© 2019 Cesare Buiatti

ESSAYS ON INTERNATIONAL MACROECONOMICS

BY

CESARE BUIATTI

DISSERTATION

Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy in Economics
in the Graduate College of the
University of Illinois at Urbana-Champaign, 2019

Urbana, Illinois

Doctoral Committee:

Assistant Professor Minchul Shin, Chair
Associate Professor Stephen Parente, Co-Chair
Professor Anne P. Villamil, University of Iowa
Assistant Professor Gregory Howard

ABSTRACT

This dissertation is composed by three chapters presenting different topics in international macroeconomics. The first chapter focuses on the role of private information in the dynamics of sovereign debt yields. I propose a model of sovereign debt and default where information is incomplete: Investors receive private noisy signals about the current state of the economy, resulting in heterogeneous information sets across the investors and the government. I show that having a large enough signal noise is a sufficient condition for uniqueness of equilibrium. The main empirical contribution of this chapter is proposing and implementing a structural estimation strategy for the private information noise. Using forecasts data about real GDP growth in the euro area, available since 1999, I measure private information noise at a quarterly frequency, by insuring that the informational structure of my model implies statistics of forecast dispersion and uncertainty consistent with those observed in the data. Private information noise in the euro area shows two interesting characteristics: It peaks during crises and it has remained persistently larger than before since the Great Recession. I calibrate my model to be consistent with observed moments of the euro area economy, and I show it is successful in accounting key untargeted statistics of sovereign spreads. By means of counterfactual exercises, I assess the impact that a change in private information noise has on the spreads' statistics. First of all, I investigate what are the implications of the private information noise having remained at the lower levels prevalent before the Great Recession. I find that spreads would have been on average lower and less volatile. Then, I simulate my model for various levels of private information noise, and this exercise uncovers a hump-shaped relation between private information noise and the average spread level, and

the spread volatility: While spread levels and volatility are initially increasing with private information noise, for large enough levels of the latter the relation becomes negative.

The second chapter, written in collaboration with Joao B. Duarte and Luis Felipe Sáenz, studies the divergence in aggregate labor productivity between Europe and the U.S., observed since the 1990's, from a sectoral standpoint. In particular, we are interested in explaining differences in sectoral labor productivity levels in the service sector, by large the most important industry in both European and U.S. economies. An issue in accounting for sectoral contributions to aggregate labor productivity is that the sectoral composition of the economy is endogenous with respect to sectoral productivity dynamics. To tackle this endogeneity, we employ a model of structural transformation portraying an economy with thirteen different sectors, corresponding to agriculture, manufacturing and eleven service industries. By means of our model, we measure cross-country comparable productivity levels in the U.S. and in eight European countries. Then, through a set of counterfactual exercises we identify which sectors are mainly accountable for the falling behind of European labor productivity. We find that the most serious problems are in the service sector, and, more specifically, in wholesale and retail trade, business services, and financial services. Finally, we decompose our measures of labor productivity in total factor productivity and contributions stemming from traditional and ICT capital. We uncover that the underperforming sectors in Europe experienced a significant fall in ICT capital endowment per hour worked with respect to the U.S. Also, total factor productivity accounts for a large share of labor productivity in these services, suggesting that the ultimate cause of the European falling behind has to be found in limitations to TFP growth.

The last chapter, written in collaboration with Luis Felipe Sáenz, is about the Ricardo Effect, namely the substitution of capital for labor in the production process at the expense of workers' income, when cheaper capital becomes available. We employ a unique plant-level longitudinal dataset for Colombian manufacturing establishments for the period 1982-1998 to document the existence and quantify the Ricardo Effect in the Colombian manufacturing industry. Moving from a theory of

production based on a CES technology, we estimate that the elasticity of substitution between capital and labor is significantly larger than one. This finding proves the existence of the Ricardo Effect. We take into account the fact that Colombia went through a profound transformation in the early 1990's, by the introduction of important market-oriented reforms. Our empirical investigation points out that these reforms did not significantly alter the elasticity of substitution between capital and labor. However, they accelerated the fall in the price of capital relative to the price of labor, which has been occurring since 1986. We extrapolate the average decline in the relative price of capital between 1986 and 1998, and, given our preferred estimate of the elasticity of substitution in the production, we conclude that the Ricardo Effect can account for half of the reduction in the labor income share observed in Colombia between 1994 and 2014.

To Emma

ACKNOWLEDGMENTS

We are living exceptional times with regard to the role of high education in society. Notwithstanding the fact that high education is increasingly central in the path to economic growth, the challenges posed by the transformation of the global economy has lead the Western society to look at the so called "experts" with a suspicion and a lack of confidence just unimaginable a decade or so ago. At the same time, the development of the social networks provided for a stage in which everyone may pretend - or genuinely believe - to be an expert. I am leaving the University of Illinois as a *philosophiae doctor* in economics, holding the highest academic degree that may be awarded in this field. To the eyes of many people, a Ph.D. qualifies me as an expert in economics. However, I think there is a big misunderstanding around the notion of being an expert. Too many people think that means having a strong opinion about any matter and an answer to any question. Perhaps, the single most important lesson that I have learned doing a Ph.D. is that, in the moment you think you are an expert as intended by the general public, you have just failed. As compared to the day in which I first entered in this university, I am impressed by the amount of knowledge I am leaving with. The merit of this goes to so many people that I can hardly mention in full, but I will nevertheless name the most significant ones in the next paragraphs. However, the knowledge I have acquired is incomparably smaller than the amount of still unanswered questions that I will bring with me. These questions are as valuable as the skills and tools to try to answer them. Both of them - the questions and the tools - are the most precious endowment that the Ph.D. has gifted me.

In the six years that I spent in Illinois, I have met and interacted with an im-

pressively talented faculty. With some of them the relation has been cordial and pleasant, with some others it has been more challenging and defying. In both cases, the intellectual benefits I have received are stunning. Among them, I want to thank in particular Stephen Parente, In-Koo Cho and Rui Zhao. I want to thank also David Albouy and Ben Marx. My advisor Minchul Shin deserves a special mention. The awareness that I could always count on him has been so important and encouraging in the most difficult times. His affability can rival only with his smartness. Finally, I want to thank also Gregory Howard and Anne Villamil, who join Minchul Shin and Stephen Parente in my doctoral committee.

Sharing my graduate student life with other extremely endowed students has been as well rewarding and instructive. Moving across the challenges and difficulties of pursuing a Ph.D. is less burdensome if you know you are not alone in that. I want to thank my study buddies Camila Henao Arbelaez, Carlos Hurtado and Julia Gonzalez, and my colleague, teacher, and co-author Joao B. Duarte. Drinking coffee, discussing the most varied topics, and working together with Luis Felipe Sáenz has been one of the most enjoyable and enriching experiences during my graduate studies. His friendship is one of the best things I am bringing with me from Illinois.

Finally, my largest and deepest gratitude goes to my family. I could have never accomplished the doctorate without the love, the friendship, and the support of my wife Andrea. She has always trusted in me and believed in my abilities, and this has been the most important encouragement I received in these years. Her advise has been fundamental when I had to take big decisions. Moreover, my time in Illinois has represented also the begin of our life together, the time in which we became adults together, and, lastly, parents. Having all of that in my life has been essential to keep me motivated and confident in myself. Also, that enabled me to look at the events in the right perspective, helping me to understand what really matters. Obtaining a Ph.D. has not been my biggest accomplishment in my life so far. That is my daughter Emma, to which this dissertation is dedicated.

TABLE OF CONTENTS

CHAPTER 1 SOVEREIGN DEBT CRISES WITH INCOMPLETE INFORMATION 1
1.1 Introduction 1
1.2 Theoretical setup 5
1.3 Empirical analysis 21
1.4 Quantitative analysis 28
1.5 Conclusion 50

CHAPTER 2 WHY IS EUROPE FALLING BEHIND? STRUCTURAL TRANSFORMATION AND SERVICES' PRODUCTIVITY DIFFERENCES BETWEEN EUROPE AND THE U.S. 52
2.1 Introduction 52
2.2 Facts on a Disaggregated Service Sector 58
2.3 Model 63
2.4 Calibration 69
2.5 Quantitative Analysis 74
2.6 Counterfactual Experiments 86
2.7 Empirical Analysis of Labor Productivity Differences in Services . . . 95
2.8 Conclusions and Discussion 110

CHAPTER 3 THE RICARDO EFFECT: EVIDENCE FROM THE MANUFACTURING LABOR INCOME SHARE IN COLOMBIA 113
3.1 Introduction 113
3.2 The Ricardo Effect: A review of the literature 115
3.3 Data 119
3.4 Setup and Empirical Strategy 126
3.5 Results 128
3.6 Conclusion 135

REFERENCES 137

APPENDIX A ADDITIONAL MATERIAL FOR CHAPTER 2 142
A.1 Data sources 142
A.2 Measurement of Sectoral Labor Productivity in Each European Country 142
A.3 Findings of the Counterfactual Experiments for Each European Country 146

CHAPTER 1

SOVEREIGN DEBT CRISES WITH INCOMPLETE INFORMATION

1.1 Introduction

In the study of economic dynamics under uncertainty, an aspect that has been often discussed and taken into account in various different contexts is what Angeletos and Lian (2016) define as incomplete information. Let us consider an economy with a stochastic state component, in which economic agents need to make optimal decisions conditional on their expectations about the uncertain event. In this case, information is incomplete if the information sets they use in forming expectations are heterogeneous. In other words, incompleteness of information arise when at least a subset of agents has some sort of private information, unknown by other agents. With incomplete information, agents differ in their subjective probability distributions over the uncertain states of the economy. As a consequence, they may have different expectations and make different choices. However, this is not a departure from the rational expectations hypothesis: Each agent's expectations are rational, in the sense that they are optimally based upon all the information the agent possesses, but this information is not the same for all the agents.

What are the implications of incomplete information for the dynamics of sovereign debt yields and the occurrence of sovereign debt crises is, to the best of my knowledge, a still unanswered question in the existing literature of sovereign debt and default. This chapter intends to fill this gap by investigating, in particular, the effects of incomplete information on the pricing of government debt. The main quantitative finding is that, for low initial values of information heterogeneity, incomplete infor-

mation has an amplification effect on sovereign spreads' level and volatility as well as on the long run probability of sovereign default. The result is obtained from a counterfactual experiment based on a model of sovereign debt with incomplete information, calibrated to the euro area economy. If we identify the degree of information incompleteness as the level of noise in private information, which I structurally estimate from forecast data of the euro area, the average level of sovereign spreads would have been 27 per cent lower and the volatility 8 per cent lower, had the noise remained at the relative low average level prevailing before the Great Recession. The probability of default would have been 25 per cent lower. However, when the extent of information incompleteness becomes very large, its effect on spread level and volatility is reverted. Intuitively, too much private information noise lead investors to pay less attention to it.

The model of sovereign debt and default that I develop inherits much of the properties of the original model proposed by Eaton and Gersovitz (1981). However, the assumption about government's credible commitment follows Cole and Kehoe (2000). In a nutshell, this assumption implies that the government must issue new debt before deciding whether to serve or not outstanding debt falling due. This specific timing of events is well known to generate the possibility of rollover crises, that is sovereign defaults triggered if the government fails to optimally issue new debt, but avoided otherwise. In the standard Cole and Kehoe (2000) model with complete information, the possibility of rollover crises is the direct consequence of the existence of two equilibria in certain states of the economy: One in which the rollover crisis takes place, the other in which it does not occur. In the existing literature, the selection between these equilibria is driven by the realization of a sunspot variable, which constitutes a reduced form representation of investors' belief coordination. In this chapter, the standard model *à la* Cole and Kehoe (2000) is modified by introducing information incompleteness, in the form of private noisy signals about the uncertain state of the economy received by prospective investors. A theoretical result deriving from the novel informational structure is that uniqueness of equilibrium is restored, under certain conditions. The mechanism leading to this finding is in a very similar spirit to the one discussed by Morris and Shin (1998). More recently, Szkup (2018)

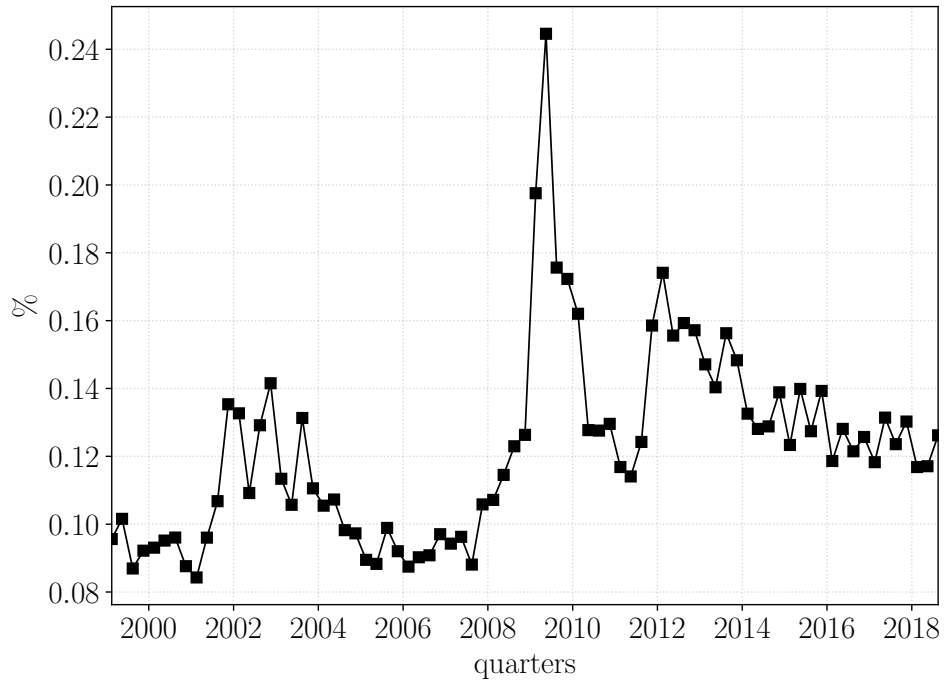
proved how private information leads to uniqueness of equilibrium in a model very similar to mine, but consisting of only two periods.

I calibrate the model to match moments of a set of financial and economic variables in the euro area. The test of the theory is whether the calibrated model can account for key statistics of sovereign spreads in the euro area, a variable not targeted in the calibration. The model is successful in accounting for the average level and volatility of spreads, and the negative correlation between spreads and real GDP growth. It is quantitatively less successful in accounting for the positive correlation between spreads and maturing debt to GDP ratio, although the model correlation is qualitatively correct. Having passed the test of the theory, the model is used as workhorse for the above mentioned counterfactual experiment that uncovers the relation between incomplete information and sovereign spreads. A central step in the present calibration exercise is measurement of private information noise – the degree of information incompleteness – in the euro area. The methodology proposed and the resulting measures represent the main empirical contribution of this chapter. One should notice that direct observation of private information is impossible by definition. On the other hand, there exists a variety of sources reporting forecast data. In this chapter, I look at forecasts about annual real GDP growth in the euro area, available through the Survey of Professional Forecasters (SPF) run by the European Central Bank (ECB). My identification assumption is that observable forecasts reflect unobservable private information in a well defined manner, based, in essence, on the informational structure of my model. This structure delivers two moment conditions binding private information noise to the dispersion and average uncertainty of forecasts, two statistics that we easily compute in the data. I measure private information noise by means of a GMM estimation strategy that combines these two moment conditions.

Figure 1.1 displays the time series of private information noise resulting from my measurement strategy, with a quarterly frequency. We notice that private information noise increases quite sharply at the inception of recessions or periods of turbulence[1], in particular during the Great Recession (2008 - 2009) and the European Debt

[1]In 2001, the U.S. entered in recession, while this did not occur in the euro area.

Figure 1.1: Private information noise



Crisis (2011 - 2012). Also, after these two events, private information noise has remained persistently higher than before. It appears that the Great Recession created a structural break in the time series, an hypothesis supported by standard statistical tests discussed in the empirical section. Exploiting the difference of pre-2008 noise mean is the main idea at the base of the counterfactual experiment.

The research presented in this chapter talks to a vast existing literature. Since the seminal theoretical research on sovereign debt of Eaton and Gersovitz (1981), Cole and Kehoe (2000), Calvo (1988), and Bulow and Rogoff (1989), many scholars have developed quantitative applications of this class of models, such as Arellano (2008), Aguiar and Gopinath (2006), Mendoza and Yue (2012), and Chatterjee and Eyigungor (2012). The methods for the quantitative analysis in this chapter follow closely the ones of these papers. Examples of quantitative applications of models *à la* Cole and Kehoe (2000), allowing for rollover crises, are given by Aguiar, Chatterjee, Cole,

and Stangebye (2017) and Bocola and Dovis (2018). The latter studies the dynamics of Italian sovereign spreads during the European Debt Crisis, the same crisis event considered here. My research is particularly close to models of sovereign debt that include some form of information and learning. In particular, Durdu, Nunes, and Sapriza (2013) studies the effect of a public signal that shifts the expectations of all the agents in the economy. Gu and Stangebye (2018) introduce costly information that investors might optimally choose whether to use or not. Models of soverign debt have been extensively used to study crisis episodes in emerging economies, while their calibration to mature economies, like in this chapter, is more recent. Bocola, Bornstein, and Dovis (2018) address what the latter implicates. This chapter is also related to empirical research on information and forecasts, such as Giordani and Söderlind (2003) and Patton and Timmermann (2010).

The rest of the chapter is organized as follows. Section 1.2 describes the model and the theoretical results, with emphasis on the characterization of the equilibrium and the effects of incomplete information on pricing government debt. Section 1.3 enters into the details of the measurement of private information noise from the ECB SPF data. Section 1.4 is about the quantitative analysis performed, from the calibration of the model and the way the model is solved numerically to the findings of the counterfactual experiment. Section 1.5 contains the conclusive remarks.

1.2 Theoretical setup

In this section, I describe a model of sovereign debt and default with incomplete information. I characterize the equilibrium of a simplified version of the model. The simplifications are necessary to obtain an analytic representation of the equilibrium, and I remove them later when the model is solved numerically for the quantitative analysis. I discuss under which conditions the equilibrium is unique. Finally, by means of a comparative static exercise, I uncover the channels through which incomplete information influences the equilibrium price of government debt.

1.2.1 Description of the model

In a small open endowment economy, there is a benevolent infinitely-lived government with full control over the endowment in each period (time is discrete). The government spends the endowment in transfers to the identical households populating the economy. The households enjoy felicity from consuming. They are risk-averse, *i.e.* they have a preference for a smooth intertemporal consumption. Since the government is benevolent, it shares the same preferences of the households: Its objective is an optimally smoothed plan of intertemporal spending. To achieve this objective, the government can borrow or save in a financial asset maturing in one period. When an outstanding debt is due to reimbursement, the government may choose to default on its obligations. In case of default, the government is temporarily excluded from future financial transactions. Moreover, it suffers an endowment loss in each period until readmission to the markets.

The counterpart of the government in the financial transactions is given by a continuum of investors. Investors are active for two periods: In the first period, they lend to or borrow from the government, and in the second period they settle their positions. Hence, in any period two cohorts of investors coexist: "Young" investors taking a position and "old" investors getting repaid or repaying. In contrast to most applications of this class of models, I do not restrict investors to be *foreigners*. This assumption is common in existing research because the vast majority of it has studied sovereign debt crisis in emerging economies, where most of government debt is indeed held by international investors. However, this is not the case for a developed economy like the euro area, the object of the present study. Therefore, I do not specify the origin of investors, and this does not change the main mechanisms of the model, as discussed by Bocola et al. (2018). They prove that, by a theoretical point of view, the relevant state variable in an advanced economy is the total amount of government debt, and distinguishing between domestically or internationally held debt does not really matter. In terms of investors' origin, the only assumption that I impose is that the government does not consider profits made by domestic investors when it chooses the optimal spending plan.

ENDOWMENTS. In each period t, the government has access to a stochastic endowment Y_t, which is log-normally distributed according to

$$y_t \equiv \log Y_t \sim \mathcal{N}(\hat{y}_t, \sigma_y^2) \qquad \forall t. \tag{1.1}$$

Let Y^t represent the history of endowment realizations up to Y_t included. I assume that $\hat{y}_t = \hat{y}\left(Y^{t-1}\right)$. Hence, the stochastic process of endowment is, in general terms, history-dependent.

Each cohort of investors has the same size W. In the second period of activity, investors receive an endowment large enough to meet any liabilities they may owe. This means that they are perceived as safe borrowers, and they are charged the international risk-free rate r on any amount they borrow.

PREFERENCES AND BUDGET CONSTRAINTS. The only action of households is consuming all the income they have. I do not explicitly specify households' utility and constraints. However, their risk aversion is reflected in the concavity of the government's utility function.

The government has preferences over the intertemporal spending plan $\{G_t\}$, with $G_t \geq 0$ in any period, represented by the additive utility function

$$U\left(\{G_t\}\right) = \sum_{t=0}^{\infty} \beta^t u\left(G_t\right),$$

with discount factor $\beta \in (0,1)$. The intratemporal utility function u is strictly increasing, concave, and satisfies the Inada conditions. Also, $u(G_t)$ tends to $-\infty$ as G_t tends to 0 from above. In any period of time in which default does not occur, government spending G_t satisfies the constraint

$$G_t \leq Y_t - B_t + q_t B_{t+1},$$

where B_t is the amount of government debt issued in period $t-1$ and due to reimbursement in period t, and $q_t B_{t+1}$ are the proceedings received by the government

from issuing new debt B_{t+1} at price q_t. A negative value of B_t or B_{t+1} is interpreted as a saving balance. On the other hand, if the government defaults in period t, the constraint on spending is given by

$$G_t \leq h\left(Y_t\right),$$

where the function h introduces an endowment loss associated with default, if $h\left(Y_t\right) < Y_t$. Notice that the latter budget constraint is in place also in any period following a default, until the government regains access to the financial markets. In any of these periods, readmission may occur with a probability λ.

Investors are risk-neutral and aim at maximizing their profit. Given that they are endowed in the second period of activity, an investor that wants to purchase units of government debt in the first period has to borrow the funds needed, and pay them back with accrued interest in the second period. In present value, the profit an investor makes by purchasing one unit of government debt is thus

$$\begin{cases} \frac{1}{1+r} - q_t & if \quad \text{repayment} \\ -q_t & if \quad \text{default} \end{cases}.$$

Informational structure. At the beginning of period t, the endowment Y_t is realized, but it is not immediately revealed to the economic agents. Both the government and the investors become aware of Y_t only at the end of the period. Initially, the agents know the history Y^{t-1} and the stochastic process of endowment defined in equation (1.1). Let $f_{t-1}(Y_t)$ denote the actual probability distribution function of Y_t, conditional on the past endowment realizations.

Before becoming aware of Y_t, each investor receives a private noisy signal about the current endowment realization. Let the signal received by investor i be

$$x_t^i = y_t + \varepsilon_t^i \qquad \varepsilon_t^i \sim \mathcal{N}(0, \sigma_x^2) \qquad \forall t.$$

In this setup, information is incomplete exactly because of the existence of these signals. Notice that, having observed their own signals, the investors "know more"

than the government, who does not receive any signal. One may justify this by saying that the government is transparent and releases all information in its hands to the public. Moreover, from a technical point of view, I do not need a government's signal to create information incompleteness: Every information set is different from each other even if one of the agents does not receive a signal.

In light of the signal received, investor i updates her beliefs about the probability distribution over the possible realizations of endowment. Applying the Bayes' theorem, investor i perceives that y_t is distributed according to

$$y_t|x_t^i \sim \mathcal{N}(\xi_t^i, \omega^2) \qquad \forall t,$$

where

$$\xi_t^i = \frac{\sigma_y^2}{\sigma_y^2 + \sigma_x^2} x_t^i + \frac{\sigma_x^2}{\sigma_y^2 + \sigma_x^2} \hat{y}_t,$$

$$\omega^2 = \left(\frac{1}{\sigma_y^2} + \frac{1}{\sigma_x^2} \right)^{-1}.$$

Let $p_{t-1}^i(y_t)$ denote the perceived probability distribution function of y_t, conditional on the endowment history and on receiving signal x_t^i. Notice that the perceived distributions differ across investors in terms of the perceived mean ξ_t^i but they have the same variance. Given that ξ_t^i is increasing in the signal received, we have that a perceived distribution induced by a signal x_t^i has first-order dominance over all the perceived distributions induced by signals smaller than x_t^i.

An important remark about the role of σ_x in the model is necessary. This parameter quantifies the noise in the informational content of every signal. Intuitively, it tells how *precise* a private signal is. For this reason, σ_x determines the weight attached to the signal in the expression of ξ_t^i, and influences the perceived uncertainty given by ω^2. At the same time, signals are drawn for infinitely many investors. Because of this, σ_x measures also the *dispersion* of signals across investors. We will often return to the duality "precision-dispersion" later in the chapter.

Debt Auction. Government debt is issued through a marginal price auction. For the sake of tractability, I impose the following institutional rule: Every investor presents a bid for one and only one unit of government debt. The auction works as it follows. The government announces the face value of the new debt B_{t+1} to be issued. Any investor i bids the price q_t^i for one unit of B_{t+1}. The units of B_{t+1} are allotted to the investors sequentially, following the descending rank of the bids $\{q_t^i\}$, until the supply of debt is exhausted. The market clearing price q_t is the price bidden by the last investor to receive one unit of B_{t+1}, the *marginal* investor. Every investor who has been allotted a unit of debt is charged q_t.

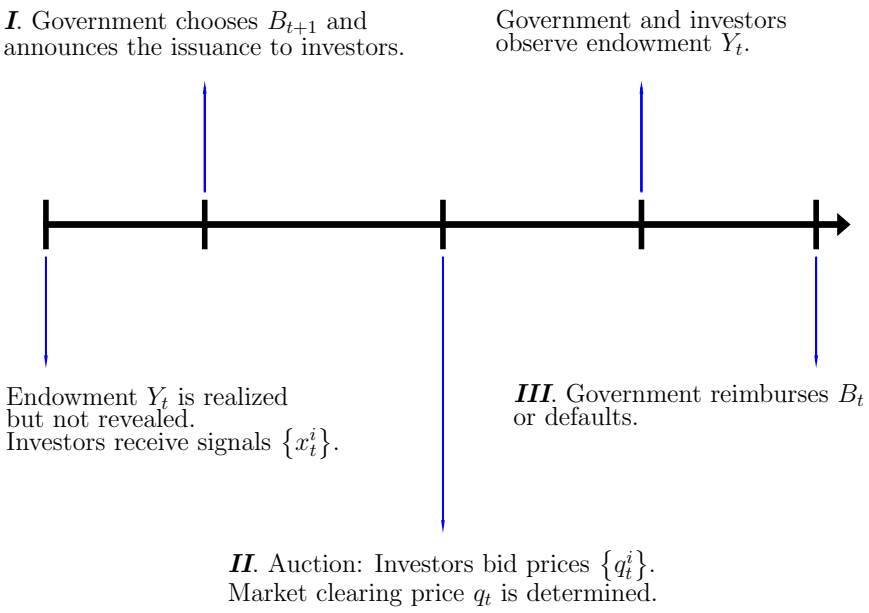
Because of the institutional rule, in each period the total demand for government debt is by construction equal to W, the measure of investors. Hence, letting χ_t represent the share of investors who bid a price at least as high as q_t, the market clearing condition for the debt auction is

$$B_{t+1} = W\chi_t \qquad \forall\, t.$$

Clearly, a technical constraint is that $B_{t+1} \leq W$.

Timing Protocol. As in Cole and Kehoe (2000), I assume that the government faces a time inconsistency. The optimality of reimbursing outstanding debt B_t depends on the outcome of the issuance of B_{t+1}, *i.e.* $q_t B_{t+1}$. Given a lack of credible commitment, the government must follow a timing of events in which the issuance of new debt occurs before the reimbursement of outstanding debt. Following the literature, I adopt the typical assumption that, if default occurs in period t, both B_t and B_{t+1} are repudiated, and the proceedings $q_t B_{t+1}$ forfeited. This implies that a default hurts not only "old" investors that hold B_t, but also "young" investors that have just purchased B_{t+1}. The flow of information intersects the timing of events. When the government issues B_{t+1}, it is unaware of Y_t, but it observes the latter before deciding whether to default or not. Investors have already received their signals when they present their bids, but they are unaware of the actual endowment level. Figure 1.2 is a graphical representation of the timing protocol. In terms of decision

Figure 1.2: TIMING PROTOCOL



making, let us define three distinct interim stages within any period t: In stage I, the government issues new debt; in stage II, investors present their bids and the market clearing price is set in the auction; in stage III, the government decides whether it reimburses the outstanding debt or not.

1.2.2 Recursive equilibrium

DEFINITION. In this model, a recursive equilibrium consists of a pricing function $q\left(Y^{t-1}, Y_t, B_t, B_{t+1}\right)$, value functions $V^D\left(Y_t\right)$ and $V^R\left(Y^{t-1}, Y_t, B_t\right)$, and policy functions $B'\left(Y^{t-1}, B_t\right)$ and $b\left(Y^{t-1}, x_t^i, B_t, B_{t+1}\right)$, such that

1. The government optimally issues $B_{t+1} = B'\left(Y^{t-1}, B_t\right)$, given outstanding debt B_t and the endowment history.

2. Any investor i optimally bids the price $q_t^i = b\left(Y^{t-1}, x_t^i, B_t, B_{t+1}\right)$ for one unit

of B_{t+1}, given outstanding debt, endowment history, and the private signal received.

3. The market clearing price $q_t = q\left(Y^{t-1}, Y_t, B_t, B_{t+1}\right)$ is the price bidden by the marginal investor in B_{t+1}, given outstanding debt, endowment history, and the current endowment Y_t.

4. By defaulting on B_t, the government receives the value $V^D(Y_t)$, which depends only on current endowment. By reimbursing B_t, the government receives the value $V^R\left(Y^{t-1}, Y_t, B_t\right)$, which depends on the current endowment, the outstanding debt, the issued debt, and the auction's proceedings.

5. The government defaults if and only if $V^D\left(Y_t\right) > V^R\left(Y^{t-1}, Y_t, B_t\right)$.

Characterization. In order to characterize the equilibrium analytically, I simplify the model in two dimensions. First, I set $h(Y_t) = Y_t$ for any level of endowment, and $\lambda = 0$. This implies that there is no direct cost, in the form of endowment loss, associated with default and that the only "punishment" for a defaulting government is a perpetual exclusion from future borrowing or lending. Second, I let the endowment to be identically distributed in any period, by imposing $\hat{y}_t = \hat{y}\ \forall\, t$. Under this simplification, endowment history is no longer a relevant state variable and this argument drops out from all equilibrium functions. These simplifications are relaxed in the quantitative analysis of Section 1.4.

I characterize the equilibrium by backward induction, starting from the interim stage III of period t. At this stage, endowment Y_t has been revealed and the issuance of new debt has been realized, *i.e.* B_{t+1} and q_t are given. The value that the government receives in case of default is given by the function V^D that solves

$$V^D(Y_t) = u\left(Y_t\right) + \beta \int V^D\left(\upsilon\right) f(\upsilon) d\upsilon,$$

and it depends on Y_t only. Given the set of predetermined variables $S_{t.III} = (B_t, B_{t+1}, q_t)$

and the just revealed Y_t, the government obtains the value

$$\tilde{V}^R(Y_t, S_{t.III}) = u\left(Y_t - B_t + q_t B_{t+1}\right) + \beta \int \max\left\{V^R\left(\upsilon, B_{t+1}\right), V^D\left(\upsilon\right)\right\} f(\upsilon) d\upsilon,$$

if it decides to reimburse B_t. Notice that $\tilde{V}^R(Y_t, S_{t.III}) = V^R(Y_t, B_{t+1})$ if and only if $S_{t.III} = \left(B_t, B'\left(B_t\right), q\left(Y_t, B_t, B'\left(B_t\right)\right)\right)$. At stage III of period t, the government defaults if and only if $V^D(Y_t) > \tilde{V}^R(Y_t, S_{t.III})$.

At this point, it is useful to introduce the following lemma.

Lemma 1. *There exists a unique threshold log endowment $\bar{y}_t \in (-\infty, \infty)$ such that*

$$\textit{Default occurs in } t \Leftrightarrow y_t < \bar{y}_t,$$

where $y_t = \log Y_t$, if and only if the variables in $S_{t.III}$ are such that $B_t > 0$ and $q_t B_{t+1} - B_t < 0$.

Proof. The partial sufficiency and necessity of $B_t > 0$ is trivial: If the government has no outstanding debt, default is impossible and no threshold exists.

When $q_t B_{t+1} - B_t < 0$, the value function $\tilde{V}^R(Y_t, S_{t.III})$ tends to $-\infty$ as Y_t tends to $-\left(q_t B_{t+1} - B_t\right)$. This derives from the properties of u, and the fact that the continuation value in $\tilde{V}^R$ does not depend on Y_t thanks to the second simplification. At the same time, $V^D(Y_t)$ is finite for $Y_t = -\left(q_t B_{t+1} - B_t\right)$. Hence, there exists a non empty set of low enough Y_t realizations for which the government chooses to default. On the other hand, it is the case that $\frac{\partial \tilde{V}^R(Y_t, S_{t.III})}{\partial Y_t} = u'\left(Y_t - B_t + q_t B_{t+1}\right)$ and $\frac{\partial V^D(Y_t)}{\partial Y_t} = u'\left(Y_t\right)$, given the simplifications. When $q_t B_{t+1} - B_t < 0$, $\tilde{V^R}$ increases faster than V^D as Y_t increases, by the concavity of u. Therefore, there exists a non empty set of high enough Y_t realizations for which the government chooses to reimburse. By the monotonicity of u, $\tilde{V}^R$ crosses V^D from below at one and only one endowment level $\bar{Y}_t$, establishing the uniqueness of the threshold. The properties of $\bar{Y}_t$ are invariant to the logarithmic transformation $\bar{y}_t = \log \bar{Y}_t$.

To prove the partial necessity of $q_t B_{t+1} - B_t < 0$, suppose *ab absurdo* that $q_t B_{t+1} - B_t > 0$. Then, $\tilde{V}^R(Y_t, S_{t.III}) > V^D(Y_t)$ for any level of Y_t, and a default will never occur. □

Lemma 1 establishes under which predetermined variables a default may occur, and for which realizations of y_t it occurs. The threshold log endowment is a function of the variables in $S_{t.III}$: $\bar{y}_t = \bar{y}(q_t, B_t, B_{t+1})$. In particular, for any B_t, B_{t+1}, the threshold is decreasing in q_t, because larger proceedings $q_t B_{t+1}$ reduce the incentives of defaulting on B_t.

At interim stage II, each investor i bids a price q_t^i for a unit of B_{t+1}. Given the bids, a market clearing price q_t for B_{t+1} is set. The government announces the amount B_{t+1} in advance, so that the set of predetermined variables is $S_{t.II} = (B_t, B_{t+1})$. Also, at this stage investors have already incorporated the signals in their beliefs. Each investor assesses the expected return from investing in one unit of B_{t+1}. An investor is reimbursed if default does not take place in the next period *and* in the current period. Conditional on no period t default, the probability of default in period $t+1$ corresponds to the probability that $V^R(Y_{t+1}, B_{t+1}) \geq V^D(Y_{t+1})$. According to the simplifications, the endowment is identically distributed in each period, implying that a signal received in period t is not informative about endowment realizations in period $t+1$. As a consequence, all agents assign the same probability to the event $V^R(Y_{t+1}, B_{t+1}) \geq V^D(Y_{t+1})$. Let

$$R_t = \frac{Prob\left(V^R(Y_{t+1}, B_{t+1}) \geq V^D(Y_{t+1})\right)}{1+r}$$

represent the expected present value return of investing in a unit of B_{t+1}, conditional on no period t default. Clearly, R_t is the same for any investor and it is a function of B_{t+1} only: $R_t = R(B_{t+1})$.

Investors are aware that at the next stage the government will opt to default if and only if $y_t < \bar{y}(q_t, S_{t.II})$. The probability assigned to this event by each agent depends on the signal received. In summary, the present value investment return expected by investor i is given by

$$v^i(q_t, S_{t.II}) = \left(\int_{\bar{y}(q_t, S_{t.II})} p^i(\upsilon) d\upsilon\right) R(B_{t+1}),$$

where p^i is investor i's perceived probability distribution over the realizations of y_t.

Notice that the expected return of any investor depends on the equilibrium price q_t set in the auction: There is an essential strategic complementarity in investors' bids, making this problem a global game. Higher q_t reduces the likelihood of default *ceteris paribus*, and this increases $v^i(q_t, S_{t.II})$. Moreover, by the first order stochastic dominance property of p^i, it is the case that

$$v^i(q_t, S_{t.II}) > v^j(q_t, S_{t.II}) \Leftrightarrow x_t^i > x_t^j,$$

for any $q_t, S_{t.II}$.

Any investor i decides the optimal bidden price q_t^i in order to maximize her expected profit, which is equal to

$$\pi_t^i(q_t, S_{t.II}) = v^i(q_t, S_{t.II}) - q_t.$$

The following propositions characterize the outcome of the auction in terms of investor's behaviors.

Proposition 1. *If investors bid according to a monotonic strategy profile, there exists a unique price q_t that clears the market.*

Proof. Without loss of generality, suppose investors' bids satisfy

$$q_t^i > q_t^j \Leftrightarrow x_t^i > x_t^j.$$

Then, the rank of the bids corresponds to the rank of the signals. By the market clearing condition, there exists a unique marginal investor who receives a signal x_t^m satisfying

$$B_{t+1} = W\Phi\left(\frac{y_t - x_t^m}{\sigma_x}\right),$$

where Φ represents the cumulative distribution function of a standard normal random variable. By the distribution of signals, $\Phi\left(\frac{y_t - x_t^m}{\sigma_x}\right)$ is the proportion of investors who bid higher prices than the marginal investor.

The price bidden by the marginal investor is the unique market clearing price q_t. □

Proposition 1 establishes a sufficient condition for the existence of a unique equilibrium, which is that investors' bid are monotonic in the signal received. The next proposition establishes a sufficient condition for the investors to follow a monotonic strategy profile.

Proposition 2. *If the standard deviation ω of the perceived probability distribution of y_t satisfies*

$$\omega > \phi\left(\frac{\xi_t^i - \bar{y}\left(q_t, S_{t.II}\right)}{\omega}\right)\left|\frac{\partial \bar{y}\left(q_t, S_{t.II}\right)}{\partial q_t}\right| R\left(B_{t+1}\right)$$

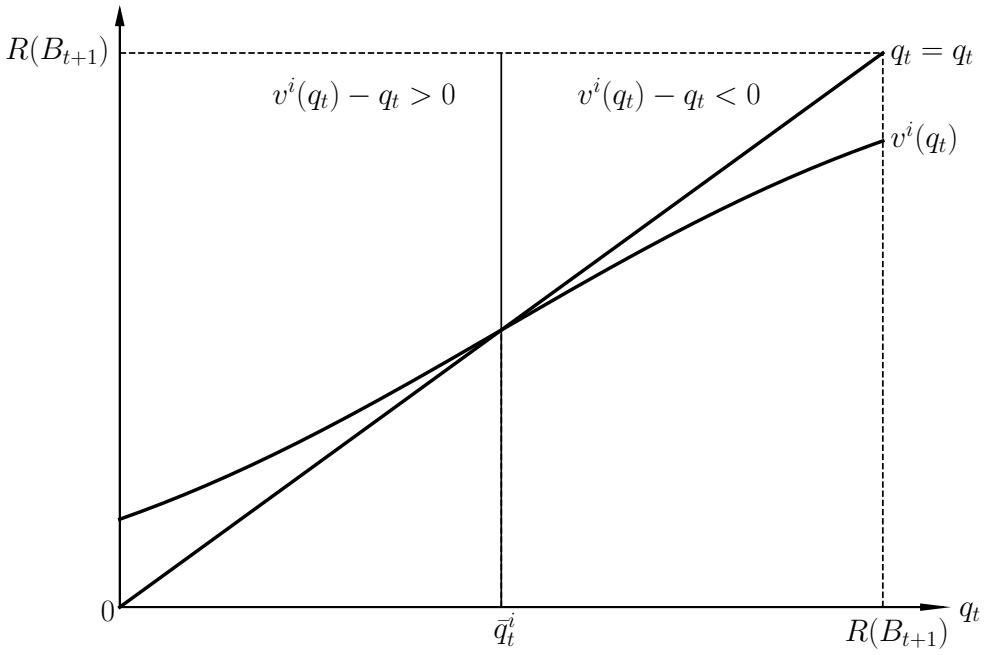
for any $q_t \in [0, R(B_{t+1})]$ given $S_{t.II}$, then investors' optimal bids are increasing in the signal received.

Proof. First, notice that the feasible values of q_t are constrained to $[0, R(B_{t+1})]$, since no investor would ever bid more than $R(B_{t+1})$ (she would expect a loss) and a negative price is meaningless. The condition of the proposition guarantees that $\frac{\partial v^i(q_t, S_{t.II})}{\partial q_t} < 1$ at any feasible value of q_t, and that the expected investment profit is decreasing in q_t. Moreover, it implies that, for any investor i, there exists a unique indifference price $\bar{q}_t^i$ at which the investor expects zero profits, *i.e.* $\pi_t^i(\bar{q}_t^i) = 0$. Figure 1.3 graphically represents a function v^i having these properties. Notice that investor i expects a loss (profit) for any q_t larger (lower) than $\bar{q}_t^i$.

It turns out that bidding $\bar{q}_t^i$ is the *unique* dominant strategy for investor i. I show this in two steps. First, notice that bidding $q_t^i = \bar{q}_t^i$ weakly dominates any other strategy. Indeed, if $q_t^i > q_t$, investor i is allotted a unit of B_{t+1}, and expects a profit since $q_t < \bar{q}_t^i$, while, if $q_t^i < q_t$, investor i does not receive any unit, but it is hedged from an expected loss, given that $q_t > \bar{q}_t^i$. Second, bidding $q_t^i = \bar{q}_t^i$ is the unique dominant strategy because any other strategy is dominated for some specific realization of q_t. For instance, consider the case in which $q_t^i > q_t > \bar{q}_t^i$: Investor i is allotted a unit of B_{t+1} on which she expects a loss. On the other hand, if $q_t^i < q_t < \bar{q}_t^i$, investor i foregoes an investment with an expected profit.

Therefore, under the condition of the proposition, any investor i bids the unique $\bar{q}_t^i$. It is straightforward to show that $\bar{q}_t^i > \bar{q}_t^j$ if and only if $x_t^i > x_t^j$. Indeed, if $x_t^i > x_t^j$, then $v^i(q_t) > v^j(q_t)$ for any q_t. Hence, $v^i(\bar{q}_t^j) - \bar{q}_t^j > v^j(\bar{q}_t^j) - \bar{q}_t^j = 0$. $v^i(\bar{q}_t^j) - \bar{q}_t^j > 0$

Figure 1.3: Expected return and profit



implies that $\bar{q}_t^i > \bar{q}_t^j$. Suppose $\bar{q}_t^i > \bar{q}_t^j$, and take any q_t' such that $\bar{q}_t^i > q_t' > \bar{q}_t^j$. Then $v^i(q_t') - q_t' > 0 > v^j(q_t') - q_t'$, and $v^i(q_t') > v^j(q_t')$. But this is the case only if $x_t^i > x_t^j$. This proves that the optimal bids are increasing in the signal received. □

The equilibrium pricing function $q\left(Y_t, B_t, B_{t+1}\right)$ delivers the market clearing price of B_{t+1}, given Y_t, B_t. A corollary from the proof of Proposition 2 is that every investor bids the unique price making her expect a zero investment profit. Therefore, under the condition of Proposition 2, the equilibrium price $q_t = q\left(Y_t, B_t, B_{t+1}\right)$ is implicitly defined by $\pi_t^m\left(q_t\right) = 0$, that is

$$q_t = \left(\int_{\bar{y}(q_t, S_{t.II})} p^m(\upsilon) d\upsilon \right) R\left(B_{t+1}\right),$$

where the superscript m indicates that the relevant probability distribution is the

one perceived by the marginal investor. This is the investor that receives the signal x_t^m solving

$$B_{t+1} = W\Phi\left(\frac{y_t - x_t^m}{\sigma_x}\right).$$

At the interim stage I, the government decides the amount of new debt B_{t+1} to issue. It is strategic, in the sense that it foresees the price at which the auction will clear for any realization of Y_t, given the outstanding debt B_t. In other words, the government knows the mapping induced by the equilibrium pricing function $q\left(Y_t, B_t, B_{t+1}\right)$. Then, the government chooses the optimal amount B_{t+1} by solving

$$\max_{B_{t+1}} \int_0^\infty \Big[u(\vartheta - B_t + q\left(\vartheta, B_t, B_{t+1}\right) B_{t+1}) + \\ +\beta \int_0^\infty \max\left\{V^R(\upsilon, B_{t+1}), V^D(\upsilon)\right\} f(\upsilon) d\upsilon \Big] f(\vartheta) d\vartheta$$

In summary, uniqueness of equilibrium is ensured by the fact that the auction of government debt delivers a unique market clearing price for any B_{t+1}, given any state Y_t, B_t. Proposition 1 presents a sufficient condition for that, and Proposition 2 states a sufficient condition to identify the unique equilibrium price as the indifference price of the marginal investor. In the rest of the discussion, I assume that the latter condition is satisfied.

1.2.3 Signal noise and the pricing of debt

At this point, I intend to explore the theoretical linkages between the level of signal noise and the pricing of government debt in the model. Based on the analytic characterization of the equilibrium, I do a comparative static exercise to study how a change in σ_x, the signal noise parameter, translates into a shift of the equilibrium price q_t. We see that the dual role of σ_x is important. Indeed, it is possible to identify two distinct channels through which σ_x affects q_t: A dispersion channel and a precision channel.

Using the functional forms of the perceived moments and probability distribution of y_t, the equilibrium price q_t of some B_{t+1}, given Y_t and B_t, is defined by the equations

$$B_{t+1} = W\Phi\left(\frac{y_t - x_t^m}{\sigma_x}\right),$$

$$q_t = \Phi\left(\frac{\xi_t^m - \bar{y}(q_t, B_t, B_{t+1})}{\omega}\right) R(B_{t+1}),$$

$$\xi_t^m = \frac{\sigma_y^2}{\sigma_y^2 + \sigma_x^2} x_t^m + \frac{\sigma_x^2}{\sigma_y^2 + \sigma_x^2}\hat{y},$$

$$\omega = \left(\frac{1}{\sigma_y^2} + \frac{1}{\sigma_x^2}\right)^{-\frac{1}{2}},$$

where Φ denotes the cumulative distribution function of a standard normal random variable. Notice that σ_x enters directly into the first equation, which is the market clearing condition, and into the third and fourth equations, defining respectively the mean and standard deviation of investor m's perceived probability distribution. Through these moments, σ_x affects indirectly the second equation, that establishes the zero expected profit condition for investor m.

The market clearing condition pins down x_t^m, the signal received by the marginal investor. To this end, given the issued amount B_{t+1} and total demand W, relevant statistics are where the distribution of signals is centered and how dispersed signals are around their mean. In essence, the role of σ_x in the first equation is quantifying the *dispersion* of signals across investors. On the other hand, in the last two equations, σ_x determines the optimal weight assigned to the signal in the perceived mean of y_t, relative to the actual mean $\hat{y}$, and how the perceived variance of y_t differs from the actual one, σ_y^2. In this regard, the role of σ_x is that of measuring the *precision* of the signal.

By means of implicit differentiation, the mathematical expression of how σ_x affects q_t is

$$\frac{\partial q_t}{\partial \sigma_x} = \Omega_t \left[-\frac{\sigma_y^2}{\sigma_y^2 + \sigma_x^2} \Phi^{-1}\left(\frac{B_{t+1}}{W}\right) +$$

$$+\frac{2\sigma_x\sigma_y^2}{\left(\sigma_y^2+\sigma_x^2\right)^2}\left(\hat{y}-x_t^m\right)-\frac{\sigma_y^2}{\sigma_x\left(\sigma_y^2+\sigma_x^2\right)}\left[\xi_t^m-\bar{y}(B_t,B_{t+1},q_t)\right]\Bigg],$$

where Ω_t is a positive term. The first term in the brackets quantifies the change in q_t caused by the fact that σ_x shifts the marginal investor's signal. Then, this term captures the impact of the signal noise on the price of debt by affecting signals' dispersion. The other two terms measure, respectively, how q_t is affected by the re-balancing of x_t^m and $\hat{y}$ in the perceived mean and by the change in the perceived variance, as a consequence of a shift in σ_x. Jointly, these two terms assess the effects of signal noise on the price of debt by changing signals' precision.

Can we reach some qualitative conclusions about the directions of these effects? With regard to the dispersion channel, the effect is negative for $B_{t+1}/W > 1/2$. In other words, for large enough debt issuance, the increase in private information dispersion associated with an increase in σ_x reduces the price at which the debt can be sold. Intuitively, this is related to the extent of "pessimism" characterizing the marginal investor, *i.e.* how low x_t^m is. The larger B_{t+1}, the more pessimist the marginal investor has to be for the market to clear. For large enough B_{t+1}, an increase in σ_x means that the marginal investor has to be even more pessimist, for the same Y_t, and this reduces the market clearing price. Notice that the empirical counterpart of B_{t+1}/W is the reciprocal of the bid-to-cover ratio. The analysis of Beetsma, Giuliodori, Hanson, and de Jong (2018) on debt auctions of main European countries suggest that the bid-to-cover ratio is less than 2 on average. Moreover, it tends to be lower during crisis, due to flight to safety phenomena. Heuristically, the condition for a negative dispersion effect seems to hold in the data.

Conclusions about the sign of the precision channel are more ambiguous, because they depend on the endogenous solutions for the marginal signal x_t^m and the threshold log endowment $\hat{y}_t$. Intuitively, this ambiguity rests on the fact that an increase in signal precision makes the marginal investor more confident about her private information. The latter acquires more influence on the investor's expectations, and the investor perceives less uncertainty around her expectations. The effect of all of this on the price of debt depends on how "positive" the private information is, and

how "optimistic" investor's expectations are. Re-balancing the perceived mean may have a positive effect on q_t if σ_x increases and x_t^m is lower than $\hat{y}$. Indeed, in this case the marginal investor attributes less importance to her relatively "bad" signal x_t^m and more importance to the relatively "good" prior $\hat{y}$. As a result, the perceived mean becomes more optimistic, and the market clearing price increases. At the same time, an increase in σ_x enlarges the perceived variance. Whether this has positive or negative effects on q_t depends if the marginal investor' posterior means is smaller or larger than the threshold $\hat{y}_t$. In the quantitative analysis, the ambiguity about the effect of signal noise on the price of debt will be clarified through the counterfactual experiment.

1.3 Empirical analysis

The main focus of this section is the measurement of private information noise in the euro area. This measure will serve as a base for the calibration of the signal noise parameter σ_x in the quantitative analysis. The measurement is based on forecasts data from the Survey of Professional Forecasters run by the European Central Bank. The strategy consists in reading these data through the lenses of the model. In essence, I maintain the informational structure of my model, and from the latter I derive the model implied dispersion and uncertainty of forecasts. Not surprisingly, these model statistics depend on σ_x. On the other hand, I can observe forecast dispersion and uncertainty in the data. By means of a generalized method of moments, I estimate the value of σ_x that minimizes the distance between the model and data statistics.

1.3.1 Data description

The ECB Survey of Professional Forecasters reports forecasts about real GDP growth, inflation, and unemployment in the euro area, on a quarterly basis. In general, survey participants are financial institutions, industrial organizations, and research institutes. For any variable of interest, various forecasting horizons are reported. The

survey has run since 1999-Q1 and it is still ongoing.

For the purpose of the present study, I consider the quarterly forecasts of real GDP growth over the four quarters since the last quarter for which an official GDP figure is available. Given that the survey is presented at the beginning of every quarter, the last GDP release refers to two quarters back. For instance, in the survey dating 2019-Q1, participants has reported their forecasts of GDP growth over 2018-Q4, 2019-Q1, 2019-Q2 and 2019-Q3. In this sense, to some extent these forecasts reflect the contemporaneous uncertainty of professionals about the current GDP growth. Each respondent discloses two pieces of information. First, she gives a point forecast of the growth rate. Second, she expresses her subjective probability distribution: Given a set of growth rate brackets coming with the survy, she assigns the probability that the growth rate will fall in any one of those brackets. Forecasts are about annual growth rates, while one period in my model corresponds to a quarter in the data. I transform the annual rates in quarterly rates by assuming homogeneous growth over the four quarters.

Other data that I use in this section are sovereign yields, sovereign debt levels and GDP of the euro area countries. The sources of these data are the Global Financial Database (for the yields) and the OECD (for all the rest).

1.3.2 Forecast dispersion and uncertainty in the data

Forecast dispersion quantifies the extent of disagreement in forecasters' opinions. In order to assess the dispersion of the observed forecasts, I compute the cross-section standard deviation of point forecasts across respondents, for each quarter from 1999-Q1 to 2018-Q3. Forecast uncertainty reflects the degree of confidence a forecaster has on her forecasts. The measurement of forecast uncertainty is based on the subjective probability distributions reported. Intuitively, an extremely confident forecaster should assign probability one to the bracket containing her point forecasts and zero probability to all the others. Hence, forecast uncertainty is related to how widespread across brackets the reported probability mass is.

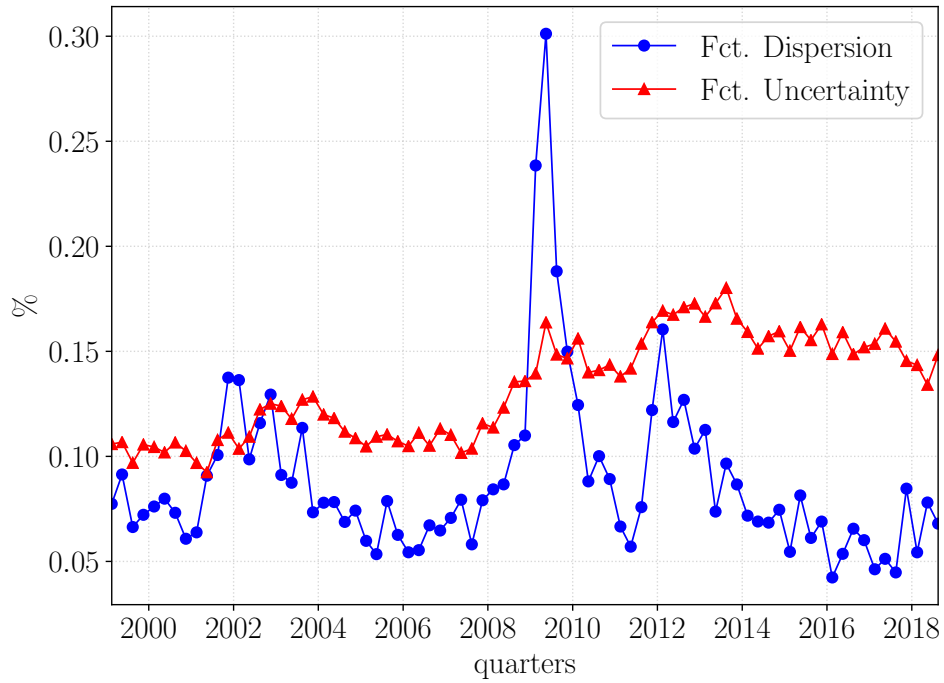
Patton and Timmermann (2010) discuss various methods to infer a measure of

dispersion from a discrete histogram-like probability distribution, as the one we have here. The strategy that I follow is slightly different, and it is based on non-parametric smoothing and Monte Carlo sampling. First, I smooth each reported probability distribution over the brackets by means of a kernel density estimator. The bandwidth of the estimator is selected in order that the probability mass assigned by the estimated density to each bracket matches the one indicated by the forecaster. For the smoothing, I have to deal with the fact that the two growth brackets at the extremes are open brackets. I assume that the mid growth rate in these bins is distant twice the length of each bracket (the same for all of them) from the defined closed bound. Second, I simulate multiple samples out of the estimated kernel density, and I compute the average standard deviation of the simulated samples. I take this moment as the measure of the uncertainty associated with any forecast reported in the survey. In conclusion, I measure quarterly forecast uncertainty by repeating this process in each quarter between 1999-Q1 and 2018-Q3, and averaging the uncertainty measures across the respondents of the quarter.

The quarterly measures of forecast dispersion and forecast uncertainty are plotted in Figure 1.4. The two series present different properties in terms of dynamics. Forecast dispersion is more volatile and it suddenly increases in periods of crises. This has occurred during the Great Recession (2008-2009), the European Debt Crisis (2011-2012), and, to a lesser extent, at the time of the 2001 U.S. recession. However, forecast dispersion reverts to lower levels in "normal" times. Dispersion deflation took place after all of those crisis episodes. On the other hand, forecast uncertainty is less volatile, but it had a behavior similar to forecast uncertainty during the European Debt Crisis and the 2001 recession: it increased and it slowly reverted to lower levels later on. However, the dynamics of the two measures around the Great Recession is very different. Forecast uncertainty went through a structural break as a consequence of the Great Recession. Indeed, the levels of uncertainty has fluctuated at levels persistently higher than before since 2008-2009. In particular, the "normal" levels of uncertainty post-European Debt Crises are significantly higher than the levels of uncertainty observed during the first years of the series.

At this point, it is interesting to explore the correlation between the dynamics

Figure 1.4: Forecast dispersion and uncertainty



of the two forecast statistics and the dynamics of sovereign spreads, which is what a model of sovereign debt and default is ultimately meant to account for. If any significant correlations are encountered, this exercise will also motivate why we should not overlook private information in studying sovereign debt. After all, forecasts are significantly influenced by private information. I explore these correlations by means of a panel regression in which the dependent variable is the sovereign spread of each country belonging to the euro area, except Germany. I define the sovereign spread as the difference between the yield paid by a country's government bond with five year maturity and the yield paid by the corresponding German government bond (this is why Germany is excluded from the sample). Among the independent variables, I include the forecast dispersion and uncertainty as well as real GDP growth and the government debt to GDP ratio, which are the typical explanatory factors of spreads in the existing empirical research on the topic. I control for country fixed effects and

for quarter effects, to control for any seasonality in the series.

The estimated correlations are reported in Table 1.1. My findings confirm that

Table 1.1: Panel estimates

Dependent: Spread 5 year maturity (basis points)				
Debt to GDP (%)	8.236***	8.375***	7.870***	8.632***
	(1.249)	(1.268)	(1.344)	(1.441)
GDP Growth (%)	−64.432***	−58.125***	−62.671***	−58.768***
	(22.844)	(21.867)	(23.485)	(22.614)
Fct. Dispersion (%)		194.179**		210.749**
		(96.052)		(87.082)
Fct. Uncertainty (%)			120.722	-80.932
			(206.725)	(188.231)
Fixed effects:				
Quarter × Country	Yes	Yes	Yes	Yes
Country	Yes	Yes	Yes	Yes
Obs.	753	753	753	753
Adj. R-sq.	0.28	0.28	0.28	0.28

Notes. $p-$value: $^{*} < 0.1$, $^{**} < 0.05$, $^{***} < 0.01$. Robust standard errors.

spreads tend to increase with the debt to GDP ratio, and they are on average lower when GDP growth is larger. These are well known stylized facts. The novelty of my analysis rests on the inclusion of forecast statistics. There is a significant positive correlation between sovereign spreads and forecast dispersion. Specifically, a one per cent increase in the dispersion of forecasts is associated with sovereign spreads that are approximately 200 basis points higher, on average. On the other hand, forecast uncertainty does not show any significant correlation with the spreads. The takeaway is that private information has the potential to help us understanding spread dynamics better, since the level of disagreement among forecasters - forecast dispersion - is likely explained by differences in the private information they have. Also, we may try to relate the empirical findings with the theoretical conclusions of Section 1.2. The model predicts that the price of debt decreases if the dispersion of

private information increases, for the levels of bid-to-cover ratio typically observed in Europe. This prediction is consistent with the observed positive correlation between forecast dispersion and spreads. With regard to the precision of private information, the model does not have clear-cut predictions of its effect on the price of debt. This is not at odds with the empirical result that spreads are uncorrelated with forecast uncertainty, a measure clearly influenced by the precision of private information.

1.3.3 Structural estimation of private information noise

The first step of the estimation strategy is taking the informational structure of the theoretical setup and deriving the expressions for forecast dispersion and uncertainty in the model. Let us consider all the investors as forecasters, and let us think at y_t as real GDP growth. Recall that any investor i believes that y_t is distributed according to

$$y_t \sim \mathcal{N}\left(\xi_t^i, \omega^2\right).$$

Hence, investor i's forecast of y_t is the perceived mean

$$\xi_t^i = \frac{\sigma_y^2}{\sigma_y^2 + \sigma_x^2} x_t^i + \frac{\sigma_x^2}{\sigma_y^2 + \sigma_x^2} \hat{y}_t.$$

The dispersion of forecasts across investors is fully accounted for by the dispersion of signals, σ_x. Hence, the cross-sectional standard deviation (sd) of forecasts is given by

$$sd\left(\xi_t^i\right) = \frac{\sigma_y^2}{\sigma_y^2 + \sigma_x^2} \sigma_x$$

in the model. On the other hand, the uncertainty associated with a forecast has been measured in the data as the average standard deviation of forecasters' subjective probability distribution over growth realizations. Clearly, its counterpart in the model is given by

$$\omega = \left(\frac{1}{\sigma_y^2} + \frac{1}{\sigma_x^2}\right)^{-\frac{1}{2}}.$$

The next step is to use a GMM strategy to estimate the value of σ_x that minimizes the distance between the model based and the observed statistics. I slightly modify the setup by allowing both σ_x and σ_y to be time-varying. Let

$$G_t = \begin{bmatrix} sd\left(\hat{\xi}_t^i\right) - \frac{\sigma_{y,t}^2}{\sigma_{y,t}^2+\sigma_{x,t}^2}\sigma_{x,t} \\ \hat{\omega}_t - \left(\frac{1}{\sigma_{y,t}^2} + \frac{1}{\sigma_{x,t}^2}\right)^{-\frac{1}{2}} \end{bmatrix},$$

where the hat superscript denotes data moments. Then, I estimate $\sigma_{x,t}$ by solving

$$\min_{\sigma_{x,t}} G_t' W G_t$$

quarter by quarter. I set W equal to a 2×2 identity matrix. Obviously, I need to assign a value to $\sigma_{y,t}$ in each quarter beforehand. I estimate an AR(1) specification on real GDP growth of the euro area since 1995-Q1, modeling the volatility as an ARCH(1) process. I set the series of $\sigma_{y,t}$ equal to the series of the conditional volatility predicted by the empirical model.

The series of σ_x measured through this strategy is plotted in Figure 1.1, in the introduction. Not surprisingly, one may recognize that the dynamics of σ_x shares the properties of both observed forecast dispersion and observed forecast uncertainty. In particular, σ_x peaked during the crises and moved down afterwards, but remained at levels persistently higher than before since 2008. In order to establish the existence of a structural break around the Great Recession in a more rigorous way, I perform statistical test to check significant differences in the mean and variance of σ_x pre- and post-2008. I rely on a Welch's t-test for equality of means and on a F-test for equality of variances. Table 1.2 summarizes the results of these tests. Both the null hypotheses that mean and variance are equal before and after 2008 are rejected at standard significance levels. This evidence supports the thesis that private information noise experienced a structural break at the time of the Great Recession.

Table 1.2: Structural break tests

	Equality of Means	Equality of Variances
Test Stat.	8.0295	2.9555
p-value	< 0.0001	0.0007

Notes. Testing equality of moments based on sub-samples pre- and post-2008. The null hypothesis of each test is that the moments are equal.

1.4 Quantitative analysis

In this section I bring the model to the data, I show that it is able to account for key statistics of the sovereign spreads in the euro area, and I study the impact of private information noise on spreads' dynamics. I start by parametrizing the functional forms of the model and by giving some details about the algorithm used for the numerical solution. Then, I explain the way parameter values are assigned. I test the theory using untargeted moments of the average sovereign spreads in the euro area between 1999-Q1 and 2018-Q3. Finally, I perform counterfactual exercises to assess the role of signal noise on sovereign spreads.

1.4.1 Parametrization and numerical algorithm

To solve the model numerically, first I need to define the parametric specification of the stochastic endowment process, the utility function, and the function h controlling the direct default costs. I assume that log endowment is serially correlated: It follows the autoregressive process of order one

$$y_t = \rho y_{t-1} + \nu_t \qquad \nu_t \sim \mathcal{N}\left(0, \sigma_y^2\right).$$

The intratemporal utility function belongs the the CRRA class:

$$u\left(G_t\right) = \frac{G_t^{1-\eta}}{1-\eta},$$

where the parameter η measures the constant relative risk-aversion of the households, which the government takes care of. For modeling the direct costs of default in terms of endowment loss, I adopt the typical assumption in this literature since Arellano (2008) that default costs are asymmetric. Specifically, let

$$h(Y_t) = \begin{cases} Y_t & if \quad Y_t \leq \bar{Y} \\ \bar{Y} & if \quad Y_t > \bar{Y} \end{cases},$$

for some threshold parameter $\bar{Y}$. The implication of this choice is that the cost of default increases with Y_t more than proportionally for endowment levels above the threshold. Hence, the government's incentive to default falls quickly when endowment realizations become larger. This helps the model replicating the stylized fact that defaults tend to occur when the state of the economy is bad.

For the numerical solution, the variables Y_t, B_t, and x_t^i are discretized: They take values from finite grids. The probability distribution induced by the AR(1) process of endowment is also discretized in a Markov chain. Because of the serial correlation in endowment, a state consists of the triplet (Y_{t-1}, Y_t, B_t). The previous period endowment is relevant to compute the transition probabilities of the Markov chain. Also, any state and level of B_{t+1} are associated with a specific marginal investor through the market clearing condition. The marginal investor's perceived probability distribution determines the equilibrium price. Hence, I build different transition probabilities for any signal that may be received by a marginal investor (recall that the signal only shifts the center of the distribution). The grid lengths for endowment, debt, and signals are respectively 25, 50, and 1000. The algorithm for the numerical solution consists of two nested iterations: The first on the pricing function and the second on the government' value functions. In practice, for any pricing function obtained in the last iteration, I iterate over the value functions until they

converge to a solution conditional on the pricing function. Then, I use the solution for the value functions to update the pricing function using the zero expected profit condition for the marginal investor, where the latter differs for any B_{t+1} and any state. I repeat this process until the pricing function reaches convergence. Notice that the numerical algorithm incorporates the equilibrium concept defined by Proposition 1 and 2: The marginal investor is pinned down by the market clearing condition given a bidding strategy increasing in signal, and the equilibrium price makes the marginal investor indifferent.

1.4.2 Calibration

As a first step, I need to define consistent empirical measures for the variables in my model. Following the existing literature, I choose real GDP as the empirical counterpart of Y_t. This choice poses an issue. In my model, the stochastic process of $y_t = \log Y_t$ is stationary. On the other hand, I run different statistical tests on log GDP of the euro area and the conclusion is that this series has a unit root, which rejects stationarity. In order to reconcile the model with the data, I normalize the variables of my model as relative to previous period endowment. Hence, the empirical counterpart of $y_t = \log Y_t - \log Y_{t-1}$ is $\log GDP_t - \log GDP_{t-1}$, the quarterly growth of real GDP. With the normalization, any debt level B_t is interpreted as an actual debt to GDP ratio. In the model, B_t is the amount of debt maturing in each quarter. Then, the correct measure is not the entire debt to GDP ratio of the euro area in a given quarter, but rather the maturing debt to GDP ratio. As reported by European Commission (2018), the average maturity of government debt in the euro area is 6.8 years, or 28 quarters. Hence, approximately 1/28 of the entire stock of debt is due to reimbursement in each quarter. Let $\hat{B}_t$ denote the quarterly value of the government debt to *annual* GDP ratio of the euro area, available from the OECD database. Assuming that annual GDP is produced homogeneously in the 4 quarters, the correct measure for B_t is $\hat{B}_t/7$. The average value of the latter is 10 per cent, indicating that in each quarter the euro area reimburses liabilities approximately amounting to 10 per cent of its GDP. Finally, given the normalization, I assume that

investors' measure W is constant in relative terms to last period endowment, rather than in absolute terms. This permits to keep interpreting B_{t+1}/W as the reciprocal of the bid-to-ask ratio. In the analysis, I compare model and data statistics of spreads. With regard to the model, I define the sovereign spread in period t as

$$spr_t = \left(\frac{1}{q_t} - 1\right) - r.$$

From the data, I build a measure of average sovereign spreads in the euro area by averaging the spread series of euro area members countries defined in Section 1.3, weighted by the country's share of euro area government debt.

In Table 1.3, I display the values assigned to the 9 parameters of my model. The calibration of 7 of them is disciplined by external statistics, and it does not involve model simulation. The parameters governing the endowment process are calibrated

Table 1.3: Calibration of parameters

Parameter	*Value*	*Target*
Endowment process		Real GDP in euro area
ρ	0.643	
σ_y	0.43%	
Government preferences		
η	2.000	Standard in literature
β	0.897	Debt service to GDP
Default costs		
λ	0.062	Length of bailout programs
$\bar{Y}/E(Y)$	0.990	Historical default probability
Risk-free interest rate		
r	0.007	Yield of German government bond
Measure of investors		
W	0.230	Bid-to-cover ratio
Signal noise		
σ_x	0.12%	ECB SPF statistics

based on the estimates of an AR(1) model fitted on quarterly real GDP growth of the euro area, between 1995-Q1 and 2018-Q3. The persistence of the process is controlled by the parameter ρ, set at 0.64, and the volatility σ_y is set equal to 0.43

per cent. The risk-free rate r is measured as the average quarterly yield of a German government bond with a five year residual maturity, between 1999-Q1 and 2014-Q4 (I cut the horizon shorter to avoid the effects of the ECB's asset purchasing program). This value is 0.7 per cent. The value 0.23 of W targets an average bid-to-cover ratio of 1.9, based on data reported by Beetsma et al. (2018), given an average maturing debt to GDP ratio of 10 per cent. The parameter λ represents the probability that the government is readmitted to the financial markets after a default. As a target, I take the average length of the financial assistance programs put in place for the bail-out of five euro area countries between 2010 and 2018. The average length of these programs has been 16 quarters. A value of λ equal to 6.2 percent insures that the average government's exclusion from markets lasts 16 periods. The risk-aversion parameter η is set equal to 2, as standard in this literature. The parameter constituting the novelty of my framework - the signal noise parameter σ_x - takes the average value of the private information noise measured in Section 1.3 and displayed in Figure 1.1, that is 0.12 percent. Finally, the calibration of the last two parameters - the discount factor β and the threshold $\bar{Y}$ - is based on model simulation, in order that model predictions match the average maturing debt to GDP ratio of 10 per cent and a default probability of 0.15 per cent. I compute the latter based on the historical data of Reinhart and Rogoff (2011), considering the default frequency of the euro area member countries since 1900 (war time defaults excluded). As a result, calibrated parameters are $\beta = 0.897$ and $\bar{Y} = 0.99E(Y)$. The performance of the model with respect to the targeted moments is reported in the upper panel of Table 1.4.

1.4.3 Test of the theory

Let us judge the model by verifying if it can account for some statistics in the data. I base the test on untargeted moments summarizing important properties of the sovereign spreads' dynamics. In particular, I consider spreads' mean and volatility, and the correlation between spreads, GDP growth, and maturing debt to GDP ratio. Table 1.4 presents a comparison between these moments as computed in the data and

as resulting from model simulation. Overall, the model is successful in accounting

Table 1.4: Moments - Simulation vs. Data

	Simulation	Data
Targeted moments		
$m(B_t)$ - Avg. Maturing Debt to GDP	9%	10%
Probability of Default	0.12%	0.15%
Other moments		
$m(spr_t)$ - Avg. Spread	11 *bps*	16 *bps*
$sd(spr_t)$ - Std. Dev. Spread	0.24%	0.18%
$\rho(spr_t, y_t)$ - Corr. Spread, GDP growth	-0.38	-0.28
$\rho(spr_t, B_t)$ - Corr. Spread, Maturing Debt to GDP	0.13	0.27

for the key statistics of sovereign spreads in the euro area. In particular, the model does a good job in terms of spreads' mean and standard deviation, by predicting an average level of spreads of 11 basis points (16 in the data) and spreads' volatility equal to 0.24 per cent (0.18 per cent in the data). With regard to the correlations, the model is slightly imprecise in quantitative terms, but it is able to capture the negative correlation between spreads and GDP growth, and the positive correlation between spreads and maturing debt to GDP ratio.

I want to highlight the volatility of spreads obtained from the simulations. A well known problem with the standard model of sovereign debt and default is that it hardly generates enough spread volatility, as compared to the data. In their chapter for the *Handbook of Macroeconomics*, Aguiar, Chatterjee, Cole, and Stangebye (2016) discuss this point with some examples, showing that the standard model underestimates the spread volatility experienced in Mexico even if it features asymmetric default costs. On the other hand, Arellano (2008) shows that this type of

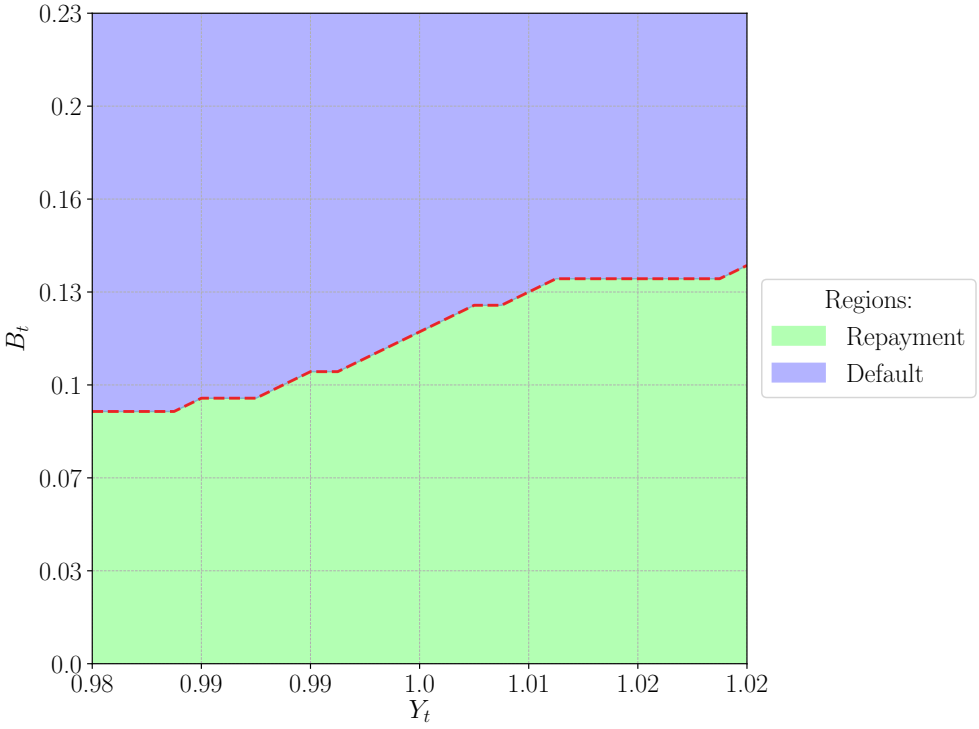
costs is sufficient to capture the spread volatility of Argentina. According to Aguiar et al. (2016), the cause of the failure has to be identified in the low volatility of endowment, reflecting the relative stability of the Mexican economy with respect to Argentina. The present volatility of endowment is also quite low, being calibrated to relatively stable advanced economies. Nevertheless, the model is able to generate a level of spread volatility in fact larger than the one observed in the data. The only component making the present model really different from the ones in the cited studies is the incomplete information. Hence, this is a preliminary indication that including noisy private information in the model amplifies the volatility of spreads. I will address this point more in depth when I discuss the counterfactual exercise.

1.4.4 Properties of the numerical solution

Before moving to the counterfactual experiment, I present and discuss some properties of the equilibrium solved numerically. Figure 1.5 visualizes a partition of the state space $\{(Y_t, B_t)\}$ depending on whether a default occurs in each state (Y_t, B_t). The third state variable, Y_{t-1}, has been averaged out using the invariant probability distribution associated with the actual Markov chain of Y_t. The properties of the default region are similar to the ones obtained in existing quantitative applications of this class of models. A default is more likely for lower levels of Y_t, *i.e.* governments tend to default during recessions, which is a fact supported by a vast empirical evidence. Also, there is a quite narrow interval of debt levels for which a default is neither excluded nor expected with certainty. These levels are between 9 per cent and 14 per cent of GDP approximately. However, this narrowness does not seem in contrast to the findings of similar research, for instance Aguiar and Gopinath (2006).

The solution of the policy function $B'(Y_{t-1}, B_t)$ is pictured in Figure 1.6, where the first state variable is again averaged out using the invariant probability. Notice that, for low levels of B_t, on average the government chooses to accumulate more debt. On the other hand, when outstanding debt becomes excessively large, the government has a strong incentive to deleverage its passive position. As a result, in the long run the maturing debt to GDP ratio is expected to fluctuate within a

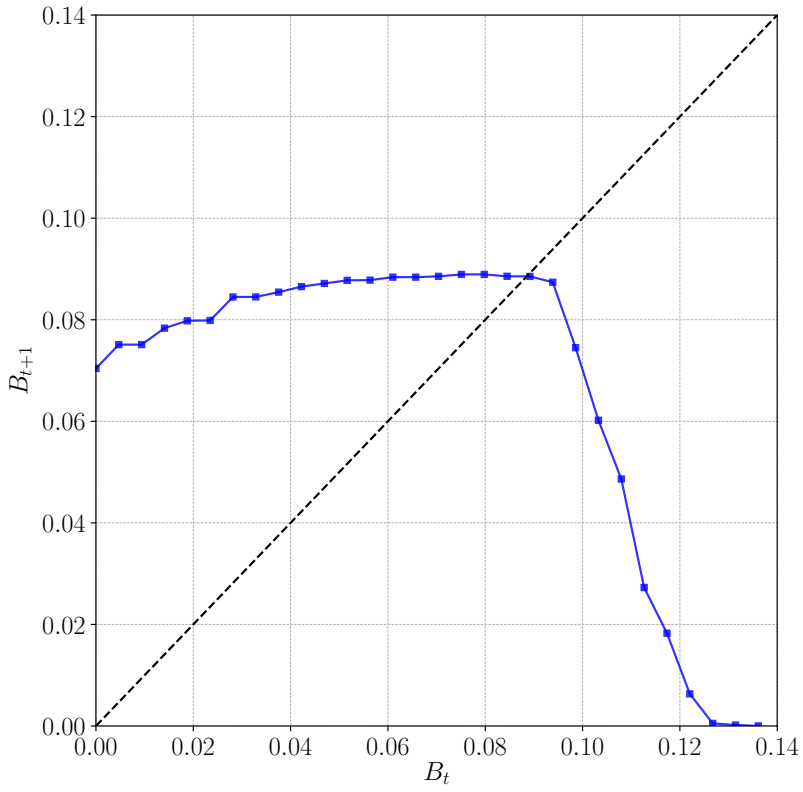
Figure 1.5: Default region



neighborhood of 9 per cent. Notice that the latter corresponds to the upper level of B_t for which there is a zero probability of default. Hence, these properties suggest that the government tends to dislike risky debt positions, probably because they are heavily discounted with lower prices by the investors. If the outstanding debt has an even small probability of default, the government scales it down toward debt levels perceived as safer. In light of this behavior of the government, it is not surprising that the model predicts low average spreads (11 bps) and low probability of default (0.12 per cent), in coherence with the data.

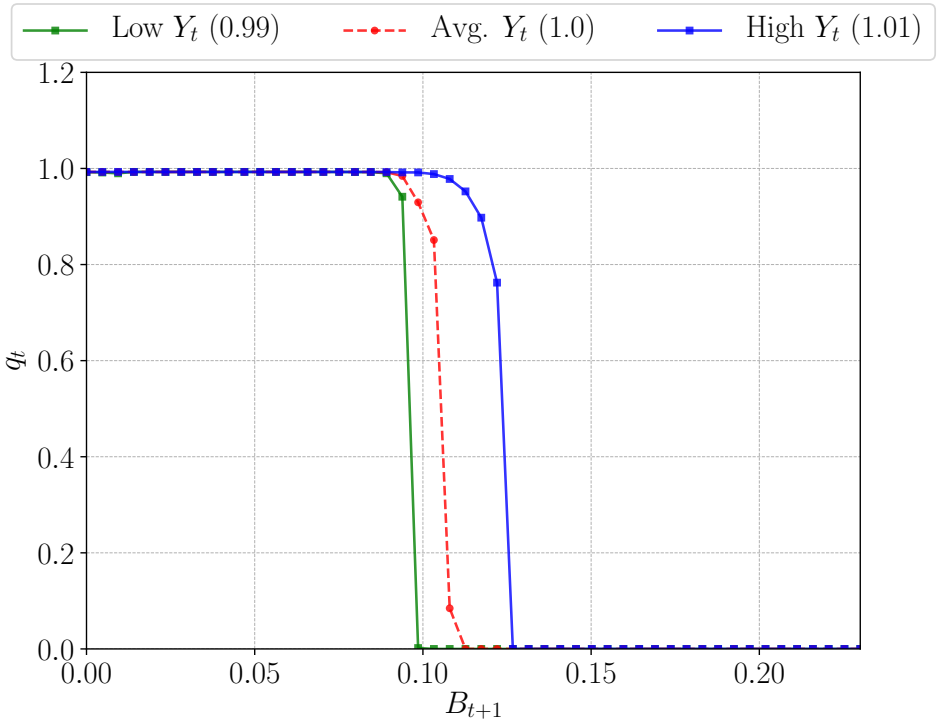
Let us consider the numerical solution of the pricing function. In the quantitative applications of the standard model, the equilibrium price of B_{t+1} is usually a function of B_{t+1} itself and Y_t, to the extent that the latter is informative about next period endowment because of serial correlation. In the present setup, the pricing function has four arguments: Y_{t-1}, Y_t, B_t and B_{t+1}. The first and last ones have the same

Figure 1.6: Policy function



role as Y_t and B_{t+1} in the traditional pricing function (Y_{t-1} matters only because of serial correlation). The novelty of this pricing function rests on the role of the arguments Y_t and B_t, absent in the traditional pricing function. Here, Y_t does not matter for the equilibrium price because of its informational content. Indeed, neither the government nor the investors are aware of Y_t when the price is set. However, Y_t determines the center of the signals' distribution, which pins down the marginal investor through the market clearing condition. Hence, Y_t has an impact on the price of debt by affecting the selection of the marginal investor's signal. The effects of Y_t on the price of debt are showed in Figure 1.7. In this figure, the pricing function is plotted against values of B_{t+1}. Again, the argument Y_{t-1} is averaged out using the invariant probability distribution. Maturing debt to GDP ratio B_t is set at

Figure 1.7: Pricing function - Changing current endowment

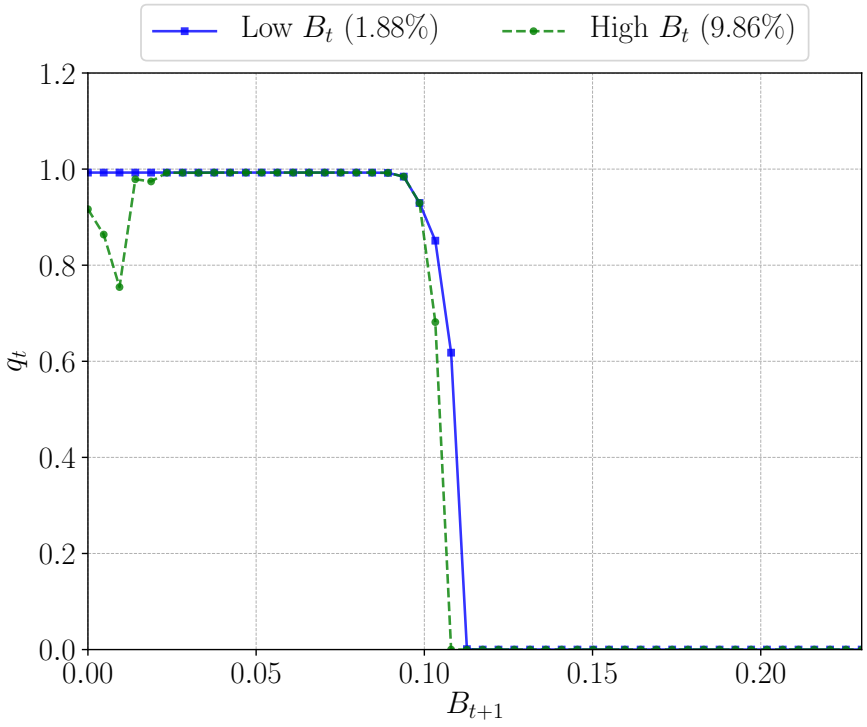


the average level of 9 per cent, and different realizations of current endowment are considered. As Y_t falls, the pricing function moves leftwards: For the same B_{t+1} the price is lower or, at best, unchanged. The intuition is straightforward. Being the center of the signals' distribution, a decrease in Y_t means that the entire mass of signals is displaced toward lower levels. On average, all investors are more pessimistic, and so does the marginal investor. On turn, this depresses the equilibrium price.

Also, outstanding debt B_t is a new argument of the pricing function. This result is driven by the assumptions *à la* Cole and Kehoe (2000) of the present model. A default may occur only *after* investors have purchased B_{t+1}, and these investors will suffer a loss if default indeed occurs. Given that the amount of B_t is key in establishing the probability of a default at time t, this becomes a relevant variable in pricing B_{t+1}. Figure 1.8 plots pricing functions similar to those of Figure 1.7, but instead I fix Y_t at its average level and I consider two different levels of B_t.

For any level of B_{t+1}, the equilibrium price is lower when the maturing debt B_t is

Figure 1.8: Pricing function - Changing outstanding debt



larger. Clearly, a larger B_t increases the chances of a default in the current period, and this reduces the equilibrium price of B_{t+1}. A more interesting result is that, for large enough B_t, the pricing function becomes non-monotonic in B_{t+1}. In this aspect, the present pricing function is remarkably different from the traditional one. In the standard models, the pricing function is monotonically decreasing in B_{t+1}, since the probability of a default on B_{t+1} is monotonically increasing in B_{t+1}. The last assertion is valid also in the present setup, as one might infer from the shape of the default region in Figure 1.5: Low values of B_{t+1} create less incentives for a default *in period* $t+1$. However, low values of B_{t+1} mean less resources and less spending in period t, and that actually creates more incentives to default in the current period. As a consequence, the price of B_{t+1} decreases. In summary, the intuition is that a safe level of B_{t+1} should be low enough to reduce the chances of a default in the next

period, but also large enough to minimize the risk of a default in the current period.

To conclude the discussion about the pricing function, I want to point out the sharp steepness over some intermediate debt levels. These levels roughly correspond to the interval of B_t with probability of default strictly in $(0, 1)$, described when discussing the default region. When B_{t+1} leaves the safe region with boundary at 9 per cent and moves to the right, the price tumbles very quickly. Investors appear really sensitive to the size of B_{t+1} and they discount a lot the risk coming with larger debt issuance. This behavior explains the government's choice of low risk debt discussed above. Facing borrowing conditions becoming quickly harsher, the government is better off by remaining at debt levels perceived as safe by the investors.

1.4.5 The amplification effect of private information noise

Having established that the model gives a good representation of the actual spread dynamics, we can use it as a lab environment for experiments. By means of a counterfactual exercise, I intend to establish if and how the level of private information noise affects the pricing of government debt, and thus the sovereign spread (recall that the model predictions derived analytically are ambiguous). In practice, I use the model to answer the following question: How different would the spread statistics have looked, had the private information noise remained at the level prevailing before the Great recession? In essence, I exploit the structural break occurred after 2008 in my measurement of private information noise in the euro area, and I explore how different the model simulations look like under the hypothesis that there is no structural break. This exercise also sheds light on the contribution of the increase in private information noise toward the spread dynamics observed during the European Debt Crisis.

In the counterfactual experiment, I simulate the model keeping all the parameter values as in the baseline model but one: Signal noise σ_x, which is set at the pre-2008 average level of 0.10 per cent. I seed the random number generator in such a way that the realizations of endowment shocks in the counterfactual simulations are identical to those used to generate the baseline simulations. Hence, any difference

between the counterfactual and baseline simulations is entirely driven by the change in σ_x. Table 2.5 shows the key spread statistics obtained from the counterfactual simulations along with the baseline statistics from Table 1.4. Under the counterfac-

Table 1.5: Moments - Simulation vs. Data

	Baseline $\sigma_x = 0.12\%$	Counterfactual $\sigma_x = 0.10\%$	Change %
$m(spr_t)$	11 *bps*	8 *bps*	-27.3%
$sd(spr_t)$	0.24%	0.22%	-8.3%
$\rho(spr_t, y_t)$	-0.38	-0.28	-26.3%
$\rho(spr_t, B_t)$	0.13	0.07	-46.1%
Probability of Default	0.12%	0.09%	-25%

tual hypothesis, the values of all statistics are slightly less than the baseline ones. In absolute terms, the differences are small, and this is not surprising: The baseline values of spread mean and volatility are already quite little in the first place, and the counterfactual shift in σ_x is also very small. Taking this into consideration, we should read the counterfactual findings in relative terms, as summarized in the third column of the table. The percent changes are not minimal. If private information noise had remained at pre-2008 levels, the euro area sovereign spreads would have been 27 per cent lower and less volatile by 8 per cent. These findings demonstrate that there exist an amplification effect on spreads, stemming from the noise in private information. Interestingly, with less noisy information, spreads are also less sensitive to GDP growth shocks and movements in the debt to GDP ratio. Indeed, in the counterfactual simulations the correlation between spread and GDP growth is 26 per cent lower and the correlation between spread and maturing debt to GDP ratio falls by 46 per cent, with respect to the baseline moments. Finally, the overall probability of default is lower when private information is less noisy: With the counterfactual σ_x, defaults are 25 per cent less frequent than in the baseline case.

In summary, the takeaway of the counterfactual experiment is that there exists an amplification effect of private information noise on sovereign spreads. With lower private information noise, sovereign spreads are on average smaller, less volatile, and less sensitive to fundamental shocks. The default probability is also lower. These conclusions have important policy implications. As an extra tool to reduce sovereign risk, a fiscal or monetary authority should focus on dissipating the noise of private information. Releasing more transparent and more reliable *public* information seems an effective strategy, to the extent that precise public information is able to "crowd out" noisy private information. The findings have implications for investing too. The profitability of investing in sovereign debt does not depend only on the realizations of random events, but also on the general disagreement about the likelihood of such events among investors. This is the essential consequence of incomplete information: Agents find themselves in a global game, characterized by strategic complementarities. Therefore, a successful investor in the sovereign market should not only gather as much private information as possible, to have the best forecast of the state of the economy, but she should also get a sense of how different her view is from that of other investors, since this is in itself a determinant of sovereign debt pricing.

1.4.6 Sensitivity of simulations to information parameters

In the next paragraphs, I intend to deepen the analysis of what happens to the spread statistics simulated by my model when the properties of information change. First, I focus on private information noise, and on its dual role of defining dispersion and precision of private information. Then, I explore a different assumption about the information which is common to both the government and the investors, specifically the knowledge of the actual stochastic process of endowment. I suppose agents' believes about endowment volatility depart from the actual one and I investigate the effect of changing in these believes on model simulations.

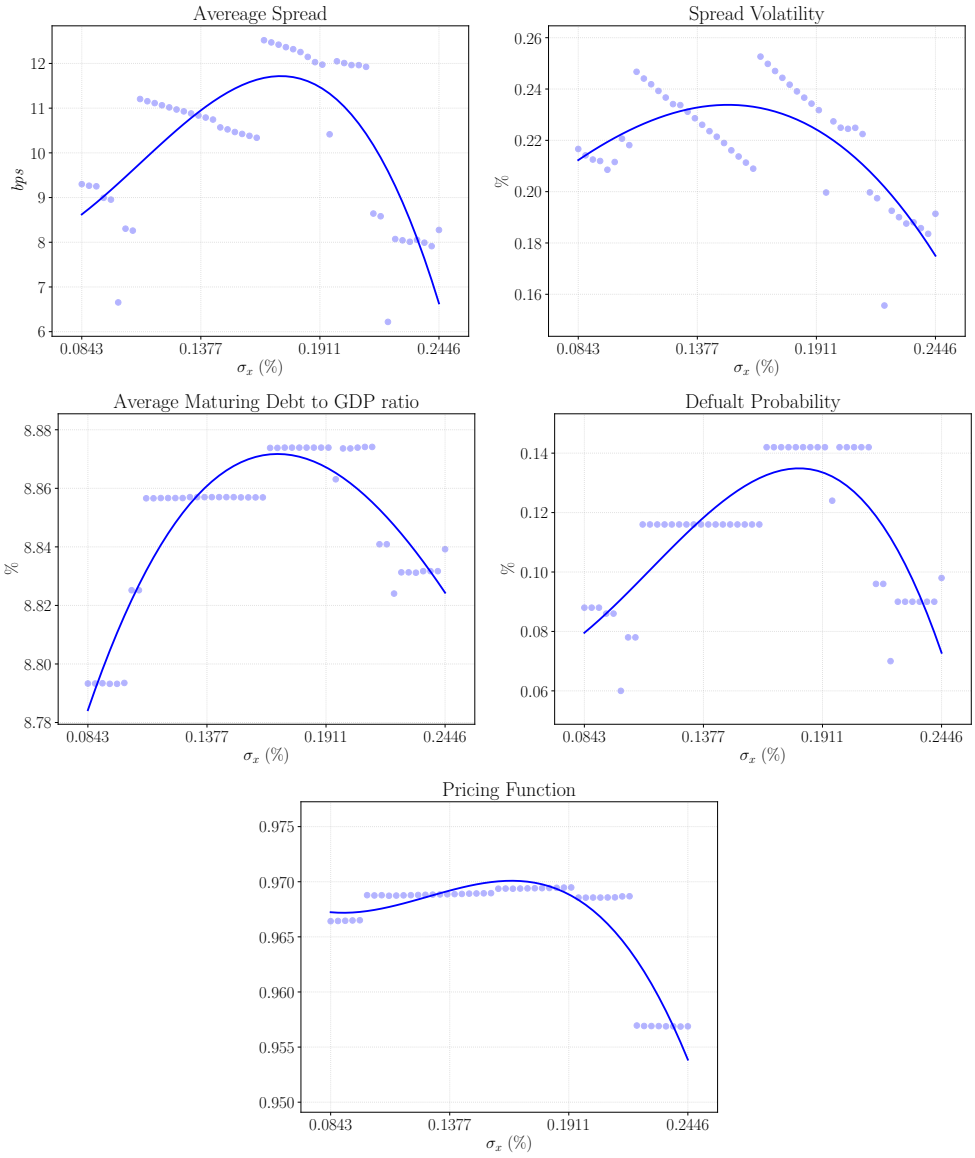
To study the sensitivity of model simulation to private information noise, I consider a grid of fifty values of σ_x, contained within the extremes of my measurement (0.084 per cent and 0.245 per cent respectively). I repeat the simulation of the model

varying the value of the parameter in question over this grid, recording the spread statistics and other variables of interest. Figure 1.9 displays the relation between the simulation results and the values of σ_x, whereby the pairs of signal noise and simulation result values are also smoothed using a polynomial approximation[2]. It appears clear that my model captures a hump-shaped relation between the level of private information noise and all the aggregates considered, which are the long run average level and volatility of sovereign spreads, the average of maturing debt to GDP ratio, and the default probability. As signal noise increases, the government holds more debt on average. Not surprisingly, this calls for a larger default probability and a higher average level of spreads. However, when σ_x becomes very large, the trend is reverted: The average stock of debt is reduced, and both spreads and default probability falls with it. Notice that spread volatility follows the same dynamic. Why does the government's borrowing decision react in such non-monotonic fashion to signal noise? A hint is to be found in the last panel of Figure 1.9, giving a representation of the equilibrium pricing function for each σ_x value in the grid considered. For each of these values, I solve numerically my model and I quantify the "average" equilibrium price by *a*) setting the previous period endowment at the unconditional mean and the outstanding debt at the average maturing debt to GDP ratio in the data, *b*) averaging out the levels of current endowment using the invariant probability distribution, *c*) computing the average of equilibrium prices over the range of debt levels obtained in the various simulations. In this way, the plot gives an idea of how signal noise influence the terms of borrowing faced on average by the government. Also in this case, we notice a slightly hump-shaped relation, which reflect the conflicting forces uncovered when discussing the comparative statics of the model. Hence, we observe that the government increases (reduces) its stock of debt whenever the terms of borrowing improve (deteriorate). Interestingly, the impact on the spread level induced by the change in the amount of debt offsets the opposite effect associated with movements in the terms of borrowing.

We may want to see whether in the data there is evidence in support of a hump-

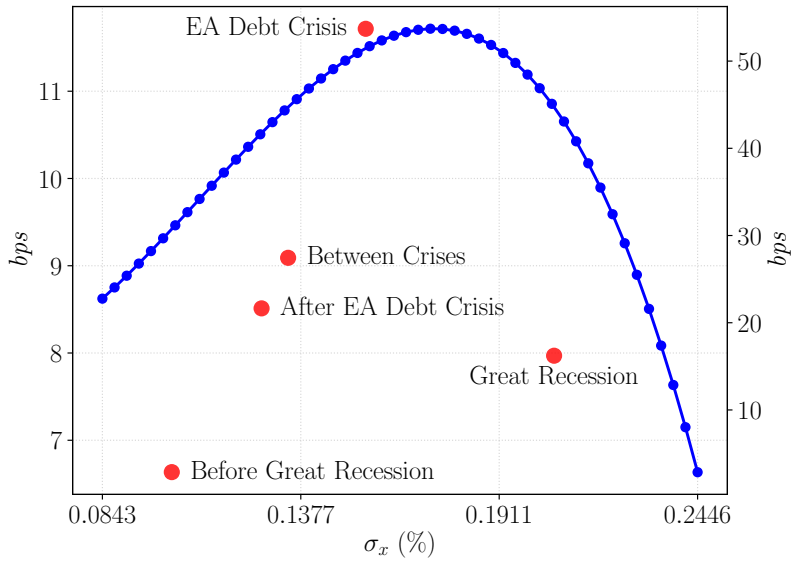
[2]The simulation findings have discontinuous patterns, due to numerical reasons. In light of this, I choose to use a smoothing technique

Figure 1.9: Signal Noise and Model Simulations



shaped relation between spread levels and private information noise. To this purpose, I compare the averages of observed spreads and measured noise for well-defined sub-periods of my data: The Great Recession, the European Debt Crisis, and the quarters before, after and in-between these events. A graphical representation of this exercise is provided for in Figure 1.10, along with the hump-shaped curve already showed in the first panel of Figure 1.9. Similar to the findings from model simulations,

Figure 1.10: Signal Noise and Spread: Model Simulations vs. Data



the sub-period averages suggest the existence of a non-monotonic relation between sovereign spreads and private information noise in the data as well. In particular, the largest spread level has been observed during the European Debt Crisis, a period characterized by a medium size of noise. On the other hand, private information noise reached its maximum during the Great Recession, when the spread has been larger than in the previous quarters but below the levels reached afterwards. In non-crisis periods, both spreads and noise levels have been generally small. However, the comparison between model simulations and the averages in the data should be taken with a grain of salt. Indeed, the former refers to equilibrium predictions of the model in the long run, while the latter are averages over short periods of time

that can hardly be seen as stationary. Nevertheless, I think that the exercise and the findings are evocative.

In Section 1.2, I have described the peculiar dual role the signal noise σ_x plays in the model. One the one hand, this parameter controls how dispersed the private signals are among the investors. On the other, it characterizes the precision of private information and, consequently, the weight investors attach to it in forming their expectations. I want to highlight the specific contribution of each of these aspects in driving the hump-shaped relation evidenced in the previous paragraph. To this end, I do the following. First, in the numerical solution I vary the value of σ_x over the grid defined above only where it affects the distribution of signals and the selection of the marginal investor, but I keep the value of σ_x as in the calibration of Table 1.3 where the moments of investors' perceived probability distributions are defined. Then, I do the opposite, changing σ_x where it affects the perceived moments and keeping it at the baseline value where it matters for the dispersion of signals.

Let us consider how σ_x affects model simulations through the dispersion channel. Figure 1.11 shows a summary of the findings, in the same fashion of Figure 1.9. We immediately see that the hump-shaped relation is not present. Instead, the average levels of spread and maturing debt to GDP ratio are monotonically associated to the extent of signal dispersion. Also in this case, it appears that the mechanism leading to these results is moved by the average terms of borrowing. As visible in the last panel, the average equilibrium price strictly increases with signals' dispersion, inducing the government to accumulate more debt. However, the level of spread falls even though the stock of debt increases, hinting that the improvement in the terms of borrowing is not offset by the extra borrowing. This is reflected also in the fact that the default probability is decreasing in signals' dispersion as well. Nevertheless, it seems that the larger stock of debt eventually generates a higher spread volatility.

When signals' precision in considered, the findings are rather different, as pictured in Figure 1.12. The average spread level has a hump-shaped relation with precision, which is mirrored by the probability of default and the spread volatility. The average terms of borrowing and the average maturing debt to GDP ratio show, with some approximation, a co-movement similar to the one described above, although on the

Figure 1.11: Signal Dispersion and Model Simulations

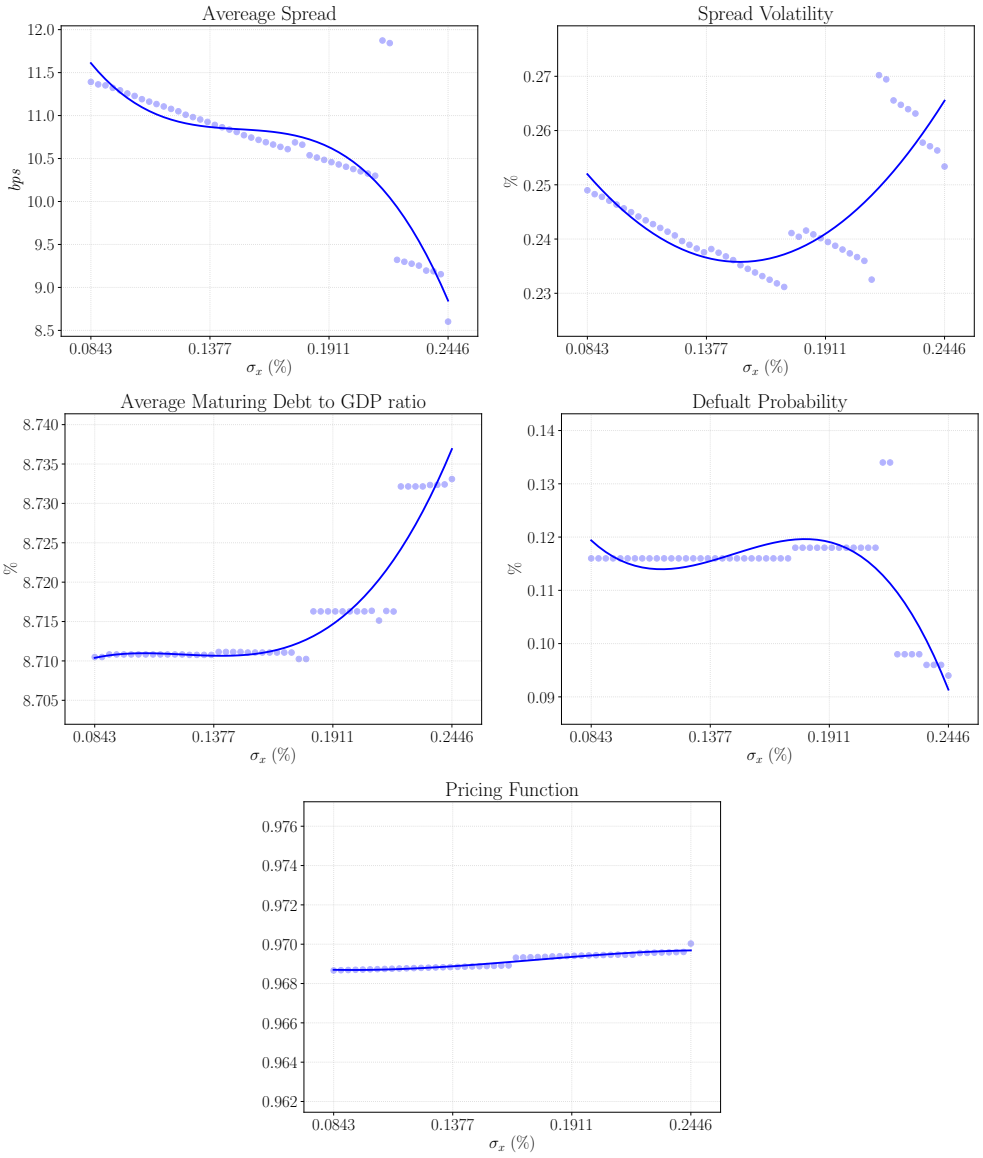
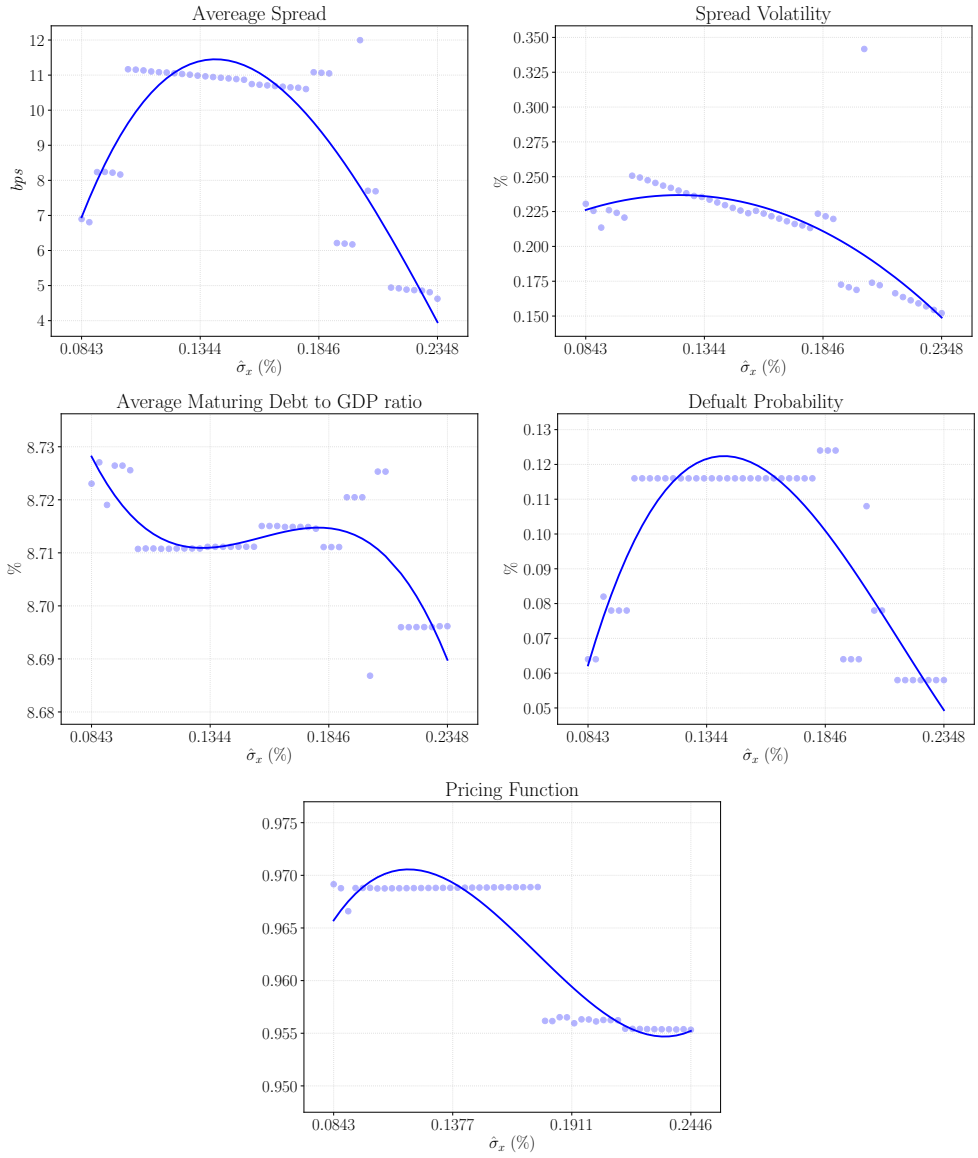


Figure 1.12: Signal Precision and Model Simulations



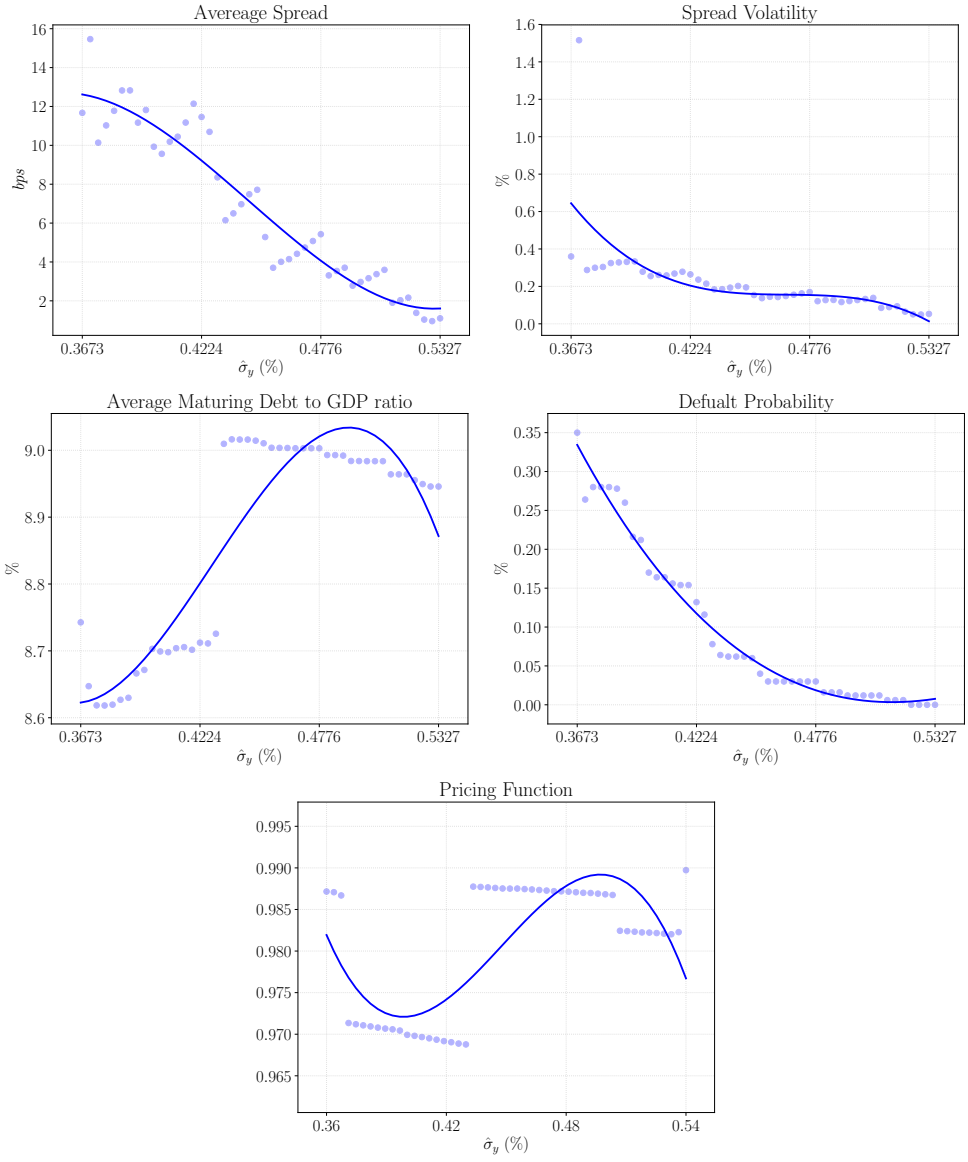
opposite direction. The equilibrium price tend to decrease as private information becomes more imprecise, and overall the government follows this decline by reducing the stock of debt. In summary, the hump-shaped sensitivity of model simulations to signal noise is a synthesis of the sensitivity to private information dispersion and precision. With regard to terms of borrowing and debt stock, the latter channels operate in opposite directions: An increase in dispersion (precision) induces better (worse) terms of borrowing and more (less) borrowing. The hump-shaped relation in the third and fourth panels of Figure 1.9 is given by the balance of these channels. On the other hand, the hump-shaped relation of average level and volatility of spread, and default probability with signal noise appears mainly driven by the precision channel, which dominates over the different effect of signals' dispersion.

At this point I change the focus of the analysis, departing from the private information noise. Instead, I entertain the possibility that both the government and the investors do not know exactly the actual probability distribution of endowment. In particular, I assume they hold the belief that endowment volatility is $\hat{\sigma}_y$, potentially different from the actual volatility σ_y. Hence, both the government and the investors have a *prior* view that log endowment y_t is distributed according to

$$y_t \sim \mathcal{N}\left(\hat{y}_t, \hat{\sigma}_y^2\right).$$

I experiment by varying the value of $\hat{\sigma}_y$ over a fifty-point grid and looking at the implications for model simulations. Notice that a change in $\hat{\sigma}_y$ modifies government's expectations about endowment realizations and investors' perceived probability distributions. However, the generation of random endowment paths in simulating the model reflects the actual endowment volatility σ_y, which is kept fixed at the value of Table 1.3. Figure 1.13 presents the findings of model simulation sensitivity to $\hat{\sigma}_y$. One can see that average level and volatility of spread and default probability are decreasing in the level of perceived endowment volatility. The borrowing terms and the average debt to GDP ratio follow an approximately hump-shaped path with respect to $\hat{\sigma}_y$, in general respecting the regularity that better (worse) terms of borrowing are associated with more (less) government debt. To reconcile the patterns of average

Figure 1.13: Prior Volatility and Model Simulations



spread and default probability with those of debt to GDP ratio and equilibrium price, it should be the case that, as initially the latter two increase, the improvement in the terms of borrowing dominates the enlargement of the stock of debt, leading to less risk of default and lower spread on average. On the other hand, when the terms of borrowing deteriorate, the debt de-leveraging by the government is strong enough to avoid an increase in default probability and spreads.

1.5 Conclusion

This chapter fills a gap in the theoretical and quantitative literature of sovereign debt and default by studying the impact of incomplete information, *i.e.* the existence of noisy private information, on the pricing of sovereign debt. The main conclusion is that the effects of the noise in private information on the average level and volatility of sovereign spreads have a hump-shaped fashion. For low initial levels, an increase in private information noise amplifies spread level and volatility, increasing the risk of default. However, when noise is very large, a further increase reduces spreads and spread volatility. This result is mainly driven by how the equilibrium pricing of government debt reacts to changes in private information noise. Intuitively, an increase in dispersion and imprecision of private information is self-defeating, because investors attach less weight to it in forming their expectations. These findings have important implications from the point of view of a government that intends to reduce its borrowing cost, and for an investor aiming at maximizing the expected return from purchasing sovereign debt.

With regard to the theoretical contributions of this chapter, I propose a way to introduce incomplete information in an otherwise standard infinite-horizon model of sovereign debt and default. I discuss under which conditions there exists a unique equilibrium, as opposed to the multiplicity of equilibria resulting in this class of models with complete information. Quantitatively, the inclusion of incomplete information enables my model to generate a level of spread volatility consistent with the data, something on which standard models with complete information has generally

failed.

From an empirical standpoint, I propose and implement a structural methodology to measure the level of private information noise from observed forecast data. Based on the ECB Survey of Professional Forecasters, this method delivers a time series of noise with interesting properties. First, private information noise spikes when the economy is in a bad state. Second, the Great Recession caused a structural break in the series: Private information noise in the euro area remained at levels persistently larger than the pre-2008 average ones. I think this latter finding is in its own very relevant, regardless the application in the present context. It should definitely deserve further research, to better understand its origin and its different implications.

In this chapter, private information is modeled in a simple reduced form: Exogenous signals, which are right on average but individually noisy. The next step may be that of modeling a more realistic process of information acquisition and learning. An ambitious but interesting extension may see not only the level of noise influencing the pricing of debt, but also the latter as well as the realization of endowment affecting on the noise itself, creating en endogenous feedback effect that might explain noise and spread co-movements.

CHAPTER 2

WHY IS EUROPE FALLING BEHIND? STRUCTURAL TRANSFORMATION AND SERVICES' PRODUCTIVITY DIFFERENCES BETWEEN EUROPE AND THE U.S.

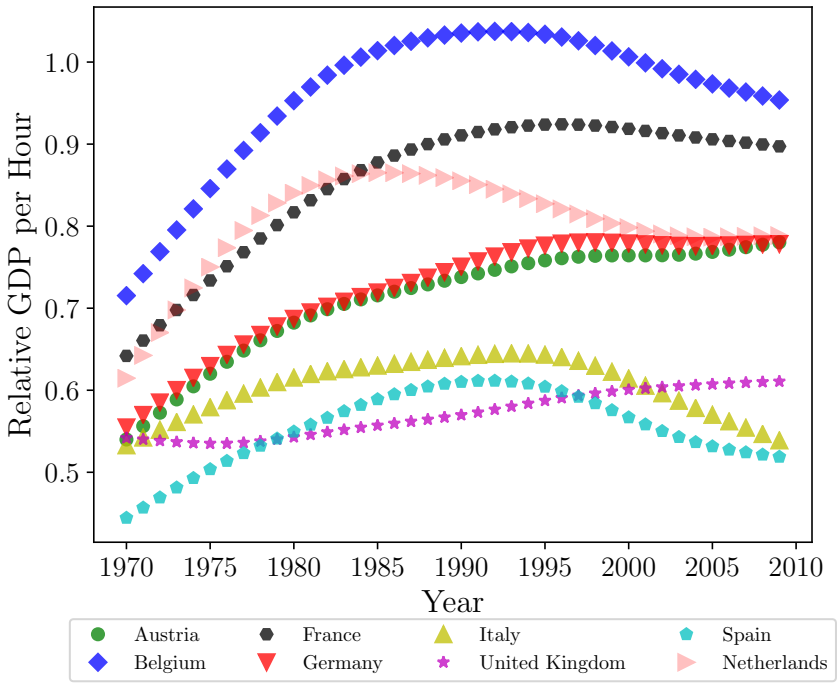
Joint work with Luis Felipe Sáenz and Joao B. Duarte

2.1 Introduction

Labor productivity in Europe has been falling behind the United States since the beginning of the 1990s, reversing a previously observed pattern of convergence between these two economies. Figure 2.1 illustrates how this process of catch-up came to a halt and later even reversed for the majority of the European countries. Average annual labor productivity (measured as GDP per hour of work) in the U.S. accelerated from 1.3 per cent in the 1970-1990 period to 1.7 per cent from 1990 to 2009 while the European countries on average experienced a labor productivity growth slowdown between these two time periods from 2.9 per cent to 1.5 per cent. The divergence is a combination of the U.S. taking off together with a European slowdown.

During this period, these economies underwent large scale sectoral reallocations of labor in a process commonly known as structural transformation (Kuznets (1957); Herrendorf, Rogerson, and Valentinyi (2014)). With Europe and the U.S. at their later stages of structural transformation (the so-called post-industrial era), labor has reallocated further away both from agriculture and manufacturing toward services. As Duarte and Restuccia (2010) suggest, through the lenses of structural transformation it is possible to conclude that the service sector is responsible for most cases of relative stagnation in aggregate productivity observed at later stages of develop-

Figure 2.1: Relative aggregate labor productivity



GDP per hour worked relative to the United States. We used GDP per capita measures from the Maddison Project to measure the PPP-adjusted aggregate labor productivity of each European economy relative to the U.S. for 1970. Then, we used the World KLEMS to compute the remainder of the time series with annualized growth rates of aggregate labor productivity.

ment since almost no other country experienced the productivity gains in the service sector witnessed in the U.S.

We believe that it is crucial to break down the service sector in order to understand the relative under-performance of Europe *vis-á-vis* the U.S. Services constitute the predominant (and growing) sector for the vast majority of advanced economies, and the lack of labor productivity gains in this sector is an increasing cause of concern for long-run economic growth. In this chapter, we put forth a theory of structural transformation and decompose the service sector into sub-sectors comparable across Europe and the U.S. to investigate how changes in labor allocations brought by changes in sectoral productivity explain the (relative) slowdown of the European

aggregate labor productivity.

First, using the World KLEMS database, we decompose services into 11 comparable sub-sectors[1]. We document that the reallocation of labor toward the various types of services has followed similar patterns both in Europe and the U.S. Motivated by these facts, we develop a theoretical model of structural transformation that combines the CES non-homothetic preferences crafted by Comin, Lashkari, and Mestieri (2015) with production functions whose unique input is labor, as in Duarte and Restuccia (2010). Our model economy includes a total of 13 sectors: agriculture, manufacturing, and the 11 service sub-sectors. We calibrate the model to account for the the U.S. development experience, and we then use it to measure comparable sectoral labor productivity levels for the 13 sectors in all the European countries of our sample. The tests for our theory are based on the model's capacity to explain the structural transformation in Europe and the U.S. as well as the relative differences in aggregate productivity. We show that the model is quantitatively able to reproduce the labor allocation in the vast majority of the sectors in all countries, and the main stylized fact presented in Figure 2.1. We perform counterfactual experiments to identify which services have been dragging down the aggregate labor productivity in Europe, and last, we empirically explore our country-sector panel measures of labor productivity levels to assess the importance of various input factors in determining the performance of services' labor productivity of Europe relative to the U.S. In particular, we decompose the levels of relative sectoral labor productivity measured with our model into the contributions stemming from sectoral physical and information and communication technology (ICT) capital to labor ratios, and sectoral total factor productivity (TFP).

Our quantitative experiment suggests substantial differences in sectoral labor productivity of services between Europe and the U.S. The European countries are in generally more productive than the U.S. in communication, education, real estate, and health services. However, the European countries are less productive in wholesale and retail trade and business services, with sectoral labor productivity levels of approximately 20 per cent of that of the U.S. Led by our counterfactual experiments,

[1]We classify these sectors according to the ISIC Rev. 3 at one digit level.

we identify wholesale and retail trade, business services, and, to a lesser extent, financial services as the sectors responsible for most of the divergence in aggregate labor productivity between Europe and the U.S. We find that if Europe had experienced the same gains in labor productivity as the U.S. in wholesale and retail trade and business services *alone* since 1990, it would have had a 3.2 per cent and a 2.4 per cent higher aggregate labor productivity in 2009, respectively. In fact, if Europe had caught up with the U.S. in the labor productivity of wholesale and retail trade and business services by 2009, the aggregate labor productivity in Europe would have been 25.8 per cent and 17.1 per cent higher, respectively. We also show that if the European financial services had caught up with the U.S. in terms of labor productivity by 2009, the gains on aggregate labor productivity would have been only 1.5 per cent.

Why the labor productivity in wholesale and retail trade and business services had a poor performance in Europe? We find that most of the productivity gap in the various services between Europe and the U.S. is accounted for by differences in sectoral TFP. This is particularly relevant in wholesale and retail trade and business services, where relative sectoral TFP represents, on average, approximately 90 per cent of relative sectoral labor productivity. These two sectors had the lowest average levels of relative sectoral TFP between 1990 and 2009 among all the services, and these levels kept falling over this period. In addition, we find that during the years of the falling behind (1990-2009) the level of physical and ICT capital endowment per hour worked in Europe, relative to the U.S., fell significantly in the service sector. This fact clearly contributed to the lower level of services' labor productivity of Europe compared to the U.S. We identify that the fall in ICT to labor ratio hurt more the productivity of wholesale and retail trade and business services. While employment in these two sectors increased between 1990 and 2009 in Europe even faster than in the U.S., the level of sectoral ICT utilization actually decreased in comparison to the U.S., dragging down the labor productivity in these services.

Our first contribution is to document comparable disagreggated services' labor reallocation and labor productivity dynamics across Europe and the U.S. We classify services industries from the World KLEMS data into eleven sectors that are com-

parable across a large set of European countries and the U.S. Thus, we extend the Timmer, Vries, and de Vries (2014) database on productivity from 5 to 11 service industries for selected European countries and the U.S. Our documentation of labor reallocation within the service sector is in line with the explanation of the rise in services due to the marketization of home production, as shown by Buera and Kaboski (2012) and also thanks to an important expansion in services oriented to businesses.

Our second contribution is to show that shift-share analysis would underestimate (overestimate) the effect of productivity gains in wholesale and retail trade (business services) on aggregate productivity. In the workhorse models of structural transformation, the labor allocation across sectors is responsive to changes in the level of income and to changes in the sectoral relative productivity. As productivity changes, shifts in the sectors' employment shares occurs endogenously. Our model accounts for these general equilibrium effects. Hence, our identification of the relevance of sectoral labor productivity levels based on model counterfactuals incorporates the endogenous changes in labor shares deriving from considering alternative productivity paths. In this respect, we argue that our approach is superior to other quantitative methods, such as shift-share analysis, in which changes in sectoral labor shares and sectoral labor productivity cannot be studied simultaneously. Indeed we show that the endogenous changes in sectors' weights, *i.e.* labor shares, resulting from our counterfactual analysis are significant, thus, we show that counterfactuals that disregard general equilibrium effects on labor shares are biased. For instance, in contrast to Timmer, Inklaar, O'Mahony, and van Ark (2011), we find that manufacturing productivity growth does not have a sizable impact on aggregate productivity, that business services had much more important role, and that financial services had a smaller role in the slowdown of the European relative aggregate labor productivity.

This chapter is related primarily to the literature of structural transformation that dates back to the works of Kuznets (1957) who documented the sweeping changes across the different industries in the process of economic development. More recent contributions to structural change build upon the works of Kongsamut, Rebelo, and Xie (2001) and Ngai and Pissarides (2007) who emphasized the role of income and sector-biased productivity channels respectively as the drivers of structural trans-

formation. Several attempts have been made to incorporate both mechanisms in a single framework, such as Buera and Kaboski (2009) and Duarte and Restuccia (2010) among many others.[2] Our chapter uses the long-run Engel curves proposed by Comin et al. (2015) to study productivity differences in the service sector in a framework where labor is the unique production input, and shows that the model is quantitatively successful in capturing the structural transformation across most sectors and countries in our sample.

The widening of the productivity gap between Europe and the U.S. that occurred in the last decades has been the focus of many past studies. The large majority of this literature has studied productivity growth, rather than levels, and relied on growth accounting techniques and shift-share analysis. van Ark, Inklaar, and McGuckin (2003), Inklaar, Timmer, and van Ark (2008), and Timmer et al. (2011) identify ICT as a main source of problems for labor productivity in Europe, providing evidence that both ICT-producing and ICT-utilizing sectors performed badly in Europe compared to the U.S. The different approach to ICT utilization between Europe and the U.S., and its effects on labor productivity, is the main point of Bloom, Sadun, and Van Reenen (2012) too. Relative to these studies, we show empirically that the lack of physical capital investment also played an important role in explaining the productivity level gaps between Europe and the U.S. The diversity in levels of sectoral productivity across countries is also treated by Lewis (2005), based on the case-study analysis of the McKinsey Global Institute. The conclusion of Lewis (2005) is that market regulations have been a much greater obstacle to competition in Europe than in the U.S., and a major factor in creating productivity differences. The role of regulation on labor productivity is also the focus of Nicoletti and Scarpetta (2003), Crafts (2006), and Cette, Lopez, and Mairesse (2016). In this chapter, given our findings on the crucial role played by TFP differences in explaining labor productivity gaps, we cannot rule out the relevance of regulation in explaining labor productivity differences found by these previous studies. However, we highlight the need of more detailed data on services regulation measures and a better understanding of how

[2]For a detailed survey of the literature of structural change see Matsuyama (2008) and Herrendorf et al. (2014).

regulation affects competition in services.

The rest of the chapter is organized as follows: Section 2.2 discusses the main stylized facts of structural transformation within services. Section 2.3 develops a simple conceptual framework that extends the structural transformation model of Comin et al. (2015) to include service sub-sectors. Section 2.4 calibrates the baseline model. Section 2.5 uses the calibrated model to measure the first period levels of sectoral productivity in Europe and tests the model predictions against the data. Section 2.6 presents the counterfactual exercises that quantify the relevance of each sector in aggregate labor productivity. Section 2.7 explores the components of services' sectoral labor productivity levels, and how they contributed to forming the productivity gap between Europe and the U.S. Finally, Section 2.8 provides the concluding remarks.

2.2 Facts on a Disaggregated Service Sector

We use the World KLEMS[3] data on hours worked and value added to document both the process of labor reallocation and the labor productivity growth of disaggregated service industries. We make use of the International Standard Industry Classification (ISIC) Rev. 3 at the two digits level to classify thirteen comparable sectors. We aggregate agriculture and manufacturing in the same way these sectors are aggregated in past studies in which the analysis is restricted to three sectors[4]. However, our data allows us to disaggregate the service sector into eleven different comparable sub-services.

Country-wise, our objective is to have the most disaggregated service sector possible comparable across the largest set of European countries and the U.S. To reach this goal given data constraints, we restrict our sample to nine countries from 1970 to 2009. The countries that meet our selection criteria in this chapter are Austria, Belgium, France, Germany, Italy, Spain, the Netherlands, the United Kingdom, and the U.S. Table 2.1 presents the most disaggregated service sectors' classification possible

[3]For more details see O'Mahony and Timmer (2009).

[4]See for instance Duarte and Restuccia (2010).

in order to have comparable measures across the European countries with the U.S. In the quantitative section and in our counterfactual experiments, for comparison purposes between Europe and the U.S. we often discuss European averages. By this we mean the average of the eight European economies weighted by their GDP size. All data are trended using the Hodrick-Prescott filter with a smoothing parameter $\lambda = 100$.

2.2.1 Service Sector Structural Transformation

Our data on labor shares from 1970 to 2009 for the European economies and the U.S. show that the employment in these economies is dominated by services, as these countries during our sample period have experienced a large reallocation of labor from both agriculture and manufacturing into services. We are interested in

Table 2.1: Sectors' classification

Code	Name	Section
`agr`	Agriculture, hunting and forestry	A
	Fishing	B
`man`	Mining and quarrying	C
	Manufacturing	D
	Electricity, gas, and water supply	E
	Construction	F
`trd`	Wholesale and retail trade	G
`rst`	Hotels and restaurants	H
`trs`	Transport and storage	I(60-63)
`com`	Post and telecommunication	I(64)
`fin`	Financial intermediation	J
`res`	Real estate activities	K(70)
`bss`	Renting and business activities	K(71-74)
`gov`	Public administration and defense	L
`edu`	Education	M
`hlt`	Health and social work	N
`per`	Other community, social and personal activities	O

Table 2.2: Structural transformation within services

	Emp. Share in Europe: 1970			Emp. Share in Europe: 2009		
	Sector	*%*	*Relative*	*Sector*	*%*	*Relative*
1	trd	13.53	0.97	bss	14.99	1.08
2	gov	6.14	0.81	trd	14.79	1.09
3	hlt	4.46	0.39	hlt	9.33	0.54
4	trs	4.08	1.18	per	6.98	1.03
5	per	4.04	1.01	gov	6.46	2.08
6	bss	3.91	0.61	rst	5.34	0.81
7	edu	3.19	0.49	edu	5.33	0.68
8	rst	3.05	0.78	trs	4.48	1.32
9	fin	2.05	0.59	fin	2.99	0.68
10	com	1.63	0.6	com	1.39	0.87
11	res	0.36	0.45	res	1.00	0.74

Services' employment shares in Europe – absolute and relative to the U.S. – for the first and last year of our sample.

documenting if there is an historical pattern in the way labor is allocated within services. Our goal then is to document the labor allocation taking place within services.

Our disaggregated services' data on labor shares suggest that with the exception of communication and government, there is a systematic rise in the labor share of all service industries; additionally, the employment in health and business services is growing faster than services as a whole.

Table 2.2 presents the average sectoral labor shares of services for the European average and how these labor shares compare relative to the U.S. in the first and last years of our sample. Between the two periods, all service industries increased their labor share, except communication in Europe and government in the U.S. In addition, we observe that the rise of the service sector was outpaced by the surge in business and health services. The business services' labor share evolved from being the sixth sector with highest labor share in the European economy to being first, despite the fact that the majority of all other service sub-sectors increased their labor share during the same period. In contrast, the labor share remained

relatively constant for some previously large service sub-sectors, such as government and wholesale and retail trade. Finally, we observe that the ratios of the labor shares of Europe relative to the U.S. increased for all service industries, indicating a convergence in the composition of the labor force within the service sector.

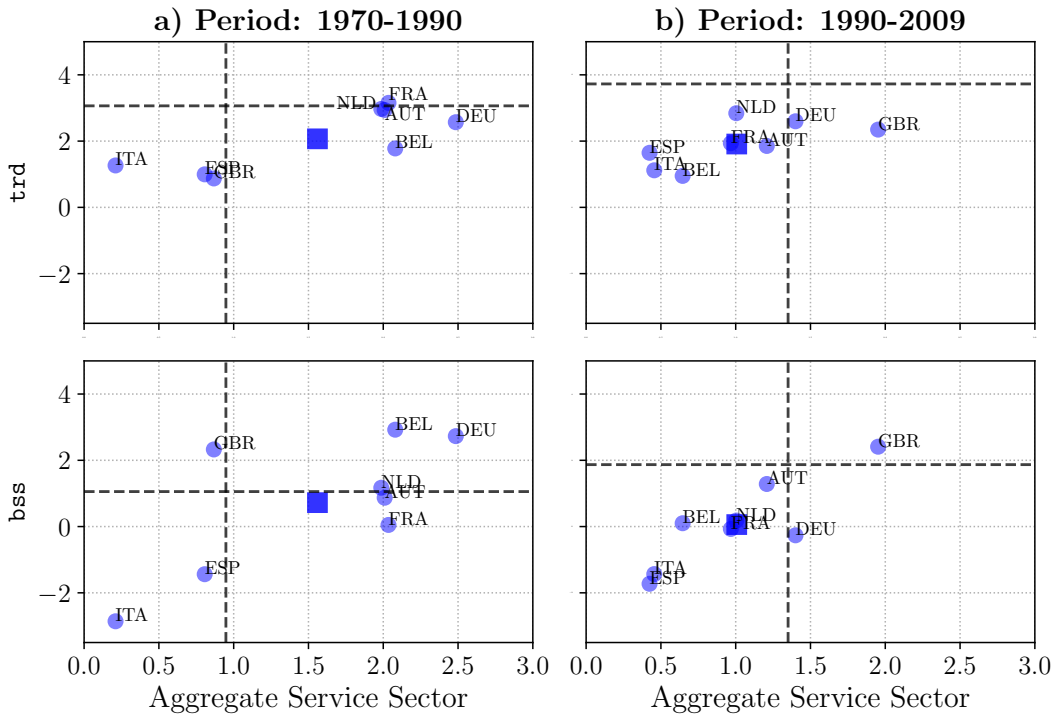
2.2.2 Services' Labor Productivity

From 1970 to 2009, the U.S. annualized labor productivity growth rate in the service sector was approximately 1.1 per cent. Except for Italy and Spain, all European countries experienced a higher growth rate than the U.S. in aggregate service labor productivity for the same period. However, simply looking at the entire sample hides two very distinct phases – one of strong catch-up (1970-1990) and another of stagnation and divergence (1990-2009). We perform a sub-sample analysis of these two periods and we find that the U.S. accelerated from approximately 1 per cent growth in aggregate services' labor productivity in the first period to 1.4 per cent in the second period. At the same time, most European countries experienced a major slowdown in services' average labor productivity between the two periods[5], with the European average growth rate in services' labor productivity falling from 1.6 per cent to 1 per cent.

The disaggregated data on labor productivity measured as real valued added per hour worked calls attention to the fact that, relative to the U.S., European countries had a significantly higher productivity growth in health and personal services while they had a significantly lower productivity growth in wholesale and retail trade and business services. Figure 2.2 compares the relative performance of the latter two sectors in the European economies and the U.S., between the two sub-sample periods. The scatter plots describe a positive correlation between the labor productivity growth rate of each of these sectors and aggregate services. In addition, wholesale and retail trade had strong gains in labor productivity and the U.S. was the leading country in this sector in both periods. Between the two periods, the U.S. accelerated

[5]One exception being the United Kingdom which accelerated even faster than the U.S. from approximately 1 per cent to 2 growth in services' average labor productivity.

Figure 2.2: Average growth in services' productivity



Scatter plots of value added per hour annualized growth rate of the aggregate service sector with the value added per hour annualized growth rate of business services and wholesale and retail trade. The horizontal lines indicate the service sectoral labor productivity growth rates observed in the United States, and the vertical line indicates the aggregate service labor productivity growth rate of the United States for both periods. The blue square marker indicates the annualized growth of labor productivity growth pairs for Europe.

from 3 per cent to almost 4 per cent while the European economies maintained the same growth rate. On the other hand, business services' productivity in the U.S. accelerated between the two periods doubling its growth rate from 1 per cent to 2 per cent, while most European countries suffered a slowdown in this sector.

2.3 Model

This section presents a model of structural transformation with agriculture, manufacturing, and 11 different services, where the process of structural transformation depends on income and price effects. We choose the number of sectors in the model to account for the same sectors explored in the previous section. The model borrows the production structure from Duarte and Restuccia (2010) and the preferences from Comin et al. (2015). By combining these two frameworks, Engel curves and heterogeneous labor productivity growth rates are sufficient to account for the structural transformation. The model does not have capital (consistent with Duarte and Restuccia (2010)), which means that there is no investment sector in this economy, and that the model has no dynamic component. Therefore, the structural transformation, namely the reallocation of labor over time across sectors, is taken as a sequence of static optimal allocations.

2.3.1 Environment

In our model economy there is an infinitely lived stand-in household of measure L that supplies labor inelastically.[6] Its only endowment is time. There are thirteen sectors, and each sector produces its good or service using labor as the unique input. In addition, labor moves freely across these sectors.

Household

The household has preferences over its consumption stream over time, but since we are not defining inter-temporal problems in our model (*i.e.* there are no savings), there is no need to formalize the structure of preferences toward the inter-temporal substitution of consumption. Therefore, following Comin et al. (2015), the intra-temporal choice problem is described by a representative household that has

[6] Alternatively, one can think of a household of measure one with and endowment of L hours each period. In this case, the definition of the measure is trivial, in spite of allowing growth of the labor force, because the structural transformation is a sequence of static choices.

preferences over the consumption of commodities (or services) produced in different sectors, represented by

$$\sum_{i \in I} \Omega_i^{\frac{1}{\sigma}} C^{\frac{\epsilon_i - \sigma}{\sigma}} c_i^{\frac{\sigma-1}{\sigma}} = 1, \tag{2.1}$$

where C is the aggregate consumption[7], I is the set of sectors of the economy, c_i is the consumption from output produced in sector i, $\sigma \in (0,1)$ is the price elasticity of substitution, $\epsilon_i \geq 1$ is the income elasticity for good i, and $\Omega_i > 0$ are constant weights for each good i, $\sum_{i \in I} \Omega_i = 1$. Notice that there are no time subscripts since the model is static. There are three main reasons[8] that support the use of this particular non-homothetic CES preference structure to explain the structural transformation in our model of 13 sectors. First, it naturally extends for any arbitrary number of sectors, which is not a feature of other types of preferences such as in Boppart (2014), Herrendorf, Rogerson, and Valentinyi (2013), and Duarte and Restuccia (2010), among many others. Second, it gives rise to heterogeneous sectoral log-linear Engel curves that are consistent with the empirical evidence (Aguiar and Bils (2015); Comin et al. (2015)). Last, the income effects on the relative consumption of sectoral goods and services do not level off as income rises, contrary to structural transformation demand-side theories that rely on Stone-Geary preferences, which is a crucial feature to account for the rise of services in the long-run. Therefore, these preferences allow the demand channel to have a strong role at later stages of development. The household's problem is defined as follows:

Household's Problem

[7] In the empirical counterpart of the model C is considered as income per capita since there are no savings in our model.

[8] There is greater detail in the exposition of other useful features of the non-homothetic preferences in Comin et al. (2015). In our chapter, we highlight the most useful ones for our particular purpose of decomposing extensively the service sector.

$$\max_{c_i} C \quad \text{s.t.} \quad \text{i)} \quad \sum_{i \in I} \Omega_i^{\frac{1}{\sigma}} C^{\frac{\epsilon_i - \sigma}{\sigma}} c_i^{\frac{\sigma-1}{\sigma}} = 1$$
$$\text{ii)} \quad \sum_{i \in I} p_i c_i \leq WL \tag{2.2}$$
$$\text{iii)} \quad c_i \geq 0,$$

where W is the wage of the household, WL reflects the total disposable income and p_i is the price of output c_i. We assume interior solutions, so the First-Order Conditions are sufficient. The optimal consumption of goods for each sector i is

$$c_i = \Omega_i \left(\frac{p_i}{P}\right)^{-\sigma} C^{\epsilon_i}, \tag{2.3}$$

and the optimal value added share of sector i is described by

$$\frac{p_i c_i}{PC} = \Omega_i^{\frac{1}{\sigma}} C^{\frac{\epsilon_i - \sigma}{\sigma}} c_i^{\frac{\sigma-1}{\sigma}}, \tag{2.4}$$

where P is the aggregate price index. Notice that the parameters ϵ_i and σ describe the income and price mechanisms of the structural transformation. Whereas ϵ_i measures the sensitivity for changes in consumption of goods from sector i with respect of changes in income, namely the Engel curve for sector i, σ reflects how sensitive the quantities demanded are toward changes in prices. For the empirical relevant case of $\sigma < 1$, where all goods are gross complements, the price effect illustrates the so-called Baumol's cost disease in which, in this context, labor is continuously allocated toward less productive sectors in the long-run.

Firms

In each period, there are 13 different goods produced in agriculture, manufacturing, and eleven types of services, as described in the previous section. There is a large number of competitive firms in each sector i that use a technology of production linear in labor described by

$$y_i = A_i l_i \qquad \forall i \in I, \tag{2.5}$$

where y_i represents the output produced by a representative firm of sector i, A_i reflects the labor productivity of the firm, and l_i is the labor input demanded by the firm, measured in labor hours. The firm in this model economy hires labor at the prevailing wage W – that is the same for each sector i since labor is perfectly mobile – and produces output with the combination of labor hours and an idiosyncratic labor productivity level for each one of the 13 representative firms. The firms' problem is described as follows:

Firms' Problem

$$\max_{l_i}\{p_i A_i l_i - W l_i\} \qquad \forall i \in I. \tag{2.6}$$

Again, if one assumes interior solutions the First-Order Conditions are sufficient to describe the optimal allocations of the firm. The optimal price is described by

$$p_i = \frac{W}{A_i} \qquad \forall i \in I. \tag{2.7}$$

Equation 2.7 shows that increases in sectoral labor productivity reduce the price of a good produced in sector i, and that increases in wages have a positive impact on prices. However, notice that wages do not change the relative prices in the economy since, by assumption, all sectors in the economy pay the same rental rate of labor. Thus, it is only through heterogeneous dynamics of the labor productivity across sectors that one gets changes in relative prices. We consider labor as the *numéraire* in our model economy and normalize its price – the wage rate W – to one, taking advantage that in our construction wages do not have sectoral implications for labor allocation. The sectoral price then is simply described as $p_i = 1/A_i \quad \forall i \in I$, and it is the inverse of sectoral labor productivity, as in Duarte and Restuccia (2010). Given the simplicity of the production technology, A_i can be considered as an exogenous reduced form measure of all of the structural factors that in reality affect labor productivity. In the empirical section we will address this issue by disentangling

the effects on the labor productivity coming solely from TFP *vis-á-vis* the effects coming through other production inputs. But for now one can think of these factors as components implicitly embedded in A_i.

Market Clearing Conditions

At each date, the market for each sectoral good and service clears

$$c_i = y_i \qquad \forall i \in I, \tag{2.8}$$

and the labor market also clears. The total demand for labor must equal the exogenous supply of labor by the household at every point in time:

$$\sum_{i \in I} l_i = L. \tag{2.9}$$

2.3.2 Equilibrium

DEFINITION: *A Competitive Equilibrium is a collection of exogenous labor productivity paths* $\{A_{i,t}\}$ *and optimal allocations* $\{c_{i,t}, l_{i,t}\}$ *such that for each period* t *and for each sector* i*:*

i) given prices, $c_{i,t}$ *allocations solve the household's optimization problem defined in 2.2;*

ii) given prices, $l_{i,t}$ *allocations solve the firm's optimization problem defined in 2.6;*

iii) market clearing conditions defined in 2.8 and 2.9 hold.

Combining equations 2.4, 2.5, 2.7 and the market clearing conditions in 2.8 one gets

$$\frac{Wl_i}{PC} = \Omega_i^{\frac{1}{\sigma}} C^{\frac{\epsilon_i - \sigma}{\sigma}} (A_i l_i)^{\frac{\sigma - 1}{\sigma}},$$

and after algebraic manipulation, we reach an expression for the sectoral labor demand

$$l_i = \left(\frac{P}{W}\right)^{\sigma} \Omega_i C^{\epsilon_i} A_i^{\sigma-1}. \tag{2.10}$$

Equation 2.10 illustrates the two main drivers of the structural transformation in our model. First, the parameter ϵ_i defines the Engel curve for sector i, and shows how this non-homotheticity affects the labor demand for each sector, linking it directly to the sector's income elasticity. Second, the parameter σ shows the relation of the price elasticity of substitution on the labor demand. As long as this parameter is smaller than one, increases in productivity will reduce the labor hours demanded in a given sector. Equation 2.10 predicts the levels of labor demand, and shows that aggregate prices and wages[9] also affect the labor demand in absolute terms, but they are not going to affect the relative labor demand, *i.e.* the structural transformation. Using the aggregate market clearing conditions in equation 2.9, the equation that defines the structural transformation is given by

$$\frac{l_i}{L} = \frac{\Omega_i C^{\epsilon_i} A_i^{\sigma-1}}{\sum_{j \in I} \Omega_j C^{\epsilon_j} A_j^{\sigma-1}}. \tag{2.11}$$

The labor share of sector i is affected by both income effects and substitution effects: as aggregate consumption rises one to one with aggregate income in our model economy, the labor share of sector i will rise if the income elasticity of demand of good i is higher relative to all other sectors and will fall if the elasticity is small relative to all other sectors. On the other hand, as labor productivity grows, the labor share of sector i will diminish relative to other sector with slower rates of labor productivity growth.

[9]Although we are normalizing the wages in this model economy, we leave them without normalization in the model exposition to illustrate that as long as labor is freely mobile, wages will not have an impact on the structural transformation.

2.4 Calibration

The parametrization involves estimating sectoral Engel curves and one price elasticity of substitution based on equilibrium conditions derived in the previous section. We use a panel for the U.S. and the European economies in our analysis to exploit variation across sectors and countries, and variation over time. This procedure assumes that preferences do not change systematically across countries during our sample period. Therefore, we can exploit the variation at this level of aggregation to pin down the Engel curves for the U.S. Next, we normalize the initial sectoral labor productivity to 1 and we calibrate the time-invariant CES weights to match perfectly the initial labor shares for each sector for the U.S. in 1970. With the calibrated model at hand, we can then feed in exogenous observable time paths of sectoral labor productivity levels to generate endogenously sectoral labor shares and aggregate labor productivity time paths.

2.4.1 Estimation of Engel Curves and the Price Elasticity of Substitution

Consider the model's prediction for the *absolute* labor demand of a sector i, as described by equation 2.10. One can define a system of labor demand for each sector i relative to manufacturing to derive the following system of *relative* labor demands

$$\frac{l_i}{l_{man}} = \frac{\Omega_i}{\Omega_{man}} C^{\epsilon_i - \epsilon_{man}} \left(\frac{A_i}{A_{man}} \right)^{\sigma - 1}.$$

Taking logs on both sides one gets

$$\log \left(\frac{l_i}{l_{man}} \right) = \log \left(\frac{\Omega_i}{\Omega_{man}} \right) + (\epsilon_i - \epsilon_{man}) \log C + (\sigma - 1) \log \left(\frac{A_i}{A_{man}} \right). \qquad (2.12)$$

From equation 2.12 we can derive the following econometric model to estimate the income and price elasticities

$$\log\left(\frac{l_{i,t}}{l_{man,t}}\right) = (1-\sigma)\log\left(\frac{A_{man,t}}{A_{i,t}}\right) + (\epsilon_i - \epsilon_{man})\log C_t + \zeta_i^c + \nu_{i,t}^c, \qquad (2.13)$$

where i denotes any sector – except manufacturing – in country c and time t. We control for fixed-effects ζ_i^c to capture time-invariant characteristics that can potentially influence our estimates. The error term of the econometric specification is $\nu_{man,t}^c$.

Estimating equation 2.13 imposes $i-1$ cross-equation restrictions for estimating one single price elasticity of substitution for the entire economy. Given the simplicity of our production function, we estimate equation 2.13 with prices predicted by the inverse of the productivity rather than with observed prices directly, because the econometric model derived from our theoretical framework is not suited for controlling for differences in technology parameters that do have a direct influence on prices.

Our identification strategy exploits within country-sector and time variation to identify the income and price elasticities. We use World KLEMS data, which is a panel disaggregated at the sector level with comparable information for the U.S., Austria, Belgium, France, Germany, the United Kingdom, Italy, Spain, and the Netherlands, from 1970 to 2009. Our measurement for the empirical counterparts of the model are as follows: Sectoral labor shares are measured by the ratio of labor hours hired in a sector to the total labor hours demanded in the economy. The sectoral labor productivity is measured with the real value added per hour worked. Finally, the aggregate consumption C is measured directly with income per capita measures since there are no savings in our model economy. Income per capita measures in real units adjusted by PPP to perform cross-country comparisons are not available in World KLEMS, so we used the Maddison Project as a source instead.

Table 2.3 presents the estimates for the price elasticity of substitution and the sectoral Engel curves relative to manufacturing. Our estimate of the price elasticity of substitution is 0.69, which is in line with the findings in the literature. The null hypothesis of a price elasticity of substitution equal to one is rejected at the 1

Table 2.3: Engel curves and price elasticity estimates

Sector	Parameter	Estimate
	$1-\sigma$	0.31^{***}
		(0.06)
agr	$\epsilon_{agr}-\epsilon_{man}$	-0.46^{***}
		(0.14)
trd	$\epsilon_{trd}-\epsilon_{man}$	0.50^{***}
		(0.08)
rst	$\epsilon_{rst}-\epsilon_{man}$	0.65^{***}
		(0.14)
trs	$\epsilon_{trs}-\epsilon_{man}$	0.55^{***}
		(0.09)
com	$\epsilon_{com}-\epsilon_{man}$	0.63^{***}
		(0.11)
fin	$\epsilon_{fin}-\epsilon_{man}$	0.71^{***}
		(0.12)
res	$\epsilon_{res}-\epsilon_{man}$	1.17^{***}
		(0.17)
bss	$\epsilon_{bss}-\epsilon_{man}$	1.76^{***}
		(0.11)
gov	$\epsilon_{gov}-\epsilon_{man}$	0.27^{***}
		(0.10)
edu	$\epsilon_{edu}-\epsilon_{man}$	0.57^{***}
		(0.10)
hlt	$\epsilon_{hlt}-\epsilon_{man}$	0.93^{***}
		(0.14)
per	$\epsilon_{per}-\epsilon_{man}$	0.72^{***}
		(0.16)
Number of observations	360	
Fixed effects	Yes	

Estimation based on World KLEMS data for Austria, Belgium, France, Germany, Italy, the Netherlands, Spain, the United Kingdom, and the United States. Clustered standard errors at the country level in parenthesis. * $p<0.10$, ** $p<0.05$, *** $p<0.01$.

per cent level, in favor a σ below one. Our estimate of the price elasticity of substitution reflects the presence of a Baumol-cost disease, in line with the analytical descriptions of Baumol (1967) and Ngai and Pissarides (2007). This means that in our framework the economy is converging to services, as in the traditional literature of structural transformation, and also that within services the economy is converging toward the least productive sectors. This is the supply side explanation of the structural transformation.

To account for the demand side, Table 2.3 illustrates the Engel curves for each sector relative to manufacturing. The null hypothesis is that the Engel curve for a given sector i is the same as the manufacturing Engel curve. This hypothesis is rejected at the 1 per cent level of significance for each sector in the economy. Consistent with the development literature, the estimate for the Engel curve in agriculture illustrates that as long as the household grows richer, the resources devoted for the consumption of agriculture grow less than proportional relative to manufacturing, whereas for all the services in the economy the consumption grows more than proportional relative to manufacturing. In addition, the estimates of the Engel curve estimate vary significantly across services. For instance, whereas the difference in the income elasticity for government relative to manufacturing is of 0.27, for real estate and business services this difference is above one.

2.4.2 Targeting the Initial Employment Shares in the U.S.

We calibrate the model by targeting the initial labor shares in 1970 for each sector in the U.S. economy. For this purpose, we normalize the initial productivity levels A_i to one in each sector. As a consequence of this normalization, the aggregate productivity is normalized to one as well, and therefore $Y/L = A = 1$. Since in our model economy the entirety of income per capita is devoted to consumption, it follows that $C = 1$ for 1970. From equation 2.11, the normalization implies that the labor shares for the initial period of the calibration are given by

Table 2.4: Parameter values

	Value	Target/Comment
Parameters		
σ	0.69	Price elasticity estimation (Table 2.3).
ϵ_{agr}	0.53	Estimate for Engel curve for `agr` (Table 2.3).
ϵ_{man}	1	Homothetic preferences for manufacturing.
ϵ_{trd}	1.50	Estimate for Engel curve for `trd` (Table 2.3).
ϵ_{rst}	1.65	Estimate for Engel curve for `rst` (Table 2.3).
ϵ_{trs}	1.55	Estimate for Engel curve for `trs` (Table 2.3).
ϵ_{com}	1.63	Estimate for Engel curve for `com` (Table 2.3).
ϵ_{fin}	1.71	Estimate for Engel curve for `fin` (Table 2.3).
ϵ_{res}	2.17	Estimate for Engel curve for `res` (Table 2.3).
ϵ_{bss}	2.75	Estimate for Engel curve for `bss` (Table 2.3).
ϵ_{gov}	1.27	Estimate for Engel curve for `gov` (Table 2.3).
ϵ_{edu}	1.57	Estimate for Engel curve for `edu` (Table 2.3).
ϵ_{hlt}	1.93	Estimate for Engel curve for `hlt` (Table 2.3).
ϵ_{per}	1.73	Estimate for Engel curve for `per` (Table 2.3).
Ω_{agr}	0.06	Labor share of sector `agr` in 1970 for the U.S.
Ω_{man}	0.30	Labor share of sector `man` in 1970 for the U.S.
Ω_{trd}	0.14	Labor share of sector `trd` in 1970 for the U.S.
Ω_{rst}	0.04	Labor share of sector `rst` in 1970 for the U.S.
Ω_{trs}	0.03	Labor share of sector `trs` in 1970 for the U.S.
Ω_{com}	0.03	Labor share of sector `com` in 1970 for the U.S.
Ω_{fin}	0.03	Labor share of sector `fin` in 1970 for the U.S.
Ω_{res}	0.01	Labor share of sector `res` in 1970 for the U.S.
Ω_{bss}	0.06	Labor share of sector `bss` in 1970 for the U.S.
Ω_{gov}	0.07	Labor share of sector `gov` in 1970 for the U.S.
Ω_{edu}	0.07	Labor share of sector `edu` in 1970 for the U.S.
Ω_{hlt}	0.11	Labor share of sector `hlt` in 1970 for the U.S.
Ω_{per}	0.04	Labor share of sector `per` in 1970 for the U.S.
Time Paths		
$\{A_{i,t}\}$	$\{\cdot\}$	$A_{i,t+1} = A_{i,t}(1 + \gamma_{A_{i,t}})$, where $\gamma_{A_{i,t}}$ is the growth rate of sectoral real value added per hour. $A_{i,t=1970} = 1$.
$\{C_t\}$	$\{\cdot\}$	$C_{t+1} = C_t(1 + \gamma_{C_t})$, where γ_{C_t} is the growth rate of real GDP per capita.

$$\frac{l_i}{L} = \frac{\Omega_i}{\sum_{j \in I} \Omega_j}.$$

Since $\sum_{j \in I} \Omega_j = 1$, the initial labor shares for each sector i are given by Ω_i. The initial labor shares values for the U.S. in 1970 are sufficient to account for the parameterization of each Ω_i so the model and the data match for the first period, by construction. Then, we compute the sectoral labor productivity time paths $\{A_{i,t}\}_{t=1970}^{2009}$ with the observed growth rates of real value added per worker, and the aggregate consumption time path $\{C_t\}_{t=1970}^{2009}$ with aggregate labor productivity growth rates, measured by the real income per capita growth. Finally, we feed these time paths in our model to derive predictions for the evolution of the employment labor shares across sectors as described by equation 2.11. Table 2.4 summarizes the parametrization of our model.

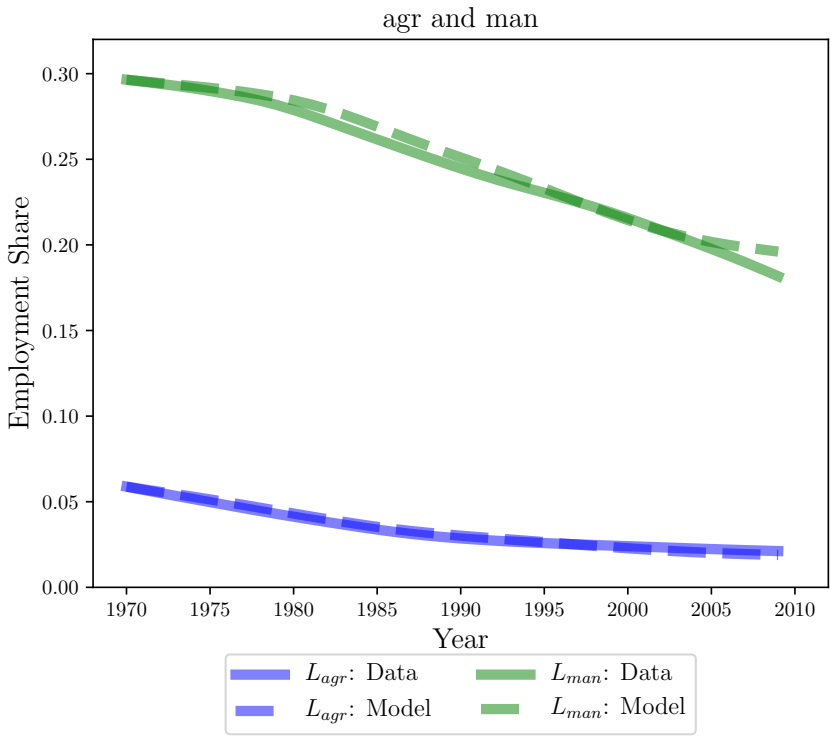
2.5 Quantitative Analysis

There are three sets of predictions that we consider as tests of whether our theory can successfully account for the structural transformation. First, the labor-share time paths generated by our model for the U.S. economy should be roughly close to their empirical counterparts in the data. Second, after recovering the initial productivity levels for each of the European economies, the model should be capable of generating labor shares roughly close for most sectors in the European countries. Third, the predicted aggregate labor productivity – namely the sum of sectoral labor productivities weighted by their participation in the labor force – should reproduce fairly close the relative aggregate labor productivity between the U.S. and Europe displayed in Figure 2.1.

2.5.1 Model's Prediction I: U.S. Structural Transformation

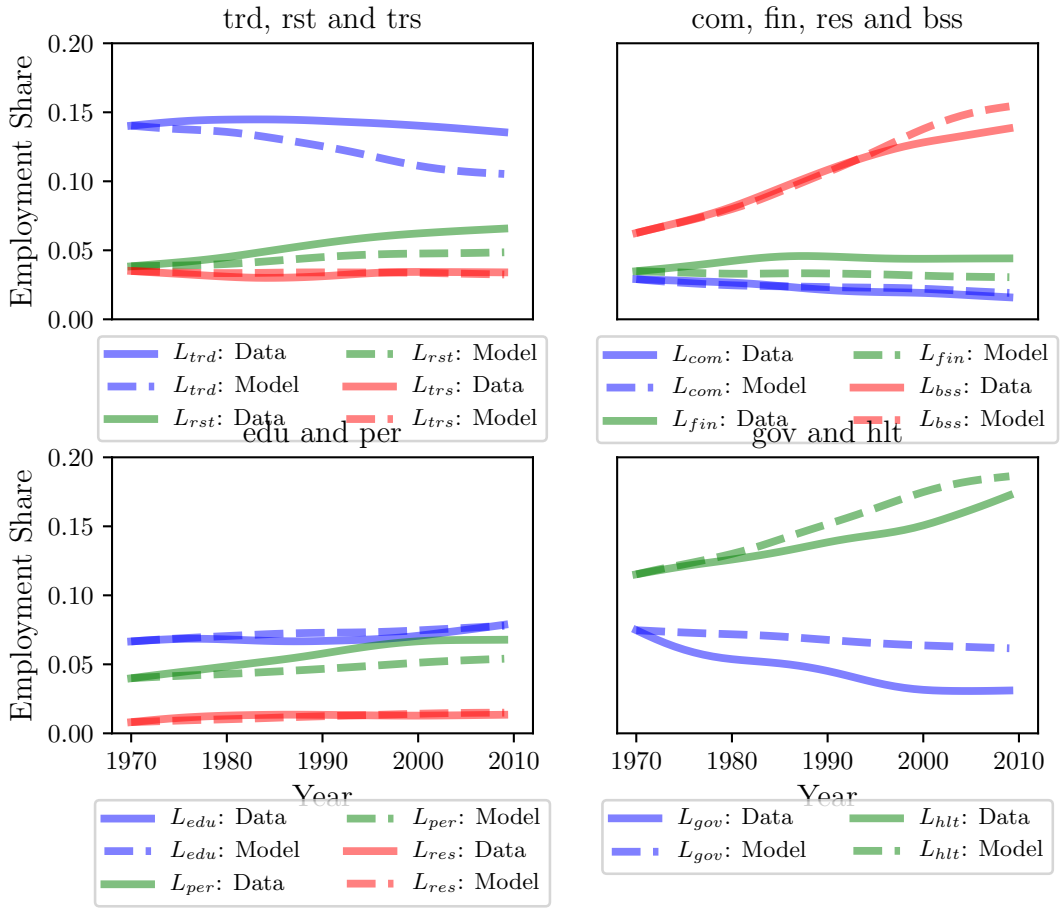
Figure 2.3 compares the predicted labor shares of our model to the U.S. data for agriculture and manufacturing. The model does a remarkably good job predicting the observed labor share paths for these two sectors during the sample period. For agriculture, the model predicts almost perfectly the decline in the labor share. Nonetheless, for 1970 most of the labor in the U.S. economy had already migrated out of agricultural activities. The model also does a good job predicting the observed de-industrialization of the U.S. economy since 1970: Whereas the observed decline of the manufacturing share of employment was from about 30 per cent in 1970 to levels

Figure 2.3: Structural transformation in the U.S. - Agriculture and manufacturing



Time paths of employment shares between 1970 and 2009, as observed in the data and predicted by our model.

Figure 2.4: Structural transformation in the U.S. - Services

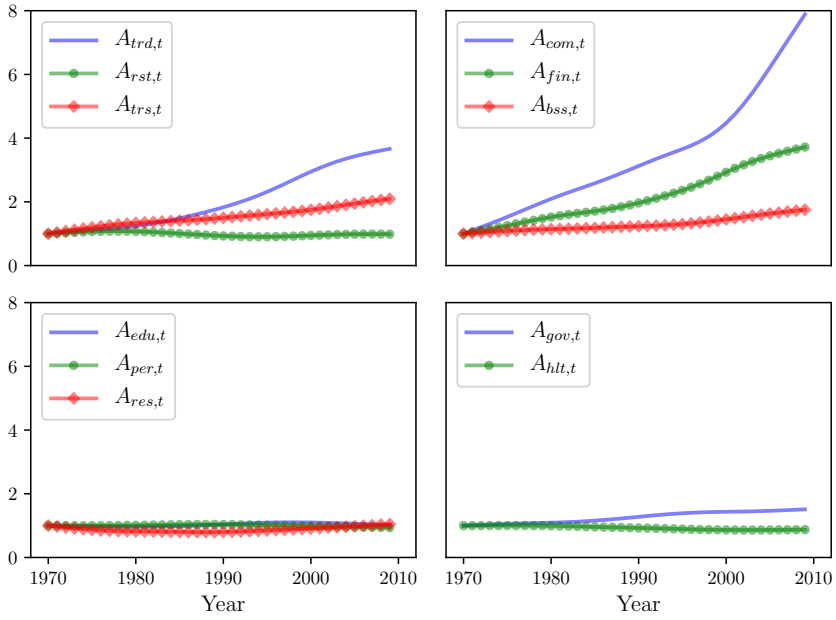


Time paths of employment shares between 1970 and 2009, as observed in the data and predicted by our model.

short of 20 per cent in 2009, the predicted decline in the manufacturing employment share is down to a level of about 21 per cent in 2009.

Figure 2.4 compares the predicted labor shares for the different services in the U.S. economy. The model does follow the labor share paths fairly close for almost every sector, including the steep rise in business services as shown in the upper right panel of Figure 2.4. The two exceptions are wholesale and retail trade and government. The upper left panel of Figure 2.4 illustrates that for wholesale and retail trade, the

Figure 2.5: Sectoral labor productivity in the U.S. - Service sector



Labor productivity is measured as the real value added per hour worked. Initial productivity levels are normalized to 1.

employment share has remained at a level close to 14 per cent during the sample period, with an observed decline of only half of a percentage point after 1990. The model, however, predicts a decline in the labor share of this sector down to a level of 10 per cent. For government (see the lower right panel in Figure 2.4) the model underpredicts its labor share's decline. Whereas the government labor share falls from above 7 per cent in 1970 to about 3 per cent in 2010, our model predicts that this share will decrease only by less than 2 per cent for the same period.

To shed more light on the model's predictions for the structural transformation within services, Figure 2.5 plots the sectoral labor productivity time paths for each service in the U.S. for the period 1970-2009. Communications, wholesale and retail trade, financial services, business services, government, and, to a lesser degree, transportation are the sectors with superior performance in labor productivity. The productivity in communications has increased by a factor of 8 from 1970 to 2009,

while the productivity has multiplied its 1970 base more than 3.5 times in wholesale and retail trade, and financial services. Transportation, business services and government also have multiplied their productivity base by a factor of 2.1, 1.7, and 1.5 respectively. The rest of the service sectors had experienced virtually no growth in their labor productivity. That is true even for sectors such as health services, whose participation in the labor force exceeded 18 per cent in 2009.

Can the evidence presented in Figure 2.5 explain why the model is not following closely the labor shares in wholesale and retail trade and government? We believe that, in spite of the simplicity of our model, the answer is yes. There are two drivers of the structural transformation in our model economy: Engel curves and heterogenous labor productivity growth rates through the price elasticity of substitution. We already showed that the income elasticity for each sector belonging to services is statistically superior to the manufacturing Engel curve. Are the income elasticities in services statistically different from each other? The answer depends on the sector. The three sectors displayed in the upper left panel of Figure 2.4 have Engel curves that are not statistically different from each other, but they are statistically lower than the Engel curves for real estate or business services. Therefore, the differences in our model predictions between wholesale and retail trade, restaurants and hotels, and transportation are to be found in the labor productivity differences. The upper left panel of Figure 2.5 shows that wholesale and retail trade has the strongest productivity growth among these three services, and therefore, according to our model, this sector should reduce its participation in the labor force. This prediction is in contrast with the observed labor shares, suggesting that in the U.S. it is not necessarily true that the labor productivity growth is shrinking the employment participation in wholesale and retail trade.

On the other hand, government does have an Engel curve significantly lower than the rest of the services with the exception of wholesale and retail trade, and it is experiencing positive productivity growth. These two forces imply in our model a decrease in the government's employment share, but both mechanisms are not sufficient to address the deployment of the labor force out of government that are evident in the U.S. data. Nevertheless, with important caveats for wholesale and

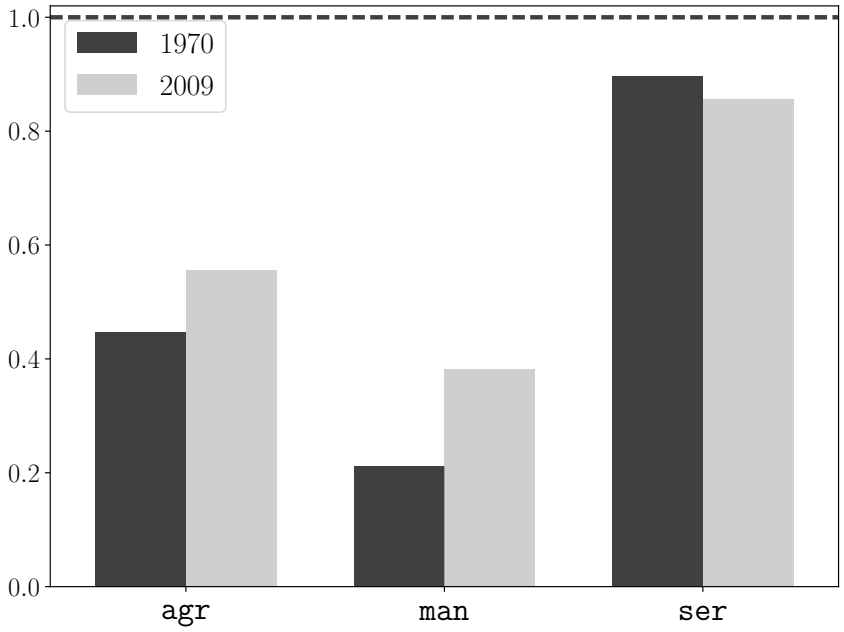
retail trade and for government, we consider that our model successfully accounts for the structural transformation in the U.S.

2.5.2 Model's Prediction II: Structural Transformation in Europe

Following Duarte and Restuccia (2010), we use our model to measure the initial productivity *levels* in Europe *vis-à-vis* the U.S. This in an important accounting step to overcome the lack of sectoral PPP-adjusted value added data. We proceed as follows: First, we use the calibrated parameters summarized in Table 2.4 to recover the productivity levels for each sector and for each European country consistent with the normalization of productivity levels in the U.S. *and* with the income level of each European country relative to the GDP per capita in the U.S. Since the U.S. income level is equal to 1 in the first period of our model (corresponding to 1970 in the data), the relative income per capita is simply the ratio of GDP per capita of each European country to the U.S. in 1970. We use the Maddison Project's GDP per capita measures since they are adjusted by PPP's, thus PPP-adjusting the initial sectoral productivity levels that our model is recovering. Then, we compute the labor productivity and income time paths with the observed growth rates of sectoral real value added per hour and real income per capita respectively, just as we did for the U.S. in the previous section. Last, with the recovered PPP-adjusted time paths, we compute the model's predictions and compare the structural transformation predicted by our model to the European data. This procedure delivers time paths that are comparable across countries, without the risk of mismeasurement due to not ideal PPP adjustments at the two digits sectoral level.

Measurement of Sectoral Labor Productivity in Europe. Figure 2.6 plots the average productivity levels measured using our calibrated model in agriculture, manufacturing, and services for Europe relative to the U.S. for 1970 and 2009 (the first and last sample periods respectively). These productivity levels are an outcome from our model needed to compute comparable productivity levels in absence of PPP-adjusted sectoral output data. First, the agricultural productivity

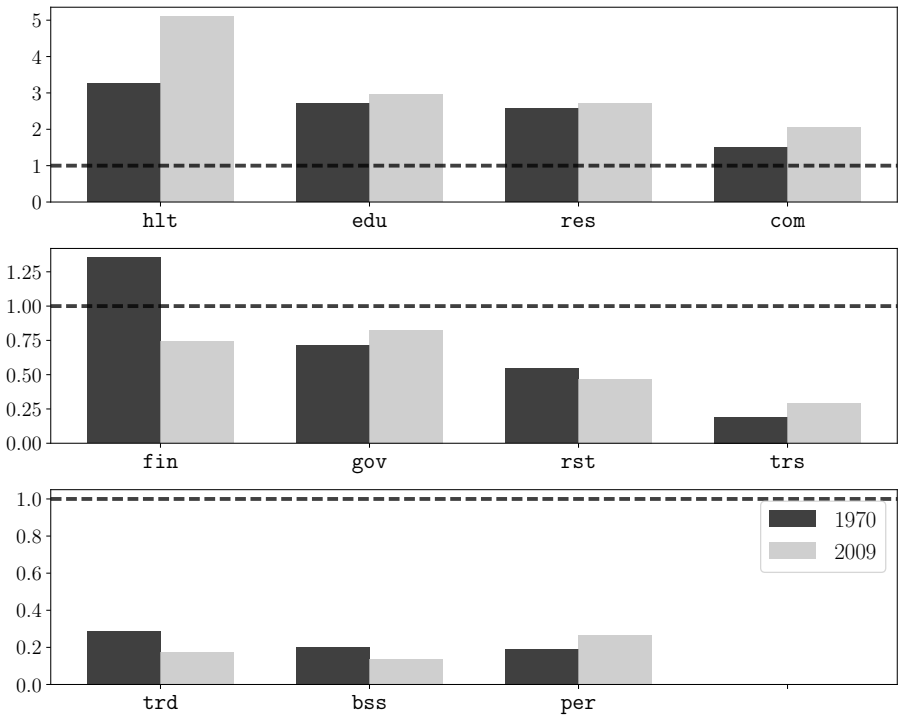
Figure 2.6: Relative labor productivity - Agriculture, manufacturing, and services



Recovered sectoral labor productivity levels in 1970 and 2009 for Europe relative to the sectoral U.S. labor productivity level. Europe stands for the average of the eight European countries' sectoral productivity levels, weighted by their national GDP.

levels recovered from our model illustrate that in 1970 the average agricultural productivity level in Europe was 45 per cent of the U.S. productivity. This gap closed partially during our sample period. By 2009, the European agricultural productivity level was 55 per cent of the U.S. agricultural productivity level, reflecting a reduction in the gap of about 20 per cent. Second, during a sample period is also evident a stronger process of convergence in the manufacturing sector. Whereas the European manufacturing productivity level was about 21 per cent of the U.S. manufacturing labor productivity, for 2009 the European manufacturing productivity was 38 per cent of the U.S. level, which represents an increase of 80 per cent. We believe that these numbers are relatively low compared to the evidence documented by Lewis (2005) for some subset of manufacturing industries, such as the automobile industry,

Figure 2.7: Relative labor productivity - Services



Recovered sectoral labor productivity levels in 1970 and 2009 for Europe relative to the sectoral U.S. labor productivity level. Sectors within services. Europe stands for the average of the eight European countries' sectoral productivity levels, weighted by their national GDP.

but the catch up in manufacturing is of similar orders of magnitude when compared to the findings of Duarte and Restuccia (2010). Third and last, the labor productivity gap in services – our object of interest – was smaller in 1970 compared to 2009. Whereas the average level for the labor productivity in the European services was 90 per cent of the U.S. services' labor productivity, for 2009 the European productivity in services represented about 86 per cent of the productivity in U.S. services. This widening is the main reason behind the recent divergence between Europe and the U.S. due to the ongoing growth of services' weight in the economy during the late stages of development.

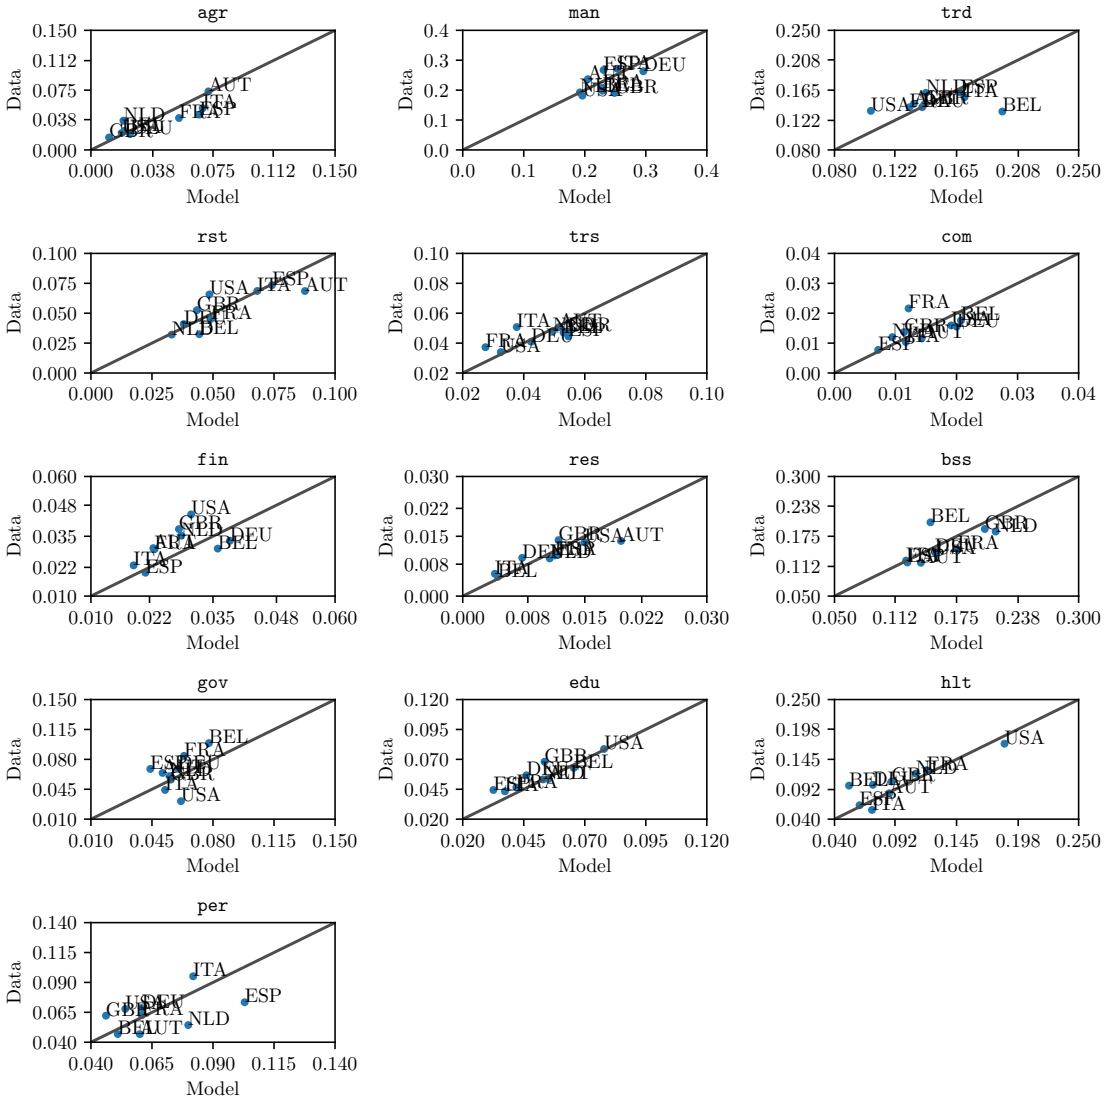
Figure 2.7 plots the initial and final productivity levels during our sample period

for each of the sectors within services in Europe relative to the United States. To the best of our knowledge, there is no independent evidence on labor productivity levels for all these 11 sector in Europe and the U.S. to compare directly the implied labor productivity levels of our model. Europe as a whole did lose ground compared to the U.S. in terms of productivity in services, but one should not infer from Figure 2.6 that *all* services were less productive in Europe compared to the U.S. Our model suggests that there are five sectors where Europe had higher productivity levels than the U.S. in 1970: Communications, financial services, real estate, education (`edu`), and health services, and with the important exception of financial services, the U.S. fell behind even further by the end of our sample period in these services.

Figure 2.7 also shows that the lower European productivity levels in services in 1970 are due to wholesale and retail trade, transportation, restaurants and hotels, business services, government, and personal services. Moreover, the gap in these sectors opened even wider by 2009 in wholesale and retail trade, restaurants and hotels, business services, and also in financial services, where the U.S. did close the productivity gap and later surpassed Europe by 2009. For instance, the productivity levels relative to the U.S. in wholesale and retail trade went from 29 per cent in 1970 down to 17 per cent in 2009. For business services the fall was from 20 per cent in 1970 down to 14 per cent in 2009. For financial services, the European productivity went from 35 per cent *above* of the U.S. level down to 74 per cent in 2009. Figure 2.7 illustrates the importance of opening services into comparable sectors between Europe and the U.S. in order to address why Europe has been falling behind with respect to the U.S. during advanced stages of development, where the service sector dominates the labor participation in the economy.

Structural Transformation Within Services in Europe. In order to address whether our model is successful in explaining the structural transformation in Europe, Figure 2.8 plots a scatter between the observed labor share for each sector in 2009 and the prediction of our model for the same period. It also plots a solid line that represents the 45 degree line starting at the origin of the y and x-axis. The closer the pair between the observed labor share (y-axis) and our model's predic-

Figure 2.8: Structural transformation in the U.S. and Europe



Sectoral employment shares in 2009, in the U.S. and the eight European countries. Each country's coordinates correspond to observed and model predicted levels.

tion (x-axis) to the 45 degree line, the more accurate our model is in capturing the

process of structural transformation.[10] Figure 2.8 illustrates that the model successfully generates sectoral employment shares roughly consistent with the data, with a few exceptions in wholesale and retail trade for the U.S. (as previously documented) and Belgium, and in personal services for Spain and the Netherlands. Nevertheless, our model succeeds overall in explaining the process of structural transformation in Europe.

2.5.3 Model's Prediction III: Aggregate Labor Productivity in Europe *vis-à-vis* the U.S.

Can our model generate the main motivating fact presented in Figure 2.1? If we consider the aggregate labor productivity level to be the weighted average of the sectoral labor productivity levles, where the weights are nothing but the labor shares of employment in each sector, *i.e.* the structural transformation, then our model's predictions can be compared directly to the evidence on aggregate labor productivity in Europe *vis-à-vis* the U.S. presented in Figure 2.1.[11] One can address the capacity of the model in generating the labor productivity ratios by using our predicted labor shares for each sector to weight the sectoral productivity levels in order to generate aggregate labor productivity time paths for each country.

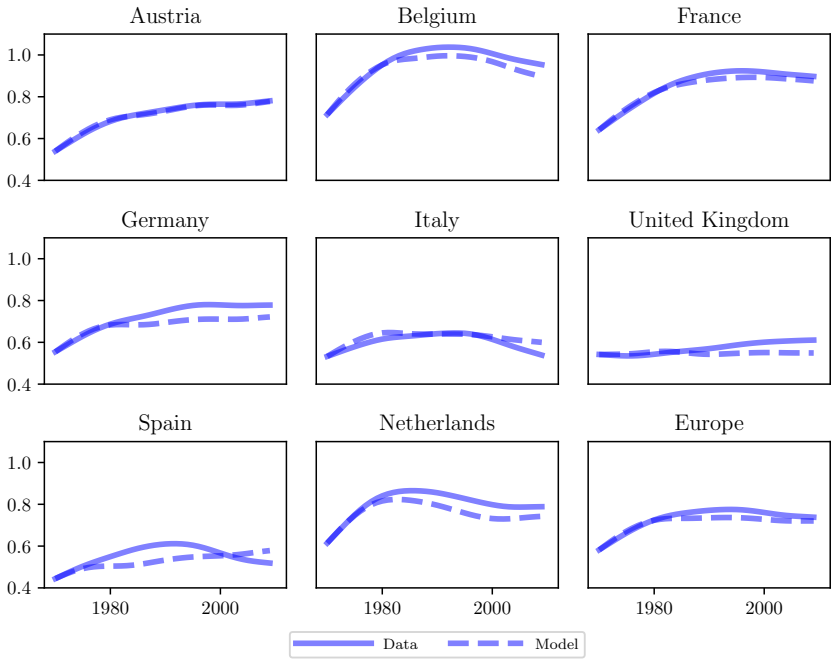
Figure 2.9 compares the model's prediction to the data for the aggregate labor productivity in each European country relative to the U.S. and for the European aggregate productivity relative to the U.S. as well.[12] After matching by construction the initial observations, the model does follow very close the observed gaps in aggregate labor productivity between Europe and the U.S., regardless on whether the

[10]Unlike the employment share in manufacturing, there are no well-defined hump-shaped patterns in the structural transformation in services. For this reason we consider that the prediction for the last observation in the sample is sufficient to assess the model's capacity to generate time paths consistent with the European structural transformation.

[11]Recall that we discipline the initial labor productivity in Europe with the relative, PPP-adjusted, income per capita measures, matching the model and the data by construction for the first period.

[12]The aggregate productivity in Europe is computed as the average of the eight European countries' aggregate productivity, weighted by their national GDP.

Figure 2.9: Relative aggregate labor productivity



Aggregate labor productivity for the European countries relative to the U.S. Europe's aggregate productivity is the average of the eight European countries' aggregate productivity, weighted by their national GDP. Model vs. Data. The model's aggregate labor productivity is the weighted average of sectoral labor productivity, where the weights are the model's predicted labor shares for each sector.

country's convergence stopped, as in France or Germany, or whether the country is falling behind the U.S., as in Belgium or the Netherlands.

In summary, we judged quantitatively the model's performance in three dimensions: i) The U.S. structural transformation, ii) the European structural transformation, and iii) the aggregate labor productivity in Europe relative to the U.S. Our exercises show that our theoretical framework is successful in accounting the participation of employment in agriculture, manufacturing, and several services in the U.S. and Europe, and it also accounts for the aggregate differences in output per hour worked between these two regions, and for each country individually. These result are reassuring that our theoretical framework is quantitatively valid, and supports

the credibility of the counterfactual experiments we expose hereafter.

2.6 Counterfactual Experiments

After illustrating the quantitative success of the theory, we proceed to use our parametrized model economy to perform a set of counterfactual experiments in order to understand the role of services sub-sectors in aggregate productivity. Our aim is to identify which sectors are largely responsible for the slowdown in European labor productivity during the last two decades relative to the United States.

2.6.1 Europe keeping the Pace with the U.S.

Our first counterfactual experiment asks what would have happened with the aggregate labor productivity in Europe had it experienced the observed sectoral productivity growth in the U.S. from 1970 to 2009, in a specific sector or group of sectors. We ask this question for each sector individually, for services as an entire sector, and for all the sectors simultaneously. More specifically, we use our model to predict the structural transformation in Europe with the observed U.S. labor productivity growth rate in each sector and compute the counterfactual aggregate productivity. Then, we compare this aggregate productivity with our benchmark prediction from Figure 2.9 to address the differences between our counterfactual scenario and the benchmark prediction for the aggregate productivity.[13] This experiment seeks to answer which sectors are responsible for the relative aggregate productivity slow down.

[13] As Figure 2.9 shows, our model is successful in predicting the dynamics for the aggregate labor productivity. One can perform this exercise by comparing the counterfactual prediction directly to the observed aggregate productivity level. We decided to compare the counterfactual scenarios to our benchmark predictions because our model successfully accounts for the aggregate labor productivity and because by comparing models' predictions we can address with certainty that the differences arise solely due to the numerical experiment. However, if one decides to compare directly to the actual data the conclusions would not change dramatically.

Table 2.5: Europe keeping the U.S. pace

	(1) 1970–2009	(2) 1990–2009
$\gamma_i = \gamma_i^{USA}$		
agr	0.4	-0.3
man	-8.9	-0.3
trd	5.1	3.2
rst	-0.5	0.0
trs	-0.2	0.2
com	-0.9	-2.9
fin	3.8	0.4
res	0.4	0.8
bss	3.0	2.4
gov	-0.1	-0.5
edu	-0.8	0.1
hlt	-5.9	-2.7
per	-0.9	-0.4
$\gamma_i = \gamma_{i,i\in\text{services}}^{USA}$	3.4	0.7
$\gamma_i = \gamma_{i,\forall i}^{USA}$	-5.3	0.2

Europe growing at the pace of U.S. in the indicated sector, for the periods 1970–2009 or 1990–2009. Percentage change of the 2009 aggregate labor productivity level. Benchmark prediction vs. counterfactual.

Table 2.5 illustrates our findings when we feed the labor productivity growth rates from 1970 to 2009 (our entire sample period) and from 1990 to 2009 (the period where Europe lagged behind). The top panel of Table 2.5 shows the results of this exercise when Europe counterfactually experiences the observed labor productivity growth rate in the U.S., in order to assess changes in aggregate labor productivity as a consequence of changes in the productivity of a single sector. Each row of the top panel represents one of the 13 sectors in our model economy.

Column (1) of Table 2.5 shows that Europe would have had an increase in aggregate labor productivity of 0.4 per cent had it experienced the U.S. productivity growth in agriculture. These modest results are not surprising. Both Europe and the U.S. are economies at advanced stages of development, with low levels for the size of agriculture in the economy even in 1970, and in steady decline since then. On the other hand, had the European countries experienced the U.S. labor productivity growth in manufacturing during our sample period, Europe as a whole would have had a lower aggregate productivity. Manufacturing is not responsible for the European underperformance *vis-á-vis* the U.S. On the contrary, it helped Europe in its path towards convergence during the first half of our sample period.

With regards to services, our counterfactual experiment suggests that the slowdown in the aggregate labor productivity comes mainly from three sectors: Wholesale and retail trade, financial services, and business services. It also suggests that Europeans experienced significantly higher productivity gains in health services.[14] During the sample period, wholesale and retail trade alone would have been responsible for an European aggregate labor productivity 5.1 per cent higher than our benchmark prediction in 2009. Financial services also would have helped to reduce the labor productivity gap had the European countries experienced the same labor productivity growth observed in this sector for the U.S. Europe as a whole would have had a labor productivity level 3.8 per cent higher than our benchmark prediction. The labor productivity would have also been higher for the European countries if they had had the U.S. labor productivity growth in business services. Our results also illustrate that Europe would have had lower aggregate productivity had it had the

[14]For the rest of the sectors the results are not large.

U.S. labor productivity growth observed in health services. It is well known that the U.S. is the advanced economy with the most expensive health sector, and our simple model shows that part of these higher costs are captured by its relatively low labor productivity in this sector.[15]

The middle and lower panels of Table 2.5 show what would have happened if Europe had experienced the productivity growth rates observed in the U.S. in all services and all sectors simultaneously, respectively. Europe would have experienced some convergence during this period if their services had experienced the U.S. labor productivity growth; the aggregate labor productivity would have been 3.4 per cent higher than our benchmark prediction for 2009. However, if *all* sectors had grown like the U.S., the gains obtained in services would have been out-weighted by a poorer performance in manufacturing, yielding an overall loss of the aggregate labor productivity of 5.3 per cent compared to our benchmark prediction in 2009.

It has been established that the aggregate productivity in Europe was converging to the U.S. before 1990, while after this year a process of either slowdown or falling behind started, depending on the country that one is considering. Our second counterfactual experiment asks what would have happened if Europe had continued with the U.S. labor productivity growth rates after 1990, which is the period when the process of convergence came to a halt. We followed the same set of exercises from the previous section, with the only difference that the U.S. growth rates that are counterfactually fed start in 1990 rather than in 1970.

Column (2) of Table 2.5 shows the results of the numerical experiments for the period between 1990 and 2009 by comparing the benchmark prediction to the counterfactual aggregate labor productivity in 2009. Whereas the results for agriculture are still negligible, the sharp drop in the aggregate labor productivity with the U.S. manufacturing labor productivity for the period 1970-2009 virtually vanishes when we feed the productivity growth rates only since 1990. This confirms our previous

[15]Nevertheless, the question of productivity in health services is one of great difficulty. Labor productivity is measured as the real value added per worker, but without a proper adjustment for quality it is difficult to address whether more health services per hour reflect more productivity in the health sector. Still, our model captures reasonably well the idea that the U.S. provides health services that are much more expensive compared to their European counterparts.

finding: Manufacturing was responsible for the catch-up observed during the 1970's and 1980's. After these years, the productivity growth in manufacturing is not as critical as before to understand the aggregate labor productivity, mainly because the weight of manufacturing has fallen due to the ongoing process of structural transformation. Wholesale and retail trade and business services continue to be of great importance to account for the European slowdown that took place after 1990. The aggregate labor productivity would have been significantly higher in Europe had it experienced the U.S. labor productivity growth in these sectors. On the other hand, financial services are no longer critical to account for the slowdown, in contrast with the counterfactual for the whole sample period, suggesting that the U.S. financial sector had a stronger labor productivity growth than the European one mainly before 1990. The results for health services are in the same direction compared to the entire sample period, but the order of magnitude of the result is about half of what it was for the 1970-2009 period, although it still represents a large distance between the benchmark and the counterfactual aggregate productivity. In addition, for the period between 1990 and 2009 a new sector emerges in which Europe appears to have overperformed the U.S. in terms of labor productivity growth: Communications.

The middle and lower panels of Table 2.5 illustrate that for the period 1990-2009, the European countries would have been modestly more productive had they had the U.S. labor productivity growth observed in the service sector. In addition, they would have been virtually the same had they had the labor productivity growth in each sector in the economy since 1990.

2.6.2 European Sectors Catching Up with the U.S. Productivity Levels in 2009

After identifying the sectors largely responsible for the European slowdown, our second set of numerical experiments ask how much the aggregate labor productivity would have grown if either wholesale and retail trade, financial services, or business services had experienced the productivity growth needed to fully catch up with the U.S. labor productivity *level* in each sector by 2009. We assume that this convergence

Table 2.6: Europe catching up with the U.S.

	Full catch up in 2009
Counterfactual:	
γ_i s.t. $A_i = A_i^{USA}$	
`trd`	25.8
`bss`	17.1
`fin`	1.5

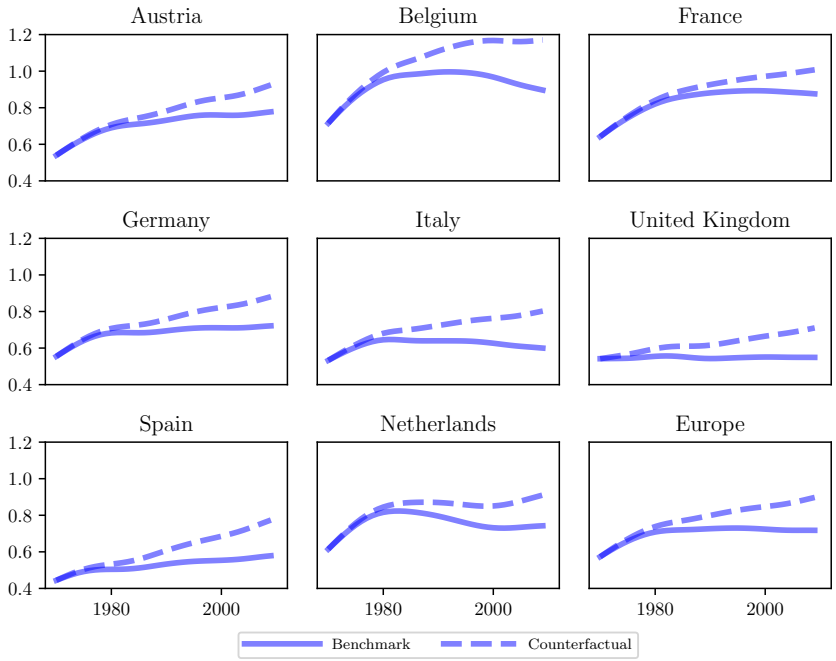
Europe catching up with the U.S. sectoral productivity *level* in 2009, in the sector indicated. Implied (annualized) growth rates under *full* catch-up in whole sale and retail trade, business services, and financial services. Percentage change of the 2009 aggregate labor productivity level (benchmark prediction vs. counterfactual).

takes place only in one sector at a time to compute the annualized growth rate consistent with the catch up to the U.S. labor productivity in the sector in question, while keeping the observed growth rates for the rest of the sectors.

Table 2.6 shows the implied change in aggregate productivity when each of these three sectors mentioned before converges to the U.S. labor productivity level in 2009.[16] Had Europe converged to the U.S. productivity level in 2009 in wholesale and retail trade or in business services, the aggregate productivity gains would have been substantial. Europe as a whole would have had an aggregate productivity level 25.8 per cent higher had it converged in wholesale and retail trade, and of 17.1 per cent had the labor productivity level converged in business services. These two sectors alone are largely responsible for the European slowdown relative to the U.S. Table 2.6 also shows that financial services is not a critical source of slowdown between Europe and the U.S. Had Europe experienced a full catch up in the labor productivity of financial services relative to the U.S. 2009 level, the aggregate productivity level would have been only 1.5 per cent higher compared to our 2009 benchmark prediction.

[16]Our model is suited to perform this numerical experiment for any sector in the economy, but for the sake of space, we decide to show only the three sectors that we identify as largely responsible for the European slowdown during the period 1970-2009.

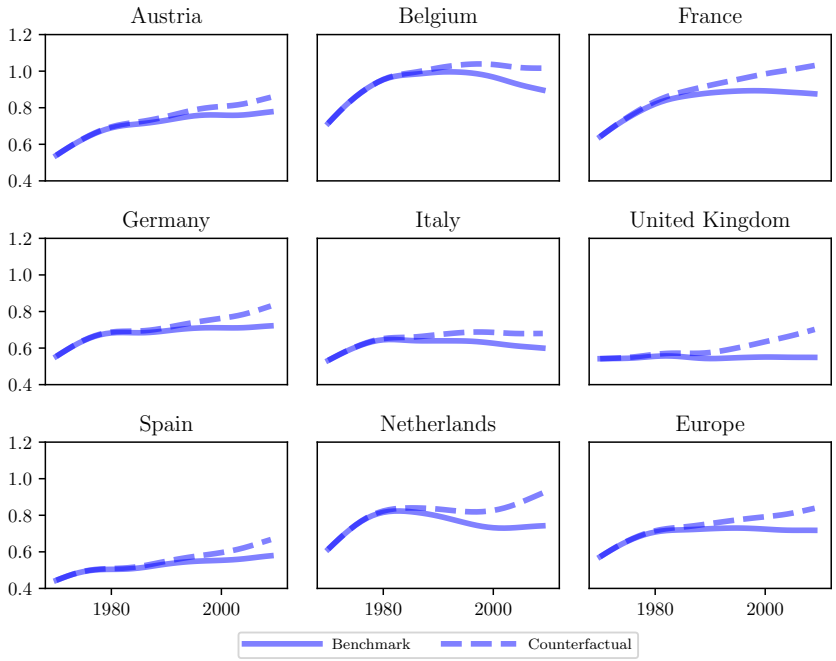
Figure 2.10: Relative aggregate labor productivity – Full catch up in wholesale and retail trade



Aggregate labor productivity in Europe *vis-á-vis* the U.S. under full catch up in the labor productivity of wholesale and retail trade sector. Benchmark prediction vs. counterfactual.

Figure 2.10 illustrates the effect of a full catch up wholesale and retail trade on the aggregate labor productivity over time, from 1970 to 2009. Had the European countries converged to the 2009 labor productivity levels in wholesale and retail trade, they would have continued their path toward convergence after 1990, with a mild deceleration in a few countries. Figure 2.10 shows that every single country in Europe would have improved its position relative to the U.S. without exception. Moreover, Austria and France would have virtually closed the labor productivity gap with the U.S., and Belgium would have surpassed the U.S. aggregate labor productivity level by 2009. The rest of the countries would have not still closed the gap, but they would have not fallen behind either, had they closed the gap in wholesale and retail trade. Europe as a whole would have closed 80 per cent of

Figure 2.11: Relative aggregate labor productivity – Full catch up in business services

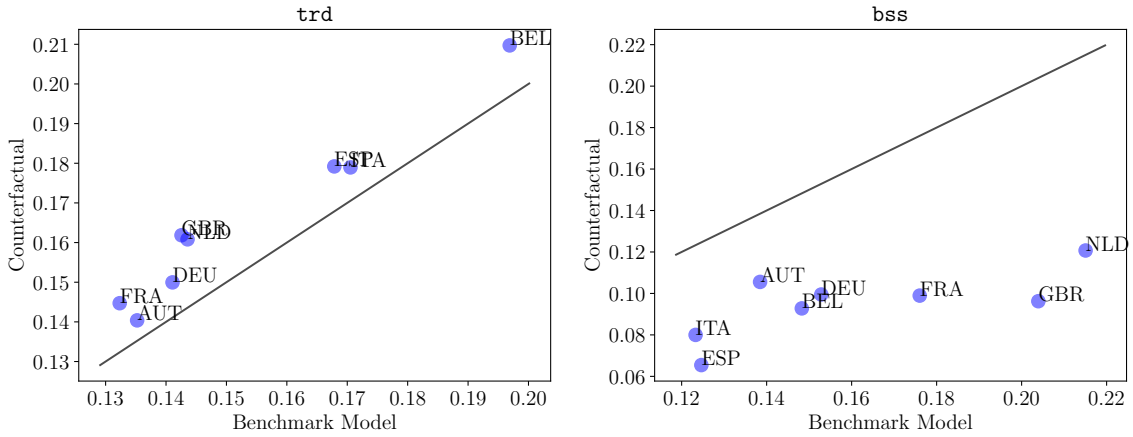


Aggregate labor productivity in Europe *vis-á-vis* the U.S. under full catch up for the labor productivity in the business services sector. Benchmark prediction vs. counterfactual.

the gap in labor productivity, if they have closed the labor productivity gap in this specific sector alone with respect to the U.S. As Lewis (2005, p. 34) puts it, "In the United States, wholesalers [...] began to consolidate their warehouses and improve the productivity of the operations in those warehouses. This change was the largest single contribution to the productivity acceleration in the U.S. economy in the late 1990's [...] not the efforts of Microsoft and Silicon Valley".

Similarly, Figure 2.11 illustrates the effect of a full catch up in business services on the aggregate labor productivity time path between 1970 and 2009. The results are qualitatively similar to our previous numerical experiment illustrated in Figure 2.10, but the magnitude of the effect from catching up in business services is much smaller compared to a full catch up in wholesale and retail trade. Still, if Europe

Figure 2.12: Labor shares under full catch up



Predicted labor shares in 2009 for whole sale and retail trade and for business services. Benchmark prediction vs. full catch-up counterfactual from Table 2.6.

had experienced a full catch up in the productivity of business services by 2009, the aggregate labor productivity would have been higher in every single country, and, with the exception of Italy, every country would have continued to close the aggregate productivity gap with respect to the U.S. after 1990, when Europe started to fall behind. Moreover, Belgium and the United Kingdom would have closed the aggregate productivity gap by catching up to the U.S. only in business services, and Europe as a whole would have closed about 60 per cent of the aggregate productivity gap with respect to the United States.

Finally, Figure 2.12 compares the 2009 labor shares of our benchmark model to the implied 2009 labor shares when Europe counterfactually experiences a full catch in either wholesale and retail trade or in business services. The solid line represents the 45 degree line starting at the origin. The purpose of this comparison is to demonstrate the importance of considering a structural transformation theory to deliver endogenously changes in the labor share as a consequence of productivity changes. This is in sharp contrast to methods of shift-share analysis – widely used in the empirical literature – where one cannot account for changes in the weight of a sector (*i.e.* the labor share) as consequence of a counterfactual change of labor

productivity. Counterfactual exercises based on a shift-share approach, as opposed to ours, would miss the change in sectoral labor shares caused by alternative sectoral labor productivity growth rates.

Figure 2.12 illustrates that if Europe had experienced a catch up in wholesale and retail trade, the weight of this sector in the economy would have been higher; The income effect brought by a full catch up in the labor productivity of this sector would need to have been stronger than the price effect in order to observe such increase in the labor shares. A shift-share analysis would underestimate the aggregate implications of this experiment significantly. On the other hand, a full catch up of the labor productivity in business services would have shrunk the participation of this sector in the economy significantly; The price effect would have dominated the Engel curve for this sector, and a shift-share analysis would overestimate the impact of this sectoral productivity change on the aggregate productivity. Moreover, a full catch up in either of these sectors would necessarily have had effects on the labor shares of *all* sectors in the economy, making our case for considering the general equilibrium effects of counterfactual changes in sectoral labor productivity stronger.

To sum up, our counterfactual experiments highlight the importance of sectoral analysis for accounting, through the lenses of a theory of structural transformation, which are the sectors responsible for the widening labor productivity gap between Europe and the U.S. After opening the service sector into 11 comparable sectors, we find that wholesale and retail trade, business services and, to a lesser extent, financial services are the sectors largely responsible for the aggregate productivity gap. We now proceed to explain these gaps empirically.

2.7 Empirical Analysis of Labor Productivity Differences in Services

In the quantitative exercise, we have measured comparable levels of sectoral labor productivity, relative to the U.S., for eight European countries. We have identified that the dynamics of labor productivity in the service sector, and in three services in

particular, had caused the fall in labor productivity, relative to the U.S., that Europe has suffered since the 1990s. What are the factors behind the differences in the labor productivity levels of services? The empirical investigation of these factors is the topic of this section. Our findings in a nutshell are that the fall in the relative labor productivity of Europe has been mainly driven by total factor productivity. The level of relative TFP has been especially low in wholesale and retail trade and business services. In these sectors, which gave the largest contribution to the falling behind, the relative gap in TFP with respect to the U.S. accounts for most of the gap in labor productivity. In addition, we document that the increase in services' employment has not been matched by a corresponding increase in the level of physical and ICT capital input. This also had a negative impact on labor productivity.

Labor productivity depends on the level of capital endowment per employment unit, known as the capital to labor ratio, and on the efficiency in combing capital and labor into the production process. The latter is usually referred to as total factor productivity (TFP). Hence, differences in sectoral capital to labor ratios and TFP levels among the European countries of our sample and the U.S. are likely to explain the productivity gap in services between the U.S. and Europe. In order to assess this claim, we start by choosing an appropriate empirical model, consistent with neoclassical production theory. In the quantitative model previously studied, sectoral production is assumed to be linear in the labor input, featuring constant returns to scale, and with a marginal product of labor corresponding to the level of labor productivity:

$$y_{i,t,c} = A_{i,t,c} l_{i,t,c},$$

where the subscripts i, t, c stand for service type, year, and country, respectively. Labor productivity $A_{i,t,c}$ is intended as a synthesis of the deeper factors just mentioned: Capital and TFP. One can think at the production function of our model as a reduced-form representation of a fully-fledged technology, in which capital and TFP are explicitly captured. Furthermore, we can distinguish between "physical" capital and "information and communication technology" (ICT) capital[17]. Let us

[17]We acknowledge that this distinction is not exhaustive. Another important aggregate to be

assume that the fully-fledged technology has the standard form of a Cobb-Douglas production function:

$$y_{i,t,c} = M_{i,t,c} k_{i,t,c}^{\alpha} s_{i,t,c}^{\beta} l_{i,t,c}^{\gamma},$$

where M stands for TFP, k for physical capital, s for ICT capital, and l for hours worked. At this stage, we assume that the fully-fledged production function is characterized by constant returns to scale, which formally requires $\alpha + \beta + \gamma = 1$. Later, we will consider the instance of departing from this assumption. The concept of labor productivity studied in this chapter is output per unit of employment. Hence, labor productivity is formally defined as $A_{i,t,c} = y_{i,t,c}/l_{i,t,c}$. Within the fully-fledged technology, this definition implies that

$$A_{i,t,c} = M_{i,t,c} \left(\frac{k_{i,t,c}}{k_{i,t,c}}\right)^{\alpha} \left(\frac{s_{i,t,c}}{l_{i,t,c}}\right)^{\beta} l_{i,t,c}^{(\alpha+\beta+\gamma-1)} = M_{i,t,c} \bar{k}_{i,t,c}^{\alpha} \bar{s}_{i,t,c}^{\beta}.$$

The last result formalizes that labor productivity is a function of TFP and the two distinct capital to labor ratios. Since the focus of the present study is labor productivity in European services relative to the U.S., we are interested in studying how differences in TFP and capital to labor ratios between the two regions relates to differences in labor productivity. We impose the assumption that the production technology is the same across countries, sectors, and years, and that variation stems only from input utilization and efficiency[18]. Therefore, labor productivity relative to the U.S. is given by

$$\frac{A_{i,t,c}}{A_{i,t,USA}} = \frac{M_{i,t,c}}{M_{i,t,USA}} \left(\frac{\bar{k}_{i,t,c}}{\bar{k}_{i,t,USA}}\right)^{\alpha} \left(\frac{\bar{s}_{i,t,c}}{\bar{s}_{i,t,USA}}\right)^{\beta}$$

or

$$\hat{A}_{i,t,c} = \hat{M}_{i,t,c} \hat{k}_{i,t,c}^{\alpha} \hat{s}_{i,t,c}^{\beta},$$

considered is, for instance, human capital. However, the available data about employees' education allow us to compute measures of sectoral human capital just for a very short sub-sample of years (from 2002 to 2009). Hence, due to the scarcity of observations, we restrict our empirical analysis to physical and ICT capital only.

[18] See Sáenz (2017) for a work considering time-varying sectoral capital intensities in production technologies.

adopting a new notation for indicating the measures relative to the U.S. ones. As a final step, we linearize the last equation by means of a logarithmic transformation:

$$\log \hat{A}_{i,t,c} = \log \hat{M}_{i,t,c} + \alpha \log \hat{k}_{i,t,c} + \beta \log \hat{s}_{i,t,c}. \tag{2.14}$$

From World KLEMS and OECD sources, we build measures of capital to labor ratios for the eleven service sectors of our study. We decompose capital into physical (land, transport equipment, machinery, and structures) and ICT (IT, communication, and software equipment). We obtain an unbalanced panel data set covering the U.S. and all the European countries of our analysis with the exception of Belgium, for which capital data are not available. The time horizons covered by the panel also vary by country, due to data availability[19]. Using the measures of capital utilization, hours worked in services, and the levels of relative productivity in the European services from our quantitative analysis, we estimate the following empirical specification:

$$\log \hat{A}_{i,t,c} = \delta_0 + \delta_1 \log \hat{k}_{i,t,c} + \delta_2 \log \hat{s}_{i,t,c} + \varepsilon_{i,t,c}. \tag{2.15}$$

δ_1 and δ_2 are least squares estimators of α and β respectively. However, we concede that unobserved characteristics of country, sectors, years, or combinations of the latter may have important effects on labor productivity, and thus not controlling for them may introduce bias in the estimates. To deal with this issue, we include dummies for fixed effects in the econometric model. We adopt an agnostic approach about the fixed effect specification to select. We estimate models with various combinations of fixed effects and we let the data guide us toward the best one, based on information criteria which weight both the goodness of fit of the model and its parsimony. Each of the first five columns of Table 2.7 displays the estimates of a different model, its fixed effect specification, and both its Akaike (AIC) and Bayesian (BIC) information criterion.

In a set of nested models, we should select the model specification with the *lowest* value of AIC or BIC. This rule leads us in preferring model (2) to model (1), con-

[19]See Appendix A for more details on the data and the sources used.

Table 2.7: Physical and ICT capital effects on labor productivity

	(1) OLS	(2) OLS	(3) OLS	(4) OLS	(5) OLS	(6) SGMM
Physical Capital	0.467***	0.454***	0.246***	0.399***	0.463***	0.484***
	(0.044)	(0.037)	(0.021)	(0.057)	(0.039)	(0.011)
ICT Capital	-0.247***	0.186***	0.050***	0.246***	0.195***	0.188***
	(0.037)	(0.026)	(0.013)	(0.040)	(0.026)	(0.007)
Fixed Effects						
Year		X	X	X	X	X
Country		X	X	X	X	X
Sector		X	X	X	X	X
Country × Sector			X			X
Year × Sector				X		
Country × Year					X	
N	1485	1485	1485	1485	1485	1485
R-squared	0.07	0.77	0.99	0.80	0.77	
AIC	5278.17	3322.93	-1500.36	3738.17	3487.76	
BIC	5294.08	3630.51	-874.59	5663.22	4267.33	

Dependent variable is log sectoral labor productivity, relative to the U.S. Physical capital is log sectoral physical capital endowment per hours worked, relative to the U.S. ICT capital is log sectoral physical capital endowment per hours worked, relative to the U.S. Standard errors are robust.

firming that including some fixed effects improves the fit of the model. Moreover, using the same rule, model (2) appears superior to models (4) and (5), and inferior to model (3). Hence, we select the fixed effects' specification of model (3) as the preferred one: It controls for unobserved characteristics of each year, country, and sector, as well as for unobserved characteristics specific of a given sector in a given country. A further concern is that the independent variables can be endogenous, i.e. not orthogonal to the OLS residuals. Following Bloom et al. (2012), we check the robustness of our findings by estimating model (3) using the system GMM method developed by Blundell and Bond (1998). The results are shown in column (6): The coefficient estimates are larger than the ones obtained by estimating model (3) by OLS, but they are qualitatively similar. We consider the results of column (3) our preferred coefficient estimates, and the rest of the discussion is based on them.

The estimated coefficients show that both physical and ICT capital endowments

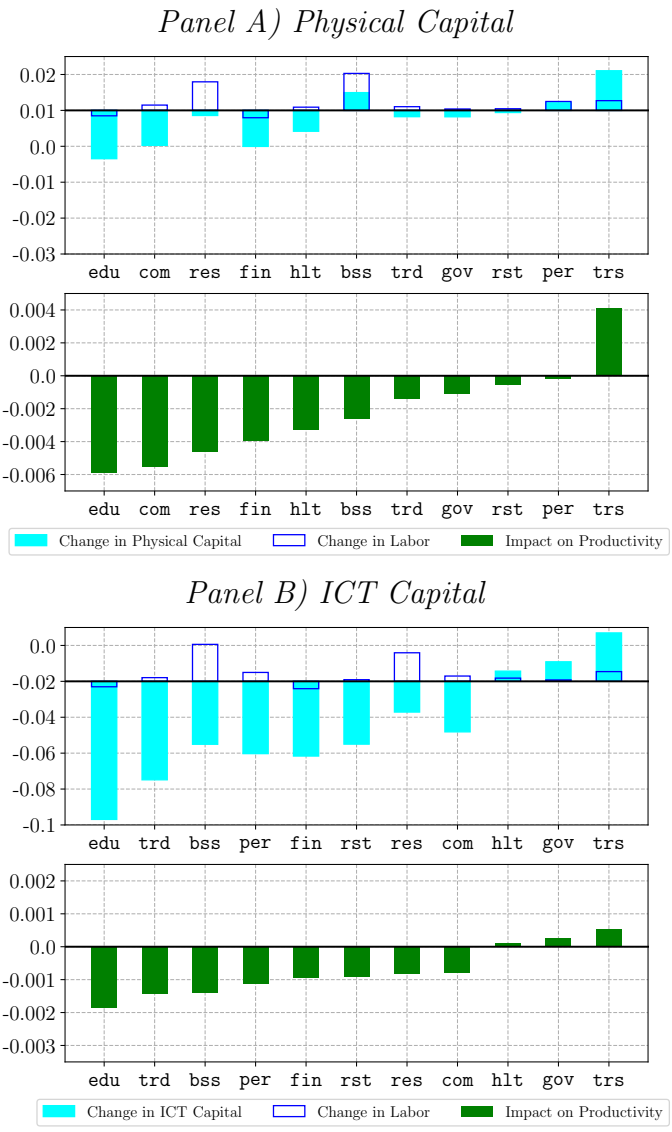
are significantly and positively associated with services' relative labor productivity. Over the years considered, the average level of labor productivity in the European service sector, relative to the U.S., increases by 0.24 per cent for a 1 per cent increase in the average level of physical capital per hour worked, relative to the U.S., and by 0.05 per cent for a 1 per cent increase in the average relative level of ICT capital per hour worked. A standard t-test rejects the hypothesis that the coefficient of ICT is not smaller than the coefficient of physical capital. This result suggests that the productivity gain from investing in physical capital is on average greater than the one obtained from increasing the ICT capital. This seems to partially contradict previous studies of the determinants of the low productivity of the European service sector. Indeed, when studying productivity growth instead of levels, and relying on growth accounting techniques and shift-share analysis, the existing literature has pointed to ICT production, adoption, and utilization as one of the major reasons of the widening gap between U.S. and European productivity. For instance, van Ark et al. (2003) identify the sectoral components of the aggregate productivity growth gap, distinguishing between contribution stemming from changes in sectoral employment shares and sectoral productivity levels. They conclude that most of the aggregate productivity gap is accounted for by so called ICT-producing sectors and ICT-using services. The former's employment share increased in the U.S. much more than in Europe, while the latter's productivity grew in Europe at a lower rate than in the U.S., suggesting that the problem of productivity in Europe comes from the fact that the European economies have not been able to take fully benefit from the ICT revolution, or at least not as much as the U.S. did. A similar conclusion is reached by Bloom et al. (2012), which identify larger returns from ICT investments enjoyed by U.S. multinationals even outside the U.S., and argue that the organizational structure of U.S. business may favor them in adapting to new technologies. Timmer et al. (2011) also contribute to the discussion by decomposing productivity growth into the contributions from changes in TFP and capital to labor ratios. They find that the productivity growth gap between Europe and the U.S. reflect mainly gaps in total factor productivity, with an important role played also by ICT capital to labor ratio. We do not think that the importance of ICT has been overestimated. However,

our empirical finding is indicative that physical capital should not be overlooked, and closing gaps also in the level of physical capital endowment can generate an important improvement in labor productivity.

From the stylized facts of structural transformation, we know that the level of employment in the service sector increased significantly both in Europe and in the U.S. over the period studied. Given the importance of capital to labor ratios for labor productivity, we ask whether the fall in European services' productivity relative to the U.S. can have occurred because of an insufficient capital accumulation to match the new employment levels. For each of the eleven services analyzed, we compute the average rate of change in employment occurred in Europe between 1990 and 2009. We compare these changes to the simultaneous average changes in sectoral physical and ICT capital levels. Figure 2.13 plots these change rates, considering the two types of capital separately.

In the lower portion of each panel, the bars report the effect on labor productivity of the difference between the change rates of employment and capital, that is how the change in the capital to labor ratio has affected labor productivity, given the estimated coefficients. For almost all the services, both the physical and ICT capital to labor ratios fell with respect to the U.S. from 1990 to 2009, with a negative impact on labor productivity. The only exceptions are given by transportation, a sector in which the physical and ICT capital endowment increased more than the employment, and by health services and government, with regard to ICT only. In all other services, the stocks of capital have not been able to keep the pace of growth in the labor allocation, or they have decreased by a far larger extent than the employment. The situation appears particularly worrisome for ICT capital. Except the sectors already mentioned, in all the services the level of ICT capital fell relative to the U.S. over the 1990-2009 horizon. In the same period, almost all the services increased their employment levels relative to the U.S. ones. Physical capital has not fallen has much as ICT, and in some sector its level has even increased. However, also in these cases the slight increase in physical capital endowment has been generally outsized by the increase in employment. This occurred to business services, one of the sectors identified as mostly accountable for the falling behind of European aggregate labor

Figure 2.13: Change rates in employment and capital stocks



Change in physical (ICT) capital is the average annual rate of change in sectoral physical (ICT) capital endowment between 1990 to 2009, averaging over the European countries of our sample. Change in labor is the average annual change in sectoral hours worked over the same period, averaging across Europe. The impact on productivity is given by the difference between physical (ICT) capital and labor rates of change, multiplied by the coefficient estimate of physical (ICT) capital to labor ratio from column 3 in Table 2.7.

productivity. Business services as well as wholesale and retail trade appear to have been particularly hit by a relative fall in their ICT endowment at the time of an increase in their level of employment. For this reason, they are among the services whose labor productivity has been mostly harmed by a fall in ICT capital to labor ratio. This finding is reconciling our analysis with the existing literature. As previous studies have pointed to ICT utilization as a major problematic area for European labor productivity, we find evidence that wholesale and retail trade and business services has suffered a severe reduction in their ICT capital to labor ratios relative to the U.S., and we have previously identified these sectors as the mostly accountable for the fall in the labor productivity of Europe.

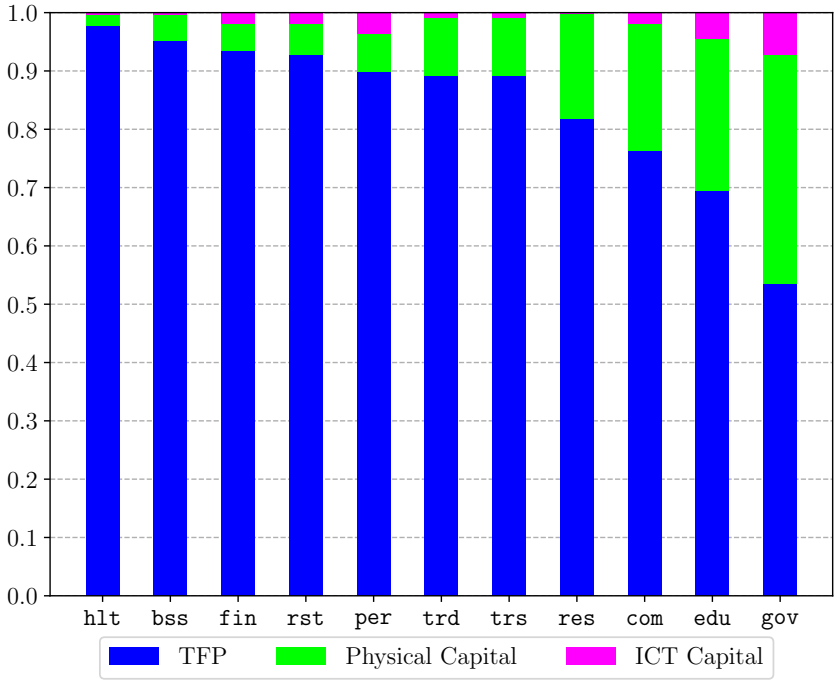
So far we have studied how services' relative labor productivity is affected by different types of capital to labor ratios. Now we turn our focus to the component of labor productivity given by the efficiency in capital and labor utilization, defined as total factor productivity (TFP). Consistent with the discussion at the beginning of this section, we obtain measures of relative TFP, in logs, using the coefficient estimates from equation 2.14:

$$\log \hat{M}_{i,t,c} = \log \hat{A}_{i,t,c} - 0.246 \log \hat{k}_{i,t,c} - 0.050 \log \hat{s}_{i,t,c}.$$

How much of the gap in labor productivity between Europe and the U.S. can be accounted for by the difference in TFP, and how much by differences in the physical and ICT capital endowment per hour worked? Figure 2.14 presents a decomposition of average sectoral labor productivity levels in Europe, relative to the U.S., based on our estimates and TFP measurement for the years between 1990 and 2009.

There is a remarkable degree of variation in the weights of each productivity component across sectors. However, we quantify that more than half of the average productivity gap between European and U.S. services is accounted for by differences in TFP. The weight of TFP is particularly high in wholesale and retail trade, business services, and financial services, accounting for approximately 90 per cent of the relative labor productivity. The decomposition gives a perspective to the discussion about the compared changes in employment and capital endowments. Although in

Figure 2.14: Decomposition of relative labor productivity



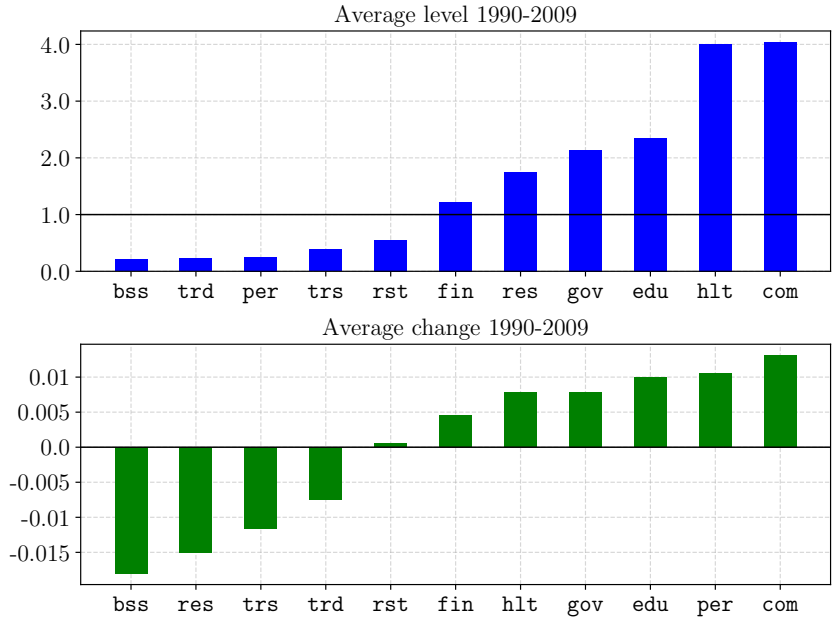
Decomposition based on average values over the European countries of our sample, between 1990 and 2009. The level of each bar corresponds to the absolute value of each element in equation 2.14, based on the coefficient estimates of column 3 in Table 2.7, divided by the absolute value of log relative sectoral labor productivity.

light of our estimates there is no doubt that a fall in the capital to labor ratios have had a negative impact on labor productivity, the role of capital endowments per hours worked appears secondary to that of TFP differences in accounting for the labor productivity gap. For most of the services in Europe - and in particular for the three at the base of the European falling behind - the issue of labor productivity differences is primarily a matter of differences in TFP.

Figure 2.15 plots the average level of total factor productivity in Europe relative to the U.S., $\hat{M}_{i,t,c}$, computed for each service over the period of the falling behind. We plot also the average rate of change of these measures over the same period.

Very marked differences across services are evident. On the one hand, we can see

Figure 2.15: Relative total factor productivity



Average sectoral relative TFP levels across the European countries of our sample.

the very low average level of TFP in business services and wholesale and retail trade in Europe, relative to the U.S., a level that also decreased for these services between 1990 and 2009. Given the above mentioned importance of TFP as a component of labor productivity in these two sectors, this finding clarifies why the productivity of business services and wholesale and retail trade has performed so poorly in Europe, dragging down the aggregate productivity of the entire economy. On the other hand, some services appear to have TFP levels extremely higher in Europe than in the U.S., and growing over the period of the falling behind.

2.7.1 Effects of Employment Levels on Labor Productivity: Decreasing Returns to Scale?

In our opinion, the fact that Europe has outperformed U.S. in some sector is not surprising. However, we quantify some average TFP levels with an order of magnitude of many times the corresponding U.S. ones, which seems puzzling. We wonder if the assumption that services' technology features constant returns to scale may lead to this unexpected measurement. If we let $\alpha + \beta + \gamma$ be, in principle, different than 1, equation 2.14 changes into

$$\log \hat{A}_{i,t,c} = \log \hat{M}_{i,t,c} + \alpha \log \hat{k}_{i,t,c} + \beta \log \hat{s}_{i,t,c} + (\alpha + \beta + \gamma - 1) \log \hat{l}_{i,t,c}, \qquad (2.16)$$

where $\hat{l}_{i,t,c}$ denotes the level of labor input in sector i of country c at time t, relative to the corresponding U.S. level. Using World KLEMS data on sectoral hours worked in the U.S. and in the European countries of our sample, as well as the data used in the previous analysis, we can estimate the following empirical model:

$$\log \hat{A}_{i,t,c} = \eta_0 + \eta_1 \log \hat{k}_{i,t,c} + \eta_2 \log \hat{s}_{i,t,c} + \eta_3 \log \hat{l}_{i,t,c} + \epsilon_{i,t,c}. \qquad (2.17)$$

The coefficient η_3 is a least squares estimator of $\alpha + \beta + \gamma - 1$, and it captures the average effect that the level of sectoral employment might have on the level of labor productivity. We estimate equation 2.17 under different specifications of fixed effects, and, as before, we let information criteria guide us. Table 2.8 contains the findings.

Also in this case, the specification controlling for year, sector, country, and sector by country effects is the preferred one, and a system GMM estimation of the same specification returns coefficient estimates in line with the OLS ones. The novelty in this empirical analysis is the coefficient estimate for the effect of log relative employment levels: It is significant and negative consistently in all the specifications considered. This is evidence that services' labor productivity, that is real output per hours worked, is negatively associated with the level of hours worked. The result is also at odds with the hypothesis that $\alpha + \beta + \gamma = 1$ and, instead, favorable to $\alpha + \beta + \gamma < 1$. Can we conclude that the production of services is characterized

Table 2.8: Capital and labor effects on labor productivity

	(1) OLS	(2) OLS	(3) OLS	(4) OLS	(5) OLS	(6) SGMM
Physical Capital	0.321***	0.026	0.108***	0.154***	-0.008	0.083***
	(0.049)	(0.032)	(0.019)	(0.042)	(0.034)	(0.015)
ICT Capital	-0.276***	0.084***	0.061***	-0.051*	0.063***	0.091***
	(0.037)	(0.021)	(0.011)	(0.029)	(0.021)	(0.010)
Labor	-0.394***	-1.860***	-0.484***	-2.257***	-1.982***	-1.771***
	(0.049)	(0.065)	(0.029)	(0.078)	(0.066)	(0.025)
Fixed Effects						
Year		X	X	X	X	X
Country		X	X	X	X	X
Sector		X	X	X	X	X
Country × Sector			X			X
Year × Sector				X		
Country × Year					X	
N	1485	1485	1485	1485	1485	1485
R-squared	0.11	0.87	0.99	0.90	0.87	
AIC	5223.50	2469.60	-1911.35	2646.71	2600.16	
BIC	5244.71	2782.49	-1280.28	4582.37	3385.03	

Dependent variable is log sectoral labor productivity, relative to the U.S. Physical capital is log sectoral physical capital endowment per hours worked, relative to the U.S. ICT capital is log sectoral physical capital endowment per hours worked, relative to the U.S. Labor is sectoral number of hours worked, relative to the U.S. Standard errors are robust.

by decreasing returns to scale? We do not think we can go that far, although this empirical result is supportive of this hypothesis. The scope of the present study is not that of estimating the definite properties of a production function for the service sector, but rather to empirically assess the relation between labor productivity and the sectoral inputs' allocation. However, we think that the evidence we find against $\alpha + \beta + \gamma = 1$ calls for deeper analysis of this aspect that might lead to further insights about the returns to scale in services. Moreover, we are not the first ones in finding empirical evidence that challenges the hypothesis of constant returns to scale: Bloom et al. (2012) also estimate a significantly negative effect of employment level on labor productivity.

How would the relevance of the main points discussed above - the impact from a fall in capital to labor ratios and the importance of TFP - change if we were to

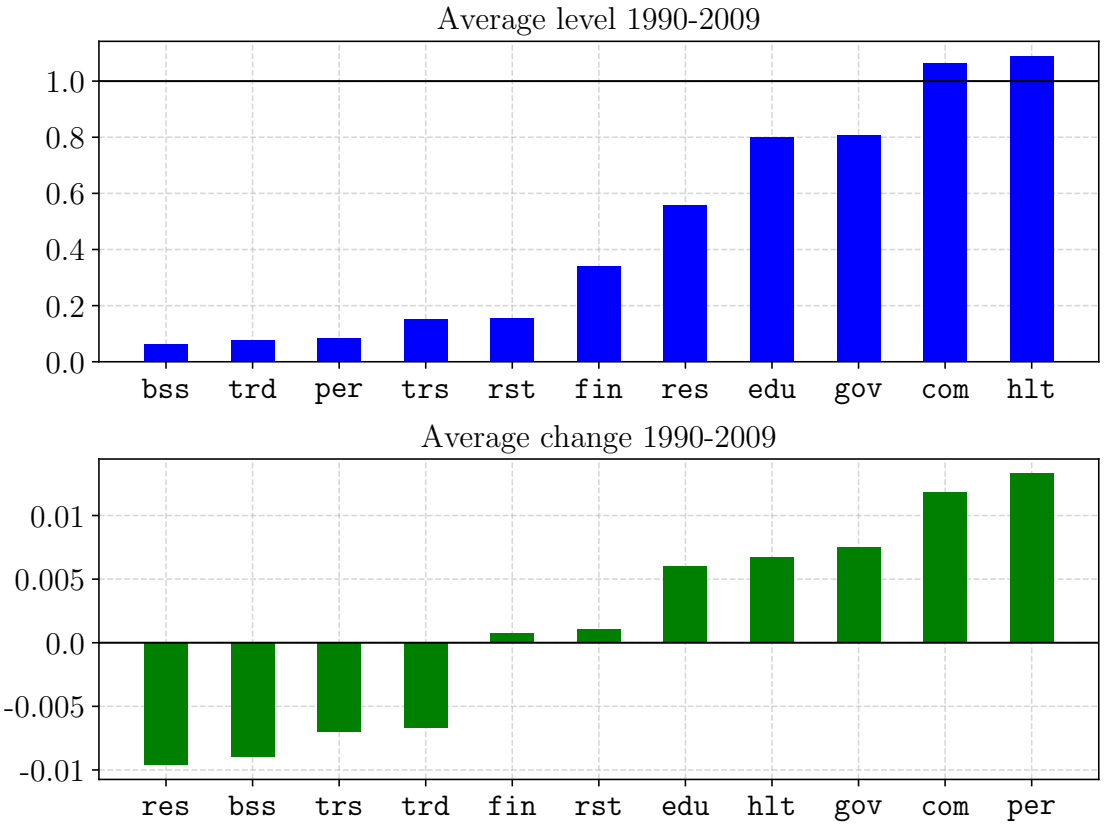
believe that services feature deceasing returns to scale? The observed reductions in sectoral physical and ICT capital to labor ratios have an even bigger negative impact on productivity under decreasing returns to scale. In order to understand this, we compute the "break-even" change in capital endowment, that is the increase in sectoral physical or ICT capital necessary to balance off an increase in sectoral employment, to the point of leaving sectoral labor productivity unchanged. Under constant returns to scale, the break-even change is exactly equal to the change in labor allocation: If capital endowment increases as much as employment, the ratio and labor productivity do not vary. With decreasing returns to scale, instead, capital must increase more than proportionally with labor. Indeed, it is not enough that the capital to labor ratio does not fall. The ratio should actually increase to compensate the negative impact on labor productivity caused by the rise in the level of employment. Hence, the fact that physical and ICT capital endowments did not match the change in labor allocation in the European services between 1990 and 2009 looks an even more serious issue if the service sector were indeed operating under decreasing returns to scale.

Including the level of employment in the empirical specification clearly changes the measurement of TFP, which is now given by

$$\log \hat{M}_{i,t,c} = \log \hat{A}_{i,t,c} - 0.108 \log \hat{k}_{i,t,c} - 0.061 \log \hat{s}_{i,t,c} + 0.484 \log \hat{l}_{i,t,c}$$

Intuitively, the new measures of TFP are likely lower than the previous ones. Indeed, the most important difference between the two measurement formulas is given by $+0.48 \log \hat{l}_{i,t,c}$. The level of hours worked in a sector of a given European country is most probably lower than the corresponding level in the U.S., by the simple reason that the labor force in the U.S. is much larger than in any single European country. Hence, $\log \hat{l}_{i,t,c}$ tends to be negative. With some approximation, this means that, when we allow for decreasing returns to scale, we subtract a potentially relevant quantity from the level of TFP that we measure under the assumption of constant returns to scale. In order to verify the soundness of this intuition, we repeat Figure 2.15 with the new measures of TFP.

Figure 2.16: Relative total factor productivity – Decreasing returns to scale



Average sectoral relative TFP levels across the European countries of our sample. Estimates from a specification allowing for decreasing returns to scale.

Figure 2.16 shows that, as expected, the average levels of sectoral TFP in Europe are smaller than the previous measures. This is particularly evident in the services with an estimated TFP larger than the U.S.. For instance, the current measures of TFP in European communication and health services are now about 10 per cent larger than in the U.S., much less puzzling than the four-fold levels previously measured. However, we want to highlight that the properties of the TFP measures for wholesale and retail trade and for business services do not seem to differ whether or not constant returns to scale are assumed. Indeed, also with decreasing returns to

scale the estimates for these two sectors portray a very small TFP level in Europe, relative to the U.S., and a reduction in this level during the period of the falling behind. Even with respect to the decomposition of labor productivity, the change in the assumption regarding the returns to scale does not modify the result that the labor productivity gap in these two services is mostly a matter of relative TFP (the weight of this component is approximately 70 per cent for wholesale and retail trade and business services). The fact that different assumptions about returns to scale lead to very different findings for some sectors and not so different results for others rises another interesting question: Are some services characterized by decreasing returns to scale more than others? We leave also this question open for future research.

2.8 Conclusions and Discussion

In this chapter, we propose a model of structural transformation that disaggregates services in order to quantitatively study the labor productivity differences between Europe and the U.S. We identify wholesale and retail trade, business services, and, to a lesser extent, financial services as the sectors that principally caused low service productivity in Europe, and ultimately lead to the divergence of European aggregate productivity from U.S. levels since the 1990's. Wholesale and retail trade has always employed a large share of labor, while business services has experienced an astonishing increase in its employment share over the period of our analysis. These patterns are similar both in the U.S. and in Europe. However, labor productivity growth in these sectors has been particularly slower in Europe than in the United States. High and/or increasing labor shares and underperforming labor productivity growth in these two sectors are at the core of the outcome uncovered by our quantitative analysis.

Having established that wholesale and retail trade and business services are the main culprits of Europe's lack of catch up with the U.S. in services labor productivity, we address empirically the components of sectoral labor productivity levels: Physical and ICT capital to labor ratios and TFP. We find that the European ser-

vices have experienced a fall in the level of capital endowment per hour worked with respect to the U.S., with negative consequences for labor productivity. Wholesale and retail trade and business services have been particularly characterized by an under-investment in ICT. Also, TFP has a very relevant role in explaining labor productivity differences. Wholesale and retail trade and business services had the lowest average levels of sectoral TFP, relative to the U.S., during the years of the falling behind, and these levels even decreased over the same period.

Which factors have led to the gap in TFP? We suspect that an important role may have been played by the different regulations of the product, capital, and labor markets in the U.S. and in Europe. The hypothesis has been discussed also in previous studies. According to Lewis (2005), stricter market regulations are the key determinant of the low productivity of services in Europe. His argument is based on an extensive analysis conducted by the McKinsey Global Institute since the 1990's, and it is substantiated by the discussion of case studies that exemplify how specific sectoral regulations can harm sectoral productivity. The conclusion of Lewis (2005) is that obstacles to the natural forces of competition are a major blow to productivity growth, and Europe has been lenient in removing them. The importance of regulation for productivity is also highlighted by Crafts (2006), who argues that the acceleration in U.S. productivity in the 1990's was possible thanks to a more flexible regulatory environment than in Europe. Nicoletti and Scarpetta (2003) state that aligning the regulatory stance of Europe to the most liberal OECD countries would substantially ameliorate European TFP growth. The negative impact of regulation on productivity has also been evidenced, more recently, by Cette et al. (2016). Unfortunately, we have not been able to find satisfactory data on the level of regulation to test empirically its importance in explaining our measures of TFP gaps. Most indexes of regulation are available only for more recent years, and do not show a significant time variation[20]. Moreover, all of them miss the crucial dimension of sectoral variation, a major focus of this work. Indeed, the available sources report country-based measures of regulation but do not capture differences, if any, in sector-specific regulation.

[20]For instance, the OECD product market regulation measures are available starting from 1998, at a five year frequency.

Some limitations we face in our chapter, particularly in data availability, highlight ways in which future research could go in unveiling labor productivity differences between Europe and the U.S. For instance, we strongly believe that, data permitting, a deeper analysis of how regulations affected labor productivity differently across sectors is an interesting topic for future research. In addition, we assume that the production technology is the same for different service types, and across countries and years. Estimating sector-specific technologies might be a relevant exercise in making more precise quantitative statements about labor productivity differences. Finally, our empirical findings suggest that the service sector might be characterized by decreasing returns to scale. We think this issue deserves further consideration in future work.

Our findings, together with the rising importance of services in the economy, imply that policies aiming at fostering aggregate labor productivity growth in European economies should be focused on wholesale and retail trade and business services, promoting investment in ICT and physical capital as well as creating an environment that facilitates a more efficient use of production inputs in these two key sectors.

CHAPTER 3

THE RICARDO EFFECT: EVIDENCE FROM THE MANUFACTURING LABOR INCOME SHARE IN COLOMBIA

Joint work with Luis Felipe Sáenz

3.1 Introduction

In the third edition of his *On the Principles of Political Economy and Taxation* (Ricardo (1821)), David Ricardo acknowledges that the replacement of human labor with machines is typically detrimental to the interests of the workers, although it improves the efficiency of the production process in general. This point has been retaken in consideration by both Keynes (1930) and von Hayek (1942)). The latter has introduced the term Ricardo Effect to denote the old Ricardo's idea that the substitution of capital for labor in the production comes at the expense of workers.

The purpose of this chapter is to provide quantitative evidence of the Ricardo Effect using a unique plant-level longitudinal dataset for Colombian manufacturing establishments for the period 1982-1998. The data requirements needed for establishing a relationship between different inputs of production are very stringent. It must include at least information of labor and capital at the plant level, which is the relevant unit of analysis, and it must vary across time since this relationship is dynamic. This is precisely the information available in the Annual Manufacturing Survey (EAM[1]) in Colombia.

Colombia is a very interesting case of study mainly for two reasons. First, the information at disposal is based on a uniquely rich and representative data for Colombian

[1]Acronyms in Spanish for "Encuesta Anual Manufacturera".

manufacturing plants, derived from yearly plant censuses over the period 1982-1998 with detailed information of physical quantities of inputs. It is the most complete source of product-level information in a nationally representative plant database in any country (Kugler and Verhoogen (2012)). Second, the Colombian experience can be considered as a "natural experiment" of exogenous shocks to the relative prices of inputs, since during the early 1990s the country underwent countrywide market-oriented reforms, and thus the data provide a clean base for comparison between pre-reform and post-reform periods (Eslava, Haltiwanger, Kugler, and Kugler (2004)).

This chapter uses the EAM data to test (i) whether there is supporting evidence of the Ricardo Effect in Colombia and (ii) whether this effect changed under a period of market-oriented reforms whose purpose was, among several others, to reduce distortions in the factor markets. To test for the existence of the Ricardo Effect, we assume a setup in which production plants operate a technology featuring constant elasticity of substitution between capital and labor. The value of this elasticity is key: The Ricardo Effect is in place if such elasticity is larger than one. Indeed, in that case, the arrival of newer, cheaper capital induces a substitution for labor and a consequent decline in the labor income share, *i.e.* the proportion of income devoted to compensate workers. This is exactly the instance described by Ricardo: A replacement of labor with capital causing an economic damage to the laboring class.

Using different specifications and estimation techniques, we estimate that the elasticity of substitution between capital and labor is significantly larger than one. The preferred estimate equals 1.85. These findings provide evidence of the existence of the Ricardo Effect in the manufacturing sector of Colombia. The market-oriented reforms do not seem to have caused any change in the elasticity of substitution between capital and labor. However, the reforms have accelerated the fall of the cost of capital relative to the cost of labor. Given our estimates of the elasticity of substitution, this fact has reinforced the decline in the labor income share. By means of a simulation exercise, we assess how much of the observed decline in the Colombian labor income share can be accounted for by the Ricardo Effect, given the observed decline in the relative cost of capital and our estimates of the elasticity of substitu-

tion. The conclusion is that the Ricardo Effects explains half of the decline of the labor income share in Colombia.

Our findings are very close to those of Karabarbounis and Neiman (2014). They study the effect of capital and labor substitution on the labor income share using a production setup with CES technology, similar to what we are doing. They estimate an elasticity of substitution between these two inputs of 1.25. Based on this parameter value, their model can explain roughly half of the worldwide decline in the labor income share. The main difference between our research and their paper is that their estimation is based on cross-country variation, while we are using longitudinal micro data.

The rest of the chapter is organized as follows. Section 3.2 provides more details about the concept of Ricardo Effects and the literature that studied it. Section 3.3 describes the data and provides a brief description of the Colombian context in light of the evidence. Section 3.4 illustrates the empirical strategy pursued in this research. Section 3.5 presents the main results of the chapter. Section 3.6 provides some concluding remarks.

3.2 The Ricardo Effect: A review of the literature

In 1821, for the third edition of his 1817 masterpiece entitled *On the Principles of Political Economy and Taxation*, David Ricardo decided to include a whole new chapter to his bestseller in which he wrote a *mea culpa* regarding his previous ideas with respect to the role of machines. Ricardo confessed that before writing "On the Machinery", the 31st chapter of his classic, he was not aware of any conflict between the interests of the laboring class and the arrival of machines to the production process. In Ricardo's own words:

> (...)I have been of opinion, that such an application of machinery to any branch of production, as should have the effect of saving labour, was a general good, accompanied only with that portion of inconvenience

> which in most cases attends the removal of capital and labour from one employment to another. Ricardo (1821, pp. 466-67)

David Ricardo devoted a whole new chapter in his *Principles* to reveal his change of opinion. His new vision was that the application of machinery could reduce labor demand (Ricardo (1821), Samuelson (1989)).[2]

Sympathetic with Ricardo's new chapter, Samuelson (1988) introduced a "simple classical model"[3] in which the invention of robots reduces the demand for labor permanently, as Ricardo predicted. Contrary to the opinion of several followers of Ricardo, Samuelson considered chapter 31st as the best single chapter of Ricardo's book. He provided a dramatic example to illustrate that the invention of robots capable of replacing the entire human labor in the production of corn will yield Ricardo's prediction: human jobs are replaced by machines. An interesting implication explained in detail by Samuelson (1988) is that if robots are relative cheaper compared to labor, even by just a small fraction, no labor will be demanded at all. Samuelson crafted this overdramatic example of robots replacing humans as a way of vindicating Ricardo's reasoning as logically feasible, at a time when his new chapter was in doubt and was considered as a logical fallacy (Samuelson (1988)). Lord Keynes also contributed in this debate coining a term to describe the unemployment created by the introduction of machines: *Technological Unemployment*. According to Keynes (1930, pp. 196),

> We are being afflicted with a new disease of which some readers may not yet have heard the name, but of which they will hear a great deal in the years to come - namely, technological unemployment. This means unemployment due to our discovery of means of economizing the use of labor outrunning the pace at which we can find new uses for labor.

[2]The interest reader should consult directly Ricardo (1821) to understand the evolution of his ideas with respect to the role of machines. The discussion of his arguments is beyond the scope of this article, mainly because a preliminary discussion of the value labor theory is imperative in order to address Ricardo's concerns related with the distribution of income.

[3]As Paul Samuelson called it himself.

The debate regarding the complementarity/substitutability between labor and capital goods in the production process today is well and alive. Burke and Rumberger (1987) compile a series of papers that address the impacts of technology on work and education in the United States and Australia. The principal questions that are addressed in the collection or papers are related with the job creation/destruction due to the increased used of new technologies, and with what kinds of jobs will be created and what kinds will be destroyed. They conclude that new technologies, especially those associated with micro-electronics, are capable of further routinizing and simplifying tasks into repetitive and machine operated-monitored functions, but also new technologies enhance the decision role of employees and potentialize the skills and education of the labor force.

Knights and Willmott (1988) consider that as long as economies are expanding, the substitution of capital for labor due to the dramatic advance in the use of new technologies is not reflected in unemployment figures instantaneously, but with the continuum arrival of new technologies, labor demand suffers, specially during times of recessions where the technological expansion is still supported by governments.

Krusell, Ohanian, Ríos-Rull, and Violante (2000) llustrate that changes in observed inputs of production can explain most of the variations in the labor skill premium from 1963 to 1991 in the United States. They identify the following puzzle: The supply of skilled labor increased significantly during this period but at the same time the skill premium, defined as the wage of skilled labor relative to that of unskilled labor, has grown considerably since 1980. They argue that with a neoclassical production function whose technology is capital-skill complementary, the puzzle is explained in terms on input variations. In short, with the development of better and cheaper capital equipment the wages of unskilled workers are (relatively) driven down since unskilled labor is competing not only with skilled employees, but with persistently cheaper and better machines.

Krusell et al. (2000) found that the substitution elasticity between unskilled labor and equipment is 1.67 whereas for skilled labor and equipment is 0.67. They also found that the skill premium is driven by changes in observed factor quantities. The supply of skilled labor puts a downward pressure to the premium, while the capital

skill complementarity effect puts an upward pressure which ultimately dominates. Hanson (2001) considers an exogenous growth model in which machines are complement to human labor when they become more productive, but also machines are substitutes for human labor by taking over jobs. The conclusion of this modeling exercise is that in spite of the complementary effects due to increases in productivity, in the end the substitution effects are dominant.

Acemoglu (2002) contributes to this debate by addressing the direction and bias of technical change, since in most situations technical change is not neutral: it benefits some factors of production more than others. He develops a workhorse to understand why technical change can be skill biased, and why new technologies introduced during the late eighteenth and early nineteenth centuries were unskilled biased. This framework provides analytically the conditions for capital and labor to be gross complements or gross substitutes based on the idea that firms can invest resources to develop technologies that complement a particular factor. Acemoglu (2002) provides an explicit micro-foundation to the complementarity/substitutability nature of technology and production inputs.

More recently, Acemoglu and Autor (2011) proposed a framework called "A Ricardian Model of the Labor Market" in which they explicitly incorporate a distinction between workers' skills and job tasks, and they allow the assignment of skills and tasks to depend on labor supplies, technologies, and task demands. They consider that the distinction between skills and tasks is critical to understand how the set of tasks that workers perform responds to changes in supplies or technology. According to Acemoglu and Autor (2011), a task is a unit of work activity that produces output while a skill is the worker's endowment of capabilities to perform various tasks. They argue that "(...) an explicit distinction between skills and tasks (...) will enable the model to allow for certain tasks to become mechanized." Acemoglu and Autor (2011, pp. 1119) Therefore, in the task-based approach, tasks are applied to produce output, and skills have an influence in output through its relation with tasks.

3.3 Data

The data used in the estimation come from the project "Plant-Level Price Indices for Output and Materials" created under a technical cooperation between the Colombian National Administrative Department of Statistics (DANE hereafter for its acronym in Spanish) and John Haltiwanger from the University of Maryland. This database have the same coverage period and most of the information that was used in Eslava et al. (2004). The information gathered is taken directly from the Colombian Anual Manufacturing Survey (EAM hereafter for its acronym in Spanish).

The EAM is an unbalanced panel that has information since 1982 of any industrial establishment in Colombia that employs ten or more employees, or whose annual output is worth more than 65 million Colombian pesos (around 35 thousand dollars) at the reference year. These reports are adjusted each year with the producers price index created by the Colombian Central Bank. The dataset of Haltiwanger's project contains information for each establishment of the manufacturing sector for the following variables: production, capital (buildings, structures, machinery, and equipment), employees (production and non-production personnel), hours worked (average hours worked per employee times number of employees per sector per year), materials (intermediate consumption), and energy consumption. Production, capital and materials are in constant thousands of pesos of 1982, whilst energy is in Kw per hour.

Eslava et al. (2004) and the technical document that accompanies the "Plant-Level Price Indices for Output and Materials" database provide detailed documentation of the construction of the variables. However, since the measurement of capital is critical for our purposes, we will explain briefly the construction of this variable. The capital stock is constructed recursively based on the following formula:

$$K_{it} = (1 - \delta)K_{it-1} + \frac{I_{it}}{D_t}$$

where K_{it} are the units of physical capital for plant i in year t, K_{it-1} are units of physical capital for plant i in year $t-1$, δ is the depreciation rate, I_{it} is the gross investment for plant i in year t, and D_t is the gross capital deflator for year

t.[4] The capital stock series only includes equipment, machinery, buildings, and structures. With the information on fixed assets reported by each plant together with depreciation rates and inflation reported to adjust fixed asset values, gross investment series for each plant are generated to compute the capital series (Eslava et al. (2004)).

The "Plant-Level Price Indices for Output and Materials" dataset contains also demand-shift shocks estimated in Eslava et al. (2004). In the estimation, we use these shocks as instruments for endogenous input levels. Eslava et al. (2004) builds these demand shocks as total output measures in *downstream* industries, which are industries satisfying two conditions: They buy at least 15 per cent of upstream production and the purchasing cost from the upstream industry represents no more than 15 per cent of total costs[5].

Other data that we use are the cost of capital and labor, from which we compute the relative price of production inputs. The webpage of the Colombian Central Bank provides the historical series for the Producer Price Index (PPI) since 1970 under several classifications. In particular, under the category "PPI by use or destination of good" the subcategory of capital goods is available. The cost of capital then is measured as the capital goods' PPI relative to the manufacturing PPI, normalized to a base of 100 for 1982.[6] Regarding labor costs, Urrutia and Ruiz (2010) present real wage series for several sectors and periods in Colombia. They provide the real wages discriminated by economic activity for the period 1980-2006. We constructed

[4]See Eslava et al. (2004) and the technical document of the construction of variables of the "Plant-Level Price Indices for Output and Materials" for the details regarding the depreciation rates, deflators, the generation of the gross investment series for each plant, and the assumptions for the initial capital stocks.

[5]See Eslava et al. (2004) for more details of the estimation of demand shocks.

[6]There is an extensive literature related with computations of capital costs in Colombia, but i) they consider the capital cost mostly in terms of the opportunity cost, ii) these calculations do not vary across plants or sectors in the manufacturing industry, and iii) the PPI of capital goods is already a major component of the capital costs in the algorithm. For the purposes of this document, using the PPI solely to construct relative costs of inputs is an approach more clean and tractable compared to using of any of the algorithms available. See Diéz, Gaitán, and Valderrama (2011) for a short literature review and summary of the methodologies related with the computation of capital costs in Colombia. In particular see the discussion in Diéz et al. (2011) regarding the lack of consensus to estimate capital costs.

the cost of labor as the industrial wages from Urrutia and Ruiz (2010) multiplied by the consumer price index (CPI) and divided by the manufacturing PPI, normalized to a base of 100 for 1982.

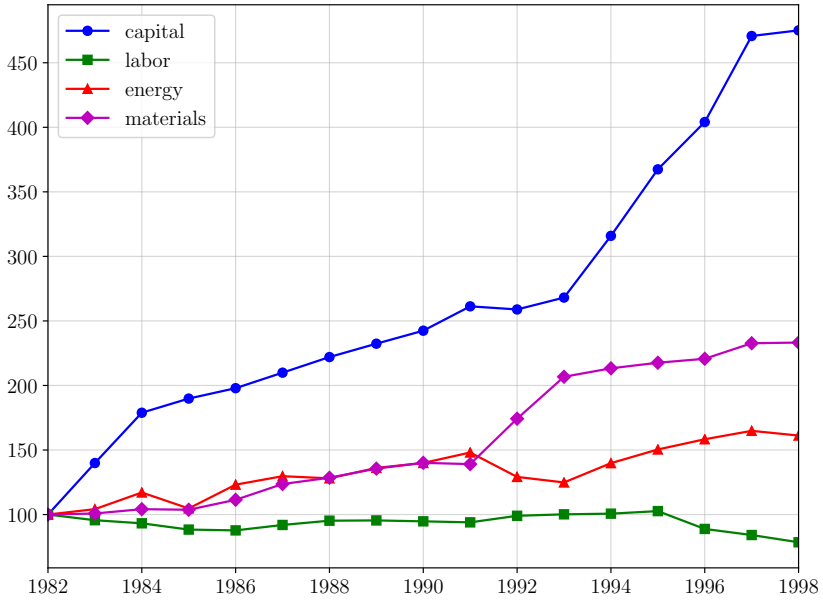
Finally, in the simulation exercise we use data about the labor income share in Colombia. The source is the Penn World Tables (PWT 9.0). In this database, the series goes back to 1950. However, before the 1990s the value of the labor share is constant, suggesting that the compilers has made an extrapolation. Hence, we take in consideration the data from 1994 onward only.

3.3.1 The market-oriented reforms

In the Colombian context, the early 1990s is a period that deserves special attention. After the infamous murder of Luis Carlos Galán, the virtual winner for the 1990 presidential elections, Cesar Gaviria won the presidency for the period 1990-1994. President Gaviria was a technocrat who worked in Galán's campaign as Chief of Staff. During his tenure several episodes marked dramatically the modern history of the country: Pablo Escobar was killed and his entire drug cartel was dismantled after years of terror; the most emblematic left-wing guerrilla group, the M-19, signed an armistice with the Colombian Government, and a new constitution in 1991 created a whole new legal environment in every level of the state. Additionally, during Gaviria's administration the Colombian economy underwent extensive structural reforms whose purpose were to enhance the role of productivity and undermine the rigidity in factor markets, with special emphasis on artificially imperfect competitive markets (Eslava et al. (2004)). In particular, dismissal costs on labor were reduced dramatically, the average tariffs fell significantly, capital markets and banking legislation were modernized, and restrictions on FDI were removed (Eslava et al. (2004)).

Can we see any changes in the time series we are using at the time of the reforms? Figure 3.1 illustrates the quantities of inputs used by the plants of our database, relative to the 1982 quantities that are normalized to 100. We can immediately notice that during Gaviria's administration both the stock of capital and the usage of materials increased significantly, at a rate larger than in the rest of the horizon. At

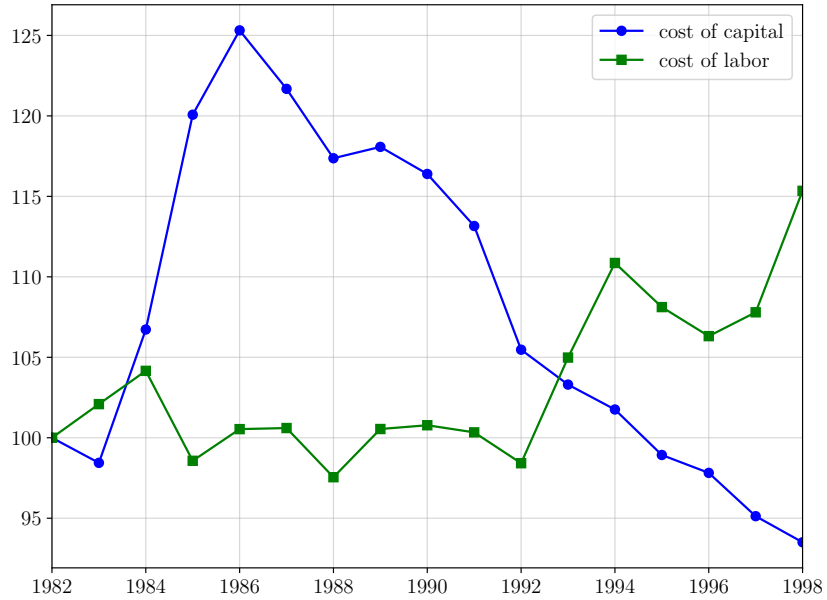
Figure 3.1: Input quantities



the same time, the amount of labor remained substantially constant until 1995, and slowly declined afterward. These facts are a clue that the reforms shift manufacturing toward a more capital intensive production form.

In Figure 3.2, we display the cost of capital and the cost of labor over the time horizon of our analysis. While the cost of capital has been steadily declining since 1986, during the reform period its fall has accelerated. This is consistent with the fact that after the reform process, the average tariffs fell, the banking sector was modernized, and the prevailing sectorial restrictions to Foreign Direct Investment were removed (Eslava et al. (2004); Edwards and Steiner (2008)). At the same time, the cost of labor jumped up in the early 1990s, after having been rather stable in the previous decade. This increment can be explained through the fact that in spite of the policies oriented to enhance the flexibility on hiring labor force as well as the reduction in hiring costs, the reform period introduced also mechanisms to provide better protection of the worker's rights, and protection to the union activity Edwards and Steiner (2008). Additionally, in 1993 a national reform increased by

Figure 3.2: Input prices



13.5 per cent the contributions of payroll to social security, where 75 per cent of these contributions were paid directly by employers Eslava et al. (2004). In conclusion, the market-oriented reforms contributed significantly to the reduction of the relative price of capital.

3.3.2 Descriptive Statistics

Table 3.1 presents the principal descriptive statistics. Capital, total employment hours, materials, energy, output, and demand shocks[7] are in logs, and the cost indexes are normalized to a base of 100 for 1982. For the period 1982-1998, the number of observations for all variables oscillates between 90 and 100 thousand, although the indexes for capital and labor costs are repeated observations of the same sector (or plant) invariant number per year in the panel. The average of capital is 8.44 with a standard deviation of 2.12. Its range is from -2.3 to 17.44 log points. The log average

[7]In section 3.4 we will describe in detail the construction of the demand shocks.

Table 3.1: Descriptive statistics - Full horizon

Variable	Mean	(Std. Dev.)	Min.	Max.	N
Capital	8.44	(2.12)	-2.3	17.44	96,232
Total Emp. Hours	10.96	(1.17)	6.68	17.81	99,102
Materials	9.89	(1.89)	-1.11	17.79	90,938
Energy	11.42	(1.93)	0	20.29	99,476
Capital Cost	108.27	(10.07)	93.51	125.32	100,114
Labor Cost	108.83	(18.84)	91.13	153.28	100,114
Materials Price Index	767.47	(962.82)	34.52	58847.97	91,540
Energy Price Index	8,394.5	(1,344,645.04)	-10,872.46	373,056,000	100,114
Output	10.68	(1.78)	-1.87	18.46	100,114
Demand Shocks	5.12	(2.65)	-1.62	32.08	100,114

Capital, total employment hours, materials, energy, output, and demand shocks are in logs, while indexes are normalized to a base of 100 for 1982.

of total employment hours is close to 11, which is about 60 thousand labor hours (employees times hours worked), with a standard deviation of 1.2. The averages for materials and energy are 9.9 and 11.4 respectively.

Regarding cost indexes, the descriptive statistics of Table 3.1 for capital cost and wages simply reflect the message of Figure 3.2 since they are nothing but time series. However, for materials and energy costs, the data has information that varies across plants. The average index for materials is 767 while for energy is 8,394. There is an important degree of dispersion in the data for these two inputs. The standard deviation for the materials price index is 962.82 while for energy is 1,344,645. This excessive volatility in energy prices is possibly explained from the fact that energy consumption is measured in Kw per hour and the bill of Kw per year, reported directly in the EAM, and the energy prices per plant can be considered on its own a measure of capital utilization. Prices of materials (and output) are constructed with Tornqvist indices where weighed average for growth in prices of materials (or products) generated by the plant are used.[8]

Last, Table 3.1 shows that for the full sample, the average output per plant was about 10.7 with a standard deviation of 1.8 with a minimum of 1.87 and a maximum

[8]See Eslava et al. (2004) and the technical document of the "Plant-Level Price Indices for Output and Materials" project for more details on plant level prices.

Table 3.2: Descriptive statistics – Pre-Reform vs. Post-Reform

Variable	Mean	(Std. Dev.)	Min.	Max.	N
Panel A. Pre-reform Period. 1982-1990					
Capital	8.21	(2.05)	-2.3	17.22	53,034
Total Emp. Hours	10.97	(1.1)	6.95	16.13	55,055
Materials	9.60	(1.85)	-1.11	17.5	51,741
Energy	11.3	(1.88)	0	20.29	54,762
Capital Cost	113.98	(9)	98.44	125.32	55,298
Labor Cost	96.49	(3.83)	91.13	102.04	55,298
Materials Price Index	307.29	(343.94)	40.69	23,109.16	52,280
Energy Price Index	7,506.85	(1,586,504.62)	0.06	373,056,000	55,298
Output	10.49	(1.67)	5.15	18.05	55,298
Demand Shocks	5.08	(2.6)	0.07	31.76	55,298
Panel B. Post-Reform Period. 1991-1998					
Capital	8.75	(2.18)	-2.13	17.44	43,198
Total Emp. Hours	10.95	(1.25)	6.68	17.81	44,047
Materials	10.25	(1.88)	0.21	17.79	39,197
Energy	11.55	(1.99)	0	20.19	44,714
Capital Cost	101.23	(6.05)	93.51	113.16	44,816
Labor Cost	124.06	(18.84)	95.26	153.28	44,816
Materials Price Index	1,380.27	(1,160.36)	34.52	58,847.97	39,260
Energy Price Index	9,489.76	(966,109.47)	-10,872.46	201,600,000	44,816
Output	10.9	(1.88)	-1.87	18.46	44,816
Demand Shocks	5.18	(2.72)	-1.62	32.08	44,816

Capital, total employment hours, materials, energy, output, and demand shocks are in logs, while indexes are normalized to a base of 100 for 1982.

18.46 log points. The average demand shock is of 5.1 log points, with a standard deviation of 2.6. The range for this shocks goes from 0.1 to 31.8.

In order to provide a first snapshot of the differences between pre- and post-reform periods in the sample, Panels A and B of Table 3.2 splits the sample between 1982-1990 (pre-reform period), and 1991-1998 (post-reform period) and provide the main descriptive statistics for each period. The capital increased from 8.2 to 8.8 log points. In 1982 thousand pesos, this is a difference of about 2,633, on average, for the period after the the reforms. The output increased in the post-reform period on average about half log point, or 18,000 thousand pesos of 1982. Table 3.2 delivers the following stylized fact: During the post-reform era, the plants on average increased

its production and its demand of capital, while the demand of workers remained stagnant. It is also noticeable that the number of observations between pre- and post-reform periods was reduced in about 10,000 observations. Even though there are 9 years in the pre-reform period and only 8 years for the Post-Reform Period, Table 3.2 suggests that some plants did not survive the new competitive environment imposed by the market-oriented reforms.

3.4 Setup and Empirical Strategy

In order to empirically investigate the existence of the Ricardo effect, we must follow a theoretical setup allowing for its existence in the first place. The typical Cobb-Douglas production function combining capital K and labor L,

$$Y = AK^{\alpha}L^{1-\alpha},$$

fails in this respect. In a competitive equilibrium, firms operating with this technology respond to a change in the relative input price by shifting the optimal ratio of input quantities *proportionally*. This implies that the substitution in input use leaves the income shares unchanged. Indeed, with a Cobb-Douglas technology, the income shares of capital and labor are fixed at α and $1-\alpha$.

Moving from the existing literature (Krusell et al. (2000), Karabarbounis and Neiman (2014)), we assume that firms produce according to a CES technology:

$$Y = A\left(\alpha K^{\frac{\sigma-1}{\sigma}} + (1-\alpha)L^{\frac{\sigma-1}{\sigma}}\right)^{\frac{\sigma}{\sigma-1}}.$$

The parameter σ captures the elasticity of substitution between capital and labor. If σ is larger than one, a change in the relative input price induces a *more than proportional* substitution of inputs in the production. Under these circumstances, the change in relative price has effects on the income shares as well: A larger portion of income is used to compensate the input that has become relatively cheaper. To see this, let R and W represent the input price of capital and labor respectively. In

a competitive equilibrium, the *ratio* of income shares satisfies the equation

$$\frac{RK}{WL} = \frac{\alpha}{1-\alpha}\left(\frac{R}{W}\right)^{1-\sigma}. \tag{3.1}$$

For σ larger than one, a reduction in the relative cost of capital (R/W decreases) generates a rise in the income share compensating the capital input (RK/WL increases). Therefore, testing for the existence of the Ricardo Effect amounts to estimating the parameter σ and verifying it is larger than one.

Using our plant-level data, we proceed to directly estimate a few variants of a CES production function. The simpler specification we estimate is

$$\log Y_{it} = a_i + \frac{\sigma}{\sigma-1}\log\left(\alpha K_{it}^{\frac{\sigma-1}{\sigma}} + (1-\alpha)L_{it}^{\frac{\sigma-1}{\sigma}}\right) + \nu_t + \varepsilon_{it}, \tag{3.2}$$

where the subscripts i and t denote plant and year respectively. Notice that we allow for plant and year fixed effects. This is equivalent to assume that total factor productivity is systematically different across plants by some unobserved characteristics. Also, in each year there are both a common productivity shock affecting all plants and an idiosyncratic productivity shock.

In the second empirical model, we introduce also energy consumption (E) and materials (M) among inputs in the production technology. Since our main concern is to study the substitution between capital and labor, we add energy and materials in a parsimonious way. We let the original CES technology in capital and labor nest an ampler Cobb-Douglas function where the new inputs are also present:

$$\log Y_{it} = a_i + (1-\beta-\gamma)\frac{\sigma}{\sigma-1}\log\left(\alpha K_{it}^{\frac{\sigma-1}{\sigma}} + (1-\alpha)L_{it}^{\frac{\sigma-1}{\sigma}}\right) +$$

$$+\beta\log E_{it} + \gamma\log M_{it} + \nu_t + \varepsilon_{it}. \tag{3.3}$$

The latter specification is missing a point well discussed in the literature: There is a strong complementarity between energy consumption and the extent of capital usage. A given stock of capital can be employed in the production process with a different degree of intensity. One may think, for instance, at a machinery that

might be used once per day or continuously over the twenty four hours, depending on the necessity of production. Most likely, the more intensely capital is used, the larger quantity of energy is consumed. We design the third specification to take into account this relation. Let us define effective capital $\hat{K}$ as the product $E \times K$: Energy consumption E serves as a measure of how intensely a stock of capital K is used. Then, the effective capital replaces capital in the technology:

$$\log Y_{it} = a_i + (1-\gamma)\frac{\sigma}{\sigma-1}\log\left(\alpha \hat{K}_{it}^{\frac{\sigma-1}{\sigma}} + (1-\alpha)L_{it}^{\frac{\sigma-1}{\sigma}}\right) + \\ +\gamma \log M_{it} + \nu_t + \varepsilon_{it}. \tag{3.4}$$

3.5 Results

3.5.1 CES Technology. NLS Estimates

Table 3.3 presents the non-linear least squares estimation of all the specifications. Columns (1) and (2) correspond to specifications (3.2) and (3.3) respectively. We estimate specification (3.4) first imposing $\gamma = 0$ (no materials in the productions function) in Column (3), and then relaxing this assumption in Column (4). Looking at the coefficient of interest, σ, the estimated elasticity of substitution between labor and capital ranges from 1.15 to 1.80. Given the tight standard errors, for any specification we reject the hypothesis that the elasticity is below one. Hence, this first set of empirical findings provides positive evidence about the existence of the Ricardo Effect in the Colombian manufacturing sector. With regard to the other parameters, capital weight α falls from 0.50 to 0.06 when the definition of capital is changed into that of effective capital. This modification does not affect the parameter capturing materials' intensity (about 0.55). Finally, when energy is considered as a separate input, its intensity is estimated at 0.08.

Table 3.3: NLS estimates

	(1)	(2)	(3)	(4)
α	0.5093***	0.5141***	0.0683***	0.0634***
	(0.0187)	(0.0250)	(0.0074)	(0.0085)
σ	1.6708***	1.8030***	1.1511***	1.1541***
	(0.0740)	(0.1266)	(0.0163)	(0.0194)
β		0.0827***		
		(0.0057)		
γ		0.5498***		0.5785***
		(0.0084)		(0.0084)
Obs.	75226	75226	73063	73063

Standard errors clustered at plant level in parentheses. $^{*}p < 0.10,^{**} p < 0.05,^{***} p < 0.01$. All regressions include plant and time effects. To control for plant fixed effects, the specifications are estimated in first difference.

3.5.2 CES Technology. 2SLS Estimates

Thinking at a realistic production process, one may acknowledge that the inputs differ in terms of their predetermination. The stock and the flow of capital are likely to be given in any period, since investment decision usually have horizons longer than one year. The extensive margin of labor is quite predetermined, at least with respect to permanent workers, but this is probably not true for temporary workers. The intensive margin of labor is definitely not predetermined as well as the usage of materials and energy consumption. Clearly, effective capital is not predetermined too. In light of that, the non-linear least squares estimates might be seriously biased by the simultaneity of productivity shock realization and non-predetermined input choices.

To address the endogeneity of labor, materials, energy, and effective capital, we propose an 2SLS strategy in the spirit of Eslava et al. (2004). In the first stage, the

Table 3.4: 2SLS estimations - First stage

	(1) $\hat{K}$	(2) L	(3) E	(4) M
Demand Shock	1.564***	1.004***	1.202***	2.575***
	(0.060)	(0.027)	(0.039)	(0.033)
Demand Shock (Lag 1)	0.749***	0.314***	0.230***	0.130***
	(0.074)	(0.033)	(0.048)	(0.040)
Demand Shock (Lag 2)	1.353***	0.293***	0.466***	0.249***
	(0.066)	(0.029)	(0.043)	(0.036)
Energy Price	-0.000*	-0.000	-0.000***	0.000
	(0.000)	(0.000)	(0.000)	(0.000)
Materials Price	-0.000***	-0.000***	-0.000***	-0.000***
	(0.000)	(0.000)	(0.000)	(0.000)
Obs.	62335	62335	62335	62335
F	786	284	273	1351

$^{*}p < 0.10,^{**}p < 0.05,^{***}p < 0.01$. All variables are in logs. The Energy Price and the Materials Price are indexes equal to 100 in 1982. Demand Shocks with different sector elasticities come from Eslava et al. (2004). All regressions include plant and time fixed effects.

endogenous inputs are instrumented by the downstream demand shocks described in Section 3.3 (current, one lag, and two lags values), the energy price index and the materials' price index. Table 3.4 illustrates the estimates of the first stage. As expected, positive shocks to downstream demand generates an increase in the utilization of all the inputs, and this positive effect persists in time, given that lagged shocks' coefficients are significant as well. An increase in the price of materials is associated with a decline in all inputs' use, while an increase in the price of energy reduces the employment of effective capital and energy only (given the scale of the measures, coefficients are extremely small). The F statistic is large for all

Table 3.5: 2SLS estimations - Second Stage

	(1)	(2)	(3)	(4)
α	0.0321***	0.1209**	0.1236***	0.0333***
	(0.0108)	(0.0469)	(0.0285)	(0.0122)
σ	2.4038**	2.5882**	1.5489***	1.8483***
	(1.0535)	(1.2822)	(0.0706)	(0.1384)
β		0.3385**		
		(0.1464)		
γ		0.1960***		0.2828***
		(0.0360)		(0.0332)
Obs.	53403	53403	53403	53403

Standard errors clustered at plant level in parentheses. $^{*}p < 0.10,^{**} p < 0.05,^{***} p < 0.01$. The values for the endogenous inputs are those predicted for at the first stage. All regressions include plant and time effects. To control for plant fixed effects, the specifications are estimated in first difference.

instrumented inputs, confirming the overall validity of the first stage.

Table 3.5 shows the 2SLS estimation of the CES technology specifications. Comparing to Table 3.3, the 2SLS point estimates of σ in specifications (3.2) and (3.3) are rather larger than the NLS estimates and abundantly above one, ranging between 2.4 and 2.6. However, the standard errors are quite large as well, and in this case we cannot reject that σ is equal to one. However, when we interpret energy as a measure of capital utilization, the 2SLS estimates of σ are both larger than the NLS estimates and significant larger than one. Hence, the estimates of specification (3.4) provide for evidence in favor of the existence of the Ricardo Effect. While point estimates of the other specifications are also supporting, they are not statistically strong enough. Regarding the other parameters, 2SLS estimates are generally lower than NLS estimates for α and γ, and larger for β.

Table 3.6: Reform effects

	(1)	(2)	(3)	(4)
	Pre-1990		*Post-1990*	
α	0.0695**	0.0025	0.1663***	0.0668**
	(0.0270)	(0.0033)	(0.0462)	(0.0265)
σ	1.7796***	3.5220**	1.4487***	1.6174***
	(0.1562)	(1.6343)	(0.0748)	(0.1175)
γ		0.4514***		0.2254***
		(0.0591)		(0.0345)
Obs.	24163	24163	29240	29240

Standard errors clustered at plant level in parentheses. $^{*}p < 0.10,^{**}p < 0.05,^{***}p < 0.01$.
Pre-1990: Sample includes observations until year 1990 included. Post-1990: Sample includes observations posterior to 1990. The values for the endogenous inputs are those predicted for at the first stage. All regressions include plant and time effects. To control for plant fixed effects, the specifications are estimated in first difference.

3.5.3 The impact of market-oriented reforms

As discussed in the previous sections, Colombia experienced a deep reform process between 1990 and 1993, aiming at liberalizing international trade and introducing flexibility in the labor market. We wonder whether these reforms caused a change in the production technology with respect to the substitutability of capital and labor. To verify this hypothesis, we simply split the sample at year 1990 and we repeat the 2SLS estimation of specification (3.4) over the sub-samples. Table 3.6 displays the findings. By comparing Column (1) to Column (3) and Column (2) to Column (4), we conclude that the elasticity of substitution between capital and labor did not significantly change in the years following the reforms. It is the case that the point estimate of σ in Column (2) is largely above the typical range of the coefficient estimates found in the previous estimations, but the standard error is also quite wide.

Hence, we cannot say that this estimate is statistically different from the others.

With these findings, we are not arguing that the market-oriented reforms did not affect the supply side of the Colombian economy. What Table 3.6 demonstrates is that the reforms did not change the manner in which capital and labor are optimally used in the production process, *for given input prices*. As we saw in Section 3.3, the relative cost of capital has fallen significantly during the period of our analysis, and in particular during the reform period. With a CES technology characterized by $\sigma > 1$, the reforms have accelerated the substitution of capital for labor and the relative decrease of the labor income share.

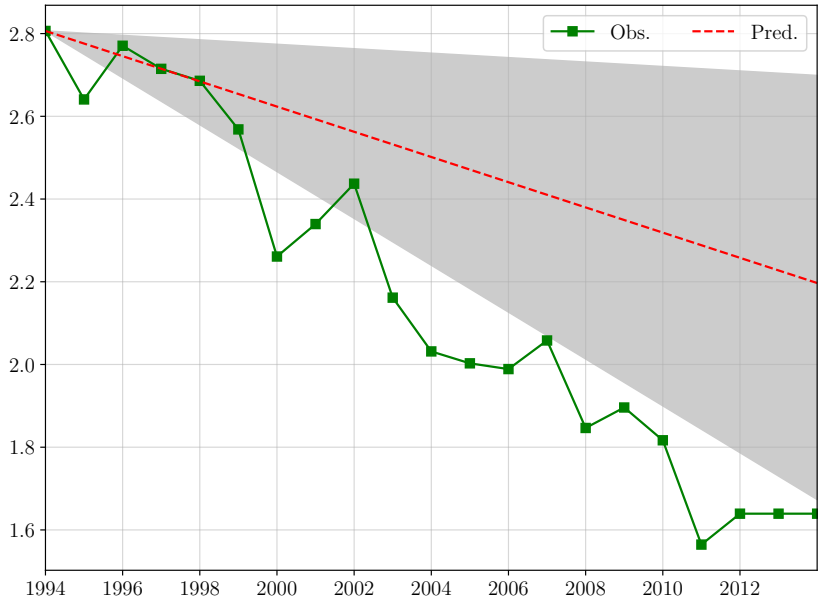
3.5.4 Simulation

The labor income share in Colombia has steadily declined between 1994 and 2014, from 74 per cent to 62 per cent of domestic income. How much of this decline can be accounted for by the Ricardo Effect? In the next paragraphs we try to answer this question. First of all, we need to make explicit some caveats regarding this exercise. We are facing limitations in data availability. On one hand, labor income share is available from 1994 only. On the other hand, our measures of input prices stop in 1998. We need to extrapolate the dynamics of the relative input price from the period of observation to the future, in order to account for the decline in the labor income share through 2014. Clearly, this requires strong assumptions of stability in the price dynamics. A second caveat has to do with the sectoral composition of the economy. While the information about the labor income share refers to the aggregate economy, the data used to estimate the CES technology's parameters belongs to the manufacturing sector only. All in all, this constraint is not as severe as it appears. In many developing countries, aggregate statistics face the weakness of missing the large existing informal sector. Informality is disproportionally present in agriculture and services, and less so in manufacturing. In this sense, aggregate statistics tend to be actually more representative of the manufacturing sector than the economy as a whole.

The simulation exercise we employ to assess the importance of the Ricardo Effect

is based on equation (3.1), expressing the relative income shares as a function of the relative input price, given the parameters of the CES technology. The relative price of capital R/W has fallen on average by 3.6 per cent per year between 1986 and 1998. We assume it kept falling at the same place through 2014. Given the estimated values of σ, we quantify the implication of this fall for the relative income shares. In

Figure 3.3: Relative income shares



Observed values from PWT 9.0 and simulated values from equation (3.1). The red dashed line is obtained setting $\sigma = 1.85$. The edges of the grey shaded area are obtained setting $\sigma = 1.15$ and $\sigma = 2.58$.

Figure 3.3, we plot the relative income share predicted through this exercise. The red dashed line is obtained by setting σ at 1.85, the value estimated in Column (4) of Table 3.5. The edges of the grey shaded area are obtained by setting σ at 1.15 and 2.58, the smallest and largest point estimates among all the different estimations performed. The green dotted line reports the actual observations of the relative income shares. In 2014, this aggregate was approximately equal to 1.6, about 60 per cent of its value in 1994. The simulation predicts a relative income share of 2.2 in

2014, or 80 per cent of the initial value. Based on this result, we conclude that the Ricardo Effect accounts for half of the decline in the labor income share observed between 1994 and 2014. If we consider the range of σ values from our estimations, the actual path of the relative income share is quite close to the one predicted if we use the largest of our point estimates of σ, and it is captured quite well until 2002. Notice that for the initial years the assumption that the relative price of capital was falling at 3.6 per cent per year is probably more accurate.

3.6 Conclusion

In this chapter, the ideas of the controversial chapter 31st of David Ricardo's master piece were tested using a unique plant level longitudinal database for the manufacturing sector in Colombia. After estimating the parameters of a CES production technology, we found that the elasticity of substitution between capital and labor is 1.85 and significantly larger than one. This proves the existence of the Ricardo Effect, *i.e.* the substitution of newer, cheaper capital for labor with the concomitant decline in the labor income share. The market-oriented reforms inaugurated in Colombia between 1990 and 1993 did not change the elasticity of substitution between capital and labor, but they induced a stronger fall in the relative cost of capital and, as a consequence, amplified the Ricardo Effect. Based on a simulation exercise, we conclude that the Ricardo Effect accounts for half of the decline of the labor income share in Colombia. As Samuelson (1989) claimed, "Ricardo was Right!"

This chapter is a positive analysis of the Ricardo Effect. Notice that no welfare consequences are addressed here. However, a class of interesting questions with welfare consequences arise from the evidence regarding labor replacement when new units of capital are demanded. Does the Ricardo Effect overall has overall positive of negative consequences for society? This remains an open question subject to further research.

This chapter estimated the Ricardo Effect only for the manufacturing sector in a developing country. Duarte and Restuccia (2010) illustrate that the labor share

in manufacturing sectors display an inverted U shape over time. It is possible that the Ricardo Effect provides an explanation for the slippery side of the labor share in manufacturing industries, but it is important to understand whether there is evidence of labor replacement in services. In particular, it would be interesting to consider whether the Ricardo Effect is an important mechanism of sectoral transformation in which labor is moving from one sector to another.

Last, a proper estimation of the dynamics of integration of the labor force in the manufacturing sector, taking into account the differences between managers and workers, could provide some light to the policy debate related with job creation through corporate tax stimulus towards investment in capital, a debate widely spread in the Colombian context.

REFERENCES

Acemoglu, D. (2002). Directed Technical Change. *The Review of Economic Studies*, *69*(4), 781–809.

Acemoglu, D., & Autor, D. (2011). Skills, Tasks and Technologies: Implications for Employment and Earnings. *Handbook of Labor Economics*, *4*, 1043–1171.

Aguiar, M., & Bils, M. (2015). "Has Consumption Inequality Mirrored Income Inequality?". *American Economic Review*, *105*(9), 2725–2756.

Aguiar, M., Chatterjee, S., Cole, H., & Stangebye, Z. (2016). "Quantitative Models of Sovereign Debt Crises". *Handbook of Macroeconomics*, *2*.

Aguiar, M., Chatterjee, S., Cole, H., & Stangebye, Z. (2017). "Self-Fulfilling Debt Crisis, Revisted: The Art of the Desperate Deal". *NBER Working Paper*, *23312*.

Aguiar, M., & Gopinath, G. (2006). "Defaultable debt, interest rates and the current account". *Journal of International Economics*, *69*, 64-83.

Angeletos, G.-M., & Lian, C. (2016). "Incomplete Information in Macroeconomics: Accomodating Frictions in Coordination". *Handbook of Macroeconomics*, *2*.

Arellano, C. (2008). "Defaul Risk and Income Fluctuations in Emerging Economies". *American Economic Review*, *98*(3), 690-712.

Baumol, W. J. (1967). "Macroeconomics of Unbalanced Growth: The Anatomy of Urban Crisis". *American Economic Review*, *57*(3), 415–426.

Beetsma, R., Giuliodori, M., Hanson, J., & de Jong, F. (2018). "Bid-to-cover and yield changes around public debt auctions in the euro area". *Journal of Banking and Finance*, *87*, 118-134.

Bloom, N., Sadun, R., & Van Reenen, J. (2012). "Americans Do IT Better: US Multinationals and the Productivity Miracle". *American Economic Review*, *102*(1), 167–201.

Blundell, R., & Bond, S. (1998). "Initial Conditions and Moment Restrictions in Dynamic Panel Data Models". *Journal of Econometrics*, *87*, 115–143.

Bocola, L., Bornstein, G., & Dovis, A. (2018). "Quantitative Sovereign Default Models and the European Debt Crisis". *NBER Working Paper*, *24981*.

Bocola, L., & Dovis, A. (2018). "Self-Fulfilling Debt Crisis: A Quantitative Analysis". *Unpublished Manuscript*.

Boppart, T. (2014). "Structural Change and the Kaldor Facts in a Growth Model with Relative Price Effects and Non-Gorman Preferences". *Econometrica*, *82*(6), 2167–2196. Retrieved from `http://dx.doi.org/10.3982/ECTA11354` doi: 10.3982/ECTA11354

Buera, F. J., & Kaboski, J. P. (2009). "Can Traditional Theories of Structural Change Fit the Data?". *Journal of the European Economic Association*, *7*(2-3), 469–477.

Buera, F. J., & Kaboski, J. P. (2012). "The rise of the Service Economy". *American Economic Review*, *102*(6), 2540-2569. Retrieved from `http://www.jstor.org/stable/41724664`

Bulow, J., & Rogoff, K. (1989). "A Constant Recontracting Model of Sovereign Debt". *Journal of Political Economy*, *97*(1), 155-178.

Burke, G., & Rumberger, R. W. (1987). *The Future Impact of Technology on Work and Education* (T. F. Press, Ed.). Philadelphia, PA: Falmer Press.

Calvo, G. (1988). "Servicing the Public Debt: The Role of Expectations". *American Economic Review*, *78*(4), 647-661.

Cette, G., Lopez, J., & Mairesse, J. (2016). "Market Regulations, Prices, and Productivity". *American Economic Review*, *106*(5), 104—108.

Chatterjee, S., & Eyigungor, B. (2012). "Maturity, indebtedness and default risk". *American Economic Review*, *102*(6), 26742699.

Cole, H. L., & Kehoe, T. J. (2000). "Self-Fulfilling Debt Crisis". *Review of Economic Studies*, *67*, 91-116.

Comin, D. A., Lashkari, D., & Mestieri, M. (2015). "Structural Change with Long-run Income and Price Effects". *NBER Working Paper No. 21595*.

Crafts, N. (2006). "Regulation and Productivity Performance". *Oxford Review of Economic Policy*, *22*(2), 186–202.

Diéz, J. M., Gaitán, S., & Valderrama, M. C. (2011). Aproximación a las metodologías de estimación del costo de capital en los proyectos de inversión. el caso colombiano. *Revista Escuela De Administración. Centro de publicaciones Universidad EAFIT*.

Duarte, M., & Restuccia, D. (2010). "The Role of the Structural Transformation in Aggregate Productivity". *Quarterly Journal of Economics*, *125*(1), 129–173.

Durdu, C. B., Nunes, R., & Sapriza, H. (2013). "News and sovereign default risk in small open economies". *Journal of International Economics*, *91*, 1-17.

Eaton, J., & Gersovitz, M. (1981). "Debt with Potential Repudiation: Theoretical and Empirical Analysis". *Review of Economic Studies*, *48*(2), 289-309.

Edwards, S., & Steiner, R. (2008). *La revolución incompleta: las reformas de gaviria*. Bogotá: Editorial Norma.

Eslava, M., Haltiwanger, J., Kugler, A., & Kugler, M. (2004). The Effects of Structural Reforms on Productivity and Profitability Enhancing Reallocation: Evidence from Colombia. *Journal of Development Economics*, *75*(2), 333 - 371.

European Commission. (2018). "Debt Sustainability Monitor 2017". *Institutional Paper 071*.

Giordani, P., & Söderlind, P. (2003). "Inflation forecast uncertainty". *European Economic Review*, *47*, 1037-1059.

Gu, G. W., & Stangebye, Z. R. (2018). "Costly Information and Heteroskedasticity in Sovereign Risk Spreads". *Unpublished manuscript*.

Hanson, R. (2001). Economic Growth given Machine Intelligence. *Journal of Artificial Intelligence Research*.

Herrendorf, B., Rogerson, R., & Valentinyi, A. (2013). "Two Perspectives on Preferences and Structural Transformation". *American Economic Review*, *103*(7), 2752–2789.

Herrendorf, B., Rogerson, R., & Valentinyi, A. (2014). "Growth and Structural Transformation". In *Handbook of Economic Growth* (Vol. 2, p. 855-941). Elsevier.

Inklaar, R., Timmer, M. P., & van Ark, B. (2008). "Market Services Productivity Across Europe and the US". *Economic Policy*, *23*(53), 150-194.

Karabarbounis, L., & Neiman, B. (2014). The Global Decline of the Labor Share. *The Quarterly Journal of Economics*, *129*(1), 61–103.

Keynes, J. M. (1930). Economic Possibilities for our Grandchildren. In *Essays in persuasion (2009)* (pp. 192–202). New York: Classic House Books.

Knights, D., & Willmott, H. (1988). *New Technology and the Labour Process*. Macmillan.

Kongsamut, P., Rebelo, S., & Xie, D. (2001). "Beyond Balanced Growth". *The Review of Economic Studies*, *68*(4), 869–882.

Krusell, P., Ohanian, L. E., Ríos-Rull, J.-V., & Violante, G. L. (2000). Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis. *Econometrica*, *68*(5), 1029–1053.

Kugler, M., & Verhoogen, E. (2012, January). Prices, Plant Size and Product Quality. *Review of Economic Studies*, *79*(1), 307-339.

Kuznets, S. (1957). "Quantitative Aspects of the Economic Growth of Nations: Ii. Industrial Distribution of National Product and Labor Force". *Economic Development and Cultural Change*, 1–111.

Lewis, W. W. (2005). *The Power of Productivity: Wealth, Poverty, and the Threat*

to Global Stability. University of Chicago Press.

Matsuyama, K. (2008). "Structural Change". In S. N. Durlauf & L. E. Blume (Eds.), *The New Palgrave Dictionary of Economics* (Vol. 2). New York: Palgrave Macmillan.

Mendoza, E. G., & Yue, V. Z. (2012). "A general equilibrium model of sovereign default and business cycles". *Quarterly Journal of Economics*, *127*(2), 889946.

Morris, S., & Shin, H. S. (1998). "Unique Equilibrium in a Model of Self-Fulfilling Currency Attacks". *American Economic Review*, *88*(3), 587-597.

Ngai, L. R., & Pissarides, C. A. (2007). "Structural Change in a Multisector Model of Growth". *The American Economic Review*, 429–443.

Nicoletti, G., & Scarpetta, S. (2003). "Regulation, Productivity and Growth: OECD Evidence". *Economic Policy*, *18*(36), 9—72.

O'Mahony, M., & Timmer, M. P. (2009). "Output, Input and Productivity Measures at the Industry Level: The EU KLEMS Database". *The Economic Journal*, *119*(538), F374–F403.

Patton, A. J., & Timmermann, A. (2010). "Why do forecasters disagree? Lessons from the term structure of cross-sectional dispersions". *Journal of Monetary Economics*, *57*, 803-820.

Reinhart, C. M., & Rogoff, K. S. (2011). "From Financial Crash to Debt Crisis". *American Economic Review*, *101*(5), 1676-1706.

Ricardo, D. (1821). *On the Principles of Political Economy and Taxation* (Third ed.). London: John Murray.

Sáenz, L. F. (2017). "Time-Varying Capital Intensities and the Hump-Shaped Evolution of Economic Activity in Manufacturing". *Unpublished Manuscript. University of Illinois*.

Samuelson, P. A. (1988). Mathematical Vindication of Ricardo on Machinery. *The Journal of Political Economy*, 274–282.

Samuelson, P. A. (1989). Ricardo Was Right! *The Scandinavian Journal of Economics*, *91*(1), 47–62.

Szkup, M. (2018). "Preventing Self-fulfilling Debt Crisis". *Unpublished Manuscript*.

Timmer, M. P., Inklaar, R., O'Mahony, M., & van Ark, B. (2011). "Productivity and Economic Growth in Europe: A Comparative Industry Perspective". *International Productivity Monitor*, *21*, 3–23.

Timmer, M. P., Vries, G. J., & de Vries, K. (2014). "Patters of Structural Change in Developing Countries". *GGDC Research Memorandum*, *149*. Retrieved from `http://www.rug.nl/research/ggdc/data/10-sector-database`

Urrutia, M., & Ruiz, M. (2010). Ciento setenta años de salarios reales en colombia. *Ensayos sobre política económica*, *28*(63), 154–189.

van Ark, B., Inklaar, R., & McGuckin, R. H. (2003). "ICT and Productivity in Europe and the United States: Where do the Differences Come From?". *CESifo Economic Studies*, *49*(3), 295–318.

von Hayek, F. A. (1942). The Ricardo Effect. *Economica*, *9*(34), 127–152.

APPENDIX A

ADDITIONAL MATERIAL FOR CHAPTER 2

A.1 Data sources

Table A.1: Data sources for sectoral physical and ICT capital

	Avail.	*Source*	*Time Horizon*
Austria	Y	EU Klems 2016 (Real Fixed Capital Stock)	1995-2009
Belgium	N		
France	Y	EU Klems 2016 (Real Fixed Capital Stock)	1978-2009
Germany	Y	EU Klems 2016 (Real Fixed Capital Stock)	2000-2009
Italy	Y	EU Klems 2016 (Real Fixed Capital Stock)	1995-2009
Netherlands	Y	EU Klems 2016 (Real Fixed Capital Stock)	2000-2009
Spain	Y	EU Klems 2016 (Real Fixed Capital Stock)	1970-2009
United Kingdom	Y	EU Klems 2016 (Real Fixed Capital Stock)	1997-2009
United States	Y	OECD (Net Fixed Capital Stock, Volumes)	1970-2009

A.2 Measurement of Sectoral Labor Productivity in Each European Country

Figures A.1 and A.2 plot the productivity levels for each sector in each country relative to the United States for the first and last sample periods. Figure A.1 shows three different patterns for agriculture, manufacturing and services. First, the agricultural productivity levels (relative to the U.S.) were either stagnant or relative higher in 1970 compared to 2009 with the exceptions of France and Germany, where minor improvements were experienced. The productivity levels are surprisingly high for

the United Kingdom, but still they show an important fall in relative productivity between 1970 and 2009. However, these differences are do play a minor role in the aggregate labor productivity because the structural transformation has reduced the agricultural labor shares dramatically for each of these countries during our sample period.

Second, European countries have been catching up with the U.S. from 1970 to 2009 in manufacturing productivity without exception, although no country reached the U.S. labor productivity during our sample period. Whereas Austria, Belgium, France and the Netherlands experienced about a two-fold increase in manufacturing productivity, the productivity growth in manufacturing was more modest in Germany, the United Kingdom, Italy and Spain.

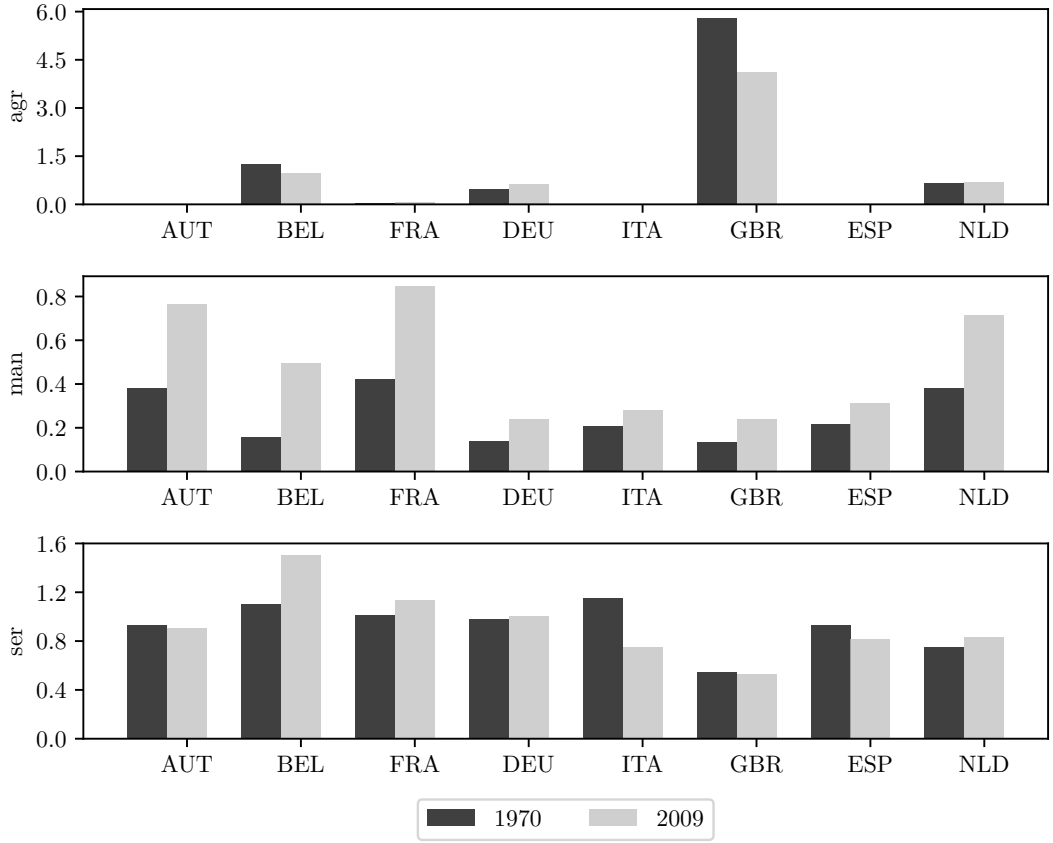
Last, with the notable exception of Belgium, no European country experienced a significant catch up in services relative to the U.S.; most countries have remained either stagnant or have experienced a decline.

Figure A.2 plots the relative labor productivity between 1970 and 2009 for each sector within services and for each European country. Within services, European countries are in generally more productive than the U.S. in telecommunications, education, and health services[1], but they are significantly less productive in wholesale and retail trade. Moreover, the productivity levels for this sector have widen out between 1970 and 2009 in every single European country.

The sector of business services in Europe is also less productive compared to the U.S. without exception, although the productivity gaps have not widened in every country. For instance, Germany and Belgium did not experienced a fall in the relative

[1]It is interesting to note that health services are much less productive in the U.S. than in Europe. In addition, productivity gap widened significantly during the sample period. The labor productivity in this sector is a source of major concern for the U.S. as it employed approximately 17 percent of the labor force in 2009. Nevertheless, the finding that Europe is more productive than the U.S. in health services, as well as in education, should be taken with some caution. Whereas in the U.S. both education and health are services mainly provided by the private sector, in most European countries education and health systems are managed by the government, and the labor hired in these two sub-sectors qualifies as public employment. This fact raises potential concerns on the extent of comparability of sectoral productivity in education and health between Europe and the U.S., even though we use our model to correct potential measurement biases in the available data.

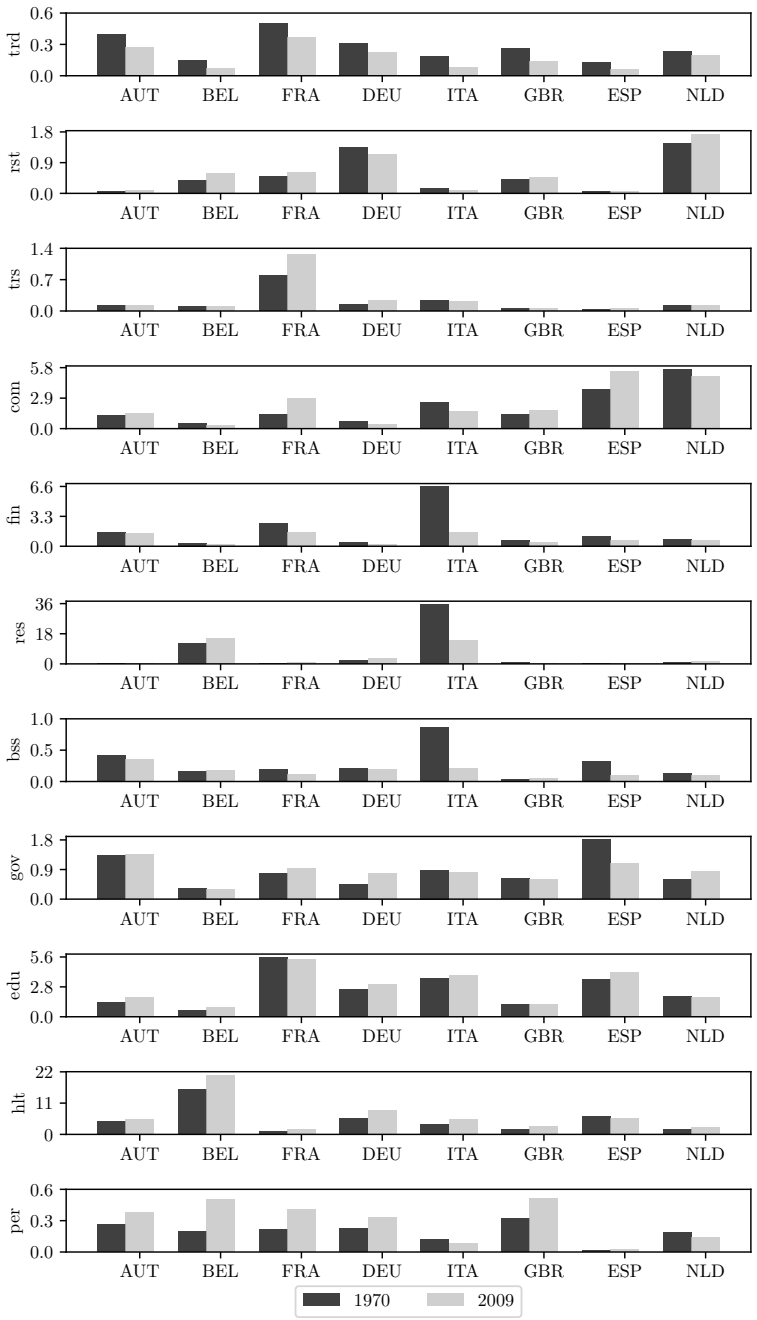
Figure A.1: Sectoral productivity levels



Recovered sectoral labor productivity levels in 1970 and 2009 for major European countries relative to the sectorial U.S. labor productivity level. Agriculture and manufacturing.

productivity, but Italy on the other hand experienced a dramatic increase in the productivity gap between 1970 and 2009 in business services relative to the U.S. The employment shares of these two sectors have been relatively large in the years of our study, hence, the levels of labor productivity in wholesale and retail trade and in business services do matter significantly for the differences in aggregate productivity between Europe and the U.S. For the rest of the service sectors the evidence is mixed across countries. An important case to highlight is financial services. Austria, France, Italy and Spain were countries with more productive financial services compared to

Figure A.2: Sectoral productivity levels



Recovered sectoral labor productivity levels in 1970 and 2009 for major European countries relative to the sectorial U.S. labor productivity level. Services.

the U.S. in 1970, and in spite of the sharp drop in productivity, they were still more productive in 2009, except for the case of Spain. Nevertheless, without exception, all countries in Europe experienced an important reduction in their productivity relative to the U.S. in financial services.

A.3 Findings of the Counterfactual Experiments for Each European Country

Table A.3 illustrates our findings when we feed the labor productivity growth rates from 1970 to 2009. The top panel of Table A.3 show the results of this exercise when a European country counterfactually experiences the observed labor productivity growth rate in the U.S., in order to assess changes in aggregate labor productivity as a consequence of changes in the productivity of a single sector. Each row of the top panel represents one of the 13 sectors in our model economy, and each column represents a European country with the exception of the last column, which represents Europe as a weighted average of the countries in our European sample.

The results for agriculture are not conclusive. Whereas some countries would have performed better such as Belgium and the Netherlands, for the rest of the European countries our model predicts that the aggregate labor productivity level would be actually lower. Nevertheless, with the exception of the Netherlands, these results have minimal implications for aggregate productivity.

On the other hand, the message for manufacturing is not ambiguous. Had the European countries experienced the U.S. labor productivity growth in manufacturing during our sample period, their aggregate labor productivity in 2009 would be lower regardless of the country. Naturally, Europe as a whole would have had a lower aggregate productivity. The upper bound of this decline is Italy, with a predicted drop of 4.2%, whereas the lower bound is Belgium, with an staggering drop of 14.7%. Manufacturing is not responsible for the European underperformance *vis-á-vis* the U.S. On the contrary, it helped Europe in its path toward convergence during our sample period.

With regards to services, our counterfactual experiment suggests that the slowdown in the aggregate labor productivity comes mainly from three sectors: wholesale and retail trade (trd), financial services (fin) and business services (bss). It also suggests that Europeans are significantly more productive in health services (hlt). Let's discuss the results of each of these four sectors in detail (for the rest of the sectors the results are ambiguous depending on the country, and the aggregate effect on labor productivity is not large).

First, during the sample period, the aggregate labor productivity in every single European country would have increased significantly had the wholesale and retail trade sector experienced the U.S. labor productivity growth in Europe. The lower bound for this prediction is for Great Britain, with an increase in aggregate labor

Table A.2: KEEPING THE U.S. PACE

	AUT	BEL	FRA	DEU	ITA	GBR	ESP	NLD	Europe
Counterfactual: $\gamma_i = \gamma_i^{USA}$									
agr	-0.1	1.3	-1.6	-1.2	-0.2	-0.2	-1.5	6.3	0.4
man	-11.0	-14.3	-12.2	-7.6	-4.2	-9.2	-4.9	-7.7	-8.9
trd	4.3	6.5	3.9	3.4	7.7	1.8	6.3	7.1	5.1
rst	-0.0	-0.1	-1.6	-1.2	0.1	-0.4	-0.3	-0.3	-0.5
trs	-1.5	-1.0	-0.5	0.6	0.9	-0.6	0.6	-0.2	-0.2
com	-1.2	1.5	-7.0	2.0	4.8	2.6	-7.0	-2.5	-0.9
fin	0.6	1.0	3.7	2.4	17.7	1.0	2.2	1.9	3.8
res	-0.3	-0.7	-0.3	-0.7	4.8	-0.5	0.4	0.3	0.4
bss	1.3	-0.2	4.0	0.6	11.7	2.6	6.6	-2.7	3.0
gov	-0.2	0.4	-1.0	-2.8	0.5	-1.8	3.5	0.4	-0.1
edu	-1.7	-1.4	0.5	-1.4	-0.9	0.4	-1.7	-0.1	-0.8
hlt	-3.7	-10.9	-5.9	-10.1	-7.3	-4.7	3.8	-8.4	-5.9
per	-1.0	-2.3	-1.8	-1.1	1.2	0.9	-1.5	-1.3	-0.9
$\gamma_i = \gamma_{i,i\in\text{services}}^{USA}$	-3.5	-7.7	-6.4	-8.6	44.7	1.2	13.3	-6.1	3.4
$\gamma_i = \gamma_{i,\forall i}^{USA}$	-14.3	-20.4	-19.5	-16.8	38.6	-8.4	5.9	-7.5	-5.3

Counterfactual sectoral productivity growth is the one experienced by the U.S. between 1970 and 2009.

productivity of 1.8%, whereas the upper bound is Italy with an increase of 7.7%. The prediction for Europe indicates that this sector alone would have been responsible for an aggregate labor productivity 5.1% higher than our benchmark prediction in 2009.

Second, financial services also would have helped to reduce the labor productivity gap had the European countries experienced the same labor productivity growth observed in this sector for the U.S. Europe as a whole would have had a labor productivity level 3.8% higher than our benchmark prediction. Furthermore, every single European country would have experienced higher aggregate labor productivity if their financial services were as productive as in the U.S., although the results for Italy are substantially higher to the rest of Europe.

Third, with the exception of Belgium and the Netherlands, the labor productivity would also be higher for the European countries if they have had the U.S. labor productivity growth in business services. Once again, the order of magnitude of this result is substantially higher for Italy compared to the rest of Europe.

Last, our results also illustrate that Europe would have had lower aggregate productivity have they had the U.S. labor productivity growth observed in health services. With the exception of Spain, every single European country would have underperformed have they had the U.S. labor productivity growth in the health sector.

Table A.3 shows the results of the numerical experiments for the period 1990-2009 by comparing the benchmark prediction to the counterfactual aggregate labor productivity. Among several differences with respect to our previous counterfactual, we would like to highlight that the results for health services are in the same direction compared to the entire sample period, but the order of magnitude of the result is about half of what it was for the 1970-2009 period, although still represent a large distance between the benchmark and the counterfactual aggregate productivity for each country, again with the exception of Spain. In addition, for the period between 1990 and 2009 a new sector emerges in which the Europeans would be worse off if they have had the U.S. labor productivity growth: Communications. With the exception of Belgium, all countries in Europe would have had lower aggregate productivity have they had the U.S. labor productivity in communications, and this difference is

Table A.3: Taking off with the U.S.

	AUT	BEL	FRA	DEU	ITA	GBR	ESP	NLD	Europe
Counterfactual: $\gamma_i = \gamma_i^{USA}$									
agr	-0.3	1.3	-0.8	-1.5	-0.2	0.3	-0.2	-1.0	-0.3
man	-2.8	-1.5	-1.4	-1.3	3.1	0.9	2.5	-1.7	-0.3
trd	4.0	4.3	4.2	2.3	4.2	1.6	3.0	2.4	3.2
rst	0.5	0.2	-0.1	-0.7	0.1	-0.2	0.7	-0.1	0.0
trs	-0.2	-0.1	0.1	0.6	0.3	0.6	0.8	-0.4	0.2
com	-0.9	1.0	-4.2	-0.1	-3.3	-6.5	-3.0	-5.9	-2.9
fin	-1.0	-0.5	1.7	1.7	1.0	0.7	0.3	-0.7	0.4
res	0.2	2.2	-0.0	0.0	2.6	0.3	0.9	0.4	0.8
bss	0.9	2.2	2.6	3.1	4.5	2.8	3.7	-0.8	2.4
gov	0.1	0.1	-0.5	-0.8	-1.1	-0.5	-0.4	-0.6	-0.5
edu	-0.4	0.4	0.9	0.2	0.1	1.5	-2.4	0.4	0.1
hlt	-1.9	-3.4	-1.2	-8.1	-1.6	-0.2	0.4	-5.7	-2.7
per	-0.5	-0.7	-0.8	-0.4	0.4	-0.3	-0.4	-0.7	-0.4
$\gamma_i = \gamma_{i,i\in\text{services}}^{USA}$	0.7	5.7	2.6	-2.3	7.4	-0.3	3.6	-11.6	0.7
$\gamma_i = \gamma_{i,\forall i}^{USA}$	-2.4	5.4	0.4	-5.2	10.6	0.9	6.0	-14.2	0.2

Counterfactual sectoral productivity growth is the one experienced by the U.S. between 1990 and 2009.

Table A.4: Catching Up with the U.S.

	AUT	BEL	FRA	DEU	ITA	GBR	ESP	NLD	Europe
Counterfactual: γ_i s.t. $A_i = A_i^{USA}$									
trd	19.0	30.6	15.0	22.1	33.8	22.7	33.8	29.4	25.8
bss	10.3	13.4	17.7	15.1	13.3	24.2	15.3	27.9	17.1
fin	-1.9	4.6	-2.3	5.8	-2.5	2.3	1.8	4.3	1.5

Counterfactual sectoral productivity growth is the one insuring full catch-up of the U.S. by the sector in 2009.

large in France, Italy, Great Britain, Spain and the Netherlands.

Table A.4 shows the implied change in aggregate productivity when the labor productivity in wholesale and retail trade, business services and financial services

converges to the U.S. labor productivity level in 2009. No European country would have experienced a reduction of its observed aggregate labor productivity have their labor productivity converged to the U.S. by 2009 in either wholesale and retail trade or in business services. Whereas the lower bound of the prediction is of 15% if France have had a catch up in whole sale and retail trade, the lower bound of the increase in aggregate labor productivity is of 10.3% for Austria have they experienced a catch up in business services.

On the other hand, financial services are not unambiguously a source of slowdown between Europe and the U.S. The last row of Table A.4 shows that have Europe experienced a full catch up in the labor productivity of financial services relative to the U.S. 2009 level, Austria, France and Italy would have had lower aggregate labor productivity. Moreover, even Germany – the most successful counterfactual scenario with an aggregate productivity 5.8% higher compared to its 2009 benchmark prediction – falls short when compared to the lower bound of the predictions for wholesale and retail trade or for business services.