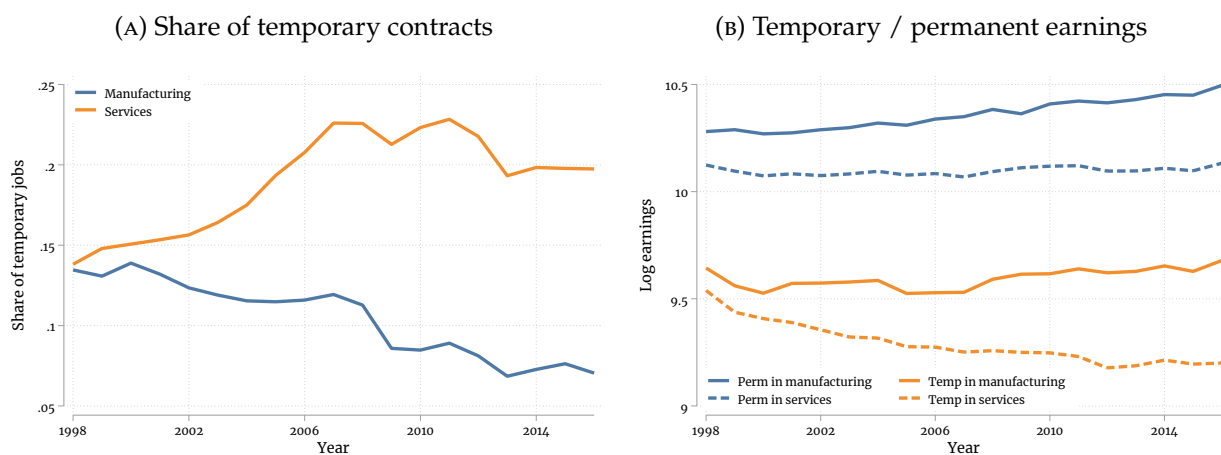
Our main coefficient of interest is β_2 in regressions about E_{ct}^M , as it allows us to evaluate the partial correlation between the share of temporary workers in services and wages in manufacturing. In both the levels and logs specification, this coefficient is negative and statistically significant: 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 per cent decrease in earnings in manufacturing, on average. A symmetric interpretation applies to coefficient β_1 in regressions about E_{ct}^S , though the corresponding estimates are not statistically significant. We interpret these results as evidence of labor market spillovers between macro-industries within a CCNL. In more mundane terms, an increase in the use of temporary contracts in services can decrease manufacturing wages, as both macro-sectors use an overlapping set of CCNLs (per Fact 2), which are regulated nationally via collective bargaining. We argue more generally that a wider diffusion of temporary contracts in a given CCNL depresses the bargaining power and thus wages of all the workers to which that specific CCNL applies, regardless of contract type; this is consistent with the results by Daruich et al. (2023).

Fact 4. *The intensity of the use of temporary contracts in manufacturing has decreased over time.*

Panel A of Figure 2 reports a relevant, and to the best of our knowledge yet unreported, fact: whereas since 2001, the share of temporary jobs in services has markedly increased, in manufacturing it has steadily declined instead (we show later there is no causal effect of Decree 368 upon this measure). Furthermore, Panel B of Figure 2 reveals that in both manufacturing and services, the total compensation of permanent workers has increased relative to that of temporary ones. Hence, while relative price considerations can explain the pattern displayed in Panel A about services, they are more difficult to reconcile with the manufacturing case. Fact 4 can be tentatively attributed to both technological¹⁴ and

¹⁴The production of durable goods does not require rapid workforce adjustments to meet fluctuations of demand, as that of non-durable goods or services. In addition, it may be easier for manufacturing firms to outsource labor through temporary employment agencies.

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



Note. Temporal evolution of selected statistic by macro-sector: share of temporary jobs in panel (A), log earnings by contract type in panel (B). Perm.: permanent [contracts]; Temp.: temporary [contracts]. Source: *Istituto Nazionale della Previdenza Sociale (INPS)*.

country-specific institutional¹⁵ characteristics of manufacturing. Exploring and quantifying the specific causes of this pattern is an interesting direction for further research.

Discussion. To motivate our empirical strategy, we offer the following cohesive summary of the four stylized facts illustrated in this subsection. In the Italian system of industrial relations, firms are largely free to choose which collective contract to apply, yet for legal and likely economic reasons, they tend to apply the same one to most of their workers. As a result of the dispersion of contract choice across firms, forces that tend to push national collective contracts in either direction (like the intensity of temporary contract usage) also affect salaries in companies that are not directly exposed to such forces. This is the case of manufacturing firms, which unlike those in services (and for idiosyncratic reasons), did not experience an increased takeup of temporary contracts in the years following the reform. Thus, as we are about to demonstrate, in manufacturing, the reform led to a clean reduction in labor costs, not accompanied by any direct consequences on the intensity of temporary contract use. This is in line with the key research question of this paper.

3 Empirical analysis

This section evaluates the causal effect of Decree 368 on Italian manufacturing. Most of our empirical estimates are based on an event study design which is described in the first of seven subsections. The ensuing subsections report the resulting estimates on TFP and other key outcomes, with two exceptions: subsection 3.5 illustrates the QDiD estimates on TFP,

¹⁵The Italian legislation traditionally provides a support scheme for manufacturing firms: the so-called *cassa integrazione*, a system of wage subsidies, which affords firms facing occasional downturns extra flexibility.

whereas subsection 3.7 provides a comprehensive discussion of the empirical evidence. Additional results, occasionally mentioned in this section, are reported in Appendix C.

3.1 Empirical strategy

We estimate the causal effect of Decree 368 on a number of outcomes of interest typically measured at the firm level. Relying on Fact 1, we define a firm as *treated* if the reform has taken effect on the CCNL that is most prevalent among its workers. Our empirical strategy exploits the staggered implementation of the reform, similarly as in other contributions (Acabbi and Alati, 2021; Daruich et al., 2023).¹⁶ As discussed in Section 2.2, the Decree took effect across CCNLs on predetermined, irregular dates. As these were easily predictable, it is plausible that firms reacted to the reform with anticipation. An event study design allows for a *prima facie* appraisal of anticipated effects via the analysis of pre-treatment differences between treated and not-yet-treated units, and their trends.

We adopt the methodology proposed by Callaway and Sant’Anna (2021)¹⁷ by estimating a number of *time-and-cohort-specific* average treatment effects on the treated (ATT):

$$ATT(g, t) = \mathbb{E} [Y_{it}(g) - Y_{it}(0) | G_i = g].$$

Here, i denotes a firm; t time measured as years; g denotes the year when the treatment occurs (“cohort”); $Y_{it}(g)$ is the value of a particular outcome Y that firm i would get at time t if it were treated on year g ; $Y_{it}(0)$ is the corresponding no-treatment counterfactual; while G_i is a categorical cohort identifier, coded as a firm i ’s treatment year. The expectation is taken over the population of firms. Because the reform eventually hits all CCNLs, we impute the counterfactuals $Y_{it}(0)$ through the outcomes of firms not treated at time t yet (“controls”). The identification of these parameters relies on a number of assumptions detailed by Callaway and Sant’Anna (2021); most notably, parallel pre-trends between treated and controls. If this assumption held unconditionally, the identification of $ATT(g, t)$ would reduce to a difference-in-long-differences comparison: thus, $ATT(2002, 2004)$ would be obtained, for example, as the average change in Y_{it} from 2001 to 2004, compared between firms treated in 2002 versus firms not yet treated in 2004.

However, it is plausible that parallel trends only hold *conditionally* on selected covariates. This would occur, for example, if the outcome of interest is TFP and CCNLs or industries whose productivity grows faster (or slower) than the average are coincidentally clustered

¹⁶We restrict the analysis to CCNL re-bargaining events that are associated with the implementation of the reform, because in manufacturing, this induced a fall in labor costs via the channel documented in Section 2.4. We do not examine “regular” episodes of CCNL re-bargaining, as these lead contractual floors to increase along inflation in a largely predictable, anticipated fashion (Faia and Pezone, 2023).

¹⁷Estimates obtained via traditional two-way fixed effects ordinary least squares specifications, which are available upon request, generally yield qualitatively similar results, though typically smaller in magnitude.

in a specific cohort. To address this, noting that the confounding factors we are concerned about are categorical industry and province identifiers, we adopt a variation of the outcome regression correction¹⁸ and estimate the $ATT(g, t)$ parameters as:

$$\widehat{ATT}(g, t) = \frac{\sum_i (\tilde{Y}_{it} - \tilde{Y}_{i(g-1)}) \mathbf{1}[G_i = g]}{\sum_i \mathbf{1}[G_i = g]} - \frac{\sum_i (\tilde{Y}_{it} - \tilde{Y}_{i(g-1)}) \mathbf{1}[G_i > t]}{\sum_i \mathbf{1}[G_i > t]},$$

where \tilde{Y}_{it} is the residual from the following regression:

$$Y_{it} = \psi_i + \eta_{s(i)t} + \lambda_{o(i)t} + \epsilon_{it},$$

while in turn, ψ_i is a firm fixed effect, $\eta_{s(i)t}$ and $\lambda_{o(i)t}$ are sector- and province-by-time fixed effects, respectively,¹⁹ and ϵ_{it} is an error term. One can easily show that if industry and geography do not affect firms from different cohorts differently in the same year (on average) this approach yields a consistent (and unbiased) estimator of the ATT parameters of interest. Standard errors are obtained via clustered bootstrap, where clusters aggregate all firms with the same three-digit sector identifier, pooling all years.

We report summaries of the estimates that we arrange graphically in an *event study* fashion. Specifically, we display the following linear combinations of the estimates, alongside their standard errors:

$$\hat{\tau}_c = \sum_g d(g) \widehat{ATT}(g, g + c),$$

where $d(g)$ is a weight specific to a cohort g which measures its frequency in the treated population. For all our outcomes of interest Y , we report results for $c \in \{-3, \dots, 3\}$.²⁰ As this paper focuses on heterogeneous effects along the productivity distribution, many figures shown in this section report $\hat{\tau}_c$ estimates for distinct subpopulations of firms, which are distinguished by their pre-reform TFP classification. In the “entry and exit” analysis from subsection 3.4, our unit of observation shifts from a firm to a “cell:” an aggregate of firms sharing the same CCNL, province, sector, and some TFP classification. Our estimation approach stays unchanged, except that standard errors are no longer clustered.

3.2 Labor market effects

We begin by examining the more direct effects of the reform on the manufacturing segment of the labor market. Specifically, we examine whether the reform has had any impact on the

¹⁸Because all confounders are categorical, we are not concerned about the specification of the outcome regression model. Conversely, inverse probability weighting is less reliable in presence of small-sized groups.


¹⁹Here, subscripts $s(i)$ and $o(i)$ denote the sector and province of firm i , respectively.

²⁰For $c = 3$, the estimation of $\hat{\tau}_3$ relies on a limited number of observations, as most CCNLs are treated in 2002 (51.07 per cent), 2003 (40.59 per cent), and 2004 (5.17 per cent); see Appendix B for graphical visualization. Hence, for most cohorts, differences observed three years after the treatment are only assessed against a small control group. This shall be taken into account when evaluating point estimates and standard errors for $\hat{\tau}_3$.

share of temporary contracts among manufacturing firms (as one would expect given the reform's content and aims) and on the labor cost of firms, measured as per-worker total earnings. The results are displayed in Figure 3.

Panel A shows that, on average, the reform did not lead to increased use of temporary contracts in the manufacturing industry: the estimated effects are all virtually zero, with a perfectly sharp trend (Panel A). While consistent with Fact 4, this is a striking result, not only because counterintuitive, but also because it contrasts with the results by Daruich et al. (2023), who show that the reform led to an increase in the use of temporary contracts across CCNLS. As we can reproduce that qualitative result among service firms, we rationalize the seeming discrepancy via Fact 2 and the unique features of (Italian) manufacturing discussed in subsection 2.4. Concurrently, labor costs underwent a significant decrease in manufacturing. Panel B depicts the estimated coefficients on per-worker labor earnings among manufacturing firms, showing a reduction of up to 5 per cent two years after the treatment, in line with other studies (Acabbi and Alati, 2021). We observe a slight pre-trend in the estimated coefficients, which might be the result of anticipating behavior by firms.

FIGURE 3: Direct labor market effects




Note. The figure reports event study coefficients $\hat{\tau}_c$ where the outcome Y_{it} is, respectively: the share of temporary contracts in a firm's workforce (Panel A), and total earnings (labor costs) per worker, in logarithms (Panel B). Confidence intervals at the 95 per cent level are obtained from clustered bootstrapped standard errors, where clusters group all firms sharing the same three-digit sector identifier across years. Source: *Istituto Nazionale della Previdenza Sociale (INPS)* and *Cerved*.

These results corroborate the discussion of the stylized facts from subsection 2.4, which is based on descriptive statistics and partial correlations, via estimates of the causal effects of Decree 368: in manufacturing, the reform did not alter the balance between permanent and temporary contracts, but nevertheless it depressed labor costs via collective bargaining spillovers. Hence, these results support our claim that this reform can be seen as a pure labor cost shifter in Italian manufacturing, which we leverage to assess its indirect effect on productivity. Note, however, that a major objective of this paper is to assess to what

extent are such indirect effects heterogeneous along the productivity distribution, and to suggest economic mechanisms to explain any difference. Should the direct labor market effects of the reform be heterogeneous along the productivity distribution, varying e.g. in their intensity or even direction, one could explain heterogeneous effects on TFP through them, rather than, say, more elaborate equilibrium arguments.²¹

FIGURE 4: Direct labor market effects, by pre-reform TFP quartiles



Note. The figure reports event study coefficients $\hat{\tau}_c$ where the outcome Y_{it} is, respectively: the share of temporary contracts in a firm’s workforce (Panel A), and total earnings (labor costs) per worker, in logarithms (Panel B). All estimates are conducted separately for the “Q1” and “Q4” subpopulations distinguished by their pre-reform TFP classification as discussed in the text. Confidence intervals at the 95 per cent level are obtained from clustered bootstrapped standard errors, where clusters group all firms sharing the same three-digit sector identifier, across years. Source: *Istituto Nazionale della Previdenza Sociale (INPS)* and *Cerved*.

To preempt this concern, it is useful to anticipate our discussion about our TFP measures. In the next section, we report estimates of causal effects on TFP measures based, for the most part, upon the methodology by Gandhi et al. (2020). To evaluate heterogeneous effects, we conduct separate estimates on different subpopulations of Italian manufacturing firms which, following Devicienti and Fanfani (2021), are distinguished by the *modal* quartile of TFP, which a firm falls into over the years leading to the reform (1996-2000), *within their own two-digits sector*. We call these groups Quartile-1 or Q1 (*ex ante* least productive firms); Q2; Q3; and Q4 (the *ex ante* most productive). We conduct this exercise for the direct labor market outcomes as well. Figure 4 reports estimates analogous to those from Figure 3 but conducted separately for the Q1 and Q4 groups.²² For both outcomes, the patterns appear very similar at the two ends of the *ex ante* productivity distribution, which is consistent with an exogenous origin of the labor market effects and which dispels *prima facie* concerns about

²¹Suppose, for example, that the reform led to an increased use of temporary contracts among firms that were already more productive before the reform, but not others. If, in addition, temporary contracts lead firms to allocate their workforce more efficiently across production tasks and over time, one would observe positive effects of TFP that are more positive on the right tail of the productivity distribution. Analogous examples involving the left tail of the distribution or heterogeneous effects on labor costs are easy to devise.


²²Q2 and Q3 are excluded from the figure to facilitate visualization.

these being direct drivers of heterogeneous effects on TFP. While the point estimates from Panel A about the share of temporary contracts in the Q1 group are positive and of sizable magnitude two years since treatment, they are too noisy to allow any firm conclusion.

3.3 Total factor productivity

We examine next the reform's effects on TFP. As the latter is notoriously difficult to measure, we conduct separate exercises about TFP estimation using the *Cerved* data, each resulting in distinct firm-level TFP measures. Our favorite one is that based on the nonparametric approach by Gandhi et al. (2020, GNR). In addition, we produce measures based on the control function, semiparametric approaches by Levinsohn and Petrin (2003, LP) as well as Akerberg et al. (2015, ACF). All TFP estimates are conducted separately across two-digit industries; each is supported by a different set of assumptions²³ (see Appendix A for details). We favor the GNR measure as it comes with the least assumptions about the functional form of production functions; the results presented in this paper treat it as our baseline. All results are qualitatively similar across the three measures.

FIGURE 5: Effects on TFP (logarithmic, multiple measures)




Note. The figure reports event study coefficients $\hat{\tau}_c$ where the outcome Y_{it} is the logarithm of one of three TFP measures (LP, ACF, and GNR) estimated as summarized in Appendix A. Confidence intervals at the 95 per cent level are obtained from clustered bootstrapped standard errors, where clusters group all firms sharing the same three-digit sector identifier, across years. Source: *Istituto Nazionale della Previdenza Sociale* (INPS) and *Cerved*.

²³We use both methods by LP and ACF to estimate gross output production functions. In this case, the identification of production function requires additional assumptions, such as *input adjustment costs* (Bond and Söderbom, 2005). These are easily motivated in the Italian setting, given the robust EPL legislation.

Figure 5 displays the average effect of Decree 368 on the three TFP measures, expressed in logarithms. The estimates are extremely similar in magnitude across the three cases and register a negative effect, in the order of 5-10 percentage points, which becomes noticeable (and statistically significant) two years after treatment. These results align with the findings by Cappellari et al. (2012), who also leverage the same reform (although they rely on a much smaller sample of CCNLs and survey data). The figure also exhibits a rather flat pre-trend, which remarkably extends up to one year post treatment, again consistently across all three measures. In light of this, one would be tempted to argue that the effect of decreased labor costs on TFP pan out with a time lag of two-three years. However, we prefer to take a cautious approach to the actual duration of this lag, due to firms' possible anticipation of the reform's direct labor market effect and because, as we are about to show, heterogeneous effects can be appreciated already one year after the treatment. Regardless of the length of the process, a lagged effect is economically plausible, since production processes are arguably slow to adjust, as are management practices.

FIGURE 6: Effects on TFP (logarithmic, GNR), by pre-reform TFP quartiles



Note. The figure reports event study coefficients $\hat{\tau}_c$ where the outcome Y_{it} is the logarithm of the GNR measure of TFP. The estimates are conducted separately for the “Q1” and “Q4” subpopulations distinguished by their pre-reform TFP classification as discussed in the text. Confidence intervals at the 95 per cent level are obtained from clustered bootstrapped standard errors, where clusters group all firms sharing the same three-digit sector identifier, across years. Source: *Istituto Nazionale della Previdenza Sociale (INPS)* and *Cerved*.

The results from Figure 5, however, hide substantial heterogeneity among firms. By following the same procedure described at the end of the previous subsection (classification of firms across groups according to their modal TFP quartile observed in each year leading to the reform; separate event study estimation by group), we uncover, in fact, divergent effects at

the two ends of the productivity distribution. Figure 6 presents the key result of this paper. Among firms that already displayed low productivity before the reform (Q1), the effect is markedly negative, coming with a reduction by ten percentage points two years after the reform, and a likely further subsequent decline (point estimates three years after the reform are remarkably close to -0.30 , but while statistically significant, they come with a sizable standard error). Conversely, already productive firms (Q4) witness a statistically significant additional rise of TFP, hovering around a five percentage points average gain since one year from the application of the reform to a firm's CCNL. Both groups display parallel pre-trends and effects that are already appreciable one year after treatment. As shown in Appendix C, we obtain qualitatively similar results whether we use the LP or ACF measures (instead of GNR) or we classify firms by their pre-reform modal *size*, rather than TFP. The appendix also reports estimates of the Q2 and Q3 groups, which are omitted here for better visualization, and that typically register values between those of the two extreme groups.

We find these results as intriguing as they are worthy of an explanation. In the rest of this section, we dissect additional empirical estimates to attempt a cohesive explanation of the economic mechanisms leading to these heterogeneous effects on TFP.

3.4 Entry and exit


The result of Figure 6 that strikes us the most at first glance is the large negative effect on the Q1 firms. To provide an explanation, we conjecture that low-productivity firms can continue their operations more easily despite competition when facing lower labor costs. Consequently, the reform should also result in decreased exit rates, conditionally on firm characteristics. In this subsection, we test this hypothesis.

We construct two auxiliary datasets providing information about the number of firms with common characteristics at any point in time between 1996 and 2016. In these datasets, the unit of observation is a "cell:" an aggregate of firms which, in any given year, share: (i) industry, (ii) province, (iii) CCNL that prevails among their workers, (iv) a certain TFP classification. The two datasets are distinguished by how TFP is classified: in one case, groups are based on the modal pre-reform quartile, as in the estimates from Figures 4 and 6; in the other case, we group firms by modal quartile *after* the reform. In each cell, we record entry and exit by counting appearances and disappearances from the firm-level panel, respectively.²⁴ Both counts are normalized by the total size of a cell in 2000. We leverage the first dataset to study the effect of the reform on firm exit and the second one to study entry: thus, we can assess to what extent the reform affects firm demographics as a function of productivity. In both cases, we perform event studies following the same

²⁴Firms persist in the MEE records as long as they maintain at least one employee in the given period.

procedure detailed in subsection 3.1, but shifting the cross-sectional unit of observation from a firm to a cell. All ATT estimates are weighted by the size of a cell.

FIGURE 7: Effects on entry and exit, by CCNL-specific aggregated “cells”



Note. The figure reports event study coefficients $\hat{\tau}_c$ where the outcome Y_{it} is the normalized of firm exits (Panel A) or entries (Panel B) within “cells” constructed as described in the text. All estimates are conducted separately for “Q1” and “Q4” subpopulations of cells. In Panel A, Q1 and Q4 refer to the modal TFP quartiles of firms *before* the reform; in Panel B, *after* the reform; TFP is estimated in both cases using the method by GNR. Estimates are weighted by cell size; confidence intervals at the 95 per cent level are obtained from bootstrapped standard errors. Source: *Istituto Nazionale della Previdenza Sociale (INPS)* and *Cerved*.

The results are reported in Figure 7. Panel A shows that, in fact, the reform led to a decline in exit rates among Q1 firms estimated between 2 and 3 percentage points two years since treatment onward, though there is no effect among Q4 firms. Conversely, Panel B displays a marked positive effect two years since treatment, amounting to 5 percentage points, for Q4 firms, whereas the corresponding estimates for Q1 do not display a uniform pattern and are not statistically significant. This analysis corroborates our original conjecture about the effect of the reform on the exit rates of low-productivity firms; in addition, it shows that the reform also enables a more forthcoming business environment for high-productivity firms, as evidenced by increased entry rates. The model that we build in Section 4 to aid the interpretation of the empirical results and welfare analysis features steady-state predictions about firm entry and the productivity exit cutoff that are in line with this evidence.

3.5 Quantile effects

Our main results from Figure 6 are obtained from an unbalanced panel of firms, due to post-reform attrition (exit). The previous subsection suggests that the negative effect in the Q1 group may be attributed at least in part to decreased exit, i.e. an extensive margin effect. It is difficult, however, to decompose the grand effect on the Q1 group between intensive and extensive margins, as we cannot observe the counterfactual TFP of the control firms that exited *because* of relatively higher labor costs; vice versa, we cannot tell which treated

firms would have exited absent the reform. Such a decomposition would help evaluate to what extent are the negative effects on TFP limited to “marginal” firms: those with a high risk of exiting, or instead involve a wider segment of the productivity distribution; an analogous question can be posed about the right tail (Q4).

To provide tentative answers to both questions at once, we can perform estimates of *Quantile Treatment Effects on the Treated* (QTTs), which we report and discuss in this subsection. For any value in the open unit segment, $u \in (0, 1)$, and for a given time t , the QTT is defined as:

$$\text{QTT}(u) = F_{\text{TFP}_{1,t}|D=1}^{-1}(u) - F_{\text{TFP}_{0,t}|D=1}^{-1}(u),$$

where $F_{\text{TFP}_{1,t}|D=1}(\cdot)$ and $F_{\text{TFP}_{0,t}|D=1}(\cdot)$ are the *unconditional* cumulative distribution functions of TFP that would be observed at time t in the population of manufacturing firms subject to the reform ($D = 1$), respectively with and without treatment. By estimating the QTT for multiple values of u we can evaluate a more nuanced impact of the reform along the entire productivity distribution. This advantage comes with costs in terms of additional assumptions and data requirements: thus, it is best to see this analysis as complementary to our main estimates of average effects restricted to subpopulations defined *ex ante*.

To estimate QTTs, we follow Callaway and Li (2019) and make two key assumptions.

QDiD Assumption 1 (Distributional Parallel Trends). *Let $\Delta\text{TFP}_{0,t} = \text{TFP}_{0,t} - \text{TFP}_{0,t-1}$. Then:*

$$\Delta\text{TFP}_{0,t} \perp\!\!\!\perp D.$$

QDiD Assumption 2 (Copula Stability). *Let $C(\Delta\text{TFP}_{0,t}, \text{TFP}_{0,t-1} | D = 1)$ be the copula between the change in untreated potential TFP and its starting level, conditional on treatment. Then:*

$$C(\Delta\text{TFP}_{0,t}, \text{TFP}_{0,t-1} | D = 1) = C(\Delta\text{TFP}_{0,t-1}, \text{TFP}_{0,t-2} | D = 1).$$

The first assumption generalizes the standard difference-in-differences parallel trends assumption, extending it from the first moment to the entire distribution of the untreated potential TFP, which is required orthogonal to treatment status. Intuitively, the extension is necessary if the parameter of interest is the change at any point in the distribution, and control firms are used to impute the counterfactuals of the treated. As shown by Fan and Yu (2012), however, alone this assumption is not sufficient. To achieve identification, Callaway and Li (2019) introduce the Copula Stability assumption,²⁵ which in this context states that among treated firms, the statistical dependence between the *level* and the *change* of TFP of the counterfactual treated firms is constant over time.²⁶ This allows one to impute the

²⁵Callaway and Li (2019) also develop versions of these two assumptions conditional on covariates.

²⁶In light of Gibrat’s Law: the empirical regularity about the statistical (in)dependence between firm size and firm growth, we find this hypothesis not too far-fetched in a firm setting.

counterfactual distribution of interest by “projecting” changes that occurred in the TFP distribution before treatment, between $t - 2$ and $t - 1$, to counterfactual changes between $t - 1$ and t . The resulting *Quantile Difference-in-Differences* (QDiD) estimator by Callaway and Li (2019) thus requires observations at least at three points in time: after treatment, before treatment, and one extra period back, which is key for the imputation step.

FIGURE 8: Quantile Treatment Effects for selected cohorts



Note. This figure presents estimates of $QTT(u)$ for $u = (.1, \dots, .9)$ and for two groups of firms: those treated in 2002 (Panel A) and those treated in 2003 (Panel B). The dependent variable is detrended TFP, estimated via the approach by Gandhi et al. (2020). Confidence intervals at the 95 percent level are constructed via bootstrap with 1,000 repetitions. Source: *Istituto Nazionale della Previdenza Sociale (INPS)* and *Cerved*.

We perform QDiD estimation on a restricted *balanced* sample of firms, focusing on detrended (i.e., purged of year effects) GNR measures of TFP. We conduct separate estimates for two cohorts of firms: those treated in 2002 and 2003, respectively (jointly, they represent almost 92 per cent of all firms). In both cases, we estimate the QTTs one year after treatment using not-yet-treated firms as controls, and restricting the sample to only those firms, both treated and control, that are observed on all three time periods required for the implementation of the QDiD estimator.²⁷ Figure 8 reports the results, one panel per cohort. The results are remarkably similar across the two groups, and display QTTs that increase monotonically along the productivity distribution. These range from a negative effect, amounting to a decrease by 5 percentage points, observed in the first decile, to an effect of slightly larger magnitude, but opposite sign, at the other extreme: the ninth decile. Most QTTs, except those estimated at the median and closely to it, are statistically significant.

²⁷These periods are selected so as to construct an appropriate time window that moves around the treatment date. Thus, for example, for the cohort treated in 2002, we estimate the effect at $t = 2003$, and we set the other two time periods required by the estimation procedure at $t - 1 = 2001$ and $t - 2 = 1999$. Similarly, for firms treated in 2003, we set $(t - 2, t - 1, t) = (2000, 2002, 2004)$. We conduct an analogous exercise to evaluate the effects on the treatment year, or two years after the treatment. The results are available upon request; they are qualitatively consistent with the main results from Figure 6. Limitations about the data and the cohort distribution (see Figure A.4 from Appendix B) prevent us from credibly stretching the QDiD estimator to evaluate effects at other time distances relative to treatment.

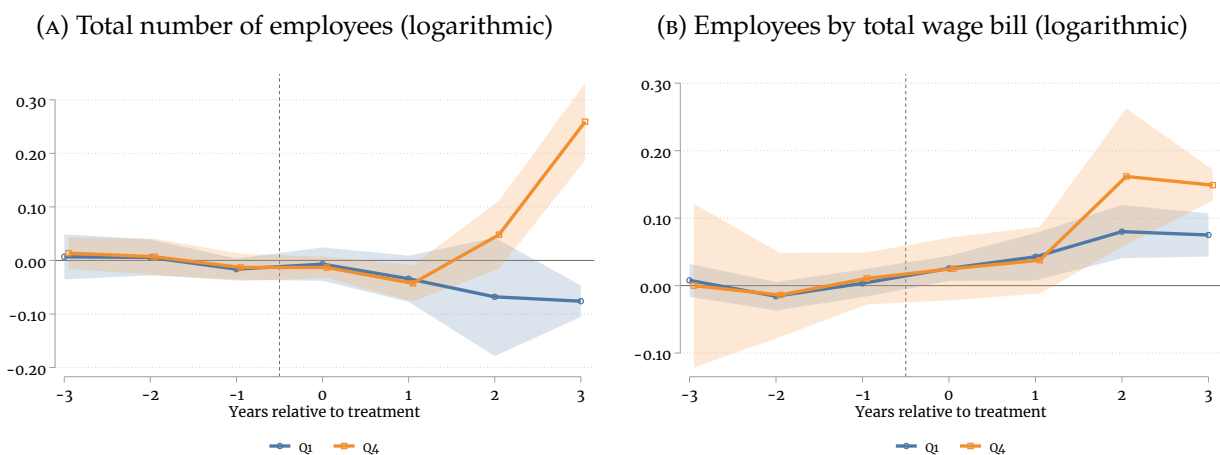
In summary, the QDiD estimates reveal that while the reform has affected the entire TFP distribution (negatively on the left, positively on the right), the consequences are more pronounced at the tails. In particular, these estimates are suggestive of intensive margin effects to the left of the median. Note, in addition, that the QDiD results are quantitatively consistent with the results about $\hat{\tau}_1$ and $\hat{\tau}_2$ from Figure 6, despite the use of a different estimator (informed by stronger assumptions) and a slightly adapted dataset.

3.6 Additional outcomes

In this subsection, we report additional results on a number of firm-level outcomes to aid the search for the mechanisms behind the main results on TFP. All ensuing ATT estimates are conducted and reported separately for the Q1 and Q4 subpopulations.

Employment. Figure 9 displays the effect of the reform on employment in both groups of firms. Panel A suggests that whereas Q4 firms have expanded, on average, following the reform, the low-productivity Q1 firms have shrunk. Consistently with a decrease in labor costs, Panel B reports that the reform led in both groups to an increase in the number of employees per unit of wage costs. This effect seems particularly pronounced for Q4 firms, whose labor/wages ratio has increased by about 16 per cent two years post treatment.

FIGURE 9: Employment effects, by pre-reform TFP quartiles




Note. The figure reports event study coefficients $\hat{\tau}_c$ where the outcome Y_{it} is the logarithm of the total number of employees in Panel A and the ratio between employees and the total wage bill in Panel B. All estimates are conducted separately for the “Q1” and “Q4” subpopulations distinguished by their pre-reform TFP classification as discussed in the text. Confidence intervals at the 95 per cent level are obtained from clustered bootstrapped standard errors, where clusters group all firms sharing the same three-digit sector identifier, across years. Source: *Istituto Nazionale della Previdenza Sociale (INPS)* and *Cerved*.

Labor productivity and capital intensity. Figure 10 displays two findings that appear puzzling at first glance; however, they can be rationalized in light of the evidence presented

so far. Panel A shows that labor productivity, defined as the ratio between deflated revenues and total workforce count, two years after the reform has declined in both groups by 10 percentage points or more. Whereas this result is expected for Q1 firms, it is counterintuitive in the Q4 case, as it contrasts with the increase in TFP from Figure 6. To reconcile these two facts, we propose this explanation: subsequent to a decrease in labor costs, high-productivity firms have expanded their workforce, as shown in Figure 9, Panel A (first-order effects); with decreasing marginal returns of labor, this leads to lower labor productivity. At the same time, there have been second-order, positive effects on TFP, though not sufficient to compensate for the first-order effect on labor productivity. While plausible, this argument does not provide a mechanism behind the increase in TFP among high-productivity firms.

FIGURE 10: Labor productivity and capital intensity, by pre-reform TFP quartiles



Note. The figure reports event study coefficients $\hat{\tau}_c$ where the outcome Y_{it} is the logarithm of: labor productivity, i.e., deflated revenues by employees (Panel A), and the capital/labor ratio (Panel B). All estimates are conducted separately for the “Q1” and “Q4” subpopulations distinguished by their pre-reform TFP classification as discussed in the text. Confidence intervals at the 95 per cent level are obtained from clustered bootstrapped standard errors, where clusters group all firms sharing the same three-digit sector identifier, across years. Source: *Istituto Nazionale della Previdenza Sociale (INPS)* and *Cerved*.

Panel B of Figure 10 reports instead results about capital intensity, measured as the ratio between a firm’s total assets and employees. The panel shows that two years after treatment occurs, the capital/labor ratio takes divergent paths across Q1 and Q4 firms: rising in the former and falling in the latter, in both cases by approximately 20 percentage points or more (in either direction). This time, the result for Q4 is intuitive, while the one for Q1 is slightly less so. In fact, a decrease in capital intensity among high-productivity firms is consistent with simple theoretical predictions about the fall of labor costs relative to the cost of capital, a pattern of workforce expansion, as well as our previous explanation for the results from Panel A. It is less immediate to explain why the capital/labor ratio increases among Q1 firms. We conjecture that a decrease in labor costs helps the survival of smaller firms, as suggested by Figure 7; to operate, however, even small firms must incur fixed capital costs. At the same time, firms of small-to-medium size that experience downturns

would still manage to survive more easily; such firms may lay off some of their workers while retaining their previously installed, slowly depreciating capital stock. Both patterns would yield, on average, a higher capital/labor ratio.

3.7 Discussion

We have shown that the Italian 368/2001 Decree, originally enacted to enhance flexibility in the labor market, had unintended effects in the manufacturing sector. In fact, the reform led to a decrease in labor costs, comparable across firms that were *ex ante* heterogeneous in their size or productivity, but divergent effects on TFP, as pre-existing differences have been magnified. Leveraging the additional results reported in this section, we now attempt an economic explanation of the mechanisms behind these divergent effects, separately for the two groups of firms around which our results are constructed and reported (Q1 and Q4). Our arguments raise additional questions that deserve future investigation.

The negative effects on the TFP of those firms with low productivity before the reform (Q1) are more easily explained via a partial equilibrium argument: thanks to lower labor costs, these firms could sustain their operations more easily in a competitive environment. This argument is especially supported by the results on firm exit in Figure 7, but it is also consistent with the pattern of the QDiD estimates in Figure 8, the results about firm size from Figure 9, and those on the capital/labor ratio in Figure 10 (as per our interpretation). We do not rule out concurrent mechanisms that are consistent with the extant empirical evidence. For example, we find it plausible that, as a second-order effect, managers and entrepreneurs of low-productivity firms decreased effort toward an incremental product, process, or organizational innovation, while the chances of firm survival had increased.²⁸ However, this hypothesis cannot be assessed with the data at hand since the *Cerved* database we can access at INPS reports neither measures of innovation nor expenditures in research and development (R&D) or other types of investment.

We find the negative effects on the TFP of those firms that were already highly productive (Q4) more puzzling. However, the empirical evidence provides a number of clues that encourage us to attempt an explanation, proceeding by exclusion. First, we rule out direct effects due to a looser regulation of the labor market, as the event study about the share of temporary contracts is flat for this group (Figure 4). Second, the explanation cannot be attributed to enhanced investment, perhaps in more modern, up-to-date equipment: this is an arguable implication of increased entry and competition (Figure 7). In fact, if anything, the capital/labor ratio decreases (Figure 10), though we cannot straightforwardly measure

²⁸A proper formulation of this hypothesis should explain why would managers and entrepreneurs have relatively fewer incentives to innovate as labor costs decrease. In an environment populated by many small firms such as the Italian one, a tentative answer can be found in entrepreneurs' *taste for firm survival*, e.g. the utility derived from the very fact of being a business owner, regardless of the firm's financial value.

investments in intangible assets such as R&D. Importantly, we note that the mechanism we are seeking is likely a byproduct of the workforce expansion coming with the reduced labor costs, as otherwise, the main result on TFP is hard to reconcile with a fall in plain labor productivity (Figure 10). We thus argue that higher-productivity firms with the ability and financial leeway to hire more workers can draw efficiency gains from the very act of expanding, as they can re-optimize their workforce across production tasks. We suggest to explore this hypothesis further in future, ideally experimental, work.

4 Welfare analysis

We present next a tractable theoretical framework that helps interpret our empirical findings and draw welfare implications from policies that affect labor market frictions.

4.1 Overview

We develop an extension on the celebrated closed-economy framework with heterogeneous firms by Melitz (2003). We facilitate exposition for readers familiar with this model by keeping a notation that closely overlaps the original. Our key innovation is the introduction of *informational financial frictions* (IFFs) in the firm entry stage. It is fair to wonder why, in an empirical paper about the interplay between the labor market and firm productivity, we choose to introduce informational asymmetries at the financing stage into a model originally built to study international trade. The answer is that we find this a convenient route to develop considerations about social welfare. In the later part of this section, we introduce an *ad valorem* distortionary tax (a “wedge”) on wages. Without IFFs, a policy that would remove such a wedge has unambiguously positive welfare consequences, since the closed Melitz economy is Pareto-optimal (Dhingra and Morrow, 2019). With IFFs, and depending on their quantitative extent, such a policy may decrease welfare, as the positive effects about available varieties and prices are offset by a fall in average productivity. Conveniently, removing the wedge also yields predictions that are close to our empirical results, especially on the left tail of the productivity distribution.

We organize the rest of this section in three parts: one about the model’s setup, one about equilibrium, and one about the welfare analysis. Technical proofs of the two propositions we formulate in this section and their related corollaries are provided in Appendix D.

4.2 Setup

We study a closed economy populated by a representative consumer whose preferences for individual goods are characterized by a Constant Elasticity of Substitution (CES). The utility function is expressed as $U^{\frac{\sigma-1}{\sigma}} = \int_{\omega \in \Omega} q(\omega)^{\frac{\sigma-1}{\sigma}} d\omega$, where Ω represents the set of

product varieties available in equilibrium, $q(\omega)$ denotes the quantity of product $\omega \in \Omega$ consumed, and $\sigma > 1$ is the elasticity of substitution. Firms, which supply the varieties in Ω , are heterogeneous in their *productivity* $\varphi(\omega) > 0$, which, importantly, we treat as exogenous; in addition, they are characterized by a linear cost function with increasing returns due to fixed costs $f > 0$. The labor demand function, also linear in quantity, is thus 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 , normalized to unity ($w = 1$).

This economy inherits the standard properties of the monopolistic competition model by Dixit and Stiglitz (1977) as extended 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 $r(\varphi)$ 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 and determined by competitive selection of firms in equilibrium. However, the details of the process differ in our model: to enter the economy, firms must secure financing for their setup from financial intermediaries. In what follows, we refer to the latter simply as “banks.”

The set of varieties Ω with associated productivity $\varphi(\omega)$ is endogenously determined by the interaction between *entrepreneurs* and *banks*. An entrepreneur is a pair (φ, θ) , where $\theta > 0$ is an individual *signal* about the true productivity φ . These two random variables are drawn from a common knowledge joint probability distribution $Q(\varphi, \theta)$; however, they are initially unobserved by all agents involved. Banks are instead 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.²⁹ Firm creation proceeds 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, firms that pay f_n and secure financing from banks set their prices and quantities; they can exit and supply zero output if, due to fixed costs, optimal profits conditional on producing are negative. Firms that stay operate in the economy until some event occurs with an exogenous probability $\delta \in (0, 1)$ forces them to exit.

²⁹This is a simplification and a normalization, as a more elaborate production function for the capital good would not significantly alter the analysis.

This model features IFFs because, 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 financial frictions (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, IFFs instead affect entry into the domestic market,³⁰ that is, firm selection in the left tail of the productivity distribution.³¹

We conduct our analysis under the following two assumptions.

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

Model Assumption 2. *Log-normality:* $Q(\varphi, \theta)$ is a cumulative bivariate (joint) log-normal distribution with standard log-normals as marginals. Let $\rho = \text{Corr}(\log \theta, \log \varphi) \in (0, 1)$.

Per Assumption 1, signals are ordered so that higher values yield conditional distributions of productivity that first-order stochastically dominate those obtained by lower values. Thus, "lower" signals imply a higher risk for banks. We introduce Assumption 2 for tractability's sake; assuming standard marginals is a mere normalization. We find the assumption that φ is log-normal realistic, not dissimilar from the Pareto assumption by Chaney (2016). Note that Assumption 1 implicitly constrains ρ to non-negative values.

4.3 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 solve the model, we thus proceed recursively and analyze how banks and entrepreneurs behave in the firm entry stage under IFFs.

Consider banks' decision on whether to finance a firm's setup or not. Since θ is the only information that banks receive about an entrepreneur, one can conveniently write the claim they demand in exchange for f_b as $b(\theta)$. Thus, when financing an entrepreneur with signal θ , a bank's expected profit is $\tilde{\pi}(\theta)b(\theta)/\delta - f_b$, where $\tilde{\pi}(\theta)$ 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 banks making zero profit in expectation. The reason is straightforward: there cannot be an equilibrium where $\tilde{\pi}(\theta)b(\theta)/\delta > f_b$ 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,

³⁰The distinction between f_n and f_b can be interpreted as a kind of liquidity constraint: of the full Melitz entry cost f_e , entrepreneurs can only pay $f_n < f_e$ upfront; $f_b < f_e - f_n$ must instead be financed by banks.

³¹Unger (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.

as they would simply deny financing to all entrepreneurs with signal values θ such that, given the corresponding share $b(\theta)$, it is $\tilde{\pi}(\theta)b(\theta)/\delta < f_b$.

By Assumption 1, these considerations imply the existence of a threshold signal that makes banks indifferent towards financing an entrepreneur or not assuming that banks capture all post-entry profits. This threshold is the smallest positive number $\theta^* \geq 0$ such that:

$$\frac{\tilde{\pi}(\theta^*)}{\delta} - f_b = 0. \quad (1)$$

We guess that a suitable interior value of θ^* exists; we verify our conjecture *ex post*. Thus, in equilibrium only those firms showing signal $\theta \geq \theta^*$ receive financing; $b(\theta^*) = 1$; and for any two signals $\theta_1 \geq \theta^*$ and $\theta_2 \geq \theta^*$, banks set shares that yield zero profits in expectation with the property that $b(\theta_1)/b(\theta_2) = \tilde{\pi}(\theta_2)/\tilde{\pi}(\theta_1)$, and $\tilde{\pi}(\theta)b(\theta) = \tilde{\pi}(\theta^*)$ for any $\theta \geq \theta^*$.³² Since (1) completely summarizes the trade-off faced by banks and the equilibrium in the capital markets, we call it (with some abuse of terminology) the Arbitrage Condition (AC). The AC 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 $v_n = \delta^{-1} \int_{\theta^*}^{\infty} \tilde{\pi}(\theta) [1 - b(\theta)] dC(\theta)$, 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, by incorporating the equilibrium in the subsequent financing subgame and solving for the value of the bank share $b(\theta)$, one obtains the following Free Entry (FE) condition in this Melitz economy with IFFs:

$$\int_{\theta^*}^{\infty} \frac{\tilde{\pi}(\theta)}{\delta} dC(\theta) - [1 - C(\theta^*)] f_b - f_n = 0. \quad (2)$$

Together with the AC in (1), this equation characterizes the economy's equilibrium. As (2) shows, entrepreneurs anticipate the probability of bearing the financing cost f_b , which they only do if they receive a signal $\theta \geq \theta^*$.

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

$$\tilde{\pi}(\theta) = \mathbb{E}_{\varphi|\theta} [\pi(\varphi)|\theta] = f \left\{ \int_{\varphi^*}^{\infty} \left(\frac{\varphi}{\varphi^*} \right)^{\sigma-1} dQ(\varphi|\theta) - [1 - Q(\varphi^*|\theta)] \right\}, \quad (3)$$


where φ^* is the threshold value of productivity below which, in equilibrium, firms find production unprofitable and exit. Note that (3) implicitly embeds a "Zero Profit Condition" (ZPC) *à la* Melitz, which is specific to a particular signal θ . Note, in addition, that both (1)

³²For the sake of exposition, we refrain from providing an extensive, formal formulation of the Bayes-Nash equilibrium of the game, inclusive of banks' strategies.

and (2) are implicitly functions of φ^* through (3). A pair of thresholds (θ^*, φ^*) completely determines the equilibrium. For convenience, we prove the existence and uniqueness of the equilibrium under Assumption 2, though other joint distributions for (θ, φ) are likely to deliver a similar result.

Proposition 1. Under Assumptions 1-2, an equilibrium pair (θ^*, φ^*) always exists and is unique. It is identified by the intersection between the set of points satisfying the AC, which is given by $\varphi^* = A (\theta^*)^\rho$ for some 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^* , which is traced out by the FE curve.

FIGURE 11: 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 (1) and the FE condition (2). The AC curve slopes upward because a higher signal threshold leads to higher productivity due to selection, and vice versa. For $\theta \leq \theta^*$, increasing productivity and signal thresholds necessitate higher profits to incentivize entry, hence they increase concurrently along the FE curve. Conversely, if $\theta \geq \theta^*$, a higher signal threshold reduces the probability of repaying the financing cost f_b to such an extent that the productivity threshold must fall to keep incentives constant. In this region, the FE curve slopes downward. The dashed lines illustrate the shift in the FE curve and the leftward rotation of the AC curve resulting from the introduction of a wedge $\nu > 0$ in the labor cost, as discussed in the text.

Figure 11 illustrates the equilibrium as the intersection between the two solid lines. The AC curve increases monotonically 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 interplay of multiple mechanisms. Like in the Melitz model, the higher the productivity threshold, the higher the profits required to motivate entry. The signal threshold, instead, affects the condition in two ways. On the one hand, a higher value of θ^* , similarly to φ^* , requires higher profits to motivate entry: the first element on the left-hand side of (2). On the other hand, it also reduces the probability of

bearing the setup cost f_b after observing the signal (second element), implicitly reducing the productivity threshold that keeps v_n constant. The first mechanism dominates at low values of θ ; the second, at higher ones. 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 and identify the threshold signal θ^* as the one where the two aforementioned mechanisms offset one another at the margin. It is instructive to study two limit cases not admitted by Assumption 2: when $\rho \rightarrow 0$ (when signals are completely uninformative) and when $\rho \rightarrow 1$ (when they perfectly predict productivity). As shown in Appendix D, both can be seen as special Melitz economies with different primitives, with φ^* being higher for $\rho \rightarrow 1$ than under $\rho \rightarrow 0$ thanks to sharper selection.

As anticipated in the overview, this model is amenable to the analysis of a distortionary wedge on wages. We specify this as an *ad valorem* tax on firms' per-period costs, including f , such that $w = 1 + v$ with $v \geq 0$, though workers still receive only one unit of numeraire. However, keeping with our interpretation of the setup costs f_n or f_b as the "capital" that firms need to operate, we leave both unaffected by the wedge (or else the analysis would be unchanged, if with a different normalization). Introducing such a wedge has the following implication on equilibrium, which we present as a corollary to Proposition 1.

Corollary 1. Introducing a wedge $v > 0$ to firms' periodic labor costs yields an equilibrium $(\theta_{(v)}^*, \varphi_{(v)}^*) \gg (\theta^*, \varphi^*)$ where both productivity and signal thresholds are strictly higher.

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 at both stages of the entry process. In the Melitz model, this effect is interpreted as 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 11.

We interpret the Italian labor market antecedent to the introduction of Decree 368 as one with stronger distortions v , directly or indirectly due to EPLs. It is thus interesting to revisit our empirical evidence in light of Corollary 1 and our model more generally. Despite its parsimony and stylized approach to firms' technology, the model remarkably predicts some of our key results from Figures 6 and 10: (i) the fall in TFP (in the model: φ) among Q1 firms, (ii) an increase of the capital/labor ratio (in the model: $(f_n + f_b)/l$) for the same group, (iii) a symmetric decrease of the capital/labor ratio among Q4 firms. In particular, (i) and (ii) follow from a looser selection of low-productivity firms, whereas (iii) is due to lower labor costs. Note that since our model, like the original one by Melitz, describes the equilibrium in steady state, its predictions must be interpreted in terms of comparative

statics. Our results about entry and exit (Figure 7) suggest a transition path consistent with the selection mechanism described by the model. Because we intentionally treat φ as exogenous in our model, the latter cannot predict the increase in TFP among Q4 firms.

4.4 Welfare

This model presents intriguing and non-trivial implications about social welfare. To evaluate the latter in this setting, one would proceed similarly to Melitz, but with two main differences. First, there are now three types of labor to be compensated: entrepreneurial (L_n), banking (L_b), and production (L_p) labor, with $L = L_n + L_b + L_p$. Second, we allow for labor market frictions, which imply that the total payment to production workers is only a fraction of the difference between firm aggregate revenues (R) and profits (Π): $L_p = (R - \Pi) / (1 + \nu)$. The expression for social welfare is, as in Melitz (2003):³³

$$\mathcal{W} = \frac{\sigma - 1}{\sigma} V^{\frac{1}{\sigma-1}} \tilde{\varphi}, \quad (4)$$

where V is the mass of active firms in equilibrium and $\tilde{\varphi}$ is their “aggregate” productivity calculated as a generalized average of order $\sigma - 1$. The relative contribution of the latter two quantities depends on the extent of IFFs as measured in our model by the parameter ρ .

We evaluate (4) in steady state. To proceed, let $\mathcal{P}_\theta^* \equiv \Pr(\theta \geq \theta^*)$ and $\mathcal{P}_\varphi^* \equiv \Pr(\varphi \geq \varphi^*)$. In equilibrium, the total remuneration of entrepreneurial labor is $L_n = V_e f_n$, while bank labor receives $L_b = \mathcal{P}_\theta^* V_e f_b$. In steady state the mass of entering firms V_e exactly compensates exits: $\delta V = \mathcal{P}_\varphi^* V_e$. Furthermore, free entry implies the following relationship:

$$\mathcal{P}_\theta^* \tilde{\pi} = \delta (\mathcal{P}_\theta^* f_b + f_n), \quad (5)$$

where $\tilde{\pi} \equiv \int_{\theta^*}^{\infty} \tilde{\pi}(\theta) dC(\theta)$ are the average profits (inclusive of bank shares) that can be expected from setting up a firm before observing the signal. Combining the above:

$$\begin{aligned} L = L_p + L_b + L_n &= V \frac{\bar{r} - \bar{\pi}}{1 + \nu} + V_e (\mathcal{P}_\theta^* f_b + f_n) \\ &= V \left(\frac{\bar{r} - \bar{\pi}}{1 + \nu} + \tilde{\pi} \frac{\mathcal{P}_\theta^*}{\mathcal{P}_\varphi^*} \right), \end{aligned} \quad (6)$$

where, as in the Melitz model, $\bar{r} = R/V$ and $\bar{\pi} = \Pi/V$ are the average equilibrium revenues and profits, respectively. To interpret this expression, consider the case where $\nu = 0$. For a fixed labor force L , the equilibrium number of firms V is proportional to the quantity $\bar{\pi} - \tilde{\pi} \mathcal{P}_\theta^* / \mathcal{P}_\varphi^*$, a measure of the difference between the expected profits conditional on passing the productivity threshold, and those conditional on passing the signal threshold,

³³As in the original monopolistic competition model by Dixit and Stiglitz (1977), social welfare equals the inverse of the price level.

hence it is nonnegative. This quantity informs the incentives that firms (and banks) face when attempting the financing stage: the higher it is, the higher the number of equilibrium entrants. In the two limit cases where either $\rho \rightarrow 0$ or $\rho \rightarrow 1$, it is $\bar{\pi} - \check{\pi} \mathcal{P}_0^* / \mathcal{P}_\varphi^* \rightarrow 0$, as signals become moot. This observation leads to our next statement.

Proposition 2. Aggregate productivity $\tilde{\varphi}$ monotonically increases with ρ ; the equilibrium number of firms V instead may be maximized for an interior value of $\rho \in (0, 1)$.

This proposition highlights the particular trade-offs behind the determination of social welfare in this model. On the one hand, more pronounced IFFs (lower ρ) lead to a looser selection of firms in the economy, and hence lower aggregate productivity $\tilde{\varphi}$. On the other hand, an excessively high value of ρ could make selection *too sharp*, leading to less product variety and higher prices. Note how this result contrasts with that by Dhingra and Morrow (2019), who show that the closed Melitz economy is Pareto-optimal: by contrast, $\rho = 1$ may not lead to welfare maximization in our model. The reason for this difference is that we introduce a two-stage entry process, where the interests of banks are not aligned with those of the representative consumers, who would rather see some financing costs f_b “wasted” if that leads to the attempted entry of some further, moderately productive firms. Note that we are unable to determine the optimal value of ρ (which maximizes \mathcal{W}) in closed form, nor to determine the conditions under which it is an interior value. We can, however, simulate welfare under different parametrizations. Such exercises show that an interior value is more likely to occur when f_b is high, as that would make banks more reluctant to finance new entrants.


An implication of this analysis is that when IFFs are especially pronounced (i.e., ρ is below the optimal value), introducing frictions to labor costs can improve social welfare.

Corollary 2. The welfare-maximizing value of ρ decreases along the labor wedge v .

A way to interpret this statement is the following: when ρ is below the optimal (interior) value, increasing labor frictions can alter market outcomes in such a way that both move towards maximal welfare. This happens as the result of two forces. On the one hand, equilibrium productivity increases (as an implication of Corollary 1). On the other hand, higher labor costs prompt fewer firms to enter. However, the former effects dominate the latter, since at the margin, the *status quo* would lead to excessive entry of low-productivity firms. Again, we are unable to produce a closed-form expression for the optimal labor frictions. However, we find it useful to illustrate our key argument with the aid of numerical calculations. Figure 12 displays the social welfare, normalized by the maximum value attainable with $v = 0$, as a function of ρ , and for three different values of the labor wedge. The figure hints that it is theoretically possible to restore the social optimum when IFFs are especially pronounced by appropriately tailoring the wedge.³⁴

³⁴A symmetric implication is that when ρ is too high, *subsidies* to firms’ labor costs, i.e., negative values of v ,

FIGURE 12: Social welfare as a function of labor frictions, illustration



Note. This figure displays the social welfare (\mathcal{W}), calculated in the model as a function of ρ , and divided by a maximum value \mathcal{W}_{\max} that can be attained for $\nu = 0$, for three different values of frictions to labor costs ν as indicated in the legend. The calculations are performed under the following parametrization: $\sigma = 2$, $L = 1$, $f = 0.2$, $f_n = 0.1$, $f_b = 1$, $\delta = 0.1$.

This analysis rules out additional effects of manipulations on firms' labor costs, which would affect welfare in either direction. On the one hand, we explicitly assumed that frictions are wasteful from the workers' point of view. It is likely, instead, that workers would derive utility from institutions, like EPL, which leads to distortions in the labor market. On the other hand, we ruled out any mechanisms that, following a decrease in labor costs, would endogenously lead to productivity gains in the right tail of the distribution, as our empirical results suggest. Inspired by Bustos (2011) and Zhelobodko et al. (2012), in a previous version of this analysis we introduced *productivity-enhancing investments* that are increasingly more convenient for high-productivity firms the lower labor costs are. However, we eventually dropped this feature of the model as it does not match the empirical findings (capital intensity falls for Q4 firms, as per Figure 10). In addition, it does not introduce interesting trade-offs. We suggest that, as future research corroborates the finding of positive effects on the TFP in the right tail, and sheds more light on the mechanisms behind them, our model and the associated welfare analysis are extended accordingly.

5 Conclusions

In this paper, we document for the first time how the TFP distribution of a large sector of the economy (manufacturing) reacts to a negative exogenous shock to labor costs. Our analysis is enabled by unique institutional conditions surrounding a particular labor market reform enacted in Italy in 2001. Designed to facilitate the use of temporary employment

would increase welfare, as they would spur more entry while reducing aggregate productivity. An internally consistent model that allows for subsidies, however, should specify how these are sourced (financed).

contracts primarily among service workers, the reform depressed average worker earnings even in manufacturing, thanks to the cross-sector transmission of collective bargaining agreements. We estimate effects that depend on a firm's position in the TFP distribution: relative to the no-treatment counterfactual, a further decline of TFP observed in the right tail is contrasted by an increase in the right tail. Thanks to additional results about other firm-level outcomes (entry, exit, capital intensity, size, labor productivity), we conclude that the effects on the left tail are best explained by a decreased competitive pressure, which facilitates the survival of low-productivity firms. Our interpretation of the right-tail effects is more speculative, and posits that more efficient firms are better equipped at implementing efficiency improvements when they can more easily expand their workforce.

We believe that each of our two main results suggests different lines for future research. Motivated by our findings on the left tail and by the welfare implications of the conceptual framework we develop in Section 4, we suggest pursuing research on the effect of labor market policies, especially those targeting low-wage workers and/or small firms, that is more conscious of equilibrium implications on other dimensions, especially firm productivity, and selection. In this regard, we find that our results align with the recent developments and findings of the literature on minimum wage (e.g. Dustmann et al., 2021). Our results on the right tail, by contrast, call for more empirical research to corroborate our findings, verify the existence of the mechanism we suggest, and possibly uncover other explanations, be they complementary or alternative to our current one. Should this finding be replicated in other observational studies, we would find it worthwhile to test it in a more controlled experimental setting, however challenging the undertaking. We are looking forward to such an opportunity.

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 2, 7, 13, and 15).
- 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 17, A.3, and A.7).
- 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 page 4).
- BHULLER, MANUDEEP, KARL OVE MOENE, MAGNE MOGSTAD, AND OLA L. VESTAD (2022)

- “Facts and Fantasies about Wage Setting and Collective Bargaining,” *Journal of Economic Perspectives*, 36 (4), 29–52, 10.1257/jep.36.4.29. (Cited on page 1).
- 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 4).
- BOND, STEPHEN AND MØANS SÖDERBOM (2005) “Adjustment costs and the identification of Cobb Douglas production functions,” in *Working Paper No. 05/04, Institute for Fiscal Studies (IFS)*, 10.1920/wp.ifs.2005.0504. (Cited on pages 17 and A.3).
- BUSTOS, EMIL (2023) “The Effect of Centrally Bargained Wages on Firm Growth,” July, 10.2139/ssrn.4405873. (Cited on page 4).
- 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 pages 3 and 34).
- 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 2, 21, and 22).
- 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 page 13).
- 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 2, 4, and 18).
- CHANEY, THOMAS (2016) “Liquidity Constrained Exporters,” *Journal of Economic Dynamics and Control*, 72, 141–154, 10.1016/j.jedc.2016.03.010. (Cited on pages 3 and 28).
- COOPER, RUSSELL, GUAN GONG, AND PING YAN (2018) “Costly Labour Adjustment: General Equilibrium Effects of China’s Employment Regulations and Financial Reforms,” *The Economic Journal*, 128 (613), 1879–1922, 10.1111/eoj.12528. (Cited on page 1).

- 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 1, 2, 7, 11, 13, and 15).
- 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 2, 4, and 16).
- DEW-BECKER, IAN AND ROBERT J. GORDON (2012) "The Role of Labor-Market Changes in the Slowdown of European Productivity Growth," *Review of Economics and Institutions*, 3 (2), 45, 10.5202/rei.v3i2.74. (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 pages 26 and 33).
- 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 27 and 32).
- 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 4).
- 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 35).
- FAIA, ESTER AND VINCENZO PEZONE (2023) "The Cost of Wage Rigidity," *The Review of Economic Studies*, 91 (1), 301–339, <https://doi.org/10.1093/restud/rdad020>. (Cited on pages 6 and 13).
- 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 21).
- FANFANI, BERNARDO (2022) "The Employment Effects of Collective Bargaining," *LABORatorio R. Revelli Working Paper no. 184*. (Cited on page 6).
- FARBER, HENRY S, DANIEL HERBST, ILYANA KUZIEMKO, AND SURESH NAIDU (2021) "Unions and Inequality over the Twentieth Century: New Evidence from Survey Data*," *The Quarterly Journal of Economics*, 136 (3), 1325–1385, 10.1093/qje/qjab012. (Cited on page 4).
- 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 16, 17, 22, and A.3).
- 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 page 10).
- 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 4).
- GOPINATH, GITA, ŞEBNEM KALEMLI-ÖZCAN, LOUKAS KARABARBOUNIS, AND CAROLINA VILLEGAS-SANCHEZ (2017) “Capital Allocation and Productivity in South Europe,” *The Quarterly Journal of Economics*, 132 (4), 1915–1967, doi.org/10.1093/qje/qjx024. (Cited on page 5).
- HOPENHAYN, HUGO A. (1992) “Entry, Exit, and firm Dynamics in Long Run Equilibrium,” *Econometrica*, 60 (5), 1127–1150, doi.org/10.2307/2951541. (Cited on page 1).
- 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 5).
- JÄGER, SIMON, SHAKED 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 page 4).
- 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 4).
- LAGOS, RICARDO (2006) “A Model of TFP,” *The Review of Economic Studies*, 73 (4), 983–1007. (Cited on page 5).
- 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 17, A.2, and A.7).
- 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 3 and 28).
- MARIMON, RAMON AND FABRIZIO ZILIBOTTI (1999) “Unemployment vs. Mismatch of Talents: Reconsidering Unemployment Benefits,” *The Economic Journal*, 109 (455), 266–291, 10.1111/1468-0297.00432. (Cited on page 5).

- MELITZ, MARC J. (2003) "The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity," *Econometrica*, 71 (6), 1695–1725. (Cited on pages 1, 3, 26, 27, 32, and A.11).
- ORTEGO-MARTI, VICTOR (2020) "Endogenous TFP, Labor Market Policies and Loss of Skills," University of California at Riverside, Department of Economics Working Papers. (Cited on page 5).
- OWEN, DONALD BRUCE (1980) "A table of normal integrals," *Communications in Statistics - Simulation and Computation*, 9 (4), 389–419. (Cited on page A.10).
- PRAGER, ELENA AND MATT SCHMITT (2021) "Employer Consolidation and Wages: Evidence from Hospitals," *American Economic Review*, 111 (2), 397–427, 10.1257/aer.20190690. (Cited on page 4).
- 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).
- SONG, ZHENG, KJETIL STORESLETTEN, AND FABRIZIO ZILIBOTTI (2011) "Growing Like China," *The American Economic Review*, 101 (1), 196–233. (Cited on page 1).
- 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 28).
- 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 pages 3 and 34).

Appendix A Data construction: details

This appendix elaborates on our choices for data cleaning and selection, as well as on the construction of some measures (TFP, entry, exit) that are central to our analysis.

A.1 Data cleaning and selection

We begin the construction of our dataset by extracting a firm-year panel dataset from the *Uniemens* database. In this dataset, we assign a unique province and industry for each observation (as a single firm might operate in more than one sector or geographical region with some branch) while keeping the observation with the highest number of employees.

We augment this panel with matched employer-employee records from *Uniemens*. As a preliminary step, we discard contracts that have lasted less than nine weeks in a year. Then, we take multiple measures in order to assign only one establishment to each worker-year. First, we resolve multiple spells *within* the same employer in a year by keeping the one that pays more. Subsequently, we resolve multiple spells over different employers within the same year by keeping the one that pays more over the entire year, and in the rare event of ties, we keep the longer-lasting spell. Lastly, we discard contracts reporting 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. In addition, for those variables that are expected to report nonnegative values (costs, revenues, purchases, assets), we replace occasional negative numbers with missing values. 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. Lastly, after estimating each TFP measure, 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 (GNR, LP, ACF) that we use.

We ultimately merge *Uniemens* with *Cerved* via unique firm identifiers. The resulting dataset provides two (highly correlated) measures of labor cost: one from balance sheets and one resulting from the aggregation of matched employer-employee records.

A.2 Deflation of balance-sheet measures

We adjust the firm-level balance sheet measures for inflation using three different price indices provided by ISTAT. First, we deflate capital measures (fixed assets and liquidity) using a purchasing power index. Secondly, we deflate revenues using an industry-specific index of final good prices, matched at the finest available digit each year. Third, we use an

imputed cost index at the industry level to adjust production inputs (net purchases and labor costs), again matching all 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 element of the matrix represents the relative weight an input has in the output costs in a given year.
2. For each sector-year pair, we construct a cost index as the resulting weighted sum of the cost indexes associated with the input sectors.
3. Each cost index is thus assigned to the best available industry-specific price index in ISTAT.

We set the base year of all three indices in 2015. In the event that industry price indices are unavailable for the 1996-1999 interval, we backcast them via the available series of the index in question. Specifically, we use an ARIMA(0, 1, 0) model augmented with external predictors (primarily the series of the lagged salary index) as we find it provides a good compromise between parsimony and out-of-sample prediction.

A.3 Estimation of TFP measures

We provide a brief summary of the three TFP measures used in this paper and introduced in the text (LP, ACF, and GNR). While we defer to the original papers, from whose authors the three acronyms are derived, for extended discussions of the estimation methodologies and the assumptions underpinning them, here we highlight some features of each approach that are most relevant for the research question and setting examined in this paper.

The notation adopted in this discussion (which is self-contained) intentionally overlaps with that found elsewhere in the paper, to comply with familiar conventions typical of the industrial organization literature on production functions. Throughout, i denotes a firm, t is time, $s(i)$ denotes the firm i 's industry (at the two digits level). In addition, for any firm i and time t , Y_{it} are deflated sales, K_{it} is capital (total assets, deflated), L_{it} is the labor force, M_{it} is the deflated cost of materials, ω_{it} is an endogenous component of productivity which leads to the transmission bias, ϵ_{it} is an exogenous (unanticipated) component, and z_{it} is the share of wages over total firm revenue.

Levinsohn and Petrin (2003, LP). Under the LP approach, we estimate a Cobb-Douglas production function with industry-specific parameters $(\beta_{K,s(i)}, \beta_{L,s(i)}, \beta_{M,s(i)})$:

$$Y_{it} = K_{it}^{\beta_{K,s(i)}} L_{it}^{\beta_{L,s(i)}} M_{it}^{\beta_{M,s(i)}} \exp(\omega_{it} + \epsilon_{it}).$$

TFP measures are obtained as the estimated residuals of the resulting log-log specification.

Estimation proceeds sector-by-sector via a two-stage control function approach: M_{it} is first used as a proxy for unobserved productivity in an initial semi-parametric stage where $\beta_{L,s(i)}$ is estimated. As shown by Akerberg et al. (2015), this approach suffers from issues of non-parametric identification, unless some specific assumptions can be maintained: for example, optimization error in the choice of the L_{it} , or if choices about L_{it} and M_{it} are performed with different timing and/or information sets. We find the latter plausible in the Italian setting, due to the particular timing of CCNL re-bargaining (see section 2.1).

Akerberg, Caves and Frazer (2015, ACF). The ACF approach is very similar to LP's, the key difference being that in the first stage, L_{it} is included in the control function (hence, $\beta_{L,s(i)}$ is only estimated in the second stage). The assumptions detailed by ACF identify a "value added" production function rid of materials; to keep our TFP measures conceptually consistent with one another, we instead employ the method by ACF to estimate the same Cobb-Douglas production function for gross output as in the LP case. Identification thus requires additional assumptions: for example, adjustment costs in the choice of flexible inputs (Bond and Söderbom, 2005). Given the existence of robust EPL legislation in Italy, we find this a tenable, realistic assumption.

Gandhi, Navarro and Rivers (2020, GNR). The GNR approach is nonparametric:

$$Y_{it} = F_{s(i)}(K_{it}, L_{it}, M_{it}) \exp(\omega_{it} + \epsilon_{it}),$$

where TFP is obtained as the residual of Y_{it} from the sector-by-sector estimation of function $F_{s(i)}(\cdot)$. The latter is, similarly to LP and ACF, also obtained from a two-stage approach. In the first stage, one estimates a function of the sector-specific labor elasticity, $D_{s(i)}(\cdot)$:

$$z_{it} = \log D_{s(i)}(K_{it}, L_{it}, M_{it}) - \epsilon_{it},$$

which is derived from the first-order conditions. The key identifying assumption is that firms adjust their choice of labor as a function of changes in labor costs in a predictable fashion, which bodes well with the setting and research question of this paper. The second stage identifies $F_{s(i)}(\cdot)$ via knowledge of $D_{s(i)}(\cdot)$, the fundamental theorem of calculus, and a polynomial function (in our case, of second degree) representing the stochastic evolution of ω_{it} . We treat both $F_{s(i)}(\cdot)$ and $D_{s(i)}(\cdot)$ as polynomials of third degree.

A.4 Exit and entry accounting

We use INPS records of a firm's registered cessation in a year to identify permanent exit events: instances where a firm permanently terminates its operation (while communicating this to the Social Security Institute). Similarly, we use the registered creation events to


identify the entry of newborn employers or establishments. Moreover, because firms appear in the INPS panel as long as they employ at least one worker (either full- or part-time; whether with a temporary or a permanent contract), we interpret disappearances from the dataset in specific years as a signal of firms' inactivity in those periods. Thus, we code a period of at least two consecutive years in which the firm does not appear in the panel as a *disappearance* event. It is important to note that disappearance events thus defined also include permanent cessation as per the official INPS records. We code *reappearance* in the panel-an indication of a firm's (re)activation-in a similar fashion.

Appendix B Selected descriptive statistics

This appendix provides additional descriptive statistics about some key variables utilized in our empirical analysis.

Sample size dynamics. Figure A.1 reports the yearly count of manufacturing firms covered in our dataset. The observed evolution is primarily due to the expansion in the number of firms' balance sheets recorded in the *Cerved* database. Our analysis is restricted to those firms that we can assign to one of the "Q" groups before the reform.

FIGURE A.1: Sample size dynamics




Note. This figure illustrates the temporal dynamics of the sample size over the sample period from 2006 to 2016. Source: *Cerved*.

Temporal evolution of TFP measures. Figure A.2 represents the time evolution of selected statistics: mean, variance, interquartile range (p75-p25), and inter-extreme range (p90-p10);

for three TFP measures (respectively the LP, ACF, and GNR measures) in our manufacturing sample. 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 in our dataset. By contrast, the interquartile and the inter-extreme ranges display a much more stable evolution over time. All observed trends appear very consistent across all three TFP measures.


FIGURE A.2: Descriptive statistics of the sample



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

Firm entry and exit. Figure A.3 depicts the dynamics of both firm entry and exit in our sample, measured at the province-by-3-digit-sector level each year. The number of exit events (including potential temporary suspensions) has increased over time; the same trend is observed when looking at rates of change. This may reflect the stagnation that occurred in Italian manufacturing over the last decades. Consistently, the number of entry events (including potential re-entries), began to decline following the 2008 financial crisis in both absolute numbers and rates of change. Both panels show that new firms continue to be established and enter the market at a higher rate than that at which they cease their activity.


FIGURE A.3: Firm dynamics in the sample



Note. This figure reports 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 firms in our manufacturing sample. "Creation:" official establishment of a firm according to INPS. "Cease:" official discontinuation of a firm according to INPS. "Panel entry [exit]:" effective appearance [disappearance] of a firm in [from] our panel, as described in Appendix A. Source: *Istituto Nazionale della Previdenza Sociale (INPS)* and *Cerved*.

Distribution of event years. Figure A.4 displays the empirical distribution of the "event years" across CCNLs; an "event year" is defined as the one when the 368/2001 Decree has become applicable to a specific CCNL. The Figure shows that for more than 90 per cent of all CNLs, the reform took effect in either 2002 (51.07 per cent) or 2003 (40.59 per cent).

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 when the 368/2001 Decree has taken effect for a particular CCNL. Source: *Istituto Nazionale della Previdenza Sociale (INPS)* and *Consiglio Nazionale dell'Economia del Lavoro (CNEL)*.

Appendix C Additional empirical results

This appendix reports additional empirical results not included in the text for the sake of exposition and conciseness, along with some brief commentaries.

Heterogeneous effects on TFP: LP and ACF measures. Figure A.5 replicates the results from 6 using the LP (Panel A) and ACF (Panel B) TFP measures, respectively. Note that the Q1-Q4 classification is based here on the LP and ACF measures, as appropriate; this contrasts with the classification used for the main results in the text, which is based on the GNR measure. The results are qualitatively similar, with two notable differences: on the one hand, the effect on Q1 three years since treatment appears to be smaller in magnitude; on the other hand, the effects on Q4 appear noisily estimated, not statistically distinguishable from zero (especially in the LP case three years since treatment).


FIGURE A.5: Effects on TFP (logarithmic, LP & ACF), by pre-reform TFP quartiles



Note. The figure reports event study coefficients $\hat{\tau}_c$ where the outcome Y_{it} is the logarithm of: the LP TFP measure (Panel A), and the ACF TFP measure (Panel B). All estimates are conducted separately for “Q1” and “Q4” subpopulations distinguished by their pre-reform TFP classification (which here is evaluated via LP or ACF estimates of TFP, as appropriate). Confidence intervals at the 95 per cent level are obtained from clustered bootstrapped standard errors, where clusters group all firms sharing the same three-digit sector identifier across years. Source: *Istituto Nazionale della Previdenza Sociale (INPS)* and *Cerved*.

Key results by pre-reform firm size classification. Figure A.6 replicates some key results of the paper: those about heterogeneous effects on per-worker labor costs (Figure 9, Panel B) as well as TFP (Figure 6) in Panels A and B respectively. However, in both cases, firms are classified by their pre-reform *size*, rather than TFP; our classification method stays otherwise unchanged (firms are assigned conditional size quartiles in each year until 2000, then the modal assignment is chosen). The results for group Q1 are very similar to the baseline, as they display a negative effect on both variables of interest. In Q4, instead, we register very noisy estimates, though the point estimates for TFP agree with the baseline.


FIGURE A.6: Key results, revisited: pre-reform classification by size



Note. The figure reports event study coefficients $\hat{\tau}_c$ where the outcome Y_{it} is the logarithm of: total labor costs per worker (Panel A), and the GNR TFP measure (Panel B). All estimates are conducted separately for “Q1” and “Q4” subpopulations distinguished by their pre-reform size classification, as discussed in this appendix. Confidence intervals at the 95 per cent level are obtained from clustered bootstrapped standard errors, where clusters group all firms sharing the same three-digit sector identifier, across years. Source: *Istituto Nazionale della Previdenza Sociale (INPS) and Cerved.*

Key results: Q2 and Q3 groups. Figure A.7 is an extended version of Figure 6 from the text, which reports results for the Q2 and Q3 groups as well. The post-treatment effects for these two subpopulations, while noisy, typically register values between those of the Q1 and Q4 groups, as expected and consistently with the QDiD estimates from Figure 8.

FIGURE A.7: Effects on TFP (logarithmic, GNR), by pre-reform TFP quartiles: all “Q” groups



Note. The figure extends Figure A.7 from the text by reporting results about the “Q2” and “Q3” subpopulations. Source: *Istituto Nazionale della Previdenza Sociale (INPS) and Cerved.*

Appendix D Additional model analysis

Analysis of Proposition 1

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

$$\begin{aligned} h &= \log \theta \\ p &= \log \varphi \\ v &= -\log \theta \\ v' &= -\log \theta + \rho(\sigma - 1) \\ z &= \frac{\log \varphi - \rho \log \theta}{\sqrt{1 - \rho^2}} \end{aligned}$$

and use asterisks to denote the values of these transformations evaluated at the corresponding threshold value of their argument(s): thus, $h^* = \log \theta^*$, $p^* = \log \varphi^*$, *et cetera* (but as an exception, $z^*(h) = (p^* - \rho h) / \sqrt{1 - \rho^2}$ is set as a function of $h = \log \theta$). 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 ; at the same time, let $\Phi_\rho(x, y)$ be the cumulative bivariate normal distribution with standard normal marginals and correlation parameter ρ , which is evaluated at a point (x, y) .

We start by elaborating expression (3) as a function of any real h :

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

Therefore, the Arbitrage Condition (1) reads:

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

with an associated implicit function $p^* = \rho h^* + a$ 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, write $\tilde{\pi} \equiv \int_{\theta^*}^{\infty} \tilde{\pi}(\theta) dC(\theta)$ as the expected joint profits that accrue to both the entrepreneur and the bank following the experimentation stage. This quantity can be expressed as a function of (h^*, p^*) :

$$\begin{aligned}
\tilde{\pi}(h^*, p^*) &= \int_{h^*}^{\infty} \tilde{\pi}(e^h) \phi(h) dh \\
&= f \int_{h^*}^{\infty} e^{(\sigma-1)(\rho h - p^*) + \frac{1}{2}(\sigma-1)^2(1-\rho^2)} \Phi\left(\frac{\rho h - p^* + (\sigma-1)(1-\rho^2)}{\sqrt{1-\rho^2}}\right) \phi(h) dh \\
&\quad - f \int_{h^*}^{\infty} \Phi\left(\frac{\rho h - p^*}{\sqrt{1-\rho^2}}\right) \phi(h) dh \\
&= f e^{\frac{1}{2}(\sigma-1)^2 - (\sigma-1)p^*} \int_{-\infty}^{-h^* + \rho(\sigma-1)} \Phi\left(\frac{-\rho v' - p^* + (\sigma-1)}{\sqrt{1-\rho^2}}\right) \phi(v') dv' \\
&\quad - f \int_{-\infty}^{-h^*} \Phi\left(\frac{-\rho v - p^*}{\sqrt{1-\rho^2}}\right) \phi(v) dv \\
&= f \left[e^{\frac{1}{2}(\sigma-1)^2 - (\sigma-1)p^*} \Phi_{\rho}(-p^* + \sigma - 1, -h^* + \rho(\sigma - 1)) - \Phi_{\rho}(-p^*, -h^*) \right],
\end{aligned}$$

where the last line follows from the analysis of the standard normal cumulative distribution's moments as in Owen (1980). Thus, write the Free Entry condition as follows:

$$\begin{aligned}
\mathcal{H}(p^*, h^*) &= e^{\frac{1}{2}(\sigma-1)^2 - (\sigma-1)p^*} \Phi_{\rho}(-p^* + \sigma - 1, -h^* + \rho(\sigma - 1)) - \\
&\quad - \Phi_{\rho}(-p^*, -h^*) - \frac{\delta f_b}{f} \Phi(-h^*) - \frac{\delta f_n}{f} = 0.
\end{aligned}$$

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

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

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

$$\begin{aligned}
\frac{\partial \mathcal{H}(p^*, h^*)}{\partial h^*} &= - \left[e^{(\sigma-1)(\rho h^* - p^*) + \frac{1}{2}(\sigma-1)^2(1-\rho^2)} \Phi\left(\frac{\rho h^* - p^* + (\sigma-1)(1-\rho^2)}{\sqrt{1-\rho^2}}\right) - \right. \\
&\quad \left. - \Phi\left(\frac{\rho h^* - p^*}{\sqrt{1-\rho^2}}\right) - \frac{\delta f_b}{f} \right] \phi(h^*),
\end{aligned}$$

which is not a monotone function of h^* . However, an analysis of this derivative shows that, for a fixed p^* , it is $\lim_{h^* \rightarrow -\infty} \partial \mathcal{H}(p^*, h^*) / \partial h^* = \lim_{h^* \rightarrow \infty} \partial \mathcal{H}(p^*, h^*) / \partial h^* = 0$; that the derivative equals exactly 0 whenever $h^* = (p^* - a) / \rho$ (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 11. Observe in addition that the line $p^* = \rho h^* + a$ can only intersect the implicit function of p^* with respect to h^* 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.

Analysis of Corollary 1

This is straightforward: as already mentioned a (and thus A) is increasing in f ; at the same time, $\partial \mathcal{H}(p^*, h^*; f) / \partial f = \delta [f_b \Phi(-h^*) + f_n] f^{-2} > 0$. Consequently, as the effective fixed cost of production shifts from f to $f(1 + \nu)$, the AC and FE curves also shift according to the pattern depicted in Figure 11. As the two curves must meet at the maximum of the FE-based implicit function of p^* over h^* , both thresholds are higher in the new equilibrium.

Analysis of Proposition 2

This analysis is split into two parts. We first show that $\tilde{\varphi}$ monotonically increases with ρ . This is straightforward given the definition of $\tilde{\varphi}$, which matches equation (9) in Melitz (2003), and that is increasing in the productivity threshold φ^* . In our model, $\tilde{\varphi}$ depends on ρ only via φ^* . Hence, to make our point, it is enough to show that φ^* is monotonically increasing in ρ . This can be attained via an analysis of the equilibrium conditions of Proposition 1 akin to the proof of Corollary 1, which we omit for brevity.

The second part of the proposition concerns the equilibrium number of entrants

$$V = L \left(\frac{\bar{r} - \bar{\pi}}{1 + \nu} + \tilde{\pi} \frac{\mathcal{P}_\theta^*}{\mathcal{P}_\varphi^*} \right)^{-1}.$$

It is helpful to first restrict the analysis to $\nu = 0$. Note that:

$$\mathcal{P}_\varphi^* \bar{\pi} - \mathcal{P}_\theta^* \tilde{\pi} = \mathcal{P}_\varphi^* \bar{\pi} - \delta (\mathcal{P}_\theta^* f_b + f_n) \geq 0:$$

this expression must be nonnegative, or else entrepreneurs would not have incentives to attempt entry. In addition, it is easy to see that this expression goes to zero either as $\rho \rightarrow 0$ or as $\rho \rightarrow 1$, since in both cases the financing stage is moot (respectively all experimenting entrepreneurs, or only the successful entrants, are financed) and thus $\mathcal{P}_\theta^* = \mathcal{P}_\varphi^*$ and $\tilde{\pi} = \bar{\pi}$. In both cases, one obtains $V = L / \bar{r}$ as in the Melitz model, though \bar{r} would take different

values at either limit, with $\bar{r}_{\rho \rightarrow 1} \geq \bar{r}_{\rho \rightarrow 0}$. Note that, adapting the analysis by Melitz:

$$\bar{r} = \sigma f \left(\frac{\tilde{\varphi}(\varphi^*)}{\varphi^*} \right)^{\sigma-1} = \frac{\Phi(\sigma-1-p^*)}{\Phi(-p^*)} \sigma f e^{-p^*(\sigma-1) + \frac{1}{2}(\sigma-1)^2}$$

which is monotonically increasing in p^* and thus in ρ (thanks to sharper selection, fewer, more productive, larger firms have higher revenue on average). This effect leads to a decrease in the number of equilibrium firms as ρ moves from 0 to 1. In our model, however, this effect is mitigated by the fact that $\bar{\pi} - \tilde{\pi} \mathcal{P}_\theta^* / \mathcal{P}_\varphi^*$ first rises, and then conceptually falls to zero again. It is thus conceptually possible, as stated in the proposition, that V is maximized by some interior value of ρ . To confirm this statement is sufficient to identify parametrizations of the model, such as that from Figure 12 for $\nu = 0$, for which this is true.

For positive values $\nu > 0$ of the wedge, the analysis stays qualitatively unchanged. One can show that:

$$\frac{\bar{r} - \bar{\pi}}{1 + \nu} = f \left\{ (\sigma - 1) \left[\frac{\Phi(\sigma-1-p^*)}{\Phi(-p^*)} \sigma f e^{-p^*(\sigma-1) + \frac{1}{2}(\sigma-1)^2} - 1 \right] - \sigma \right\}$$

where f is *not* the effective fixed cost faced by firms (which is $f(1 + \nu)$ instead). Also, this expression is monotonically increasing in p^* and thus in ρ . By contrast, $\tilde{\pi} \mathcal{P}_\theta^* / \mathcal{P}_\varphi^*$ decreases for low-to-moderate values of ρ . Hence, frictions to the labor cost ν do not prevent the existence, under appropriate parametrizations, of interior values of ρ that maximize V .

Analysis of Corollary 2

We provide a heuristic argument motivated by our prior analysis. Suppose that $\nu^* = 0$ and ρ^* is at an interior value that maximizes social welfare. As $\nu > 0$ *infinitesimally* increases to a positive value, $\tilde{\varphi}$ increases (per Corollary 1) while V decreases (fewer firms survive in equilibrium). By construction, these updated values of $\tilde{\varphi}$ and V will result in lower social welfare. Thus, for a given small $\nu > 0$, slightly lowering the measure of IFFs to a value $\rho < \rho^*$ leads to lower $\tilde{\varphi}$ and higher V : both constituent components of social welfare move towards the original optimum. If the decrease in IFFs is large enough though, $\rho \ll \rho^*$, the effects on $\tilde{\varphi}$ and V would exceed those that restore the social optimum, as average productivity in particular decreases. This implies that if an interior welfare-maximizing value of ρ exists, it decreases monotonically in ν .

Analysis of limit cases

In this section, we illustrate two limit cases of our model, not admitted under Assumption 2: $\rho \rightarrow 0$ and $\rho \rightarrow 1$. These can be viewed as two Melitz economies with different primitives. This analysis helps build better intuition about our conceptual framework.

If $\rho \rightarrow 0$, signals cannot predict productivity anymore, hence $\theta^* \rightarrow 0$ and all entrepreneurs are financed in equilibrium, even if they exit the economy after the revelation of their true productivity. In this case, the free entry condition reduces to:

$$\frac{(1 - Q(\varphi_0^*))}{\delta} \bar{\pi}_0 - f_e = 0.$$


Adopting a notation close to the one by Melitz, $f_e = f_b + f_n$; $\bar{\pi}_0$ is the expected profit conditional on surviving selection; and $Q(\varphi)$ is the marginal cumulative distribution of φ . The ZPC here is $\bar{\pi}_0 = f k(\varphi_0^*)$, where $k(\varphi^*) = [\bar{\varphi}(\varphi^*) / \varphi^*]^{\sigma-1} - 1$ like in the original Melitz model, and $\bar{\varphi}(\varphi^*)$ is the generalized average of order $\sigma - 1$ of the productivity φ of surviving firms, which equals to $\bar{\varphi}$ in (4).

If $\rho \rightarrow 1$ instead, signals perfectly predict post-entry survival and all entrepreneurs that obtain financing f_b by banks would operate in the economy. Because of perfect competition between banks, the setup cost f_b is repaid by all entering firms δf_b units per period, which effectively adds up to firms' per-period fixed costs. Thus, the free entry condition becomes:

$$\frac{(1 - Q(\varphi_1^*))}{\delta} \bar{\pi}_1 - f_n = 0,$$

while the ZPC here is $\bar{\pi}_1 = (f + \delta f_b) k(\varphi_1^*)$. Because IFFs disappear, in this scenario, Pareto optimality is restored. One can show analytically that the two productivity thresholds φ_0^* (for $\rho = 0$) and φ_1^* (for $\rho = 1$) can be ordered as $\varphi_0^* < \varphi_1^*$, as selection improves.

FIGURE A.8: Welfare analysis: graphical intuitions



Note. Panel A: analysis of the two extreme cases: $\rho \rightarrow 0$ (continuous lines) and $\rho \rightarrow 1$ (dashed lines) as Melitz economies with different primitives. Panel B: the effect of a wedge on wages in the Melitz model; the dashed line is the new ZPC curve obtained by raising $\nu > 0$.

The equilibria obtained under the two extreme cases are graphically represented on the $(\varphi^*, \bar{\pi})$ plane in Figure A.8, Panel A. The latter shows that moving from full IFFs ($\rho \rightarrow 0$) to no IFFs ($\rho \rightarrow 1$) leads to a higher productivity threshold φ^* ; however, the consequence on $\bar{\pi}$ (representing the equilibrium incentives for entrepreneurs to attempt entry) depend on the functional form of the FE and ZPC curves. The analysis of these two limit cases also helps appreciate the productivity consequences of frictions to labor costs, as per Corollary 1. To illustrate, consider the $\rho \rightarrow 0$ case: adding a strictly positive wedge $\nu > 0$ does not affect the FE condition directly, while the ZPC becomes $\bar{\pi} = (1 + \nu) f_k(\varphi_0^*)$. Graphically, this implies an outward shift of the ZPC curve as shown in Panel B of Figure A.8, leading to a higher productivity threshold (from φ^* to $\varphi_{(\nu)}^{**}$ in the Figure). The wedge ν can possibly be tailored to make the resulting productivity threshold equal to that of the “full information,” efficient outcome obtained with $\rho = 1$.