**Title**
Essays on Institutions and Resource Allocation



**Author**
Pellegrino, Bruno

**Publication Date**
2020

UNIVERSITY OF CALIFORNIA

Los Angeles

# Essays on Institutions and Resource Allocation

A dissertation submitted in partial satisfaction of the
requirements for the degree Doctor of Philosophy
in Management

by

Bruno Pellegrino

2020

ABSTRACT OF THE DISSERTATION

Essays on Institutions and Resource Allocation

by

Bruno Pellegrino

Doctor of Philosophy in Management

University of California, Los Angeles, 2020

Prof. Hugo Hopenhayn and Nico Voigtländer, Co-Chairs

The dissertation of Bruno Pellegrino is approved.

Andrew Atkeson

Edward Leamer

Stavros Panageas

Romain Wacziarg

Hugo Hopenhayn, Committee Co-Chair

Nico Voigtländer, Committee Co-Chair

University of California, Los Angeles

2020

I dedicate my dissertation to my family in Italy.

# Contents

# List of Figures

# List of Tables

# VITA

Bruno Pellegrino

## EDUCATION

Candidate of Philosophy (December 2016) in Management, University of California in Los Angeles, Los Angeles, California.

Master of Science (July 2009) in Finance and Economics, London School of Economics.

Bachelor of Arts in International Economics and Management (July 2008), Bocconi University.

## ACADEMIC EMPLOYMENT

Graduate Teaching Assistant, Anderson School of Management, University of California in Los Angeles, January 2015–present. Responsibilities include: assisting professors with the preparation and presentation of MBA courses, grading, and tutoring.

Research Assistant to Romain Wacziarg, Anderson School of Management, University of California in Los Angeles, Summer 2015. Research activities include data cleaning, management, and analysis.

Research Assistant to Luigi Zingales and Christian Leuz, Booth School of Business, University of Chicago, 2012-2014. Research activities include data collection, analysis and writing.

## ACADEMIC AWARDS

Xavier Dreze Prize, UCLA Anderson (2020).

European Economic Association, Best Job Market Paper Award (2019).

Academic Scholarship, Graduate Division, University of California in Los Angeles, 2018–2019 (five years).

Academic Fellowship, Graduate Division, University of California in Los Angeles, 2014–2018 (one year).

## PUBLISHED WORK

Social Capital and Informal Contracting: Experimental Evidence. Economics Bulletin (2019).

**TALKS AND CONFERENCE PRESENTATIONS**

**2020:** NBER Summer Institute, University of Maryland Smith, University of Cambridge, Northwestern Kellogg (Finance), Northwestern Kellogg (MEDS), University of Southern California (Economics Dept.), Texas A&M Mays, Bank of Italy, Stockholm IIES, Einaudi Institute for Economics and Finance (EIEF), EPFL-Swiss Finance Institute, Bocconi University, Society for Economic Dynamics (*Invited Session, scheduled*)

**2019:** Finance, Organizations and Markets (Los Angeles), European Economic Association (Manchester) AEA: American Economic Association (Atlanta - *session organizer*)

**2018:** European Economic Association (Köln), Econometric Society European Meeting (Köln), International Association for Applied Econometrics (Montreal), Econometric Society Asia Meeting (Seul), Bank of England/ESCoE Economic Measurement Conference (London), UC Berkeley-UCLA Political Economy Workshop (Santa Barbara)

**2017:** London Business School (Trans-Atlantic Doctoral Conference), University of Chicago Booth

**ACADEMIC SERVICE**

Referee: Econometrica, Journal of the European Economic Association, International Tax and Public Finance.

# Chapter 1

# Product Differentiation, Oligopoly and Resource Allocation

Bruno Pellegrino, UCLA[1]

[1] I gratefully acknowledge financial support from the Price Center for Entrepreneurship at UCLA Anderson. I am indebted to my dissertation advisors Hugo Hopenhayn, Nico Voigtländer, Andy Atkeson, Ed Leamer, Stavros Panageas and Romain Wacziarg for their invaluable advice and guidance. I benefited from valuable comments from David Baqaee, Ariel Burstein, Kurt Mitman, Oleg Itskhoki, Tommaso Valletti and NBER Summer Institute participants. I thank John Asker, Alberto Bisin, Michela Giorcelli, Paola Giuliano, Thomas Groesbeck, Sebastian Ottinger, Nicola Persico, Maria Lucia Yanguas for thoughtful writing and presentation advice. Finally, I acknowledge outstanding data collection assistance by Michele Schincaglia. This paper has never been submitted for peer review to any publication.

Industry concentration and corporate profit rates have increased sharply in the United States over the past two decades. This paper investigates the welfare implications of economic activity concentrating within a few firms that hold market power. I develop a general equilibrium model that features granular firms that compete in a network game of oligopoly, alongside a continuum of atomistic firms with free entry. To capture the degree of product differentiation among the oligopolists, I introduce a Generalized Hedonic-Linear (GHL) demand system. I show how to identify this demand system using a publicly-available dataset that measures product similarity among all public corporations in the US. Using my model, I estimate a large deadweight loss from oligopolistic behavior, equal to 11% of the total surplus produced by public firms. This loss would increase to 20% if all these firms were allowed to collude. The distributional effects of oligopoly are quantitatively important as well: under perfect competition, consumer surplus would double with respect to the oligopolistic equilibrium. I also estimate that the deadweight loss has increased by at least 2.5 percentage points since 1997. The share of surplus that accrues to producers as profits also has increased. Finally, I show how the dramatic rise in startups' proclivity to sell off to incumbents (rather than go public) may have contributed to these trends.

## 1.1 Introduction

Industry concentration, markups, and profit rates have all increased in the United States during the past two decades (Grullon et al., 2018; De Loecker et al., 2020). This fact has spurred important public debates over whether these trends reflect a generalized oligopolization of U.S. industries and whether a revised antitrust policy is warranted (Khan, 2018; Werden, 2018). While standard price theory arguments suggest that the welfare implications of these trends might be significant, interpreting these trends presents an imposing methodological challenge. The study of market power has traditionally resided within the domain of Empirical Industrial Organization (EIO). Yet, there is a consensus that these trends are macroeconomic in nature: standard EIO methodologies are unfeasible, as they require data that is not available for more than a handful of industries (Syverson, 2019).

This paper investigates the welfare consequences of increasing industry concentration in the United States. I address the existing methodological challenges by introducing a novel general equilibrium model with two types of firms: a finite set of *granular* firms that behave as oligopolists and a continuum of atomistic producers that behave competitively and can freely enter and exit. To model product market competition among the oligopolists, I use a hedonic demand system, which I estimate using the data set recently developed by Hoberg and Phillips (2016). This dataset provides measures of product similarity for all pairs of publicly-traded corporations in the U.S.. The empirical implementation of the model allows me address the following question: how have consumer surplus and the welfare costs of oligopoly evolved as a consequence of increased industry consolidation during this period?

Using my novel theoretical framework, I show that the increased concentration of US industries over the past twenty years was accompanied by an increase in oligopoly power, as measured by: 1) an increase in the deadweight losses induced by oligopolistic behavior; 2) a decline in the share of total surplus that accrues to consumers. My methodology also allows to me associate these trends with another well-known stylized fact: the dramatic rise in takeovers of startups that began in the mid 1990s, and which coincided with the well-known secular decline in Initial Public Offerings (IPOs) (Kahle and Stulz, 2017).

Economists have long been concerned with market power. Since the 1980s, the EIO literature has been developing a conceptual "toolkit" that researchers and antitrust enforcement practitioners have used to analyze market power within industries (Einav and Levin, 2010). The EIO approach requires the researcher to first understand the structure of product market rivalries in an industry: a firm's ability to price above marginal cost depends critically on the intensity of competition from firms that produce similar products. As a consequence, this literature has shown how oligopoly power is inextricably linked to the notion of product differentiation: to measure market power in an industry with $n$ firms, the economist effectively needs to first estimate $n^2$ cross-price demand elasticities—one for each pair of rivals. In industry studies, this is usually achieved by using a hedonic demand system (Berry, Levinsohn and Pakes, 1995).

The current resurgence in market power and antitrust research, however, has a distinctive macroeconomic angle. Because we do not observe output volume, prices, or product characteristics for a sufficiently large cross-section of industries, the EIO approach cannot be directly applied in a macroeconomic context. This challenge is compounded by the problem that, even at the macro level, product-market rivalry is not well approximated by industry classifications. Industry classifications (such as NAICS) tend to be based on similarities in the production process, not on the degree of product substitutability. In other words, they are appropriate for estimating production function, but they are unreliable when it comes to measuring the cross-price elasticity of substitution between products. In addition, the very concept of industry/sector is more fluid than (macro)economists have traditionally tended to assume. While industry classifications are static, larger companies (those more likely to have market power) move frequently from one industry to another, and have been shown to strategically manipulate their industry classification—a phenomenon that has been dubbed *industry window dressing* (Chen et al., 2016).

Despite these challenges, the macroeconomics literature has made significant progress in incorporating market power into general equilibrium models: Baqaee and Farhi (2020, henceforth BF) have recently shown how to approximate the welfare costs of markups, under minimal assumptions, using the cross-sectional distribution of markups. This approach—by design—is agnostic about the origin of the observed variation in markups: its advantage is that it captures all observed variation in markups (and therefore all sources of inefficiency); its downside is that it does not model how the observed dispersion in markups originates in the first place. Therefore, a separate theory of markups formation is required to simulate changes in market structure.

This study breaks new ground by providing a theory of firm size and profitability that generalizes the Cournot oligopoly model to differentiated products and hedonic demand, and embeds it in a general equilibrium model. The objective of my model, rather than capturing all sources of variation in markups, is to isolate the variation in firm size and markups that can be reliably attributed to product market rivalry. Through this approach, I can quantify the contribution of each individual producer to aggregate welfare, and I can study the general equilibrium effects of events that are relevant to antitrust policy, such as mergers or the entry of additional firms.

To achieve this, my theoretical model dispenses with the notions of industry and sector altogether, building instead on the tradition of hedonic demand (Lancaster, 1966; Rosen, 1974). Thus, I can link the cross-price elasticity of demand between all firms in the economy to the fundamental attributes of each firm's product portfolio. Each firm's output is modeled as a bundle of characteristics that are individually valued by the representative consumer. The cross-price elasticity of demand between two firms depends on the characteristics embedded in their output. If the product portfolios of two companies contain similar characteristics, the cross-price elasticity of demand between their products is high. The result is a rather different picture of the product market: not a collection of sectors, but a network, in which the distance between nodes reflects product similarity and strategic interaction between firms.

The main assumptions of my model are: (1) the representative consumer holds a linear-quadratic hedonic utility (a generalization of Epple, 1987); (2) firms compete à la Cournot[2]; (3) the marginal cost function is linear in output. Based on these assumptions, the firms in my model play a linear-quadratic game over a weighted network, a type of potential game that has been extensively studied in the micro theory literature (see Ballester, Calvó-Armengol and Zenou, 2006; Ushchev and Zenou, 2018).

This is the first paper to show how to derive the network Cournot model starting from a hedonic utility specification, to embed the game in a general equilibrium framework and to take the model to the data in a structural way. I use a recently-developed data set (Hoberg and Phillips, 2016, henceforth HP) that provides measures of product similarity for every pair of publicly traded firms in the United States. These product-similarity scores—which are based on a computational-linguistics analysis of the firms' regulatory 10-K forms—give rise to a continuous, high-dimensional representation of the product space. My model maps these bilateral similarity scores to an $n \times n$ matrix of cross-price demand elasticities. Moreover, because HP's similarity scores are time-varying (yearly observations since 1997), my model is unique in that the degree of product substitution between firms is allowed to change over time.

Perhaps even more importantly, the empirical implementation of my model does not require any proprietary or confidential data, and is computationally tractable. The two datasets it requires are Compustat (which is purchased by most economics departments and business schools), and the HP's cosine similarity data, which the authors have made publicly-accessible through an online repository.[3]

I use my model to compute the (static) deadweight loss from oligopoly and to simulate changes in total surplus and consumer surplus for a number of policy counterfactuals. I find that the welfare costs of oligopoly are sizable. By moving to an allocation in which firms price at marginal cost (that is, in which they behave as if they were atomistic players in a perfectly competitive market), total surplus would rise by approximately 10.7 percentage points; consumer surplus would double, partly due to total surplus being reallocated from producers to consumers. By computing a separate counterfactual that only rectifies allocative distortions (markups are equalized, rather than eliminated, and labor supply is assumed to be inelastic), I can determine that a significant share of the welfare losses from oligopoly—about 6 percentage points of the aforementioned 10.7—occur by way of factor misallocation. In other words, the deadweight losses are driven not only by an underutilization of inputs, but also by a suboptimal mix of goods being produced. I also simulate a counterfactual in which all firms in the economy are owned by a single producer that implements a collusive equilibrium. Under this scenario,

[2] I also study the Bertrand case in the Online Appendix.

[3] See hobergphillips.tuck.dartmouth.edu

total surplus would drop by about one tenth: with some degree of abstraction, we can think of this estimate as an upper bound to the welfare benefits of Antitrust. Also, in this monopolistic/collusive equilibrium consumer surplus would decrease by about 38%, due partly to surplus being reallocated from consumers to producers.

By mapping my model to the data for a period of 21 consecutive years, I can investigate the welfare consequences of the observed trends in concentration and markups between 1997 and 2017. I find that the share of surplus appropriated by companies in the form of oligopoly profits has increased from about 50.7% (in 1997) to nearly 56% (in 2017). When fixed costs (such as capital and overhead) are subtracted from profits and total surplus, this increase becomes significantly steeper: from 17% in 1997 to 28% in 2017. This result is robust to different measurements of fixed costs and intangible capital, and suggests that the increase in the profit share of surplus is not justified by larger fixed costs.

The welfare costs of oligopoly have also increased over this period. In terms of total surplus, the gap between the oligopolistic equilibrium and the first best has increased from 8.2% (in 1997) to 10.7% (in 2017). The resulting effect on the consumer could be best described as a double whammy: less surplus is produced overall (as a percentage of the surplus that *could* be produced), and less of the diminished surplus is allocated to the consumer in equilibrium. Thus, another unique contribution of this work is the ability to dig deeper into the distributional implications of the rise in industry concentration.

Finally, I use the counterfactual-building capabilities of the model to better understand the causes of rising oligopoly power. In particular, I study the effects of the dramatic secular shift in the type of venture capital (VC) exits observed in the past 20 years:[4] in the early 1990s, most VC-backed startups (80%–90%), if successful, would exit through IPOs. Today, the near entirety (about 94%) of the successful VC exits conclude with the startup being acquired by an incumbent. I find that this shift accounts not only for the secular decline in the number of public corporations in the United States (from about 7,500 in 1997 to about 3,500 in 2017) but also for the measured increase in the welfare costs of oligopoly, as well as the rising profit share of surplus. Overall, my results suggest that increased concentration and markups resulted in sizable welfare losses and affected how surplus is shared between producers and consumers.

This paper aims to bridge the new EIO literature (Einav and Levin, 2010) with two recent and growing branches of macroeconomics that use micro-data.

The first is the literature on networks (Acemoglu, Carvalho, Ozdaglar and Tahbaz-

[4] In the entrepreneurial finance literature, an "exit" is the termination of a VC investment and should not be confused with a business termination. If the VC investor exits with an IPO, that event marks the entry of that firm in the universe of public firms, not an enterprise death.

Salehi, 2012; Acemoglu, Ozdaglar and Tahbaz-Salehi, 2017; Carvalho and Tahbaz-Salehi, 2019; Carvalho and Grassi, 2019; Baqaee and Farhi, 2020; Carvalho, Nirei, Saito and Tahbaz-Salehi, 2020). I contribute to and expand this literature, which has mostly focused on input-output networks, by considering a different type of network: that of product market rivalry relationships. Bloom, Schankerman and Van Reenen (2013) have previously explored this type of network in a seminal empirical study of R&D spillovers. This paper develops a formal theory of product market rivalry and oligopoly and embeds it in a general equilibrium environment.

The second is the recent macro literature on markups, concentration and superstar firms (De Loecker, Eeckhout and Unger, 2020; Autor, Dorn, Katz, Patterson and Van Reenen, 2020; Edmond, Midrigan and Xu, 2018). This paper builds on and expands this body of work by incorporating hedonic demand as well as new data. These features allow me to go beyond markups and concentration, and to create a rich, high-dimensional representation of the competitive environment. In my model, firms differ not only by their productivity, but also by their products' characteristics; as a consequence, each firm has a distinct set of competitors that changes over time, as firms update their business description in their regulatory filings.

This paper also connects the recent literature on market power to the secular decline in the number of public companies (Kahle and Stulz, 2017; Doidge et al., 2018) and IPOs (Bowen, Frésard and Hoberg, 2018; Gao, Ritter and Zhu, 2013). My model provides an avenue to quantify the effects of these phenomena on the intensity of product market competition.

The rest of the paper is organized as follows. In Section 1.2, I present my theoretical model. In Section 1.3, I present the data used throughout the empirical part of the paper (including HP's data set) and show how it can be mapped to the model. In Section 1.4, I present empirical results. In Section 1.5, I discuss a number of extensions and robustness checks. In Section 1.6, I present my conclusions and discuss how my findings can inform the current debate on market power and antitrust policy.

## 1.2 A Theory of Imperfect, Networked Competition

In this section, I present a general equilibrium model in which firms produce differentiated products and compete à la Cournot. For expositional purposes, I start by laying out the basic model that only includes granular oligopolistic firms. After characterizing the equilibrium of this model economy and outlining a series of counterfactuals of interest, I extend the model (in subsection 1.2.6) by adding a continuum of perfectly-competitive atomistic firms.

### 1.2.1 Basic Setup: the Generalized Hedonic Linear (GHL) Demand System

There are $n$ firms, indexed by $i \in \{1, 2, ..., n\}$ that produce differentiated products. Following the tradition of hedonic demand in differentiated product markets (Lancaster, 1966; Rosen, 1974), I assume that consumers value each product as a bundle of characteristics. The number of characteristics is $k + n$ .

There are two types of characteristics. The first $k$ characteristics are common across all goods and are indexed by $j \in \{1, 2, ..., k\}$, while the remaining $n$ characteristics are idiosyncratic (that is, they are product-specific and cannot be imitated by other products) and therefore have the same index $i$ as the corresponding product. The scalar $a_{ji}$ is the number of units of common characteristic $j$ provided by product $i$. Each product is described by a $k$-dimensional column vector $\mathbf{a}_i$ , which I assume (without loss of generality) to be of unit length – formally:

$$\mathbf{a}_i = \begin{bmatrix} a_{1i} & a_{2i} & \dots & a_{ki} \end{bmatrix}'$$

$$\text{such that} \qquad \sum_{j=1}^{k} a_{ji}^2 = 1 \quad \forall\, i \in \{1, 2, ..., n\}$$

The assumption that $\mathbf{a}_i$ is of unit length amounts to a normalization assumption. For every product, we need to pick an output volume metric (kilograms, pounds, gallons, etc.). The normalization consists in picking the volume unit so that each unit is geometrically represented by a point on a $k$-dimensional hypersphere. The vector $\mathbf{a}_i$ therefore provides firm $i$'s coordinates in the space of common characteristics. We can stack all the coordinate vectors $\mathbf{a}_i$ inside a $k \times n$ matrix that we call $\mathbf{A}$:

$$\mathbf{A} = \begin{bmatrix} \mathbf{a}_1 & \mathbf{a}_2 & \cdots & \mathbf{a}_n \end{bmatrix} = \begin{bmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ a_{k1} & a_{k2} & \cdots & a_{kn} \end{bmatrix}$$

Let $q_i$ be the number of units produced by firm $i$ and consumed by the representative agent, which we write inside the $n$-dimensional vector $\mathbf{q}$:

$$\mathbf{q} = \begin{bmatrix} q_1 & q_2 & \cdots & q_n \end{bmatrix}'$$

A vector $\mathbf{q}$ that specifies, for every firm, the number of units produced is called an *allocation.* I assume that there exists a representative agent. Consistent with the hedonic demand literature, the consumer's preferences are defined in terms of the total units of characteristics, combined linearly from different products. Letting $x_j$ being the total units of characteristic $j$, we have:

$$x_j = \sum_i a_{ji} q_i$$

Hence, geometrically, the matrix $\mathbf{A}$ projects the vector of units of goods purchased $\mathbf{q}$ onto the space of common characteristics:

$$\mathbf{x} = \mathbf{A}\mathbf{q} \tag{1.1}$$

With regard to the $n$ idiosyncratic characteristics, I assume that each unit consumed of good $i$ provides exactly one unit of its corresponding idiosyncratic characteristic, hence we can just write $q_i$ as the units of idiosyncratic characteristic provided the consumption of good $i$.

The representative agent's preferences are described by a utility function that is quadratic in the common characteristics ($\mathbf{x}$) and in terms of the idiosyncratic characteristics (equal to the output vector $\mathbf{q}$); the agent's preferences also incorporate a linear disutility for the total number of hours of work supplied ($H$):

$$U(\mathbf{x}, \mathbf{q}, H) \stackrel{\text{def}}{=} \alpha \cdot \sum_{j=1}^{k} \left( b_j^x x_j - \frac{1}{2} x_j^2 \right) + (1-\alpha) \sum_{i=1}^{n} \left( b_i^q q_i - \frac{1}{2} q_i^2 \right) - H$$

where $b_j^x$ and $b_i^q$ are characteristic-specific preference shifters. In linear algebra notation:

$$U(\mathbf{x},\mathbf{q},H) \stackrel{\text{def}}{=} \alpha\left(\mathbf{x}'\mathbf{b}^x - \frac{1}{2}\cdot\mathbf{x}'\mathbf{x}\right) + (1-\alpha)\left(\mathbf{q}'\mathbf{b}^q - \frac{1}{2}\cdot\mathbf{q}'\mathbf{q}\right) - H \tag{1.2}$$

The parameter $\alpha$ determines the utility weight that is assigned to common characteristics. Hence, $\alpha$ governs the degree of *horizontal differentiation* among products. This utility specification is a generalization of the preferences used by Epple (1987). In addition to introducing idiosyncratic characteristics, I make leisure the outside good: that allows me to close the model and make it general equilibrium.

I denote by $h_i$ the labor input acquired by every firm, so that the labor market clearing condition is:

$$H = \sum_i h_i$$

I assume (without loss of generality) that labor is the numéraire of this economy (the price of one unit of labor is 1\$), therefore $h_i$ is also the total variable cost incurred by firm $i$. Firm $i$ produces output $q_i$ using a quasi-Cobb Douglas production function:

$$q_i = k_i^\theta \cdot \ell(h_i)$$

where $k_i$ is the capital input (fixed) and the function $\ell(\cdot)$ is such that firm $i$'s technology can be described by the following quadratic total variable cost function:

$$h_i = c_i q_i + \frac{\delta_i}{2} q_i^2 \tag{1.3}$$

where $c_i$ and $\delta_i$ depend on $k_i$. MC and AVC denote, respectively, the marginal cost and the average variable cost:

$$\text{MC}_i = c_i + \delta_i q_i; \qquad \text{AVC}_i = c_i + \frac{\delta_i}{2} q_i$$

For some of the empirical analysis, I will later also consider fixed costs ($f_i$). Firm $i$'s total cost function will then become:

$$\text{TC}_i = f_i + c_i q_i + \frac{\delta_i}{2} q_i^2$$

The representative consumer buys the goods bundle $\mathbf{q}$ taking $\mathbf{p}$ (the vector of prices) as given. Moreover, I assume that the representative consumer is endowed with the shares of all the companies in the economy. As a consequence, the aggregate profits are paid back to them. Their consumption basket, defined in terms of the unit purchased $\mathbf{q}$, has to respect the following budget constraint:

$$H+\Pi \;=\; \sum_{i=1}^{k} p_i q_i$$

Notice that for now we have defined aggregate economic profits $\Pi$ to include all non-labor compensation (which equates to assuming that $f_i$ is sunk). We will later consider a narrower metric of profits from which fixed costs $\left(F \overset{\text{def}}{=} \sum_i f_i\right)$ are netted out.

### 1.2.2 Equilibrium

To streamline notation, let us define:

$$b_i \;\overset{\text{def}}{=}\; \alpha \sum_j a^x_{ji} x_j + (1-\alpha)\, b^q_i$$

or, in linear algebra notation:

$$\mathbf{b} \;\overset{\text{def}}{=}\; \alpha\, \mathbf{A}'\mathbf{b}^x + (1-\alpha)\,\mathbf{b}^q \tag{1.4}$$

Then, plugging equation (1.1) and (1.4) inside equation (1.2), we obtain the following Lagrangian for the representative consumer:

$$\mathscr{L}(\mathbf{q},H) \;=\; \mathbf{q}'\mathbf{b} - \frac{1}{2}\mathbf{q}'\left[\mathbf{I}+\alpha\left(\mathbf{A}'\mathbf{A}-\mathbf{I}\right)\right]\mathbf{q} - H - \lambda\left(\mathbf{q}'\mathbf{p} - H - \Pi\right)$$

The choice of labor hours as the numéraire immediately pins down the Lagrange multiplier $\lambda = 1$. Then, the consumer chooses a demand function $\mathbf{q}(\mathbf{p})$ to maximize the following consumer surplus function:

$$S(\mathbf{q}) \;=\; \mathbf{q}'(\mathbf{b}-\mathbf{p}) - \frac{1}{2}\mathbf{q}'\left[\mathbf{I}+\alpha\left(\mathbf{A}'\mathbf{A}-\mathbf{I}\right)\right]\mathbf{q} \tag{1.5}$$

Let us now define the concept of *cosine similarity.* We call the dot product $\mathbf{a}_i'\mathbf{a}_j$ the *cosine similarity* between $i$ and $j$.

The rationale for this nomenclature is that – geometrically – $\mathbf{a}_i'\mathbf{a}_j$ measures the cosine of the angle between vectors $\mathbf{a}_i$ and $\mathbf{a}_j$ in the space of common characteristics $\mathbb{R}^k$. Hence, the cosine similarity ranges from zero to one. Because, by definition:

$$\left(\mathbf{A}'\mathbf{A}\right)_{ij} \;=\; \mathbf{a}_i'\mathbf{a}_j$$

the matrix $\mathbf{A}'\mathbf{A}$ contains the *cosine similarities* between all firm pairs. A higher cosine similarity implies that two products provide a more overlapping mix of characteristics, and this reflects in patterns of product substitution: if $\mathbf{a}_i'\mathbf{a}_j > \mathbf{a}_i'\mathbf{a}_k$, an increase in the supply of product $i$ leads to a larger decline in the marginal utility of product $j$ than it does on the marginal utility of product $k$.

Figure 1.1 helps visualize this setup for the simple case of two firms—1 and 2—competing in the space of two common characteristics A and B. As can be seen in the figure, both firms exist as vectors on the unit circle (with more than three characteristics, it would be a hypersphere instead). The cosine similarity $\mathbf{a}_i'\mathbf{a}_j$ captures the width of the angle $\theta$. An increase in the cosine of the angle $\theta$ (a lower angular distance) implies a lower angular distance, a more overlapping set of common characteristics.

We can streamline the notation further by defining:

$$\boldsymbol{\Sigma} \stackrel{\text{def}}{=} \alpha \left(\mathbf{A}'\mathbf{A} - \mathbf{I}\right)$$

then the demand and inverse demand functions are given by:

$$\text{Aggregate demand}: \quad \mathbf{q} = \left(\mathbf{I} + \boldsymbol{\Sigma}\right)^{-1}\left(\mathbf{b} - \mathbf{p}\right) \tag{1.6}$$

$$\text{Inverse demand}: \quad \mathbf{p} = \mathbf{b} - \left(\mathbf{I} + \boldsymbol{\Sigma}\right)\mathbf{q} \tag{1.7}$$

Notice that the quantity sold by each firm may affect the price of the output sold by every other firm in the economy (unless the matrix $\boldsymbol{\Sigma}$ is null). The derivative $\partial p_i / \partial q_j$ is proportional to $\mathbf{a}_i'\mathbf{a}_j$, the product similarity between $i$ and $j$. The closer these two firms are in the product characteristics space, the larger is this derivative in absolute value. Because $\mathbf{A}'\mathbf{A}$ is symmetric, we have $\partial q_i / \partial p_j = \partial q_j / \partial p_i$ by construction. My rationale for using a linear demand is discussed at length in the Online Appendix.

In terms of elasticities, we have:

$$\text{Inverse cross} - \text{price elasticity of demand}: \quad \frac{\partial \log p_i}{\partial \log q_j} = -\frac{q_j}{p_i} \cdot \sigma_{ij} \qquad \forall\, i \neq j \tag{1.8}$$

$$\text{Cross} - \text{price elasticity of demand}: \quad \frac{\partial \log q_i}{\partial \log p_j} = -\frac{p_j}{q_i} \cdot \left(\mathbf{I} + \boldsymbol{\Sigma}\right)^{-1}_{ij} \tag{1.9}$$

It is worth stopping to inspect equation (1.9) more closely. The first thing can notice is that the cross-price demand elasticities depend on the inverse $\left(\mathbf{I} + \boldsymbol{\Sigma}\right)^{-1}$. This implies that, while cosine similarities are positive by construction, it is entirely possible for goods to be complements. This property of the model is discussed at length in Section 1.5.

Next, let us consider the case $i = j$, where (1.9) simply becomes the own residual demand elasticity. The first major difference between the GHL demand system and CES is that, while in CES the own demand elasticity is equal to a constant, here the own demand elasticity is an equilibrium object (as it depends on $\mathbf{q}$) and will generally differ among firm pairs. This implies that, unlike CES, this demand system produces heterogenous markups. In fact, we can see that two forces drive cross-sectional differences in market power across firms. The more familiar one is the incomplete passthrough from marginal cost to prices: that is, larger firms (high $q_i$) charge higher markups. The second force, which is instead a feature of hedonic demand models, is asymmetric product differentiation. That is, firms that produce "unique" products, as measured by the term $(\mathbf{I} + \boldsymbol{\Sigma})^{-1}_{ii}$, face a less elastic residual demand.

Next, I define the economic profits $\pi_i$ as follows:

$$\begin{aligned}\pi_i(\mathbf{q}) &\overset{\text{def}}{=} p_i(\mathbf{q}) \cdot q_i - h_i \\ &= q_i(b_i - c_i) - \left(1 + \frac{\delta_i}{2}\right) q_i^2 - \sum_{j \neq i} \sigma_{ij} q_i q_j\end{aligned}$$

Firms compete à la Cournot: each firm $i$ strategically chooses its output volume $q_i$ by taking as given the output of all other firms. By taking the profit vector as a payoff function and the vector of quantities produced $\mathbf{q}$ as a strategy profile, I have implicitly defined a *linear-quadratic network game* (Ballester, Calvó-Armengol and Zenou, 2006, henceforth BCZ). The reason is that the matrix $\Sigma$ can be conceptualized as the adjacency matrix of a weighted network: in this specific instance, it is the network of product market rivalry relationships that exists among the firms, based on the substitutability of their products.

Linear-quadratic network games belong to a larger class of games known as "potential games" (Monderer and Shapley, 1996): the key feature of potential games is that they can be described by a scalar function $\Phi(\mathbf{q})$, which we call the game's *potential*. The potential function can be thought of, intuitively, as the objective function of the *pseudo-planner* problem that is solved by the Nash equilibrium allocation. The potential function is shown below, together with the aggregate profit function $\Pi(\mathbf{q})$ and the aggregate welfare function $W(\mathbf{q})$:

$$\begin{aligned} \text{Aggregate Profit}: \quad \Pi(\mathbf{q}) &= \mathbf{q}'(\mathbf{b}-\mathbf{c}) & -\mathbf{q}'\left(\mathbf{I}+\frac{1}{2}\boldsymbol{\Delta}+\boldsymbol{\Sigma}\right)\mathbf{q} \\ \text{Cournot Potential}: \quad \Phi(\mathbf{q}) &= \mathbf{q}'(\mathbf{b}-\mathbf{c}) & -\mathbf{q}'\left(\mathbf{I}+\frac{1}{2}\boldsymbol{\Delta}+\frac{1}{2}\boldsymbol{\Sigma}\right)\mathbf{q} \\ \text{Total Surplus}: \quad W(\mathbf{q}) &= \mathbf{q}'(\mathbf{b}-\mathbf{c}) & -\frac{1}{2}\cdot\mathbf{q}'(\mathbf{I}+\boldsymbol{\Delta}+\boldsymbol{\Sigma})\mathbf{q} \end{aligned} \tag{1.10}$$

$$\text{where} \qquad \boldsymbol{\Delta} \stackrel{\text{def}}{=} \begin{bmatrix} \delta_1 & 0 & \cdots & 0 \\ 0 & \delta_2 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & \delta_n \end{bmatrix}$$

The three functions in equation (1.10) are visually similar to each other; they only differ by the scalar weight applied to the quadratic terms. The Cournot potential $\Phi$ is somewhat of a hybrid between the aggregate profit $\Pi$ and the total surplus $W$: the diagonal elements of the quadratic term are the same as the aggregate profit function, while the off-diagonal terms are the same as the aggregate surplus function. By maximizing the potential $\Phi(\mathbf{q})$, we find the Cournot-Nash equilibrium. I shall assume all these three functions are concave. Because the oligopolists in this model will be actual firms in the data (who produce positive output by definition) we can look directly at the unique internal solution. The Cournot-Nash equilibrium of the game described above is $\mathbf{q}^{\Phi}$ – the maximizer of the potential function $\Phi(\cdot)$:

$$\mathbf{q}^{\Phi} \stackrel{\text{def}}{=} \arg\max_{\mathbf{q}} \Phi(\mathbf{q}) = (2\mathbf{I}+\boldsymbol{\Delta}+\boldsymbol{\Sigma})^{-1}(\mathbf{b}-\mathbf{c}) \tag{1.11}$$

The derivation of the potential function, as well as the proof that its maximizer $\mathbf{q}^{\Phi}$ is the genuine Nash equilibrium, appear in the Online Appendix. Equation (1.11), which characterizes the Cournot-Nash equilibrium, tells us which factors determine the size of each firm in equilibrium. The diagonal matrix $\boldsymbol{\Delta}$, which contains the slopes of the marginal cost functions, captures economies of scale. $\boldsymbol{\Sigma}$ is the adjacency matrix of the network of product rivalries. $\mathbf{b}$ and $\mathbf{c}$ are, respectively, the demand and supply function intercepts. Hence, $(b_i - c_i)$ is simply the marginal surplus of the very first unit produced by firm $i$; also, $b_i$ can be interpreted as a measure of vertical product differentiation (quality).

BCZ show that another way to interpret equation 1.11 is as a measure of *network centrality* – specifically, that developed by Katz (1953) and Bonacich (1987). The intuition is that firms that are more "isolated" in the network of product similarities face less product market competition and behave more as monopolists. These centrality measures are a recurring feature of the literature on network in macroeconomics (see Carvalho, Nirei, Saito and Tahbaz-Salehi, 2020; Carvalho and Tahbaz-Salehi, 2019). I discuss this Nash-Bonacich linkage more in detail in the Online Appendix.

The discrepancy between the potential function and the total-surplus function implies that the network Cournot game delivers an equilibrium allocation that is not socially-optimal. A benevolent social planner can theoretically improve on the market outcome for two reasons. First, they can coordinate output choices across firms; second, they can internalize consumer surplus.

### 1.2.3 Separability of Consumer Surplus

Next, I investigate the problem of how to measure surplus appropriation in this model. I show that under GHL demand, consumer surplus has a desirable property, which I call *additive separability*: it means we can attribute a certain share of the consumer surplus to each firm. I will use this separability property to propose a measure of oligopoly power that varies by firm, and which I will be able to link to surplus appropriation. Assume that the allocation $\mathbf{q}$ maximizes the consumer utility given the price vector $\mathbf{p}$. We say that the consumer surplus $S(\mathbf{q})$ is *additively separable* if it can be written as the sum over the set of firms of some function $s(b_i, q_i, p_i)$ that only depends on the triple $(b_i, q_i, p_i)$. That is:

$$S(\mathbf{q}) = \sum_i s(b_i, q_i, p_i)$$

The consumer surplus function $S(\mathbf{q})$ from equation *(1.5)* is additively separable. Noting that the inverse demand function can be rearranged as $(\mathbf{I} + \mathbf{\Sigma})\mathbf{q} = \mathbf{b} - \mathbf{p}$, we can write equation (1.5) as:

$$\begin{aligned} S(\mathbf{q}) &= \mathbf{q}'(\mathbf{b} - \mathbf{p}) - \frac{1}{2}\mathbf{q}'(\mathbf{I} + \mathbf{\Sigma})\mathbf{q} \\ &= \mathbf{q}'(\mathbf{b} - \mathbf{p}) - \frac{1}{2}\mathbf{q}'(\mathbf{b} - \mathbf{p}) = \frac{1}{2}\mathbf{q}'(\mathbf{b} - \mathbf{p}) \end{aligned} \tag{1.12}$$

the last term can be rewritten in summation form as $\sum_i \frac{1}{2} q_i (b_i - p_i)$. By substituting $p_i$ for the inverse demand function, we obtain the firm-level consumer surplus, which attributes to each firm $i$ a certain share $s_i$ of the consumer surplus $S(\mathbf{q})$:

$$s_i(\mathbf{q}) \stackrel{\text{def}}{=} \frac{1}{2}\left(q_i^2 + \sum_{j \neq i} \sigma_{ij} q_i q_j\right)$$

We can then also define a firm-level total surplus function, which specifies for every firm $i$ a certain share $w_i$ of the total surplus $W(\mathbf{q})$:

$$
\begin{aligned}
w_i(\mathbf{q}) &\stackrel{\text{def}}{=} \pi_i(\mathbf{q}) + s_i(\mathbf{q}) \\
&= q_i(b_i - c_i) - \frac{1}{2}\left[(1+\delta_i)\, q_i^2 + \sum_{j \neq i} \sigma_{ij} q_i q_j\right]
\end{aligned} \tag{1.13}
$$

### 1.2.4 Oligopoly Power and Surplus Appropriation at the Firm-Level

The canonical oligopoly model with perfectly-substitutable products establishes the Herfindahl-Hirschmann Index (HHI) as a measure of market power. The reason for that is that the HHI relates the (market share-) weighted average firm-level inverse demand elasticity to the industrywide inverse demand elasticity. Let $Q = \sum_i q_i$. Then:

$$
\begin{aligned}
\frac{\partial \log p}{\partial \log q_i} &= \frac{\partial \log p}{\partial \log Q} \cdot \frac{q_i}{Q} \\
\frac{q_i}{Q} \cdot \frac{\partial \log p}{\partial \log q_i} &= \frac{\partial \log p}{\partial \log Q} \cdot \left(\frac{q_i}{Q}\right)^2 \\
\sum_i \frac{q_i}{Q} \cdot \frac{\partial \log p}{\partial \log q_i} &= \frac{\partial \log p}{\partial \log Q} \cdot \mathrm{HHI}
\end{aligned} \tag{1.14}
$$

where $\mathrm{HHI} = \sum\limits_i \left(\frac{q_i}{Q}\right)^2$ is the Herfindahl-Hirschmann Index. The lemma above illustrates the fact that the reason why the HHI is informative about the residual demand elasticity is that the individual *market shares* are themselves informative about the demand elasticity of firm $i$—this fact is frequently overlooked or forgotten. The first line of equation (1.14) evinces this: the ratio of the inverse demand elasticities for firm $i$ and the industry as a whole is simply the market share of firm $i$. Hence, if we wanted to derive a firm-level counterpart of the HHI index, it would simply be the market share of firm $i$.

Let us now return to the network Cournot model and define the following statistic for firm $i$. I define $\omega_i$, the *weighted market share* of firm $i$, as follows:

$$
\omega_i \stackrel{\text{def}}{=} \frac{q_i}{q_i + \sum_j \sigma_{ij} q_j}
$$

Notice that, under homogenous products ($\sigma_{ij} = 1\, \forall\, i, j$) this is simply the market share of firm $i$ . It is possible to show that the ratio of firm profits $\pi_i$ to consumer surplus $s_i$ proportional to the similarity-weighted market share $\omega_i$ of firm $i$. In the Cournot-Nash equilibrium allocation, the ratio of profits to consumer surplus for firm i is proportional to its weighted market share - specifically:

$$\frac{\pi_i}{s_i} \quad = \quad (2+\delta_i)\,\omega_i \tag{1.15}$$

See the Online Appendix. Therefore, in the Network Cournot model, the similarity-weighted market share $\omega_i$ replaces the HHI as a firm-level measure of market power that accounts for product differentiation. As is the case for the HHI, the similarity-weighted market share is an equilibrium object—an endogenous outcome of the Cournot game played by the oligopolists.

The identity in Lemma (1.2.4) reflects the fact that, in my model, there are no clearly-defined industry boundaries. This is also the case in the real world: if we consider antitrust lawsuits for example, a major object of litigation is the market's definition. Defendants (alleged monopolies) have an incentive to define the relevant market broadly, while plaintiffs have an incentive to define the relevant market narrowly.

In my model, firms exist in a continuous space of product characteristics. Hence, there is no uniquely-defined peer group that we can compare each firm to. To understand how dominant firm $i$ is, we need to compare its market share vis-à-vis every other firm in the economy, weighting each of them by their distance in the product space.

The Herfindahl Index can be seen as a special case of the weighted market share: as a measure of surplus appropriation, it is only valid in the special case where similarity scores are dichotomous (implying sharp industry boundaries) and firms are exchangeable ($b_i - c_i$ is constant across firms).

### 1.2.5 Market Structure Counterfactuals

A key application of my theoretical model is to study how welfare statistics - such as total surplus - respond to changes in market structure. What that means in practice is that, having made the required assumption about what are the rules of the game played by the firms (we have assumed a Cournot equilibrium), we can then consider counterfactuals in which the same firms play by a different set of rules. In this sub-section, I define four of these counterfactuals: each of these counterfactuals coincides the solution to a specific maximization problem.[5]

[5] The closed-form expressions for the output vector $\mathbf{q}$ which I provide below assume an internal solution. For my empirical analysis, I also compute a numerical solution that is subject to a non-negativity constraint on $\mathbf{q}$ and I verify it is approximately equal to the unconstrained solution (error $< 0.1\%$ for the total surplus function in Perfect Competition). The non-negativity constraint binds for very few firms.

The first counterfactual that I consider is *perfect competition*: firms act as atomistic producers, and price all units sold at marginal cost. The *Perfect Competition* allocation $\mathbf{q}^W$ is defined as the maximizer of the aggregate total surplus function $W(\mathbf{q})$ :

$$\mathbf{q}^W \overset{\text{def}}{=} \arg\max_{\mathbf{q}} W(\mathbf{q}) = (\mathbf{I} + \mathbf{\Delta} + \mathbf{\Sigma})^{-1}(\mathbf{b} - \mathbf{c}) \tag{1.16}$$

The second counterfactual that I consider is called *Monopoly*: it represents a situation in which one agent (that does not internalize consumer surplus) has control over all the firms in the economy and maximizes aggregate profits. The *Monopoly* allocation is defined as the maximizer of the aggregate profit function $\Pi(\mathbf{q})$:

$$\mathbf{q}^\Pi \overset{\text{def}}{=} \arg\max_{\mathbf{q}} \Pi(\mathbf{q}) = (2\mathbf{I} + \mathbf{\Delta} + 2\mathbf{\Sigma})^{-1}(\mathbf{b} - \mathbf{c}) \tag{1.17}$$

This allocation can be alternatively conceptualized as an economy with no antitrust policy, where firms have unlimited ability to coordinate their supply choices.

While the *Monopoly* counterfactual is an interesting limit case, using the model we can also study the welfare impact of mergers and collusion among specific firms.

When it comes to modeling mergers and collusions, the I.O. literature has used multiple approaches. Following Baker and Bresnahan (1985), I choose model mergers and collusion interchangeably as *coordinated pricing.* That is, I assume that the merger or the collusion does not affect the product range offered by the merging/colluding enterprises; instead, a single agent determines the output of the merging firms to maximize the joint profits.[6] Consider, without loss of generality, a merger or collusion between companies $\{1, 2, ..., m\}$; then, partition the matrix $\mathbf{\Sigma}$ by separating the first m rows and columns as follows:

$$\mathbf{\Sigma} = \begin{bmatrix} \mathbf{\Sigma}_{11} & \mathbf{\Sigma}_{12} \\ \mathbf{\Sigma}_{21} & \mathbf{\Sigma}_{22} \end{bmatrix}$$

The post-merger equilibrium allocation maximizes the following modified potential function:

$$\Phi(\mathbf{q}) = \mathbf{q}'(\mathbf{b} - \mathbf{c}) - \mathbf{q}'\left(\mathbf{I} + \frac{1}{2}\mathbf{\Delta}\right)\mathbf{q} - \frac{1}{2} \cdot \mathbf{q}' \begin{bmatrix} 2\mathbf{\Sigma}_{11} & \mathbf{\Sigma}_{12} \\ \mathbf{\Sigma}_{21} & \mathbf{\Sigma}_{22} \end{bmatrix} \mathbf{q} \tag{1.18}$$

See the Online Appendix. The maximizer of the re-defined $\Phi(\mathbf{q})$, which corresponds

[6] This approach is particularly convenient in this setting, where products are highly differentiated. The alternative to following this approach would be to make heroic assumptions about the nature and characteristics of the hypothetical product produced by the combined entity.

to the post-merger equilibrium allocation, is:

$$\mathbf{q}^{\Phi} = \left(2\mathbf{I} + \boldsymbol{\Delta} + \begin{bmatrix} 2\boldsymbol{\Sigma}_{11} & \boldsymbol{\Sigma}_{12} \\ \boldsymbol{\Sigma}_{21} & \boldsymbol{\Sigma}_{22} \end{bmatrix}\right)^{-1} (\mathbf{b} - \mathbf{c})$$

That is, to simulate the new equilibrium following a merger or a collusion among existing firms, one only needs to amend the potential function by doubling the off-diagonal quadratic terms corresponding to the merging firms. It is easily verified that when all firms are merged, $\Phi(\mathbf{q})$ simply becomes the aggregate profit function $\Pi(\mathbf{q})$, and the equilibrium allocation converges to the Monopoly counterfactual (equation 1.17).

Another interesting counterfactual is one in which resources are allocated efficiently but the labor supply is fixed. That is, the social planner maximizes the aggregate surplus function subject to the constraint of using no more labor than in the observed Cournot equilibrium. I define the resource-efficient counterfactual $\mathbf{q}^H$ as the solution to the following constrained maximization problem:

$$\mathbf{q}^H \stackrel{\text{def}}{=} \arg\max_{\mathbf{q}} W(\mathbf{q}) \quad \text{s.t.} \quad H(\mathbf{q}) = H\left(\mathbf{q}^{\Phi}\right)$$

Setting up the Lagrangian and using $(1-\mu)$ as the Lagrange multiplier, we find that the resource-efficient counterfactual takes the form:

$$\mathbf{q}^H = (\mathbf{I} + \mu\boldsymbol{\Delta} + \boldsymbol{\Sigma})^{-1} (\mathbf{b} - \mu\mathbf{c}) \tag{1.19}$$

where $\mu$ solves:

$$H\left(\mathbf{q}^H(\mu)\right) = H\left(\mathbf{q}^{\Phi}\right)$$

The Lagrange multiplier term $\mu$ turns out to be the common markup charged by all firms in the resource-efficient counterfactual. The Resource-efficient counterfactual $\mathbf{q}^H$ equalizes markups across firms. Let all firms price at a constant markup $\mu$ over marginal cost:

$$p_i = \mu \cdot \mathrm{MC}_i$$

expanding the expression for the marginal cost and the equilibrium price we have:

$$\mathbf{b} - (\mathbf{I} + \boldsymbol{\Sigma})\,\mathbf{q} = \mu\,(\mathbf{c} + \boldsymbol{\Delta}\mathbf{q})$$

rearranging the equation above we obtain (1.19). Because this counterfactual uses the same amount of labor as the observed equilibrium, by comparing welfare in this allocation to the first-best we can effectively disentangle the welfare costs of monopoly into two components: misallocation and factor-suppression. We can also interpret this counterfactual as the deadweight loss in an alternative model where the supply of labor is completely inelastic. Notice that when this allocation is not constrained by the labor supply (the Lagrange multiplier $1-\mu$ is zero), the common markup is one (firms price at marginal cost) and the resource-efficient allocation coincides with perfect competition.

The counterfactuals considered thus far do not account for how a firm's incentives to participate in the market are affected by the intensity of competition. When the market moves from Cournot competition to (say) Bertrand [7] or perfect competition, the resulting lower profits might insufficient to cover fixed costs, and therefore too low to justify a firm's continued existence. If this is the case, perfect competition may not be a realistic benchmark in the long-run: this is the classical criticism of static welfare analysis.

Next, I construct an "efficient" allocation that takes into account (to the extent possible in a static model) these dynamic incentives. The starting point is again a benevolent social planner, to which we are adding a constraint, in the form of a participation condition on the firms' side: firms have to be able (on average) to recover their fixed costs ($F$) at the optimum.[8] The Second-Best Allocation $\mathbf{q}^{2nd}$ is defined as the solution to the following constrained maximization problem:

$$\mathbf{q}^{2nd} \quad \stackrel{\text{def}}{=} \quad \arg\max_{\mathbf{q}} W(\mathbf{q}) \quad \text{s.t.} \quad \Pi(\mathbf{q}) \geq F$$

where

$$F \stackrel{\text{def}}{=} \sum_i f_i$$

Setting up the Lagrangian of this problem and imposing $\lambda$ as the Lagrangian multiplier, we find that the resource-efficient counterfactual takes the form:

$$\mathbf{q}^{2nd} \quad = \quad \left[\frac{1+2\lambda}{1+\lambda}\cdot(\mathbf{I}+\mathbf{\Sigma})+\mathbf{\Delta}\right]^{-1}(\mathbf{b}-\mathbf{c})$$

Assuming that the constraint binds at the optimum, the Lagrange multiplier $\lambda$ solves:

$$\Pi\left(\mathbf{q}^{2nd}(\lambda)\right) \;=\; F$$

[7] The Bertrand model is covered in the Online Appendix.

[8] There are two reasons why I consider a constraint on aggregate profits rather than individual profits ($\pi_i \geq f_i$). The first is that such individual constraint is already violated by many firms in the observed (Cournot) equilibrium. The reason is that a individual constraint would make the optimization problem numerically intractable, since we would need to solve for thousands of constraints (one for each firm in the model).

As the constraint is relaxed ($\lambda \to 0$), this counterfactual allocation converges to the first-best. When the constraint becomes arbitrarily tight ($\lambda \to \infty$), it converges to the *Monopoly* allocation. In addition to the counterfactuals considered above, which admit closed-form solutions, we can simulate the introduction or the removal of granular firms. The latter can be trivially implemented by computing an allocation where a firm's output is constrained to be zero. In order to simulated instead the introduction of new firms, we require additional assumptions or data. Namely, in order to simulate the introduction of an additional firm (let us label it firm zero), we would need to know the value of $(b_0 - c_0)$, as well as the firm's similarity to every other firm in the economy ($\mathbf{a}_{i0}$). One such counterfactual is considered in Section (1.4).

### 1.2.6 Adding a Continuum of firms with Endogenous Entry

Next, I show how to expand the model to include a continuum of atomistic firms that behave competitively and can enter and exit exogenously. This is an important extension of the model for two reasons. First, it allows to incorporate firms for whom we do not observe product similarity data (foreign and private firms). Second, it allows to incorporate entry and exit in an otherwise static model. The idea is that we can model these unobserved companies as atomistic firms.

The key to tractably integrating these atomistic firms is an aggregation result. I describe these atomistic firms through a productivity distribution: the set of active firms will be characterized by a productivity cut-off value, in the style of Hopenhayn (1992).

Next, I show that these atomistic companies can be aggregated into a representative firm: variation in the size of such representative firm reflect both the intensive and the extensive (entry/exit) margin of production of the atomistic firms. I index this representative firm as $i = n + 1$, effectively adding a row and a column to the matrices $\mathbf{A}'\mathbf{A}$ and $\mathbf{\Delta}$ and adding one dimension to the vector $\mathbf{b}$. Assume that there is a mass one of potential entrants that are indexed by a productivity parameter $z \in (\underline{z}, \infty)$ and that produce a homogeneous good using the following quadratic cost function:

$$h(z) = \frac{\delta(z)}{2} \cdot q^2(z)$$

with $\underline{z} > 0$ and

$$\delta(z) = \frac{1}{z}$$

Assume also that the firms face cost of entry equal to one unit of labor (without loss of generality) and that the probability density of type-$z$ potential entrants is given by

$$f(z) = \frac{\beta - 1}{z^{\beta+1}}$$

implying that z follows a Pareto distribution with shape parameter $\beta$ and scale parameter $\underline{z} \stackrel{\text{def}}{=} [(\beta - 1)/\beta]^{\frac{1}{\beta}}$.[9] Then, as the parameter $\beta$ converges down to 1, the cost function of the corresponding aggregate representative firm is approximated by

$$h_{n+1} = \frac{q_{n+1}^2}{2}$$

where and $h_{n+1}$ and $q_{n+1}$ are, respectively, the labor input and the output of the representative firm, and the productivity cutoff for entry converges to $z_{\min} = \frac{1}{q_{n+1}}$. Because employment and revenues are proportional to $z$, it follows that, if the assumptions above are respected, both the revenues and employments distribution of firms must approximate a Pareto distribution with shape parameter $\beta = 1$, sometimes called a Zipf Law.

Although this might look like a knife-edge assumption, it is not. It is a well-documented empirical regularity that the size distribution of firms closely approximates a Pareto distribution with shape parameter $\beta = 1$. This stylized fact was confirmed to hold for both the employment *and* the revenue distribution of US firms by Axtell (2001), using Census micro-data.

Because the representative firm behaves competitively, its first order condition will differ from those of granular firms $\{1, 2, ..., n\}$. The latter maximize individual profits:

$$\pi_i'(q_i) = 0 \qquad \text{for} \quad i = 1, 2, ..., n$$

The representative firm, on the other hand, prices at marginal cost, and therefore it maximizes total surplus:

$$W'(q_i) = 0 \qquad \text{for} \quad i = n + 1$$

We can write the full system of first order conditions in linear algebra notation as:

$$0 = \begin{bmatrix} \mathbf{b}^{(n)} - \mathbf{c}^{(n)} \\ b_{n+1} - c_{n+1} \end{bmatrix} + \left( \begin{bmatrix} 2\mathbf{I} & \mathbf{0} \\ \mathbf{0} & 1 \end{bmatrix} + \mathbf{\Sigma} + \mathbf{\Delta} \right) \begin{bmatrix} \mathbf{q}^{(n)} \\ q_{n+1} \end{bmatrix}$$

[9] While revenues and employees for US firms follow approximately a Zipf Law, this distribution has the undesirable property that its mean (and therefore $q_{n+1}$ and $h_{n+1}$) grows unboundedly as $\beta \to 1^+$. This particular choice of the scale parameter ensures that $q_{n+1}$ and $h_{n+1}$ integrate to a finite number as $\beta \to 1^+$.

where $c_{n+1} = 0$, $\delta_{n+1} = 1$ and the superscript $(n)$ identifies the sub-vector corresponding to the granular firms. A simpler way to rewrite this set of equations is

$$0 = \mathbf{b} - \mathbf{c} - (\mathbf{I} + \mathbf{G} + \mathbf{\Sigma} + \mathbf{\Delta})\,\mathbf{q}$$

where $\mathbf{G}$ is a diagonal matrix that identifies granular firms – that is, whose diagonal elements equal 1 for firms 1 to $n$ and to 0 for firm $n+1$:

$$\mathbf{G} = \begin{bmatrix} 1 & 0 & \cdots & 0 & 0 \\ 0 & 1 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \cdots & 1 & 0 \\ 0 & 0 & \cdots & 0 & 0 \end{bmatrix}$$

The potential function for the model that includes the representative firm is:

$$\Phi(\mathbf{q}) = \mathbf{Q}(\mathbf{b} - \mathbf{c}) - \frac{1}{2}\mathbf{q}'(\mathbf{I} + \mathbf{G} + \mathbf{\Sigma} + \mathbf{\Delta})\,\mathbf{q}$$

and the equilibrium quantity vector is:

$$\mathbf{q}^{\Phi} = (\mathbf{I} + \mathbf{G} + \mathbf{\Delta} + \mathbf{\Sigma})^{-1}(\mathbf{b} - \mathbf{c})$$

## 1.3 Data and Identification/Calibration

In this section, I outline the data used to estimate the model in Section 1.2. Additional details are provided in the Online Appendix, which also contains a table which summarizes the model mapping and identification.

### 1.3.1 Firm Financials

My data source for firm financials is the Compustat database, which I access via the Wharton Research Data Services (WRDS) platform. From this database, I extract information on firm revenues, Costs of Goods Sold (COGS), Selling General and Administrative (SGA) costs, R&D expenditures and Property Plant and Equipment (PPE).

I follow (De Loecker, Eeckhout and Unger, 2020, henceforth DEU) in mapping accounting revenues to model revenues, COGS to variable costs, and in computing an estimate of fixed costs costs ($f_i$):

$$f_i \leftarrow \text{SGA}_i + \text{Property Plant \& Equipment}_i \times \text{User Cost of Capital} \tag{1.20}$$

### 1.3.2 Text-Based Product Similarity

The key data input required to estimate the model presented in Section 1.2 is the matrix of product similarities $\mathbf{A}'\mathbf{A}$. The empirical counterpart to this object is provided by Hoberg and Phillips (2016, henceforth HP).

HP created a publicly-available database that provides product cosine similarities for the universe of public corporations in the United States. This dataset originates from a computational linguistics analysis of regulatory forms 10-K. HP's cosine similarity scores are time-varying. A complete matrix of similarity scores (one score for every pair of public firms) is provided for every year, beginning in 1997.

The 10-K is a mandatory form that is filed by American public corporations with the U.S. Securities and Exchange Commission on a yearly basis. Item 1 of the 10-K contains a long and detailed description of the product or service sold by the company. HP's product cosine similarities are constructed by comparing these textual product descriptions.

I briefly outline the construction of this dataset. HP start by building a vocabulary of 61,146 words that firms use to describe the characteristics of their products.[10] Based on this vocabulary, HP produce, for each firm $i$, a vector of word frequencies $\mathbf{o}_i$. Each of component of this vector is equal to the number of occurrences of a specific word, from HP's vocabulary, inside firm $i$'s 10-K product description:

$$\mathbf{o}_i = \begin{bmatrix} o_{i,1} \\ o_{i,2} \\ \vdots \\ o_{i,61146} \end{bmatrix}$$

Similar to the model in Section 1.2, this vector is then normalized (divided by the

[10] I report here verbatim the methodology description from the original paper by Hoberg and Phillips (2016):"[...] *In our main specification, we limit attention to nouns (defined by Webster.com) and proper nouns that appear in no more than 25 percent of all product descriptions in order to avoid common words. We define proper nouns as words that appear with the first letter capitalized at least 90 percent of the time in our sample of 10-Ks. We also omit common words that are used by more than 25 percent of all firms, and we omit geographical words including country and state names, as well as the names of the top 50 cities in the United States and in the world.* [...]"

Euclidean norm). We have thus obtained the empirical counterpart of $\mathbf{a}_i$:

$$\mathbf{a}_i = \frac{\mathbf{o}_i}{\|\mathbf{o}_i\|}$$

finally, all $\mathbf{a}_i$ vectors are dot-multiplied to obtain $\mathbf{A}'\mathbf{A}$:

$$\mathbf{A}'\mathbf{A} = \begin{bmatrix} \mathbf{a}_1'\mathbf{a}_1 & \mathbf{a}_1'\mathbf{a}_2 & \cdots & \mathbf{a}_1'\mathbf{a}_n \\ \mathbf{a}_2'\mathbf{a}_1 & \mathbf{a}_2'\mathbf{a}_2 & \cdots & \mathbf{a}_2'\mathbf{a}_n \\ \vdots & \vdots & \ddots & \vdots \\ \mathbf{a}_n'\mathbf{a}_1 & \mathbf{a}_n'\mathbf{a}_2 & \cdots & \mathbf{a}_n'\mathbf{a}_n \end{bmatrix}$$

Hence, to the extent that the word frequencies are a good proxy for product characteristics, the resulting matrix is the exact empirical counterpart to $\mathbf{A}'\mathbf{A}$ — the matrix of cross-price effects in my theoretical model. The fact that all publicly traded firms in the United States are required to file a 10-K form makes the HP data set unique in that it covers the near entirety (97.8%) of the Compustat universe.

HP use these cosine similarity scores to produce a dynamic industry classification, called TNIC, which they extensively validate: one way they do so, in the published paper that outlines their methodology, is by exploiting another dataset - CapitalIQ. This dataset provides dummy variables for a sub-set of Compustat firm pairs which identify product market rivalry relationships; they are based on most recent corporate filings as well as other sources (no time variation is available in this dataset). HP show that TNIC outperforms SIC and NAICS in predicting competitor pairs in CapitalIQ. I refer the reader to HP's original paper for more information.

Since their introduction in 2011, HP's industry classifications have been swiftly adopted across the empirical corporate finance literature. Indeed, for a variety of applications, using traditional industry classifications is no longer considered standard. A major reason for this shift is that, while industry classifications have been traditionally used (for lack of better alternatives) to capture product market competition, it is well-known that they are based on the concept of *production process* similarity, not product similarity.[11].

Consistently with this definition, the I.O. and Antitrust literature have limited the use of industry classifications to the estimation of production functions,[12], and have instead embraced hedonic models for the estimation of demand systems.

One reason why this shift has not occurred in the macroeconomics literature is that

[11] See the following Bureau of Labor Statistic Guide

[12] DEU's method to estimate markups, for example, uses production function estimates for NAICS industries.

estimating the hedonic demand systems that are standard in the EIO literature requires data on prices and physical quantities that is not available in a macroeconomic setting. Moreover, such demand systems can become computationally intractable when the number of firms becomes large. A key methodological contribution of this paper is to provide a new demand system (the GHL demand system) that is highly tractable, scalable and that can be estimated using revenue and cost data for the entire Compustat universe, leveraging HP's publicly-available product similarity data.

There are other factors that differentiate HP's database from traditional industry classifications. While NAICS and SIC are binary (firms are either in the same industry or different industries), HP's database also provides continuous similarity scores ranging from zero to one, thus accommodating the inherent fuzziness of product market rivalries. While NAICS and SIC are seldom updated, HP's similarity scores are updated yearly. While NAICS and SIC are arbitrarily assigned (Chen et al., 2016 show that firms strategically manipulate their industry classifications), HP's similarity scores are rule-driven and incentive-compatible: executives face civil and criminal liability for misrepresenting company information in SEC filings.

Though there are other data sets have a network structure that are related to product market similarity[13], they all have the following shortcomings; (a) they are either directly or indirectly based on industry classifications; and (b) they lack sufficient coverage of the Compustat universe.

I begin my empirical analysis by producing a bi-dimensional visualization of HP's dataset. The challenge in doing so is that each firm exists in a space of characteristics that has as many dimensions as there are words in the vocabulary that HP used to create their similarity data set ($\sim 61,000$). To create a bidimensional visualization of the product space, I use the algorithm of Fruchterman and Reingold (1991, henceforth FR), which is widely used in network science to visualize weighted networks[14].

The result of this exercise is Figure 1.2: every dot in the graph is a publicly traded firm as of 2004. Firm pairs that have a high cosine similarity score appear closer; they are also joined by a thicker line. Conversely, firms that are more dissimilar are not joined, and tend to be more distant. The product space is manifestly uneven: some areas are

[13] Bloom, Schankerman and Van Reenen (2013), for example, use Compustat segment data.

[14] The algorithm models the network edges as particles, letting the nodes dynamically arrange themselves on a bidimensional surface as if they were particles subject to attractive and repulsive forces. One known shortcoming of this algorithm is that it is sensitive to the initial configurations of the nodes, and it can have a hard time uncovering the cluster structure of large networks. To mitigate this problem, and to make sure that the cluster structure of the network is displayed correctly, before running FR I prearrange the nodes using a different algorithm, OpenOrd, (Martin et al., 2011) which was explicitly developed for this purpose.

significantly more densely populated with firms than others. Also, the network displays a pronounced community structure: large groups of firms tend to cluster in certain areas of the network.

In the Online Appendix, I show that the visualization is not an artifact of dimensionality reduction or measurement error: notwithstanding the dimensionality reduction, a remarkable degree of overlap exists between the macro-clusters of this network and broad economic sectors. In addition, this exercise allows me to independently validate HP's product similarity data.

### 1.3.3 Identification of Output, Prices and Cost Intercept

All of the unobserved variables in the model are identified subject to two parameters: $\alpha$, which controls the degree of horizontal differentiation (and therefore the elasticity of substitution) between goods, and the matrix $\mathbf{\Delta}$, which controls returns to scale. I will first show how all model objects are identify conditional on these two parameters and then show how to calibrate these two parameters using external data.

I estimate the matrix $\mathbf{A}'\mathbf{A}$ using Hoberg and Phillips (2016)'s cosine similarity data. This in turns (conditional on $\alpha$) provides the matrix $\mathbf{\Sigma}$. We can then identify real output $q_i$ from revenues and total variable cost data as:

$$q_i \;=\; \sqrt{\frac{\pi_i}{1+\delta_i/2}} \qquad \text{if} \quad i \le n \tag{1.21}$$

If the model includes an aggregate competitive firm, the identification of output for the aggregate firm $n+1$ will be different. Specifically, the marginal cost pricing condition ($p_i = \mathrm{MC}_i$) implies that:

$$q_i \;=\; \sqrt{\pi_i + h_i} \qquad \text{if} \quad i = n+1 \tag{1.22}$$

Where $(\pi_{n+1} + h_{n+1})$ is measured as the Gross Value Added of private and foreign firms, which I compute using the OECD Trade in Value Added (TiVA) Dataset. After identifying $q$, we can then pin down the vector of prices and and the cost function intercepts:

$$p_i \;=\; \frac{p_i q_i}{q_i} \qquad c_i \;=\; \frac{h_i}{q_i} - \frac{\delta_i}{2} q_i$$

Finally, I identify the demand intercept $b_i$ using equation (1.11):

$$\mathbf{b} \;=\; (2\mathbf{I} + \mathbf{\Delta} + \mathbf{\Sigma})\,\mathbf{q} + \mathbf{c}$$

or, in the presence of an aggregate competitive firm:

$$\mathbf{b} \;=\; (\mathbf{I}+\mathbf{G}+\boldsymbol{\Delta}+\boldsymbol{\Sigma})\,\mathbf{q}+\mathbf{c}$$

### 1.3.4 Calibration of $\alpha$ and $\boldsymbol{\Delta}$

The last step that is required in order to take the model to the data is to calibrate the scalar $\alpha$ and the diagonal matrix $\boldsymbol{\Delta}$. Let us start from the latter. To calibrate each diagonal element $\delta_i$, we use the fact that the markup (price-marginal cost ratio) of firm $i$ can be written as a function of observables (revenues, total variable costs) and $\delta_i$. Hence, the markup $\mu_i$ is identified given $\delta_i$. The markup $\mu_i$ is equal to:

$$\mu_i \;\stackrel{\text{def}}{=}\; \frac{p_i}{\mathrm{MC}_i} \;=\; \frac{(2+\delta_i)\cdot p_i q_i}{2\cdot h_i+\delta_i\cdot p_i q_i} \tag{1.23}$$

See the Online Appendix. DEU compute the revenue-weighted average markup for the same universe of companies. My strategy for calibrating $\boldsymbol{\Delta}$ is to target DEU's revenue-weighted average markup. The detailed methodology for calibrating $\boldsymbol{\Delta}$ is outlined in detail in the Online Appendix.

To calibrate $\alpha$, we rewrite equation (1.6) as:

$$\left|\frac{\partial \log p_i}{\partial \log q_j}\right| \;=\; \alpha\cdot \mathbf{a}_i'\mathbf{a}_j\,\frac{q_j}{p_i} \quad \forall\; i\neq j \tag{1.24}$$

By calibrating $\boldsymbol{\Delta}$, we have already pinned $q_i$ and $p_i$. The matrix of product similarity of Hoberg and Phillips (2016) provides the empirical counterpart to $\mathbf{A}'\mathbf{A}$. Hence, the matrix of equilibrium cross-price demand elasticities is identified given $\alpha$.

My strategy for calibrating $\alpha$ is to target target microeconometric estimates from the Industrial Organization literature. I obtain, for a number of firm pairs, estimates of the cross-price demand elasticity from empirical IO studies that estimate the demand function econometrically. These estimates of the cross-price demand elasticity are then manually matched to to the corresponding firm pair in Compustat. Finally, for each firm pair, I can obtain an estimate of $\alpha$ by rearranging equation (1.24):

$$\hat{\alpha}_{ij} \;=\; \left|\frac{\partial \log p_i}{\partial \log q_j}\right| \Big/ \left(\mathbf{a}_i'\mathbf{a}_j\,\frac{q_i}{p_i}\right) \tag{1.25}$$

In the absence of mis-specification and measurement error, all these estimates $\hat{\alpha}_{ij}$ would yield the same value. Instead, what I obtain in this case is a range of estimates. I calibrate $\alpha$ to the median value among these estimates, which is 0.05. The full methodology is presented in the Online Appendix, where I also discuss how the model fits non-targeted moments in the data.

## 1.4 Empirical Findings

In this section, I present the results of the estimation of my model. My baseline estimates reflect the model that only including granular firms (Compustat). In the next section, I discuss the robustness of my estimates to the inclusion of private and foreign firms as a representative competitive firm (with free entry).

### 1.4.1 Welfare Statics

My first empirical exercise is to compute total surplus and its breakdown into profits and consumer surplus. This is done for both the observed equilibrium (which is assumed to be a Nash-Cournot equilibrium) and the counterfactuals considered in Section 1.2. These estimates are all shown in Table 1.1.

I estimate that the (publicly-traded) firms in my sample earn an aggregate economic profit of \$5 trillion and produce an estimated total surplus of \$9.1 trillion. Consumer surplus is therefore estimated to be about \$4 trillion. About 55% of the total surplus produced is appropriated by the companies in the form of oligopoly profits. For context, the GDP of U.S. corporations in the same year (2017) is \$11 trillion. The difference between GDP and total surplus is that total surplus does not include the value of labor input but it does include the value of inframarginal consumption. GDP, on the other hand, includes the value of labor input but not the inframarginal value of consumption.[15]

The first counterfactual I consider, *Perfect Competition*, appears in the second column. The comparison between the Cournot-Nash allocation suggests that the welfare costs of oligopoly are significant. In this allocation, aggregate surplus is significantly higher – \$10.2 trillion – hence, the deadweight loss amounts to about 11% of the total surplus.

Although the implications of oligopoly for Pareto efficiency are significant, even more significant are the *distributional* implications. Because firms price at marginal cost, a much larger share of the surplus goes to the consumer: \$8.2 trillion, more than double than in the Cournot allocation. This amounts to 80% of the total surplus. This is a novel empirical finding that is made possible by estimating the GHL demand system.

[15] In this model, the labor supply is perfectly elastic, therefore each unit of labor is paid exactly its marginal disutility and there is no inframarginal value of leisure.

**Figure 1.1:** Example Product Space: Two Firms, Two Characteristics

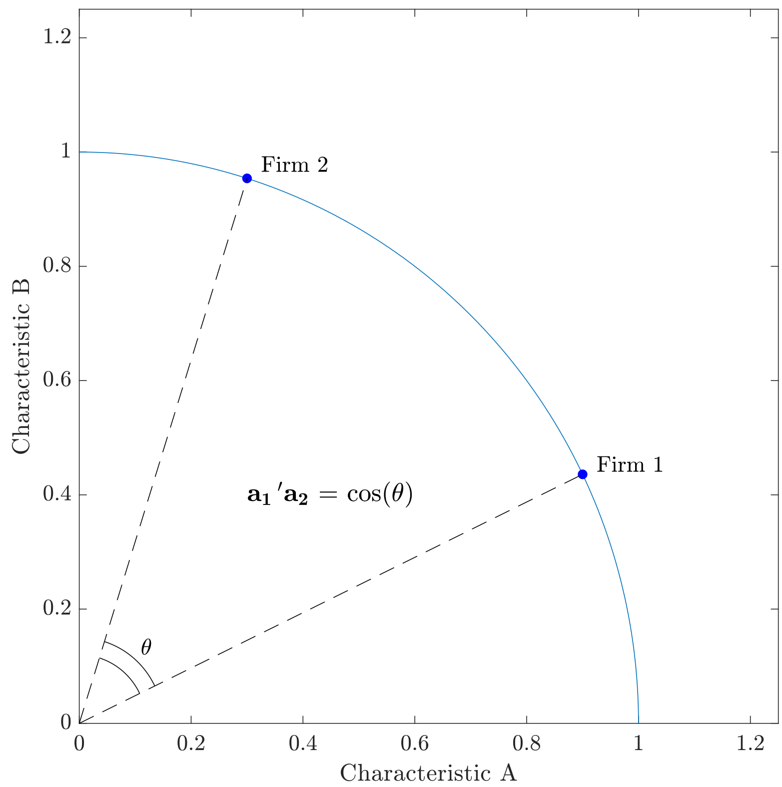


FIGURE NOTES: The following diagram exemplifies the hedonic demand model, for the simple case where there are only two product characteristics (A and B) and only two competitors (1 and 2). Each firm exists as a vector on the unit hypersphere of product characteristics (in this example, we have a circle). The dot product $\mathbf{a}_i'\mathbf{a}_j$ equals the cosine of the angle $\theta$. The tighter the angle, the higher the cosine similarity, and the larger (in absolute value) the inverse cross-price elasticity of demand.

**Figure 1.2:** Network Visualization of the Hoberg-Phillips Dataset

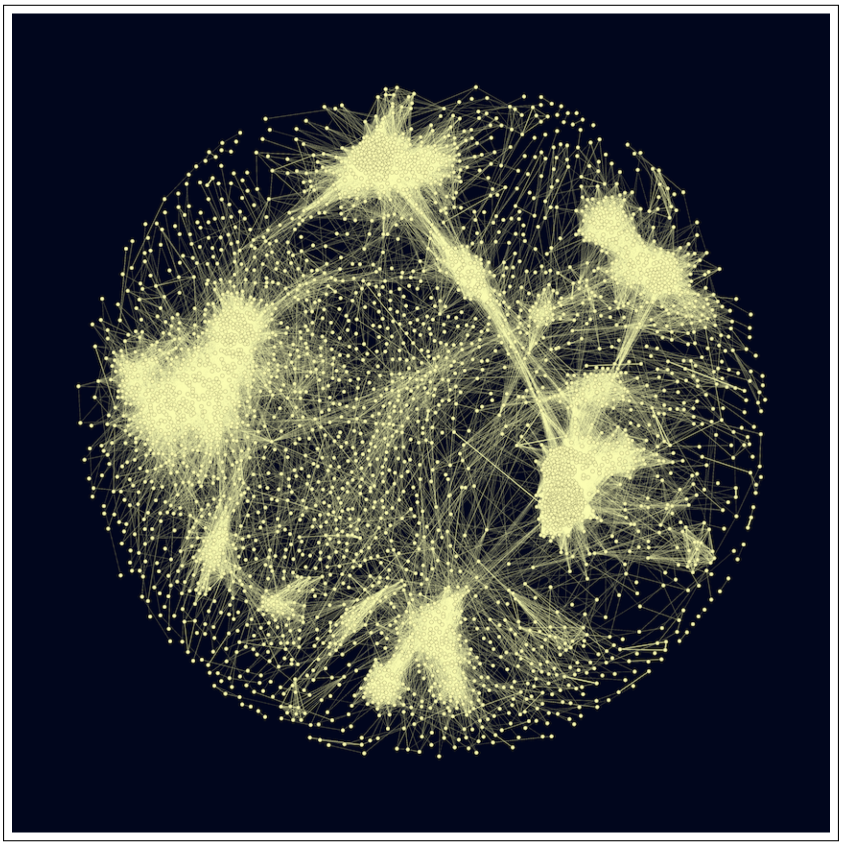

Figure Notes: The following diagram is a two-dimensional representation of the network of product similarities computed by Hoberg and Phillips (2016), which is used in the estimation of the model presented in Section 1.2. The data covers the universe of Compustat firms in 2004. Firm pairs that have thicker links are closer in the product market space. These distances are computed in a space that has approximately 61,000 dimensions. To plot this high-dimensional object over a plane, I applied the gravity algorithm of Fruchterman and Reingold (1991), which is standard in social network analysis.

**Table 1.1:** Welfare Statics (2017)

| | *Scenario* | *Cournot-Nash* (1) | *First-Best* (2) | *Monopoly* (3) | *Resource-Efficient* (4) | *Second-Best* (5) |
|---|---|---|---|---|---|---|
| **Welfare Statistic** | **Variable** | $\mathbf{q}^{\Phi}$ | $\mathbf{q}^{W}$ | $\mathbf{q}^{\Pi}$ | $\mathbf{q}^{H}$ | $\mathbf{q}^{2nd}$ |
| Total Surplus (US$ trillions) | $W(\mathbf{q})$ | 9.086 | 10.208 | 8.183 | 9.869 | 10.017 |
| Aggregate Profits (US$ trillions) | $\Pi(\mathbf{q})$ | 5.043 | 1.995 | 5.673 | 3.135 | 3.910 |
| Consumer Surplus (US$ trillions) | $S(\mathbf{q})$ | 4.043 | 8.213 | 2.510 | 6.735 | 6.106 |
| Total Surplus / Perfect Competition | $\frac{W(\mathbf{q}^{\Phi})}{W(\mathbf{q}^{W})}$ | 0.890 | 1.000 | 0.802 | 0.967 | 0.981 |
| Aggregate Profit / Total Surplus | $\frac{\Pi(\mathbf{q}^{\Phi})}{W(\mathbf{q}^{\Phi})}$ | 0.555 | 0.195 | 0.693 | 0.318 | 0.390 |
| Consumer Surplus / Total Surplus | $\frac{S(\mathbf{q}^{\Phi})}{W(\mathbf{q}^{\Phi})}$ | 0.445 | 0.805 | 0.307 | 0.682 | 0.610 |

Table Notes: The following table shows my estimates of aggregate profits, consumer surplus and total surplus in each of the counterfactuals scenarios presented in Section 1.4.


32

The next counterfactual I analyze, the *Monopoly* counterfactual, appears in the third column: it represents a scenario in which all firms are controlled by a single decision-maker that coordinates supply choices. In this allocation, aggregate surplus is significantly lower than in the Network Cournot equilibrium allocation: \$8.2 trillion. Despite the decrease in aggregate welfare, profits are significantly higher: \$5.7 trillion. Consequently, consumer surplus is reduced to just \$2.5 trillion, a mere 33% of the total.

Next, I consider the *Resource Efficient* counterfactual, in which the social planner maximizes total surplus subject to not changing overall labor usage. In this scenario, markups across firms have been equalized, but not eliminated. By removing all dispersion in markups, this counterfactual targets the malallocative effects of concentration.

The total surplus produced in this counterfactual is \$9.9 trillion, about 3.3 percentage points lower than in perfect competition, and \$800 billion higher than the observed Cournot-Nash equilibrium. Most of the surplus produced – \$6.9 trillion, 68% of the total – goes to the consumer; profits are reduced to \$3.1 trillion. Because labor is fixed, all the welfare gains with respect the Cournot equilibrium come from *reallocation* of labor. Hence, an important take-away from this counterfactual is that a large share of the inefficiencies from oligopoly are driven by resource misallocation. A different way to say this is that the dispersion in markups (caused by oligopolistic competition with differentiated products) matters at least as much as the level of markups in determining the overall deadweight loss.

The last counterfactual I consider is the *Second-Best,* in which a benevolent social planner maximizes aggregate surplus subject to an aggregate participation constraint (profits must cover fixed costs on average). In this counterfactual, total surplus is very close to the level achieved by the perfectly competitive outcome: \$10 trillion. The main difference is the surplus split: the consumer receives \$6.1 trillion (two trillion less than under perfect competition but two more than under Cournot), while total profits amount to \$3.9 trillion.

One takeaway of these findings is that oligopoly affects consumers are affected through two channels: it increases the dispersion of markups, generating resource misallocation which raises the deadweight loss; it also increases the level of markups, which in turn affects how surplus is shared between producers and consumers.

### 1.4.2 Time Trends in Total Surplus and Consumer Surplus

The data used for this paper, HP's Compustat, is available as far back as 1997. By mapping my model to firm-level data, year by year, I can produce annual estimates of the welfare metrics presented above. This allows me to study the welfare implications of the rising concentration of US industries. Most importantly, because my model leverages HP's time-varying product similarity data, these estimates account for how the product offering of US public firms changed over time. This is another contribution of this study.

In Figure 1.3, I plot aggregate consumer surplus $S$ (the dark area), and profits $\Pi$ (the light area) for every year between 1997 and 2017. The combined area represents total surplus $W(\mathbf{q})$. I also plot, on the right axis (dotted black line), profits as a share of total surplus $\Pi(\mathbf{q})/W(\mathbf{q})$.

The graph shows that the total surplus produced by US public corporations has nearly doubled between 1997 and 2017 from \$4.6 trillion to \$9.1 trillion between 1997 and 2017. Profits have increased more-than-proportionally with respect to consumer surplus – from about \$2.2 trillion to about \$5 trillion. Consumer surplus increased instead from \$2.2 trillion in 1997 to about \$4 trillion i 2017. As a consequence, the profit share of surplus has increased from about 50% of total surplus to nearly 56%. The consumer appears to capture a decreasing share of the surplus generated by public companies.

In Figure 1.4, I plot, over the same period, the percentage gain in total surplus from moving from the competitive equilibrium $\mathbf{q}^{\Phi}$ to the first best $\mathbf{q}^{W}$. This is the deadweight loss from oligopolistic behavior, and is plotted as the darker line. Both series experienced upward trends that mimic that of profit share of surplus: the total surplus gains have increased from 8.5% (in 1997) to the current level of 11% (in 2017). This suggests that oligopoly generates increasing Pareto-inefficiencies.

To investigate the impact of fixed costs on these results, I plot, in the same figure, the percentage difference in total surplus between the Cournot equilibrium and the *Second Best* (the light line): it increases from 6.4 percentage points (in 1997) to 9.3 percentage points (in 2017). In other words, when fixed costs are taken into account, the Pareto-inefficiencies generated by oligopolistic competition start from a lower level (as is to be expected, mathematically) but increase more sharply over time (+45% over the period)

Overall, my findings are consistent with the interpretation that the increasing concentration in U.S. industries over the past few decades reflect a generalized increase in oligopoly power that has negatively affected total surplus and, particularly, consumer welfare.

### 1.4.3 The Role of Entry Costs

The increase in concentration that has occurred across US industries is reflected by the stark decline in the number of public companies: their number has decreased from about 7,500 in 1997 to about 3,500 in 2017 (Kahle and Stulz, 2017).

One force that could potentially reconcile these facts with the increasing profits and deadweight loss is an increase in the costs of entry . While entry costs cannot be observed directly, we can use the model presented in Section 1.2 to obtain a proxy of the entry costs incurred by a marginal entrant (here we are treating each IPO as an entry). The idea is to ask how large should unobserved entry costs need to be, in order to deter an entrant that resembles a typical granular firm.

Thus far, I have kept the number of granular firms fixed, while allowing the number of atomistic firms to adjust endogenously. Now I am going to consider the entry problem of a marginal granular firm indexed by $i = 0$ that has quality-adjusted productivity $(b_0 - c_0)$ and fixed costs $(f_0)$ that are equal to the median of the population of incumbents. We can use the model to compute the economic profits $\pi_0$ conditional on entry. Supposing that firm 0 faces an entry cost of $e_0$ , and that firm 0 enters the market if $\pi_0 - f_0 \geq e_0$ , the net profits $(\pi_0 - f_0)$ earned by the marginal entrant conditional on entry provide a lower bound for the entry cost $e_0$.

Because in my model firms exist in a space of product characteristics, in order to make this measurement I also need to impute a cosine similarity $\mathbf{a}_0'\mathbf{a}_i$ between the potential entrant 0 and every other firm in the dataset. I use a measure of product cosine similarity computed by Hoberg, Phillips and Prabhala (2014). This cosine similarity is computed, for every Compustat company, against a generic venture capital-backed startup. It is based on the startups' product descriptions provided in the VenturExpert dataset (which covers VC-backed startups). This measure is available for the period 1997-2008.[16]

Figure 1.5 presents the resulting estimate of the implied entry cost, normalized at the level of 1997. Between 2001 and 2008, it increases by around 60%. What this exercise tells us is that incentives for this hypothetical startup to enter have increased; hence, from a marginal analysis perspective, the model suggests that constraints or disincentives to enter should also have increased over this period. This occurs at approximately the same time as the number of new IPOs collapses.

This exercise obviously does not shed any light on the reasons why this decline has occurred. This will be the object of my investigation in the next sub-section.

### 1.4.4 Startup Takeovers as a Driver of Concentration

One interesting and puzzling aspect of the decline in IPOs is that it does not appear to be related to the decline in the startup rate that has been observed in the broader economy (Decker, Haltiwanger, Jarmin and Miranda, 2014). Far from declining, the number of

[16] I thank the authors for retrieving and sharing this data, which is not part of the publicly-available HP data.

startups that are backed by Venture Capital (VC), which make up the majority of startups that eventually become public companies, has boomed over this period.

Figure 1.6, panel B, displays the number of Venture Capital exits in the United States by year and type, for the period 1985-2017[17]. In the diagram, I split VC exits between IPOs and acquisitions. It is very clear from that graph that, while at the beginning of the 1990s the vast majority of VC exits were IPOs, starting from the mid-90s there has been a dramatic shift toward acquisitions. One implication of this observation is that – from a simple accounting standpoint – the decline in IPOs was not driven by a dearth of startups. Instead, the reason why IPOs have decreased is that they have been largely replaced by acquisitions.

To what extent can this secular shift from IPOs to acquisitions account for the increasing profit share of surplus, and the rising welfare costs of oligopoly measured in Section 1.4? We can investigate this next question using the counterfactuals developed in subsection 1.2.5.

Specifically, I construct a counterfactual in which I add granular firms into the model. For each firm $i$ entering the Compustat-HP dataset, I add $(N_i - 1)$ new firms to the model. These firms are "similar" to $i$ in the sense they share the same value of $(b_i - c_i)$ as well as the same coordinates in the space of common characteristics $(\mathbf{a}_i)$; they also exit the sample in whichever year firm $i$ exits the dataset. However, they are not perfect substitutes to $i$, due to the presence of idiosyncratic characteristics.

$N_i$ is determined so that, in this counterfactual, the ratio of IPOs to acquisitions remains constant after 1997. Specifically, if we define $\mathrm{IR}_t$ to the ratio of IPOs to total VC exits at time $t$, we define $N_i$ as:

$$N_i = \begin{cases} \frac{\mathrm{IR}_{1997}}{\mathrm{IR}_t} & \text{if } i \text{ went public at time } t \\ 1 & \text{otherwise} \end{cases}$$

Figure 1.7 shows the result of this counterfactual exercise: it plots the difference in consumer surplus between the Cournot equilibrium and the *Perfect Competition* counterfactual, under two alternative scenarios for the set of active granular firms. The lighter line shows the baseline case: consistent with the findings of subsection 1.4.2, the percentage gap in consumer surplus increases from 42% to 51%, reflecting both a larger deadweight loss and the larger share of total surplus accruing to producers in the form of profits.

The darker line shows the counterfactual scenario in which the ratio of IPOs to acqui-

[17] This data is sourced from the National Venture Capital Association (NVCA).

sitions stays constant after the year 1996. Under this alternative scenario, the increase in consumer surplus gap is significantly less pronounced, ending at about 43.5% in year 2017. This reflects a more muted increase in the deadweight loss, as well as a *decrease* in the profit share of surplus. These calculations suggest that, from a quantitative point of view, the explosion of acquisitions of startups observed after the mid-90s is likely to have contributed to some of the measured welfare trends.

This counterfactual obviously comes with caveats. Most importantly, what I propose is a *proximate* explanation for the increasing oligopoly power: my model cannot be used to study the reasons for the secular shift from IPOs towards acquisitions. Also, the results from this counterfactual cannot be read as causal evidence.

That being said, these results do however complement a recent empirical literature that has focused on the anti-competitive effects of startup acquisitions. Wollmann (2019), for example, argues that startup acquisitions may have been used by large corporations to engage in what he calls *stealth consolidation*: the idea is that the majority of startup acquisitions fall under the reporting threshold for merger review by antitrust authority. Because startup acquisitions rarely undergo merger review as a result, large companies may be able to use startup acquisitions to engage in monopolization with little risk of attracting antitrust scrutiny. In another study, Cunningham, Ederer and Ma (2018) provide evidence from the pharmaceutical industry of what they call *killer acquisitions*: specifically, they provide evidence that a significant share of the acquisitions of startups by drug-makers are driven by the motive of suppressing the development of new products that might pose a competitive threat to their existing products.

## 1.5 Robustness and Extensions

### 1.5.1 Private and Foreign Firms, Endogenous Entry

I verify that my main results are not sensitive to the inclusion of private and foreign firms. Based on the aggregation result derived in subsection 1.2.6, I am able to include non-Compustat companies by adding to the model a representative firm that acts competitively and whose size reflects the endogenous selection into entry of atomistic players.

In order to implement this version of the model, I need to assume a cosine similarity between this representative firm, which I label $n + 1$, and every other firm $i = 1, 2, ..., n$. I assume that the cosine similarity between $i$ and $n + 1$ is simply equal to the average cosine similarity between $i$ and every other firm – formally:

$$\mathbf{a}_i'\mathbf{a}_{n+1} = \frac{1}{n-1}\sum_{j\neq i}\mathbf{a}_i'\mathbf{a}_j$$

My empirical results are only slightly changed as a consequence of this modification: corporate profits , as a percentage of the surplus produced by public firms, increase from 50% in 1997 to 55.5% in 2017. The deadweight loss increases from 7.8% in 1997 to 10.3% in 2017.

In order to validate the approach of modeling non-Compustat firms as atomistic, it is worth investigating whether Compustat becomes an increasingly-large or small share of GDP over time. I compute an estimate of the value added by Compustat companies, and investigate how it changes over time as a percentage of corporate business GDP. Reassuringly, I find that this percentage does not trend either positively or negatively over the period considered: it is 46.8 in 1997 and 47.4 in 2017, with a standard deviation of 5.9 percentage points over the period.

### 1.5.2 Fixed Costs

Thus far, I have defined aggregate profits ($\Pi$) and total surplus ($W$) *gross* of fixed costs ($F$). Next, I want to study how would Figure 1.3 and 1.4 change if we subtracted $F$ from $\Pi$ and from $W$. In other words, I want to investigate whether the higher economic profits are somehow justified by higher fixed costs.

In the Online Appendix, I reproduce both these figures after redefining $\Pi$ and $W$ to be computed net of fixed costs. By comparing Figure 1.3 and 1.4 with Figures **??** and **??**, we can see that my core empirical results are unaffected by how fixed costs are incorporated in the analysis. The most remarkable difference in the findings that we observe is that when we net fixed costs, is that the profit share of surplus increases much more dramatically over the period 1997-2017: from 11% in 1997 to 22% in 2017.

### 1.5.3 Intangible Capital

With regard to the estimation of fixed costs, there has been some debate in the literature about how Selling, General & Administrative (SGA) costs should be treated from an accounting standpoint. This item, as presented in Compustat, includes miscellaneous costs that are not directly linked to production (see Traina, 2018). It also includes R&D expenditures.

While it is generally understood that these are not variable costs, it is also not entirely clear that this cost item is simply overhead. Eisfeldt and Papanikolaou (2013) have argued that SGA partly embeds investments in *intangible capital*, and therefore should not be treated as overhead but capitalized.

Based on this argument, Peters and Taylor (2017) have developed a measure of intangible capital for Compustat: they treat R&D expenditures, plus 30% of the remaining portion of SGA, as investment in intangible capital. They then computed the firm-level intangible capital stock by applying a perpetual inventory model. If we choose to capitalize, rather than expense, these putative investments in intangible capital, we then obtain the following alternative measure of fixed costs:

$$f_i = (\text{SGA}_i - \text{R\&D}_i)\times 0.7 + (\text{Property Plant \& Equipment}_i + \text{Intangible Capital}_i)\times \text{User CoC} \tag{1.26}$$

Changing the definition of fixed costs does not affect my measurements in 1.3 nor the deadweight loss from 1.4, since fixed costs do not enter these measures. They do however affect the distance from the second-best, which is shown in 1.4, as well as my additional analyses from the Online Appendix, which I have discussed above. In the same Online Appendix, I show that that my results are virtually unaffected by how I account for these potential investments in intangible capital (capitalized v/s expensed).

### 1.5.4 Labor Supply Elasticity

My utility specification implies a perfectly-elastic labor supply. Next, I investigate how my empirical results change if made the opposite assumption: a completely-inelastic labor supply function.

By definition, profit as a share of total surplus would be the same. The deadweight loss would instead become the welfare difference between the Cournot equilibrium and the *Resource-Efficient* counterfactual, which we previously described in subsections in Section 1.2.5. As can be seen in Table 1.1, this welfare difference is smaller than the deadweight loss. Intuitively, this is because the labor supply (by definition) cannot respond to the removal of the oligopolistic distortions.

I compute this alternative measure of the deadweight loss (the percentage difference in total surplus between Cournot and *Resource Efficient*) over the period 1997-2017. I find that my core empirical results carry through: the level of this "alternative" deadweight loss is 5.2% in 1997, and it increases to 7.9% by 2017. In other words, the level of the deadweight loss is lower if we assume a fixed labor supply (as should be expected), but it increases more sharply (by half) over the 20-year period.

### 1.5.5 Multi-product Firms (Diversification vs. Differentiation)

Like most of the existing macroeconomics literature, this paper does have access to product-level data on characteristics. Hence, while the model does reasonably well in capturing measured variation in markups across firms (see the Online Appendix), and it is conceivable that my model could be taken to product-level given more granular data, it is important to underline that my empirical exercises have nothing substantive to say about how the market power of a firm may vary across individual product markets (say, how Apple's market power in the smartphone industry may be higher or lower than in the personal computer market). Nevertheless, it is worth discussing the conditions under which my model may apply to multi-product firms, given that the extensive presence, in Compustat, of multi-product firms.

Suppose that there are still $n$ firms and $k$ characteristics, but now the $n$ firms produce a total of $m \geq n$ products. The same product might be produced by multiple firms and the same firm may produce more than one product. The vector of units produced for each good is now the $m$-dimensional vector $\mathbf{y}$. Similarly to matrix $\mathbf{A}$ in Section 1.2, matrix $\mathbf{A}_1$ transforms units of products into units of characteristics:

$$\mathbf{x} \quad = \quad \mathbf{A}_1 \mathbf{y} \tag{1.27}$$

Because firms are diversified, each firm now produces a basket of goods: instead of representing the number of units produced of each product, the vector $\mathbf{q}$ now represents the number of baskets produced by each firm. The matrix $\mathbf{A}_2$ projects quantity indices for each basket/firm onto units of products supplied:

$$\mathbf{y} \quad = \quad \mathbf{A}_2 \mathbf{q} \tag{1.28}$$

Now I put together the previous two equations. Letting $\mathbf{A} = \mathbf{A}_1 \mathbf{A}_2$, I have

$$\mathbf{x} \quad = \quad \mathbf{A}_1 \mathbf{y} \quad = \quad \mathbf{A}_1 \mathbf{A}_2 \mathbf{q} \quad = \quad \mathbf{A} \mathbf{q}$$

The relationship above demonstrates how the linear hedonic structure of the model makes the model generalizable to multi-product firms. The intuition is that, if the output of a certain firm $i$ is not a single product, but rather a basket of products, one can equivalently project the basket quantity index $q_i$ onto the characteristics space in two steps (by projecting it first onto goods and then onto characteristics), or in one single step (using the composite projection matrix $\mathbf{A}$).

The limitation of this multi-product interpretation of the model is that, while firms can change their supply $q_i$, the vector $\mathbf{a}_i$ must stay fixed. What this means is that while firms may produce more than one product and scale up or down the quantity of the basket of products produced, they must keep producing the products in constant quantity ratios. An intuitive way to say this is that the limitation of firm-level data is that it does not allow to study the reallocation of resources *within* firms, but only *across* firms. This implies that my estimates of the deadweight loss are likely to be conservative: if firms have different degree of market power in different markets, this will generate *within-firm* variation in markups. My model does not capture the additional welfare gains that could be realized if we removed *within-firm* dispersion in markups.

It is important to emphasize that this limitation is not specific to this paper, but it is endemic in the literature. While my linear demand specification does not fully address it, I claim that the GHL demand system handles multi-product firms significantly better than CES preferences.

Additionally, suppose we had product-level data (including similarity scores for individual products), so that we could relax the assumption that $\mathbf{A}_2$ is fixed (firms can change output ratios among products); also, suppose that each product within the firm is produced by a distinct *plant*. The model would then need to account for the fact that plants within the same firm coordinate their output choices, as firms do in the counterfactual from equation 1.2.5.[18] Hence, one argument in favor of using firm-level data is that doing so is a tractable way of modeling plant coordination within firms.

Finally, another concern is how does diversification affect the measured intensity of competition across firms. If, for example, firms become more diversified, would the model erroneously interpret the resulting change in cosine similarity as an increase in market power? The answer to this question is *it depends:* the effect of diversification of measured competition is ambiguous. It is possible to construct examples where diversification leads to higher, lower, or unchanged cosine similarity. In the Online Appendix, I construct three such examples to illustrate this argument.

### 1.5.6 Complements

Because the matrix $\mathbf{\Sigma}$ is non-negative by construction, the marginal utility from one unit of product $j$ is always non-increasing in $q_i$ - formally:

$$\frac{\partial^2 S}{\partial q_i \partial q_j} = -\sigma_{ij} \leq 0 \quad \forall\, i \neq j \tag{1.29}$$

[18] A number of papers in the literature map plant-level data to model firms, ignoring the fact that firms that are rationally-managed should coordinate output choices across plants. The obvious rationale for this approach is tractability.

In light of equation (1.29), it is tempting to jump to the conclusion that all products are by construction substitutes and that no pair of products are complements. That conclusion is, however, incorrect.

To understand why, we need to recall the textbook definition of substitution and complementarity. Two goods $(i, j)$ are:

$$\text{Complements} \quad \text{if} \quad \frac{\partial q_i}{\partial p_j} < 0 \qquad \text{Substitutes} \quad \text{if} \quad \frac{\partial q_i}{\partial p_j} > 0 \tag{1.30}$$

We intuitively expect this derivative to have the opposite sign of that in equation (1.29). In the case of CES, this intuition is correct. In the case of my model, however, this intuition fails. This is a consequence of the fact that the cross-price demand elasticity depends on the inverted matrix $(\mathbf{I} + \mathbf{\Sigma})^{-1}$, not on $\mathbf{\Sigma}$ itself. If $\mathbf{\Sigma}$ is not symmetric (here it is not) the off-diagonal elements of $-(\mathbf{I} + \mathbf{\Sigma})^{-1}$ will generally include positive as well as negative elements. This implies that, in my implementation of the model, many producer pairs are strategic complements.[19]

For example, if we compute the vector of cross-price derivatives for car manufacturer General Motors in 2017, we will find that it includes several negative elements (i.e. complements), mostly corresponding to energy and consumer finance companies. This makes sense: intuitively, we expect higher oil prices, loan rates or insurance premia to adversely affect the residual demand for cars.

Hence, despite the property of the model described by equation (1.29), my model does indeed produce strategic complementarity. I argue that one of the strengths of the network Cournot model is that its asymmetric nature allows to capture a rich set of strategic interactions.

### 1.5.7 Limitations and Future Work

This model (like every other macroeconomic model) has limitations and it does leave out certain aspects of market power that might be relevant to the current debate on antitrust policy.

For example, one important assumption that I make in my model (in order to retain tractability) is that all firms are final goods firms. In other words, input-output linkages between individual firms are not part of the model. This might potentially lead to

[19] It is fairly easy to come up with counter-examples: consider the following, with three goods (1,2,3) and three common characteristics (A,B,C). If goods 1 and 3 load entirely on characteristic A and C, respectively, and good 2 loads equally on all three characteristics, then it can be verified that goods 1 and 3 are strategic complements.

underestimating the welfare costs of oligopoly, if input-output linkages result in double marginalization. One interesting extension of the model would be to modify the firm's production function to allow firms to use other firm's output as inputs. In order to take such extended model to the data, granular input-output data on firm-to-firm relationships would be required.

Another important restriction in my model is that it treats the firms' position in the product characteristics space as fixed. While the assumption that product characteristics are exogenous is standard in the demand estimation literature (Berry et al., 1995; Nevo, 2001), it renders models less suitable to make predictions about the long-run effects of policies that aim to reduce market power. This is because, given enough time, firms may be able to endogenously change their product portfolios. A natural direction for future research is to endogenize the firms' position in the characteristics space. Some progress on this front has already been made in the IO literature (see Fan, 2013; Wollmann, 2018). It is not obvious currently how this could be achieved in my setting, given that HP's database only provides similarity data ($\mathbf{A}'\mathbf{A}$) and not individual characteristics ($\mathbf{A}$).

Finally, another force that is left out of this paper is labor market power. Concentration may lead not only to oligopoly power in output markets: it may also lead to oligopsony in input markets. Because the model that I presented does not speak to this channel, my estimates of income distribution and aggregate efficiency do not include the effects of the labor market power, which might also have increased over this period.

With additional labor market data, it is definitely possible to apply the methods developed in this paper to a labor market setting. Cosine similarities can easily be constructed for textual descriptions of job vacancies: if we see labor input as a differentiated good, we can then develop and estimate a labor market version of the model presented in Section 1.2. In the Online Appendix, I propose two ways of reframing my model to study labor market power: the first considers *workers*' monopoly power (workers with unique characteristics command higher wages); the latter looks at oligopsony (firms that utilize unique inputs are able to charge higher markups).

## 1.6 Conclusions

In this study, I have presented a new general equilibrium model of oligopolistic competition with hedonic demand and differentiated products, with the objective of measuring the welfare consequences of rising oligopoly in the United States from 1997 to 2017. To estimate my model I used a data set (recently developed by Hoberg and Phillips, 2016) of bilateral product similarity scores that covers all public firms in the United States on a yearly basis. Through the lens of my model, these similarity scores are used to retrieve the cross-price elasticity of demand for every pair of publicly traded firms.

My measurements suggest that oligopoly has a considerable and growing effect on aggregate welfare. In particular, I estimate that, if all publicly traded firms were to behave as atomistic competitors, the total surplus produced by this set of companies would increase by 11 percentage points. Consumer welfare would increase even more dramatically—it would more than double—as surplus would be largely reallocated from producers to consumers. I find that most of the deadweight loss caused by oligopoly (7.9 percentage points) can be attributed to resource misallocation—that is, a significant share of the deadweight losses could theoretically be recovered by a benevolent social planner, even if we assumed labor to be inelastically supplied.

I also find evidence of large potential welfare losses from collusive behavior (or gains from antitrust): consolidating firm ownership in the hands of one producer that induces firms to collude would depress aggregate surplus by about 10 percentage points. Consumer surplus would suffer even more, with a projected decrease of about 38 percentage points. Overall, my analysis of firm-level data suggests that there is evidence of sizable welfare distortions due to oligopoly power.

By mapping my model to firm-level data for every year between 1997 and 2017, I find that, while both the profits earned by U.S. public corporations and the corresponding consumer surplus have increased over this period, profits have increased at a significantly faster pace: consequently, the share of surplus appropriated by firms in the form of oligopoly profits has increased substantially (from 50% to 55.5%). Consistent with this finding, I estimate that the welfare costs of oligopoly, computed as the percentage increase in surplus that is obtained by moving to the competitive outcome, have increased (from 8.5% to 11%). Overall, my estimates are consistent with the hypothesis that the observed secular trends in markups and concentration have resulted in increased welfare losses, particularly at the expense of the consumer.

The model allows me to compute a number of novel counterfactuals that are highly relevant for current issues in antitrust policy, and to shed light on the alleged oligopolization of U.S. industries. I have shown that a potential contributor to the measured trends might lie in the secular decline in IPOs and the surge in takeovers of VC-backed startups. Through the lens of my model, this shift can quantitatively account for a large share of the measured increase in the deadweight loss from oligopoly as well as the larger share of surplus accruing to producers.

This paper contributes—both methodologically and empirically—to a growing literature in macroeconomics and fin‘nce that is devoted to incorporating heterogeneity, imperfect competition and Industrial Organization methods in general equilibrium models. In particular, it shows that combining firm financials with measures of similarity based on natural-language processing of regulatory filings offers a promising avenue to model product differentiation and imperfect substitutability at the macroeconomic level: it affords the opportunity to impose a less arbitrary structure on the degree of substitution across sectors and firms.

While this paper makes a theoretical contribution first, the welfare measurements enabled by my model presented add to a growing body of empirical work on the effects rising market power (De Loecker, Eeckhout and Unger, 2020) and the anti-competitive effects of startup acquisitions (Cunningham et al., 2018; Wollmann, 2019). One potential policy implication of my findings is that, while antitrust agencies tend to focus most of their merger review work on mergers between large incumbents, acquisitions of VC-backed startups may also have important implications for competition. These findings strengthen the case for increasing the antitrust oversight of these transactions, which are remarkably frequent and may provide an avenue for large corporations to buy out nascent competition with a little risk of undergoing merger review.

**Figure 1.3:** Total Surplus of US public firms (1997-2017)

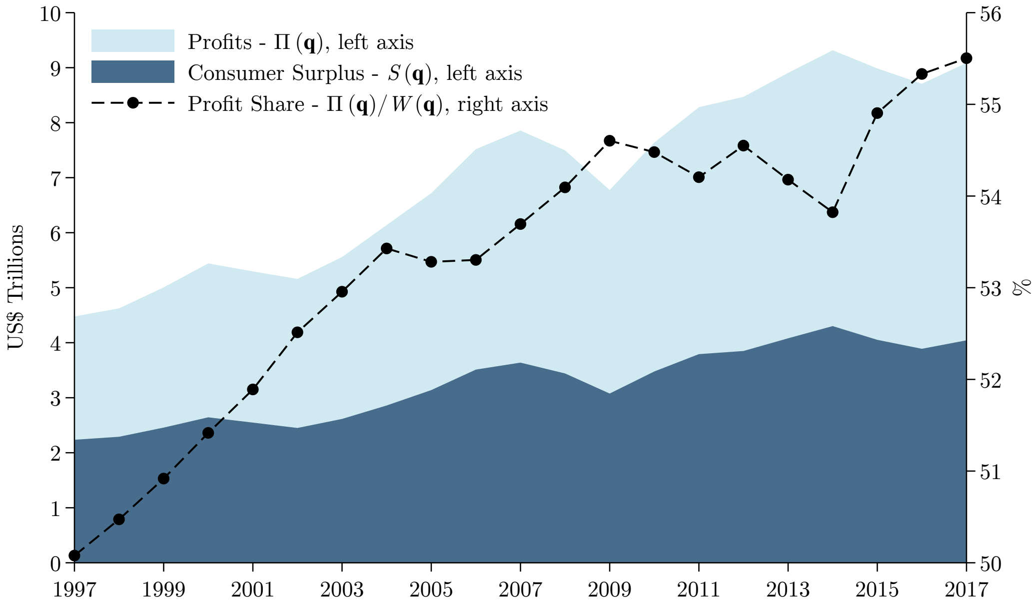


Figure Notes: The figure above plots the evolution, between 1997 and 2017, of aggregate (economic) profits $\Pi(\mathbf{q})$, aggregate consumer surplus $S(\mathbf{q})$ and total surplus $W(\mathbf{q})$, as defined in the model from Section 1.2. Profits as a percentage of total surplus ($\Pi/W$, black dotted line) are shown on the right axis. These statistics are estimated over the universe of the US publicly-listed corporations. These surplus measures are gross of fixed costs. The Online Appendix replicates this graph using surplus measures that are net of fixed costs.

**Figure 1.4:** Deadweight Loss from Oligopoly (1997-2017)

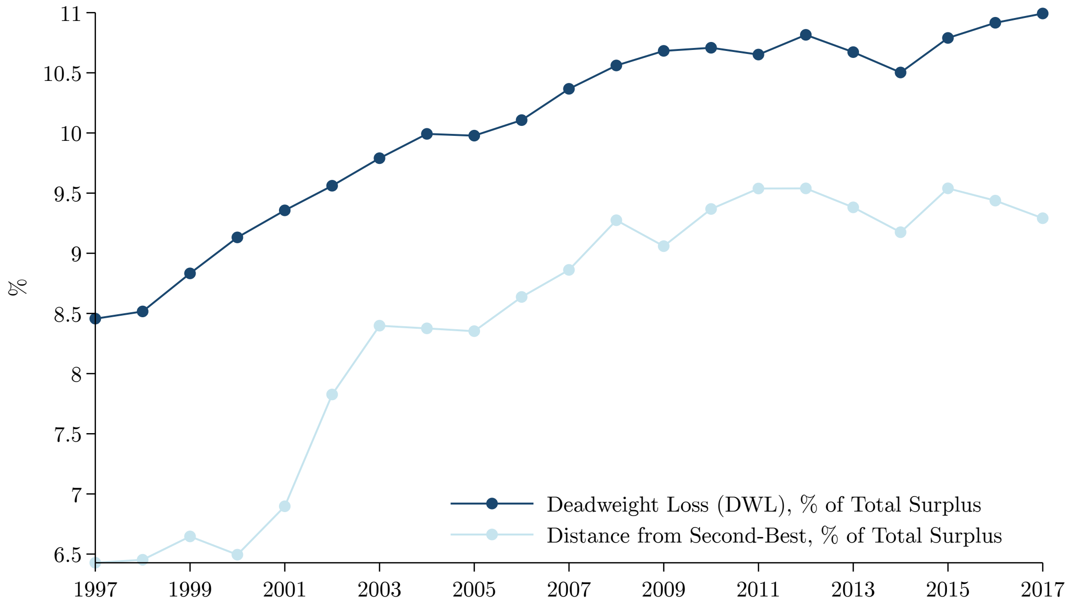


FIGURE NOTES: The following figure plots the estimated deadweight loss (DWL) from oligopoly, between 1997 and 2017. The lighter line is the traditionally-defined DWL - the % difference in total surplus between the Cournot equilibrium and the First-Best scenario, while the darker line is the % difference between the Cournot equilibrium and the Second-Best scenario as defined in Section 1.2. These surplus measures are gross of fixed costs. In the Online Appendix, I replicate this graph using surplus measures that are net of fixed costs.

**Figure 1.5:** Implied Cost of Entry for a VC-backed Startup

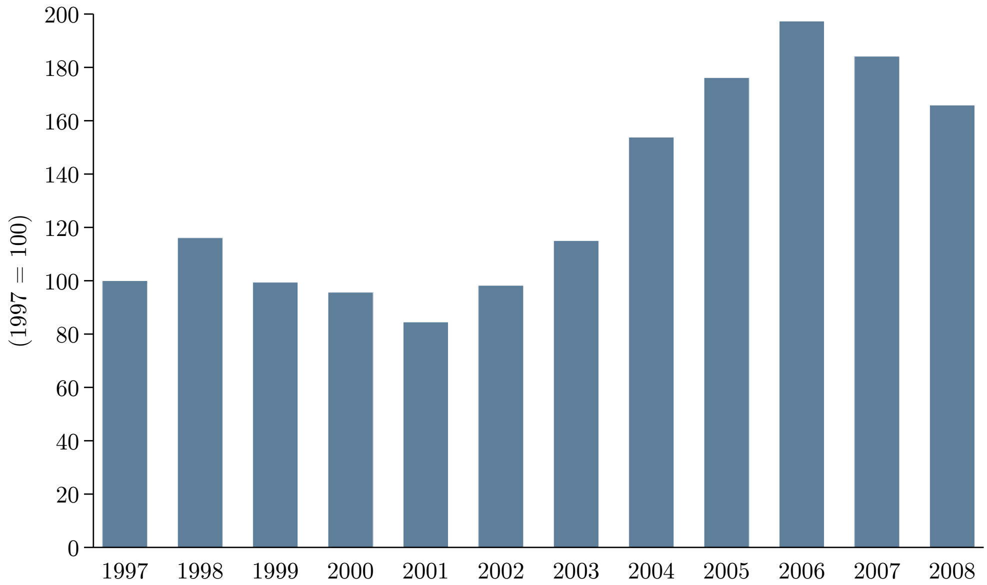


FIGURE NOTES: the figure above plots the implied cost of entry for a VC-backed startup with median quality-adjusted productivity $(b_i - c_i)$ and fixed cost $(f_i)$, as implied by the model in Section 1.2. The similarity scores for the marginal entrant are computed by Hoberg, Phillips and Prabhala (2014), using VenturExpert product descriptions.

**Figure 1.6:** Venture Capital Startup Exits by Type

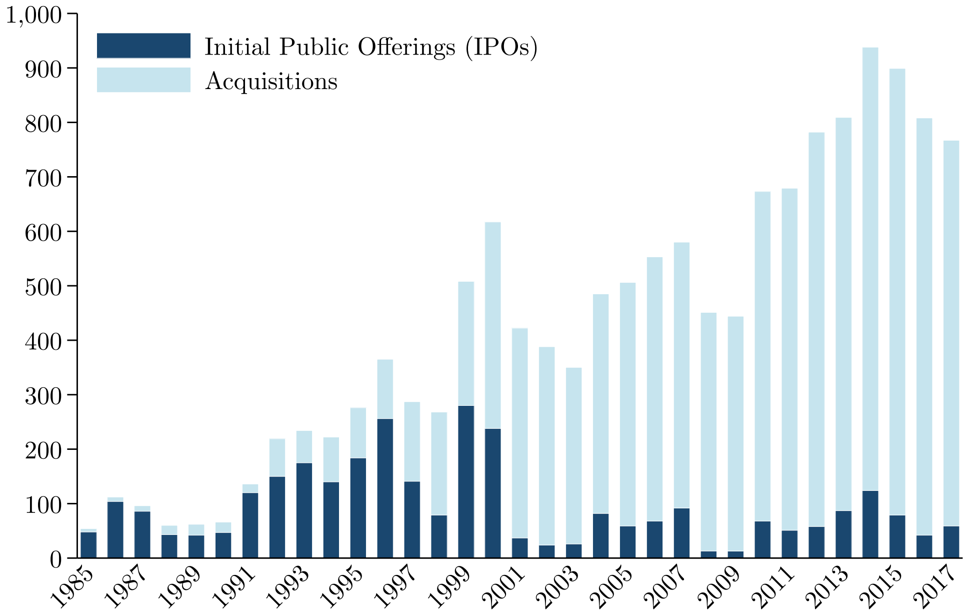


FIGURE NOTES: the figure above plots the number of successful venture capital exits in the United States by year and type (Initial Public Offering v/s Acquisition). The data is sourced from the National Venture Capital Association (NVCA).

**Figure 1.7:** Consumer Surplus, % difference from Perfect Competition

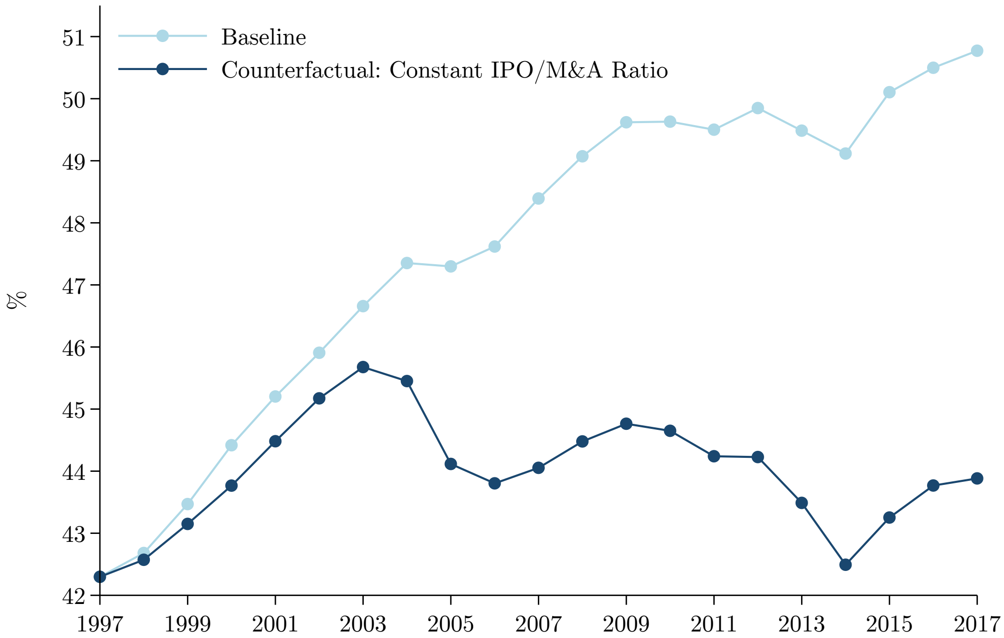


Figure Notes: The following figure plots the percentage difference in consumer surplus between oligopolistic competition and perfect competition, between 1997 and 2017. This is defined as the percentage increase in consumer surplus from the Cournot Equilibrium to Perfect Competition. The observed equilibrium value (light line) is plotted against a counterfactual scenario (darker line) in which the ratio of IPOs to startup acquisitions remains constant after 1997.

# Chapter 2

# Measuring the Economic Cost of Red Tape

Bruno Pellegrino, UCLA[1]

Geoffery Zheng, UCLA

[1] Bruno Pellegrino gratefully acknowledges financial support from the Price Center for Entrepreneurship at UCLA Anderson and travel sponsorship from International Association for Applied Econometrics. We thank Hugo Hopenhayn, Andy Atkeson, Chang-Tai Hsieh, Nico Voigtländer, Simone Lenzu, Benjamin Moll, Rodrigo Pinto, as well as seminar participants at University of Chicago, UCLA, UC Berkeley, LBS, and the summer Econometric Society/IAAE meetings, for helpful comments and feedback. We also thank Carlo Altomonte, Tommaso Aquilante and the Bruegel staff for their assistance in using the EFIGE dataset.

Survey evidence from around the world suggests that regulations and bureaucracy can impose a significant cost on firms and discourage investment. In this paper, we estimate the effect of red tape on aggregate output and investment across 85 countries using a dynamic general equilibrium model that features firm-level heterogeneity in the impact of red tape. To estimate the effect of red tape on individual firms' investment, we augment country-level data with a unique combination of firm-level survey data and balance sheets. Our key innovation is to explicitly model the firms' decisions to report obstacles to growth due to red tape in a survey. We show that the distortionary impact of bureaucracy manifests as a shift in the distribution of the marginal revenue product of capital, for firms that report being constrained by red tape. Measuring this shift allows us to infer, through our model, the effect of red tape on aggregate output. Bureaucracy reduces aggregate output through two channels: it reduces aggregate investment and it misallocates factors of production between firms. We find that the economic cost of red tape varies widely across countries, and that it may account for up to 7% of the observed cross-country dispersion in output per employee.

## 2.1 Introduction

One of the objectives of economics as a discipline is to inform legislation that is conducive to economic growth. Regulations are a key policy tool to address market failures. Even more importantly, they establish those property rights and contracting institutions that form the foundation of a well-functioning market economy (Williamson, 2000b).

Regulations, however, also impose a burden on businesses and citizens, who incur monetary and non-monetary costs to learn about and comply with them (Posner, 1975; Gray, 1987). Regulations (pretty much by design) restrict the agents' choice sets: they can therefore distort the allocation of resources; their enforcement also absorbs valuable resources. As a result, benevolent legislators should ideally balance the costs and benefits when deciding whether to increase the regulatory burden.

Because dealing with bureaucracy can impose significant costs on businesses, economic cooperation organizations, such as the Organisation for Economic Co-operation and Development, the International Monetary Fund, and the World Bank, routinely encourage member countries to cut red tape, as a way to improve a country's business environment, encourage investment and boost aggregate output. A whole literature in economics (Djankov, La Porta, Lopez-de Silanes and Shleifer, 2002, henceforth DLLS) has emerged in the early 2000s with the objective of measuring cross-country differences in regulatory burden.

Measuring the economic cost of bureaucracy is crucial to gain a better understand of the trade-off that legislators incur when deciding whether to impose new regulations. But how do we measure the cost of red tape? In this paper, we present a novel theoretical and empirical analysis to address this question.

Our analysis starts with some simple empirical observations. The upper panel of Figure 2.1 displays GDP per employee in 2011 US dollars (as measured by the Penn World Tables), plotted against a measure of entry regulations, which we obtained from the well-known dataset of DLLS, and which we use as a proxy of red tape. The graph shows an evidently-strong negative correlation between these two variables. In the lower panel we replace GDP per employee with capital stock per employee. We find a very similar relationship, except for a much steeper slope – suggesting that the correlation between income and regulation is possibly driven by capital accumulation.

These graphs only present an intruiguing correlation, of course – not a causal relationship. Yet, these correlations motivate what we perceive is an important question: is there an economic mechanism driving (at least partially) this correlation? How do we model and estimate the impact of regulations in terms of aggregate output?

Previous research has attempted to shed light on this question using econometrics, and by proposing various sources of exogenous variation in regulatory quality to estimate the causal effect of regulations on growth (Hall and Jones, 1999; Acemoglu et al., 2001b). While the balance of evidence suggests that indeed regulations adversely impact growth, there's a dearth of work linking the effect of red tape on growth and investment in a structural model that can produce quantitative estimates of its aggregate cost.

In this paper, we attempt to shed light on this old question by adopting the novel approach of leveraging survey micro-data and combining dynamic macroeconomic modelling with firm-level data analysis.

We start by developing a parsimonious general equilibrium model with heterogenous firms that are diversely (and adversely) affected by red tape. We model the impact of bureaucracy as a shadow tax on capital, in the style of Hsieh and Klenow (2009a). Unlike previous papers studying resource misallocation, we attempt to make a specific statement regarding the nature of these distortions. We do so thanks to the fact that we have access to firm-level survey data (from the EFIGE dataset) that allows to partially identify cross-sectional variation in the impact of red tape. We propose a dynamic model with an

**Figure 2.1:** GDP and Capital per Employee v.s. Entry Regulations

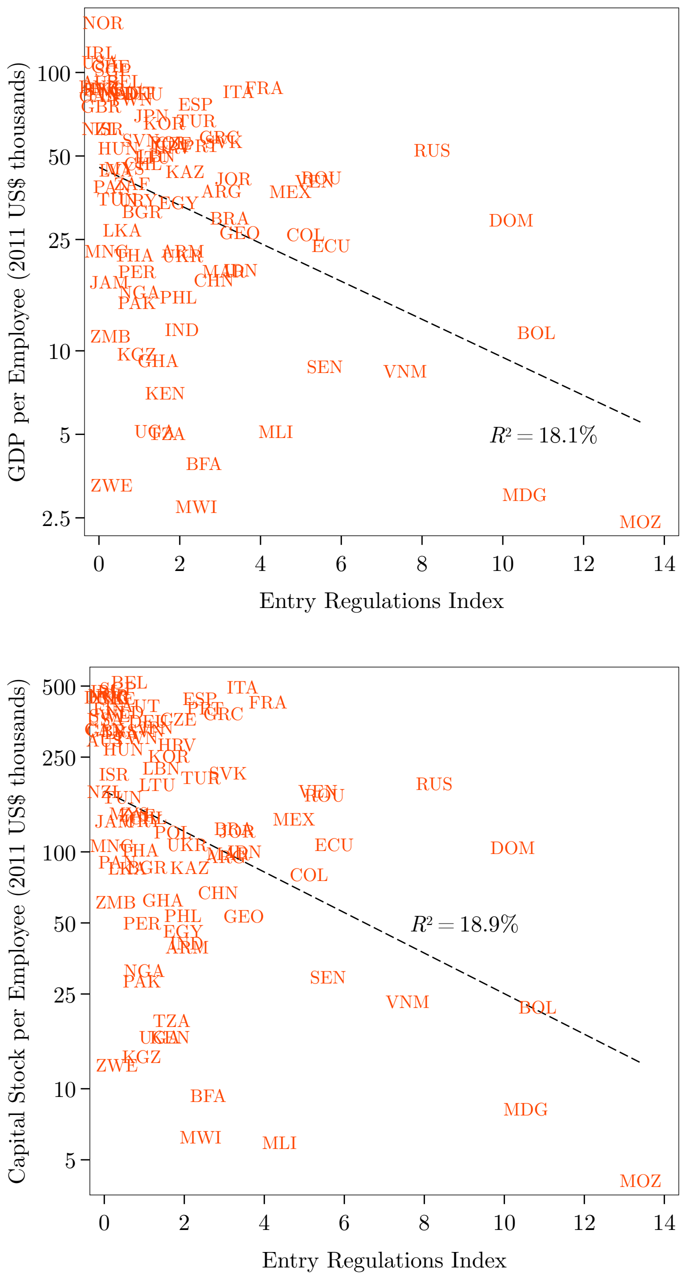


Figure Notes: The figure above plots GDP per employed person (upper panel) and Capital Stock per employed person (lower panel) in 2011 US$ thousands, against an index of entry regulations computed from the dataset of Djankov, La Porta, Lopez-de Silanes and Shleifer (2002). Each observation is a country. The variables plotted on the vertical axis, which uses a log scale, were obtained from the Penn World Tables v9.1.

endogenous saving rate, which allows us to disentangle two margins of the potentially-adverse effect of bureaucracy on output. The first is under-investment: by imposing a positive shadow tax on capital, red tape discourages investment across all firms. The second channel is capital misallocation: cross-sectional heterogenity in the impact of red tape distorts the allocation of capital among firms.

Our model possesses the desirable feature of producing closed-form formulas for both the percetage loss in aggregate output, and for the percentage change in aggregate TFP induced by red tape. The latter isolates the misallocation channel.

Both these formulas depend on a parameter that characterizes the cross-sectional distribution of these wedges. This implies that, in order to obtain an estimate of the output loss, we have to make a parametric assumption about the probability distribution of the shadow taxes induced by bureaucracy over the set of active firms in the economy (we make the assumption that the shadow taxes are Pareto-distributed).

To identify the shape parameter of the distribution of the wedges (the key unknown parameter in our model), we perform a firm-level regression analysis, leveraging a unique dataset that combines firm-level balance sheet data from the Bureau Van Dijk with survey data. The survey prompts firms to report whether red tape imposes a significant constraint on the company's growth. We take the radical approach of explicitly modelling the firms' decision to report growth constraint in the survey. We then show how to identify the shape of the shadow tax distribution using moments of the combined survey/balance sheet data.

The intuition behind our approach is that, to the extent that the survey data correctly identifies firms that are most adversely affected by red tape, we should observe a "shift" if the cross-sectional density of the marginal revenue product of capital (MRPK). By effectively measuring the magnitude of this shift, we can identify the probability distribution of shadow taxes.

In the final part of our empirical analysis we use DLLS's extensive data, which covers 85 countries in total, to retrieve an estimate of this parameter for each of those countries. Based on our model, we then estimate for each of those the GDP loss induced by red tape. We find that there is wide variation among countries in the impact of red tape: it ranges from as low as 0.76% of GDP in Australia to as high as 33% of GDP in Mozambique (although the estimates for right outlier countries should be taken with some caution). Interestingly, the GDP loss appears to reach significant levels even in relatively well developed countries such as Italy (2.25%) and France (3.03%).

We find that the effect of Bureaucracy on the capital misallocation are highly non-linear: except for a limited number of countries where the estimated impact of Bureaucracy is largest, the aggregate TFP loss (which identifies the misallocation channel) is generally a small share of the GDP loss. However, for a few countries in the right tail of the distribution (such as Russia or Venezuela), the capital misallocation can be quite significant, accounting for up to a third of the overall GDP loss.

Our paper contributes to the literature studying cross-country differences in institutions and growth from a new angle. We focus in this Section on only the papers most proximate to ours, and refer readers to handbook chapters for a more wholistic survey of the literature.

## 2.2 Model

In this section we present a parsimonious dynamic general equilibrium model that incorporates heterogeneous firm-level capital distortions due to bureaucracy.

The economy features a representative agent with the following utility function:

$$\sum_{t=0}^{\infty} \beta^t U(C_t)$$

$U(\cdot)$ increasing, concave and twice differentiable. $C_t$ is the consumption of final good at time $t$ and $\beta$ is the subjective discount rate. The representative agent is endowed with labor units $L$ for period, which they supply inelastically at a wage rate $w$.

There is a final good producing firm that produces output $Y_t$ using a CES technology and taking inputs $y_{it}$ from $i \in [0,1]$ final good-producing firms:

$$Y_t = \left( \int_0^1 y_{it}^{\frac{\sigma-1}{\sigma}} \, \mathrm{d}i \right)^{\frac{\sigma}{\sigma-1}}$$

The final good firm behaves competitively, hence the price of the final good is:

$$P_t = \left( \int_0^1 p_{it}^{1-\sigma} \, \mathrm{d}i \right)^{\frac{1}{1-\sigma}}$$

and $p_i$ is the price of input $i$. The intermediate good firms use Cobb-Douglas production function:

$$y_{it} = z_{it} \, k_{it}^{\alpha} \, \ell_{it}^{1-\alpha}$$

where $k_{it}$ is capital and $\ell_{it}$ is labor, which is assumed to be the numeraire. The multiplicative term $\left(\frac{1-\alpha}{\alpha}\right)$ in front of capital simply shifts the distribution of $z_i$ . It does not affect the model other than simplifying the algebra. Each firm $i$ rents capital and labor from the representative agent at prices $r_t$ and 1, respectively, so that firm $i$'s profits are:

$$\pi_{it} \;=\; p_{it}y_t - r_t k_{it} - w_t \ell_{it}$$

Following Hsieh and Klenow (2009a, henceforth HK), we assume that the firm maximizes a distorted profit function:

$$\tilde{\pi}_{it} \;=\; p_{it}y_{it} - e^{\tau_i} r_t k_{it} - w_t \ell_{it}$$

where $\tau_i$ is the (unrealized) shadow tax on capital imposed by bureaucracy/red tape. This implies that, in equilibrium, the shadow tax of bureaucracy shows up in the data as excess profits. We assume that this shadow tax is invariant over time. Choosing to model

Let the aggregate capital and labor supply be:

$$K_t \;=\; \sum_i k_{it} \qquad \text{and} \qquad L = \sum_i \ell_{it}$$

Capital depreciates at rate $\delta$ from period to period. Aggregate consumption is equal to income minus the required investment to bring capital to $K_{it+1}$ in the next period.

$$\underbrace{P_t Y_t}_{\text{Nominal GDP}} \;=\; P_t\left(C_t + \underbrace{K_{t+1} - \delta K_t}_{\text{Real Investment}}\right) \;=\underbrace{r_t K_t + w_t L + \Pi_t}_{\text{Nominal GNI}}$$

where $(1-\delta)$ is the rate of depreciation of capital. The Euler equation is:

$$\beta\, U'\left(C_{t+1}\right)\left(\frac{r_{t+1}}{P_{t+1}} + \delta\right) \;=\; U'\left(C_t\right)$$

We look for the steady state equilibrium. Then the Euler equation yields the following long-run equilibrium real interest rate:

$$\frac{r}{P} \;=\; \frac{1}{\beta} - \delta \tag{2.1}$$

The first order conditions for firm $i$ are:

$$\text{MRPK}_{it} \stackrel{\text{def}}{=} \frac{\sigma-1}{\sigma}\alpha \cdot \frac{p_{it}y_{it}}{k_{it}} \;=\; e^{\tau_i} r_t \tag{2.2}$$

$$\text{MRPL}_{it} \stackrel{\text{def}}{=} \frac{\sigma-1}{\sigma}\left(1-\alpha\right)\frac{p_{it}y_{it}}{\ell_{it}} \;=\; w_t$$

They yield the following steady-state capital-labor ratio:

$$\frac{k_i}{\ell_i} = \frac{\alpha}{1-\alpha} \cdot \frac{w}{re^{\tau_i}}$$

We can then write output in terms of labor alone:

$$y_i = z_i \ell_i \left( \frac{\alpha}{1-\alpha} \cdot \frac{w}{re^{\tau_i}} \right)^{\alpha}.$$

Because demand is isoelastic firms price at a constant markup over their marginal cost:

$$p_i = \frac{\sigma}{\sigma - 1} c_i$$

where $c_i$ is the marginal cost, gross of the shadow capital wedge

$$c_i = \frac{1}{z_i} \left( \frac{re^{\tau_i}}{\alpha} \right)^{\alpha} \left( \frac{w}{1-\alpha} \right)^{1-\alpha}$$

The labor requirement is:

$$\begin{aligned} \ell_i &= \left( \frac{1-\alpha}{\alpha} \cdot \frac{re^{\tau_i}}{w} \right)^{\alpha} \frac{y_i}{z_i} = \left( \frac{1-\alpha}{\alpha} \cdot \frac{re^{\tau_i}}{w} \right)^{\alpha} \frac{1}{z_i} p_i^{-\sigma} \left( \frac{Y}{P^{-\sigma}} \right) \\ &= \left( \frac{1-\alpha}{\alpha} \cdot \frac{re^{\tau_i}}{w} \right)^{\alpha} \frac{1}{z_i} \left[ \frac{\sigma}{\sigma-1} \cdot \frac{1}{z_i} \left( \frac{re^{\tau_i}}{\alpha} \right)^{\alpha} \left( \frac{w}{1-\alpha} \right)^{1-\alpha} \right]^{-\sigma} \left( \frac{Y}{P^{-\sigma}} \right) \end{aligned} \tag{2.3}$$

because the aggregate labor supply $L = \sum_i \ell_i$ is fixed we can drop the constant terms when computing labor shares:

$$\frac{\ell_i}{L} \propto \left( \frac{z_i}{e^{\alpha \tau_i}} \right)^{\sigma - 1}$$

Then, normalizing the labor force to one ($L = 1$) we have the following firm-level labor demand:

$$\ell_i = \frac{z_i^{\sigma-1} e^{\alpha(1-\sigma)\tau_i}}{\int_0^1 z_i^{\sigma-1} e^{\alpha(1-\sigma)\tau_i} \, \mathrm{d}i}$$

We therefore obtain the following expression for the steady-state equilibrium output of firm $i$ :

$$y_i = \left( \frac{w}{r} \right)^{\alpha} \frac{z_i^{\sigma} e^{-\alpha\sigma\tau_i}}{\int_0^1 z_i^{\sigma-1} e^{\alpha(1-\sigma)\tau_i} \, \mathrm{d}i}$$

We aggregate output across firms to obtain GDP:

$$Y = \left( \frac{w}{r} \right)^{\alpha} \left[ \int_0^1 z_i^{\sigma-1} e^{\alpha(1-\sigma)\tau_i} \, \mathrm{d}i \right]^{\frac{1}{\sigma-1}} \tag{2.4}$$

To find the steady-state factor price ratio $(w/r)$ we solve for the CES price index (the GDP deflator):

$$\begin{aligned} P &= \frac{\sigma}{\sigma-1}\left(\int_0^1 c_i^{1-\sigma}\right)^{\frac{1}{1-\sigma}} \mathrm{d}i \\ &= \frac{\sigma}{\sigma-1}\cdot\left(\frac{r}{\alpha}\right)^{\alpha}\left(\frac{w}{1-\alpha}\right)^{1-\alpha}\left(\int_0^1 z_i^{\sigma-1}e^{\alpha(1-\sigma)\tau_i}\,\mathrm{d}i\right)^{\frac{1}{1-\sigma}} \end{aligned} \tag{2.5}$$

Multiplying each side of this equation by the respective sides of equation (2.1) and rearranging we find:

$$\kappa\left(\frac{w}{r}\right)^{1-\alpha} = \left(\int_0^1 z_i^{\sigma-1}e^{\alpha(1-\sigma)\tau_i}\,\mathrm{d}i\right)^{\frac{1}{\sigma-1}}$$

where we define the constant $\kappa$ as:

$$\kappa \overset{\text{def}}{=} \frac{\sigma}{\sigma-1}\left(\frac{1}{\alpha}\right)^{\alpha}\left(\frac{1}{1-\alpha}\right)^{1-\alpha}\left(\frac{1}{\beta}-\delta\right)$$

Plugging inside the steady-state GDP equation (2.4) we obtain:

$$Y = \kappa^{\frac{\alpha}{\alpha-1}}\left[\int_0^1 z_i^{\sigma-1}e^{\alpha(1-\sigma)\tau_i}\,\mathrm{d}i\right]^{\frac{1}{(\sigma-1)(1-\alpha)}}$$

Notice that we can re-write the term in parentheses as an expectation:

$$Y = \kappa^{\frac{\alpha}{\alpha-1}}\left[\mathbb{E}\left(z_i^{\sigma-1}e^{\alpha(1-\sigma)\tau_i}\right)\right]^{\frac{1}{(\sigma-1)(1-\alpha)}}$$

Following HK, we make the assumption that $\tau_i$ is statistically independent of productivity $(z_i \perp \tau_i)$. This allows us to separate expectations:

$$Y = \kappa^{\frac{\alpha}{\alpha-1}}\left[\mathbb{E}\left(z_i^{\sigma-1}\right)\right]^{\frac{1}{(\sigma-1)(1-\alpha)}}\cdot\left[\mathbb{E}\left(e^{\alpha(1-\sigma)\tau_i}\right)\right]^{\frac{1}{(\sigma-1)(1-\alpha)}} \tag{2.6}$$

To solve the second expectation explicitly we need to make a parametric assumption on the distribution of $\tau_i$. We assume that $\tau_i$ is exponentially distributed. This is equivalent to saying that $\exp(\tau_i)$ follows a Pareto distribution:

$$\tau_i \sim \text{Exp}(\lambda) \quad\Longleftrightarrow\quad e^{\tau_i} \sim \text{Pareto}(\lambda, 1)$$

where $\lambda$ is the shape parameter. In equation (2.6), the second expectation in square brackets is then simply the $\alpha(1-\sigma)^{th}$ moment of a Pareto-distributed variable, and has a known closed-form solution. Under our parametric assumption, we therefore have the following intuitive expression for aggregate output (per employee):

$$Y = \underbrace{\kappa^{\frac{\alpha}{\alpha-1}} \left[\mathbb{E}\left(z_i^{\sigma-1}\right)\right]^{\frac{1}{(\sigma-1)(1-\alpha)}}}_{\substack{\text{Undistorted GDP} \\ \text{(per employee)}}} \cdot \underbrace{\left[\frac{\lambda}{\lambda+\alpha(\sigma-1)}\right]^{\frac{1}{(\sigma-1)(1-\alpha)}}}_{\substack{\text{Cost of Red Tape} \\ \text{(between 0 and 1)}}} \tag{2.7}$$

The equation above implies that, given a parametrization for the triple $(\sigma, \alpha, \lambda)$, we can compute the percentage loss of GDP attributable to red tape. While the literature offers very clear guidance on how to parametrize $\sigma$ and $\alpha$, the obvious challenge is to obtain country-level values of $\lambda$. This is the focus of the next section, where firm-level data becomes essential to our empirical exercise.

There are two mechanisms contributing to lowering aggregate output: the first is a lower savings rate, as bureaucracy lowers the ex-post return on capital. The second effect is a loss in aggregate productivity due to capital misallocation, which is induced by heterogeneity in bureaucracy wedges (Hsieh and Klenow, 2009a).

We can isolate the "misallocation" channel by first computing the firm-level demand for capital:

$$k_i = \frac{\alpha}{1-\alpha} \cdot \left(\frac{w}{r}\right) \cdot \frac{z_i^{\sigma-1}\, e^{[\alpha(1-\sigma)-1]\tau_i}}{\int_0^1 z_i^{\sigma-1}\, e^{\alpha(1-\sigma)\tau_i}\, \mathrm{d}i}$$

We then aggregate and separate expectations to obtain the steady-state supply of capital:

$$K \propto \frac{\alpha}{1-\alpha} \cdot \left(\frac{w}{r}\right) \cdot \frac{\mathbb{E}\left\{e^{[\alpha(1-\sigma)-1]\tau_i}\right\}}{\mathbb{E}\left\{e^{\alpha(1-\sigma)\tau_i}\right\}}$$

We finally define Total Factor Productivity (TFP):

$$\text{TFP} \overset{\text{def}}{=} \frac{Y}{K^{\alpha}L^{1-\alpha}} \propto \left\{\mathbb{E}\left[e^{\alpha(1-\sigma)\tau_i}\right]\right\}^{\frac{1}{\sigma-1}} \frac{\left\{\mathbb{E}\left[e^{\alpha(1-\sigma)\tau_i}\right]\right\}^{\alpha}}{\left\{\mathbb{E}\left[e^{[\alpha(1-\sigma)-1]\tau_i}\right]\right\}^{\alpha}} = \frac{\left[\frac{\lambda}{\lambda+\alpha(\sigma-1)}\right]^{\frac{1+\alpha(\sigma-1)}{\sigma-1}}}{\left[\frac{\lambda}{\lambda+1+\alpha(\sigma-1)}\right]^{\alpha}} \tag{2.8}$$

the latter term is the loss in TFP induced by red tape as a function of $\lambda$, which isolates the capital misallocation effect of bureaucracy on aggregate output. This TFP loss can be computed separately given an estimate of $\lambda$.

## 2.3 Survey Data and Identification

In this section, we present our strategy to recover the parameter $\lambda$, which determines the cross-sectional distribution of the red tape wedges $\tau$. Central to our approach is the use of firm-level survey data, which we combine with balance sheet information to reveal the cross-firm variation in $\tau$.

In particular, suppose that there are a number of different countries $c$, and each is characterized by a different wedge distribution parameter $\lambda_c$. Firm $i \in c$ is asked to report the burden imposed by red tape in an indicator variable $D_i$. We assume that this variable $D_i$ takes value 1 if the shadow tax of red tape $\tau_i$ overcomes a certain positive reporting threshold $T$ – formally:

$$D_i \stackrel{\text{def}}{=} \mathbb{I}\left\{\tau_i > T\right\}$$

A key assumption that we make here is that this threshold is uniform across countries, so that the "survey" dummy variable $D_i$ is comparable across countries.

Let $\pi_c$ be the percentage of firms in country $c$ that report bureaucracy as a constraint to growth:

$$\pi_c \stackrel{\text{def}}{=} \mathbb{P}\left(D_i = 1 \middle| \, i \in c\right)$$

Figure (2.2) outlines this setup visually.

Then, assuming that the wedges are exponentially distributed with shape parameter $\lambda_c$, we have

$$\pi_c \;=\; \exp\left(-\lambda_c T\right)$$

hence we identify the product $\lambda_c T$ as:

$$\lambda_c T \;=\; -\log\left(\pi_c\right)$$

This implies that, conditional on identifying the threshold parameter $T$, $\pi_c$ is a sufficient statistic for $\lambda_c$.

To identify the threshold parameter $T$, we use firm $i$'s first order condition for capital (equation 2.2), which we augment with an error term to allow for measurement error, in order to capture unobservable variation in MRPK that is unrelated to red tape which we are not modelling explicitly:

$$\log \widehat{\text{MPRK}}_i \;=\; \log r + \tau_i + \varepsilon_i$$

where $\widehat{\mathrm{MPRK}}_i$ is the marginal revenue product of capital as measured in the data.

We use an econometric approach because, while it is clearly impossible for our theoretical model to explicitly capture all potential sources of MRPK dispersion, we want to refrain from attributing all observed variation in MRPK to red tape.

Define, for a generic variable $v_i$, the following difference in conditional expectations:

$$\beta_c \overset{\text{def}}{=} \mathbb{E}\left(\log \widehat{\mathrm{MPRK}}_i \middle| D_i = 1, i \in c\right) \text{-}\mathbb{E}\left(\log \widehat{\mathrm{MPRK}}_i \middle| D_i = 0, i \in c\right)$$

furthermore, assume:

$$\varepsilon_i \perp D_i$$

Then, given our previous assumption that $\tau_i$ follows an exponential distribution, the conditional "gap" in MRPK between firms that report $D_i = 1$ and those that report $D_0 = 0$, is equal to:

$$\begin{aligned}
\beta_c &= \mathbb{E}\left(\tau_i \middle| \tau_i > T \text{ and } i \in c\right) - \mathbb{E}\left(\tau_i \middle| \tau_i < T \text{ and } i \in c\right) \\
&= \frac{T \cdot \exp\left(\lambda_c T\right)}{\exp\left(\lambda_c T\right) - 1} = \frac{T}{1 - \pi_c}
\end{aligned} \tag{2.9}$$

The intuition is that – conditional on firms reporting truthfully – firms that report being constrained by bureaucracy ($\tau_i > 0$) should display a higher marginal revenue product of capital. The width of this gap depends on 1) the reporting threshold $T$; 2) the shape parameter $\lambda_c$, which in turn affects the reporting frequency $\pi_c$.

To estimate the reporting threshold $T$, we run a regression of $\log \mathrm{MRPK}_i$ on the dummy $D_i$, pooling data from different countries. The resulting slope coefficient $\beta$ will be equal to the weighted average of this value across countries:

$$\beta = \sum_c \omega_c \cdot \beta_c = \sum_c \omega_c \cdot \frac{T}{1 - \pi_c}$$

where $\omega_c$ is the share of firms in country $c$:

$$\omega_c = \mathbb{P}\left(i \in c\right)$$

Our key identifying assumption that the reporting threshold is constant across countries allows us to take $T$ out of the summation.

$$\beta = T \cdot \sum_c \frac{\omega_c}{1 - \pi_c}$$

Rearranging this equation allows us to finally identify the reporting threshold $T$:

$$T = \frac{\beta}{\sum_c \omega_c (1-\pi_c)^{-1}} \tag{2.10}$$

Once $T$ is identified and $\alpha$ and $\sigma$ are calibrated, the reporting frequency $\pi_c$ becomes a sufficient statistic for the GDP loss due to red tape in country $c$. The plan for our empirical analysis is therefore to : 1) Regress $\log \mathrm{MRPK}_i$ on $D_i$ to estimate $\beta$; 2) Use equation (2.10) to obtain $T$; 3) Use $\pi_c$ and $T$ to obtain estimates of $\lambda_c$; 4) Compute reallocation gains using equations (2.7) and (2.8).

Next, we discuss the data we use for our empirical analysis.

## 2.4 Data

### 2.4.1 Firm-level Data: EFIGE

To measure the impact of bureaucracy on investment at the firm-level, we use the EU-EFIGE/Bruegel-UniCredit dataset, which is a firm-level database. The dataset contains data for a representative sample of 14,759 manufacturing firms from seven European countries (Austria, France, Germany, Hungary, Italy, Spain, UK).

The dataset is comprised of two parts. The first part is cross-sectional response data from the EFIGE executives survey, which was conducted by the think tank Bruegel in early 2010: firms were asked questions about a wide range of topics, including their organizational structure, ownership, workforce, international activities, and financing. The second part is a firm/year panel of firm financials (including turnover, assets, interest expenditure, profit and labor costs) for the period 2001-2014 merged from the Amadeus dataset, by the Bureau van Dijk.

We use a dummy variable that encodes the firms' answer to a specific question from the EFIGE survey. The specific question is:

*E6. Indicate the main factors preventing the growth of your firm:*

- ☐ *financial constraints*
- ☐ *labour market regulations*
- ☐ *legislative or bureaucratic restrictions*
- ☐ *lack of management and/or organizational resources*
- ☐ *lack of demand*
- ☐ *other*

This is a multiple-choice question, and our main explanatory variable, the dummy variable *Red Tape.* equals one if the firm corresponding to the observation ticked answer three.

We also encode firms ticking answer two as an additional control variable, which we call *Labor Regulations.*

Because the survey asked firms about their activities in 2009, we use 2009 balance sheet data for our firm-level analysis. However, given that 2009 was a recession year, we also check that our main econometric results carry through when we use 2008 data.

The survey portion of the EFIGE dataset comes with sampling weights to ensure the representativeness of the survey sample. Weighting ensures that the in-sample distribution of firms over industries and size classes matches the population's.

Unfortunately, while the weigths guarantee representativeness of the survey portion of the dataset, there is a well-known sample selection issue affecting the balance sheet portion of the dataset. The Amadeus database, which is the source of firm financials, has known issues of coverage and sample selection (Kalemli-Ozcan et al., 2015). Specifically, firm financials appear to be missing, for certain countries (Austria, Germany, the UK) in a non-random way.

We are able to address this issue thanks to the fact that the stratification variables (employment size and NACE 2-digit industry) belong to the survey part of the dataset, and are therefore available for all the firms in the sample, regardless of whether BvD financial data for the corresponding firm is available. This allowed us to devise our own weights, which are computed so that, after reweighteding our sample reflects the within-country distribution of the *population* of firms over ISIC rev. 3 broad sectors and employment size classes (10-19,20-49,50-250,250+). We are able to obtain the population distribution from the OECD Structural and Demographic Business Statistics (SDBS) dataset.

We use these weights in our robustness checks. Our baseline results hold when we use our weights that account for sample selection in BvD. In our analysis, we also address sample selection in with a second method, using the fact that France, Hungary, Italy and Spain are virtually free from the sample selection problem (coverage in these countries is nearly 100%). Our baseline estimates also hold when we exclude Austria, Germany and the UK from the Sample.

Our main dependent variable, the Marginal Revenue Product of Capital (MRPK), is computed as follows:

$$\widehat{\text{MRPK}}_i = \frac{\sigma - 1}{\sigma} \cdot \alpha \cdot \frac{\text{Value Added}_i}{\text{Fixed Assets}_i}$$

where Value Added can be computed as either revenues less intermediate input costs, or as the sum of capital and labor compensation (EBITDA+Labor Costs).

Following the literature, we calibrate $\sigma = 3$ and $\alpha = 1/3$. We use the *log* of MRPK as a dependent variable; hence, as long as $\sigma$ and $\alpha$ do not vary sistematically within countries/sectors these calibrated values have no effect on our empirical estimates of $\beta$, because they are absorbed by the regression fixed effects.

### 2.4.2 Country-level Data

In order to validate the EFIGE survey data, as well as to extend our analysis to a larger set of countries that are not covered in the EFIGE dataset, we use the international data on regulatory barriers to entry, compiled by DLLS. The data set covers 85 countries and quantifies the difficulty of forming new firms in each country.

We use this dataset because it makes a legitimate attempt to measure red tape across countries with an objective, comparable methodology. The downside is that these variables are not a perfect match for our model and firm-level data, as neither the model nor the EFIGE survey focus specifically on new entrants. In order to work with this dataset, we shall make the explicit assumption that countries that impose more severe constraints on new entrants *also* impose more severe constraints on incumbents. Under this assumption, DLLS's dataset can still serve as a useful proxy. A different way to say this is that we shall use country-level entry regulations as a proxy for business regulations in a more general sense.

We focus on the three measures introduced in DLLS's dataset: (i) the number of procedures, (ii) screening time (measured in steps to register) , and (iii) registration cost (measured in days). Red tape can potentially affect all three of these measures, and so we combine the three measures into a single Entry Regulations Index for our study.

To construct the Entry Regulations Index, we take the first principal component of the logarithms of the three measures. The principal component is a natural choice of methodology because it captures common variation across the three factors. In this way, we can be agnositc as to whether Red tape disproportionately specific measures, and instead focus on the cross-country variation in bureaucratic constraints on entry. Principal component analysis is designed to be done on unbounded variables, whereas the measures in DLLS are positive by nature of their construction. Thus, we take logarithms of the three measures to make the data amenable to our analysis. The three measures are highly correlated within-country, so that first principal component captures over 90 percent of the cross-country variation in the data.

In Figure 2.3 compare, for the 7 countries included in EFIGE, the percentage of firms reporting Bureaucracy as a significant constraint to firm growth, with the corresponding Entry Regulations Index from DLLS's dataset. The number of observations in the plot is small, but the correlation between these two variables is nearly perfect (the $R^2$ is 98.5%). This is important for two reasons: first, it is reassuring that EFIGE's survey data correlates with an objective, well-established statistic; second, this tight relationship allows us to predict $\pi_c$ (the percentage of firms reporting bureaucracy as a constraint) out-of-sample, and this allows us to compute the GDP losses due to red tape not only for the 7 countries in EFIGE, but for all 85 countries included in DLLS.

## 2.5 Empirical Results

### 2.5.1 Firm-Level results within the EFIGE sample

For our sample of European firms in the EFIGE sample, we demonstrate that, consistent with our model, firms which report being constrained by bureaucratic red tape in the EFIGE survey also exhibit higher average Marginal Revenue Product of Capital (MRPK) in the BvD financials database. We show this visually in Figure 2.4, which plots density estimates of log MRPK conditional on firms' survey responses. The density curve corresponding to constrained firms is plotted in a dashed line, while the unconstrained firms' density curve lies beneath the shaded area in the figure. Visual inspection suggests that the MRPKs of the constrained firms are higher, and we quantify this with regression evidence in Table 2.1. Our specification is

$$\log \text{MRPK}_i \;=\; \gamma_c + \varsigma_s + \text{Red Tape}_i \beta_1 \cdot + \mathbf{x}_i \beta_2 + \varepsilon_i$$

where $\text{Red Tape}_i$ is the EFIGE survey indicator for bureaucratic constraints, and $\mathbf{x}_i$ is a vector of firm $i$, sector $s$, and country $c$ characteristics, which we describe in more detail for each regression. In Column (1), we report the regression results absent any controls, and find a statistically significant positive coefficient. When we augment the regression to include country and sector fixed effects, shown in Columns (2) through (4), the sign of the coefficient remains positive but the magnitude decreases to roughly 0.07. This implies that the MRPK of a constrained firms is 7 percent higher than that of an otherwise similar unconstrained firm. This magnitude remains stable controlling for addition firm characteristics, such as Firm Age in column (5).

Because red tape in the general sense is also likely to correlate with labor regulations, we also want to make sure that what we are not inadvertedly capturing labor regulations; for this reasons, in column (6) we add, as a control, the dummy variable *Labor Regulations*, which indicates whether the same firm indicated being constrained by labor regulations in the same survey question.

Using our model, we can use our regression coefficient estimates, along with country-level data on reporting frequencies, to estimate the $T$ parameter from our model using

Equation (2.10). We use our point estimate of $\beta = 0.077$, which implies a value for $T$ of 0.055.

### 2.5.2 Computation of GDP loss

Table 2.2 presents our main empirical results. To extend our analysis to the broader counter sample, we estimate a logit model to relate the DLLS Entry Regulations Index to a country-level reporting frequency $\pi_c$. The Entry Regulations Index is the first principal component of three measures capturing the number of procedures, screening time, and registration cost that a start-up must bear before it can operate legally. The first principal component loads equally on all three measures and captures over ninety percent of the cross-sectional variation in the difficult of entry. As seen in Figure 2.3, the relationship between the Entry Regulations Index and $\pi_c$ is close to logarithmic for low levels of $\pi_c$. Therefore, the logit model is a natural choice for this setting because of it's continuous support as well as its bounded range. Given that the EFIGE sample was conducted in countries with relatively low values of this Index, we take care to ensure that the inferred reporting frequencies remain bounded, even for countries with high Entry Regulation Indices. Our specification is:

$$\log \frac{\pi_c}{1-\pi_c} = a + b \, \text{Entry Regulation}_c + \varepsilon_c$$

where $\text{Entry Regulation}_c$ denotes the Entry Regulations Index extracted from Djankov et al. (2002). We first estimate estimate parameters $a$ and $b$ in the subset of countries in the EFIGE survey dataset and then use those fitted parameters to impute country-level reporting frequencies for the larger DLLS sample.

Figure 2.5 plots results of this logit model. In the left panel, we plot the empirical values of the Index against the reporting frequency $\pi_c$, as well as the predictions of the estimated model. In the right panel, we illustrate where the out-of-sample countries lie along this fitted logit curve. Using the fitted values of reporting frequencies $\hat{\pi}_c$, we can compute country-level parameter $\beta_c$ using only our estimate of $T$ obtained from the micro-level evidence in the EFIGE plus BvD sample. For each country in the DLLS sample, we present estimated output gains, computed using Equation 2.7, as well as the output gains arising solely from improvements in aggregate productivity, computed using Equation (2.8). Totaling up our estimates, we estimate that global output could be increased by over \$2.1 trillion through the removal of these taxes. Furthermore, we see that, while large percentage gains accrue to relatively countries with high DLLS Entry Regulations Indices, the bulk of the gains in global GDP arise from large, moderately constrained economies improving their institutions. Global output would increase by over \$1 trillion if just two countries, China and the Russian Federation, improved their economies by eliminating bureaucratic regulations affecting capital. Furthermore, we estimate global output would increase by roughly \$1.4 trillion if all countries, rather than eliminating entry regulations completely, reduced them to a level comparable to that of the United States.

By comparing the overall gain in GDP to the gains arising from increases in aggregate productivity, we see that a relatively small, but still economically large, portion of GDP gains are attributable to improving the allocation of resources. For example, roughly one quarter of th GDP gains we estimate for the Russian Federation stem from reductions in capital misallocation, corresponding to an increase in annual GDP of approximately 4.5%, or \$160 billion. As seen from Equations (2.7) and (2.8), the decomposition of outgains into gains stemming from capital accumulation versus gains from reallocation are non-linear, and larger gains of reallocation accrue to countries with larger amounts of red tape. For example, in the case of Portugal, the relative share of reallocation gains is an order of magnitude smaller than for the Russian Federation.

### 2.5.3 Robustness and Sensitivity Analysis

We next discuss several factors that could affect the reliability of our estimates and our strategy to address this robustness concerns. We also discuss how sensitive our estimates of the loss in GDP are to variations in our estimates of the coefficient $\beta$, which summarizes the firm-level impact of red tape on MRPK.

The first source of concern is sample selection. In particular, the balance sheet portion of our firm data is sourced from Bureau van Dijk (BvD). This dataset has known issues of sample selection for a number of countries. Some of them (Austria, Germany, UK) are included in the EFIGE survey. If selection in the sample occurs on variables that correlate with Bureaucracy and such variables and the correlation between MRPK and Bureaucracy varies conditionally on these variables, our estimate for $\beta$ will be biased.

We can show our estimates are robust to controlling for selection on size and sector. This is because we observe the distribution of firms for sector and employment class variables in the OECD SDBS dataset. In Online Appendix A, we discuss how we use this data to create a weghting scheme that we use in Online Appendix B to run a weighted regression that controls for sample selection, ensuring that the analysis sample is representative along these two dimensions.In this specification the estimate for $\beta$ is shown to increase only slightly.

In order to address more broadly selection on unobservables, we also use a more aggressive specification in which we eliminate from our sample all firms from the three countries

(Austria, Germany and UK) that are affected by the sample selection problem. In this specification, the estimate for $\beta$ is only slightly smaller.

Another shortcoming of our sample is that the survey was run at the beginning of 2010, following what was a major recession year for the EU. We we do not use time-variation in our analysis to identify $\beta$, our estimates might not be representative of an expansion period if there is a cyclical component to the correlation between MRPK and Red Tape. We therefore run a specification in which we compute MRPK using data from 2008: this is the year in which demand peaked across the EU (the recession only reached the Europe from the US in early 2009). While we do expect a lower estimates due to the fact that we are using data from a year that is farther removed in time from the survey data, the estimate from this specification (0.066) is only slightly different from our benchmark estimate, showing that sample selection based on time also does not seem to be driving our estimates.

The next issue we want to address is that of production function mis-specification. Our computation of MRPK is based on the assumption of a Value Added based production function (i.e. value added is the measure of nominal output). To ensure that the correlation between MRPK and Red Tape is not sensitive to this assumption, we compute MRPK for an output-based production function (where revenues are the numerator). We re-run our benchmark regression, and again show that our estimate or $\beta$ is only marginally affected.

The final conceptual issue we want to address is the possibility that our general equilibrium model itself might be mis-specified. We wonder how our estimates might change if Red Tape was not be best modeled as a tax on capital, but rather on output. If this was the case, the effect of the shadow tax would *not* only manifest on MRPK, but also of markups. In order to account for this possibility, we compute markups using the method of De Loecker and Warzynski (2012) and we re-run our main regression analysis using markups as the left-hand side variables in place of MRPK.[2] We show that Red Tape does not shift the conditional distribution of markups significantly, which demonstrates how the effect of red tape is likely more accurately modeled as a shadow tax on capital rather than as a shadow tax on output.

Additionally, we present estimates computed under alternative values of $\beta$ in Online Appendix C. Rather than use an intermediate value of $\beta$ from Table 2.1, we use a higher and lower value and repeat the procedure of estimating output and TFP. When we assume a $\beta$ of 0.10, corresponding to a shadow tax of 10% between the average constrained and unconstrained firm, our estimate of global output gains increases from \$2 trillion to \$2.6

[2] This specification accounting for the critique of (Bond, Hashemi, Kaplan and Zoch, 2020).

trillion. Symmetrically, a value 0.06 for $\beta$ would lead us to predict output gains of \$1.7 trillion. We emphasize two takeaways from this robustness analysis. Firstly, the gains of eliminating red tape and over-regulation are significant even if shadow taxes are twenty percent lower than our preferred point estimate. Secondly, the gains are roughly unit elastic in the shadow tax, so that a twenty percent change results in approximately a twnety percent change in the expected output gains.

Finally, we present estimates computed using an alternative data source in Online Appendix C. We make use of the World Bank Enterprise Survey, which contains establishment level survey and financial data. We make use of survey questions "j.30c" and "l.30a", which solicit the degree to which licensing and labor regulations, respectively, impede business operations. The survey allows respondents to rank the severity of the impediment on a scale from 0 through 4. We define dummy variables analogous to those in the EFIGE dataset by treating survey responses of 3 and 4, denoting "Major obstacle" and "Very severe obstacle", as indicating that the firm is contrained by bureaucratic and labor regulations. Estimates derived from World Bank data are not directly comparable to those derived from EFIGE data due to inconsistencies between the two surveys. For one, the World Bank survey quesiton that we use to construct our indicator of bureaucratic constraints asks only about licensing regulations, rather than being an all-inclusive category for bureaucratic regulation as in the EFIGE survey. Another difference is that the World Bank data survey structure differs from that of EFIGE. In the World Bank survey, firms are sequentially asked about a list of twenty potential factors impeding business operations. In the EFIGE survey, firms are asked to name major factors and the indicator variables denote identified contraints. Thus, there is a methodological distinction between prompted and unprompted identification of institutional constraints. Despite these differences, we find that the economic costs of red tape are still economically significant when estimated using this alternative dataset.

## 2.6 Conclusions

We study the impact of bureaucratic regulations impeding the allocation of and investment in capital. We introduce a model featuring heterogeneous firms constrained by firm-specific taxes on capital, and compute the aggregate impact of these taxes on both output as well as total factor productivity. We then take our model to the data by linking firm-level survey data to financial performance in the combined Amadeus-EFIGE dataset. This enables us to identify which firms are relatively constrained, as well as observe the impact of these constraints on their performance.

A key advantage of our approach is that we can then take our estimates from the

Amadeus-EFIGE dataset and extend them to a large international panel of 85 countries using the DLLS Entry Regulations Index. We find that cross-country variation in the DLLS Index is closely related to cross-country heterogeneity in the reporting of bureaucratic constraints. We map the DLLS Index to estimated output and productivity gains using our Amadeus-EFIGE sample estimates and find that global GDP would be over $2 trillion higher absent these bureaucratic regulations. Our model allows us to decompose the increase in GDP into two distinct channels: (i) the effects of improved allocation of resources, and (ii) the effects of increased capital investment. We find that both effects are significant and that the majority of the gains arise through the latter channel.

Our results highlight the costs of over-regulation, as well as the importance of institutions for growth. In the cross section of nations, we identify a number of countries which can significantly improve their annual economic output by curtailing excessive regulation on entry. We hope that these estimates guide future policy work into ameliorating burdens on new entrants.

Finally, our goal in this study is to focus on and quantify the effects of a specific friction, bureaucratic restrictions on capital. Due to this specificity, we use only variation in MRPK attributable to these frictions when extrapolating from the EFIGE sample estimated parameters to the larger DLLS sample. A number of other constraints have been identified in the literature, and we hope that future work extends our analysis with new data to address the impact of these on the global economy.

**Figure 2.2:** Graphic Illustration of the Firms' Survey Reporting Decision

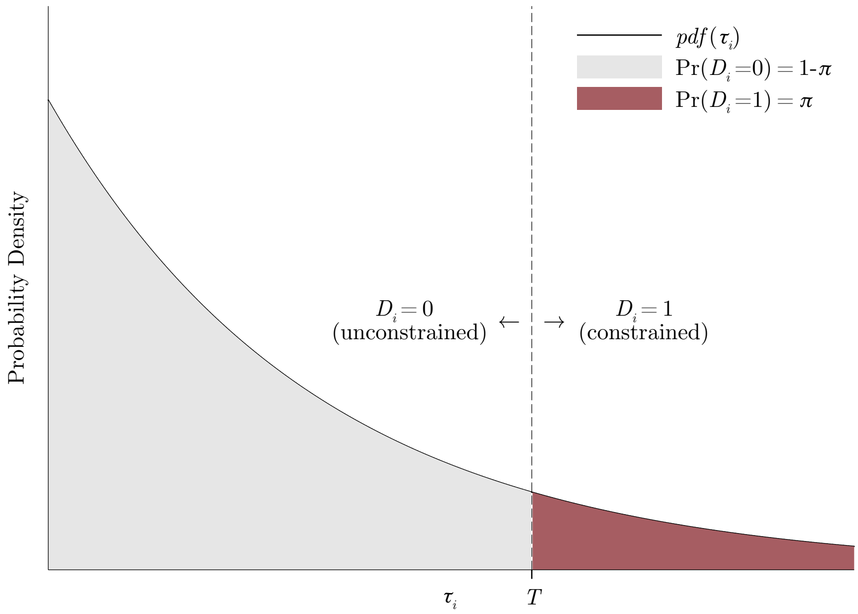


FIGURE NOTES: The above diagram exemplifies firm $i$'s survey reporting decision.

**Figure 2.3:** Regulations and Survey Data

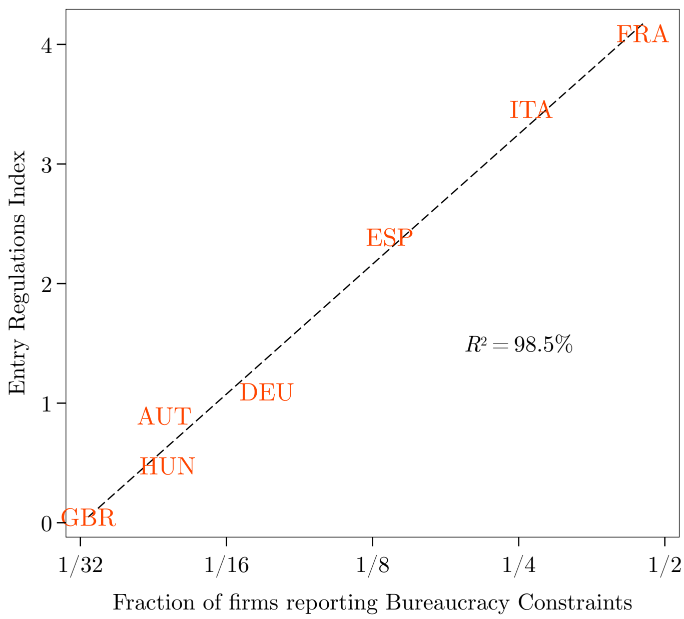


FIGURE NOTES: The figure above plots the probability of a firm reporting bureaucracy as a constraint in the EFIGE survey, by country, against an index of regulatory burden computed from the dataset of Djankov, La Porta, Lopez-de Silanes and Shleifer (2002). The dotted line is a fitted regression line.

**Figure 2.4:** Conditional Distribution of MRPK

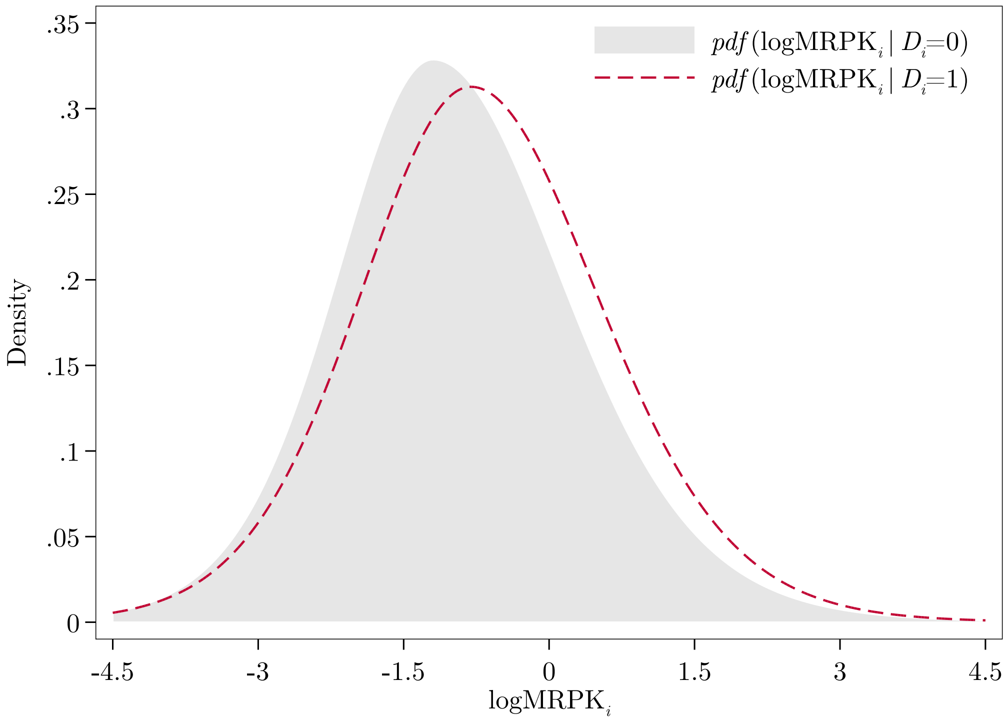


Figure Notes: The figure above plots parametric estimates of the conditional density of the log of Marginal Revenue Product of Capital (MRPK), conditional on the value of the survey dummy in which firms may report bureaucracy as a significant constraint to growth. The grey area is the estimated density for firms that do *not* report bureaucracy as a constraint to growth ($D_i = 0$). The dotted dark line is the estimated density for firms that *do* report bureaucracy as a constraint to growth ($D_i = 1$). The fitted distribution is a 4-parameter Skewed T distribution, which allows for skewness and fat tails.

**Table 2.1:** Marginal Revenue Product of Capital and Red Tape: Regression Analysis

Dependent Variable: $\log \text{MRPK}$

| | (1) | (2) | (3) | (4) | (5) | (6) |
|---|---|---|---|---|---|---|
| Red Tape | 0.250*** | 0.266*** | 0.059* | 0.077** | 0.080** | 0.065* |
| | (0.036) | (0.035) | (0.035) | (0.034) | (0.034) | (0.038) |
| log Age | | | | | - 0.124*** | |
| | | | | | (0.020) | |
| Labor Regulations | | | | | | 0.029 |
| | | | | | | (0.037) |
| $R^2$ | 0.007 | 0.058 | 0.153 | 0.206 | 0.212 | 0.207 |
| Observations | 6,895 | 6,895 | 6,895 | 6,895 | 6,875 | 6,895 |
| Country Fixed Effects | | ✓ | | ✓ | ✓ | ✓ |
| Sector Fixed Effects | | | ✓ | ✓ | ✓ | ✓ |

Table Notes: The table above presents Ordinary Least Squares estimates for the following linear regression model:

$$\log \text{MRPK}_i \;=\; \gamma_c + \varsigma_s + \text{Red Tape}_i\, \beta_1 \cdot + \mathbf{x}_i \beta_2 + \varepsilon_i$$

Where $\text{MRPK}_i = \frac{\sigma-1}{\sigma}\alpha\frac{\text{Value Added}_i}{\text{Fixed Assets}_i}$ and $\alpha$ is calibrated to 1/3. $\text{Red Tape}_i$ is the EFIGE survey dummy in which firms indicate being constrained by red tape. $\gamma_c$ and$\varsigma_s$ are, respectively, country and sector fixed effects and $\mathbf{x}_i$ is a vector of control variables. Robust standard errors in parentheses: $^{*}p < .1; ^{**}p < .05; ^{***}p < .01$

**Figure 2.5:** Probability of Firm Reporting Bureaucracy as a Constraint

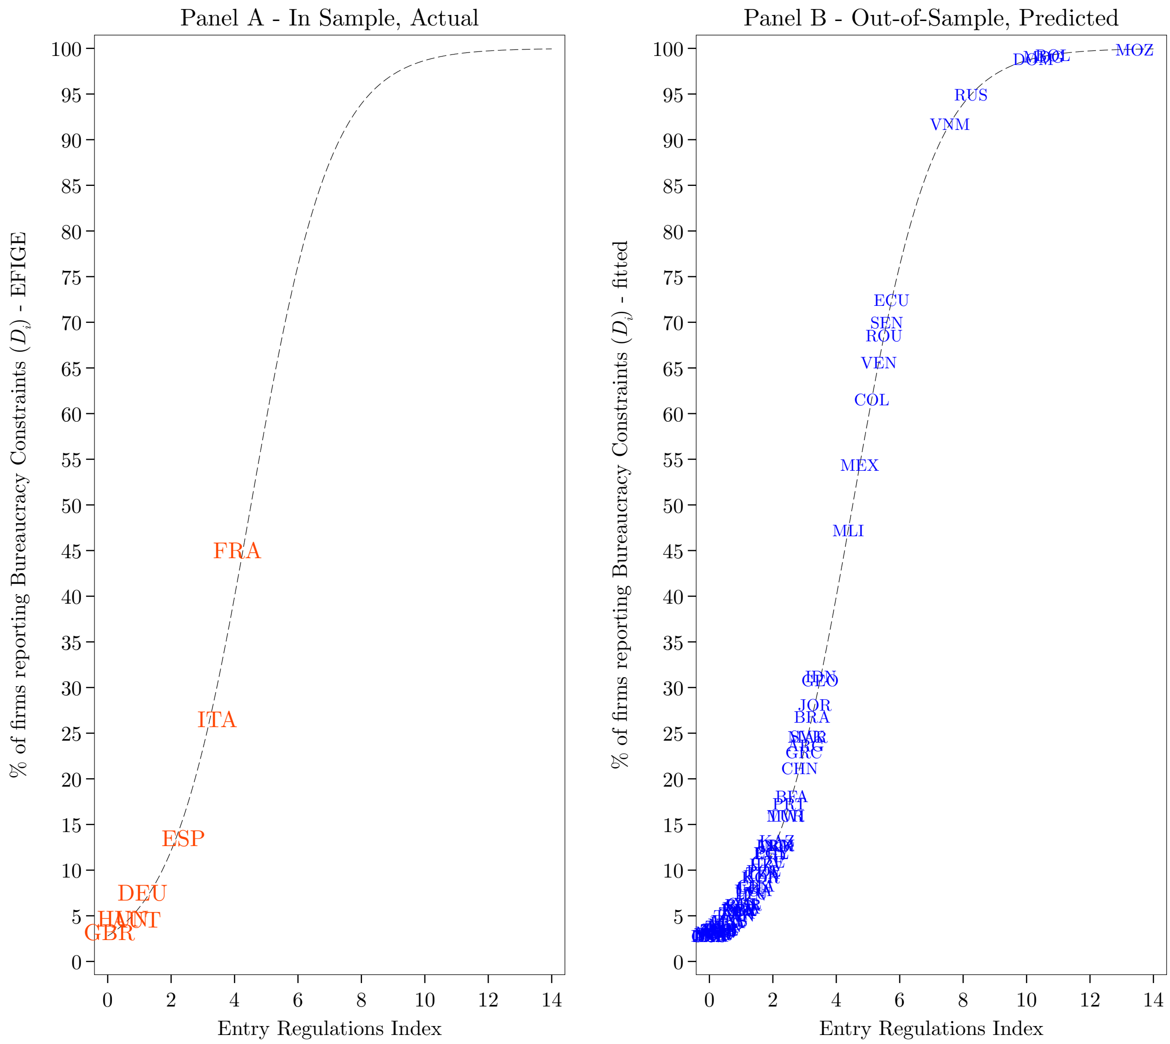


Figure Notes: The figure above plots the probability of a firm reporting bureaucracy as a constraint, by country, against an index of regulatory burden computed from the dataset of Djankov, La Porta, Lopez-de Silanes and Shleifer (2002, DLLS). The left panel shows the actual percentage of firms whose management indicates Bureaucracy as a major constraint to the firm's growth, as recorded by the EFIGE survey, on the left axis. The dotted line is a fitted logit curve. The right panel shows the probability predicted - out of sample - by the logit fit. These are the countries that are covered by the DLLS dataset but not in EFIGE.

**Table 2.2:** Estimated GDP losses from Red Tape

| Country | GDP Loss from Red Tape US$bln | %GDP | Δ%TFP |
|---|---|---|---|
| Russia | 760.8 | 21.33 | 4.52 |
| China | 244.1 | 1.74 | 0.03 |
| United States | 118.1 | 0.76 | 0.01 |
| Mexico | 78.1 | 4.22 | 0.19 |
| India | 77.1 | 1.31 | 0.02 |
| Vietnam | 73.4 | 17.18 | 2.98 |
| France | 73.0 | 3.03 | 0.10 |
| Brazil | 60.7 | 2.04 | 0.05 |
| Indonesia | 49.9 | 2.30 | 0.06 |
| Japan | 48.3 | 1.04 | 0.01 |
| Italy | 48.1 | 2.25 | 0.06 |
| Dominican Rep. | 34.6 | 29.68 | 8.61 |
| Germany | 34.4 | 0.99 | 0.01 |
| Venezuela | 30.4 | 5.82 | 0.36 |
| Colombia | 28.2 | 5.16 | 0.28 |
| Romania | 22.6 | 6.38 | 0.43 |
| Turkey | 21.7 | 1.48 | 0.02 |
| Spain | 21.5 | 1.47 | 0.02 |
| Korea, Rep. | 18.1 | 1.14 | 0.01 |
| United Kingdom | 17.2 | 0.77 | 0.01 |
| Bolivia | 16.6 | 30.97 | 9.36 |
| Argentina | 13.3 | 1.87 | 0.04 |
| Egypt, Arab Rep. | 11.5 | 1.28 | 0.02 |
| Canada | 11.1 | 0.76 | 0.01 |
| Ecuador | 11.0 | 7.25 | 0.56 |
| Poland | 10.2 | 1.18 | 0.02 |
| Madagascar | 10.1 | 30.43 | 9.04 |
| Thailand | 8.3 | 0.93 | 0.01 |
| Taiwan, China | 8.1 | 0.91 | 0.01 |
| Nigeria | 8.1 | 0.96 | 0.01 |
| Australia | 7.9 | 0.76 | 0.01 |
| Mozambique | 7.6 | 32.93 | 10.56 |
| Pakistan | 7.4 | 0.95 | 0.01 |
| Philippines | 7.0 | 1.27 | 0.02 |
| Netherlands | 6.4 | 0.85 | 0.01 |
| Ukraine | 5.9 | 1.32 | 0.02 |
| South Africa | 5.7 | 0.92 | 0.01 |
| Malaysia | 5.0 | 0.87 | 0.01 |
| Greece | 4.9 | 1.83 | 0.04 |
| Kazakhstan | 4.8 | 1.34 | 0.02 |
| Morocco | 4.5 | 1.92 | 0.04 |
| Portugal | 4.1 | 1.54 | 0.03 |
| Switzerland | 3.9 | 0.81 | 0.01 |

| Country | GDP Loss from Red Tape US$bln | %GDP | Δ%TFP |
|---|---|---|---|
| | (*continued*) | | |
| Belgium | 3.7 | 0.88 | 0.01 |
| Czech Republic | 3.5 | 1.23 | 0.02 |
| Chile | 3.3 | 0.98 | 0.01 |
| Austria | 3.3 | 0.94 | 0.01 |
| Sweden | 3.1 | 0.78 | 0.01 |
| Norway | 3.1 | 0.78 | 0.01 |
| Peru | 2.9 | 0.95 | 0.01 |
| Singapore | 2.6 | 0.82 | 0.01 |
| Hong Kong SAR | 2.5 | 0.78 | 0.01 |
| Slovak Republic | 2.4 | 1.93 | 0.04 |
| Senegal | 2.4 | 6.68 | 0.47 |
| Denmark | 1.9 | 0.76 | 0.01 |
| Hungary | 1.8 | 0.85 | 0.01 |
| Israel | 1.8 | 0.80 | 0.01 |
| Ireland | 1.7 | 0.77 | 0.01 |
| Sri Lanka | 1.7 | 0.87 | 0.01 |
| Jordan | 1.7 | 2.12 | 0.05 |
| Finland | 1.6 | 0.78 | 0.01 |
| Kenya | 1.3 | 1.15 | 0.01 |
| Tanzania | 1.2 | 1.17 | 0.01 |
| Ghana | 1.2 | 1.09 | 0.01 |
| Bulgaria | 1.1 | 0.98 | 0.01 |
| New Zealand | 1.1 | 0.76 | 0.01 |
| Croatia | 1.0 | 1.21 | 0.02 |
| Mali | 1.0 | 3.49 | 0.13 |
| Tunisia | 0.9 | 0.84 | 0.01 |
| Georgia | 0.9 | 2.28 | 0.06 |
| Lebanon | 0.8 | 1.08 | 0.01 |
| Uganda | 0.7 | 1.07 | 0.01 |
| Lithuania | 0.7 | 1.05 | 0.01 |
| Uruguay | 0.6 | 0.95 | 0.01 |
| Slovenia | 0.5 | 0.97 | 0.01 |
| Panama | 0.5 | 0.82 | 0.01 |
| Zambia | 0.4 | 0.81 | 0.01 |
| Armenia | 0.3 | 1.32 | 0.02 |
| Burkina Faso | 0.3 | 1.58 | 0.03 |
| Latvia | 0.3 | 0.83 | 0.01 |
| Malawi | 0.2 | 1.48 | 0.02 |
| Kyrgyz Republic | 0.2 | 0.94 | 0.01 |
| Mongolia | 0.2 | 0.79 | 0.01 |
| Zimbabwe | 0.2 | 0.81 | 0.01 |
| Jamaica | 0.2 | 0.80 | 0.01 |

# Chapter 3

# Diagnosing the Italian Disease

Bruno Pellegrino, UCLA[1]
Luigi Zingales, University of Chicago

Italy's aggregate productivity abruptly stopped growing in the mid-1990s. This stop represents a puzzle, as it occurred at a time of stable macroeconomic conditions. In this paper, we investigate the possible causes of this "disease" by using sector and firm-level data. We find that Italy's productivity disease was most likely caused by the inability of Italian firms to take full advantage of the ICT revolution. While many institutional features can account for this failure, a prominent one is the lack of meritocracy in the selection and rewarding of managers. Unfortunately, we also find that the prevalence of loyalty-based management in Italy is not simply the result of a failure to adjust, but an optimal response to the Italian institutional environment. Italy's case suggests that familism and cronyism can be serious impediments to economic development even for a highly industrialized nation.

[1] Bruno Pellegrino gratefully acknowledges financial support from the Price Center for Entrepreneurship at the UCLA Anderson School of Management. Luigi Zingales gratefully acknowledges financial support from the Stigler Center at the University of Chicago Booth School of Business. We thank Carlo Altomonte and Tommaso Aquilante for providing us with additional data from the EFIGE dataset and for useful discussions. We are grateful to Andrei Shleifer, Raffaella Sadun, Dale Jorgenson, Paola Giuliano, Romain Wacziarg, Christian Dippel, Erik Brynjolfsson and seminar participants at Chicago Booth, UCLA Anderson, Harvard Business School, Berkeley Haas, and Purdue for their helpful comments and feedback.

## 3.1 Introduction

In the mid-1990s, Italy's labor productivity growth slowed down almost to a halt, while its total factor productivity (TFP) started to decline. While in most Western countries TFP growth slowed down after the Great Financial Crisis (Fernandez-Villaverde and Ohanian, 2018), Italy experienced a *decline* in TFP and experienced it before the Great Financial Crisis.

Italy's productivity growth disease is not only key to Italy's ability to sustain its debt and to remain a member of the Euro Area, it is also an economic puzzle, at odd with standard growth theory: how can firms unlearn what they have learned? A possible explanation is provided by Gopinath et al. (2017), who claim that it is the aggregate TFP that declined, not the firm-level one, due to a misallocation of investments. This explanation is very convincing for Spain, which experience a large increase in investments (gross fixed capital formation (GFCF) went from 21% of GDP in 1994 to 31% in 2007) and a healthy growth in labor productivity. It is less convincing for Italy, where corporate investments did not experience a similar growth and labor productivity barely increased.

A decline in TFP could be explained by some major institutional turmoil, but during the period 1996-2006, Italy benefited from the most stable economic and political environment since the early 1960s: low and stable interest rates, low and stable inflation, it even enjoyed the most long-lasting governments since WWII.

This is not to say that the Italian institutional environment was great. Italy lags behind other developed countries across many dimensions. While these deficiencies might be able to explain why it is less productive overall, they cannot easily account for the drop in TFP that occurred in the mid-90s, since these deficiencies were present in the 1950s and 1960s, when Italy was considered an economic miracle, and persisted in the 1970s and 1980s, when Italy continued to have GDP and productivity growth above the European average. In order to explain Italy's drop in TFP it is necessary to identify a significant deterioration of the institutional environment, or some institutional factor which did not matter before 1995, and then became a major driver of competitiveness in later years.

To find such a factor we resort to the existing literature. Italy's TFP started decreasing at the same time that US labor productivity started accelerating. In a seminal contribution, Bloom, Sadun and Reenen (2012a) have attributed the US's acceleration to the impact of the ICT revolution, which occurred in the same period. American firms were able to take advantage of ICTs thanks to meritocratic management practices, which have been shown to be strongly complementary with ICT capital (Bresnahan et al., 2002). We also know that Italy stands out, among developed countries, for the scarce diffusion of these practices: Bandiera et al. (2008) documented this fact extensively using a combination of survey and administrative data, and linked it with certain cultural traits that are particularly prevalent in Italy. Can the interaction of the ICT revolution and lack of meritocracy explain Italy's productivity decline?

The first challenge in testing this hypothesis is theoretical. Under the standard growth accounting framework (Jorgenson et al., 1987, 2005), a complementarity between managerial practices and ICT would impact only labor productivity, not TFP: any effect of the complementarity would be factored in the contribution of ICT capital to output growth, leaving TFP unaffected.

To address this challenge, we build a simple model to show that in the presence of spillovers from ICT investments, this conclusion does not hold and a complementarity between management and ICT impacts TFP growth. Most importantly, we show that this effect of the complementarity on TFP growth does not invalidate the rest of the growth accounting framework. Therefore, we can use the standard growth accounting framework developed in EU KLEMS data to test whether this complementary can explain Italy's TFP decline.

Consistent with our hypothesis, we find that TFP grew faster in more ICT-intensive sectors in countries where firms are more likely to select, promote, and reward people based on merit, as captured by a measure we derived from answers to the World Economic Forum (WEF) expert survey. This effect explains between 66 and 73% of Italy's TFP growth gap.

Since a country's level of meritocracy in the business sector is correlated with many other institutional characteristics (quality of government, ICT infrastructure, size of the shadow economy), aggregate data alone cannot rule out other possible interpretations. For this reason, we probe deeper with a firm-level dataset (the Bruegel-Unicredit EFIGE dataset). The advantage of this data is that – besides accounting information – it contains also a rich questionnaire shedding light on firms' managerial practices and ICT use. The main drawback is that it does not contain estimates of ICT capital investments, separately from the rest of the investments, as EU KLEMS does.

Using responses to the EFIGE questionnaire, we construct a firm-level measure of meritocratic management which reflects the firm's actual organizational practices. In constructing this index, we follow previous work by Bloom et al. (2012b) and Bandiera et al. (2008). The EFIGE data confirm the stylized fact that Italian firms are particularly likely to select and reward their managers based on loyalty and family considerations, rather than performance.

The firm-level data exhibits the same patterns as the KLEMS sectoral data: TFP grows faster in more meritocratic firms in sectors where the contribution of ICT investments to GDP growth is larger. This result holds after controlling for country and sector fixed effects.

Using EFIGE survey data, we can investigate directly whether the effect of meritocratic management on TFP growth is mediated by a more intensive use of ICTs. Consistent with Garicano and Heaton (2010), we find that more meritocratic firms indeed utilize ICTs more. As for TFP growth, this correlation is stronger in sectors where ICTs have the largest impact on output growth.

These findings raise a further question: Why does Italy lag behind in the adoption of meritocratic management practices? Bandiera et al. (2008) juxtapose meritocratic management to loyalty-based management, i.e., a system where managers are selected, rewarded, and promoted based on their loyalty to the firm's owner and/or their belonging to his extended family. This practice may not be just the result of lack of sophistication, but the optimal response to an environment where loyalty to the owner is more important than competency. In order to pay a bribe to a government official or to evade taxes without being caught, a company needs the loyalty of its top management. Among developed countries, Italy stands out for its level of corruption and tax evasion. Thus, a possible explanation is that, at the onset of the ICT revolution, Italy found itself with a managerial class that was perfectly suitable for its domestic environment, but incapable of taking full advantage of the newly available technologies.

To test this hypothesis, we exploit another feature of the EFIGE survey: firms are asked to indicate the main impediments to their growth. We look at three major sources of external constraints: access to finance, labor market regulation, and bureaucracy. In the overall sample, meritocratic firms are less likely to experience any of these constraints. Yet, this positive effect disappears for Italian firms. In particular, loyalty-based management seems to provide a relative advantage in overcoming financial and bureaucratic constraints in Italy. This result is consistent with judicial evidence of bribes paid to obtain credit and bypass bureaucratic constraints.[2]

We are certainly not the first to point out Italy's productivity slowdown. In fact, it is so well known as to have become an international problem in the aftermath of the Eurozone crisis (see, for example, the 2017 IMF Country Reports on Italy). Yet, there is a dearth of data-based explanations. A notable exception is (Calligaris et al., 2016),who find an increase over time of the left tail of firms with low productivity, which they interpret as an increase in the misallocation of production factors. They do not explain,

[2] See for example, https://www.cronachemaceratesi.it/2018/05/03/il-giudice-su-bianconi-condotta-gravemente-infedele-verso-banca-marche/1097711/ , http://espresso.repubblica.it/inchieste/2019/04/10/news/mose-tangenti-galan-1.333599 .

however, why this misallocation started to arise only in the mid 90s, nor can they explain why this TFP gap is concentrated in ICT-intensive sectors. Similarly, García-Santana et al. (2019) attributes Spanish TFP decline during the same period to an increase in capital misallocation, but they find that this misallocation is particularly severe where connections with public officials are more important for business success. By contrast, we find that TFP decline is particularly pronounced in ICT-incentive sectors.

We are also not the first ones to point to Italy's delay in the adoption of ICT: Bugamelli and Pagano (2004) use micro data from the mid- to late 1990s to show that, in Italy firms need to undergo major reorganization in order to adopt ICT.Milana and Zeli (2004) were the first to correlate these delays with sluggish aggregate productivity growth in the years 1996-99. Their channel is the lower level of ICT investment. Hassan and Ottaviano (2013) use the same channel to explain the slowdown in Italian TFP growth. In our analysis, while we confirm that lower investment is part of the problem, we show that the reduced productivity of such investments is indeed even more important. Schivardi and Schmitz (2017) build on our findings to construct a model that explains productivity differences between Germany and Italy.

The rest of the paper proceeds as follows. Section 3.2 describes our data. In Section 3.3, we explore the possible structural causes for the lack of productivity growth using sector-level data. In Section 3.4, we explore the robustness of our econometric estimates. In Section 3.5, we discuss alternative explanations. In Section 3.6, we analyze firm-level data. Section 3.7 provides suggestive evidence of why, in Italy, loyalty prevails over merit in the selection and rewarding of managers. In Section 3.8, we conclude.

## 3.2 Data and measurement

### 3.2.1 Growth accounting by country and sector

Our main source of sector-level data is the EU-KLEMS structural database (O'Mahony and Timmer, 2009). It contains harmonized measures of value added, capital, labor, total factor productivity and input compensation at the two-digit ISIC level for 25 European countries, as well as Australia, South Korea, Japan, and the United States, accounting for approximately half of the world's GDP. This level of disaggregation allows us to control for country and sector level confounders using fixed effects. It also allows us to study the interaction between country-specific factors and industry-specific factors. Data is available, depending on the country, from as far back as 1970.

Multiple releases of this dataset are available. Based on the degree of harmonization, and the need to merge this dataset with external data, we have chosen to use the 2011 release, which is based on the ISIC rev3.1 sector definition. Based on data availability, we use data from 1984 to 2006. We stop at 2006 because we are interested in the pre-financial crisis period, yet this choice is also convenient because in 2008 there is a structural break in the EU KLEMS data.

The dataset provides industry-level growth accounting. One of its key advantages is the ability to quantify separately the impact of ICT assets and non-ICT assets. In other terms, EU KLEMS breaks down value added growth at constant prices into: 1) TFP growth, 2) the contribution of ICT capital (computers, communication equipment, software...); 3) the contribution of non-ICT capital (land, buildings, machinery...); 4) the contribution of hours worked; 5) the contribution of human capital.

EU KLEMS measures the growth in "labor services" as the weighted average of the growth of hours worked by different worker categories, where the weights are given by the compensation share of each worker's category (age, sex, and skill level). Concordantly, human capital growth is defined as the difference in growth rates between labor services and unweighted hours worked.

We now proceed to summarize the EU KLEMS growth accounting methodology. Assume that, for every country $c$, sector $s$, time $t$, there exists a representative firm that produces output $Y$ (measured as value added at constant prices) by combining capital $K$ and labor $L$ using a generic production function $F$ :

$$Y_{cst} = A_{cst} \cdot F_{cst}\left(K_{cst}, L_{cst}\right)$$

where $A$ is the firm-level total factor productivity. Capital itself is broken down into two different types: ICT capital, and non-ICT capital:

$$K_{cst} = K_{cst}\left(K^{I}_{cst}, K^{N}_{cst}\right) \quad .$$

Similarly, there are $J$ different categories of workers, which differ by demographic factors, skill level, and so on. The total labor input is a combination of the hours worked by the different categories of workers

$$L_{cst} = L_{cst}\left(N^{1}_{cst}, N^{2}_{cst}, ..., N^{J}_{cst}\right)$$

where the total hours worked is defined as:

$$N_{cst} = \sum_{j=1}^{J} N^{j}_{cst} \quad .$$

Let $P$, $R^I$, $R^N$, $W^j$ be, respectively, the prices of output, ICT capital, non-ICT capital and type-$j$ labor and define the following notation for the natural logarithm of a generic variable $X$:

$$x_{cst} := \log X_{cst}$$

under the assumption of constant returns to scale and competitive markets, we have

$$P_{cst} Y_{cst} = R_{cst} K_{cst} + W_{cst} L_{cst}$$

where the sector-level price indices $W$ and $R$ are defined implicitly by:

$$R_{cst} K_{cst} = R^I_{cst} K^I_{cst} + R^N_{cst} K^N_{cst} \tag{3.1}$$

$$W_{cst} L_{cst} = \sum_{j=1}^{J} W^j_{cst} N^j_{cst} \quad . \tag{3.2}$$

As shown by Jorgenson et al. (1987, 2005), we can then obtain the sector-level growth of Total Factor Productivity (TFP) from the following equation:

$$\frac{dy_{cst}}{dt} = \underbrace{\frac{da_{cst}}{dt}}_{\text{TFP growth}} + \underbrace{\left(1 - \frac{W_{cst} L_{cst}}{P_{cst} Y_{cst}}\right) \frac{dk_{cst}}{dt}}_{\text{Capital Contribution}} + \underbrace{\left(\frac{W_{cst} L_{cst}}{P_{cst} Y_{cst}}\right) \frac{d\ell_{cst}}{dt}}_{\text{Labor Contribution}} \tag{3.3}$$

where we have used the first-order condition:

$$MRPL_{cst} = P_{cst} \cdot \frac{\partial Y_{cst}}{\partial L_{cst}} = W_{cst} \tag{3.4}$$

which, after multiplying both sides by $L/PY$ yields:

$$\frac{\partial Y_{cst}}{\partial L_{cst}} \cdot \frac{L_{cst}}{Y_{cst}} = \frac{W_{cst} L_{cst}}{P_{cst} Y_{cst}} \quad . \tag{3.5}$$

The same result holds for capital. Also, the growth of capital and labor inputs can be further decomposed as:

$$\frac{dk_{cst}}{dt} = \left(\frac{R^I_{cst} K^I_{cst}}{R_{cst} K_{cst}}\right) \frac{dk^I_{cst}}{dt} + \left(\frac{R^N_{cst} K^N_{cst}}{R_{cst} K_{cst}}\right) \frac{dk^N_{cst}}{dt} \tag{3.6}$$

$$\frac{d\ell_{cst}}{dt} = \sum_{j=1}^{J} \left(\frac{W^j_{cst} N^j_{cst}}{W_{cst} N_{cst}}\right) \frac{dn^j_{cst}}{dt} \quad . \tag{3.7}$$

Based on these equilibrium relationships, the yearly growth of log value added at the sector level is decomposed, in the EU KLEMS database, into the sum of the following five flow variables:

$$\text{ICT Contribution}_{cst} \quad := \quad \left[\!\left[1 - \frac{W_{cst}L_{cst}}{P_{cst}Y_{cst}}\right]\!\right] \left[\!\left[\frac{R^I_{cst}K^I_{cst}}{R_{cst}K_{cst}}\right]\!\right] \Delta k^I_{cst} \tag{3.8}$$

$$\text{Non} - \text{ICT Contribution}_{cst} \quad := \quad \left[\!\left[1 - \frac{W_{cst}L_{cst}}{P_{cst}Y_{cst}}\right]\!\right] \left[\!\left[\frac{R^N_{cst}K^N_{cst}}{R_{cst}K_{cst}}\right]\!\right] \Delta k^N_{cst} \tag{3.9}$$

$$\text{Human Capital Contribution}_{cst} \quad := \quad \left[\!\left[\frac{W_{cst}L_{cst}}{P_{cst}Y_{cst}}\right]\!\right] \Delta\left(\ell_{cst} - n_{cst}\right) \tag{3.10}$$

$$\text{Hours Worked Contribution}_{cst} \quad := \quad \left[\!\left[\frac{W_{cst}L_{cst}}{P_{cst}Y_{cst}}\right]\!\right] \Delta n_{cst} \tag{3.11}$$

$$\Delta \log \text{TFP}_{cst} \quad := \quad \Delta y_{cst} - (3.8) - (3.9) - (3.10) - (3.11) \;=\; \Delta a_{cst} \tag{3.12}$$

where the delta () symbol represents taking the one-period time difference operator and the thick bracket $[\![\cdot]\!]$ represents taking the average between the beginning and end-of-period values of a variable.

As shown in Table 3.2, Panel A, over the 1996-2006 period, $\log TFP$ has an annual average of 0.012, and ranges from -0.292 to 0.20 (with a standard deviation of 0.0364). *ICT Contribution* has an annual average of 0.005, a standard deviation of 0.006 and ranges from -0.005 to 0.055. *Non-ICT Contribution* has an annual average of 0.008, a standard deviation of 0.013 and ranges from -0.028 to 0.095.

For further information on the EU KLEMS dataset, please refer to O'Mahony and Timmer (2009).

### 3.2.2 Trade data

In order to measure the impact of competition from China across countries and sectors, we need trade data by origin country, origin sector, destination country, destination sector and year. The output concept and industry classification must be compatible with EU KLEMS (value added and ISIC rev 3.1, respectively). The only database that satisfies all of these requirements is the World Input-Output database (WIOD), by Timmer et al. (2015). For each year, the dataset contains data for $41^2$ country pairs $\times$ $28^2$ sector pairs combinations, for a total of 1,317,904 observation in any given year.

We start by computing the "China Shock" in sector $s$ to destination market $m$ as

$$\text{China Shock}_{smt} := \left[\!\left[ \frac{Y_{\text{China},smt}}{\sum_{c \neq \text{China}} Y_{csmt}} \right]\!\right] \cdot \Delta \log Y_{\text{China},smt}$$

where $m$ identifies the country/sector of destination of the export. $Y_{cmst}$ is the export in value added (at constant prices) of country $c$, sector $s$, into destination market $m$ at time $t$. Note that the growth of Chinese export is multiplied by the market share of Chinese export vis-a-vis all its competitors in destination market $m$. The derivation and the rationale for this variable are explained more in detail in the Online Appendix.

Then, for every country $c$ sector $s$ we compute *China Exposure* as the weighted average of *China Shock* in sector $s$ across all destination markets

$$\text{China Exposure}_{cst} = \sum_m \left[\!\left[ \frac{Y_{csmt}}{\sum_m Y_{csmt}} \right]\!\right] \cdot \text{China Shock}_{smt} \quad .$$

Notice that what makes *China Shock* specific to country $c$ is the weighting, given by the share that destination market $m$ represents $c$.

We aggregate across 41 destination countries (including a "Rest of the world" aggregate) and 23 destination sectors, implying that every observation of *China Exposure* is the result of taking a weighted average of 943 WIOD data points.

Summary statistics for *China Exposure* are also presented in Table 3.2, Panel A: it has a mean of 0.012, a standard deviation of 0.021 and it ranges from -0.001 to 0.193. For further information on the WIOD dataset, please refer to Timmer et al. (2015).

We also use a similar dataset, the OECD-WTO Trade in Value Added (TiVA) dataset, to compute the following metric of openness to international trade.

$$\text{Trade Openness}_{cst} = \frac{E_{cst} + I_{cst}}{Y_{cst}}$$

where $E$ and $I$ are, respectively, exports and imports in value added. The reason why we use a different database for this variable is that the TiVA dataset, unlike WIOD, provides country/sector-level estimates of *total* exports, imports and value added (the WIOD does not). At the same time, it does *not* provide a detailed breakdown of trade by destination country and sector, which we do require in order to compute *China Exposure*. The variable *Trade Exposure* has mean 0.897, standard deviation 0.849 and it ranges from 0.017 to 8.116.

### 3.2.3 Country-level variables

We present here variables that vary at the country level. Summary statistics are presented in Table 3.2, Panel B.

To measure the extent to which firms select, promote, and reward people based on merit, we construct a variable called *Country Meritocracy*. It is built using response data from the WEF Global Competitiveness Report Expert Opinion Survey (2012). We compute the variable as the average numerical answer to the following three questions: 1) "In your country, who holds senior management positions?" [1 = usually relatives or friends without regard to merit; 7 = mostly professional managers chosen for merit and qualifications]; 2) "In your country, how do you assess the willingness to delegate authority to subordinates?" [1 = not willing at all – senior management makes all important decisions; 7 = very willing – authority is mostly delegated to business unit heads and other lower-level managers]; and 3) "In your country, to what extent is pay related to employee productivity?" [1 = not at all; 7 = to a great extent]. The reason we opted to construct our own measure of meritocratic management, is that the pool of countries for which similar measures are already available (Bandiera et al., 2008; Bloom et al., 2012b) does not overlap with the EU KLEMS sample. Using an alternative variable would shrink the size of our sector-level dataset by 38% or more, resulting in insufficient country-level variation to identify the desired effect.

*Country Meritocracy* has a mean of 4.683 and a standard deviation of 0.635. Italy has the lowest value: 3.387. Sweden has the highest: 5.504. This variable has the obvious downside of being perception-based and we do not want our empirical results to hinge on its specific construction. Unfortunately, we do not have access to data sources that allow us to compute an alternate measure of meritocratic management at the country level. We do, however, have access to a firm-level dataset, which allows us to gauge meritocratic management practices more objectively and granularly. This data is discussed in detail in 3.2.5.

As the main measure of regulatory protection of employment we use the composite index *Employment Laws* developed by Botero et al. (2004), which captures difficulty of hiring, rigidity of hours, difficulty of redundancy, and redundancy costs: it has a mean value of 0.535, a standard deviation of 0.201, and it ranges from 0.164 (Japan) to 0.745 (Spain). Italy has a value of 0.650.

Because we do not want our results to rely on the specific variable chosen to quantify this effect, we use an alternative measure for robustness: the OECD Employment Protection Legislation (EPL) composite index (version 1). This variable has a panel structure and is available for different countries with different start dates. We average it across the years included in the post-1995 sample period (1996-2006) for which it is available; if the earliest available year is after 2006, we use the earliest available datapoint. *Employment Protection* has mean 2.153, standard deviation 0.747 and it ranges from 0.260 (USA) to 3.310 (Czech Republic). Italy has a value of 2.76.

The variable *ICT Infrastructure* is a sub-index of the Networked Readiness Index, published by the World Economic Forum (2012); it measures the quality of ICT infrastructure that different countries have in place and is constructed by combining country-level data on mobile network coverage, the number of secure internet servers, internet bandwidth, and electricity production. *ICT Infrastructure* has mean 5.894, standard deviation 0.708 and it ranges from 4.317 (Hungary) to 6.904 (Sweden). Italy has a value of 4.779.

To control for cross-country differences in the quality of management training, we use the variable *Management Schools,* which is also derived from response data from a question of the WEF executive opinion survey: "In your country, how do you assess the quality of business schools?" [1 = extremely poor – among the worst in the world; 7 = excellent – among the best in the world]. It has a mean of 5.109, a standard deviation of 0.645, and it ranges from 3.963 (Czech Republic) to 6.121 (Belgium). Italy has a value of 4.792.

Finally we also use the variable *Shadow Economy,* an estimate of size of the shadow economy, as a share of GDP, computed country-by-country by Schneider (2012): it has mean 0.172, standard deviation 0.055 and it ranges from 0.086 (USA) to 0.270 (Italy).

In Section 3.5, we take into account the effect of variation in institutional quality across countries and time. To do so, we use two indicators from the World Bank's Worldwide Governance Indicators (WGI): Rule of law and Control of Corruption. We use the changes in these variables (*Rule of Law* and *Control of Corruption*, respectively) over the period 1996-2006. *Rule of Law* has mean 0.002, standard deviation 0.021 and it ranges from -0.063 (Italy) to 0.023 (Ireland). *Control of Corruption* has mean -0.003, standard deviation 0.020 and it ranges from -0.034 (Czech Republic) to 0.027 (Japan). Italy has a value of 0.010.

One important caveat about these measures is that they are standardized within years: they do not therefore carry, in theory, cardinal meaning, but only ordinal meaning. We believe nonetheless that they are suitable for our analysis, for two reasons. Firstly, analysis by Kaufmann et al. (2006) finds "no systematic time-trends" in these indicators. Secondly, Acemoglu et al. (2006) argue that a country's distance from the technological frontier depends on the relative, rather than absolute quality of its institutions.

Nevertheless, for robustness, we also use a distinct, non-composite measure of the quality of government, computed by Chong et al. (2014), that is expressed in levels. This last variable, which is based on the length of time needed to get back a letter sent to a fictions address in a foreign country, we call *Govt Inefficiency*: it has mean 94.3, standard deviation 42.0 and ranges from 16.2 (USA) to 173.4 (Italy): a higher value corresponds to lower quality of public sector output.

### 3.2.4 Sector-level variables

We could not find an existing measure of how much each sector is dependent on government regulation and intervention. Thus, we constructed one by counting news in major economics and financial news outlets from Dow Jones' Factiva News Search database. We exploit the fact that, in this database, news are tagged by sector and topic. To construct our variable, we build a correspondence table between ISIC rev 3.1 (EU KLEMS's sector definition) with Factiva's industry tags.

The variable *Govt Dependence* is defined, for each sector $s$, as the number of news articles having "Government Contracts" or "Regulation/Government Policy" as topic, as a percentage of the total news articles for sector $s$. We consider the universe of articles from Dow Jones, the Financial Times, Reuters, and the Wall Street Journal published from 1984 to 2017. The value of this variable, for each sector, is displayed in Figure 3.4. It has mean 0.045, standard deviation 0.024 and it ranges from 0.020 (Basic Metals) to 0.126 (Agriculture, Forestry and Fishing).

In order to capture variation in the need for labor force mobility across sectors, we use mass layoff rates in US industries, computed by Bassanini and Garnero (2013) using data from Current Population Survey (CPS) displaced workers supplements covering the period 2000–2006 . The variable *US Layoff* has mean 0.052, standard deviation 0.017 and it ranges from 0.022 (Utilities) to 0.090 (Textiles and Apparel).

### 3.2.5 Firm-level data

For the firm-level analysis of Section 3.6, we use the EFIGE (European Firms in a Global Environment) dataset, developed by Altomonte and Aquilante (2012). The dataset covers 14,759 manufacturing firms from seven European countries (Austria, France, Germany, Hungary, Italy, Spain, and the United Kingdom).

In addition to balance sheet information obtained from the Amadeus-BvD databank, this dataset contains response data from a survey undertaken in 2010 that covers a wide range of topics related to the firms' operations. In particular, this survey contains questions about managerial practices, which allows us to compute a measure of firm-level meritocracy. Specifically, the questions are: 1) "Can managers make autonomous decisions in some business areas?" 2) "Are managers incentivized with financial benefits?"

3) "Has any of your executives worked abroad for at least one year?" 4) "Is the firm not directly or indirectly controlled by an individual or family-owned entity? If it is, was the CEO recruited from outside the firm?" 5) "Is the share of managers related to the controlling family lower than 50%?" . We construct the variable *Firm Meritocracy* by summing the number of affirmative answers to the above questions: it has mean 1.554, standard deviation 1.272, and it ranges from 0 to 5. The average value for Italian firms is 1.07.

Our firm-level metric of meritocratic management overlaps conceptually with the *People Management* practices score by Bloom and Van Reenen (2007): therefore, if it is correctly constructed, we would expect the two measures to correlate. In the Online Appendix, we show that this is indeed the case.

Similarly, the survey asks whether a firm's management uses: 1) IT systems for internal information management; 2) IT systems for e-commerce; and 3) IT systems for management of the sales/purchase network. We construct the variable *ICT Usage* as the sum of the affirmative answers to these questions: it has mean 1.262, standard deviation 0.935, and it ranges from 0 to 3.

The survey also collected information on the constraints faced by firms, by asking managers which of the following (non-mutually exclusive) factors prevent the growth of their firms: 1) financial constraints, 2) labor market regulation, 3) legislative or bureaucratic restrictions, 4) lack of management and/or organizational resources, 5) lack of demand, and 6) other. Firms are also offered the option to say that they face no constraints. To measure these constraints, we create three dummy variables that represent, respectively, whether the firm chooses the first (34.1% of the firms in EFIGE), and/or second (39.2%), and/or third option (20.8%).

In order to corroborate our findings from aggregate EU KLEMS series, we need to build a firm-level measure of TFP. Unfortunately, there is no internally-consistent way to do this. The reason is that, in response to different data availability constraints and methodological challenges, the macroeconomics and the industrial organization (IO) literatures have developed widely different approaches to compute TFP, which are not consistent with each other (Foster et al., 2016).

The "macro" approach, exemplified by KLEMS, is to use aggregate value added at constant prices as the measure of output, assume perfect competition, and obtain production function parameters from the share of labor compensation in aggregate value added. IO economists, on the other hand, use (deflated) firm revenues or gross output as the output concept and assume imperfect competition to recover production function parameters from firm-level data.

Our data does not allow to resolve this debate. The best we can do, given our data, is to compute TFP at the firm-level, using a methodology that mimics as closely as possible the one used by EU KLEMS, and studying its robustness. In Sub-section 3.6.5, we discuss why this methodology is problematic if we wish to relax the assumption of perfect competition, and show how we can use EFIGE data to compute an alternative firm-level TFP growth series under the assumption of monopolistic competition. We will use this alternative TFP measure to investigate the robustness of our econometric results at the firm level to violations of the perfect competition assumption.

Our baseline, EU-KLEMS consistent measure of TFP at the firm level is given by the following formula:

$$\Delta \log \mathrm{TFP}_{it} = \Delta y_{it} - \left(1 - \frac{W_{cst}L_{cst}}{P_{cst}Y_{cst}}\right)\Delta k_{it} - \left(\frac{W_{cst}L_{cst}}{P_{cst}Y_{cst}}\right)\Delta \ell_{it} \quad . \tag{3.13}$$

Firm-level value added is computed as EBITDA+labor costs (which implies the same intermediate input definition as EU KLEMS), deflated using the EU KLEMS sector-level value added price index. Firm-level labor input is given by labor costs, deflated using the EU KLEMS sector-level price index. The firm-level capital stock is measured as Fixed Assets (lagged), deflated using sector-level GFCF price indices from the OECD Structural Analysis (STAN) dataset.[3]

BvD accounting data in the EFIGE dataset is available beginning in 2001: therefore (in order to avoid using data from the crisis) our firm-level TFP growth will be computed for the period 2001-2007. The resulting variable $\Delta \log TFP_{2001-2007}$ has mean 0.002, standard deviation 0.073 and it ranges from -2.116 to 1.916.

The dataset also contains information on the firms' workforce composition. We use the percentage of employees with a university degree and the percentage of employees with temporary employment contracts. The variable *Employees with Degree* has mean mean 0.094, standard deviation 0.134; the variable *Temporary employees* has mean 0.256 and standard deviation 0.385.

[3] We use OECD StAn capital deflators because capital deflators for France and Hungary are not provided directly in the EU KLEMS dataset (due to confidentiality constraints). OECD StAn is the most similar database to EU KLEMS and uses the same sector definition.

## 3.3 Evidence from sector-level data

### 3.3.1 Decomposing labor productivity growth by country

In order to understand Italy's low labor productivity growth, for each country $c$ and sector $s$, we decompose the log growth of GDP per hour worked during 1996-2006, following the EU KLEMS methodology. Subtracting the growth of hours worked from both sides of (3.12) and using the constant returns to scale assumption, we obtain the following decomposition of labor productivity (GDP/hour worked):

$$\underbrace{\Delta\left(y_{cst}-n_{cst}\right)}_{\text{LP growth}} = \underbrace{\Delta a_{cst}}_{\text{TFP growth}} + \underbrace{\left[\!\left[1-\frac{W_{cst}L_{cst}}{P_{cst}Y_{cst}}\right]\!\right]\cdot\frac{R^I_{cst}K^I_{cst}}{R_{cst}K_{cst}}\cdot\Delta\left(k^I_{cst}-n_{cst}\right)+}_{\text{ICT Contribution}} \tag{3.14}$$
$$+\underbrace{\left[\!\left[1-\frac{W_{cst}L_{cst}}{P_{cst}Y_{cst}}\right]\!\right]\cdot\frac{R^N_{cst}K^N_{cst}}{R_{cst}K_{cst}}\cdot\Delta\left(k^N_{cst}-n_{cst}\right)}_{\text{Non}-\text{ICT Contribution}}+\underbrace{\left[\!\left[\frac{W_{cst}L_{cst}}{P_{cst}Y_{cst}}\right]\!\right]\cdot\Delta\left(\ell_{cst}-n_{cst}\right)}_{\text{Human Capital Contribution}}\ .$$

This decomposition is shown for each of the countries in EU KLEMS in Table 3.3 and, graphically, in Figure 3.2. The sector weight is the same (1/23) for all sectors in each country, in order to sterilize the effect of differences in specialization across countries.

Italy has by far the lowest labor productivity growth over the 11 year period: 5% vs an average of 33% for all other EU-KLEMS countries. The only other country with a single-digit labor productivity growth is Spain (9%). During the same period Sweden saw its labor productivity soar by 49%. When we decompose labor productivity growth in its four components we find that most of the action is in the residual (the TFP). For Italy, changes in labor composition added 1.3 percentage points to labor productivity growth (versus an average of 3.4%). Investment in non-ICT capital contributed 7.9 percentage points (versus an average of 9.9%). ICT capital investments contributed 2.5 percentage points versus an average of 5.5%. Based on OECD aggregate data, Hassan and Ottaviano (2013) attribute the low labor productivity growth in Italy to low ICT investments. These figures seem to suggest that ICT investments only played a secondary role. The overwhelming share of Italy's labor productivity growth gap remains unexplained, absorbed into TFP: Italian TFP *shrank* by 6.8% during this period, while for the average country it grew by 14.2%, amounting to a gap of 21 percentage points.

Overall, this analysis suggests that very little of Italy's gap in labor productivity growth can be explained by a failure to accumulate capital or to improve the skill mix of the labor force, or by the sectoral composition of its economy. Italy's slowdown appears to be overwhelmingly driven by its lag in TFP growth. This result is not specific to Italy.

The countries that do better in terms of labor productivity growth (Hungary, Austria, Sweden) are also the same that do better in terms of TFP growth. The same is true for the two laggards (Italy and Spain). Thus, we need to explain why Italian TFP growth fell behind. This is what we will try to do next.

### 3.3.2 Decomposing output growth by sector

In Table 3.4 we perform the same decomposition by sector. Not surprisingly, the sectors experiencing the greatest labor productivity growth tend to be the most high-tech sector, while the laggards tend to be services or brick-and-mortar sectors.

The variance across sectors is much larger than across countries: the fastest growing sector, electrical equipment (30 to 33) experienced a labor productivity growth during the period of 88%. In the second one, Post and Telecommunication, labor productivity grew by 73%. By contrast, real estate and business services (70 to 74) and fuel production (23) showed a decline in labor productivity.

By and large, the observed differences in labor productivity growth are mostly driven by differences in TFP growth.

### 3.3.3 Productivity growth during the ICT revolution

We observed that high-tech sectors grew more both in labor productivity and TFP than low-tech ones. Similarly, if we exclude Hungary (which is still catching up), we observe that richer countries (like Sweden and Austria) grew more than poorer ones (like Spain and Italy), contrary to what traditional growth models would predict. Most of these differences seem to be driven by variation in TFP growth. What can explain these patterns?

The mid-1990s marked the beginning of the ICT revolution. One of the unique characteristics of ICT capital investment is the strong complementarity with organisational capital (Brynjolfsson and Hitt, 2003; Brynjolfsson et al., 2002). Consistent with this hypothesis, Bloom et al. (2012a) show that differences in management style between Europe and the United States can explain why labor productivity growth in the Old Continent fell behind the U.S. one after 1995. Is it possible that similar differences within Europe can explain our observed patterns? If so, could this help explain Italy's TFP drop?

Before we move on to investigate this hypothesis further, however, we need to ask one question. Why would the effect of ICT/management complementarities show up in TFP, rather than in the contribution of ICT capital? If the marginal productivity of ICT capital varies systematically across firms or countries according to managerial practices (and these are constant over time) then this should be reflected by the compensation share of ICT capital. To see why this is the case, consider a simplified version of the model presented in Bloom et al. (2012b), in which the production function varies at the sector level and the output concept is value added. Managerial capital is captured by the unobserved input $M$, which we assume for simplicity to vary across countries, and which has the effect of increasing the output-ICT capital elasticity:

$$Y_{cst} = A_{cst} \cdot M_c \cdot \left(K_{cst}^I\right)^{\alpha_{cst}^{KI} + \sigma M_c} \left(K_{cst}^N\right)^{\alpha_{cst}^{KN}} \left(L_{cst}\right)^{1 - \sigma M_c - \alpha_{cst}^{KI} - \alpha_{cst}^{KN}}$$

the first order condition for ICT capital is

$$MRPK_{cst}^I = \left(\alpha_{cst}^I + \sigma M_c\right) \frac{P_{cst} Y_{cst}}{K_{cst}^I} = R_{cst}^I$$

implying:

$$\frac{R_{cst}^I K_{cst}^I}{P_{cst} Y_{cst}} = \left(\alpha_{cst}^I + \sigma M_c\right)$$

the contribution of ICT capital to output growth equals

$$\text{ICT Contribution}_{cst} = \left(\alpha_{cst}^I + \sigma M_c\right) dk_{cst}^I = \frac{\partial y_{cst}}{\partial k_{cst}^I} dk_{cst}^I$$

and TFP growth is given as before, by:

$$\Delta \log \text{TFP}_{cst} = da_{cst}$$

hence, the complementarity between ICTs and management style is captured by *ICT Contribution* and does not affect TFP growth.

For ICT-management complementarities to have an impact on TFP growth, we need to expand the growth accounting framework. We do so by assuming externalities in ICT capital accumulation. While there are other modeling choices (see 3.3.5) that could account for the observed correlations this is, in our view, the simplest and most parsimonious way to allow ICT capital to affect TFP growth.

### 3.3.4 Modeling externalities

Let us start with the simplest version of the firm-level production function with externalities à la Romer (1986) which we assume for simplicity to be a Cobb-Douglas function:

$$Y_{it} = A_{it} \cdot \left(K_{it}^I\right)^{\alpha_{cst}^{KI}} \left(K_{it}^N\right)^{\alpha_{cst}^{KN}} \left(L_{it}\right)^{\alpha_{cst}^L}$$

where $A$ depends on the country/sector-level accumulation of ICT capital $\left(K^I_{cst}\right)$:

$$A_{it} = \overline{A}_{it} \left(K^I_{cst}\right)^{M_i \cdot \alpha^{KI}_{cst}} \quad .$$

$M$ is a country-level parameter that reflects country differences in the adoption of meritocratic management practices and $\overline{A}_{it}$ is the exogenous component of TFP. Bloom et al. (2012a) and Garicano and Heaton (2010) assume that there are complementarities between meritocratic management and ICT capital at the *firm* level. We assume a similar complementarity between meritocratic management and ICT capital at the *aggregate* level. For example, a firm that compensates management according to performance can benefit more from electronic data that suppliers and customers generate when they digitize their production process. Note that the magnitude of this externality depends on how ICT-intensive a firm's production process is, as proxied by the elasticity $\alpha^{KI}$. In the context of the previous example, this assumption implies that the impact of having digitized customers and suppliers is greater if you are more digitized yourself.

If we assume, as EU KLEMS does, that $\alpha^{KI}_{cst}$ is constant over time, then TFP growth at the firm level can be written as:

$$\Delta \log \text{TFP}_{it} = \Delta \overline{a}_{it} + M_i \cdot \underbrace{\alpha^{KI}_{cst} \cdot \Delta k^I_{cst}}_{\text{ICT Contribution}_{cst}} \quad . \tag{3.15}$$

At the EU KLEMS level, we do not observe capital as a stock, but only in changes; as a result, we are going to estimate equation (3.15) in changes. Furthermore, we do not observe firm-level TFP, but sector level-TFP. Finally, we don't observe Meritocracy at the firm level, but only a country-level proxy: the variable *Country Meritocracy*, which is described in Section 3.2. We therefore estimate the following relationship:

$$\Delta \log \text{TFP}_{cst} = \Delta \overline{a}_{cst} + M_c \cdot \text{ICT Contribution}_{cst} \quad . \tag{3.16}$$

Since we don't know the nature of the relationship between the "true" country-level meritocracy $M_c$ and the observed proxy *Country Meritocracy*, we assume the following linear relationship:

$$M_c = \beta_1 + \beta_2 \cdot \text{Country Meritocracy}_c \quad . \tag{3.17}$$

Substituting equation (3.17) into (3.16), we obtain the following regression specification, which we implement in long-term differences (as in Brynjolfsson and Hitt, 2003):

$$\begin{aligned} \Delta \log \text{TFP}_{cs} = {} & \gamma_c + \varsigma_s + \beta_1 \cdot \text{ICT Contribution}_{cs} \\ & + \beta_2 \cdot (\text{ICT Contribution}_{cs} \times \text{Country Meritocracy}_c) + \varepsilon_{cs} \end{aligned} \tag{3.18}$$

where the term $\gamma_c$ is a country fixed effect and $\varsigma_s$ is a sector fixed effect. In other terms, if there are externalities in ICT adoption, the EU KLEMS total factor productivity growth rate should be positively correlated with an interaction term, which is equal to the product of a country-level measure of meritocratic management and the contribution of ICT capital to value added growth.

### 3.3.5 Myopia as an alternative mechanism

Externalities are not the only way in which TFP growth might be dependent on the contribution of ICT capital. A simpler explanation could be based on the failure of firms to recognize the complementarities between ICT and organizational capital. Since there is a discussion even among economists about whether these complementarities exist, it might be reasonable to assume that firms ignore them in their maximization process.

If firms ignore these complementarities, they equalize

$$P_{it} \cdot \frac{\partial Y_{it}}{\partial K_{it}^I} = \alpha_{cst}^I \frac{P_{it} Y_{it}}{K_{it}^I} = R_{it}^I$$

the result would be an under-investment in $K^I$ and a residual TFP which incorporates the effect of complementarities.

Because the externalities model and the myopic model are observationally equivalent with respect to productivity trends, we will not try to disentangle the two empirically in this paper. We recognize, however, that the policy implications of these two models are not the same.

### 3.3.6 Identification

Before ICT revolution the contribution of ICT capital to growth was negligible. Thus, the ICT-Meritocracy interaction (computed in the *post* period) should not predict TFP growth in pre-treatment period (1985-1995). This is a test of the "parallel trend" assumption.

Even if the parallel trend assumption is satisfied, we are concerned that ICT capital growth, the main component of *ICT Contribution*, may depend on sector-level productivity growth. This would be the case in a simple neoclassical growth model, where the rate of capital accumulation along a balanced growth path is directly proportional to the growth of aggregate productivity. If country and sector fixed effects fail to control for this effect and there is a structural break in the capital accumulation process around 1995, it is possible that the OLS estimates of $\beta_2$ might capture not just the causal effect of

Meritocracy and ICT on TFP growth, but also the directionally opposite effect of TFP on ICT capital accumulation.

We rule out this possibility using two different strategies. The first exploits differences and similarities between ICT capital and non-ICT capital. If indeed factors tend to accumulate at a higher rate in sectors where TFP grows faster, this should conceivably affect non-ICT capital as well as ICT capital. If, instead, there is an effect of ICT capital accumulation on TFP that is mediated by meritocratic management, we would expect TFP growth to correlate with the interaction of *Country Meritocracy* and *ICT Contribution*, but not with the interaction of *Country Meritocracy* and *Non-ICT Contribution*. We incorporate this intuition into our econometric analysis by using *Non-ICT Contribution* as a placebo treatment. If indeed *ICT Contribution* is as good as exogenous, then we expect the same regression analysis to *not* yield a statistically significant result when *Non-ICT Contribution* is used in its place.

In a separate specification, we use the same intuition to construct an instrument for *ICT Contribution*:

$$Z^{ICT}_{cst} = \left[\!\left[1 - \frac{W_{cst}L_{cst}}{P_{cst}Y_{cst}}\right]\!\right] \left[\!\left[\frac{R^{I}_{cst}K^{I}_{cst}}{R_{cst}K_{cst}}\right]\!\right] \Delta\left(k^{I}_{cst} - k^{N}_{cst}\right) \tag{3.19}$$

which is identical to its endogenous counterpart, except for the fact that the growth of ICT capital is here replaced by the differential rate of accumulation of ICT capital vis-à-vis non-ICT capital. For the exclusion restriction to hold, it is necessary that, conditional on country and sector fixed effects, faster technical progress does not differentially affect ICT and non-ICT capital accumulation.

The second way we address for endogeneity of *ICT Contribution* is to instrument it with its own cross-country, sector-level average:

$$\text{ICT Contribution Avg}_{st} = \frac{1}{18}\sum_{c} \text{ICT Contribution}_{cst} \quad . \tag{3.20}$$

By shutting down all within sector/cross country variation in this variable, and controlling for country and sector fixed effects, we eliminate any possible direct reverse causation of TFP on ICT. Hence, endogeneity could bias our results only if the *interaction* of *ICT Contribution* (averaged across countries) and *Country Meritocracy* correlates with the regression residuals. Such a correlation could arise if TFP growth causes, *simultaneously*, the adoption of meritocratic management practices as well as ICTs and if *Country Meritocracy* (which is measured in 2012) captures some of the 1996-2006 time variation in management practices due to TFP growth.

While we are not particularly concerned about such temporal variation in *Country Meritocracy* (the WEF surveys on which this variable is based do not exhibit substantial time variation), we still want to account for this potential source of endogeneity by finding an instrumental variable for *Country Meritocracy.*

The ideal instrument would pre-dispose a country to the adoption of meritocratic management, but not change in response to the country's TFP dynamics in 1996-2006. We think that *Judicial Inefficiency* (Djankov et al., 2003), which measures the average number of days to enforce a contract (specifically, the eviction of a tenant or the collection of a bounded check), represents such an instrument.

*Judicial Inefficiency* influences meritocratic management practices because enforcement based on family or personal ties is a substitute for legal enforcement. Thus, when legal enforcement is very slow, firms rely more on loyalty-based bonds rather than formal contracts. Since judicial efficiency is a slow-moving institutional variable, it is also unlikely to react to changes in TFP growth or in firm-level meritocratic management over the period we analyze (1996-2006).

We find that *Judicial Inefficiency* is a strong predictor of *Country Meritocracy* and that *Judicial Inefficiency* interacted with *ICT Contribution* (averaged across countries) is a strong predictor of the interaction between *Country Meritocracy* and *ICT Contribution.* When *Judicial Inefficiency* is included as an interaction effect in a regression in which *Country Meritocracy* is also present (that is, when it is used as a control variable rather than as instrument) *Country Meritocracy* always dominates. While we understand all too well that this is not a test of the exclusion restriction, we hope this check might offer readers some reassurance that the effect of the ICT revolution is truly mediated by meritocracy, and not by judicial efficiency.

In summary, our IV specification uses three instrumental variables for *ICT Contribution* and its interaction with *Country Meritocracy*: the first is $Z^{ICT}$ (which is based on the differential rate of accumulation of ICT capital vis-à-vis non-ICT capital); the second is the interaction $Z^{ICT}\times$ *Judicial Inefficiency*; the third is the interaction *ICT Contribution Avg* $\times$ *Judicial Inefficiency.* We use this combination of instruments because we do not wish to rely only on the difference between ICT and Non-ICT capital, and because our preferred instrument for *ICT Contribution*, *ICT Contribution Avg*, is collinear with sector fixed effects. In addition, by using these three instruments together we over-identify the model and thus we can perform a Sargan-Hansen test, which provides us with a useful diagnostic of whether the relevant exclusion restrictions

$$\mathbb{E}\left(\varepsilon_{cs}\cdot\begin{bmatrix} Z^{ICT}_{cs} \\ Z^{ICT}_{cs}\times\text{Judicial Inefficiency}_c \\ \text{ICT Contribution Avg}_s\times\text{Judicial Inefficiency}_c\end{bmatrix}\right)=0$$

are satisfied in the data.

### 3.3.7 Sector-level TFP growth regressions

The estimation results for the specification in equation (3.18) are shown in Table 3.5. Column 1 shows the OLS estimates when *ICT Contribution* as well as the interaction *ICT Contribution* $\times$ *Country Meritocracy* are used as explanatory variables. In this specification, we find that the interaction coefficient is positive and statistically significant at the 5% level. The two coefficients can be interpreted in the following way: the combined effect of Meritocracy and ICT can either boost or dampen the impact of ICT investments on output, depending on whether the sum of the baseline and interaction coefficient is positive or negative. For a typical country (*Country Meritocracy* = 4.7 to 5) the effect of ICT Contribution is neither dampened not amplified, as the interaction effect is approximately offset by the baseline coefficient on *ICT Contribution*. For Italy, who has the lowest meritocracy score (*Country Meritocracy* = 3.4) the effect of ICT investments on TFP (3.4 $\times$ 1.1 - 5.2 $\approx$ -1.4) entirely offsets their direct effect on output growth, implying that ICT investments in Italy are effectively not contributing to growth. For a country like Sweden, which has very high meritocracy (= 5.5), the indirect effect of ICT investments on output growth mediated by TFP is nearly as large (5.5 $\times$ 1.1 - 5.2 $\approx$ +.85) as the direct effect, implying that the total effect of ICT investment on output growth is nearly double the baseline effect.

In column 2, we perform a "placebo" regression, using *Non-ICT Contribution* in place of *ICT Contribution*. Contrary to the previous specification, the interaction of this variable with *Country Meritocracy* does not appear to predict TFP growth across countries and sectors: the interaction coefficient is negative and not statistically significant. In column 4, we perform an Instrumental Variable regression, using the variables presented in equations (3.19) and (3.20) as instruments for *ICT Contribution*. The IV coefficient for the interaction of *ICT Contribution* and *Country Meritocracy* is positive, statistically significant, and quantitatively close to the OLS estimate. We also present an under-identification test statistic (Kleibergen-Paap): it rejects the null hypothesis that the first-stage coefficients are jointly zero. The Sargan-Hansen test and the Wu-Hausman test yield p-values way above rejection thresholds, which we take as a reassurance that there are no “red flags” of endogeneity in our analysis.

In column 4, we test the parallel trend assumption by using, as dependent variable, the growth of TFP in the period 1985-1995 instead of 1996-2006. The coefficient estimates for *ICT Contribution* and its interaction with *Country Meritocracy* are statistically and economically insignificant: this suggests that, in our empirical design, the parallel trend assumption is satisfied.

In column 5, we perform an additional check: we use lagged values of both TFP growth and ICT Contribution. The intuition for this specification is that, if there really exists a complementarity effect between ICT and Meritocracy, we would expect to find an interaction coefficient of the same magnitude it in the pre-period as well. While we do not want ICT intensiveness in the post-treatment period (1996-2006) to predict TFP growth in the pre-treatment period (this would be a violation of the parallel trend assumption tested in Column 4), we would expect the interaction coefficient to be broadly unchanged when we lag *both* the left-hand side variable as well as the right-hand side variable. We also expect the standard errors to increase, as in the pre-treatment period ICT capital made a very small contribution to output growth across all sectors (there was no treatment)[4]. Consistently with the hypothesis, we find a coefficient of nearly exactly the same magnitude as in column 1 that is statistically insignificant, with a standard error of twice the magnitude.

In Figure 3.3, we summarize these results graphically: we sort countries according to their value of *Country Meritocracy* [5], and sectors according to their (cross-country, post-1995) average value of *ICT Contribution.* i.e. how much growth in value added is attributable to higher ICT investments. We divide both countries and sectors into terciles: we label the top country tercile of meritocracy as “High Merit” and the bottom tercile as “Low Merit”; concordantly, we label the top tercile of sectors as “High ICT” and the bottom tercile as “Low ICT”. Then, we sort countries/sectors into four groups: “High/High”, “High/Low”, “Low/High”, “Low/Low”. For each of these groups we compute the cross-country, median TFP growth during the period 1985–2006. We then plot the four TFP indices so obtained, using 1995 as the base year.

As we can see from Figure 3.3, before 1995 TFP growth was fairly similar across all four groups. By contrast, after 1995 there is a clear pecking order. High-ICT sectors in high-meritocracy countries grow the fastest (19.4% cumulatively). Then, low-ICT sectors in low-meritocracy countries (12.3%). Third come the low-ICT sectors in high-meritocracy

[4] We wish to credit an anonymous referee for making this nifty remark.

[5] We exclude countries for which there is no TFP data before 1995 (Czech Republic, Hungary and Slovenia), so that graph shows the same countries before and after 1995

countries (9.8%) and last the high ICT sectors in low-meritocracy, with barely positive growth (5.3%).

### 3.3.8 Magnitude of the Effect

How much of the Italian TFP growth gap can be explained by smaller ICT externalities due to lack of meritocracy? To answer this question, we compute a counterfactual based on how higher Italy's TFP would be in 2006 had Italy had a Country Meritocracy of 5, which would place it at the same level of Germany and Japan (but lower than US, UK and the Scandinavian countries).

The key input to this counterfactual calculation is the magnitude of the interaction coefficient for *ICT Contribution* and *Country Meritocracy.* We start from the value of 1.1, our baseline estimate in Table 5 column 1. Because our proxy of meritocratic management is measured with error, we consider this estimate conservative (it is attenuated by the measurement error). We multiply this interaction effect times the average *ICT Contribution,* times the gap in *Country Meritocracy*, and sum the gap in *ICT Contribution.* We obtain that Italy would have grown 13.3 percentage points more in TFP. In other words, in a conservative scenario, the interaction of ICT and meritocracy can explain 63% of Italy's missing TFP. If we input a more aggressive coefficient, to account for the attenuation bias, we obtain a counterfactual additional growth in TFP of 16.1 percentage points. In this scenario, we are able to explain up to 77% of Italy's missing TFP growth.

## 3.4 Robustness

### 3.4.1 Potential confounders of meritocracy

Because meritocracy correlates, at the country level, with many other institutional variables, we want to make sure that the observed effect is truly due to meritocracy and not to other factors. To this purpose, in Table 3.6, we regress TFP growth across countries and sectors on a batch of potential confounders of *Country Meritocracy*, interacted with *ICT Contribution.* In particular, we use measures of *ICT Infrastructure*, the quality of *Management Schools* computed by the World Economic Forum, as well as estimates of the size of the *Shadow Economy* computed by Schneider (2012), all interacted with the ICT capital contribution. These variables are described in detail in Section 3.2.

*ICT Infrastructure* and *Shadow Economy* (columns 2-3) don't seem to have a significant impact on TFP growth, when interacted with *ICT Contribution. Management Schools*, on the other hand, is borderline significant (10%, column 4). In column 5, we re-introduce *Country Meritocracy* and control for all these other interactions. Only the interaction *ICT Contribution* × *Country Meritocracy* remains statistically significant.

In the Online Appendix, we reproduce this table using an alternative batch of potential confounders, which include an alternative measure of management training based on the number of GMAT score reports received by business schools in each country, an estimate of average firm size by the OECD, and estimates of human capital from Barro and Lee (2013).

### 3.4.2 Small sample size and measurement of meritocracy

Two obvious weaknesses of our sector-level analysis are the small size of the dataset and the fact that we have to resort to a perception-based measure of meritocratic management. We address both of these shortcomings by augmenting our analysis with a firm-level dataset in Section 3.6.

Additionally, our firm-level data can be used to validate the World Economic Forum measure of *Meritocracy*. Based on Bloom et al. (2012a)'s insight that it is the location of the firm's ownership that determines the ability to leverage ICT, we average the firm-level measure of meritocracy from EFIGE at the headquarter country-level. We can then examine the correlation of *Country Meritocracy* and *Firm Meritocracy* across 44 countries. This relationship can be seen in Figure 3.6. The $R^2$ of this regression is 64.3%, which suggests that *Country Meritocracy* is an acceptable proxy for our sector-level regression analysis.

Another important consideration is how both our measures of meritocratic management relate to the *People Management* practices score of Bloom and Van Reenen (2007). Because only 11 countries overlap across their sample and ours, we cannot use their measure in our sector-level analysis. We also cannot use it in our firm-level analysis because the firm-level data is anonymized. Nonetheless, we are able to carry out a cross-sectional validation exercise for both of our variables. In the Online Appendix, we show that both *Country Meritocracy* and *Firm Meritocracy* correlated strongly with the WMS People Management score: this provides further reassurance that these variables measure what we intend them to.

### 3.4.3 Emerging Europe and Italy

We want to exclude the possibility that our results are entirely driven by Italy, which has by far the lowest Meritocracy score among the countries in our sample. In the Online

Appendix, we repeat our estimation without Italy. Not only does the coefficient remain statistically significant, but its magnitude is very similar to the one estimated in Table 3.5. We do the same for our China regressions.

Moreover, our sample includes three developing European countries - Czech Republic, Hungary and Slovenia - for which no growth accounting data is available before 1995. Hence, in Online Appendix A, we show that our results are robust to the exclusion of these countries.

### 3.4.4 Mismeasurement of the production function

In the EU-KLEMS framework, sector-level input expenditures are used to estimate production function elasticities. This approach has drawbacks that are well documented.

In the last twenty years, significant advances have been made in production function estimation that leverage firm-level data: econometric techniques have been introduced that account for sample selection and simultaneity in the production function (see for example Olley and Pakes, 1996; Levinsohn and Petrin, 2003; Wooldridge, 2009). Unfortunately, these approaches are not implementable in our setting: this is the case both for our sector-level data and our firm-level data. The reason is that we do not observe the input of ICT capital at the firm level. Hence, we are forced to rely on EU KLEMS output/capital elasticities.

One reason we make our analysis robust to mismeasurement of the production function parameters is by using the instrumental variable specification. However, in order to have an additional safeguard against this possibility, in the Online Appendix we investigate the robustness of our estimates from Table 3.5 to mismeasurement of parameters of the production function. In particular, we worry about how non-constant returns to scale and mismeasurement of the output/capital elasticities might bias our measures of Total Factor Productivity growth. We argue, and subsequently provide evidence by using the GMM framework, that if such mismeasurement exists, it is small and does not undermine our estimates.

## 3.5 Alternative explanations

### 3.5.1 Capital and labor misallocation

In this section, we want to consider alternative explanations for Italy's dismal productivity growth. The main competing explanation is trade integration.

Recent work by Gopinath et al. (2017) shows that Spain (and more broadly Southern Europe) has experienced a decrease in the efficiency of capital allocation as a consequence of the massive capital investments triggered by the decline in the real interest rate that followed the introduction of the common currency. This misallocation, in turn, led to sizable losses in aggregate TFP. We investigate whether this mechanism could provide, at least quantitatively, an alternative explanation for Italy's (as well as Spain's) productivity decline.

The explanation is very convincing for Spain, where gross fixed capital formation (GFCF) went from 21% of GDP in 1994 to 31% in 2007. It is less plausible for Italy, where GFCF went from 18% of GDP in 1994 to 21% in 2007.

### 3.5.2 The China shock

The second alternative hypothesis that we want to consider is the China schock. China's entry in the WTO in 2001 threatened Italy's market share in global manufactures (Tiffin, 2014), precisely at the time when Italy had given up exchange rate flexibility by joining the euro. Contemporary trade theory (see Melitz, 2003) suggests that trade liberalization should have a positive impact on productivity, since it favors the downsizing of less productive firms and the reallocation of factors towards more productive ones. However, this might not necessarily have been the case for countries, such as Italy, in which labor regulation might have hindered such reallocation. It is indeed possible for a sector's productivity to decrease in the wake of a demand shock if the firms operating in that sector are unable to adjust their scale in response to the shock. In other words, while in the US, where there are fewer labor markets frictions, competition from Chinese products resulted in significant displacement of manufacturing workers and productivity gains (Pierce and Schott, 2016), in Italy the effect might have been reversed, causing sizable productivity losses with moderate effects on employment.

In order to test this hypothesis, we regress TFP growth across countries and sectors on a proxy of the magnitude of the China shock (*China Exposure*). The result of estimating this specification are presented in Table 3.7, Column 1. As expected, we find a positive, albeit not statistically significant effect of *China Exposure* on productivity growth. The economic significance of this coefficient can be described as follows: if competition from China causes value added in a country/sector to drop by 10%, we expect TFP to rise by about 0.4% as a consequence.

If the impact of the China shock on TFP growth is mediated by labor regulation, we should find that the positive effect of *China Exposure* on TFP growth is reverted for countries that make it difficult to reallocate labor by granting a lot of regulatory protection to employees. To capture this in our regression specification, in Column 2 we interact *China Exposure* with a measure of labor market employment protection. As our primary measure, we use a composite index of employment law strictness from Botero et al. (2004). As an alternative measure of employment regulations we use, in column 3, OECD's Employment Protection Legislation index. The resulting regression equation is:

$$\begin{aligned}\Delta \log \mathrm{TFP}_{cs} = \gamma_c + \varsigma_s + \beta_1 \cdot \mathrm{China\ Exposure}_{cs} \\ + \beta_2 \cdot (\mathrm{China\ Exposure}_{cs} \times \mathrm{Employment\ Laws}_c) + \varepsilon_{cs} \quad .\end{aligned} \tag{3.21}$$

The regression intercept is allowed to vary across countries and sectors through the inclusion of fixed effects; there is no time variation in the variables because we use long-term differences/averages. The results of these regressions are presented in columns 2 and 3. Both interaction effects are statistically insignificant: moreover, the effect is positive, contrary to what would be needed to explain Italy's slowdown.

The penetration of Chinese exports could be itself the result of low TFP growth in the country of destination. By averaging the China shock across countries-of-destination in the construction of *China Exposure*, we mitigate this concern. To further alleviate endogeneity concerns about *China Exposure*, we use an instrumental variable. Our instrument, like *China Exposure*, is also a weighed average of the effect of the variable *China Shock* across destination markets; however, it differs from *China Exposure* in that it excludes the domestic market from the domain of summation:

$$Z_{cst}^{China} = \sum_{m \neq (c,s)} \left[\!\left[ \frac{Y_{csmt}}{\sum_{m \neq (c,s)} Y_{csmt}} \right]\!\right] \cdot \mathrm{China\ Shock}_{smt} \quad .$$

In column 4, we carry out the instrumental variable regression, obtaining similar results. The estimated coefficients become larger in absolute terms (0.243 for the baseline coefficient and 1.086 for the interaction with *Employment Laws*). The p-value for the Kleibergen-Paap test is below 0.01, suggesting that the first stage is strong. The Wu-Hausman test yields a p-value of over 0.044, somewhat confirming our suspicion that *China Exposure* might be endogenous.

One potential concern is that the China shock might have impacted all sectors equally, resulting in insufficient within-country variation to identify the effect of interest. Empirically, this does not appear to be an issue. By computing the ratio of the 75th percentile to the 25th percentile of *China Exposure*, we find that there is significant heterogeneity: the country/sector at the 75th percentile of the distribution is 8 times as exposed to demand shocks from China as the country/sector lying at the 25th percentile. Furthermore, if there was not enough country/sector variation, it would be impossible for the

China shock to explain the Italian disease, because the shock would have hit all countries equally.

In sum, these findings suggest that, between 1995 and 2006, productivity tended to grow faster, not slower, in countries/sectors that were more exposed to competition from China. This effect does not appear to reverse for countries with strong regulatory protection of workers, regardless of the measure used. Hence, the hypothesis that competition from China (combined with domestic labor market rigidity) caused Italy's slowdown does not find support in the data.

### 3.5.3 Labor market regulation

Some commentators (Calligaris et al., 2016) attribute Italy's TFP drop to an increase in productivity dispersion within different size/area groups, due – among other factors – to a friction in reallocating capital and labor.

The evidence in Table3.7 suggests that the China shock alone cannot explains this drop in productivity. This lack of findings, however, might be due to the fact that China's entry into the WTO is not the only possible reason why factors might need to be reallocated. In order to test the labor reallocation hypothesis more broadly, we adopt an alternate variable to gauge the sectorial need for labor reallocation: *US Layoff Rate.* It is defined as the rate of mass layoffs in US industries, computed by Bassanini and Garnero (2013) using data from the CPS biennial displaced workers supplement. The rationale for using this variable, similar to that of the financial dependence metric used in Rajan and Zingales (1998), is that we know United States to have minimal labor market distortions. By using this variable, we aim to capture the technological demand for labor reallocation.

In Table 3.7, column 5, we interact this variable with country-level *Employment Laws.* As expected, the coefficient is negative, suggesting slower TFP growth in countries with rigid laws in sectors where the need for reallocation is high; this effect is, however, quantitatively small and not statistically different from zero.

### 3.5.4 The Eurozone accession

As the ICT revolution gained footing, Italy and other European countries adopted a common currency, the Euro, preventing competitive devaluation. This restriction might have affected Italian exports due to the fact that Italian exports had greatly benefited from competitive devaluation in the past.

In the short term, a decrease in external demand for Italian products can adversely affect productivity through several channels. First, there is a scale effect. A reduction in export volumes can slow down or reverse firm growth. This, in turn, might adversely impact TFP gains through the scale elasticity and by stopping learning-by-doing. Second, a decrease in external demand for Italian products has a negative impact on the profitability of Italian firms. To the extent firms are liquidity constrained, this reduction in profitability can also lead to a reduction in investments in R&D and new technologies, slowing down not only labor productivity but also TFP growth. The third potential channel is labor adjustment costs. In the absence of growth in internal demand, a decrease in external demand forces Italian firms to cut back production, at least temporarily. If firms cannot easily lay off workers in response to this shock, productivity will drop, the more so the harder it is to lay off workers (i.e., the stronger employment protection is). All these negative effects should be short term. In the long term, if there is a permanent drop in demand for Italian products, firms will eventually adjust or close. If they adjust, they will probably be forced to increase productivity. If they close, the least productive firms will close first, increasing the average productivity simply through a compositional effect. Thus, the predictions for the long term are the opposite. While it is hard to imagine that 10 years are still the short term, we should let the data speak. If this were the case, the sectors that would be more affected would be those more open to trade at the beginning of the period and the countries that would be more affected are those with stricter labor protection laws.

In Table 3.8, column 1 we regress TFP growth on *Trade Openness* (defined in Section 3.2) as well country and sector fixed effects. We find the effect of *Trade Openness* to be economically and statistically indistinguishable from zero. In columns 2, we add an interaction term with *Employment Laws*. We find that the interaction term is negative and borderline significant (at 10% confidence level) giving some credence to Euro hypothesis. In column 3, we add to this specification our key explanatory variables *ICT Contribution* as well as it interaction with *Country Meritocracy*. The interaction term *ICT Contribution* × *Country Meritocracy* has a positive, statistically significant coefficient that is very similar in magnitude to the one obtained in Table 3.5. Interestingly, also the interaction between *Trade Openness* and *Employment Laws* remains negative and statistically significant, suggesting that the two explanations are orthogonal. In column 4, we test the robustness of this latter result by replacing *Employment Laws* with its OECD-supplied counterpart *Employment Protection*. We find that the interaction term *ICT Contribution* × *Country Meritocracy* remains positive and statistically significant, while the interaction *Trade Openness* × *Employment Protection* is found to be statistically and economically insignificant.

In sum, while we cannot reject the Euro hypothesis, the evidence in favor of it is weak and it doesn't seem to undermine the ICT-based one.

### 3.5.5 Labor market reforms and shadow employment

Starting from 1997, the Italian government passed a series of legislative measures that regulated certain categories of temporary and part-time work; these include the well-known "Biagi Law", the "Pacchetto Treu" as well as Law "2002 n.189", which allowed for the regularization of illegal work of non-EU immigrants. The aim of these regulations was, at least in part, to reduce shadow employment and increase official employment.

Some observers, notably Krugman (2012) in a New York Times column, suggested that this might have biased employment growth statistics upwards for Italy, bringing down Italy's productivity: according to this theory, Italy's productivity slowdown might be nothing more than a statistical artifact.

Unfortunately, we are unable to determine whether or to what extent this effect is present in the EU KLEMS labor input time series. However, we can present two pieces of evidence which suggest that, if this effect exists, it cannot account but for a small fraction of Italy's productivity growth gap.

First, recent empirical analysis of matched Italian employer-employees data (see Daruich et al. 2018) determined that the increase in aggregate employment as a result of the reforms was minimal. Second, the Italian Statistics Institute (Istat) has been computing estimates of the incidence of undeclared work since the early 90s. We recovered these estimates for the years 1992, 1997 and 2003, from a statistical document that Istat (2005) produced for a parliamentary commission. These estimates allow us to perform a back-of-the-envelope calculation of the potential effect that these regulations might have had, based on conservative assumptions.

According to Istat's estimates, the incidence of undeclared workers as a percentage of total employment was 13.4% in 1992, 14.8% in 1997, and then again 13.4% in 2003. We make the conservative assumptions that 1) this effect is totally missed by EU KLEMS employment data; 2) shadow employment would have grown between 1997 and 2003 by the same percentage amount it did between 1992 and 1997, had the labor market reforms not been passed. Then, employment growth has been overestimated by, at most, 2.8%. By multiplying this by an assumed labor elasticity of 2/3, we obtain an upper bound to TFP underestimation of 1.9%, which is trivial vis-à-vis Italy's 21.1% TFP growth gap.

### 3.5.6 An institutional decline?

An alternative explanation to Italy's productivity decline is that Italy experienced, over the 1996-2006 period, a decline in the quality of its institutions. Over this period, in fact, Italy recorded the sharpest decline in "Rule of Law" (one of the Worldwide Governance

Indicators) within our sample (Gros, 2011). If Italy's government is the real culprit of the TFP drop, we should observe that the sectors more dependent on regulations and government inputs should experience a sharper TFP drop.

We don't lack country-level indicators of government effectiveness (e.g., La Porta et al., 1999), but we do lack a measure of sectoral dependence on government inputs. As a source of country-level variation, we use the change in the World Bank's *Rule of Law* and *Control of Corruption* scores. To measure how much each sector is dependent on the government, we compute our own measure of sectoral government dependence. Specifically, we count news articles using the Factiva news search engine. The variable is defined, for each sector, as the ratio of total news counts having "government" as the topic to total news for that sector. Figure 4 shows how this variable varies across EU KLEMS sectors. This measure has been validated by both Akcigit et al. (2017) and Giordano et al. (2015), who find a positive correlation between the variation in public sector efficiency across Italian provinces and the level of value added per employee.

In Table 3.9, column 1, we regress TFP growth on the interaction between *Government Dependence* and *Rule of Law*. We find that these variables have no significant effect on TFP growth. Similarly, we verify that our results from Table 3.9 are not sensitive to how we measure variation in institutional quality: in columns 2 and 3, we show that there is no substantial difference in the results when we use, $\Delta$*Control of Corruption, Judicial Inefficiency* or *Govt Inefficiency* (Chong et al., 2014; Djankov et al., 2003) as alternative measures of institutional quality.

In column 4, we include all four interaction effects, as well as *ICT Contribution* $\times$ *Country Meritocracy*. None of our measures of institutional quality, interacted with *Government Dependence*, have a statistically or economically significant effect on productivity growth, while the effect of the interaction term *ICT Contribution* $\times$ *Country Meritocracy*, is positive, statistically significant and broadly unchanged in magnitude with respect to our regression in Table 3.5.

Based on this additional analysis, we conclude that our main results are unlikely to be driven by an omitted variable linked to institutional quality.

## 3.6 Evidence from firm-level data

An even better way to ensure that our findings from Section 3.3 are not spurious is to try and corroborate them using firm-level data. To this purpose, we use Bruegel's EFIGE, which allows us to compute a firm-level measure of meritocratic management (*Firm Meritocracy*, see Section 3.2). Besides varying at the firm level, this measure has the advantage of reflecting factual information about firm characteristics, as opposed to perceptions, allowing us to rule out that the effects we measured in the aggregate data are driven by mismeasurement of the variable *Country Meritocracy*.

EFIGE also contains data that allows us to test the effect of labor market frictions on growth and to control in a much more careful way for human capital.

The use of this dataset, however, comes with limitations too. First, we do not observe ICT Capital at the firm level. More importantly, at the firm level specification (3.18) does follow from an accounting decomposition, but it is applied just by analogy. Last but not least, the dataset contains too few country/sector clusters to reproduce the IV specification we used for sector-level data. As a consequence, the firm level results should be interpreted as corroborating evidence, rather than formal tests of our hypothesis.

### 3.6.1 Firm Level Meritocracy

Figure 3.5 shows the distribution of firm-level meritocracy variable by country. Notice that for Italy the distribution of this variable across firms is significantly skewed to the right with respect to other countries in our sample. Almost half of the Italian firms in our sample score zero. Thus, Italy is an outlier according to this measure too.

Similarly, Figure 3.6 shows that *Firm Meritocracy* is highly correlated with *Country Meritocracy*, providing more credence to the sector-level results.

### 3.6.2 TFP growth regressions

As explained in Section 3.2, we compute annual growth rates in firm-level TFP in a way consistent with the EU KLEMS methodology, by using firm financials from the Amadeus-BvD dataset, for the period 2001-2007. If indeed meritocratic management mediates the productivity-enhancing effects of ICT adoption, we should observe, at the firm level, the same qualitative effect that we estimated in Section 3.3.

In Table 3.10, column 1, we reproduce a similar specification as in Table 3.5, column 1 using firm-level data. One difference with respect to the sector-level analysis is that sector-level TFP growth is replaced by firm-level TFP growth, and that *Country Meritocracy* is now replaced by *Firm Meritocracy*. In addition, the greater number of degrees of freedom allow us to control not just for country and sector fixed effects but for country-by-sector fixed effects. This allows us to control for potential reverse-causation of TFP on ICT capital accumulation even better than we did in our sector-level analysis with EU KLEMS data[6]. Because *Firm Meritocracy* is not absorbed by country $\times$ sector fixed

[6] In the reported table, we do not control for firm size since our only consistently-available measure of

effects, it is now also included as a standalone variable.

The estimates obtained from the EFIGE firm-level regressions are presented in Table 3.10. The interaction effect of *ICT Contribution* and *Firm Meritocracy* is positive and statistically significant, mimicking our findings with sector-level KLEMS data.

One of the advantages of the EFIGE firm-level dataset is that we can control more granularly for the effect of labor market frictions on growth. We do so by using the variable *Labor Frictions*, which is described in Section 3.2. Surprisingly, the coefficient of this variable is positive (not negative as expected) albeit not statistically significant.

At the firm level, one important determinant of the absorption of ICT is the amount of human capital per employee. We can control for this factor because EFIGE provides the share of employees who are college graduates. We add this variable, as well as its interaction with *ICT contribution*, as a control, in columns 3. Unsurprisingly, the variable *Employees with Degree* has a positive and statistically significant effect on TFP growth. However, when interacted with *ICT contribution*, it has a negative, statistically non-significant coefficient. Most importantly, inserting this variable does not change the effect of the interaction term between *Firm Meritocracy* and *ICT Contribution*.

### 3.6.3 Temporary workers and gerontocracy in the firm

The Italian labor market reforms of the late 90's and early 2000's, which we previously mentioned in subsection 3.5.5, might have contributed to Italy's productivity slowdown through a different channel. Daveri and Parisi (2015, henceforth DP), suggest that these reforms had the effect of increasing the incidence of temporary employment contracts , which in turn reduced the firms' incentives to invest in training. According to DP, this effect, combined with the elevated age of Italian CEOs, limited the ability of Italian firms to innovate, ultimately causing the productivity slowdown.

This hypothesis can potentially threaten identification in our econometric analysis if meritocratic management correlates with either the age of the CEO or the proclivity to use temporary employment contracts. We account for this alternative hypothesis by adding the percentage of temporary workers and the age of the firm's CEO (which EFIGE measures in decades) as control variables, to the regression of Table 3.10, column 3. To make sure that Firm Meritocracy is not actually capturing the effect of neither of these variables, we also interact them with *ICT Contribution*. The results are shown in the adjacent column 4.

---

size at the firm level is observed at the end of the panel, and therefore could be influenced by cross-firm differences in productivity growth. Nevertheless, to make sure that our results are robust, we repeat this set of regression in the Online Appendix, by controlling for the number of employees.

In contrast with the findings of DP, we find that the percentage of temporary workers does not have a statistically significant effect on productivity growth. CEO age has actually a positive and statistically significant effect on productivity growth. The estimated impact of meritocracy and its interaction with *ICT Contribution* remains broadly unchanged. Provided that our controls are not impacted by significant measurement error, we can therefore reasonably exclude that our findings of 3.6.2 are confounded by the effects described by DP.

### 3.6.4 ICT usage regressions

Because Italy does not appear to under-invest significantly in ICT capital, our argument is that its productivity slowdown is due to a lower ability to exploit these technologies. Using firm-level data, we can test whether this interpretation is consistent with the data. We do this by computing the variable *ICT Usage*, a firm-level score (ranging from 0 to 3) of the extent to which ICT technologies are utilized by the firm's management. The construction of this variable is outlined in Section 3.2.

In Table 3.11, column 1, we estimate an Ordered Probit regression of *ICT Usage* on the same set of explanatory variables as in Table 3.10. If the joint effect of *Firm Meritocracy* and *ICT Contribution* is mediated by the effective integration of ICTs in the firm's management, we would expect their interaction to predict higher values of *ICT Usage*.

We find that more meritocratic firms tend to use ICT more. This effect is more pronounced in sectors where the contribution of ICT capital was larger. Both effects are statistically significant. Based on these estimates, when a firm in a typical sector increases its level of meritocracy from 0 to 5, it doubles its probability of attaining a high level of *ICT Usage* (2 or 3), from 26.6% to 52%. The effect is even stronger in the more ICT-intensive sectors.

In Table 3.11, columns 2-3, we add, as a control variable, the percentage of employees with a college degree. This variable has a positive and statistically significant effect on *ICT Usage*, but its interaction with *ICT Contribution* does not. The coefficients of *Firm Meritocracy* and *ICT Contribution*, as well as their interaction, remain substantially unchanged.

In column 3, we add *CEO Age* and *Temporary Employees* as additional control variables, together with their interaction with *ICT Contribution*. The coefficients for these variables are not statistically different from zero, with the exception of the interaction term *Temporary Employees* $\times$ *ICT Contribution*, which has a p-value just below 10. The sign of the coefficient, however is the opposite of what we would expect given DP (positive rather than negative).

### 3.6.5 Imperfect competition, revenue and output productivity

In Sub-section 3.2.5, we warned that, while our firm-level measure of TFP is consistent with EU KLEMS methodology, it is susceptible to violations of the assumption of perfect competition. This is because we deflate value added using a sector-level index. If markets are not perfectly competitive and firms charge a markup, our measure of TFP will capture idiosyncratic variation in firm-level prices. As a consequence, it will be akin to revenue-based productivity (TFPR). This is problematic, because TFPR is known to capture a variety of factors that are unrelated to actual productivity (TFPQ), such as firm-level distortions (Hsieh and Klenow, 2009b)[7].

In order to make sure that our firm-level econometric results are not reliant on the assumption of perfect competition, we want to build an alternative (robust) firm-level measure of TFP growth that may not necessarily be consistent with the EU KLEMS approach. The simplest way to do that is to use firm-level output deflators. Unfortunately, while the EFIGE dataset contains plenty of information about management, workforce and IT usage, it falls short of providing firm-level price data. To correct for firm-level variation in prices, we therefore resort to an insight of De Loecker (2011) which allows us to do so by using sector-level prices alone: this requires imposing some structure on demand.

We follow the predominant practice in the literature and assume a CES demand system, which yields the following firm-level demand function:

$$Y_{it} = Y_{cst} \left( \frac{P_{it}}{P_{cst}} \right)^{-\sigma}$$

where the parameter $\sigma$ is the elasticity of substitution, and $Y_{cst}$ and $P_{cst}$ are the country/sector-level output and price indices. Rearranging this demand function yields the following expression for the (estimated) real log output growth at the firm level:

$$\Delta \hat{y}_{it} = \Delta y_{cst} + \frac{\sigma}{\sigma - 1} \left[ \Delta \log \left( P_{it} Y_{it} \right) - \Delta \log \left( P_{cst} Y_{cst} \right) \right] \tag{3.22}$$

which we can compute using firm-level value added in conjunction with sector-level value added (volume and price indices) from the EU KLEMS dataset.

Our dataset does not allow us to estimate the elasticity of substitution $\sigma$, therefore we use the conservative approach of inputing low values of $\sigma$ (similar to Hsieh and Klenow,

[7] Specifically, we worry about the possibility that our firm-level results might be biased if our TFP measure incorporates variations in markups, which would then become an omitted variable in the regression . The bias on the main coefficients would be positive if more meritocratic firms increased their markups (rather than their physical productivity), in more ICT-intensive sectors.

2009b). Notice that, as $\sigma$ becomes large (demand approaches perfect competition), output growth in equation (3.22) converges to our baseline TFP measure (value added deflated using sector-level price indices).

We use this estimate of firm-level output growth to compute an alternative measure of TFP:

$$\Delta \log \mathrm{TFP}_{it} = \Delta \hat{y}_{it} - \left(1 - \frac{W_{cst}L_{cst}}{P_{cst}Y_{cst}}\right)\Delta k_{it} - \left(\frac{W_{cst}L_{cst}}{P_{cst}Y_{cst}}\right)\Delta \ell_{it} \tag{3.23}$$

input volumes and shares for capital and labor are the same as in equation (3.13).

In the Online Appendix, we present alternative estimates of our regression of Table 3.10 where the dependent variable is TFP computed according to equation (3.23), using the values of the elasticity of substitution $\sigma = 5$ and $\sigma = 3$, which account for substantial deviations from perfect competition. Because we implement the regression in long-term differences, we can reasonably assume that short-term demand shocks are being averaged out. We also present additional estimations, in which we use a similar TFP measure, computed using the gross output concept (rather than value added).

### 3.6.6 Sample Selection in BvD-Amadeus

The EFIGE dataset is built out of the stratified sample of firms that received the EFIGE survey. It is equipped by its authors with sampling weights which ensure that when we use survey data EFIGE is representative of the population of manufacturing firms.

By contrast, when EFIGE is matched with firm financials obtained from the Amadeus dataset, it inherits the sample selection issues of Amadeus. To address this problem, every time financial information is employed, we use the methodology developed by Pellegrino and Zheng (2017) to generate new sampling weights that make the sample representative.

However, in order to completely rule out the possibility that our results are driven by sample selection, we replicate our firm-level TFP regressions in the Online Appendix by dropping the three countries for which sample selection might be an issue (Austria, Germany and the UK), and show that the regression estimates are virtually unchanged. For all other countries, sample coverage is close to 100%.

## 3.7 Distortions to competition and meritocracy in the firm

When we look at the decade ending in 1995, it appears that this loyalty-based management style had no negative consequences on Italy's TFP growth. By contrast, with the advent of the ICT revolution, the lower ability of the loyalty-based system of translating ICT investments into productivity seems to have cost Italy between 13 and 16 percentage points of TFP growth (see Sub-section 3.3.8).

If this is the case, why did Italian firms fail to adopt superior managerial techniques? To be more specific, how can we explain the persistence of the loyalty model of management in Italy, given its cost in terms of lack of TFP growth?

One explanation could be hysteresis. To use a metaphor from genetics, up to the 1980s the loyalty-based management style was simply a neutral mutation. When the advantages of meritocracy came about, Italian firms were slow to adapt. This explanation has the advantage of containing the hope that, in the long run, the adaptation will take place, even absent policy interventions.

A more rational (but less optimistic) interpretation is that in Italy, even today, there are some advantages to adopting the loyalty-based management system which offset (or partially offset) the inability to fully exploit the ICT revolution. If this were the case, then convergence in the long run might not occur without a policy intervention.

But what are the advantages of loyalty-based management? Caselli and Gennaioli (2005, 2013), for example, argue that allocating power to cronies rather than talented managers can be individually efficient (while socially inefficient) in the presence of credit frictions and/or lack of product market competition. An alternative explanation is that loyalty-based management might better function in environments where legal enforcement is either inefficient or unavailable. Among developed countries, Italy stands out for its lack of competition in the banking sector, its inefficient legal system (the average time to enforce a contract, as measured by Djankov et al. (2003) is 638 days, nearly 2.5 times the cross-country average) and for the diffusion of tax evasion and bribes (in 2017, it ranked 60th in Transparency International's Corruption Perceptions Index, behind every other country in our sample).

Thus, a reasonable hypothesis is that, at the onset of the ICT revolution, Italy found itself with the optimal level of management for its institutions, but the worst possible type for taking advantage of this revolution. To corroborate this hypothesis, we need to find a way to measure the differential benefit of being loyalty-based in Italy.

To this end, we use another set of variables from the EFIGE survey. Specifically, we use the firms' answers to a multiple-choice question in which they are asked to identify the main factors constraining the firm's growth. We focus on three most cited constraints, namely: financial constraints, labor regulation, and bureaucracy. In Table 3.12, we estimate, using a probit model, the conditional probability that the firm encounters each of these constraints. Beside sector fixed effects, the key explanatory variables are the firm's of meritocratic management score, and its interaction with a dummy for Italy.

As expected, more meritocratic firms face fewer constraints (of any kind). However, this effect is not present in Italy. The interaction between the meritocracy index and the Italy dummy is very similar in magnitude, but opposite in sign, to the baseline coefficient of meritocracy. Interestingly, this interaction effect for Italy is significant for financial constraints and bureaucratic constraints, but not for labor market constraints. This difference makes a lot of sense. Loyal management can exchange favors with banks and bypass bureaucracy through political connections or bribes, but finds it more difficult to overcome the constraints that labor regulation puts on growth. These results are obviously not hard proof that loyalty-based management is advantageous in Italy, but they are consistent with this hypothesis. Moreover, our results resonate with the findings of Akcigit et al. (2017), which focus on another channel that appears to worsen business dynamics in Italy - namely, the role of political connections.

We know from Demsetz (1983) that more efficient firms tend to grow larger. Thus, if meritocratic firms tend to be more efficient in other countries, but not in Italy, we expect firm size to be positively correlated with meritocratic practices at the firm level in general, but not in Italy. We test this hypothesis in column 4 of Table 3.12. We find that indeed, on average, there is a positive and statistically significant correlation between firm size meritocratic practices at the firm level. When we interact meritocracy with the Italian dummy, however, the coefficient is negative and statistically significant. Thus, in Italy meritocratic practices are less correlated with firm size than in the rest of the sample. In fact, we add the two coefficients, we find that in Italy meritocracy and size are not significantly positively correlated. In sum, in Italy loyalty-based management seems to pay off.

## 3.8 Conclusions

Economists have long tried to identify the institutional causes of economic development (e.g., Acemoglu et al., 2001a). By and large, this analysis has treated institutions as enabling conditions (i.e., enforcement of property rights) rather than as inputs in the production function. One consequence of this choice is that the institutional factors enabling development are independent of the technology used.

Our diagnosis of the Italian disease suggests we should start studying how technological change interacts with pre-existing institutions. As technological change can be skill-biased, it can also be biased in favor of certain institutional arrangements (e.g., more formalized). If we accept Williamson's (2000) characterization of institutions as slow-moving, a higher technological demand for certain institutional features will not immediately produce the desired institutional change. Thus, transparency-biased technological change will foster growth in countries with more transparency and meritocracy, and will delay it in countries where informality and cronyism prevail.

In this paper we argue that the ICT revolution represents an example of transparency-biased technological change, which favors countries with more meritocratic institutions and more objective incentive-schemes. We show that – given existing management practices – this bias can go a long way towards explaining not only the Italian productivity slowdown, but the productivity slowdown of Southern Europe in general. We also show that in Italy, loyalty-based management is not necessarily a leftover of the past. Even today, un-meritocratic managerial practices provide a comparative advantage in the Italian institutional environment.

We do not attempt here to prescribe a medicine for the Italian disease. Entrepreneurs lose out collectively from an environment that is less prone to the adoption of new technology; yet, they lose out individually from adopting transparency-biased technologies when their peers are not. Given this conundrum and the fact that country's institutions are intrinsically hard to change, it appears that Italy serves as a cautionary tale of the importance of building institutions that aren't simply appropriate at one historical juncture, but that are also attuned to the pace of technological progress.

**Figure 3.1:** Aggregate labor productivity in selected countries (1974-2016)

This chart shows GDP per hour worked for USA, Germany, France and Italy in 1974-2016 in 2010 US$.

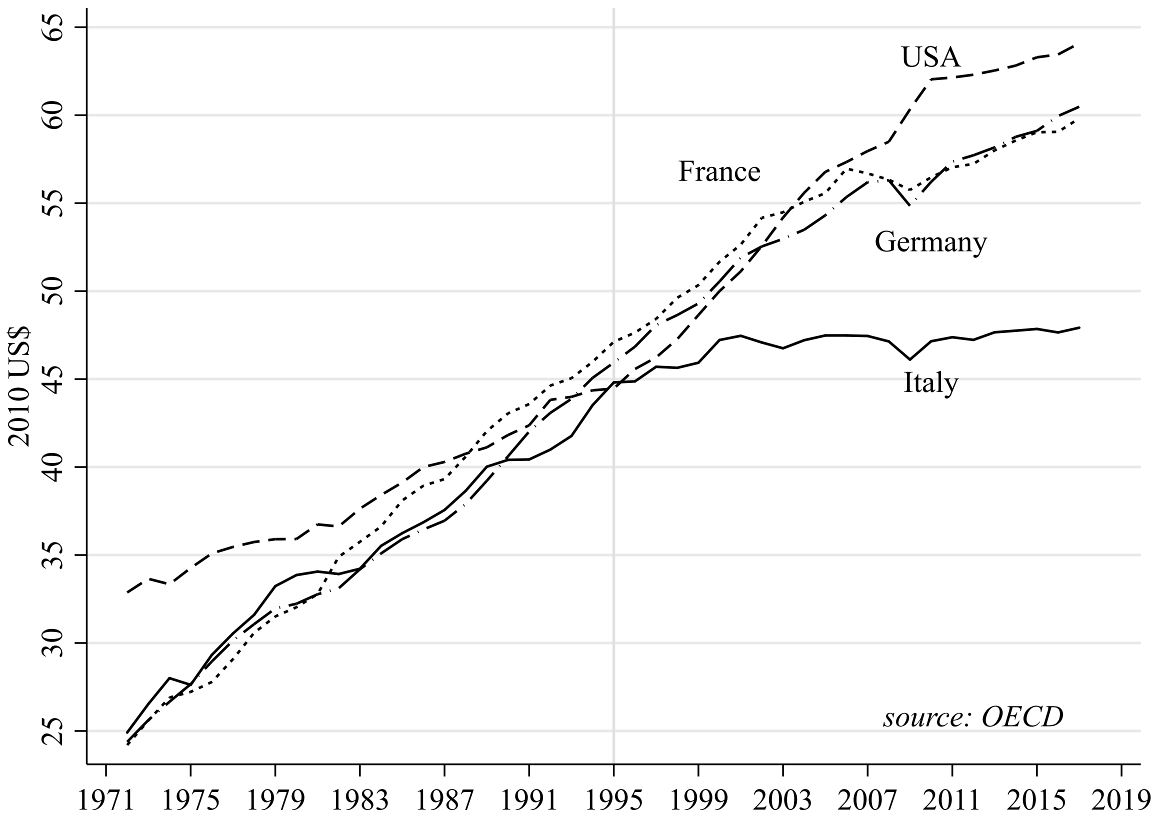

**Figure 3.2:** Decomposition of labor productivity growth (unweighted, 1996-2006)

This chart shows the breakdown of log growth in GDP per hour worked at constant prices between 1996 and 2006 into its four components: TFP growth and the contributions of ICT capital, non-ICT capital and labor composition. For this chart we use industry-level data in the business sector. Growth across sectors is unweighted, in order to factor out the sectoral composition of the economy.

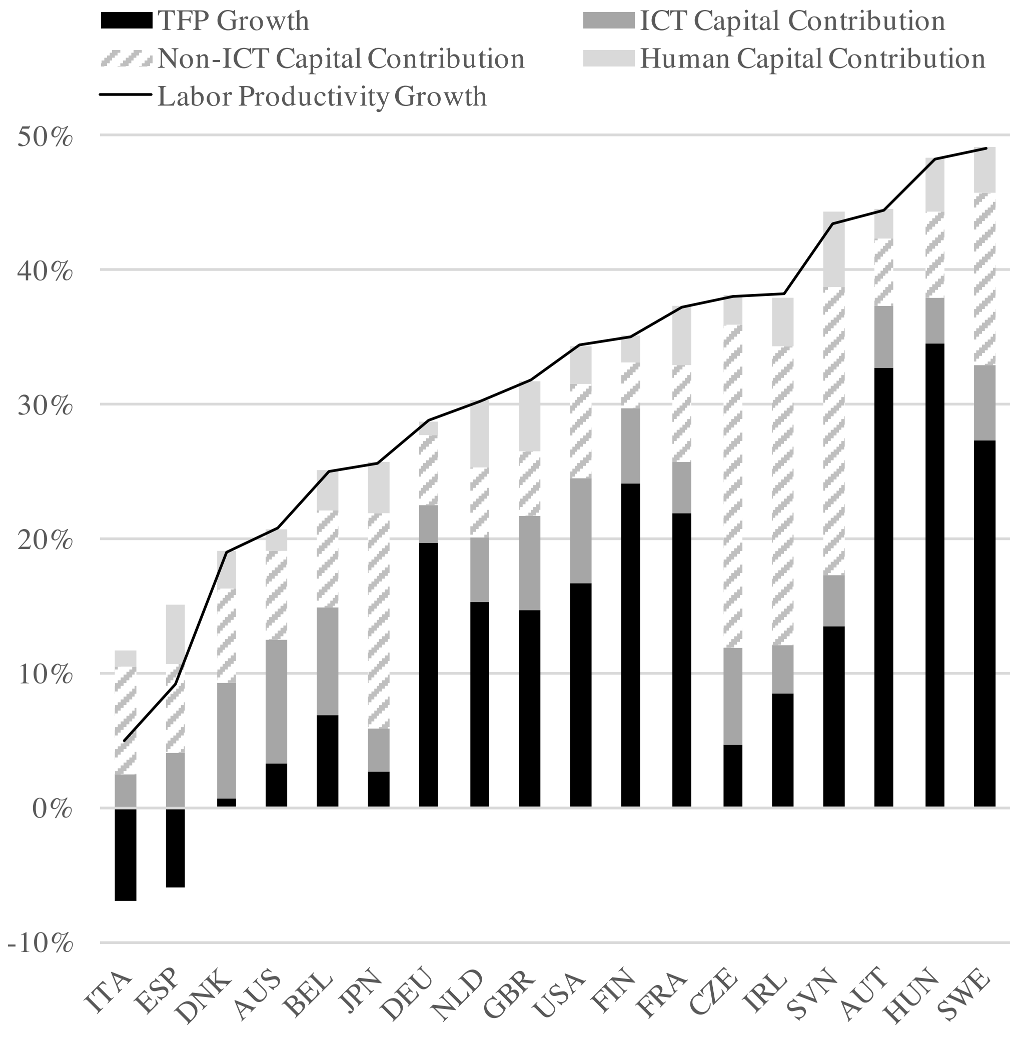

**Figure 3.3:** Productivity growth by country Meritocracy and sector ICT intensiveness

This figure displays the evolution of TFP estimates, indexed at 1995, from the EU KLEMS database for different country/sector groups. We sort high-Meritocracy versus low-Meritocracy countries (top tercile versus bottom tercile based on our country-level measure of meritocracy) and high ICT intensiveness versus low ICT intensiveness sectors (top eight versus bottom eight sectors based on the sector-level, cross-country average contribution of ICT capital to output growth in 1996–2006). We take the median TFP growth rate for each group/year. giving equal weight to all country/sectors. Czech Republic, Hungary and Slovenia are excluded since there is no TFP data for these countries before 1995.

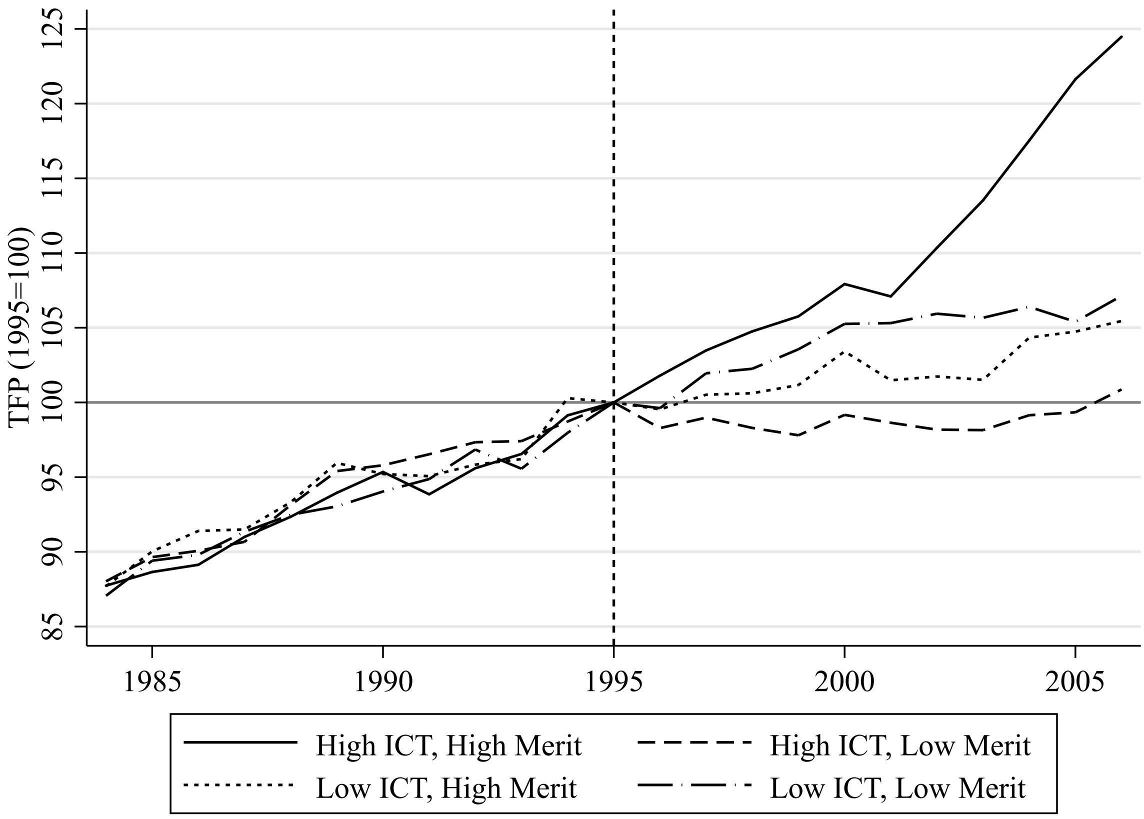

**Figure 3.4:** Government dependence scores

This chart depicts values of the variable *Govt Dependence*, built using news count data from Dow Jones' Factiva News Search service. We exploit the Factiva topic and industry "tags". *Govt Dependence* is defined, for each sector, as the share of news articles having the topic tag "Government Contracts" or "Regulation/Government Policy". We use all news articles from Dow Jones, the Financial Times, Reuters, and the Wall Street Journal published from January $1^{st}$ 1984 to December $31^{st}$ 2017.

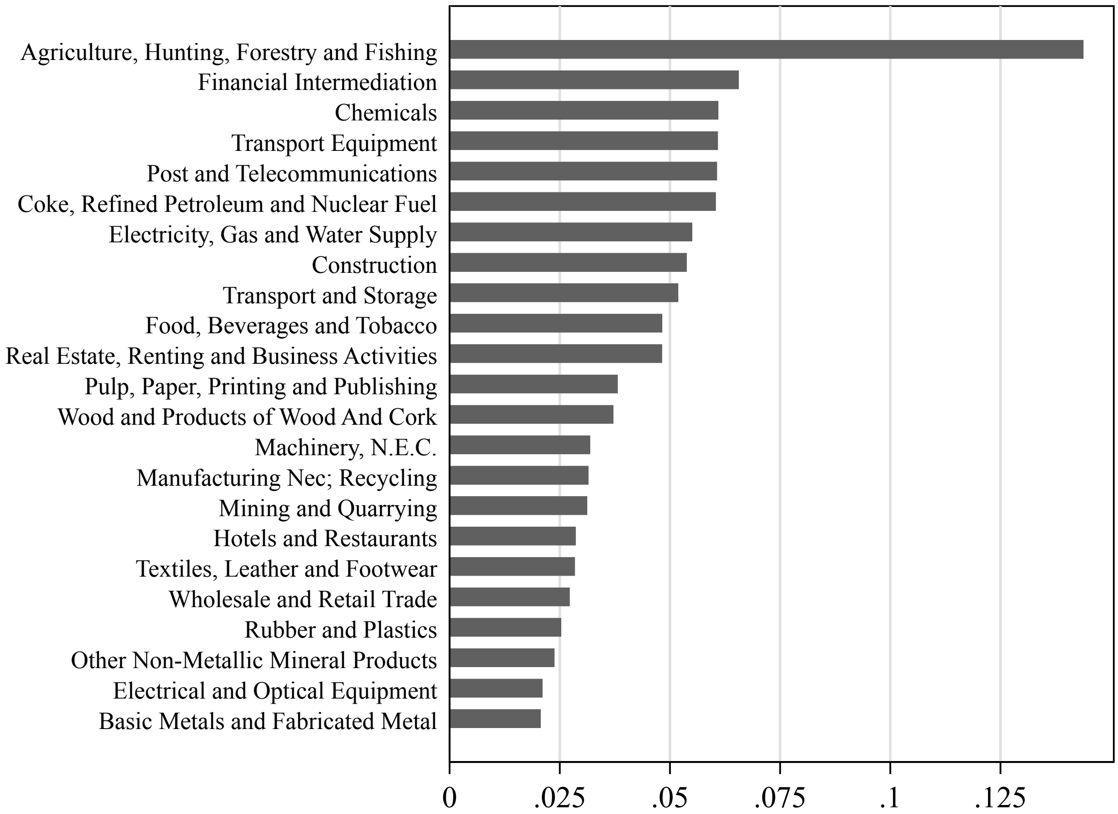

**Figure 3.5:** Distribution of firm-level Meritocracy

The figure below displays histograms, by countries and for the whole sample, of firm-level meritocracy. Observations are weighted using the sampling weights of the EFIGE survey in order to obtain consistent population estimates of the distribution of the Meritocracy index.

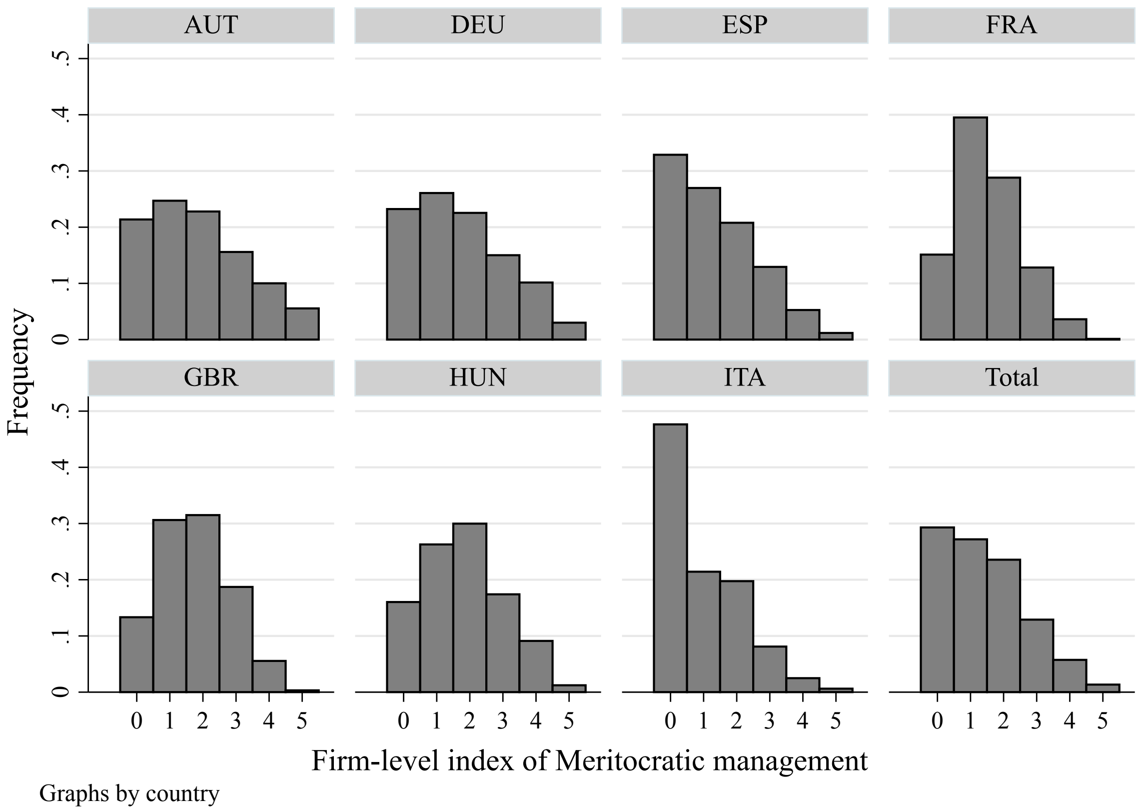

**Figure 3.6:** Firm-level and country-level Meritocracy

The following figure plots of our country-level measure of meritocratic management, derived from WEF surveys, against our firm-level meritocratic management metric, constructed from firm-level EFIGE survey data. The latter is averaged at the level of the country of headquarters. To account for the fact that all the firms in our sample are operating in Austria, France, Germany, Hungary, Italy, Spain or the UK, the Firm-level score is adjusted by including a dummy variable for these 7 countries on the right hand side of the regression equation here depicted. The effect of the dummy is summed to these firms' meritocracy score. Countries that are represented by fewer than 10 firms in the EFIGE dataset are excluded.

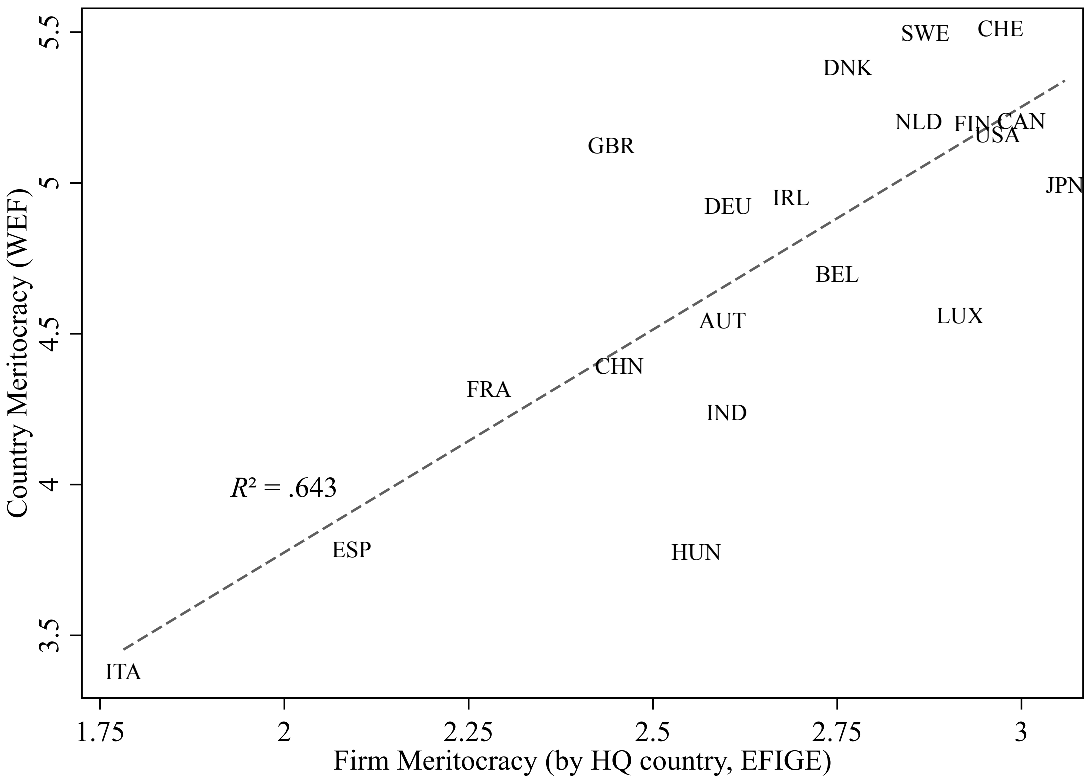

**Table 3.1:** Variables Descriptions

| Variable | Description | Source |
|---|---|---|
| *Bureaucratic Frictions* | Dummy equal to one if the firm selects "Bureaucracy/Government Regulation" when prompted to "indicate the main factors that hamper the growth of your firm." | Bruegel-Unicredit EU-EFIGE Dataset |
| *CEO Age* | Age of current CEO/company head in years, grouped into seven categories: $<$25, 26-35, 36-45, 46-55, 56-65, 66-75, $>$75. | Bruegel-Unicredit EU-EFIGE Dataset |
| *China Exposure* | Predicted effect of China exports growth on domestic output, by country and sector. Computed assuming that the effect of China export growth is symmetric across all competitor countries. See Section 2 for derivation. | World Input/Output Database |
| *Country Meritocracy* | Average of three Global Competitiveness Report Expert Surveys (2012): 1) "In your country, who holds senior management positions?" [1 = usually relatives or friends without regard to merit; 7 = mostly professional managers chosen for merit and qualifications]; 2) "In your country, how do you assess the willingness to delegate authority to subordinates?" [1 = not willing at all – senior management makes all important decisions; 7 = very willing – authority is mostly delegated to business unit heads and other lower-level managers]; and 3) "In your country, to what extent is pay related to employee productivity?" [1 = not at all; 7 = to a great extent]. | World Economic Forum, 2012 |
| *Employees with degree* | (Firm-reported) Share of the firm's workforce that are university graduates. If the percentage of employees with a college degree is not reported, but the absolute level is reported, we compute the percentage ourselves from the absolute figures, dividing the number of employees with degree by the total number of employees. | Bruegel-Unicredit EU-EFIGE Dataset |

| | | |
|---|---|---|
| *Employment Laws* | Composite Index of Strictness of Employment Laws. Obtained by Botero et al. (2004) combining measures of difficulty of hiring, rigidity of hours, difficulty of redundancy, and redundancy costs (in weeks of salary). | Botero et al. (2004) |
| *Financial Constraints* | Dummy equal to one if the firm selects "Financial Constraints" when prompted to "indicate the main factors that hamper the growth of your firm." | Bruegel-Unicredit EU-EFIGE Dataset |
| *Firm Meritocracy* | Takes on integers 0–5. It is the sum of the affirmative answers to the following questions: 1) "Can managers make autonomous decisions in some business areas?" 2) "Are managers incentivized with financial benefits?" 3) "Has any of your executives worked abroad for at least one year?" 4) "Is the firm not directly or indirectly controlled by an individual or family-owned entity? If it is, was the CEO recruited from outside the firm?" 5) "Is the share of managers related to the controlling family lower than 50%?". If the percentage of managers affiliated with the controlling family is not reported, we use 1 minus the percentage of managers not affiliated with the controlling family (if this is reported). If this is also missing, but the absolute levels are reported, we compute the percentage ourselves from the absolute figures. | Bruegel-Unicredit EU-EFIGE Dataset |
| *Government Dependence* | Ratio of government-related news to total sector news in a pool of articles from Dow Jones, Financial Times, Reuters, and the Wall Street Journal from 1984 to 2017. We define as government-related news items that have at least one of the following subject tags in the Factiva news database: 1) government policy/regulation, 2) government aid, 3) government contracts. | Factiva News Search |
| *Government Inefficiency* | Average number of days needed for the authors of Chong et al. (2014) to get back a letter sent to an inexistent address in a certain country. | Chong et al. (2014) |

| | | |
|---|---|---|
| *ICT Contribution* | Average yearly contribution of ICT (Information and Communication Technologies) capital to value added growth in 1996–2006. It is defined as the two-period average compensation share of capital in value added (estimated by subtracting labor compensation from value added) times the ICT assets share of capital compensation (estimated using current rental prices), times the rate of growth in ICT capital (estimated through a perpetual inventory model). | EU KLEMS |
| *ICT Infrastructure* | Infrastructure component of the 2012 Networked Readiness Index. It is computed by the World Economic Forum using country data on mobile network coverage, the number of secure internet servers, internet bandwidth, and electricity production. | World Economic Forum, 2012 |
| *ICT Usage* | Sum of "YES" answers to the following three EFIGE survey questions on whether the firm has access to/uses: 1) IT systems for internal information management; 2) IT systems for e-commerce; 3) IT systems for management of the sales/purchase network | Bruegel-Unicredit EU-EFIGE Dataset |
| *Judicial Inefficiency* | Estimate of the number of days required to enforce a contract. Average of the estimate for "cashing a bounced check" and "evicting a tenant". | Djankov et al. (2003) |
| *Non-ICT Contribution* | Average yearly contribution of non-ICT (Information and Communication Technologies) capital to value added growth in 1996–2006. It is defined as the two-period average compensation share of capital in value added (estimated by subtracting labor compensation from value added) times the non-ICT assets share of capital compensation (estimated using current rental prices), times the rate of growth in non-ICT capital (estimated through a perpetual inventory model). | EU KLEMS |
| *Labor Frictions* | Dummy equal to one if the firm selects "Labor Market Regulation" when prompted to "indicate the main factors that hamper the growth of your firm." | Bruegel-Unicredit EU-EFIGE Dataset |

| | | |
|---|---|---|
| *US Layoff Rate* | Mass layoff rates for US sector. Computed by Bassanini and Garnero (2013) using the CPS biennial Displaced Workers Supplement (2000–2006, even years). | Bassanini and Garnero (2013) |
| *Management Schools* | Average of Global Competitiveness Report Expert Survey (2012): "In your country, how do you assess the quality of business schools? [1 = extremely poor – among the worst in the world; 7 = excellent – among the best in the world]" | World Economic Forum, 2012 |
| *Shadow Economy* | Shadow Economy, percent of GDP (average in 1999–2006). Estimated by the authors using a latent variable, Multiple Indicators Multiple Causes (MIMIC) model. | Schneider (2012) |
| *Temporary Employees* | (Firm-reported) Percentage of employees which, in 2008, have worked for the firm with a fixed-term contract. | Bruegel-Unicredit EU-EFIGE Dataset |
| *Trade Openness* | Sector-level exports (Domestic value added embodied in foreign final demand) plus imports (Foreign value added embodied in domestic final demand), divided by value added. All variables measured in 1995 in millions US$. | OECD-WTO TiVA Dataset |
| *Control of Corruption* | Average yearly change in Control of Corruption Index, from the Worldwide Governance Indicators (time series sourced through the Quality of Government OECD dataset) | World Bank |
| *logTFP* | Average log growth of total factor productivity growth over a certain period: 1996-2006 for sector-level data and 2001-2007 for firm-level data, unless otherwise noted. It is estimated as the residual growth in value added at constant prices after subtracting the contributions of capital and of the labor services (see Section 3.2 for more information). For firm-level data, we use output/input elasticities and deflators for added value and labor input from the EU KLEMS dataset, as well as capital deflators from the OECD Structural Analysis (StAn) dataset. | sector-level: EU KLEMS firm-level: Bruegel-Unicredit EU-EFIGE, EU KLEMS and OECD. |
| *Rule of Law* | Average yearly change in Rule of Law Index, from the Worldwide Governance Indicators (time series sourced through the Quality of Government OECD dataset) | World Bank |

**Table 3.2:** Descriptive statistics

We present here summary statistics for our main variables, sorted by their level of variation (firm, country, sector). Additional variables (used for robustness tests) are presented in the the Online Appendix.

Panel A: Variables that vary across countries and sectors (1996-2006)

| Variable | Obs | Mean | SD | Min | Max |
|---|---|---|---|---|---|
| China Exposure | 414 | 0.012 | 0.021 | -0.001 | 0.193 |
| ICT Contribution | 414 | 0.005 | 0.006 | -0.005 | 0.055 |
| Non-ICT Contribution | 414 | 0.008 | 0.013 | -0.028 | 0.095 |
| Trade Openness | 414 | 0.897 | 0.849 | 0.017 | 8.116 |
| $\Delta$logTFP$_{96-06}$ | 414 | 0.012 | 0.036 | -0.292 | 0.204 |

Panel B: Variables that vary across countries

| Variable | Obs | Mean | SD | Min | Max |
|---|---|---|---|---|---|
| Country Meritocracy | 18 | 4.683 | 0.635 | 3.387 | 5.504 |
| Employment Laws | 18 | 0.535 | 0.201 | 0.164 | 0.745 |
| Employment Protection | 18 | 2.153 | 0.747 | 0.260 | 3.310 |
| Firm Size | 17 | 18.129 | 10.284 | 6.183 | 39.289 |
| Govt Inefficiency | 18 | 94.256 | 41.955 | 16.200 | 173.400 |
| ICT Infrastructure | 18 | 5.894 | 0.708 | 4.317 | 6.904 |
| Management Schools | 18 | 5.109 | 0.645 | 3.963 | 6.121 |
| Shadow Economy | 18 | 0.172 | 0.055 | 0.086 | 0.270 |
| $\Delta$Control of Corruption | 18 | -0.003 | 0.020 | -0.034 | 0.027 |
| $\Delta$Rule of Law | 18 | 0.002 | 0.021 | -0.063 | 0.023 |

Panel C: Variables that vary across EU KLEMS sectors

| Variable | Obs | Mean | SD | Min | Max |
|---|---|---|---|---|---|
| Govt Dependence | 23 | 0.045 | 0.024 | 0.020 | 0.126 |
| US Layoff Rate | 20 | 0.052 | 0.017 | 0.022 | 0.090 |

Panel D: Variables that vary across firms

| Variable | Obs | Mean | SD | Min | Max |
|---|---|---|---|---|---|
| Bureaucratic Frictions | 12,444 | 0.208 | 0.406 | 0.000 | 1.000 |
| CEO Age | 14,701 | 4.254 | 1.038 | 1.000 | 7.000 |
| Employees with degree | 14,749 | 0.094 | 0.134 | 0.000 | 1.000 |
| Financial Frictions | 12,444 | 0.341 | 0.474 | 0.000 | 1.000 |
| Firm Meritocracy | 14,205 | 1.554 | 1.272 | 0.000 | 5.000 |
| ICT Usage | 14,756 | 1.262 | 0.935 | 0.000 | 3.000 |
| Labor Frictions | 12,444 | 0.190 | 0.392 | 0.000 | 1.000 |
| Temporary employees | 14,640 | 0.256 | 0.385 | 0.000 | 1.000 |
| $\Delta$logTFP$_{01-07}$ | 9,880 | 0.004 | 0.150 | -2.301 | 2.355 |

**Table 3.3:** Decomposition of labor productivity growth, by country

This table presents the breakdown, at the country level, of the log growth in GDP per hour worked at constant prices between 1996 and 2006 into its four components: TFP growth and the contributions of ICT capital, non-ICT capital and human capital. For this table, we use industry-level data in the business sector. Growth across sectors is unweighted, in order to factor out the sectoral composition of the economy.

| Country | TFP Growth | ICT Capital Contribution | Non-ICT Capital Contribution | Human Capital Contribution | Labor Productivity Growth |
|---|---|---|---|---|---|
| AUS | 3.4% | 9.2% | 6.6% | 1.6% | 20.8% |
| AUT | 32.7% | 4.7% | 4.9% | 2.3% | 44.5% |
| BEL | 7.0% | 7.9% | 7.3% | 2.9% | 25.1% |
| CZE | 4.7% | 7.2% | 24.1% | 2.1% | 38.1% |
| DEU | 19.7% | 2.9% | 5.1% | 1.1% | 28.8% |
| DNK | 0.6% | 8.6% | 7.1% | 2.8% | 19.1% |
| ESP | -6.0% | 4.1% | 6.7% | 4.4% | 9.2% |
| FIN | 24.2% | 5.5% | 3.5% | 2.0% | 35.1% |
| FRA | 22.0% | 3.7% | 7.3% | 4.5% | 37.3% |
| GBR | 14.6% | 7.1% | 4.8% | 5.2% | 31.8% |
| HUN | 34.6% | 3.3% | 6.4% | 4.1% | 48.3% |
| IRL | 8.5% | 3.6% | 22.2% | 3.7% | 38.2% |
| ITA | -6.8% | 2.5% | 7.9% | 1.3% | 5.0% |
| JPN | 2.6% | 3.3% | 16.0% | 3.8% | 25.7% |
| NLD | 15.3% | 4.8% | 5.2% | 5.1% | 30.3% |
| SVN | 13.5% | 3.8% | 21.4% | 5.6% | 43.4% |
| SWE | 27.4% | 5.5% | 12.8% | 3.4% | 49.0% |
| USA | 16.7% | 7.8% | 7.0% | 2.8% | 34.4% |
| Average ex.Italy | 14.2% | 5.5% | 9.9% | 3.4% | 32.9% |
| Difference vs. Italy | 21.1% | 3.0% | 2.0% | 2.1% | 28.0% |

**Table 3.4:** Decomposition of labor productivity growth, by sector

This table presents the breakdown, at the sector level, of the log growth in GDP per hour worked at constant prices between 1996 and 2006 into its four components: TFP growth and the contributions of ICT capital, non-ICT capital and human capital. For this table, we use industry-level data in the business sector. Growth across sectors is unweighted, in order to factor out the sectoral composition of the economy.

| Sector | Code | TFP Growth | ICT Capital Contribution | Non-ICT Capital Contribution | Human Capital Contribution | Labor Productivity Growth |
|---|---|---|---|---|---|---|
| Agriculture, Hunting, Forestry and Fishing | 01t05 | 26.6% | 0.8% | 8.6% | 3.4% | 39.4% |
| Mining and Quarrying | 10t14 | -1.0% | 2.6% | 19.8% | 1.7% | 23.1% |
| Food, Beverages and Tobacco | 15t16 | -0.9% | 3.5% | 10.6% | 3.6% | 16.7% |
| Textiles, Leather and Footwear | 17t19 | 17.5% | 2.8% | 8.6% | 5.5% | 34.5% |
| Wood and Products of Wood And Cork | 20 | 18.8% | 2.5% | 6.2% | 3.7% | 31.3% |
| Pulp, Paper, Printing and Publishing | 21t22 | 13.8% | 8.3% | 11.5% | 3.7% | 37.3% |
| Coke, Refined Petroleum and Nuclear Fuel | 23 | -41.6% | 6.1% | 28.5% | 3.9% | -4.8% |
| Chemicals and Chemical Products | 24 | 16.3% | 4.7% | 21.6% | 2.7% | 45.4% |
| Rubber and Plastics | 25 | 26.4% | 2.5% | 8.8% | 3.6% | 41.4% |
| Other Non-Metallic Mineral Products | 26 | 20.2% | 3.9% | 9.5% | 3.4% | 37.0% |
| Basic Metals and Fabricated Metal | 27t28 | 14.2% | 2.5% | 4.4% | 3.6% | 24.9% |
| Machinery, Nec | 29 | 25.4% | 4.4% | 10.5% | 4.0% | 44.3% |
| Electrical and Optical Equipment | 30t33 | 63.8% | 7.3% | 12.4% | 4.1% | 87.9% |
| Transport Equipment | 34t35 | 30.8% | 3.7% | 9.9% | 3.6% | 48.0% |
| Manufacturing Nec; Recycling | 36t37 | 12.5% | 2.8% | 5.4% | 4.8% | 25.6% |
| Electricity, Gas and Water Supply | 40t41 | 10.6% | 4.7% | 18.9% | 1.4% | 35.4% |
| Construction | 45 | -4.8% | 1.7% | 2.8% | 2.5% | 2.3% |
| Wholesale and Retail Trade | 50t52 | 16.0% | 6.0% | 7.4% | 2.5% | 31.8% |
| Hotels and Restaurants | 55 | -10.5% | 2.3% | 3.4% | 2.2% | -2.4% |
| Transport and Storage | 60t63 | 3.2% | 4.9% | 6.8% | 2.3% | 17.2% |
| Post and Telecommunications | 64 | 36.0% | 22.8% | 11.9% | 2.9% | 73.5% |
| Financial Intermediation | 65t67 | 17.4% | 16.5% | 0.2% | 3.9% | 38.0% |
| Real Estate, Renting and Business Activities | 70t74 | -10.9% | 4.9% | -2.6% | 2.2% | -6.7% |

**Table 3.5:** Sector-level TFP-ICT Regressions

This table displays estimation results of Ordinary Least Squares (OLS) and Instrumental Variable (IV) regressions of sector-level total factor productivity growth from the EU KLEMS dataset on *ICT Contribution*, *Non-ICT Contribution* and interaction terms. In all regressions, the left-side variable is log TFP growth, averaged over a 11-year period. In columns 1-3, TFP growth is computed over the period 1996–2006; in columns 4-5, it is computed over the 1985-1995 period. Each observation represents a country-sector. *Country Meritocracy* varies at the country level. $\log TFP$, *ICT Contribution* and *Non-ICT Contribution* vary at the country/sector level.

| | (1) $\Delta \log TFP_{96-06}$ OLS | (2) $\Delta \log TFP_{96-06}$ OLS | (3) $\Delta \log TFP_{96-06}$ IV | (4) $\Delta \log TFP_{85-95}$ OLS | (5) $\Delta \log TFP_{85-95}$ OLS |
|---|---|---|---|---|---|
| ICT Contribution | -5.247** | | -5.411*** | 0.414 | -7.277 |
| | (2.151) | | (1.285) | (6.171) | (4.943) |
| ICT Contribution × Country Meritocracy | 1.094** | | 0.977*** | 0.030 | 1.167 |
| | (0.510) | | (0.325) | (1.201) | (1.044) |
| Non-ICT Contribution | | 0.525 | | | |
| | | (2.171) | | | |
| Non-ICT Contribution × Country Meritocracy | | -0.077 | | | |
| | | (0.444) | | | |
| $R^2$ | 0.350 | 0.339 | 0.344 | 0.153 | 0.162 |
| Kleibergen-Paap underid. test P-value | | | 0.000 | | |
| Sargan-Hansen overid. test P-value | | | 0.406 | | |
| Wu-Hausman exogeneity test P-value | | | 0.351 | | |
| Observations | 414 | 414 | 414 | 345 | 345 |
| Lagged ICT Contribution ('85-'95) | | | | | ✓ |
| Country Fixed Effects | ✓ | ✓ | ✓ | ✓ | ✓ |
| Sector Fixed Effects | ✓ | ✓ | ✓ | ✓ | ✓ |

Robust Standard Errors in Parentheses     * p<.10, ** p<.05, *** p<.01

**Table 3.6:** Sector-level TFP-ICT Regressions with additional country-level covariates

This table displays estimation results of Ordinary Least Squares (OLS) regressions of sector-level total factor productivity growth from the EU KLEMS dataset on *ICT Contribution* and interaction terms. In all regressions, the left-side variable is log TFP growth, averaged over the 1996–2006 period. Each observation represents a country-sector. *Country Meritocracy*, *ICT Infrastructure*, *Shadow Economy* and *Management Schools* vary at the country level. $\log TFP$ and *ICT Contribution* vary at the country/sector level.

| | (1) | (2) | (3) | (4) | (5) |
|---|---|---|---|---|---|
| | $\Delta \log TFP_{96-06}$ | $\Delta \log TFP_{96-06}$ | $\Delta \log TFP_{96-06}$ | $\Delta \log TFP_{96-06}$ | $\Delta \log TFP_{96-06}$ |
| | OLS | OLS | OLS | OLS | OLS |
| ICT Contribution | -5.247** | -5.461 | 0.552 | -3.278* | -10.540 |
| | (2.151) | (3.359) | (1.016) | (1.676) | (6.545) |
| ICT Contribution × Country Meritocracy | 1.094** | | | | 2.272** |
| | (0.510) | | | | (0.910) |
| ICT Contribution × ICT Infrastructure | | 0.876 | | | -0.424 |
| | | (0.599) | | | (0.818) |
| ICT Contribution × Shadow Economy | | | -4.480 | | 12.270 |
| | | | (4.699) | | (8.774) |
| ICT Contribution × Management Schools | | | | 0.612* | 0.004 |
| | | | | (0.364) | (0.405) |
| $R^2$ | 0.350 | 0.345 | 0.340 | 0.344 | 0.355 |
| Observations | 414 | 414 | 414 | 414 | 414 |
| Country Fixed Effects | ✓ | ✓ | ✓ | ✓ | ✓ |
| Sector Fixed Effects | ✓ | ✓ | ✓ | ✓ | ✓ |
| Robust Standard Errors in Parentheses | | | * p<.10, ** p<.05, *** p<.01 | | |

**Table 3.7:** Sector-level TFP-China regressions

This table displays estimation results of Ordinary Least Squares (OLS) and Instrumental Variable (IV) regressions of sector-level total factor productivity growth from the EU KLEMS data set on *China Exposure*, *US Layoff Rate* and interaction terms. In all regressions, the left-side variable is log TFP growth, averaged over a the 1996–2006 period. Each observation represents a country-sector. *Employment Laws*, and *Employment Protection* vary at the country level. *US Layoff Rate* varies at the sector level. $\log TFP$ and *China Exposure* vary at the country/sector level.

| | (1) $\Delta \log TFP_{96-06}$ OLS | (2) $\Delta \log TFP_{96-06}$ OLS | (3) $\Delta \log TFP_{96-06}$ OLS | (4) $\Delta \log TFP_{96-06}$ IV | (5) $\Delta \log TFP_{96-06}$ OLS |
|---|---|---|---|---|---|
| US Layoff Rate × Employment Laws | | | | | -0.082 (0.375) |
| China Exposure | 0.040 (0.060) | 0.034 (0.183) | -0.059 (0.271) | 0.243 (0.480) | |
| China Exposure × Employment Laws | | 0.012 (0.325) | | 1.086 (1.193) | |
| China Exposure × Employment Protection | | | 0.047 (0.120) | | |
| $R^2$ | 0.337 | 0.337 | 0.337 | 0.204 | 0.409 |
| Kleibergen-Paap underid. test P-value | | | | 0.002 | |
| Wu-Hausman exogeneity test P-value | | | | 0.044 | |
| Observations | 414 | 414 | 414 | 414 | 360 |
| Country Fixed Effects | ✓ | ✓ | ✓ | ✓ | ✓ |
| Sector Fixed Effects | ✓ | ✓ | ✓ | ✓ | ✓ |

Robust Standard Errors in Parentheses * $p<.10$, ** $p<.05$, *** $p<.01$

**Table 3.8:** Sector-level TFP-Trade Openness regressions

This table displays estimation results of Ordinary Least Squares (OLS) regressions of sector-level total factor productivity growth from the EU KLEMS dataset on *Trade Openness* and interaction terms. In all regressions, the left-side variable is log TFP growth, averaged over the 1996–2006 period. Each observation represents a country-sector. *Employment Laws* and *Employment Protection* vary at the country level. $\log TFP$, *Trade Openness* and *ICT Contribution* vary at the country/sector level.

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| | $\Delta \log TFP_{96-06}$ | $\Delta \log TFP_{96-06}$ | $\Delta \log TFP_{96-06}$ | $\Delta \log TFP_{96-06}$ |
| | OLS | OLS | OLS | OLS |
| ICT Contribution | | | -5.841*** | -5.785*** |
| | | | (1.993) | (1.920) |
| ICT Contribution × Country Meritocracy | | | 1.219** | 1.206** |
| | | | (0.500) | (0.486) |
| Trade Openness | -0.002 | 0.018** | 0.020** | 0.017 |
| | (0.006) | (0.008) | (0.008) | (0.020) |
| Trade Openness × Employment Protection | | | | -0.009 |
| | | | | (0.011) |
| Trade Openness × Employment Laws | | -0.039* | -0.041** | |
| | | (0.020) | (0.021) | |
| $R^2$ | 0.338 | 0.359 | 0.376 | 0.361 |
| Observations | 414 | 414 | 414 | 414 |
| Country Fixed Effects | ✓ | ✓ | ✓ | ✓ |
| Sector Fixed Effects | ✓ | ✓ | ✓ | ✓ |
| Robust Standard Errors in Parentheses | * p<.10, ** p<.05, *** p<.01 | | | |

**Table 3.9:** TFP-government effectiveness regressions

This table displays estimation results of Ordinary Least Squares (OLS) regressions of sector-level total factor productivity growth from the EU KLEMS dataset on *ICT Contribution* and interaction terms. In all regressions, the left-side variable is log TFP growth, averaged over 1995–2006 period. Each observation represents a country-sector. *Country Meritocracy*, *Rule of Law*, *Control of Corruption*, *Control of Corruption* vary at the country level. $\log TFP$ and *Trade Openness* vary at the country/sector level.

| | (1) | (2) | (3) | (4) | (5) |
|---|---|---|---|---|---|
| | $\Delta\log TFP_{96-06}$ | $\Delta\log TFP_{96-06}$ | $\Delta\log TFP_{96-06}$ | $\Delta\log TFP_{96-06}$ | $\Delta\log TFP_{96-06}$ |
| | OLS | OLS | OLS | OLS | OLS |
| ICT Contribution | | | | | -4.954** |
| | | | | | (2.245) |
| ICT Contribution × Country Meritocracy | | | | | 1.045** |
| | | | | | (0.525) |
| Govt Dependence × ΔRule of Law | 1.512 | | | | 1.305 |
| | (1.939) | | | | (2.845) |
| Govt Dependence × ΔControl of Corruption | | 2.145 | | | 1.018 |
| | | (2.597) | | | (2.794) |
| Govt Dependence × Govt Inefficiency | | | 0.001 | | 0.002 |
| | | | (0.001) | | (0.001) |
| Govt Dependence × Judicial Inefficiency | | | | -0.000 | -0.000 |
| | | | | (0.000) | (0.000) |
| $R^2$ | 0.337 | 0.337 | 0.338 | 0.340 | 0.356 |
| Observations | 414 | 414 | 414 | 414 | 414 |
| Country Fixed Effects | ✓ | ✓ | ✓ | ✓ | ✓ |
| Sector Fixed Effects | ✓ | ✓ | ✓ | ✓ | ✓ |

Robust Standard Errors in Parentheses * p<.10, ** p<.05, *** p<.01

**Table 3.10:** Firm-level productivity regressions

This table displays estimation results of ordinary least squares (OLS) regressions of firm-level total factor productivity growth computed using Amadeus data in the EFIGE dataset. In all regressions, the left-side variable is log TFP growth averaged over 2001–2007. Every data point is a firm. The variable ICT Contribution, which comes from the EU KLEMS dataset, varies at the country/sector level. The explanatory variable Firm Meritocracy ranges from 0 to 5, and is constructed using firm-level information from the EFIGE survey. The variable CEO Age is categorical: a unit increment represents a 10-year increase in the age of the firm's CEO. The variables Employees with Degree and Temporary Employees are expressed as a percentage of the firm's labor force and are part of the EFIGE survey response data. Labor Frictions is a dummy that varies at the firm level. Observations are weighted to ensure that the regression sample is representative.

| | (1)<br>$\Delta$logTFP$_{01-07}$<br>OLS | (2)<br>$\Delta$logTFP$_{01-07}$<br>OLS | (3)<br>$\Delta$logTFP$_{01-07}$<br>OLS | (4)<br>$\Delta$logTFP$_{01-07}$<br>OLS |
|---|---|---|---|---|
| Firm Meritocracy | -0.001<br>(0.002) | -0.001<br>(0.002) | -0.002<br>(0.002) | -0.002<br>(0.002) |
| Firm Meritocracy × ICT Contribution | 2.181***<br>(0.695) | 2.123***<br>(0.687) | 2.355***<br>(0.724) | 2.413***<br>(0.730) |
| Employees with degree | | | 0.055**<br>(0.023) | 0.057**<br>(0.023) |
| Employees with degree × ICT Contribution | | | -8.445<br>(8.163) | -8.522<br>(8.175) |
| CEO Age | | | | 0.004**<br>(0.002) |
| CEO Age × ICT Contribution | | | | -1.204<br>(0.837) |
| Temporary employees | | | | -0.001<br>(0.008) |
| Temporary employees × ICT Contribution | | | | -0.298<br>(2.854) |
| Labor Frictions | | 0.002<br>(0.004) | | |
| $R^2$ | 0.034 | 0.038 | 0.035 | 0.036 |
| Observations | 9,486 | 7,309 | 9,482 | 9,437 |
| Country × Sector Fixed Effects | ✓ | ✓ | ✓ | ✓ |

Robust Standard Errors in Parentheses * p<.10, ** p<.05, *** p<.01

**Table 3.11:** Firm-level ICT Usage regressions

This table displays estimation results of ordered probit regressions of firm-level ICT Usage, from the EFIGE survey (2009). In all regressions, the left-side variable is a firm-level measure of ICT usage, which ranges from 0 to 3 and which we compute using information from the EFIGE survey. The variable ICT Contribution, which comes from the EU KLEMS dataset, varies at the country/sector level. The explanatory variable Firm Meritocracy ranges from 0 to 5, and is constructed using firm-level information from the EFIGE survey. The variable CEO Age is categorical: a unit increment represents a 10-year increase in the age of the firm's CEO. The variables Employees with Degree and Temporary Employees are expressed as percentage of the firm's labor force and are part of the EFIGE survey response data. Observations are weighted to ensure that the regression sample is representative.

| | (1)<br>ICT Usage<br>O.Probit | (2)<br>ICT Usage<br>O.Probit | (3)<br>ICT Usage<br>O.Probit |
|---|---|---|---|
| Firm Meritocracy | 0.127***<br>(0.013) | 0.113***<br>(0.013) | 0.112***<br>(0.013) |
| Firm Meritocracy × ICT Contribution | 13.078**<br>(5.177) | 12.358**<br>(5.244) | 12.170**<br>(5.276) |
| Employees with degree | | 0.770***<br>(0.119) | 0.811***<br>(0.121) |
| Employees with degree × ICT Contribution | | -29.676<br>(33.180) | -31.024<br>(33.318) |
| CEO Age | | | 0.011<br>(0.014) |
| CEO Age × ICT Contribution | | | -5.174<br>(6.694) |
| Temporary employees | | | -0.047<br>(0.068) |
| Temporary employees × ICT Contribution | | | 47.695*<br>(24.359) |
| Observations | 14,204 | 14,196 | 14,058 |
| Country × Sector Fixed Effects | ✓ | ✓ | ✓ |

Robust Standard Errors in Parentheses * p<.10, ** p<.05, *** p<.01

**Table 3.12:** Meritocracy and Misallocation

In this table, we regress firm-level variables from the EFIGE survey on our firm-level metric of meritocratic management and its interaction with a dummy for Italian firms. Column (1)-(3) present Probit estimates of a regression in which the dependent variables are firm-level dummies representing the firms' answers to the question "Indicate the main factors preventing the growth of your firm" from the EFIGE survey (firms may indicate more than one choice). Column (4) presents OLS estimates of a regression in which the dependent variable is the logarithm of the number of employees of the company. The explanatory variable Firm Meritocracy ranges from 0 to 5, and is constructed using firm-level information from the EFIGE survey. "Italy" is the dummy variable identifying Italian firms. Observations are weighted to ensure that the regression sample is representative.

| | (1) | (2) | (3) | (4) |
|---|---|---|---|---|
| | Financial constraints | Labor Frictions | Bureaucratic Frictions | log Employees |
| | Probit | Probit | Probit | OLS |
| Italy | -0.135 | 0.364 | 0.242 | 0.114* |
| | (0.213) | (0.450) | (0.399) | (0.049) |
| Firm Meritocracy | -0.059** | -0.090** | -0.075*** | 0.288*** |
| | (0.027) | (0.042) | (0.026) | (0.020) |
| Firm Meritocracy × Italy | 0.063** | 0.059 | 0.075*** | -0.114*** |
| | (0.028) | (0.043) | (0.028) | (0.020) |
| $R^2$ | | | | 0.161 |
| Sector Fixed Effects | ✓ | ✓ | ✓ | ✓ |
| Standard errors clustering variable | Country | Country | Country | Country |

Robust Standard Errors in Parentheses * p<.10, ** p<.05, *** p<.01

# Bibliography


Acemoglu, D., P. Aghion, and F. Zilibotti (2006): "Distance to frontier, selection, and economic growth," *Journal of the European Economic Association*, 4, 37–74.

Acemoglu, D., V. M. Carvalho, A. Ozdaglar, and A. Tahbaz-Salehi (2012): "The network origins of aggregate fluctuations," *Econometrica*, 80, 1977–2016.

Acemoglu, D., S. Johnson, and J. A. Robinson (2001a): "The colonial origins of comparative development: An empirical investigation," *American economic review*, 91, 1369–1401.

——— (2001b): "The Colonial Origins of Comparative Development: An Empirical Investigation," *American Economic Review*, 91, 1369–1401, 10.1257/aer.91.5.1369.

Acemoglu, D., A. Ozdaglar, and A. Tahbaz-Salehi (2017): "Microeconomic origins of macroeconomic tail risks," *American Economic Review*, 107, 54–108.

Akcigit, U., S. Baslandze, and F. Lotti (2017): "Connecting to Power: Political Connections, Innovation, and Firm Dynamics."

Altomonte, C., and T. Aquilante (2012): "The EU-EFIGE/Bruegel-Unicredit dataset,"Technical report, No. 2012/13. Bruegel working paper.

Autor, D., D. Dorn, L. F. Katz, C. Patterson, and J. Van Reenen (2020): "The fall of the labor share and the rise of superstar firms," *The Quarterly Journal of Economics*, 135, 645–709.

Axtell, R. L. (2001): "Zipf Distribution of US Firm Sizes," *Science*, 1818–1820.

Baker, J. B., and T. F. Bresnahan (1985): "The gains from merger or collusion in product-differentiated industries," *The Journal of Industrial Economics*, 427–444.

Ballester, C., A. Calvó-Armengol, and Y. Zenou (2006): "Who's who in networks. Wanted: The key player," *Econometrica*, 74, 1403–1417.

Bandiera, O., L. Guiso, A. Prat, and R. Sadun (2008): *Italian managers: fidelity or performance? (in T. Boeri, A. Merlo, and A. Prat. The ruling class: Management and politics in modern Italy)*: Oxford University Press.

Baqaee, D. R., and E. Farhi (2020): "Productivity and misallocation in general equilibrium," *The Quarterly Journal of Economics*, 135, 105–163.

Barro, R. J., and J. W. Lee (2013): "A new data set of educational attainment in the world, 1950–2010," *Journal of development economics*, 104, 184–198.

Bassanini, A., and A. Garnero (2013): "Dismissal protection and worker flows in OECD countries: Evidence from cross-country/cross-industry data," *Labour Economics*, 21, 25–41.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile prices in market equilibrium," *Econometrica: Journal of the Econometric Society*, 841–890.

Bloom, N., R. Sadun, and J. V. Reenen (2012a): "Americans Do IT Better: US Multinationals and the Productivity Miracle," *American Economic Review*, 102, 167–201.

Bloom, N., R. Sadun, and J. Van Reenen (2012b): "The organization of firms across countries," *Quarterly Journal of Economics*, 127, 1663–1705.

Bloom, N., M. Schankerman, and J. Van Reenen (2013): "Identifying technology spillovers and product market rivalry," *Econometrica*, 81, 1347–1393.

Bloom, N., and J. Van Reenen (2007): "Measuring and explaining management practices across firms and countries," *Quarterly Journal of Economics*, 122, 1351–1408.

Bonacich, P. (1987): "Power and centrality: A family of measures," *American journal of sociology*, 92, 1170–1182.

Bond, S., A. Hashemi, G. Kaplan, and P. Zoch (2020): "Some Unpleasant Markup Arithmetic: Production Function Elasticities and their Estimation from Production Data,"Technical report, National Bureau of Economic Research.

Botero, J. C., S. Djankov, R. L. Porta, F. Lopez-De-Silanes, and A. Shleifer (2004): "The Regulation of Labor," *Quarterly Journal of Economics*, 119, 4.

Bowen, D. E. I., L. Frésard, and G. Hoberg (2018): "Technological disruptiveness and the evolution of IPOs and sell-outs."

Bresnahan, T. F., E. Brynjolfsson, and L. M. Hitt (2002): "Information technology, workplace organization, and the demand for skilled labor: Firm-level evidence," *Quarterly Journal of Economics*, 117, 339–376.

Brynjolfsson, E., and L. M. Hitt (2003): "Computing productivity: Firm-level evidence," *Review of Economics and Statistics*, 85, 793–808.

Brynjolfsson, E., L. M. Hitt, and S. Yang (2002): "Intangible assets: Computers and organizational capital," *Brookings Papers on Economic Activity*, 2002, 137–198.

Bugamelli, M., and P. Pagano (2004): "Barriers to Investment in ICT," *Applied Economics*, 36, 2275–2286.

Calligaris, S., M. Del Gatto, F. Hassan, G. I. Ottaviano, F. Schivardi et al. (2016): "Italy's productivity conundrum. a study on resource misallocation in Italy,"Technical report, Directorate General Economic and Financial Affairs (DG ECFIN), European Commission.

Carvalho, V. M., and B. Grassi (2019): "Large firm dynamics and the business cycle," *American Economic Review*, 109, 1375–1425.

Carvalho, V. M., M. Nirei, Y. Saito, and A. Tahbaz-Salehi (2020): "Supply chain disruptions: Evidence from the great east japan earthquake," *The Quarterly Journal of Economics, forthcoming.*

Carvalho, V. M., and A. Tahbaz-Salehi (2019): "Production networks: A primer," *Annual Review of Economics*, 11, 635–663.

Caselli, F., and N. Gennaioli (2005): "Credit constraints, competition, and meritocracy," *Journal of the European Economic Association*, 3, 679–689.

——— (2013): "Dynastic management," *Economic Inquiry*, 51, 971–996.

Chen, H., L. Cohen, and D. Lou (2016): "Industry window dressing," *The Review of Financial Studies*, 29, 3354–3393.

Chong, A., R. LaPorta, F. L. de Silanes, and A. Shleifer (2014): "Letter Grading Government Efficiency," *Journal of the European Economic Association*, 12, 277–299.

Cunningham, C., F. Ederer, and S. Ma (2018): "Killer acquisitions," *Available at SSRN 3241707.*

Daruich, D., S. Di Addario, and R. Saggio (2018): "The Effects of Partial Employment Protection Reforms: Evidence from Italy."

Daveri, F., and M. L. Parisi (2015): "Experience, innovation, and productivity: empirical evidence from Italy's slowdown," *ILR Review*, 68, 889–915.

De Loecker, J. (2011): "Product differentiation, multiproduct firms, and estimating the impact of trade liberalization on productivity," *Econometrica*, 79, 1407–1451.

De Loecker, J., J. Eeckhout, and G. Unger (2020): "The rise of market power and the macroeconomic implications," *The Quarterly Journal of Economics*, 135, 561–644.

De Loecker, J., and F. Warzynski (2012): "Markups and firm-level export status," *American economic review*, 102, 2437–71.

Decker, R., J. Haltiwanger, R. Jarmin, and J. Miranda (2014): "The role of entrepreneurship in US job creation and economic dynamism," *Journal of Economic Perspectives*, 28, 3–24.

Demsetz, H. (1983): "The structure of ownership and the theory of the firm," *The Journal of law and economics*, 26, 375–390.

Djankov, S., R. La Porta, F. Lopez-de Silanes, and A. Shleifer (2002): "The regulation of entry," *The Quarterly Journal of Economics*, 117, 1–37.

Djankov, S., R. L. Porta, F. L. de Silanes, and A. Shleifer (2003): "Courts," *Quarterly Journal of Economics*, 453–517.

Doidge, C., K. M. Kahle, G. A. Karolyi, and R. M. Stulz (2018): "Eclipse of the public corporation or eclipse of the public markets?" *Journal of Applied Corporate Finance*, 30, 8–16.

Edmond, C., V. Midrigan, and D. Y. Xu (2018): "How Costly Are Markups?"Technical report, National Bureau of Economic Research.

Einav, L., and J. Levin (2010): "Empirical industrial organization: A progress report," *Journal of Economic Perspectives*, 24, 145–62.

Eisfeldt, A. L., and D. Papanikolaou (2013): "Organization capital and the cross-section of expected returns," *The Journal of Finance*, 68, 1365–1406.

Epple, D. (1987): "Hedonic prices and implicit markets: estimating demand and supply functions for differentiated products," *Journal of political economy*, 95, 59–80.

Fan, Y. (2013): "Ownership consolidation and product characteristics: A study of the US daily newspaper market," *American Economic Review*, 103, 1598–1628.

Fernandez-Villaverde, J., and L. Ohanian (2018): "The Lack of European Productivity Growth: Causes and Lessons for the US,"Technical report, Penn Institute for Economic Research, Department of Economics, University of ....

Foster, L., C. Grim, J. Haltiwanger, and Z. Wolf (2016): "Firm-Level Dispersion in Productivity: Is the Devil in the Details?" *American Economic Review*, 106, 95–98.

Fruchterman, T. M., and E. M. Reingold (1991): "Graph drawing by force-directed placement," *Software: Practice and experience*, 21, 1129–1164.

Gao, X., J. R. Ritter, and Z. Zhu (2013): "Where have all the IPOs gone?" *Journal of Financial and Quantitative Analysis*, 48, 1663–1692.

Garicano, L., and P. Heaton (2010): "Information Technology, Organization, and Productivity in the Public Sector: Evidence from Police Departments," *Journal of Labor Economics*, 28, 167–201.

Giordano, R., S. Lanau, P. Tommasino, and P. Topalova (2015): "Does Public Sector Inefficiency Constrain Firm Productivity; Evidence from Italian Provinces,"Technical report, International Monetary Fund.

Gopinath, G., Ş. Kalemli-Özcan, L. Karabarbounis, and C. Villegas-Sanchez (2017): "Capital allocation and productivity in South Europe," *The Quarterly Journal of Economics*, qjx024.

Gray, W. B. (1987): "The cost of regulation: OSHA, EPA and the productivity slowdown," *The American Economic Review*, 77, 998–1006.

Gros, D. (2011): "What is holding Italy back?" *Vox. Research-based policy analysis and commentary.*

Grullon, G., Y. Larkin, and R. Michaely (2018): "Are US industries becoming more concentrated?" *Forthcoming, Review of Finance.*

Hall, R. E., and C. I. Jones (1999): "Why Do Some Countries Produce So Much More Output Per Worker Than Others?" *Quarterly Journal of Economics*, 83–116.

Hassan, F., and G. Ottaviano (2013): "Productivity in Italy: The great unlearning," *VoxEU. org*, 30.

Hoberg, G., and G. Phillips (2016): "Text-based network industries and endogenous product differentiation," *Journal of Political Economy*, 124, 1423–1465.

Hoberg, G., G. Phillips, and N. Prabhala (2014): "Product market threats, payouts, and financial flexibility," *The Journal of Finance*, 69, 293–324.

Hopenhayn, H. A. (1992): "Entry, exit, and firm dynamics in long run equilibrium," *Econometrica: Journal of the Econometric Society*, 1127–1150.

Hsieh, C.-T., and P. J. Klenow (2009a): "Misallocation and Manufacturing TFP in China and India," *The Quarterly Journal of Economics*, 124, 1403–1448.

——— (2009b): "Misallocation and manufacturing TFP in China and India," *Quarterly Journal of Economics*, 124, 1403–1448.

Istat (2005): "L'economia sommersa e il lavoro non regolare," in *Audizione del Presidente dell'Istituto nazionale di statistica, Luigi Biggeri presso le Commissioni riunite V Commissione "Programmazione economica, bilancio" del Senato e V Commissione "Bilancio" della Camera.*

Jorgenson, D. W., F. M. Gollop, and B. M. Fraumeni (1987): "Productivity and US Economic Growth."

Jorgenson, D. W., M. Ho, and K. Stiroh (2005): "Productivity, Volume 3: Information Technology and the American Growth Resurgence,"Technical report, The MIT Press.

Kahle, K. M., and R. M. Stulz (2017): "Is the US public corporation in trouble?" *Journal of Economic Perspectives*, 31, 67–88.

Kalemli-Ozcan, S., B. Sorensen, C. Villegas-Sanchez, V. Volosovych, and S. Yesiltas (2015): "How to construct nationally representative firm level data from the ORBIS global database,"Technical report, National Bureau of Economic Research.

Katz, L. (1953): "A new status index derived from sociometric analysis," *Psychometrika*, 18, 39–43.

Kaufmann, D., A. Kraay, and M. Mastruzzi (2006): *Governance matters V: aggregate and individual governance indicators for 1996-2005* Volume 4012: World Bank Publications.

Khan, L. (2018): "The new Brandeis movement: America's antimonopoly debate."

Krugman, P. (2012): "More about Italy," *New York Times*, 28.

La Porta, R., F. Lopez-de Silanes, A. Shleifer, and R. Vishny (1999): "The quality of government," *Journal of Law, Economics, and Organization*, 15, 222–279.

Lancaster, K. J. (1966): "A New Approach to Consumer Theory," *Journal of Political Economy*, 74, 132–157.

Levinsohn, J., and A. Petrin (2003): "Estimating production functions using inputs to control for unobservables," *The review of economic studies*, 70, 317–341.

Martin, S., W. M. Brown, R. Klavans, and K. W. Boyack (2011): "OpenOrd: an open-source toolbox for large graph layout," *Visualization and Data Analysis 2011*, 7868, 786806.

Melitz, M. J. (2003): "The impact of trade on intra-industry reallocations and aggregate industry productivity," *Econometrica*, 71, 1695–1725.

Milana, C., and A. Zeli (2004): "Productivity Slowdown and the Role of ICT in Italy," *The Economic Impact of ICT*, 261–277.

Monderer, D., and L. S. Shapley (1996): "Potential games," *Games and economic behavior*, 14, 124–143.

Nevo, A. (2001): "Measuring market power in the ready-to-eat cereal industry," *Econometrica*, 69, 307–342.

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

O'Mahony, M., and M. P. Timmer (2009): "Output, input and productivity measures at the industry level: The EU KLEMS database," *Economic Journal*, 119, F374–F403.

Pellegrino, B., and G. Zheng (2017): "What is the extent of Misallocation?".

Peters, R. H., and L. A. Taylor (2017): "Intangible capital and the investment-q relation," *Journal of Financial Economics*, 123, 251–272.

Pierce, J. R., and P. K. Schott (2016): "The surprisingly swift decline of US manufacturing employment," *American Economic Review*, 106, 1632–1662.

Posner, R. A. (1975): "The Social Costs of Monopoly and Regulation," *The Journal of Political Economy*, 807–827.

Rajan, R. G., and L. Zingales (1998): "Financial Dependence and Growth," *American Economic Review*, 88, 559–586.

Romer, P. M. (1986): "Increasing returns and long-run growth," *Journal of Political Economy*, 94, 1002–1037.

Rosen, S. (1974): "Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition," *The Journal of Political Economy*, 82, 34–55.

Schivardi, F., and T. Schmitz (2017): *The ICT Revolution and Italy's Two Lost Decades*: IGIER WP.

Schneider, F. (2012): "The shadow economy and work in the shadow: what do we (not) know?".

Syverson, C. (2019): "Macroeconomics and Market Power: Context, Implications, and Open Questions," *Journal of Economic Perspectives*, 33, 23–43.

Tiffin, A. (2014): "European Productivity, Innovation and Competitiveness: The Case of Italy," working paper, IMF.

Timmer, M. P., E. Dietzenbacher, B. Los, R. Stehrer, and G. J. Vries (2015): "An illustrated user guide to the world input–output database: the case of global automotive production," *Review of International Economics*, 23, 575–605.

Traina, J. (2018): "Is aggregate market power increasing? production trends using financial statements," working paper, Chicago Booth Stigler Center.

Ushchev, P., and Y. Zenou (2018): "Price competition in product variety networks," *Games and Economic Behavior*, 110, 226–247.

Werden, G. J. (2018): "Back to School: What the Chicago School and New Brandeis School Get Right."

Williamson, O. E. (2000a): "The new institutional economics: taking stock, looking ahead," *Journal of economic literature*, 38, 595–613.

——— (2000b): "The New Institutional Economics: Taking Stock, Looking Ahead," *Journal of Economic Literature*, 38, 595–613, 10.1257/jel.38.3.595.

Wollmann, T. G. (2018): "Trucks without bailouts: Equilibrium product characteristics for commercial vehicles," *American Economic Review*, 108, 1364–1406.

——— (2019): "Stealth consolidation: Evidence from an amendment to the Hart-Scott-Rodino Act," *American Economic Review: Insights*, 1, 77–94.

Wooldridge, J. M. (2009): "On estimating firm-level production functions using proxy variables to control for unobservables," *Economics Letters*, 104, 112–114.