



HAL
open science

Essays on the impact of international trade and labor regulation on firms

Gabriel Smagghue

► **To cite this version:**

Gabriel Smagghue. Essays on the impact of international trade and labor regulation on firms. Economics and Finance. Institut d'études politiques de paris - Sciences Po, 2014. English. NNT: 2014IEPP0022 . tel-03512807

HAL Id: tel-03512807

<https://theses.hal.science/tel-03512807>

Submitted on 5 Jan 2022

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Institut d'Etudes Politiques de Paris
ECOLE DOCTORALE DE SCIENCES PO
Programme doctoral en économie
Docteurat en Sciences économiques

Essays on the Impact of International Trade and
Labor Regulation on Firms

Gabriel SMAGGHUE

Thèse dirigée par

Thierry Mayer, Professeur des Universités, Sciences Po Paris

Soutenue à Paris, le 3 Juillet 2014

Jury :

M. Holger Breinlich - Professor of Economics, Essex University
M. Lionel Fontagné - Professeur des Universités, Université Paris 1 - *Rapporteur*
M. Philippe Martin - Professeur des Universités, Sciences Po Paris
M. Thierry Mayer - Professeur des Universités, Sciences Po Paris
M. Marc Melitz - Professor of Economics, Harvard University - *Rapporteur*

Remerciements

Une thèse est un objet vivant fait d'espoir, de doute, d'effort, de plaisir, de souffrance, de progrès, de régressions (beaucoup de régressions!). Ma thèse, quant à elle, aura eu une belle et longue vie. J'aimerais ici remercier ceux qui y ont contribué.

Mes premiers remerciements vont à ceux qui m'ont encadré et ce faisant ont participé à me faire aimer la recherche. Merci à Lionel Fontagné pour m'avoir accompagné en douceur dans mes premiers pas de chercheur en commerce international. Merci à Farid Toubal pour son soutien inconditionnel et amical depuis le début. Ma gratitude va évidemment aussi à Thierry Mayer pour son encadrement efficace et pour m'avoir toujours encouragé à élargir mes champs de recherche. Merci enfin à Marc Melitz pour avoir participé à rendre mon année à Harvard si enrichissante. Son amour du français, sa gentillesse et sa simplicité ont été précieux pour me mettre en confiance dans ce contexte intimidant!

Si la thèse est parfois vue comme une entreprise solitaire, j'ai eu la chance de l'agrémenter de plusieurs collaborations de recherche. Ici je veux remercier mes coauteurs et amis, Julien, Mathieu et Paul. Les gars, j'espère que ça a été autant un plaisir pour vous que pour moi. Mais surtout j'espère qu'il y a beaucoup d'autres papiers à venir.

J'ai ensuite une pensée pour tous mes collègues doctorants de Sciences Po avec qui j'ai pu partager un café, une référence d'article, des doutes et surtout beaucoup de bons moments. Dilan, ma "godmother", merci pour ton coaching permanent. Je ne te reprocherai que de ne pas m'avoir attendu pour finir ta thèse. Marion, merci pour ta bonne humeur, tes nombreuses initiatives pour la collectivité. Bravo pour tes astuces déployées pour contourner mes silences. Merci à Thomas d'avoir toujours suppléé Marion sur ce dernier point. Merci à Camille de me préparer un canapé pour m'accueillir à Barcelone. Merci à Maxime pour avoir dérivé tant de modèles économétriques directement sous mes yeux. Tu m'apprendras? Merci aux autres collègues de job market, Gong, Sarra pour ne pas avoir diffusé les vidéos de la formation. François, merci de venir à ma thèse, c'est un ordre! Merci à Lisa pour son écoute et sa bienveillance. Merci à Liza d'avoir fait avancer nos droits en défiant les horaires de fermeture de Sciences po. Merci à Jérôme pour cette riche année passée à Somerville au son des "gladiateurs".

Merci aux plus jeunes, Camilo, Jamil, Emanuele, Clément, d'avoir stimulé les plus vieux en étant aussi brillants. Merci enfin aux ATER et autres post doc, Mathieu, Morgane, Xavier, de nous faire sentir moins vieux.

Durant ma dernière année de thèse, j'ai eu la chance de visiter le département d'économie d'Aix-Marseille. En sus du soleil, j'ai pu profiter d'un environnement de travail paisible, dynamique et amical. Ces atouts ont été particulièrement appréciables en temps de job market. Je veux donc remercier Pierre-Philippe et Cécilia pour avoir rendu cela possible. Merci à Guy Laroque de m'avoir soutenu dans ma candidature. Merci beaucoup à Marc de m'avoir aidé à figoler mes présentations et pour m'avoir secoué lorsque les deadlines approchaient. Merci à Bruno et Tanguy pour nos déjeuners conviviaux. Merci à Linas, mon co-bureau, pour tous ces kilomètres parcourus. Merci aussi aux colocs d'Aix d'avoir été aussi faciles à vivre.

Une thèse est une course de fond. Merci aux amis pour l'énergie transmise tout au long du parcours. Merci aux amis d'enfance, Claire, Pierric, Pierre, Nico. Vous voir me redonne toujours un peu de mon âme d'enfant. Merci aux amis du lycée, Mehdi, Arnaud, Jean, Mériam, Sadia, Coralie. Merci d'avoir suivi tous ces parcours différents qui sont autant de richesses. Merci aux amis de l'ENS, Jérémie, Thomas, Sebastien, Maxime. On se voit au mariage. Merci aux collègues d'athlé, Thomas, Jibril, Arthémon, Adel, pour ces parenthèses qu'étaient les entraînements à la fin des dures journées de recherche. Merci aux amis de Cambridge, Bo, Jack, Nilton, Mitra, Vania. A quand le prochain séjour dans le New Hampshire?

Je veux dédier cette thèse à ma famille. Son soutien et son amour inconditionnels ont été indispensables à mon équilibre.

Merci enfin à Yohan qui me rend plus fort chaque jour.

Contents

	Page
General Introduction	vii
A Exporting firms: a minority of major importance	viii
B Thesis' questions, conclusions and contributions	xi
C A closer look at the thesis	xvi
C.1 The quality of products exported by French firms: measurement and response to low-cost competition	xvi
C.2 Do firms reallocate sales across destinations in response to unexpected demand shocks?	xviii
C.3 Size dependent labor regulation and capital labor substitution at the micro and macro level	xx
1 A new Method for Quality Estimation using Trade Data: An Application to French firms	1
1 Introduction	3
2 An Empirical Model with Demand for Quality	6
2.1 Consumer Preferences	6
2.2 Firms' Demand and Quality as a Demand Shifter	8
3 Quality Estimation Strategy	9
3.1 Dealing with Price Endogeneity	10
3.2 Discussion of the Identification	13
4 Data and Demand Estimation Results	15
4.1 Data	15
4.2 Estimation procedure	18
4.3 Pooled Industries Results	19
4.4 Demand Estimation by Industry	22
5 Analysis of Estimated Quality	23
5.1 A consistency test: Estimated versus Expert Assessed Quality	25
5.2 Correlation with firm-level variables	25
5.3 How well do Prices proxy for Quality?	26
6	27

6.1	Identification strategy	27
6.2	Results	29
7	Conclusion	31
	Appendices	33
A	A Simple Model of Endogenous Quality with Imported Inputs	33
A.1	Technology	33
A.2	Prices, Quality, Import Shares and RER's	35
A.3	Other Implications of the Model	38
B	Data Trimming	39
C	Descriptive Statistics	40
D	Robustness checks	41
E	List of Low-Wage Countries	42
2	The Interdependence of Firms' Export Decisions	43
1	Introduction	45
2	A Simple Model of Exporting Firm with Demand Uncertainty	48
3	Empirical Methodology: A Structural Diff-in-diff Approach	51
4	Data	53
4.1	Why Champagne during the 2000-2001 US crisis?	53
4.2	Data Sources	54
5	Stylized Facts	57
6	Results	59
6.1	Baseline Results	59
6.2	Discussion and Robustness	61
7	Conclusion	65
	Appendices	67
A	Additional Graphs	67
B	Deriving the Structural Expression of the Treatment Effect	67
3	Size-dependent Regulation and Factor Income Distribution	71
1	Introduction	73
2	Regulatory Background and Motivating Evidence	75
2.1	The French Regulation	75
2.2	Data	76
2.3	Motivating evidence	78
3	Theory: Heterogeneous Firms and Legal Threshold	81
3.1	The Set-up	82
3.2	The problem of the firm	83
3.3	Empirical Predictions	87
4	Estimation Strategy	93

4.1	Estimation of Counterfactual Factor Demands	93
4.2	Estimation of Capital Intensity Response	95
4.3	Mapping Counterfactual Factor Demands to Structural Parameters . .	95
5	Results	96
6	Counterfactual Analysis	98
6.1	The Structure of the General Equilibrium	98
6.2	Results	101
7	Conclusion	101
	Appendices	103
A	Appendix Tables	103
B	Proof Proposition 4	103
4	General Conclusion	105
	References	108

General Introduction

A Exporting firms: a minority of major importance

The analysis of firms' exporting activity has been very active over the last two decades. Thanks to access to new data sources, we now know much better which firms export and how international trade impacts them. This burgeoning line of research is unified around a simple and far-reaching conclusion: the behavior of exporting firms has central implications for macroeconomics and international trade. Let me detail that conclusion.

Exporting firms and the business cycle Exporting firms exhibit a number of outstanding features. Knowledge of these features is necessary to understand how individual exporters may influence aggregate outcomes. In this sense, perhaps the most important attribute of exporting firms is their size: exporters are larger than non-exporters. Bernard & Jensen (2004) find for instance that in the US in 2002, exporters on average have a workforce three times larger than non-exporters.¹ However, for exporting firms to have an aggregate impact, not only should they be large but also sufficiently numerous. And since few firms export (e.g., Bernard & Jensen (2004) report that 4% of US firms export in 2002), it is not clear a priori whether this last condition is verified. Gabaix (2011) provides us with precious insights on that point. He derives analytical conditions on the size distribution of firms under which large firms' fluctuations impact the business cycle. Then he shows that real world firm size distribution meets these conditions. In other words, he concludes that large firms are numerous enough to induce aggregate fluctuations. Since most of the largest firms export (In France, I find that 95 firms of export among the top 100 firms in terms of domestic sales), Gabaix's result imply that macroeconomic questions can be clarified by looking at the behavior of exporting firms.

The size premium of exporters also has implications for the export performance of an economy. Namely, it suggests that firms must reach a critical size to be able to export. Based on this insight, common wisdom often points to cross-country differences in firm size distribution to explain differences in the aggregate success on export markets. Many studies for instance explain Germany's outstanding export performance over past decade by the big share of large SME's in its economy.² It results that the evaluation of policies which potentially distort the size distribution of firms (e.g., any taxation scheme progressive with firm size) should take into account implications for international trade.

Exporting firms and the aggregate productivity Another remarkable feature of exporting firms is their productivity. For the US, Bernard & Jensen (2004) report that, conditional on size, exporters have a 10% larger labor productivity. This finding is consistent throughout countries. Intuitively, the fact that exporting firms are more productive opens

1. See Mayer & Ottaviano (2008) for analogous facts on European firms.

2. See for instance Fontagné et al. (2008) for a thorough comparative analysis of French and German export performance. Cancé (2009) also refers to the deficit of medium-sized firms and the excess of micro-sized firms in France relative to Germany to explain why Germany fares better on export market.

the possibility that a reallocation of production factors towards them is inductive of growth. This insight has major policy implications. In a seminal paper, Melitz (2003) shows in theory how, in the presence of productivity heterogeneity across firms within an industry, trade liberalization generates a gain of aggregate productivity. In the model, this effect arises because opening to trade leads non-exporting (and unproductive) firms to shrink due to foreign competition, while productive firms expand by entering foreign markets. This trade-induced reallocation is consistent with empirical findings from trade liberalization episodes, as evidenced by Pavcnik (2002) and Trefler (2004). Melitz (2003)'s result is remarkable as it points out to a new source of gain from trade, on top of the gain from sectoral specialization and the gain from access to new varieties analyzed respectively in the classical theory of trade and in the "new" trade theory (Krugman 1980, Helpman 1981).

The literature on the misallocation of resources suggests that the implications of the exporter productivity premium extend beyond trade policies. In an influential contribution, Hsieh & Klenow (2009) show that any policy which creates heterogeneity in the prices faced by individual producers (induced for instance by firm-specific tax rates) can generate a misallocation of resources across firms and thus lead to a sizeable decrease in output and measured total factor productivity (TFP). Restuccia & Rogerson (2008) complete this result by showing that this loss of TFP is larger when most taxed firms are also the most productive ones. A corollary to these findings is that any policy taxing exporting firms relatively more should foster a large loss of aggregate efficiency because exporting firms are precisely the most productive. Conversely, any policy subsidizing exporting firms could make the exporting sector grow at the cost of a minimal aggregate TFP loss because the induced misallocation of resources would be compensated by the fact that exporting firms are relatively more productive.

Exporting firms and inequalities The singularity of exporting firms also shapes the distribution of welfare. This is in particular the case because exporting firms supply higher quality goods.³ Since the production of such goods is intensive in skilled workers Verhoogen (2008), policies favoring exporting firms may increase the skill premium. The quality premium of exporters may also be determinant for cost of living inequalities across consumers. Faber (2012) describes how trade liberalization, by lowering the cost of importing (high quality) inputs from foreign firms, lowers the relative price of high quality final goods. Since the expenditure share of quality goods is increasing with income⁴, trade liberalization disproportionately benefits rich consumers.

Another reason why exporting firms are determinant for inequalities is because they pay higher wages, *conditional on workers' observable skills*. The exporters' wage premium is documented in the cross-section of firms by Eaton, Kortum, Kramarz & Sampognaro (2011). In a theoretical framework tailored to replicate this empirical fact, Helpman et al. (2010) show

3. Evidence on the quality premium of exporters can be found, inter alia, in Verhoogen (2008), Baldwin & Harrigan (2011) and Crozet et al. (2012).

4. There is abundant evidence on the non-homotheticity of preferences in terms of quality. See for instance Faber (2012) and Handbury (2012).

that trade liberalization increases within-skill group wage inequalities by inflating the weight of exporting firms, which pay the highest wages, in the economy.

Exporting firms and international shock transmission How does trade liberalization interconnect national economies? This question is essential, *inter alia*, to predict the economic path of a country which decides to open up to trade. Part of the answer lies in the behavior of exporting firms. Exporters are indeed at the interface between home and foreign economies. As they respond to changes in foreign markets, they act as a gateway through which external shocks diffuse at home.

An example is technology adoption. Because they are exposed to the technology used by their competitors on foreign markets, or because their trade partners provide them with blueprints, exporting firms may indeed facilitate the international diffusion of new technologies (Grossman & Helpman 1991, Westphal 2002).

Exporting firms also play a role in the international diffusion of demand shocks. When foreign economies slow down, firms' revenues abroad shrink and the domestic product falls. This traditional view implies that exporting firms synchronize international business cycles. However, a recent and fast-growing literature claims that the way firms respond to local demand shocks might be more complex. In fact, there is evidence that firms' sales are interdependent across markets. In other words, it seems that firms adjust their sales to their different destination markets as soon as a single one is hit by a demand shock. Although this literature is not yet unified as to the sign or the origin of such interdependence at the firm-level, it is consistent about the following message: more work should be done on the behavior of exporting firms to understand international business cycle co-movements.

New insights on the origin and functioning of exporters: a simple structural approach The empirical findings described above demonstrate the substantial progress achieved recently in our knowledge of exporting firms. Because these findings came with first order aggregate implications, they call for enhanced efforts to further fathom exporters. My dissertation aims to contribute to this research agenda. More specifically, my work is organized around two broad questions.

First, *how do exporting firms adjust to shocks in foreign markets?* By nature, exporters and non-exporters are exposed to different problems. Exporters' problems are more complex because, unlike non-exporters, they need to take decisions regarding many destination markets simultaneously. Problems posed to exporters are also more unstable since they are hit by shocks in many markets.

In workhorse trade models though (see, for instance, Krugman 1980, Melitz 2003), an exporting firm behaves very much like a collection of independent firms who each serves a single market. Decisions of a firm across markets are only tied together through firms' attributes (e.g., productivity) which are exogenous to shocks faced by firms in their destination markets (e.g., demand shocks, entry of new competitors). In that sense, these models have a limited

ability to convey the specific challenges faced by exporters.

My purpose is to empirically assess existing theory by looking at the way exporters adjust to shocks in foreign markets. Namely, in chapter 1, I test whether firms adjust the quality of their exports when low cost competition gets fiercer in foreign markets. In chapter 2, I investigate the way exporters adjust their sales and prices to their different destinations when an unexpected demand shock hits a single market.

Second, *what is the origin of large (exporting) firms?* If the literature makes it clear that exporting firms are a very special sub-sample of the population of firms, we still know little about the origin of these specificities. Why is it that some firms are big enough, productive enough or supply goods with high enough quality to become exporters? The answer to this question has to do with the origin of firms' performance in general and does not necessarily involve trade. In chapter 3, I empirically analyze the micro and macro impact of size-dependent labor regulation. By shifting up the cost of large firms, these regulations are suspected of hindering firms' growth and to distort firms size distribution.

Although this dissertation is mostly empirical, it heavily relies on theory. My use of theory serves two purposes. First, theory guides the identification of the parameters of interest. In some cases, it gives an economic sense to the statistical assumption underlying the estimation. In others, it even drives the choice of the specification to be estimated. Second, theory provides an economic interpretation to my statistical results. The fact that the output from my estimations correspond to clearly defined economic concepts makes it possible to incorporate them into rich theoretical frameworks to better assess the scope of my results.

B Thesis' questions, conclusions and contributions

I now present the main questions, conclusions and contributions which emerge from this dissertation.

Main questions

1. Recent literature has established that product quality is essential to explain international trade patterns, whether at the micro or macro level. The growing focus on the quality dimension of trade has also shed new light on the impact of trade on inequalities. The importance of these findings calls for a thorough understanding of quality determinants at the firm-level. However, this research agenda faces a substantial challenge: quality is unobserved. In chapter 1, jointly written with Paul Piveteau (PhD candidate at Columbia University), we propose a new solution to estimate the quality of products at the firm-level using trade data.
2. How do exporting firms coordinate their sales across markets? How do firms rebalance their sales across their destination portfolio when one destination is hit by a demand shock? Most of the "new trade theory" models assume that firms maximize their profits independently in each country they serve. Therefore, the answer to these questions

matters for our understanding of exporting firms. It is also potentially of importance to understand how exporting firms shape the international diffusion of shocks. In chapter 2, jointly written with Julien Martin (UQAM) and Mathieu Parenti (UCL), we empirically test the interdependence of firms' export decisions.

3. Labor regulations which are more binding to large firms than to small firms are widespread. An illustration is the obama care which constrains employers with more than 50 employees to purchase health insurance for their workers. These regulations are often suspected to be inefficient because they deter firms from growing. In chapter 3, I structurally estimate the size response of firms to the elements of the regulation which bind to firms above 50 employees in France. I also look at the way firms adjust their capital-labor regulation in response to the regulation. Finally, I assess the way these firm-level responses feedback at the aggregate level through the aggregate productivity and the factor income distribution.

Main conclusions

1. The first conclusion is methodological. The quality of exports at the firm-level can be obtained from the residual of an export demand function, once prices have been controlled for. The estimation of this demand function requires information on four variables: prices and sales at the firm-product level, imports by firm-source country and real exchange rates by country (herein RER). The interaction of firm-specific importing shares with RER on imports delivers a firm-specific cost shifter which can be used to instrument prices in a demand function at the firm-level. The standard diagnostics suggest that this instrument is valid. Moreover, the firm-level estimates of product quality that we obtain pass many consistent test. Overall, this suggests that the method delivers a performant measure of product quality at the firm-level.
2. In theory and as we show in chapter 2, firms' sales are substitutes across markets in the presence of demand uncertainty. This result is empirically verified for Champagne wine export during the 2000-2001 US economic crisis. There is indeed evidence that individual Champagne exporters reallocate their surplus of sales from the US toward other destinations markets. Evidence of reallocation towards domestic market is less clear.
3. Labor regulation leads large firms to shrink and to substitute capital for labor. These firm-level responses translate into a loss of aggregate productivity by distorting the allocation of market shares across producers. The negative impact on TFP is magnified by the fact that large firms are more productive. When it comes to factor income distribution, there are two opposite effect of size-dependent regulations. First effect is within firms: regulated firms substitute capital for labor which tend to redistribute aggregate income away from labor. Second effect is between firms: large firms are relatively more capital intensive. The market share reallocation induced by the regulation therefore

shifts up the relative demand for labor. Preliminary results suggest that the latter effect dominates: size-dependent labor regulation shifts up the labor share.

The first contribution of the thesis is on the measure of quality at the firm-level.

Most of the empirical literature on firm-level quality and trade uses prices as a measure for quality.⁵ This practice prevails because (i) prices are convenient (they are directly present in the data) and because (ii) it seems intuitive that firms producing higher quality goods charge higher prices. However, a firm may charge high prices for reasons unrelated to the quality of its products. It may be for instance that the firm has high unit costs of production or that it has a large market power from which it can make a large price-cost margin. Consequently, prices are conceptually limited in their ability to proxy quality and call for a more direct measure. We develop such a measure in chapter 1.

Our approach to estimate quality is tightly linked to Khandelwal (2010). A solution, developed by Khandelwal (2010) at the aggregate level, is to estimate the quality of products exported by countries from the demand side. The insight to his identification is the following: if two countries A and B sell a same good at a same price, and A sells more than B, then country A supplies varieties of higher quality. This simple approach does not easily extend to the firm-level, though. The main reason is that it entails to estimate a demand function at the firm-level. This requires to have access to an instrument for prices at the firm-level, i.e. a variable which impacts individual production costs, but which is independent from non-price determinants of the demand addressed to the firm (e.g., the quality of its products).

Such variable is hard to find. A couple of contributions have access to firm-specific cost shifters (Gervais 2011, Roberts et al. 2012). Yet, the limit to these cost shifters as instrument for prices is that they rely on micro-variables, such as physical productivity or wages, which are likely to be functions of the demand addressed to a firm. For instance, high wages may signal that the firm produces high quality products by employing highly skilled workers. By contrast, in Chapter 1, we propose to use real exchange rate fluctuations on firms' imports as an instrument for prices. Real exchange rates act as a cost shifter because when the currency of a country from where a firm sources its inputs goes up, the cost of these inputs goes up as well from the point of view of the firm. Moreover, since real exchange rates are a macro-variable, there are unlikely to be endogenous to the behavior of individual producers.

An appealing feature of our solution to estimate firm-level product quality is that it requires widely available information. It can be for instance applied to any firm-level customs dataset, such dataset being now available for a large set of countries.⁶ Another interesting aspect of our approach is that it delivers a measure of product quality which may vary over time and across countries within a firm. As a consequence, it makes it possible to study a

5. More exactly, literature uses unit values (i.e. the ratio of the value and the volume of a trade flow) as a proxy for prices.

6. For instance, in the "Exporter Dynamics Database", the world bank has gathered firm-level export data on 45 countries, for the most part over the period 2003-2009. See Cebeci et al. (2012) for details on the content and the access procedure to the database.

wealth of questions regarding the way firms adjust the quality of their products in space and time.

The second contribution regards the way exporting firms react to shocks in foreign markets. I consider in particular two types of shocks: unexpected demand shocks and the entry of low-cost competitors.

When it comes to the response to demand shocks, we show in chapter 2 that when Champagne wine exporters are hit by an unexpected negative demand in one of their destinations, they reallocate their sales toward the rest of their destination portfolio. While most existing literature relies on correlations to document the interdependence of firms' sales across markets (see Nguyen & Schaur 2010, Ahn & McQuoid 2012, Liu 2012, Vannoorenberghe 2012) , we are the first to evidence firm-level substitutability across markets through a causal approach.

Berman et al. (2011) is the only existing contribution taking a causal stand to assess the interdependence of firms' sales. Interestingly, they find complementarity across markets. The divergence of their findings and ours certainly hinges on sample differences. While they look at the whole manufacturing sector, we restrict to the Champagne industry which for technical reasons (storage costs are high, production takes time) and as predicted by our model, is most likely to exhibit sales' substitutability. The combination of Berman et al. (2011) results with ours suggest that if complementarity dominates on average, substitutability dominates in sectors with large storage costs and long production processes.

The evidence of substitutability changes our understanding of exporting firms. In particular, it is consistent with a model of the firm where demand uncertainty is costly because it leads firms to produce the wrong amount of output. In this set-up, there is a hedging gain from trade liberalization to firms: access to foreign markets with unperfectly correlated demand shocks allows firms to smooth their total sales by reallocating their output surplus in one market toward markets with relatively more favorable demand realization.

Another important implication of our findings on firms' reaction to demand shocks is relative to the international transmission of shocks. The traditional view is that exporting firms contribute to the international synchronization of business cycles because exporting firms sell less when foreign economies slow down. We show that this mechanism can be dampened by firms' reallocation strategies: when foreign economies slow down, exporting firms reallocate their sales at home and thus (partially) maintain their total sales.

When it comes to the response to competition in foreign markets, we show in chapter 1 that firms tend to increase the quality of their exports when competition from low-wage countries gets fiercer. This result contributes to the empirical literature on firm-level quality and trade. Although there is a broad literature looking at the way firms adjust the quality of their products (or the skill of their workers) to a change in trade exposure in the domestic economy (Verhoogen 2008, Bloom et al. 2011, Brambilla et al. 2012, Amiti & Cameron 2012), we are the first to look at the way firms respond when their destination markets get more

accessible to competitors. From that respect, the closest contribution to ours is Martin & Méjean (2011). The authors describe how low-wage competition pushes up the aggregate quality of French exports by reallocating market shares toward highest quality firms. While they assume that individual qualities are constant over time and focus on the between firm aspect of quality upgrading, we are able to look at the way quality moves within firms.

We also get cross-industry results consistent with existent fact at the aggregate level. Indeed, we show that firms from industries which are more vertically differentiated (i.e. industries where producers can differentiate in quality from their competitors) increase more the quality of their exports when low-cost competition rises. This is consistent with Khandelwal (2010) who finds that industries which employment and output shrink the most due to low-cost competition are the least vertically differentiated ones. In fact, our result suggests that the better resilience of vertically differentiated industries come from the fact that in such industries, producers are better able to contain low-cost competition by climbing up the quality ladder.

Finally, our findings on quality upgrading have potential implications for trade liberalization. They suggest that when a trade partner opens up its market to new competitors, this may spill over the domestic economy through the quality response of exporting firms. The quality upgrading operated by exporters may indeed increase the skill premium since the production of quality is skill-intensive (Verhoogen 2008) as well as cost of living inequalities because rich consumers demand more quality (Handbury 2012).

The third contribution of this thesis is to shed new light on the impact of size-dependent labor regulation. Policies which generate heterogeneity in prices faced by individual producers are widespread. Size-dependent labor regulations is an example of such policies. By constraining more the employment conditions for large firms, such regulations implicitly shift up the cost of labor faced by large firms relative to small firms. The literature has consistently pointed to the loss of aggregate productivity that such policies may cause by distorting the allocation of productive resources across firms. By contrast, in chapter 3, I analyze the impact of size dependent policies on the distribution of income between capital and labor. More specifically, I carry out my analyses in the context of the 50 employees threshold in the French labor regulation.

I show a combination of elements which are determinant to understand the response of factor income distribution. The first element is within firms. I show that regulated firms respond to labor regulation by increasing their capital-labor ratio. This mechanism tend to redistribute income towards capital. The second element is in the cross-section of firms: regulated (large) firms are more capital intensive than unregulated (small) firms. Therefore, by reallocating market shares toward small (labor intensive) firms, the regulation may induce a reallocation of income toward labor.

It should be clear that my approach in chapter 3 is not comprehensive enough to deliver the expression of the optimal policy. In particular, I do not consider the direct gain of the

regulation to workers. Yet, I highlight a new aggregate outcome which is impacted by the regulation and so which should factor in the design of the optimal policy: factor income distribution.

C A closer look at the thesis

C.1 The quality of products exported by French firms: measurement and response to low-cost competition

context The analysis of the quality dimension of international trade is a vibrant field of research. It has shed new light on the patterns of trade as well as on the impact of trade on inequalities. However, the pursue of this research agenda is hindered because quality is hard to measure at the firm-level.

Solutions to estimate quality from micro-data were developed in empirical industrial organization. However, these solutions require product characteristics to be observed and such information is limited in trade data. In fact, customs dataset only provide the category of a good in the nomenclature. For instance, French customs data specify whether an exported car is new or used as well as the power of its engine, but it says nothing about its size, maximum speed, number of doors or its overall robustness. Therefore, existing literature has mostly relied on prices as a proxy quality. If it seems intuitive that higher quality firms charge higher prices, high prices may also signal high production costs or a large price-cost margin. These confounding factors call for a more direct measure of quality.

question In chapter 1, we propose a new solution to estimate the quality of products at the firm-level using trade data. Besides this methodological question, as an application of our method, we also look French firms adjust the quality of their exports to the low-cost competition they face in foreign markets.

methodology The first step to estimate quality is to define it. Here, we follow Hallak & Schott (2011) and define quality as being any attribute of a good, whether tangible or not, which raises consumers' valuation of it. Accordingly with this definition, we put functional assumptions on consumers' preferences. Namely, we assume that in each destination, the representative consumer has CES preferences over the vertically differentiated varieties of a good. In this demand system, the quality of a variety shows up as a utility shifter. This shifter is the parameter of interest in our estimation.

Consumers' preferences deliver a demand function such that the log of a firms' revenue in a market is a linear function of the log of prices, of a market-specific term and of the firm-specific quality. Therefore we back out quality as the residual of a regression of log revenues on log prices and a market fixed effect. Because of the market fixed effect, the quality we estimate is relative to the mean quality in each market.

The main challenge to estimate the demand function is to deal with the fact that prices are potentially endogenous to quality (e.g., quality is costly to produce and so high quality goods are more expensive). To cope with that identification issue, we propose a new instrument variable for prices. Namely, we construct a *weighted mean of the real exchange rate faced by firms on their imports*. Weights are equal to the share of a country in the imports of a firm at the initial date of the sample.

To be valid, our instrument should be correlated to prices (the rank condition) and uncorrelated to quality (the exclusion condition). The rank condition is verified because when the currency of a source countries goes up, the cost of imported inputs goes up and, as the firm passes on its extra costs to its consumers, export prices go up as well. As to the exclusion condition, it is verified because real exchange rate fluctuations have macro determinants and are therefore exogenous to the quality of a individual firms. Yet, it might be that firms adjust the quality of their exports to real exchange fluctuations on imports. We investigate that possibility by adding controls for this mechanism in our specification. They leave results unchanged.

Once we have quality in hand, we still need a measure of low-cost competition to assess the way quality responds to it. Here we follow Martin & Méjean (2011) and measure low-cost competition as the share of low-wage countries in the total imports of a good by a country. From this product-country measure, we construct a firm-product-country measure by averaging low-wage competition faced by a firm across its different destination markets. Finally we regress time varying quality over a the firm-specific low-wage competition, a set of time dummies and a firm fixed effect. The identification comes from the co-movements in the relative fluctuations of competition and quality in the cross-section of firms.

results We apply our instrumental variable strategy to French customs data for the period 1995-2007. This dataset covers almost the universe of French imports and exports. In the data, a trade flow is characterized by a firm identifier, a country, a 8-digit product code and a year. For each flow, both the value and volume are reported.

The different outputs of our quality estimation are convincing that our approach performs well. First, the instrument is strong and corrects the price-elasticity of demand in a direction consistent with the suspected positive correlation between quality and prices. Second, the value of the price elasticity we obtain is comparable to values obtain in the industrial organization literature. Third, using Sutton's measure of vertical differentiation, we show that our demand elasticity estimates are significantly smaller in more differentiated sector. This is consistent with the results obtained by Broda & Weinstein (2006) using aggregate trade data. Last, the quality measure we estimate for the exports of Champagne wine strongly and positively correlate with the expert assessed quality of Champagne used in Crozet et al. (2012).

The other set of results regards the impact of low-wage competition on firms. We show that firms upgrade the quality of their exports when low-cost competition gets fiercer. Interestingly,

this reaction occurs with a four years lag, suggesting that adjusting quality takes time. As to the magnitude of the effect, we find that a ten percentage points increase in the penetration rate of low-wage countries (this is approximately the variation observed in France's main destination markets over the period of the sample) induces a 2% quality upgrade.

C.2 Do firms reallocate sales across destinations in response to unexpected demand shocks?

context Mainstream trade theory assumes that firms' sales decisions are independent across markets. In other words, shocks faced by a firm in one destination market have no impact on that firm's sales in other markets. This assumption has implications for the international diffusion of shocks. Indeed, even if shocks are independent across markets, then may still spread internationally through firms' decisions, provided that the independence assumption is not verified.

question Intuitively, one reason why firms' decisions might be interdependent across markets is because firms reallocate their sales to mitigate the impact of demand shocks. In chapter 2, we amend standard models to deliver this mechanism and we test the predictions of our theory.

methodology We develop a simple model of a multi-destinations exporting firm facing an isoelastic demand function in each destination. We depart from standard models by assuming that at the end of each period, firms produce their output before the demand shifter in each market is realized. Moreover, firms do not have access to a storage technology. At the end of a period, firms observe demand shifters and solve a simple problem of allocation of their total output across their destination portfolio. We show that optimal sales to a given destination are decreasing in the realization of demand in other markets. This is because firms reallocate their sales towards markets which consumers have the highest willingness to pay.

In a second step, we bring the model to the data and test its main predictions. We rely on a difference-in-differences (DD) regression approach. The scenario to our identification is one where a group of firms serve a market A and only a subset of those, the treated group, simultaneously serve a second market B . When a negative demand shock hits market B , and assuming that firms are "comparable" across groups, our model predicts that sales in market A should grow more at the time of the shock. We use our model to structurally derive the DD regression that we estimate. This allows us to get a structural expression of the treatment effect as well as to put usual DD identifying assumptions in terms of models' parameters. Thus, our main identifying assumption is that firms cost shocks and demand shocks in B at the time of the shock are on average equal across groups.

The demand shock we use to test our model is the drop in demand in the US triggered by the 2000-2001 economic recession. This demand shock is of interest for several reasons. First of all, this shock was unexpected. Second, to observe reallocation we need a shock in

a relatively large market. The US are the first non European trade partner of France. One concern might be that the 2000-2001 crisis was not circumscribed to the US and contaminated markets where we try to identify a reallocation effect. However, the structural expression of the treatment effect tells us that our approach does not require the demand shock to be zero at the time of the financial crisis, outside the US. It simply requires the shock to be relatively larger in the US. This is consistent with the idea that some elements which led to the recession were specific to the US (e.g., the 11/9 terrorist attack), and it is confirmed in our data.

In our analysis, we decide to focus on the Champagne industry as it presents several advantages. First, Champagne is a highly traded good. Second, in most cases, exporters of Champagne are also Champagne producers. Third, the US are the main export market for Champagne. Fourth, producing Champagne takes time which is a crucial assumption in our theory. Fifth, storing Champagne bottles is very costly which limits temporal reallocation. Last, Champagne has to be produced in France which rules out potential offshoring reaction following demand shocks.

To conduct our analysis, we make use of three datasets. The first, called BRN, reports information on firms' domestic production and firms characteristics. The second is provided by French customs. It reports bilateral value and quantity exported by French firms at a very low level of disaggregation. The third dataset is the expert assessed rating of Champagne established by Juhlin (2008) and merged to French customs data by Crozet et al. (2012). This last source provides us with a quality measure which we use as a control in our DD approach. It also allows us to identify Champagne producers in the BRN dataset. The merge of these dataset results in a panel of 196 Champagne producers over the period 1995-2007.

results Graphical inspection reveals that non-US exports follow the same trend across groups outside the period of the shock. This validates the DD approach. We find that in 2000, non-US exports by firms from the treated group (i.e. firms serving the US) grew 40% more than non-US exports by firms from the control group. The sign of the treatment effect is consistent with a reallocation effect.

In order to assess whether the magnitude of the coefficient is in line with our theory, we use the structural expression of the treatment effect. This expression involves the price elasticity of demand, as well as the individual realized and anticipated share of non-US sales. Our strategy is to calibrate these different elements in order to see if they deliver a calibrated treatment effect close to the estimated one. The price elasticity can be easily calibrated and the realized share of non-US sales is observable. To calibrate the anticipated share of non-US sales, we assume that firms have extrapolative anticipations: they expect the share to grow at the same rate between $t - 1$ and t as between $t - 2$ and $t - 1$. This analysis suggests that our treatment effect is close to four times too large.

The size of this effect might be explained by two confounding factors. First, treated firms are higher quality firms and so they sell to different individual consumers with potentially different income shocks. It follows that the large estimated treatment effect might result

from treated firms facing a smaller demand drop in 2000, at the time of the crisis. Second, treated firms are bigger and thus certainly less financially constrained. Since the 2000 US crisis started as a financial crisis, it is possible that the financing conditions of firms from the control group deteriorate more than conditions of treated firms. This story also implies a positive treatment effect. Inclusion of controls for these mechanisms in the DD regression barely change the coefficient which is convincing that at least part of the treatment effect is indeed due to a reallocation effect.

C.3 Size dependent labor regulation and capital labor substitution at the micro and macro level

context In many countries, labor regulation is more binding to small firms than to large firms. According to the literature on resource misallocation, such regulatory designs may have a large negative impact on the aggregate productivity. This effect is due to the fact that large firms are more productive on average. Therefore, when a size-dependent regulation is introduced and the relative production costs of larger firms goes up, market shares get reallocated toward smaller (unproductive) firms and the aggregate productivity falls.

This insight has implications beyond the question of the aggregate productivity. It more generally points toward the importance of differences across small and large firms in shaping the aggregate impact of size-dependent regulations. For instance, factor intensity is another well-established dimension along which small and large firms differ. By increasing the weight of small (labor intensive) firms, size-dependent regulations may therefore impact factor income distribution. I study this question in chapter 3.

question Chapter 3 is built around following questions: how do size-dependent labor policies impact the distribution of income between capital and labor? Do both factors lose from the loss of aggregate productivity induced by the legal constraint or can the loss be compensated for by a redistribution effect? How does the redistribution effect decompose between a within-firm effect (regulated firms substitute capital for labor) and a between-firm effect (regulation reallocates market shares toward small and labor intensive firms)?

methodology We address these questions in the context of the 50 employees threshold in the French legislation. As a firm located in France exceeds 50 employees, it has to meet a set of additional constraints (e.g, it has to set up a works council, a committee for health and working conditions, it has to agree with its employees on a profit sharing rule, etc.). The sharp discontinuity in French firms' regulatory environment makes the 50 threshold an ideal set-up to identify firms' response to the regulation.

My work proceeds in three steps. In a first step, I use a census of French firms to describe perturbations observed in the firm size distribution and in firms' capital-labor ratio in the neighborhood of the 50 employees threshold. The dataset contains firm-level information on

the workforce size, the wage bill and the main elements of the balance sheets for the universe of French firms above 20 employees and for the period 1995-2007.

In a second step, I lay down a simple partial equilibrium model with heterogeneous firms and a regulation discontinuity. The model is static and builds upon the closed economy version of Melitz (2003). In the model, firms produce out of capital and labor through a CES production function and they face an isoelastic residual demand function. As in Melitz (2003), firms differ along a single dimension, namely productivity. However, unlike Melitz (2003), productivity differences across firms may be factor biased. In other words, the relative productivity of capital to labor may differ across firms. As to the regulation, I model it as a per period unit tax on labor which applies to firms above a workforce size threshold. The model delivers observable predictions which are qualitatively consistent with the patterns in firms' behavior observed around the 50 threshold in France. I use model's predictions regarding firm size distribution and the capital labor ratio to structurally estimate the implicit tax rate of the regulation as well as the capital-labor elasticity of substitution.

The third and last step consists in embedding my partial equilibrium model in a general equilibrium set-up in order to implement counterfactual experiments. The goal is to characterize the impact of the regulation on the aggregate productivity and the factor income distribution. In the model, the labor tax is fully redistributed to workers. I use my estimates to calibrate the tax equivalent of the regulation as well as the capital-labor elasticity of substitution. Most of remaining parameters are calibrated by constraining the model to match aggregate moments from the data.

results The first set of results regards the pattern of the firm size distribution and of the firm-level relationship between capital-labor ratio and size. I identify three main facts. First, firm size distribution exhibits an excessive mass just below 50 employees (I estimate that over the interval between 40 and 49 employees, the excessive mass stands for 20% of firms). The fact that firms bunch below the threshold is consistent with some firms cutting back their size in order to avoid the cost of the regulation. Second, there is a positive discontinuity at 50 in the capital-labor ratio. This second fact is consistent with regulated firms substituting capital for labor in order to mitigate the cost of the regulation. Third, there is a positive relationship between size and capital intensity.

My model of the firm is able to replicate these facts. My structural estimates imply that the regulation shifts up the unit wage by more than 7%. However, since labor costs are only a fraction of total production costs, the regulation only increases by 3% the production costs of firms around the threshold. As to the other delivery of my structural estimation, the capital-labor elasticity of substitution, I get a value of 0.6 which is consistent with recent results in the literature.

The last set of results regards counterfactual exercises. Here, the focus is on the *effect of a full-deregulation* on the real wage and the real interest rate. In the model, the sign of these effects is a priori ambiguous. On the one hand, labor loses from the between-firm

substitution effect. Namely, market shares flow toward large (and capital intensive) firms when the regulation is deactivated. On the other hand, labor gains due to the fact that the deactivation of the regulation optimizes the allocation of resources across firms and so decreases the cost of living. As to capital, on the one hand it loses from the fact that in the absence of a regulation, large firms substitute labor for capital. And on the other hand, capital gains from the improved allocation of resources. Calibration is therefore necessary to establish which effect dominates for each production factor. Preliminary results suggest that labor loses from deregulation while capital gains.

Chapter 1

A new Method for Quality Estimation using Trade Data: An Application to French firms

Abstract

This paper presents a novel instrumental variable strategy to estimate time-varying product quality at the micro level. Our method implies the estimation of firm-level demand functions. To deal with the endogeneity of prices, we construct a new firm-specific instrument, based on variations in exchange rates combined with firm-specific import shares. By shifting a firm's costs, this instrument generates firm-specific price variations independent from demand shocks. Our approach delivers consistent estimates of the price elasticity of demand from a limited number of assumptions. Higher quality is then assigned to products with higher sales conditional on prices. We implement our method on French customs data and we assess its reliability through correlations with alternative measures of quality. We use our estimates to measure the quality response of French firms to low-wage competition on foreign markets. This important question requires time-varying quality and thus can not be addressed with existing quality estimations. Using import penetration rates as a measure of competition, our results suggest that firms upgrade quality when facing increased low-wage competition.

1 Introduction

The quality produced by firms has deep implications for aggregate outcomes in international trade. A recent literature has emphasized the key role of quality in shaping the patterns of trade.¹ The analysis of firm-level quality has also delivered new insights into the impact of globalization on inequalities.² The importance of these findings calls for a thorough understanding of quality determinants at the firm-level. However, the main challenge raised by this research agenda is that quality is unobserved.

In this paper, we propose a new empirical methodology to estimate time-varying product quality at the firm-level. Our method is robust to unobserved product characteristics. It requires information on a firm's output prices and sales as well as on its imports by country. The fact that our strategy identifies time-varying quality makes it suitable to analyze the way firms adjust their quality in response to changes in their economic environment. We implement our method on French customs data as it has been extensively used to analyze firm-level quality. We assess the reliability of our estimation by correlating the estimated quality of exports to other characteristics of the firm, including alternative measures of quality. Then, we take advantage of our measure to look at the within-firm quality response to low-cost competition on foreign markets. This important question requires a time-varying quality measure and cannot be addressed without our method.

The paper's main contributions can be summarized as follows. First, to motivate our estimation, we propose a simple system of demand for a vertically differentiated good. The demand system features nested CES preference, each nest being specific to a quality level. Secondly, we present a new firm-specific instrument for prices which allows us to consistently estimate demand functions at the firm-level. This instrument is robust to time varying quality and unobserved product characteristics. From the demand function, we are able to estimate the quality of French exports at the firm-product-destination-year level. Third, using our estimates, we show that prices, the most conventional proxy for quality, are significantly more informative on quality in vertically differentiated sectors than in homogeneous sectors. Finally, we identify the quality response of French exports to low-cost competition on foreign markets. We find evidence that firms upgrade their quality when competition increases.

The methodology developed in this paper is motivated by the observation that leading approaches to estimate quality at the firm-level require information on product characteristics.³ Unfortunately, such information is absent from most nationally-representative firm-level datasets. Consequently, existing literature has heavily relied on prices as a quality measure. This proxy is conceptually problematic because prices are polluted by differences in production costs.⁴ Recent advances on the estimation of quality have been achieved by Gervais

1. Quality has shed new line on the pattern of trade at the country-level (Schott 2004, Hallak 2006, Feenstra & Romalis 2012, Hallak & Schott 2011) as well as at the firm-level (Crozet et al. 2012, Gervais 2011, Manova & Zhang 2012).

2. Recent contributions on firm-level quality, trade and inequality include Faber (2012), Brambilla et al. (2012) and Verhoogen (2008).

3. Influential estimations of demand systems with vertically differentiated goods and observed product characteristics include, e.g., Berry et al. (1995), Nevo (2000) and Handbury (2012) .

4. At the country-product level, Khandelwal (2010) shows that the ability of prices to proxy quality is particularly low for homogeneous goods. We get a similar result at the firm-level.

(2011) and Roberts et al. (2012). These works estimate firm-level demand fundamentals in frameworks where information on product characteristics is very limited. However, they obtain identification at a cost: quality is assumed constant over time, within the firm. By contrast, our approach is robust to both unobserved product characteristics and time-varying quality.⁵

Our work proceeds in several steps. First, we present our estimation. We reveal quality from the demand side. Namely, higher quality is assigned to products with higher sales conditional on prices. Our identification of quality therefore implies the estimation of demand functions at the firm-product level. We derive the expression of these demand functions from a simple model of demand for quality. The model features nested CES preferences, each nest being specific to a quality level. The main challenge we face to estimate demand functions is to deal with the endogeneity of prices. Prices are likely to be correlated to demand shocks, because quality is costly to produce.⁶ To address this problem, we construct a novel instrument for prices, exploiting fluctuations in exchange rates. These fluctuations, interacted with firm-specific import shares, shift a firm's costs of importing goods. As the firm passes importing cost variations on to its consumers, the instrument generates firm-specific export price and sales variations. These variations are arguably exogenous to unobserved demand shocks (e.g., quality shocks) and allow us to identify the price-elasticity of exports.⁷ Quality is then identified from the residual variation of demand, once price variations have been controlled for.

The strategy which consists in backing out quality from the estimation of a demand system is present throughout the literature. Our methodological contribution is in the way we deal with the endogeneity of prices. Similarly to us, Hallak & Schott (2011) and Khandelwal (2010) rely on an instrumental variable approach to identify quality at the country-product level using trade data. To be applied at the firm-product level, their methods require an instrument for prices which varies across firms. We provide such an instrument. At the firm-level, Gervais (2011) and Roberts et al. (2012) also estimate quality by instrumenting prices. However, these studies use instruments, respectively physical productivity and wages, which are questionable if quality varies over time, within the firm. By contrast, our instrument is robust to time-varying quality. Also at the firm-level, Khandelwal et al. (forthcoming) construct quality by calibrating price-elasticity with estimates from Broda & Weinstein (2006). The relevancy of these price-elasticities estimates is open to question as they were obtained from country-level data. By opposition, we estimate price elasticity from the same data that we use to construct quality estimates.

As a second step, we implement our procedure using French customs data spanning from

5. To the best of our knowledge, Khandelwal et al. (forthcoming) propose the only time-varying measure of product quality at the firm-level, with unobserved product characteristics. However, their measure, unlike ours, is not estimated. We detail this point below, when presenting our methodology.

6. See, e.g., Hallak & Sivadasan (2011), Johnson (2012) and Kugler & Verhoogen (2012) for trade models where quality is costly and endogenous at the firm-level.

7. The use of exchange rates as an instrument for prices connects our estimation to Berman et al. (2012) and Amiti et al. (forthcoming). These studies empirically analyze the firm-level pass-through from exchange rates to export prices. However while both works are interested in the heterogeneity of the pass-through across firms, we only use the effect of exchange rates on export prices as a first stage to a demand function estimation.

1995 to 2010. On top of quality, our instrumental variable approach delivers a series of by-products which we use to assess its reliability. First, we find that the import-weighted exchange rate, our instrument, is strongly and positively correlated to export prices charged by firms. This is consistent with the assumption we make to motivate the instrumentation, namely that exchange rates shift a firm’s production costs. Another output of the estimation is the price elasticity of demand. In order to evaluate the ability of our instrument to correct for the endogeneity of prices, we estimate this elasticity both via ordinary least squares and two stages least squares. Our instrumental variable procedure affects the estimates of price-elasticities consistently with a correction of a simultaneity bias: while ordinary least squares estimates deliver a low (in absolute value) price-elasticity (0.8), the instrumental variable approach produces estimates consistent with the existing studies in the industrial organization literature, ranging from 1.8 to 2.4, depending on the specification. In order to further assess the reasonableness of our price elasticity estimates, we rely on cross-industry comparisons. In line with evidence at the country-product level, we find that demand is significantly more elastic in more homogeneous sectors.⁸

Then we investigate the properties of our quality estimates by running correlations with existing measures of quality at the firm-level. A natural benchmark is provided by Crozet et al. (2012). This study uses one of the very few “direct” measure of firm-specific quality present in the literature. More specifically, Crozet et al. (2012) observe ratings attributed by an expert to a sample of French Champagne producers. We compare these ratings with our estimated quality of exported Champagne and find a positive and strongly significant correlation. Prices, the most popular proxy for quality in the literature, are another natural benchmark with which to compare our quality estimates. We find that prices are positively and significantly correlated to quality. This is true across firms as well as over time within a firm. However, this correlation is significantly smaller for more homogeneous sectors. Using Sutton (2001)’s sectoral measure of vertical differentiation, we find that for “mineral products”, the least vertically differentiated product category, prices are approximately 3 times less elastic to quality than for “chemicals and allied industries”, the most differentiated product category. This means that prices are informative on quality, but less so in more homogeneous sectors.

Finally, to illustrate the scope of our method, we employ our quality estimates to analyze the way firms respond to low-cost competition in foreign markets. The recent increase in the participation of low-wage countries in international trade has had a large impact on manufacturing industries in developed economies. For instance, Autor et al. (2013) show how manufacturing workers in the United States have been hurt by the increasing penetration of Chinese goods on the american market. Relatedly, Khandelwal (2010) provides evidence that the impact of low-wage competition has been significantly larger in industries with shorter quality ladders. Closer to our question, Martin & Méjean (2011) show that low-cost competition induces a reallocation of market shares towards higher quality firms which ultimately results in a rise of aggregate quality. By contrast, we look at the within-firm quality response to low-wage competition. This question requires a time-varying measure of quality and as such can not be addressed with existing firm-level quality estimations with trade data.

8. See Broda & Weinstein (2006).

Our identification strategy consists in correlating the dynamics of low-cost competition in foreign markets with the dynamics of the product quality supplied by French firms to these markets. To this end, we first compute the penetration of low-wage countries at the country-product-year level using the trade dataset BACI. Then, for each destination market and each firm, we construct a measure of the low-cost competition faced by the firm in the rest of the world. This measure varies across firms within a market since firms serve different destinations. We identify the quality response to competition from the firm-specific dynamics in this rest-of-the-world measure of competition. This identification strategy assumes that there is a positive correlation in the quality of a good supplied by a firm across destinations. Intuitively, we assume that within the firm, the quality adjustment due to competition in one destination spills over the quality served to other destinations.

Using this identification strategy, our results suggest that low-cost competition induces quality upgrading within the firm. Interestingly the response of quality takes time to occur. More specifically, the quality of a firm raises by 1.2% four years after a 10 percentage point increase of the low-wage countries' penetration. We find no significant response before four years. It suggests that upgrading quality requires slow adjustments within the firm. In addition, we find that quality upgrading is more pronounced in more vertically differentiated industries. These results contribute to the literature on the relationship between firm-level quality and trade exposure. While existing studies mostly focus on firms from developing countries (see, e.g., Verhoogen 2008, Brambilla et al. 2012, Khandelwal et al. forthcoming), our results suggest a new channel through which firms from developed countries can mitigate the impact from low-wage competition.

This paper is structured as follows. In the next section, we derive a simple model of demand with vertically-differentiated goods. In section 3, we present our novel instrumental strategy, describe its validity and its positioning relative to existing methods. Section 4 applies our methodology and demonstrates its effectiveness. In section 5, we describe the quality estimates we obtain through correlations to alternative measures. In section 6, we investigate the impact of low-cost competition on within-firm quality adjustments. Finally, section 7 concludes.

2 An Empirical Model with Demand for Quality

In this section, we consider a simple demand structure for a vertically differentiated good. The model delivers an expression of the demand function at the firm-level in which quality is a component of the demand shifter. In the next section, we bring this demand function to the data.

2.1 Consumer Preferences

Let us consider a global economy composed of a collection of markets. Each market is populated with a mass of symmetric buyers and a mass of heterogeneous firms trading a differentiated good. Within a market, each firm produces a single variety of the good and has a monopoly on this variety. Varieties of a good are differentiated both horizontally and

vertically.⁹

Buyers combine varieties through a nested CES aggregator. At the lower level, buyers associate varieties horizontally within each quality nest. At the upper level, buyers bundle varieties along the quality ladder. The preference of a representative buyer is:

$$\begin{cases} \tilde{X}_{m,t}(q) &= \left[\int_{\Omega_{m,t}(q)} x_{v,t}^{\frac{\sigma-1}{\sigma}} dv \right]^{\frac{\sigma}{\sigma-1}} \\ X_{m,t} &= \left[\int (q\tilde{X}_{m,t}(q))^{\frac{\rho-1}{\rho}} dq \right]^{\frac{\rho}{\rho-1}} \end{cases} \quad (1.1)$$

With, $X_{m,t}$, $\tilde{X}_{m,t}(q)$ and $x_{v,t}$ respectively total consumption, consumption of quality q and consumption of variety v by the representative buyer in market m at date t .¹⁰ $\Omega_{m,t}(q)$ is the set of varieties with quality q supplied to market m , at date t . σ and ρ are respectively the elasticity of substitution across varieties within a nest and between nests. As is standard, we assume $\sigma > 1$ to ensure that mark-ups are finite under monopolistic competition. We also make the natural assumption that varieties are closer substitutes within a quality nest than across nests: $\sigma \geq \rho$.

In equation (1.1), quality is a utility shifter. As such, quality is an index containing any intangible (e.g. reputation, quality of the customer service, brand name) and tangible characteristic of a variety that is observed by consumers. It is a measure of the overall appeal of a variety to a consumer.¹¹

The use of Dixit-Stiglitz preferences is very standard in all leading heterogeneous trade models with vertical differentiation.¹² The novelty of our demand system lies in its quality-nested structure.¹³ With CES preferences, the absence of quality nests implies that two varieties are equally substitutable within qualities and between qualities. If this unrealistic property might not be an issue in some applications, it shows particularly undesirable in the present paper. In fact, we are interested in identifying the impact of the competition from low-wage countries on firms' demand. As low-cost firms are most likely low quality firms, it is possible that their entry impacts more strongly low quality incumbents. Our demand system allows for this possibility.^{14 15}

9. In our empirical application, we define a “variety” as a firm-country-product combination and a “market” as country-product combination.

10. Remark that $x_{v,t}$ is not indexed by m . This is because we define a “variety” v as being specific to a market m .

11. We adopt a broad definition of quality, in line with the literature on quality and trade. Some contributions have adopted a more conservative terminology to refer to the same economic concept. For instance, Roberts et al. (2012) refer to the variety-specific utility shifter as a “demand index” when Foster et al. (2008) rather opt for “demand fundamental”.

12. see, e.g., Verhoogen (2008), Baldwin & Harrigan (2011), Kugler & Verhoogen (2012), Johnson (2012), Crozet et al. (2012). By contrast, some contributions (e.g., Khandelwal 2010, Roberts et al. 2012) assume logit demand systems. Under mild conditions, CES and logit preferences deliver similar aggregate demand curves (see Anderson et al. 1987).

13. If some trade and quality papers may use a nested demand structure (Khandelwal 2010), a nest is defined as a good, not a quality level.

14. The entry of new competitors in a specific quality-nest q shifts down all firms' demand curve. However, the shift is larger for incumbent firms in nest q as they are closer substitutes to new varieties.

15. The nested structure of our demand system also has implications for the quality response of firms to low-quality competition. Since nests act as a hedge against competition shocks in other nests, firms might want to adjust the entry of new low-quality competitors by upgrading quality in order to escape most affected

2.2 Firms' Demand and Quality as a Demand Shifter

In market m , at date t , consumers allocate their total expenditure, $E_{m,t}$, across varieties, in order to maximize their utility (1.1). The aggregate demand function faced by a variety v at date t is

$$r_{v,t} = p_{v,t}^*{}^{1-\sigma} q_{v,t}^{\rho-1} \tilde{P}_{m,t}(q_{v,t})^{\sigma-\rho} P_{m,t}^{\rho-1} E_{m,t} \quad (1.2)$$

With $r_{v,t}$ the sales of variety v . $p_{v,t}^*$ is the price of variety v faced by consumers of market m . Namely, $p_{v,t}^*$ is the CIF (Cost Insurance Freight) price labeled in market m 's currency. $q_{v,t}$ is the quality of variety v . $\tilde{P}_{m,t}(q)$ is the quality q -specific price index, in market m , at date t . The dependence of $\tilde{P}_{m,t}(q)$ on q , m and t stands for the fact that the density of competitors might differ along the quality ladder as well as across markets and over time. $P_{m,t}$ is the aggregate price index.¹⁶

We assume that exporting from home country H (home country is "France" in the application) involves iceberg trade costs common to all firms serving a market. So CIF and FOB are linked by following relationship:

$$p_{vt}^* = e_{dt}^{-1} \tau_{mt} p_{vt} \quad (1.3)$$

With e_{dt} the direct nominal exchange rate from home currency to destination d 's currency (home currency is Euro in the application), τ_{pdt} the iceberg trade cost ($\tau_{pdt} \geq 1$) and p_{fpt} the FOB price in home currency. Plugging (1.3) and log-linearizing, we can re-express (3.2) as follows:

$$\log r_{vt} = (1 - \sigma) \log p_{vt} + \lambda_{vt} + \mu_{mt} \quad (1.4)$$

with

$$\begin{cases} \lambda_{vt} \equiv (\rho - 1) \log q_{vt} + (\sigma - \rho) \log \tilde{P}_{mt}(q_{vt}) \\ \mu_{mt} \equiv \log \left(\frac{\tau_{Hdt}}{e_{Hdt}} \right)^{1-\sigma} + (\rho - 1) \log P_{mt} + \log E_{mt} \end{cases}$$

Equation (1.4) is the one we bring to the data. One can see from (1.4) that the demand shifter of a firm contains firm specific and market specific arguments. Since the latter are not informative on quality at the firm level, the estimation developed in this paper identifies $\lambda_{v,t}$, the firm-specific part of the demand shifter. As $\lambda_{v,t}$ is the demand shifter of a firm, cleaned out from market-level components, one should think of it as a demand shifter deviated from the average demand shifter on a market. It follows that $\lambda_{v,t}$ will not be suited to analyze the

necks. Section A formalizes that intuition. Section 6 provides empirical evidence supporting that prediction.

16. Quality-specific and aggregate price indices verify:

$$\tilde{P}_{m,t}(q) = \left(\int_{\Omega_{m,t}(q)} p_{v,t}^*{}^{1-\sigma} dv \right)^{\frac{1}{1-\sigma}}$$

$$P_{m,t} = \left[\int \left(\frac{\tilde{P}_{m,t}(q)}{q} \right)^{1-\rho} \right]^{\frac{1}{1-\rho}}$$

aggregate quality of a market.

$\lambda_{v,t}$ is a function of quality. Holding $\tilde{P}_{m,t}$, the derivative of $\lambda_{v,t}$ with respect to $\log q_{v,t}$ is $(\rho - 1)$. The larger ρ , the smaller the love for variety of consumers along the quality ladder, the more sensitive consumers are to quality when they allocate their expenditure across qualities. In the limit case $\rho = \infty$, only the top best quality variety is demanded. ρ can therefore be seen as consumers' valuation for quality. The fact that this valuation is finite explains that low qualities are ever consumed. As in Khandelwal (2010), our identification does not allow us to separate consumers' valuation for quality from quality itself. In other words, we will measure quality in unit of sales rather than in units of utility.

$\lambda_{v,t}$ is not only a function of quality. It also depends on the quality specific price index $\tilde{P}_{m,t}(q_{v,t})$. This price index can vary along the quality ladder within a market. As such, $\tilde{P}_{m,t}(q_{v,t})$ can act as a confounding factor when it comes to using $\lambda_{v,t}$ to rank firms in quality.¹⁷ Suppose for instance that competition is tougher in the upper part of the quality ladder ($\tilde{P}_{m,t}(q)$ is decreasing in q). In that case, high quality firms enjoy a larger demand (conditional on price) because their products are more appealing. But this larger demand is dampened by higher competition. In that case, λ could be a decreasing function of quality.

In appendix A, we propose a simple model with endogenous quality where this competition effect does not dominate and $\lambda_{v,t}$ should be increasing with q . The idea behind this intuition is that if firms optimally choose their quality and if quality is costly- as assumed in models with endogenous quality- then firms need to have a positive return to quality ($\lambda_{v,t}(q)$ increasing in q) to find it profitable to invest in quality.¹⁸ It follows that $\lambda_{v,t}$ can be used to rank firms in quality within a market.

In this section we have derived the expression of a variety's demand function and made explicit its dependence on quality. Next section proposes an method to identify these demand functions.

3 Quality Estimation Strategy

In this section, we describe the novel estimation strategy we implement to identify the quality of exports at the firm-product-destination-year level, using customs data. Identification of quality is achieved by estimating demand equation (1.4). The typical challenge of demand identification lies in the fact that prices are endogenous to demand shocks. In order to deal with price endogeneity, we present a novel instrument, obtained by interacting firm-specific importing shares with real exchange rates. Our instrument is appealing as it remains

17. Existing papers on quality and trade implicitly assume the intensity of competition to be constant along the quality ladder. It is clear in models with CES demand systems: existing contributions do not have quality nests which implies that competition is summarized by the aggregate price index which is common to all firms, irrespective of their quality. It is also the case in papers with nested logit preferences, for the reason that nests are not along the quality ladder. It follows that existing estimations of quality are also polluted with competition variations along the quality ladder. Our paper makes this pollution explicit and carefully discusses its implications.

18. Let us illustrate this point through a simple example. Consider two firms, A and B . Firm A has a larger λ than firm B . If quality is costly and endogenous and firms maximize their profits, then it has to be that A has a larger quality than B , otherwise firm B would make more sales at a lower average cost by supplying the same quality as A .

valid in the presence of time-varying (unobserved) quality. As such, our demand estimation strategy has a broader scope of application than most existing ones.¹⁹

The high dimensionality of the dataset employed here requires some definitions. Hereafter, index f will identify a firm, p a eight-digit product category, d a destination country, and t a year. A ‘market’ will be defined as a destination \times product category \times year combination and a ‘variety’ as a firm \times product category \times destination. Re-written following the dimensions of our data, equation (1.4) is

$$\log r_{fpdt} = (1 - \sigma) \log p_{fpdt} + \lambda_{fpdt} + \mu_{pdt} \quad (1.5)$$

In equation (1.5), only r_{fpdt} and p_{fpdt} are observable to the econometrician. Our strategy identifies the price-elasticity σ , a market fixed effect μ_{pdt} and our parameter of interest: the variety-specific demand shifter λ_{fpdt} . Prices are presumably correlated with the demand shifter of the firm, λ_{fpdt} , through endogenous quality choices, or any internalization of demand shocks by the firm. Next subsections respectively present and discuss our instrumental strategy to fix this endogeneity issue.

3.1 Dealing with Price Endogeneity

In our setup, the endogeneity of prices comes from two mechanisms. First of all, we face a well-known simultaneity problem as prices are likely to be correlated with demand shocks. This might be due to the strategic behaviors of firms (i.e. firms adjust their mark-ups to demand shocks); but also to the presence of vertical differentiation within a product category. In the latter case, correlation between the price of a product and its level of demand would come from high quality firms passing on the cost of quality to consumers. This endogeneity channel leads ordinary least squares to underestimate the price-elasticity of demand, σ .

A second source of endogeneity, more specific to international trade data, comes from the construction of prices. Because prices are not directly observed, we follow the standard practice and use unit values as a proxy for prices. Unit values are obtained by dividing the value of a shipment by the quantity shipped. The use of this proxy may generate an attenuation bias due to the measurement error contained in the price variable.²⁰

Existing Methods Existing literature has used different empirical strategies to deal with price endogeneity. In particular, the literature in Industrial Organization has developed estimation procedures with instruments for prices. For instance, Berry et al. (1995) use competitors’ product characteristics, Hausman (1996) and Nevo (2000) use product’s price on other markets, while Foster et al. (2008) rely on estimated physical productivities. However, these instruments are not valid in the presence of unobserved vertical differentiation.²¹ As

19. We review existing firm-level demand estimations in section 3.1.

20. This attenuation bias will certainly be magnified by the flow fixed effects we use in our estimation. In fact, in the time series of a trade flow, the measurement error may represent a larger share of the variation of unit values than in the cross-section.

21. Berry et al. (1995), Hausman (1996) and Nevo (2000) all study specific markets, for which they clearly observe different varieties of a good, as well as their characteristics, reducing the possibility for unobserved

a consequence, these instruments are not usable in our context. Indeed, trade data contain no product characteristic, except for the category in the product classification. Despite a narrow definition of these categories (8-digit CN classification present in our data has around 8,000 positions), there is still a wide scope for (unobserved) vertical differentiation within each category.

Methods for demand estimation with trade data exist at the country-level. Khandelwal (2010) and Hallak & Schott (2011) use IV approaches. Their strategy are not suited to firm-level demand estimation as their instruments vary at the market level, not across firms within a market. Feenstra (1994) and Broda & Weinstein (2010) respectively develop and refine a very influential demand estimation using country-level trade data. Their identification exploits the heteroskedasticity of supply and demand shocks. Although their strategy could be applied to firm-level trade data, it involves an orthogonality assumption between demand and supply shocks which is likely to be violated in the presence of vertical differentiation (e.g., quality is costly).

Literature on demand estimation with trade data is scarcer at the firm-level. Roberts et al. (2012) and Gervais (2011) use firms' wages and physical productivities as instruments for prices. These instruments are only valid if product quality is constant over time within the firm. For instance, if a firm upgrades its quality, it might need more workers per physical unit of output. In that case physical productivity is (negatively) correlated to quality and OLS estimate of σ is biased downward. The assumption that product quality is time-invariant is not sustainable in the present paper as our goal is precisely to identify within-firm quality variations induced by low-wage countries competition. Khandelwal et al. (forthcoming) construct a firm-level quality measure by calibrating a CES demand system with price-elasticity estimates from Broda & Weinstein (2006). Conceptually, this approach raises two concerns. First, it implicitly inherits the identifying assumptions from Broda & Weinstein (2006). We explained above that these assumptions are problematic in the presence of vertical differentiation. Second, Broda & Weinstein (2006) estimates are obtained from country-level data. Elasticity may differ at the micro and the macro level (see, e.g., Imbs & Méjean (2008) or Chetty (2012)) and so generate biases in estimated firm-level quality.

Because existing method do not lend themselves to our exercise, we develop a new instrumental strategy, robust to unobserved quality differences within product categories.

A Novel Instrument for Prices at the Firm-level The approach developed in this paper takes advantage of the information coming from the importing activity of exporters. Exporters which do not import do not identify our parameters. We use real exchange rates faced by importing firms to instrument prices of exported goods. The basic idea is that real exchange rate shocks on a firm's imports are cost shocks. As the firm passes these cost shocks through to its export prices, sales adjust and the demand function is identified. Appendix A formalizes this mechanism. In order to generate firm-specific exchange rate shocks, we take advantage of the fact that the spatial structure of imports varies across firms

quality differences. In a different setup, Foster et al. (2008) and Handbury (2012) estimate demand functions for a wide range of products, but either restrict their analysis to homogenous products or use barcode-level data, which rule out the possibility of unobserved quality differences.

To gain insight into the identification, let us study the example of two firms selling in a same market. One firm imports from the United States, while the other imports from Europe. An appreciation of the dollar would induce an increase of the export price of the former, leaving unchanged the price of the latter. The response of these firms' relative sales to the change in their relative prices identifies the price-elasticity of demand. This example also conveys the intuition of our main identifying assumption: relative real exchange rate shocks across firms should be exogenous to relative demand shocks. Next subsection discusses this assumptions. It acknowledges situations where it is likely to be violated and adjusts the econometric specification accordingly.

Formally, our instrument is the import-weighted real exchange rate of a firm f at time t :

$$\overline{RER}_{ft} = \sum_s \omega_{0sf} \times \log(\text{rer}_{st}) \quad (1.6)$$

where ω_{0sf} is the share of goods imported by firm f , from source country s , at the initial date of the sample,²² and rer_{st} is the real exchange rate from a firm's own country (i.e. France in our application) to country s at time t . The exchange rate rer_{st} is defined using direct quotation, such that an increase of this variable implies larger costs for a firm. Moreover, the real term is computed using CPI indices. The formula of rer_{st} is:

$$\text{rer}_{st} = \text{er}_{st} \frac{\text{CPI}_{s,t}}{\text{CPI}_{\text{France},t}}$$

The ability of this instrument to predict export prices might be reduced for firms that imports and exports goods from the same destination.²³ Consider for instance a French firm importing from and exporting to the United States. An appreciation of the dollar should have little impact on the price of its input: such a firm should be able to implement hedging strategies out of its sales to the US to absorb exchange rate shocks. Consequently, we create a second instrument taking into account the degree of potential hedging of a firm. The idea is to interact importing and exporting weights for a same country. Formally, the second instrument is:

$$\overline{RER}_{ft}^h = \sum_s \omega_{0sf} \times \omega_{0sf}^{\text{exp}} \times \log(\text{rer}_{st}) \quad (1.7)$$

where $\omega_{0sf}^{\text{exp}}$ is the exporting weight of a firm toward destination s . We expect the pass-through from the RER on imports to export prices to be decreasing with \overline{RER}_{ft}^h . The inclusion of this second instrument will improve the strength of our first stage and therefore generate more accurately estimated exogenous price variations.

We conclude the presentation of the instruments with three remarks. First, the instrument is orthogonal to measurement errors on unit values as its construction does not involve information on exports. Therefore, our instrumental strategy deals with the measurement

22. In next section, we come back on the importance of using initial weights to compute the import-weighted exchange rate.

23. We thank Amit Khandelwal for pointing out to us this mechanism.

errors problem existing when estimating demand functions using unit values.

Second, similar instruments have been used in a series of recent international trade contributions (see Brambilla et al. 2012 or Bastos et al. 2012). In these papers, the export-weighted exchange rate generates exogenous change in firms' destination portfolio. In our case, the import-weighted average exchange rate creates exogenous firm-specific cost shifters due to the mechanical increase of the price of imported inputs.

Lastly, we are not the first paper looking at the pass-through from the cost of imported input to export prices. Amiti et al. (forthcoming) and Berman et al. (2012) run the same type of regression using respectively Belgian and French customs data. However, the motivation for their analysis differs greatly from ours. While, they are interested in the heterogeneity of the pass-through across firms, we only use the effect of exchange rates on export prices as a first stage to a demand function estimation. Moreover, we look at exchange rate movements on the import side, trying to control for exchange rate movements on the export size. Their focus is on the exchange rate on exports.

3.2 Discussion of the Identification

There are a few mechanisms that could affect the exogeneity of the instrument. First of all, the instrument is constructed from import shares, which are potentially endogenous to quality. Put simply, higher quality firms most likely import from countries with a stronger currency from where they can import higher quality inputs (In appendix A, we derive a model in which the spatial structure of a firm's imports depends on the quality it produces). So we expect the instrument to be positively correlated to quality in the cross-section of firms. This would induce the price elasticity to be biased toward zero. To fix this problem, we add variety-specific fixed effects (as defined above, a variety is a firm \times product category \times destination combination) to our demand estimation. As a result, identification is in the time series of a variety. Since the instrument is constructed using initial import shares, its time series variations are fully driven by (firm-specific) exchange rates dynamics and not contaminated by (endogenous) import share dynamics.

Another potential problem comes from the dual impact of exchange rates variations on firms' performances. While a change in exchange rates can increase input prices, it can also affect the competitiveness of firms on foreign markets. This is a concern to us as it suggests that our instrument could be correlated to a firm's demand shifter. In reality, this is not an issue with the structural demand equation we consider. Since we assume that price elasticity is common across firms, we also implicitly assume that the impact of a nominal exchange rate shock on a firm's demand curve is the same to all firms. So the impact of exchange rates on the demand shifter will be absorbed by market fixed effects. Because firms certainly face heterogeneous demand elasticities in practice, we provide several robustness checks in the appendix in order to further demonstrate the innocuity of this problem.. We proceed by excluding products sold by a firm that contemporaneously import from and export to a market (see appendix D, table D.1). Price elasticity estimates exhibit little sensitivity to sample variations. This is suggestive that the 'competitiveness' mechanism is unlikely to affect our results.

A last potential threat to the identification could come from the fact that exchange rate variations directly cause quality adjustments.²⁴ Bastos et al. (2012) show that an exchange rate shock may induce a firm to upgrade its quality if it improves its competitiveness in rich destination markets. In appendix A, we propose a model which predicts a symmetric effect on the import side. In the model, source countries produce inputs of different qualities. When an exchange rate shock makes imports from high (low) input quality countries more affordable, a firm upgrades (downgrades) the quality of its imported inputs, and output quality adjusts accordingly.

There is no reason to think that quality changes resulting from exchange rates variations are correlated with the instrument in a systematic way. A priori, an increase in $\overline{RE\bar{R}}_{ft}$ can equally result from the appreciation of the currency of a rich source country as of the currency of a poor source country. So the sign of the bias on price-elasticity, if any, is unclear. However, we take a conservative approach and neutralize the effect of exchange rates on quality by adding controls to the estimation. Namely, we incorporate the import weighted average GDP per capita of the firm as well as the export weighted average GDP per capita to the demand equation. The formula of these controls is:

$$\overline{gpc}_{ft}^{\text{exp}} = \sum_s \omega_{sft}^{\text{exp}} \times \log(\text{gdpc}_{st}) \quad (1.8)$$

$$\overline{gpc}_{ft}^{\text{imp}} = \sum_s \omega_{sft}^{\text{imp}} \times \log(\text{gdpc}_{st}) \quad (1.9)$$

These terms aim to capture quality adjustments following changes in the set of countries the firm imports from and exports to. The implicit assumption here is that GDP per Capita proxies the quality of inputs supplied by a country.²⁵ In the mechanism described above, exchange rates are suspected to affect quality only through an impact on a firm's spatial structure of imports. Controlling for that structure of exports thus implies that the instrument is orthogonal to the demand residual. The model presented in appendix A provides a theoretical foundation to these controls.

Consistently with above discussion, our preferred econometric specification is:

$$\begin{cases} \log r_{fpdt} = (1 - \sigma) \log p_{fpdt} + \alpha_2 \overline{gpc}_{ft}^{\text{exp}} + \beta_2 \overline{gpc}_{ft}^{\text{imp}} + \delta_{fpd}^{(2)} + \mu_{pdt}^{(2)} + \varepsilon_{fpdt}^{(2)} \\ \log p_{fpdt} = \gamma_1 \overline{RE\bar{R}}_{ft} + \gamma_2 \overline{RE\bar{R}}_{ft}^h + \alpha_1 \overline{gpc}_{ft}^{\text{exp}} + \beta_1 \overline{gpc}_{ft}^{\text{imp}} + \delta_{fpd}^{(1)} + \mu_{pdt}^{(1)} + \varepsilon_{fpdt}^{(1)} \end{cases} \quad (1.10)$$

With $\delta^{(i)}$, $i = 1, 2$ and $\mu^{(i)}$, $i = 1, 2$ respectively flow and market fixed effects. $\varepsilon^{(i)}$, $i = 1, 2$ are residual terms. And the identifying assumption is:

$$\mathbb{E} \left[\begin{pmatrix} \overline{RE\bar{R}} \\ \overline{RE\bar{R}}^h \end{pmatrix} \cdot \varepsilon^{(2)} \right] = 0$$

24. We thank Juan Carlos Hallak and Maria Bas for pointing out to us these potential threat to the identification.

25. In line with this assumption, Schott (2004) shows evidence that rich countries specialize in the export of high quality goods.

Namely, conditional on controls, $\overline{gpdc}^{\text{imp}}$ and $\overline{gpdc}^{\text{exp}}$, and fixed effects, demand shocks are uncorrelated to instruments.

In system (1.10), demand equation is identical to structural demand equation (1.5) where the flow-specific demand shifter, λ_{fpdt} , is identified by:

$$\lambda_{fpdt} = \alpha_2 \overline{gpdc}_{ft}^{\text{exp}} + \beta_2 \overline{gpdc}_{ft}^{\text{imp}} + \delta_{fpd}^{(2)} + \varepsilon_{fpdt}^{(2)} \quad (1.11)$$

We demonstrate the effectiveness of our method in section 4.3.

4 Data and Demand Estimation Results

Previous section has described the estimation of quality. We now apply this procedure to French exporters using French customs data. We start by describing the data we use, and provide descriptive statistics showing that they suit our exercise. Then, we report results on price elasticity. The estimates obtained from our empirical procedure are systematically larger, in absolute values, than corresponding OLS estimates. This is strongly suggestive that the use of our IV estimation corrects endogeneity biases described in section 3.1. Finally, we estimate quality levels by applying our methodology separately to different types of goods. We document the relevancy of our quality estimates through correlations with firm-level characteristics and existing measures of quality.

4.1 Data

We exploit firm-level trade data from the French customs. These data provide a comprehensive record of the yearly values and quantities exported and imported by French firms from 1995 to 2010. Trade flows are disaggregated at the firm, country and eight-digit product category of the combined nomenclature (CN).²⁶ Imports and exports are reported separately.

Information on quantities in trade data is known to be noisy. Because these noisy quantities are used to compute unit values, and could generate spurious correlations between quantities and prices, the data are cleaned along various dimensions. First, we drop quantities equal to one or two, since we suspect them to be subject to rounding errors or to be poorly reported by firms. Secondly, we drop prices which variations are “suspiciously” large between years, destinations, and relatively to competing products.²⁷ Finally, because of changes in the HS classification across years, we apply the algorithm described in Pierce & Schott (2012) in order to obtain well-defined and time invariant product categories.

there are really two samples: exporter-importer and exporter. Exporter-importer is sometimes reduced to 1995 importers such that flow and market fixed effects are identified.

26. Only annual values which exceeds a legal threshold are included in the dataset. For instance, in 2002, this threshold was 100,000 euros. This cutoff is unlikely to affect our study since, this same year, the total values of flows contained in the dataset represented roughly 98% of the aggregated estimates of French international trade.

27. Appendix B provides the details of the cleaning procedure.

Size of the Dataset After this cleaning procedure, the size of the dataset remains large, with more than 3 millions flows recorded every year. Table 1.1 displays the number of observations, varieties, and firms present in the exporting and importing tables of the dataset. 1995 importers are firms for which our instrument is constructed. Last column therefore gives us the actual size of the sample used to implement the estimation. We are left with approximately 45% of total observations, but since 1995 importers are large exports, their exports stand for almost 70% of the sample total.

Table 1.1: Size of the dataset

	Exports of Exporters	Imports of Importers	Exports of Importers	Exports of 1995 Importers
# Observations	29,102,408	37,804,366	25,583,171	13,257,803
# Varieties	5,144,897	10,108,471	4,074,342	1,799,738
# Firms	419,624	404,344	167,692	68,255
Share of Total Exports	100%	-	98%	67.5%

Notes: An observation is a trade flow at the firm, nc8 product, destination, year level. First column (respectively second column) contains the number of export (import) flows in the dataset. Third column reports the number of exporting flows for which importing flows are also reported for the same firm, at the same year. Fourth column reports the number of exporting flows for which importing flows are also reported for the same firm in 1995. A variety is a firm-product pair.

The estimation of demand functions requires a lot of variations in the data. First of all, due to the presence of market fixed effects, we need the instrument to vary across firms exporting to a given market. To get that, the set of source countries must differ across suppliers of a market. Secondly, because flow-specific and market-specific fixed effects will be used in the estimated equations, we need firms to serve a given product-destination for many years and product-destinations to be served by many firms simultaneously. Table 1.2 provides information about the distribution of the number of observations along different dimensions of the data. In this table, we provide statistics for the exports of the set of firms used to identify demand equations: firms being simultaneously importers in 1995 and exporters.²⁸ From table 1.2, one can see that the median importing-exporting firm sells in four different product category, to three different destinations, and imports from four different countries. The median flow (a firm-product-destination combination) is present for three years in the sample which means that flow fixed effects are identified for more than half the observations. Symmetrically, the median market is served by two firms so the market fixed effect is identified for at least half the sample. The last row is certainly the most important as it reports the number of ‘identifying observations’, i.e. the number of observations for which both the market and the flow fixed effect are identified. It appears that the price-elasticity is effectively estimated from close to 10 millions observations.

Descriptive Statistics on the Instrument The instrument crosses two informational sources: import shares and real exchange rates. Figure 1.1 reports the 1995-2010 evolution

²⁸. In section 6, where we look at the effect of low-wage countries’ competition, we use the entire set of exporters in our dataset. Similar statistics about this different set of firms can be found in table C.1 of appendix C

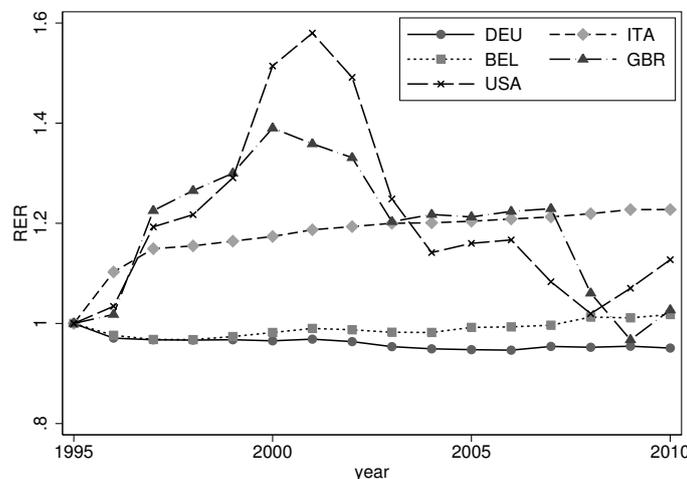
Table 1.2: Descriptive Statistics for the Exports of 1995 Importers

		p5	p25	p50	p75	p95	Mean
# Products	<i>by firm-year</i>	1	1	3	9	29	7.4
# Destinations	<i>by firm-year</i>	1	1	4	9	36	9.7
# Source countries	<i>by firm-year</i>	1	2	4	7	15	5.3
# Products	<i>by firm-country-year</i>	1	1	1	3	10	3.3
# Destinations	<i>by firm-product-year</i>	1	1	1	2	9	2.5
# Years	<i>by flow</i>	1	1	3	7	14	4.8
# Flows	<i>by market</i>	1	1	2	4	15	4.2
# Identifying Observations		9,336,602					

Notes: These statistics are from firms being exporters and importers simultaneously in our sample. A ‘flow’ is a combination of a firm, a product and a destination. A ‘market’ is a combination of a product, a destination and a year. An ‘Identifying Observation’ is a firm-product-destination-year combination which belongs to a multiple-year flow and a multiple-firm market.

of real exchange rates for the top 5 countries in terms of total French imports. After 1999, Real-exchange-rate movements of countries in Euro zone after 1999 are due solely to different rates of inflation.

Figure 1.1: RER 1995-2010-Top Source Countries



Notes: Real exchange rates are calculated as $e_{Euro,st} \times \frac{CPI_{st}}{CPI_{France,t}}$ where $e_{Euro,st}$ is the direct nominal exchange rate from Euro to j 's currency at date t . CPI is the consumer price index. After 1999, Real-exchange-rate movements of Euro zone countries are due solely to different inflation rates. 1995 real exchange rates are normalized to one.

The instruments are constructed from 1995 import shares and intend to proxy a firm's exchange rate shocks over the period. One concern is that if import shares vary a lot over time, then the instrument would be a bad proxy for real exchange rate shocks occurring towards the end of the period. This might generate a weak instrument issue. In table 1.3 we report how

import share autocorrelate over time within a firm. Even though the correlation of import shares at date t with 1995 import shares, they remain high and strongly significant. This is convincing that the instrument should not lose too much statistical power over time.

Table 1.3: Persistence of Import Shares over Time

Year t	Correlation Import Shares ($\omega_{fs1995}, \omega_{fst}$)	N
1995	1.000	185,277
1996	0.850	120,282
1997	0.795	105,671
1998	0.761	97,060
1999	0.717	89,930
2000	0.691	83,164
2001	0.676	75,518
2002	0.658	69,734
2003	0.643	64,937
2004	0.630	61,449
2005	0.611	57,496
2006	0.604	54,418
2007	0.589	51,651
2008	0.585	49,079
2009	0.577	45,568
2010	0.563	44,044

Notes: This table reports the autocorrelation of firm-country import shares over time. We drop observations for non-exporters. All correlations are significant at 1%

4.2 Estimation procedure

The validity of the identification strategy requires to take into account market-specific effects, as well as flow-specific effects, in order to make the identification within-firm, within-market. This implies to estimate two large sets of fixed effects along these two dimensions. Estimation of linear equations with two sets of high-dimensional fixed effects and an unbalanced panel is cumbersome. Indeed, because the panel is unbalanced along these two dimensions, the two sets of fixed effects are not orthogonal. Consequently, our variables need to be simultaneously projected on these two sets of fixed effects, and we cannot rely on successive projections. In order to do so, we rely on the algorithm developed in Guimaraes & Portugal (2010). This algorithm aims to demean the variables of interest along the two sets of fixed effects. Parameters of interest are then estimated using demeaned variables.

The estimation procedure is carried out in two stages: a first stage regresses the logarithm of the export prices on the firm-specific weighted exchange rate, and the sets of fixed effects. In some specifications, we will introduce the average gdp per capita of destinations and source countries as controls. Moreover, we will also run separate specifications where we introduce the

second instrument, taking into account the degree of hedging.²⁹ In order to predict variations in the logarithm of export sales, the second stage uses price predictions as an explanatory variable, as well as the other exogenous variables used in the first stage. Therefore, the parameter on price measures the price-elasticity of demand. Standard errors are adjusted to take into account this two-stages procedure. Moreover, we cluster standard errors at the market level, in order to avoid misleading inference that could potentially come from correlated demand shocks between varieties on a market.

4.3 Pooled Industries Results

In order to describe the effectiveness of the instrumental strategy, we will first present results when estimating an unique price-elasticity. The first stage of the estimation procedure shows that the instruments employed are strong enough, and impact export prices in a way consistent with economic theory. Then, we report the results of the second stage. Instrumentation corrects estimated coefficients as expected which provides support for the relevancy of our instrumental variable strategy.

First stage To build the instrument, we theorized that (i) exchange rate variations on imports input prices and (ii) that input prices impact output prices. As a preliminary test to our instrumental strategy, we test the first part of this causal chain. To do this, we regress the unit value of imports over the real exchange rates. A price is defined at the most disaggregate level: it corresponds to a firm-source country-CN 8 product category and a year. Firm-source-product fixed effects are added to the regression. Results are reported in table 1.4. As expected, real exchange rates significantly and positively impact input prices.

Table 1.4: Pass-through from Exchange-rates to Import Prices

<i>log Import Price_{fst}</i>	
log RER_{st}	0.3072*** (58.98)
Flow Effects	YES
N	15,477,199
R-square	0.003

Notes: A flow is a Prod×Dest×Year combination. *t* statistics in parentheses. *** p<0.01

Then we turn to the first stage per se. Table 1.5 shows that our instruments are strongly correlated with export prices, the endogenous variable. It presents the results of the first stage for four different specifications. Columns (1) and (2) only use the contemporaneous average exchange rate, \overline{RER}_{ft} , as a predictor of export prices. The difference between these two columns lies in the inclusion of the term capturing the potential quality adjustment following changes in GDP per capita from destination or source countries: $\overline{gpc}_{ft}^{\text{exp}}$ and $\overline{gpc}_{ft}^{\text{imp}}$. In

²⁹ The instrument strength will become a potential issue when reducing the number of observations by applying this estimation procedure separately to a number of industries.

columns (3) and (4) the specification is augmented with the second instrument that takes into account the degree of hedging, \overline{RER}_{ft}^h .

Table 1.5: First stage results

	<i>log price export</i>			
	(1)	(2)	(3)	(4)
\overline{RER}_{ft}	0.087*** (0.005)	0.092*** (0.005)	0.11*** (0.005)	0.12*** (0.005)
\overline{RER}_{ft}^h			-0.31*** (0.023)	-0.32*** (0.024)
$\overline{gdp}_{ft}^{\text{exp}}$		0.007*** (0.001)		0.007*** (0.001)
$\overline{gdp}_{ft}^{\text{imp}}$		0.012*** (0.001)		0.012*** (0.001)
N	9 336 602	9 124 226	9 336 602	9 124 226
partial F-stat	326.5	341.4	267.53	273.6

Notes: Dependent variable is the logarithm of the price of the exported good, at the firm \times nc8 \times destination \times year level. \overline{RER}_{ft} is the import-weighted exchange rate for a firm, based on its importing shares in the first year of the sample. \overline{RER}_{ft}^h is the import \times export weighted exchange rate for a firm, based on its importing and exporting shares in the first year of the sample. $\overline{gdp}_{ft}^{\text{exp}}$ is the average GDP per capita of the destinations of the firm. $\overline{gdp}_{ft}^{\text{imp}}$ is the average GDP per capita of the sources countries of the firm. Partial F-statistics are computed excluding the average GDPs per capita. Firm \times Prod \times Dest and Prod \times Dest \times Year fixed effects included in all regressions. Market-level clustered standard errors in parentheses. * p<0.1, ** p<0.05, *** p<0.01

Three main results emerge from table 1.5. First of all, the sign of the instruments' coefficients is consistent with the theoretical predictions. An increase in the average exchange rate faced by the firm is positively correlated with the price of its exported output. As an average effect, we find an elasticity of 0.9 between imported exchange rates and output prices. Moreover, we see that our second instrument is also consistent with the theory. Firms whose exporting shares are correlated with importing shares are less affected by exchange rate changes since; potentially because they are able to hedge these variations away. Secondly, the coefficients for the GDP per capita per variables are also consistent with theory. As predicted in Bastos et al. (2012), following an increase in the average GDP per capita of its destinations, a firm should upgrade its product, generating a positive impact on prices. Similarly, the average gdp per capita of source countries is positively correlated with output prices, suggesting a positive correlation between the quality of the output and the characteristics of the countries the firm imports goods from. One can notice that the introduction of these two terms does not affect the relationship between the instrument and output prices. As mentioned earlier, there is no reason to think that the cost shifter generated by exchange rates variations should induce systematic changes in the quality choices made by firms. Therefore, it is not surprising to see that these two controls do not affect the strength of the instruments. Finally, we also

observe that our set of instruments display a strong correlation with exported prices. With partial F-statistics ranging from 267 to 341, weak instruments are not an issue here.

Second stage After checking the validity of the first step, we use prices predicted by our set of instruments as an exogenous variable in the demand equation. We estimate the demand equation using the four different specifications displayed in table 1.5. Moreover, in order to assess the effectiveness of our instrumental strategy, we compare our result to a specification using OLS, that does not address the endogeneity problem. Results are displayed in table 1.6. We number columns so that second stage specifications have the same number as corresponding first stage specification in table 1.5. In addition, column (0) presents the results of the OLS specification.

Table 1.6: Second stage results

	<i>Log Export Sales</i>				
	(0) OLS	(1) IV	(2) IV	(3) IV	(4) IV
Log price $(1 - \hat{\sigma})$	0.18*** (0.002)	-1.38*** (0.20)	-1.35*** (0.18)	-0.86*** (0.14)	-0.82*** (0.13)
$\overline{gdp}_{ft}^{\text{exp}}$			0.15*** (0.003)		0.15*** (0.003)
$\overline{gdp}_{ft}^{\text{imp}}$			0.027*** (0.003)		0.021*** (0.002)
Instrument	.	Single	Single	Hedg.	Hedg.
N	13 938 430	9 336 602	9 124 226	9 336 602	9 124 226

Notes: Dependent variable is the logarithm of export sales, at the firm \times nc8 \times destination \times year level. Log price is the prediction from the first stage. $\overline{gdp}_{ft}^{\text{exp}}$ is the average GDP per capita of the destinations of the firm. $\overline{gdp}_{ft}^{\text{imp}}$ is the average GDP per capita of the sources countries of the firm. Firm \times Prod \times Dest and Prod \times Dest \times Year fixed effects included in all regressions. Market-level clustered standard errors in parentheses, adjusted for the two stages estimation procedure. *** p<0.01

This table demonstrates the effectiveness of our instrumental strategy. The coefficient for the OLS regression in column (0) is biased due to simultaneity and measurement errors problems. Whereas measurement errors drive the estimate toward zero, the simultaneity problem generates a positive bias on the estimation of the elasticity. These predictions are confirmed with a positive coefficient of 0.18 for the OLS specification. By contrast, when using our sets of instrumental variables, the estimates for the price coefficient is lower, ranging from -0.82 to -1.38. This implies estimates of the price-elasticity of demand $(-\hat{\sigma})$ ranging from -1.82 to -2.38. that are consistent with the recent findings in the literature.³⁰ Moreover, we can see that the coefficients on the variables $\overline{gdp}_{ft}^{\text{exp}}$, $\overline{gdp}_{ft}^{\text{imp}}$ are also consistent with the

30. Recent papers estimating firm-level demand functions include Nevo (2000), who finds estimates between -2.2 and -4.2 in the cereal industry, Dubé (2004) who gets estimates between -2.11 and -3.61 in the soft drinks industry. Some recent studies estimate firm-level price-elasticities for several industries. Foster et al. (2008) obtains a mean estimate of -2.41 with eleven homogeneous industries, Handbury (2012) finds a mean of -1.97 with 149 industries, and Gervais (2011) a median of -2.11 with 504 products.

theory, since they reveal that products sourced and supplied to richer countries are of better quality (i.e. they are more demanded, conditional on price). Finally, it is noteworthy that the estimates are quite consistent across specifications, even though the specifications with two instruments seem to generate slightly bigger coefficients.³¹

Estimating a single coefficient for all industries shows that instrumenting affects estimates in a direction consistent with a correction of the simultaneity bias in demand equations. However, in order to infer quality measures from these demand equations, we separately apply this method to different product categories.

4.4 Demand Estimation by Industry

In this section, we describe the results obtained by replicating the instrumentation strategy separately for fifteen product categories.³² We use the set of instruments displayed in column (4) of table 1.5. As a way to make our first stage as strong as possible, this specification includes the instrument taking into account the degree of hedging as well as both GDP per capita control variables.

Product-specific price-elasticity estimates The results of this procedure are displayed in table 1.7. For each product category, we report the IV and OLS estimates of the price-elasticities of demand, as well as the F-statistics of the first stage of the instrumental variable procedure.

We can see from table 1.7 that, for almost all industries, the IV coefficient is higher in absolute values than the OLS coefficients. This is consistent with our instrument correcting with the simultaneity bias that links quality and prices in demand equation. While the OLS estimates seem strongly driven by measurement errors, our IV estimates are all negative, and in a range consistent with the existing literature in Industrial Organization. As an outlier, the first product category, related to Animal Products records a very large and positive elasticity, that is due to a large imprecision of this parameter. Because of this, we will exclude this product category for the rest of the paper. Overall our estimates range from -0.80 to -5.98 for vertically-differentiated goods, Foster et al. (2008) estimate price-elasticities from -0.52 to -5.93 for eleven homogenous goods. More generally, we can see that our procedure performs poorly in industries that primarily rely on natural resources: Indeed, low F-stat are recorded for Animal, Vegetable, Mineral and Wood Products. This is consistent with our identification strategy since it is likely that these industries mainly rely on domestic inputs and therefore, won't be affected by exchange rates variations.

As a way to assess the reasonableness of our price elasticity estimates, we correlate them to Sutton (2001)'s measure of vertical differentiation. Our expectation is that in vertically differentiated sectors, consumers are more sensitive to quality and less to prices. As shown by

31. This phenomenon could be related at a first sight to a weak instrument bias: while single-instrument IV regression are median-unbiased in the case of weak instruments, this is not true for specifications with more than one instrument (as in (3) and (4)). Therefore, those specifications could be affected by IV bias from weak instrument. However, this possibility is ruled out by the strength of our instruments in all specifications.

32. Unfortunately, when estimating at a more disaggregated level of the product classification, the number of observations per product category decreases and our instruments become weak in an important subset of product categories.

Table 1.7: Price-elasticity estimates ($-\sigma$) for different product categories

Product categories	OLS		IV		F-stat
	Coef	SE	Coef	SE	
<i>Animal Products</i>	-0.83	(0.015)	13.3	(20.3)	1.17
<i>Textiles</i>	-0.69	(0.004)	-0.80***	(0.14)	331.9
<i>Metals</i>	-0.81	(0.006)	-0.87*	(0.46)	22.3
<i>Vegetable Products</i>	-0.81	(0.011)	-0.93	(2.27)	1.91
<i>Foodstuffs</i>	-0.95	(0.007)	-0.97	(0.81)	11.8
<i>Machinery, Electrical</i>	-0.85	(0.004)	-1.23***	(0.29)	40.8
<i>Wood, Wood products</i>	-0.79	(0.007)	-1.27	(1.08)	2.93
<i>Chemicals and Allied</i>	-0.90	(0.006)	-1.51***	(0.63)	12.9
<i>Plastics, Rubbers</i>	-0.86	(0.008)	-2.27***	(0.68)	12.6
<i>Miscellaneous</i>	-0.76	(0.005)	-2.72***	(0.57)	9.75
<i>Transportation</i>	-0.71	(0.012)	-3.20***	(0.56)	23.85
<i>Stone, Glass</i>	-0.82	(0.009)	-4.55***	(1.03)	4.93
<i>Mineral Products</i>	-0.81	(0.022)	-4.75***	(1.80)	2.30
<i>Footwear, Headgear</i>	-0.72	(0.013)	-4.84***	(1.36)	3.6
<i>Raw Hides, Skins, Leather...</i>	-0.77	(0.010)	-5.98***	(0.86)	8.08

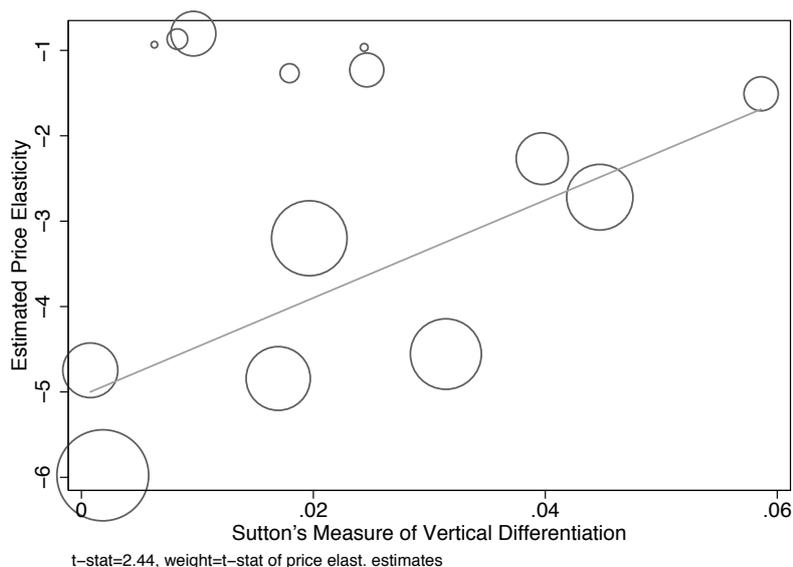
Notes: Each row corresponds to a product category for which the demand equation is estimated. The IV specifications use the average exchange rates as instruments $\overline{RE\overline{R}}_{ft}$, in addition to the hedging term, $\overline{RE\overline{R}}_{ft}^h$, and the two gdp per capita controls $\overline{gpc}_{ft}^{\text{exp}}$ and $\overline{gpc}_{ft}^{\text{imp}}$. Last column provides the value of the partial F-statistic of the first stage of the 2SLS procedure. Firm \times Prod \times Dest and Prod \times Dest \times Year fixed effects are included in all regressions. Standard errors are clustered at the market level. * p<0.1, ** p<0.05, *** p<0.01

figure 1.2, the demand faced by exporters of vertically differentiated products is significantly more elastic. This is consistent our prediction.

5 Analysis of Estimated Quality

Once demand functions have been estimated, we simply pick up quality using (1.11). As a first way to describe our estimates of quality, in table 1.8 we provide a decomposition along several dimensions. Here, it is important to remember that this quality measure is obtained at the firm \times product category \times destination \times year level. Moreover, quality is defined relatively to the average quality in the market. Therefore, it defines a position over the quality ladder in a market, rather than an absolute quality which can be compared across markets. One can see from table 1.8 that the dispersion of quality is well predicted by variety-specific effects. Indeed, half of this quality dispersion is captured by time-invariant variety-specific effects, and two thirds by time-variant variety fixed effect. From this table, it seems that the quality level of a product is strongly correlated across destinations for a specific good. We will rely on this evidence that quality choices are made at the variety level, when identifying quality

Figure 1.2: Price Elasticity Versus Vertical Differentiation



Notes: Each circle corresponds to a product category, i.e. a 1-digit position of the HS classification. The size of a circle is proportional to the absolute value of the t-statistics on $1 - \sigma$. The x-axis is Sutton (2001)'s measure of vertical differentiation, i.e. the share of R&D and advertising expenditures in a sector's total sales. It is computed at the ISIC 4-digit rev. 2 level, by Kugler & Verhoogen (2012), using US data. To compute it at the product category level, we average Sutton's measure over the different ISIC 4-digit rev. 2 a product category belongs to, and we weight by total exports over the period. The y-axis is equal to estimated price-elasticity. The line is the predicted value of a weighted OLS regression of price-elasticity over Sutton's measure. Weights are the absolute value of the t-statistics on $1 - \sigma$. "Animal Products" excluded from the regression. The relationship is not significant when it is included.

upgrading in a destination from competition shocks in other destinations served by a variety.

Table 1.8: Variance Decomposition of the quality measure

	<i>Quality</i> $\lambda_{f\text{pdt}}$			
Firm FE	✓			
Firm×Prod FE		✓		
Firm×Year FE			✓	
Firm×Prod×Year FE				✓
R^2	0.17	0.51	0.23	0.69

Notes: Each column describes the regression of our quality measure from table 1.7 on a different set of fixed effects.

We can see that there is still a wide variation across destinations. Controlling for Firm×Prod×Year FE, we can predict 69 percents of the variation of our quality measure. Moreover, half of the variation between firms are captured by firm×product fixed effects.

5.1 A consistency test: Estimated versus Expert Assessed Quality

In order to assess the relevancy of our measure, we relate it to the only objective quality measure that we are aware of in the literature. Crozet et al. (2012) take advantage of expert ratings for Champagne to analyze the importance of quality in explaining international trade flows at the firm-level. These expert assessed ratings (initially from Juhlin (2008)) are expressed in number of stars ranging from 1 to 5, one being the lowest quality. We non-parametrically regress our revealed measure of quality for Champagne exports over the number of stars.

We correlate These expert assessed ratings to our revealed measure quality for Champagne exports.³³ The quality of Champagne is described by a number of stars, one being the lowest quality, and five the highest. We nonparametrically regress our measure of quality for Champagne producers to their respective assigned quality. Results are presented in table 1.9.

Table 1.9: Estimated Versus Expert Assessed Quality of Champagne Exports

	<i>Quality $\lambda_{f\text{pdt}}$</i>
2 Stars	0.060*** (0.006)
3 Stars	0.112*** (0.006)
4 Stars	1.245*** (0.005)
5 Stars	1.421*** (0.007)
N	31,116
R-squared	0.136

Notes: Dependent variable is estimated quality $\lambda_{f\text{pdt}}$. We match Champagne producers with the ratings used in Crozet et al. (2012), initially coming from Juhlin (2008). A larger number of star means a higher expert assessed quality. We drop non-Champagne exports of Champagne producers. Robust standard errors in parentheses.
*** $p < 0.01$

We can see from table 1.9 that our measure of quality is monotonically increasing with the number of stars assigned by Juhlin (2008). Even though Champagne is a specific good in many dimensions, and cannot assess the overall quality of our measure, this is reassuring about the relevancy of our estimated measure of quality.

5.2 Correlation with firm-level variables

In order to further improve our understanding of the characteristics of our quality measure, we relate its estimated value to firms' characteristics. We merge our estimated qualities with

33. We thank the authors for sharing their data

firm-level data from France. We use the dataset BRN, that covers all French firms with revenue larger than 763 Keuros, and is constructed from reports of French firms to the tax administration.³⁴ Therefore, we are able to inspect how our quality measure is able to explain firm characteristics such as prices or average wages. Table 1.10 inspects these correlations using the number of employees of the firm, its average wage, and our estimates of quality.

Table 1.10: Estimated Quality and Firm characteristics

		<i>Log average wage</i>		
	(5)	(6)	(7)	(8)
Est. quality	0.008*** (0.001)	0.009*** (0.001)	0.006*** (0.001)	0.007*** (0.001)
log(employment)		-0.010*** (0.004)		-0.010*** (0.004)
log(export price)			0.032*** (0.003)	0.032*** (0.003)
N	13 468 383	13 468 383	13 468 383	13 468 383

The results displayed in table 1.10 are reassuring about the relevancy of our measure. First of all, we can see in panel B that our estimated measure of quality is again a predictor of the wages paid by the firm. Interestingly once estimated quality is controlled for, the sign of the wage-size correlation is significantly negative. Without controlling for quality, the literature finds a positive relationship between size and wages (see Verhoogen (2008)). The theoretically explanation is that (i) wages contain information on workers' quality and (ii) firms employing higher quality workers are larger. In these theories, the positive relationship may be turned to negative as soon as firms are assumed vertically homogeneous. Therefore, we interpret the fact that we get a negative coefficient on employees as a signal that estimated quality is a good control for workers' quality.

5.3 How well do Prices proxy for Quality?

As a last way to analyze the properties of our measure of quality, we look at the relationship between estimated quality and export prices. This is an important point since prices have been extensively used in the literature as a proxy for quality. The problem is that prices are supposedly also a function of a firm's production cost. Therefore in sector with little vertical differentiation, prices should poorly capture differences in demand fundamentals, whether across firms or over time. To test this intuition, we regress (log) prices over estimated quality and we allow the slope of the relationship to depend on Sutton's measure of vertical differentiation.

Results for this exercise are reported in table 1.11. First finding is that there is a positive relationship between estimated quality and prices in all sector. However, the slope of that relationship is significantly steeper in more vertically differentiated industries, consistently

34. This dataset has been widely used in the literature (see Eaton, Kortum & Kramarz (2011) or Berman et al. (2012) for instance).

Table 1.11: Relationship between Prices and Quality across Sectors

	$\log \text{Export Price}_{f\text{pdt}}$	
	(1)	(2)
Quality $\lambda_{f\text{pdt}}$	0.033*** (0.000)	0.025*** (0.001)
Quality $_{f\text{pdt}} \times \text{Sutton}_p$	1.260*** (0.005)	1.048*** (0.018)
Market Effects	YES	YES
Flow Effects	NO	YES
N	13 542 905	13 542 905
R-squared	0.845	0.983

Notes: Dependent variable is the logarithm of exports unit value at the firm \times nc8 \times destination \times year level. ‘Sutton’ is the share of advertising and R&D expenditures in a US sector’s sales. It is computed at the 4 digit level of ISIC-rev 4 classification by Kugler & Verhoogen (2012). A flow is a firm \times nc8 \times destination combination. A market is a nc8 \times destination \times year combination. Market-level clustered standard errors in parentheses. *** $p < 0.01$

with the intuition presented above. This is true whether we look in the cross-section of a market (column (1)) or in the dynamics of a flow (column (2)). To get a sense of the magnitude of the differences in slope across sectors, let us compare the quality-elasticity of prices between “mineral products” and “chemical and allied”, respectively the least and the most vertically differentiated product categories. In “mineral products”, the quality-elasticity of prices is approximately 0.035 when it is about 0.11 in “chemical and allied”. This means that prices are three time less informative on quality for “mineral products” than for “chemical and allied”.

6 The Quality Response to Low-Wage Countries’ Competition

In this section, we exploit our measure of quality to document the quality response of French firms to low-cost competition. We start by describing this identification strategy in the next section, and then display the results of the estimation.

6.1 Identification strategy

Following Bernard et al. (2009), we define low-wage countries’ competition (LWC) as the share of imports from countries with a GDP per capita inferior to 5% of French GDP per capita. More specifically, LWC is constructed from bilateral trade dataset BACI, according to following formula:

$$LWC_{pdt} = \frac{I_{idt}^{\text{low}}}{I_{idt}}, \quad (1.12)$$

Where I_{idt}^{low} is country d ’s imports of 6-digit HS product i from low-wage countries at date

t . Respectively, I_{idt} is country d 's total imports of product i at date t . In equation (1.12), p is an 8-digit CN product position which belongs to 6-digit HS category i .³⁵

A natural way to identify the within-firm quality response to LWC would be to regress the dynamics of the quality measure, λ_{fpdt} , over the dynamics of LWC_{pdt} . Since LWC does not vary across firms within a market, this approach would amount to looking at the impact of LWC over the mean quality of exports in a market. The problem is that our measure of quality is defined relatively to the average quality in a market. So its market-level mean is normalized to zero and is constant over time. As a consequence, identification requires variation in low-cost competition across firms, within a market.

In order to generate such variation, we make use of the information on multi-destinations exporters. Within a market, firms differ in the other markets they serve simultaneously. Therefore, for any given market, we can construct a measure of the competition faced by a firm-product variety in the rest of the world. Let LWC_{fpdt}^{ROW} be that measure and let t_{0fp} be the first year when variety fp is observed in the sample. LWC_{fpdt}^{ROW} verifies:

$$LWC_{fpdt}^{ROW} = \frac{\sum_{d' \in \mathcal{D}_{0fp}, d' \neq d} r_{0fpd'} \times LWC_{pd't}}{\sum_{d' \in \mathcal{D}_{0fp}, d' \neq d} r_{0fpd'}}$$

With \mathcal{D}_{0fp} the product-specific set of destinations of a variety $f-p$ at date t_{0fp} . r_{0fpd} is the sales of variety fp in destination d , at initial date t_{0fp} .

In the cross-section of a market, a variety with a higher LWC^{ROW} faces a fiercer low-wage competition in the rest of the world. Our identification strategy consists in correlating the dynamics of LWC_{fpdt}^{ROW} with the dynamics of λ_{fpdt} . Since competition shocks we exploit occur in a market different from the quality adjustments we intend to identify, our identifying assumption is that quality variations are correlated across destinations within a variety. In the extreme case where a variety is served with a same quality in all destinations, our strategy would capture the exact impact of a local competition shock on local quality. In general, the effect we estimate will be discounted for the fact that qualities do not perfectly co-move across destinations.

Our econometric specification is:

$$\lambda_{fpdt} = \sum_{\tau=0}^5 \beta_{\tau} LWC_{fpd,t-\tau}^{ROW} + FE_{fpd} + FE_{pdt} + u_{fpdt} \quad (1.13)$$

With FE_{fpd} a set of flow fixed effects and FE_{pdt} a set of market fixed effects. Model (1.13) identifies the effect of competition on quality, up to a five years lag. FE_{pdt} controls for the fact that competition in the rest of the world could be correlated to local competition shocks. Flow fixed effect FE_{fpd} controls for the average quality of a flow over the period. Flow fixed effects are included because in the cross-section of a market, quality might be correlated to LWC_{fidt}^{ROW} through the self-selection of firms into export markets over quality. For instance, high quality firms might self-select into markets with stronger low-wage competition. The inclusion of flow-fixed places the estimation in the dynamics of a trade flow. As we use initial export

35. Documentation about BACI can be found in Gaulier & Zignago (2010)

shares to construct LWC^{ROW} , its dynamics is not driven by some (endogenous) reallocation of exports.

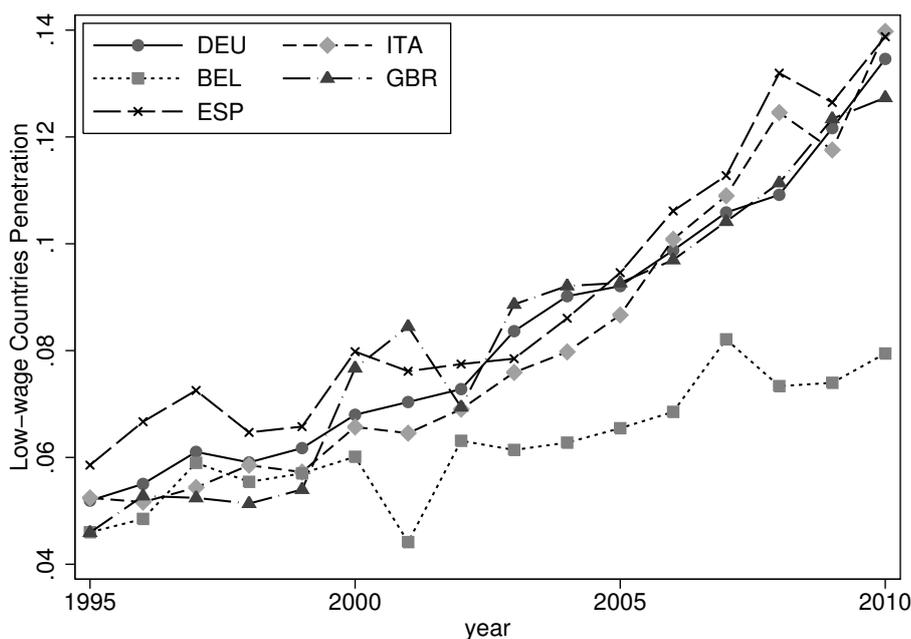
Given our fixed effect specification, our identifying assumption is that the relative dynamics of LWC^{ROW} across firm-product-destination flows, within a product-destination market are exogenous to relative quality shocks. Next subsection presents our results.

6.2 Results

In this subsection, we show the results obtained by estimating variants equation (1.13). Specifications will differ in the dependent variable we consider as well as in the number of lags we estimate.

Before turning to the core results, we report some descriptive elements on low-wage competition. In figure 1.3, we plot the penetration of low-wage countries by year in the top five largest destination markets for French exporters. Unsurprisingly, the penetration has increased steadily over the period and consistently across destinations. Our identification will come from the difference in the dynamics across countries.

Figure 1.3: Low-wage Countries' Penetration 1995-2010-Top Source Countries



Results from our main specification are reported in table 1.12. To obtain table 1.12, we have run specification (1.13) first by including each lag of rest-of-the-world competition separately and then by including all lags together. In order to make regressions comparable, we use a same sample of firms for which we observe at least five lags of quality. Overall, 1.12 suggests little quality upgrading. Only LWC_{t-4}^{ROW} is significant at 10%, when included alone. Coefficient 0.1738 associated to LWC_{t-4}^{ROW} means that after 4 years, a 10 percentage point increase in the rest-of-the-world low-cost penetration causes the firm to quality to increase by 2%. We think of the fact that the effect of competition takes time to occur as being indicative of the fact that we are indeed capturing some quality upgrading. Indeed, our measure of

quality is revealed from the demand faced by a firm. No matter the way the firm produces higher quality, it seems reasonable to think that it takes time to result into larger sales.

Table 1.12: Rest-of-the-world low-wage Competition and Quality Upgrading.

	<i>Dep. variable: Quality $\lambda_{f,p,d,t}$</i>						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
LWC_t^{ROW}	0.0296 (0.35)						0.0409 (0.49)
LWC_{t-1}^{ROW}		-0.0690 (-0.81)					-0.0868 (-1.02)
LWC_{t-2}^{ROW}			0.0654 (0.75)				0.0585 (0.67)
LWC_{t-3}^{ROW}				0.1232 (1.34)			0.0922 (1.01)
LWC_{t-4}^{ROW}					0.1738* (1.84)		0.1432 (1.52)
LWC_{t-5}^{ROW}						0.1298 (1.39)	0.1023 (1.09)
Observations	911,601	911,601	911,601	911,601	911,601	911,601	911,601
R^2	0.934	0.934	0.934	0.934	0.934	0.934	0.934

t statistics in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

In order to gain confidence into the fact that the effect we capture in table 1.12 is indeed a quality upgrading response to competition, we now interact our measure of competition we a sectoral measure of the vertical differentiation. Our prediction is that the effect of competition should be larger for more vertically differentiated sectors as firms from homogeneous sectors can not adjust their quality. This prediction is confirmed in table 1.13. More specifically, we see that for that the interaction term is significant for the third and fourth lag. This confirms the fact that the effect we identified in 1.12 is driven by firms from vertically differentiated sectors. This is reassuring that we are actually capturing a quality response.

As a robustness check, we re-run specification (1.13) by taking prices, instead of quality, as a dependent variable. The results are reported in table 1.14. This robustness test is motivated by the assumption that quality is costly to produce.³⁶ Therefore, a quality upgrading should also show up through increased prices. This is the message from table 1.14. In addition, we can notice that if the response on prices takes time, it occurs sooner the quality response in table 1.12. Combined, these results suggest that the quality upgrading operated by firms takes time to translate into larger sales (conditional on prices).

All in all, these results are very suggestive that firms upgrade their quality when the penetration of low-wage countries go up.

36. See, e.g., Hallak & Sivadasan (2011), Johnson (2012) and Kugler & Verhoogen (2012) for trade models where quality is costly and endogenous at the firm-level.

Table 1.13: Is Quality Upgrading more Significant in more Vertically Differentiated Sectors?

	<i>Dep. variable: Quality $\lambda_{f,p,d,t}$</i>						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
LWC_t^{ROW}	0.0006 (0.00)						-0.0004 (-0.00)
$LWC_t^{ROW} \times \text{Sutton}$	1.8691 (0.47)						2.4517 (0.62)
LWC_{t-1}^{ROW}		-0.0276 (-0.19)					-0.0191 (-0.13)
$LWC_{t-1}^{ROW} \times \text{Sutton}$		-2.9749 (-0.75)					-3.6405 (-0.91)
LWC_{t-2}^{ROW}			-0.0144 (-0.10)				0.0254 (0.17)
$LWC_{t-2}^{ROW} \times \text{Sutton}$			-0.7124 (-0.17)				-2.4117 (-0.55)
LWC_{t-3}^{ROW}				-0.1583 (-1.07)			-0.1474 (-1.00)
$LWC_{t-3}^{ROW} \times \text{Sutton}$				9.3622** (2.33)			8.5479** (2.08)
LWC_{t-4}^{ROW}					-0.1186 (-0.76)		-0.1187 (-0.75)
$LWC_{t-4}^{ROW} \times \text{Sutton}$					9.5673** (2.27)		8.8398** (2.01)
LWC_{t-5}^{ROW}						0.1297 (0.87)	0.1811 (1.20)
$LWC_{t-5}^{ROW} \times \text{Sutton}$						-1.3316 (-0.32)	-4.3254 (-1.03)
Observations	728394	728394	728394	728394	728394	728394	728394
R^2	0.939	0.939	0.939	0.939	0.939	0.939	0.939

t statistics in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

7 Conclusion

A recent literature has evidenced that product quality has implications for key economic outcomes such as firms' profitability or welfare inequalities. These findings make it crucial to understand the determinants of quality at the firm-level. In this paper, we have provided a necessary tool to pursue this research agenda. Namely, we have proposed a novel strategy to estimate time-varying quality at the firm-level. Our strategy is robust to unobserved vertical differentiation. It only requires firm-product information on prices, sales and imports by country.

We identify quality by estimating a demand function at the firm-product level. Quality is obtained as a residual of demand, once prices have been controlled for. In order to deal with the endogeneity of prices in the demand function, we construct a new firm-specific instrument.

Table 1.14: How Do Prices Adjust to Competition?

	<i>Dep. variable: log export prices</i>						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
LWC_t^{ROW}	0.0060 (0.32)						0.0042 (0.22)
LWC_{t-1}^{ROW}		0.0240 (1.27)					0.0169 (0.90)
LWC_{t-2}^{ROW}			0.0461** (2.33)				0.0359* (1.82)
LWC_{t-3}^{ROW}				0.0532*** (2.61)			0.0375* (1.84)
LWC_{t-4}^{ROW}					0.0688*** (3.33)		0.0514** (2.47)
LWC_{t-5}^{ROW}						0.0726*** (3.42)	0.0624*** (2.93)
Observations	911601	911601	911601	911601	911601	911601	911601
R^2	0.992	0.992	0.992	0.992	0.992	0.992	0.992

t statistics in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

This instrument interacts variations in exchange rates with firm-specific importing shares. We implement our estimation on French customs data and get a number of elements supporting the reliability of our approach.

As a first application to our method, we compare (export) prices, a widely used proxy for quality, with our export quality estimates. We find a positive and significant relationship between quality and prices, however, this relationship is weaker in more homogeneous sectors. These results hold in the cross-section as well as in the dynamics of a firm. Our findings calls for a cautious use of prices to measure quality.

Finally, we use estimated quality, along with information on low-wage countries penetration rates to identify the quality response of firms' exports to low-wage countries' competition. Our results suggest that firms upgrade their quality when competition intensifies. This result is important for policy analysis as it reveals a new channel through which exporting firms can mitigate the effect of low-wage competition.

Appendix

A A Simple Model of Endogenous Quality with Imported Inputs

In the section, we develop a partial equilibrium model with heterogeneous firms, endogenous quality, and imported inputs. The model builds upon the first variant of Kugler & Verhoogen (2012)'s model: quality only impacts variable production costs, not fixed costs. We extend the original model by assuming that production is obtained by combining a set of imported inputs rather than just a single input. The main purpose of this simple model is to ground theoretically the validity of our instrument for prices. The model formalizes the relationship between the RER's faced by a firm on its imports and its export price and hence validates the rank condition. As to the exogeneity of the instrument, the model predicts that importing shares are endogenous to quality and thus suggests that importing shares should be set constant in the instrument, which is what we do in the estimation. Moreover, the model delivers a mechanism through which quality could be endogenous to RER's on imports. This potential endogeneity of the instrument can be neutralized by controlling for a sufficient statistic also provided by the model.

In addition to its predictions on the validity of the instrument, the model delivers implications on other issues. First, we show that in the model, we can perfectly rank varieties in quality using λ . Second, on the quality response to low-cost competition, the model predicts that firms in the lower end of the quality ladder should upgrade their quality to escape competition from new entrants.

A.1 Technology

As in the model of demand developed in section 2, the unit of analysis is a variety of a differentiated final good³⁷. A variety is produced by combining inputs from different sources. For each input, a firm must decide the quality and the number of physical units involved in the production of a variety. These decisions impact the volume and the quality of the output. This production process is thus described by two functions: one for physical production, another one for the production of quality.

The physical production function is:

37. In this model, a firm is a collection of independent production lines, each line producing a variety.

$$x_{v,t} = \varphi_{v,t}^a \left(\sum_{s \in \mathcal{S}_v} \gamma_{v,s} [z_{s,v,t}]^{\frac{\kappa-1}{\kappa}} \right)^{\frac{\kappa}{\kappa-1}} \quad (1.14)$$

With $x_{v,t}$ the physical output and $z_{s,v,t}$ the quantity of input from source s involved in the production of variety v . κ is the elasticity of substitution across inputs. $\gamma_{v,s}$ is the weight of input from s in the production of v ($\sum_{s \in \mathcal{S}_v} \gamma_{v,s} = 1$). $\varphi_{v,t}$ is what Kugler & Verhoogen (2012) refer to as “capability”. As it appears in (1.14), $\varphi_{v,t}$ is of the same nature as total factor productivity: it shifts up output conditional on inputs. However, unlike physical productivity, and as will be formalized below, capability also plays a role in the production of quality. Parameter a is simply the capability-elasticity of physical output. We assume $a > 0$.

\mathcal{S}_v is the set of source countries of a firm. We take \mathcal{S}_v as fixed and given. Our sense is that making \mathcal{S}_v endogenous and varying (by assuming fixed export costs for instance) would not change the main qualitative insights of the model.

Inputs are vertically differentiated. Input quality determines output quality through following function:

$$q_{v,t} = \left[\frac{1}{2} \left(\varphi_{v,t}^b \right)^\theta + \frac{1}{2} \left(\min \{ q_{I,s,v,t} \}_{s \in \mathcal{S}_v} \right)^\theta \right]^{\frac{1}{\theta}} \quad (1.15)$$

With $q_{v,t}$ output quality and $q_{I,s,v,t}$ input s quality. The production of quality is CES in capability and in the quality of imported inputs. The innovation with respect to Kugler & Verhoogen (2012) is that a firm must decide on the quality not of a single input but of many inputs. Here we assume that different input qualities combine through a Leontieff production function. This specification is convenient as it boils the quality choice of a firm down to picking a single quality throughout all its imports. A more flexible CES form would leave our main qualitative results unaltered.³⁸

We assume $\theta < 0$ so that input quality and capability are complementary. This means that the quality upgrading obtained from a marginal increase in the quality of inputs is larger for high capability firms/varieties. This structure leads high capability firms to produce higher quality goods. Parameter b simply drives the elasticity of quality production to capability: a higher b gives a larger incentive to higher φ firms to produce high quality goods. We assume $b > 0$.

The last technology assumption regards the price of inputs. In each country, the input is produced from labor under perfect competition and constant returns to scale. Unit labor requirements are a power function of input quality. As a result, the price of input from source s with quality q_I is:

$$p_{I,s,t}^*(q_I) = w_{s,t} q_I^{\beta_s} \quad (1.16)$$

$p_{I,s,t}^*(q_I)$ is the FOB (Free on board) price of input with quality q_I labelled in s 's currency.

38. In next subsection, we discuss the fact that allowing for more substitutability across qualities plays in favor of the validity of our instrument. In that sense, the Leontieff specification is conservative.

$w_{s,t}$ is the unit wage rate in s . β_s is the elasticity of input price to quality in source s . One should think of β_s as the relative price of high to low quality in country s . As evidenced by Schott (2004), rich countries have a comparative advantage in the production of high quality goods. In terms of the model, it means that β_s is larger for poor countries. The key implication of specification (1.16) is that the optimal spatial allocation of a firm's imports depends on a firm's quality: high quality firms import high quality inputs from low β (rich) countries.

Imports of inputs involves iceberg costs. The CIF cost of an input s with quality q_I , labelled in home currency (one should think of home as France, consistently with the empirical application) is:

$$p_{I,s,t}(q_I) = e_{H,s,t} \tau_{H,s,t} p_{I,s,t}^*(q_I)$$

With $e_{H,s,t}$ the direct nominal exchange rate between home and s and $\tau_{H,s,t}$ the iceberg trade cost between home and s ($\tau \geq 1$).

Next subsection solves the optimal price, import shares and quality of the firm and draws the implications for the validity of our instrument for prices.

A.2 Prices, Quality, Import Shares and RER's

The rank condition: RER and prices A variety v faces demand (3.2). We assume that competition is monopolistic so that firms charge a constant mark-up over their marginal cost:

$$p_{v,t} = \frac{\sigma_g}{\sigma_g - 1} mc_{v,t}$$

We obtain the expression of the marginal cost of a firm (conditional on output quality) as follows. First, we use the fact that, due to the Leontieff assumption, a firm imports a single input quality. So one can invert (1.15) to get input quality as a function of output quality. By plugging this relationship into (1.16), we get input prices as a function of output prices. Finally, minimizing the production cost of a firm subject to (1.14) gives

$$mc_{v,t}(q) = \varphi_{v,t}^{-a} \left(\sum_{\mathcal{S}_v} \gamma_{v,s}^\kappa \left[\tau_{H,s,t} e_{H,s,t} w_{s,t} \left(2q^\theta - (\varphi_v^b)^\theta \right)^{\frac{\beta_s}{\theta}} \right]^{1-\kappa} \right)^{\frac{1}{1-\kappa}} \quad (1.17)$$

The marginal cost of a firm is simply a CES index of CIF import prices. Equation (1.17) formalizes the idea that prices are endogenous to quality, hence the need to instrument prices when estimating demand functions. Thankfully, equation (1.17) also provides us with a candidate instrument for prices: RER's on imports, which in terms of the model we interpret as $e_{H,s,t} w_{s,t}$. Equation (1.17) says that $e_{H,s,t} w_{s,t}$ affects output prices and thus verifies the rank condition. To be a valid instrument, the average RER on imports should also be orthogonal to quality q . We verify this in next paragraph by analyzing optimal quality.

Exogeneity Condition: Import Shares Our instrument is an import weighted average RER at the firm level. In the estimation, we set weights constant as there is a concern that they are endogenous to a firm's quality. The present model formalizes this intuition and hence justifies the use of constant weights. The expression of optimal import weights, conditional on quality is

$$\omega_{s,v,t}(q) = \frac{\gamma_{v,s}^\kappa (e_{H,s,t} \tau_{H,s,t} w_s q^{\beta_s})^{1-\kappa}}{\sum_{s' \in \mathcal{S}_v} \gamma_{v,s'}^\kappa (e_{H,s',t} \tau_{H,s',t} w_{s',t} q^{\beta_{s'}})^{1-\kappa}}$$

where $\omega_{s,v,t}$ is the share of source s in imports from variety v .

This weight is a function of quality. To better understand the way a firm sets its weights, let us write the elasticity of a weight to quality:

$$\frac{\partial \log \omega_{s,v,t}(q)}{\partial \log q} = - \left(\beta_s - \sum_{\mathcal{S}_v} \beta_s \omega_{s,v,t}(q) \right) \quad (1.18)$$

Expression (1.18) has an intuitive interpretation. When a firm upgrades its quality, it reallocates its imports towards sources in which the relative cost of quality, β_s , is low, relative to the average cost in its source portfolio, $\sum_{\mathcal{S}_v} \beta_s \omega_{s,v,t}(q)$. It follows that high quality firms import from countries with low β (i.e. developed countries, according to Schott (2004)). If the RER of a source s is correlated to its β_s , as is very likely, then it follows that the average RER of a firm is correlated to its quality, through its import shares: high quality firms import from developed countries, which have strong currencies. It is therefore necessary to fix import weights, as we do in the estimation, to guarantee the exogeneity of the instrument.

Exogeneity Condition (continued): RER and quality The optimal quality of a firm maximizes profit function:

$$\pi_{v,t}(q) = \frac{1}{\sigma} p_{v,t}^*(q)^{1-\sigma} q_{v,t}^{\rho g - 1} \tilde{P}_{m,t}(q_{v,t})^{\sigma g - \rho g} P_{m,t}^{\rho g - 1} E_{m,t} \quad (1.19)$$

we assume that exporting involves iceberg costs, so the CIF price labelled in m 's currency, $p_{v,t}^*(q)$, verifies

$$p_{v,t}^*(q) = e_{H,m,t}^{-1} \tau_{H,m,t} p_{v,t}(q)$$

It follows that the first order condition on quality is:

$$(\sigma - 1) \underbrace{\left[\frac{\sum_{\mathcal{S}_v} \beta_s \omega_{s,v,t}(q)}{1 - \left(\frac{\varphi_{v,t}^b}{2q} \right)^\theta} \right]}_{\text{Quality-Elasticity of marg. costs.}} = \underbrace{(\rho - 1) + \frac{\partial \log \tilde{P}_{m,t}(q)}{\partial \log q}}_{\text{Quality-Elasticity of the Demand Shifter exp}(\lambda)} \quad (1.20)$$

To choose their optimal quality, firms operate a quality-cost trade-off. From equation (1.20) it appears that the optimum is reached when a firm equalizes the quality-elasticity of its demand shifter to the quality-elasticity of its production costs.

Equation (1.20) implicitly defines optimal quality. It appears that optimal quality is a function of importing shares $\omega_{s,v,t}$. The rationale for that prediction hinges on the Leontieff assumption on the quality of the basket of inputs. When a firm decides to upgrade its quality, it must increase the quality imported from its whole input basket. By how much the cost of its input basket goes up as a consequence depends on the import weighted average elasticity of input prices to quality: $\bar{\beta}_{v,t}(q) = \sum_{\mathcal{S}_v} \beta_s \omega_{s,v,t}(q)$.

Importing shares are also a function of RER's. This is very intuitive: firms minimize their production cost by importing from weak currency sources. Consequently, when a RER shock occurs, firms adjust their importing share which as a result impacts their perceived relative cost of quality $\bar{\beta}_{v,t}(q)$ and eventually leads the firm to adjust its quality.

To make this mechanism more practical, consider the example of a firm importing from a developing country with a high β , say China, and from a developed country with a low β , say the USA. If yuan appreciates, then the firm reallocates its imports towards the USA, it decreases the quality-elasticity of its production costs and so the firm upgrades its quality.

The crucial implication of this discussion is that quality is potentially endogenous to RER shocks.³⁹

If this questions the validity of our instrument, note that the sign of the bias which would result from the relationship between RER and quality is unclear. To see this, let us consider previous example again. Here, the firm faces a positive cost shock (yuan appreciates) and simultaneously upgrades its quality. This suggests that the price elasticity obtained through our IV estimation is biased towards zero. Now take a symmetric situation where the dollar appreciates instead of the yuan. Then the firm reallocates its imports towards China and downgrades its quality. This case would rather suggest a price elasticity biased away from zero.

Equation (1.20) also predicts that conditional on $\bar{\beta}_{v,t}$, quality is exogenous to RER's. In terms of our estimation, this means that our instrument is valid once $\bar{\beta}_{v,t}$ is controlled for in

39. How does this result depend on the Leontieff assumption in the production of quality? Intuitively, if the firm could combine the quality of its inputs through a CES function with strictly positive elasticity of substitution, it could concentrate its imports of quality from a country with a low β , and import large physical amounts of low quality inputs from the rest of the world. Therefore, the cost of upgrading its quality would be driven by the β of the source from which it imports quality, and not from its all input portfolio. It follows that the reallocation of its physical imports induced by a RER shock would have little impact on its choice of quality. The Leontieff specification therefore is the most challenging for our instrument as it is the case where quality is the most endogenous to RER. In that sense it is a conservative assumption.

the estimation. As we think of β_s as a measure of development of a country, a natural proxy for $\bar{\beta}_{v,t}$ is the import weighted average income per capita of a firm. In section 5, we show that our price elasticity estimates are robust to whether we control on not for $\bar{\beta}_{v,t}$. This consistent with the idea that the sign of the bias, if any, is not clear theoretically.

A.3 Other Implications of the Model

The Demand Shifter as a measure of Quality In section 2, we show that the demand shifter $\lambda_{v,t}$ is a measure of quality, polluted with a quality-specific demand shifter, $\tilde{P}_{m,t}(q)$. In this paragraph we show following proposition:

Proposition 1. *If $\lambda_{v,t} > \lambda_{v',t}$, then $q_{v,t} > q_{v',t}$*

Proposition 1 simply says that despite the presence of the quality price index, $\lambda_{v,t}$ can be used to rank varieties in quality, within a market. This is a crucial result regarding the scope of application of our measure of quality.

To demonstrate proposition 1, let us consider two varieties A and B such that A has a larger demand shifter than B : $\lambda_A > \lambda_B$ (We drop time subscript in this paragraph). Assuming that firms maximize their profit implies:

$$\frac{\pi_B(q_A)}{\pi_B(q_B)} \leq 1 \quad (1.21)$$

From (1.20), it is straightforward to get:

$$\frac{\pi_B(q_A)}{\pi_B(q_B)} = \frac{\exp(\lambda_A)}{\exp(\lambda_B)} \left(\frac{mc_B(q_A)}{mc_B(q_B)} \right)^{1-\sigma_g} \quad (1.22)$$

Plugging (1.22) into (1.21) and re-arranging gives

$$\left(\frac{mc_B(q_A)}{mc_B(q_B)} \right)^{1-\sigma_g} \leq \frac{\exp(\lambda_B)}{\exp(\lambda_A)} < 1$$

Where second inequality is an assumption of the proof.

Using the fact that $mc_B(q)$ is strictly increasing and that $\sigma_g > 1$ we get

$$q_A > q_B$$

which concludes the proof.

Low-Quality Competition and Quality Upgrading In our empirical application, we are interested in the within-firm effect of low-cost competition on quality. In this paragraph, we show that our model of the firm has predictions on this question. In terms of the model, one can think of low-cost competition as an entry of firms in the lower end of the quality ladder.

The direct effect of this entry is a decrease of the low-quality price index, and, as result, an increase in the elasticity of the quality price index with quality, $\frac{\partial \log \tilde{P}_{m,t}(q)}{\partial \log q}$, in the lower end of the quality ladder. Put simply, the entry of low-quality competition makes competition relatively weaker in high-quality nests. This creates an incentive for firms to upgrade their quality in order to escape low-cost competition.

First order condition (1.20) formalizes this insight: $\frac{\partial \log \tilde{P}_{m,t}(q)}{\partial \log q}$ is a component of the return to quality. When it goes up, as is the case in the lower part of the quality ladder when low-quality competitors enter the market, incumbents optimal quality goes up.⁴⁰

B Data Trimming

Data on quantities are known to be subject to measurement errors, which could lead to spurious relationships between quantities and prices (computed by dividing values with quantities). Because variations across prices are less subject to idiosyncratic variations than quantities, we clean the data, based on their computed prices, following three dimensions.

- Observations are dropped for prices for which variations across times differ from a factor five or more. Formally, observations are dropped if $\frac{p_{f p d t}}{p_{f p d t-1}} > 2$ or $\frac{p_{f p d t}}{p_{f p d t-1}} < \frac{1}{2}$
- Observations are dropped for prices which differ from a factor five or more from the mean across all destinations. Formally, observations are dropped if $\frac{p_{f p d t}}{p_{f p \bullet t}} > 2$ or $\frac{p_{f p d t}}{p_{f p \bullet t}} < \frac{1}{2}$
- Extreme quantiles of the price distributions are censored: for each market (product \times destination \times year), observations below the 1st percentile, and beyond the 99th percentile are dropped.

Finally, for several observations, quantities are displayed in different units than weight. We convert these units in weight by regressing weights on units at the product \times year level. Therefore, we are able to back-up the weight equivalent of these units.

40. Here we run a partial equilibrium analysis. Since all incumbents face the shock in the price index schedule, they all face the same incentive to upgrade their quality. As a result, in general equilibrium, quality-specific price indices may react to the entry of new competitors indirectly, through the quality response of incumbents. Our intuition is that this general equilibrium effect should dampen, but not reverse the quality upgrading response of firms.

C Descriptive Statistics

Table C.1: Descriptive Statistics for all exporters

		p5	p25	p50	p75	p95	Mean
# Products	<i>by firm-year pair</i>	1	1	2	5	21	5.7
# Destinations	<i>by firm-year pair</i>	1	1	2	4	18	4.5
# Products	<i>by firm-country-year comb.</i>	1	1	1	2	9	2.9
# Destinations	<i>by firm-product-year comb.</i>	1	1	1	2	8	2.3
# Years	<i>by flow</i>	1	1	3	6	13	4.3
# Flows	<i>by market</i>	1	1	2	4	20	5.7

Notes: A ‘flow’ is a combination of a firm, a product and a destination. A ‘market’ is a combination of a product, a destination and a year.

D Robustness checks

Table D.1: Robustness checks

	Base (1)	No hedging (2) (3)		Long diff. (4) (5)		No crisis (6)
<i>First stage:</i>						
\bar{RER}_{ft}	0.092*** (0.005)	0.083*** (0.012)	0.089*** (0.008)	0.071*** (0.007)	0.099*** (0.008)	0.083*** (0.006)
$\bar{gdp}_{ft}^{\text{exp}}$	0.007*** (0.001)	0.009*** (0.002)	0.010*** (0.001)	0.004*** (0.001)	0.008*** (0.002)	0.006*** (0.001)
$\bar{gdp}_{ft}^{\text{imp}}$	0.012*** (0.001)	0.006*** (0.001)	0.008*** (0.001)	0.008*** (0.001)	0.012*** (0.001)	0.011 (0.001)
<i>Second Stage:</i>						
Log(Price)	-1.35 (0.18)	-1.30 (0.51)	-1.67 (0.33)	-2.38 (0.31)	-1.93 (0.26)	-0.89 (0.22)

Notes: Specification (1) is the baseline from column (3), table 1.5. Columns (2) and (3) dropped importers who export to the same country: specification (2) does it for a given year, specification (3) for any year in the sample. Specifications (4) and (5) respectively use 3 and 5 years differences instead of flow fixed effects. Finally, specification (6) drops years posteriors to 2007 to avoid the role played by the trade collapse phenomenon. All specifications use the gdp per capita controls in the second stage, even though the results are not displayed.

E List of Low-Wage Countries

Table E.1: Low-Wage Countries

Angola	Djibouti	Lao People's Rep.	Rwanda
Armenia	East Timor	Lesotho	Senegal
Azerbaijan	Eritrea	Liberia	Sierra Leone
Bangladesh	Ethiopia	Madagascar	Solomon Islands
Benin	Gambia	Malawi	Sri Lanka
Bhutan	Georgia	Mali	Sudan
Bolivia	Ghana	Mauritania	Tajikistan
Burkina Faso	Guinea	Moldova, Rep. of	Tanzania, United Rep of
Burundi	GuineaBissau	Mongolia	Togo
Cambodia	Guyana	Mozambique	Turkmenistan
Cameroon	Haiti	Nepal	Uganda
Central African Republic	India	Nicaragua	Ukraine
Chad	Indonesia	Niger	Uzbekistan
China	Iraq	Nigeria	Viet Nam
Comoros	Kenya	Pakistan	Yemen
Congo	Kiribati	Papua New Guinea	Zambia
Ivory Coast	Kyrgyzstan	Philippines	Zimbabwe

Notes: A low-wage country is defined as a country which GDP per Capita in 2002 is inferior to 5% of the French one in 2002.

Chapter 2

The Interdependence of Firms' Export Decisions

Abstract

This paper provides a simple mechanism through which firms' sales are interdependent across countries. While most of the "new trade theory" models assume that firms maximize their profits independently in each country they serve, we show that this assumption can not hold as soon as the production process features uncertainty. Since firms discover consumers' tastes after their production, they may *reallocate* their production towards countries with a higher demand shock. We derive some predictions of this reallocation effect on prices, output and sales across firms' destinations. In order to test our results, we use the unexpected fall of demand for Champagne in the US in 2001. Based on the reactions of French exporters, we find convincing support for our results.

1 Introduction

In 2009, French consumers enjoyed an unusually large offer of Champagne bottles at remarkably cheap prices. At the time, analysts explained the drop in Champagne prices by the dramatic deterioration in demand conditions in the US induced by the 2008 economic crisis.¹ It was argued that Champagne producers had reallocated the production that they expected to sell in the US toward markets less impacted by the crisis, including the domestic market.²

This anecdote illustrates the crux of our paper: price and quantity decisions of an exporting firm in a destination market may not only depend on demand conditions in this precise destination market but also on demand conditions in the rest of the world. Interestingly, the anecdote is also suggestive as to the origin of the spatial interdependence. The reallocation of sales from the US to other countries arises because, at the time they produced, Champagne exporters had not anticipated the negative shock that would hit the US in the following year. This mechanism is potentially not specific to the Champagne industry: in most industries, production takes time so output cannot adjust instantaneously to demand shocks.

In the present paper, we develop a model where exporting firms base their production decisions on the anticipation of future demand. When facing non expected local demand shocks, firms reallocate their output across destinations. Interestingly, this prediction is absent from most of the “new trade theory”. Textbook models of trade (see, e.g., Krugman 1980, Eaton & Kortum 2002, Melitz 2003) rely on a series of assumptions (e.g., markets are segmented, marginal cost are constant, fixed costs are market specific, demand is certain at the time of production) which make a firm’s profit function additively separable across markets. As a consequence, firms maximize their profits independently on each market.

Empirically assessing whether firms sales are interdependent across markets can therefore be seen as a test of mainstream trade theory which is of importance for our understanding of exporting firms. The way sales interact across markets also has potential implications for the relationship between international trade and business cycle co-movements. Standard theory predicts that international trade increases the impact of foreign demand shocks on the domestic economy as firms make part of their sales abroad.³ Introducing interdependence of firms sales across markets complicates the picture as foreign demand shocks also impact firms sales at home indirectly, via their impact on firms exports. Depending on whether exports and sales at home are substitutes or complement, firms sales interactions might weaken or strengthen the impact of trade on international business cycle co-movements.

The baseline of our model is to introduce uncertainty in the production process as a source of spatial interdependence in firms’ export decisions.⁴ The main result of our simple theory is that when demand is not certain at the time of production, a demand shock hitting a firm in one market (i) negatively impacts the firm’ sales in other markets and (ii) positively

1. See, e.g., Faljaoui (2009).

2. The UK and China were also cited as alternative destination markets.

3. There is also an international specialization channel which rather asynchronizes cycles when opening to trade. See Gruben et al. (2002) for an explanation and an empirical assessment of these different channels.

4. The presence of uncertainty in manufacturing has generated a lot of works in operational research (see Mula et al. (2006) for a review).

impacts prices charged by the firm in other markets. While our model accounts for the reallocation of sales across countries, we are silent on the reallocation across time. Abstracting from this dimension means we focus on firms or industries with prohibitive storage costs.⁵ Consistency with our model therefore imposes to consider industries with high inventory costs in the empirical analysis. This motivates, among other things, the choice of the Champagne industry.

In a second step, we bring the model to the data and test its main predictions. We rely on a difference-in-differences (DD) regression approach. The scenario to our identification is one where a group of firms serve a market A and only a subset of those, the treated group, simultaneously serve a second market B . When a negative demand shock hits market B , and assuming that firms are “comparable” across groups, our model predicts that sales in market A should grow more at the time of the shock, while prices should grow less. We use our model to structurally derive the DD regression that we estimate. This allows us to get a structural expression of the treatment effect as well as to put usual DD identifying assumptions in terms of models' parameters. Thus, our main identifying assumption is that firms cost shocks and demand shocks in B at the time of the shock are on average equal across groups.

The demand shock we use to test our model is the drop in demand in the US entailed by the 2000-2001 economic recession. This demand shock is of interest for several reasons. First of all, this shock was unexpected. Second, to observe reallocation we need a shock in a relatively large market. The US are the first non European trade partner of France. One concern might be that the 2000-2001 crisis was not circumscribed to the US and contaminated markets where we try to identify a reallocation effect. However, the structural expression of the treatment effect tells us that our approach does not require the demand shock to be zero at the time of the financial crisis, outside the US. It simply requires the shock to be relatively larger in the US. This is consistent with the idea that some elements which led to the recession were specific to the US (e.g., the 11/9 terrorist attack), and it is confirmed in our data.

In our analysis, we decide to focus on the Champagne industry as it presents several advantages. First, Champagne is a highly traded good. Second, in most cases, exporters of Champagne are also Champagne producers. Third, the US are the main export market for Champagne. Fourth, producing Champagne takes time which is a crucial assumption in our theory. Fifth, storing Champagne bottles is very costly which limits temporal reallocation. Last, Champagne has to be produced in France which rules out potential offshoring reaction following demand shocks.

To conduct our analysis, we make use of three datasets. The first, called BRN, reports information on firms' domestic production and firms characteristics. The second is provided by French customs. It reports bilateral value and quantity exported by French firms at a very low level of disaggregation. The third dataset is the expert assessed rating of Champagne established by Juhlin (2008) and merged to French customs data by Crozet et al. (2012). This last source provides us with a quality measure which we use as a control in our DD approach. It also allows us to identify Champagne producers in the BRN dataset. The merge

5. Such high costs of storage can be rationalized either by assuming that it is difficult to store the good for technical reasons (perishable goods) or because the opportunity cost of storage is too high (a firm has to sell its good to repay its loans).

of these dataset results in a panel of 196 Champagne producers over the period 1995-2007.

We estimate a treatment effect on three firm-level outcome variables: the sales and prices of non-US exports and the sales at home. The treatment group is the group of Champagne producers firms that were exporting in the US before 2000. The basic idea of the test is that following a negative unexpected shock in the US, *ceteris paribus*, a firm that planned to sell in the US should increase its sales in other markets in a greater amount than a firm that did not planned to serve the US market.

We find evidence of a reallocation effect on non-US exports. More specifically, firms which were present in the US before 2000 decreased 40 percent less their exports to the rest of the world between 1999 and 2000 than firms which did not export to the US before 2000. The sign of this coefficient is consistent with our theory: treated firms reallocated their surplus of sales from the US to the rest of the world. We use the structural interpretation of the treatment effect to evaluate its magnitude. Our theory requires that firms under-estimate their non-US sales in 2000 to obtain a positive treatment effect. The larger the under-estimation, the larger the treatment effect. Our approach consists in using the structural expression of the treatment effect to compute the anticipation error implied by the model. We compare these anticipation errors to those obtained from simply assuming that firms make extrapolative anticipations. This analysis reveals that the treatment effect is too large in the sense that it implies an average anticipation error ranging from 3.25 to 3.6 times the average anticipation error implied by extrapolative anticipations.

Thus, we investigate potential biases explaining a large estimated treatment effect. We consider two channels. The first one is on the demand side: treated and untreated firms differ in terms of quality, therefore the large treatment effect might result from (high quality) treated firms facing a more favorable demand shock during the crisis. The second channel is on the supply side. 2000 crisis was a financial crisis so that untreated firms, which are smaller, might have experienced a larger deterioration of their financing conditions. This could also bias the estimated treatment effect upward. For both channels, we run a DD regression augmented with a full set of year dummies interacted with a control for the channel under interest. The controls are respectively Juhlin (2008) quality ratings and the working capital requirement ratio, which proxies the degree of liquidity constraint at the firm-level. These alternative specifications do not significantly modify the estimated treatment effect. This suggests that the treatment effect captures a reallocation effect.

In addition to works cited so far, our paper is related to different strands of the existing literature. There is a small, yet rapidly growing, number of papers empirically looking at the interaction between firms' sales across markets. Similarly to us, most of them find some substitution across markets. However, they do not derive their evidence from an exogenous source of demand variation, rather from correlations so that no causal evidence of substitution between markets has been established so far. Liu (2012) and Ahn & McQuoid (2012) find a negative correlation between the growth of exports (or the change of export status) and the growth of domestic sales at the firm level. Nguyen & Schaur (2010) and Vannoorenberghe (2012) show that the shape of the relationship between the firm-level volatility of revenues and openness is consistent with home sales and exports being substitutes.

To our knowledge, Berman et al. (2011) is the only study attempting to analyze the causal relationship from export to sales at home. Incidentally, they are also the only one to find firms' sales to be complementarity across markets. We contribute to that literature by demonstrating for the first time that firms substitute sales across foreign markets.

In terms of modeling, our paper is close to Evans & Harrigan (2005). The authors develop a model of outsourcing with demand uncertainty. Their model implies that outsourced production locates nearby the consumption market when demand is volatile, rather than remote low wage countries when demand is less volatile. While the scope of the two papers are different, our models are very close. Namely, both point out the importance of demand uncertainty and introduce the lag between the decision to serve the demand and the actual service of the demand as a key feature.

The industrial organization literature has extensively studied the impact of demand uncertainty of firms' decisions Leland (1972). Here we consider the case where firms cannot adjust their production when they learn the demand, while prices, as well as the distribution of its production across-countries, are flexible. Besides, the sequential process we assume at the firm level is related to numbers of papers focusing on investment. In the literature the sequentiality reflects the timing between the investment in production capacity that is done under uncertainty and the production realized once demand is known (Dixit & Pindyck 1996)e.g.. Rather than investing in production capacity in the first period, we assume that firms produce in the first period.

Last, in the empirical part we make use of firm level data provided by French customs. This database starts being increasingly used (Eaton, Kortum & Kramarz 2011, Berthou & Fontagne 2008, Crozet et al. 2012)e.g.. Our focus on the Champagne industry directly relates our work to Crozet et al. (2012).

The rest of the paper is organized as follows. Section 2 develops a simple model of exporting firms with demand uncertainty. Section 3 presents our difference-in-differences approach to test the model. Section 2 describes and motivates the use of trade and production data on Champagne exporters. Section 5 provides stylized facts on the US crisis. Section 5 reports the result from the DD regression. Section 7 concludes.

2 A Simple Model of Exporting Firm with Demand Uncertainty

We develop a partial equilibrium model with heterogeneous exporting firms. We depart from traditional trade models by assuming that demand is uncertain at the time when output is produced. The model aims at delivering predictions on the way a demand shock in one market affects firm' sales in other market. The model is silent about the extensive margin of trade as we assume firms' destination portfolio, \mathcal{D}_i , to be exogenous and constant over time, irrespective of demand shock realizations.⁶

6. Although it might be that a demand shock in one market induces firms to enter and exit other markets, we abstract from this question in the paper. This choice is motivated by the fact that there is very little entry and exit for the sample of firms and the set of markets we examine in the empirical section.

Let firm j face an isoelastic demand function

$$q_{ijt} = \beta_{ijt} p_{ijt}^{-\varepsilon} \quad (2.1)$$

, with q_{ijt} and p_{ijt} respectively the quantity and price of exports from firm i to country j at date t , β_{ijt} the demand shifter corresponding to that trade flow and $\varepsilon > 1$ the price elasticity of demand.⁷

At each period t , the timing of the model is as follows. First, firms determine their optimal amount of output Q_{it} . Then consumer tastes $\beta_{it} \equiv \{\beta_{ijt}\}_{i \in \mathcal{D}_i}$ are realized and firms decide on the optimal allocation $\mathbf{q}_{it} \equiv \{q_{jit}\}_{i \in \mathcal{D}_i}$ of the output across destinations. Firms can not store output so they solve a static problem. In the presence of storage, firms could adjust demand shocks through an intertemporal reallocation of their output. Since the focus of the paper is on the spatial rather than intertemporal interdependence of sales, we ignore the question of inventories in the model. This omission implies that model's predictions should be tested on industries with large storage cost (hence our choice of the Champagne industry, as explained further in the empirical section).

Let us now solve the optimal policy of the firm backward.

Optimal allocation of the Output At the end of a period, a firm maximizes the sum of its sales across destinations

$$\sum_{j \in \mathcal{D}_i} \beta_{ijt} p_{ijt}^{1-\varepsilon},$$

subject to the constraint that the total quantity shipped is at most equal to its output

$$Q_{ijt} \geq \sum_{j \in \mathcal{D}_i} \tau_j \beta_{ijt} p_{ijt}^{-\varepsilon} \quad (2.3)$$

,

with \mathcal{D}_j the set of destinations of firm i and $\tau_j > 1$ an iceberg trade cost. This form of trade costs implies that a firm should ship τ_j unit of output for 1 unit to reach the final consumer in j . Under CES preferences, the marginal revenue is always strictly positive so that constraint (2.3) is always binding, irrespective of Q_{it} and demand shifter realizations.⁸

The expression of the optimal price charged by a firm is

7. One could derive demand function (3.2) from a general equilibrium model with monopolistic competition and CES preferences with elasticity of substitution ε . Letting Y_{jt} be the income of the representative consumer in market j , b_{ijt} be the taste for the variety supplied by firm i to market j and \tilde{P}_{jt} be the quality-adjusted price index in market j , we obtain a demand function similar to Baldwin & Harrigan (2011):

$$\begin{cases} q_{ijt} &= \beta_{ijt} p_{ijt}^{-\varepsilon} \\ \beta_{ijt} &= \frac{Y_{jt}}{(b_{it} \tilde{P}_{jt})^{1-\varepsilon}} \end{cases} \quad (2.2)$$

So a higher income Y_{jt} , a higher perceived quality b_{it} or a higher price index \tilde{P}_{jt} shift consumers' demand upwards.

8. Under alternative preferences, it might be optimal for the firm not to exhaust its output. It is for instance the case under linear preferences if the root of the marginal revenue function is smaller than total output Q_{it} .

$$p_{ijt}(Q_{it}) = \frac{\varepsilon}{\varepsilon - 1} \lambda_{it}(Q_{it}) \tau_j \quad (2.4)$$

where $\lambda_{it}(Q_{it})$ is the lagrange multiplier associated to the output-clearing condition (2.3), i.e. the marginal value of output at the end of period. It is a function of total output Q_{it} because a larger output makes constraint (2.3) less binding.

To understand equation (2.4), note that at the end of a period, the cost of shipping a unit of output to market i is not a production cost, since the output has already been produced, but an opportunity cost of shipping one less unit to other destinations. Equation (2.4) states that the firm charges a constant mark-up $\frac{\varepsilon}{\varepsilon-1}$ over this opportunity cost $\lambda_{it}(Q_{it})$.

Notice that the firm charges the same free-on-board price $p_{it}(Q_{it})$ in all markets. We obtain the expression of this price as a solution to (2.3),

$$p_{it}(Q_{it}) = \left(\frac{\sum_{i \in \mathcal{D}_i} \tau_j \beta_{ijt}}{Q_{it}} \right)^{\frac{1}{\varepsilon}} \quad (2.5)$$

. Plugging the expression of optimal prices above into the demand function, we get the expression of optimal sales:

$$r_{ijt}(Q_{it}) \equiv p_{ijt}(Q_{it}) \times q_{ijt}(Q_{it}) = \beta_{ijt} \tau_j^{1-\varepsilon} \left(\frac{Q_{it}}{\sum_{j' \in \mathcal{D}_i} \tau_{j'} \beta_{ij't}} \right)^{\frac{\varepsilon-1}{\varepsilon}} \quad (2.6)$$

. Let $\rho_{ijt} \equiv \frac{r_{ijt}}{\sum_{j' \in \mathcal{D}_i} r_{ij't}}(Q_{it})$ be the share of sales to market j in the total sales of firm i at date t . ρ_{ijt} verifies

$$\rho_{ijt} = \frac{\beta_{ijt} \tau_j^{1-\varepsilon}}{\sum_{j' \in \mathcal{D}_i} \beta_{ij't} \tau_{j'}^{1-\varepsilon}}$$

. Interestingly, ρ_{ijt} is independent of Q_{it} . This implies in particular that, all things equal, the distribution of sales whether one assumes that output is fixed before demand is realized or after, as traditionally assumed. In other words, even if demand uncertainty modifies the impact of a local demand shock on the level of sales in all markets (as we show below), it has no impact on the relative sales across markets.

Optimal Output At the beginning of the period, firms maximize their expected profit as a function of their total output. Firms produce under constant returns to scale so the profit function is

$$\pi_{it}(Q) = \mathcal{E}_{it} \left[\sum_{j \in \mathcal{D}_i} r_{ijt}(Q_{it}) \right] - c_{it} \cdot Q_{it} \quad (2.7)$$

, with c is the marginal cost of production and $\mathcal{E}_{it}[\cdot]$ is firm i 's expectancy operator at the beginning of period t with respect to the joint distribution of demand realizations β_{it} . Taking

expressions of optimal sales (2.6) we get optimal production

$$Q_{it} = \left(\frac{\varepsilon}{\varepsilon - 1} c \right)^{-\varepsilon} \left(\mathcal{E}_{it} \left[\sum_{i \in \mathcal{D}_i} \tau_j \beta_{ijt} \right]^{\frac{1}{\varepsilon}} \right)^{\varepsilon} \quad (2.8)$$

. In a very standard fashion, optimal output is decreasing in marginal cost and increasing in the expected global demand term $\sum_{i \in \mathcal{D}_i} \beta_{ijt}$. Remark that the expectancy operator applies to a concave function of the global demand term. Therefore optimal production is decreasing in the variance of $\sum_{i \in \mathcal{D}_i} \beta_{ijt}$.

Predictions on the Impact of a Local Shock Let a “demand shock” be a variation in the realization of β_{it} leaving expectations $\mathcal{E}_{it}[\cdot]$, and hence total output, unchanged.

Proposition 2. (THE REALLOCATION EFFECT): *A local demand shock (i.e. a positive shock on a single market):*

- negatively impacts a firm’s sales to other markets.
- positively impacts prices charged by a firm to all markets.

Proof. For any j, j' , $j \neq j'$, $\frac{\partial r_{ij't}(Q_{it})}{\partial \beta_{ijt}} < 0$. For any j , $\frac{\partial p_{it}(Q_{it})}{\partial \beta_{ijt}} > 0$ □

Next subsection describes our empirical approach to test this proposition.

3 Empirical Methodology: A Structural Diff-in-diff Approach

We will be testing the reallocation effect predicted by proposition (2) through a DD approach. The scenario to our DD identification is one where a demand shock hits market A (and potentially, to a lesser extent, market B and C) at date t^* . We are interested in using the impact of this shock on firms’ sales and prices in market B to identify the reallocation effect described in proposition (2). Our estimator compares the export dynamics in market B of firms serving market B but not A (the control group), to firms serving both markets A and B (the treated group).

We now derive the structural expression of the estimator. This allows us to get a clear prediction as to the determinants of the magnitude and the sign of the treatment effect that we estimate. It is also of interest to interpret the standard DD identifying assumptions in terms of model’s parameters. Let T (for “Treated”) be the set of firms serving both A and B and let NT (for “Non-Treated”) be the set of firms serving B but not A . We will test the model by estimating DD model

$$\begin{cases} r_{iBt} = \exp(\alpha_r T_i \times S_t + \gamma_r T_i + \lambda_{rt} + \pi_r X_{it} + \eta_r) u_{rit} \\ p_{iBt} = \exp(\alpha_p T_i \times S_t + \gamma_p T_i + \lambda_{pt} + \pi_p X_{it} + \eta_p) u_{pit} \\ E[u_{rit}|t, T_i, X_{it}] = 1; E[u_{pit}|t, T_i, X_{it}] = 1 \end{cases} \quad (2.9)$$

, with $T_i \equiv \{i \in T\}$ a dummy variable equal to 1 if i is in T , $S_t \equiv \{t = t^*\}$ a dummy variable for the year of the shock, λ_t a set of year dummies, X_{it} a vector of control variables and (u_{rit}, u_{pit})

residuals. We are interested in imposing distributional assumptions on model's parameters which give rise to specification (2.9). The first assumption regards firms' anticipations:

Assumption 1. (EXPECTATIONS): Let $\mathcal{P}_{ijt}(\cdot)$ be the probability distribution of β_{ijt} , from the point of view of firm i , at the beginning of period t . At the beginning of any period $t \neq t^*$, firms perfectly anticipate the realization of demand shifters β_{it} : $\mathcal{P}_{ijt}(\beta_{ijt}) = 1$. At the beginning of date t^* , firms predict the realization of β_{ijt} with an error factor $1 + \delta_{ij}$: $\mathcal{P}_{ijt^*}((1 + \delta_{ij})\beta_{ijt^*}) = 1$.

Assumption 1 amounts to considering that the shock occurring at date t^* is unexpected to firms so that they make an anticipation error δ_{ij} . If the shock is negative in j , δ_{ij} will be positive as firms over-estimate the realization of the demand shifter. This assumption that the shock is unanticipated motivates our choice of the 2000-2001 US economic crisis to estimate model (2.9). Indeed, the crash of the bubble on IT technologies as well as the September 11th terrorist attack, which both triggered an economic slowdown in the US, were arguably unexpected by economic agents. Although very stylized, our assumptions on expectations make it possible to easily derive a structural expression of the treatment effect without losing any of the qualitative features of a model with more general distributional assumptions on anticipations.

Next assumption regards cost and demand shifter dynamics:

Assumption 2. (COMMON TREND):

$$E \left[\beta_{iBt} \tau_B^{1-\varepsilon} \left(\frac{\varepsilon}{\varepsilon - 1} c_{it} \right)^{1-\varepsilon} \middle| t, T_i, X_{it} \right] = \exp \left(\gamma_r T_i + \tilde{\lambda}_{rt} + \pi_r X_{it} + \eta_r \right) \quad (\text{i})$$

$$E \left[\frac{\varepsilon}{\varepsilon - 1} c_{it} \tau_B \middle| t, T_i, X_{it} \right] = \exp \left(\gamma_p T_i + \tilde{\lambda}_{pt} + \pi_p X_{it} + \eta_p \right) \quad (\text{ii})$$

$$E \left[\frac{\sum_{j \in \mathcal{D}_i, j \neq A} (1 + \delta_{ij}) \beta_{ijt^*} \tau_j^{1-\varepsilon}}{\sum_{j \in \mathcal{D}_i, j \neq A} \beta_{ijt^*} \tau_j^{1-\varepsilon}} \middle| T_i, X_{it} \right] = \exp(\tilde{\delta}_{-A}) \quad (\text{iii})$$

Part (i) and (ii) of assumption 2 essentially state that the process of the demand shifter in B , β_{iBt} , and of the marginal cost c_{it} , are determined by the sum of a time-invariant group effect and year effects which are common across groups T and NT . It implies in particular that the shock should not impact demand shifters and marginal costs in different proportions across groups. Part (iii) of assumption 2 establishes that the firm-specific demand shock at date t^* , average over all markets except for A , should be mean-independent of the group and the controls. This independence assumption is necessary to guarantee that differences in price and sales dynamics across groups at the time of the shock (which we use to estimate the reallocation effect) are not caused by cross- group differences in demand shocks.

Both previous assumptions deliver following proposition

Proposition 3. (STRUCTURAL TREATMENT EFFECT): Under assumptions 1 and 2

$$\begin{cases} \exp(\alpha_r) = E \left[\left(1 + \rho_{iAt^*} \frac{\delta_{iA} - \delta_{i-A}}{1 + \delta_{i-A}} \right)^{\frac{\varepsilon-1}{\varepsilon}} \middle| i \in T, X_{it} \right] \\ \exp(\alpha_p) = E \left[\left(1 + \rho_{iAt^*} \frac{\delta_{iA} - \delta_{i-A}}{1 + \delta_{i-A}} \right)^{-\frac{1}{\varepsilon}} \middle| i \in T, X_{it} \right] \end{cases} \quad (2.10)$$

, with $\delta_{i-A} \equiv \frac{\sum_{j \in \mathcal{D}_i, j \neq A} \beta_{ijt^*} \delta_{ij}}{\sum_{j' \in \mathcal{D}_i, j' \neq A} \beta_{ij't^*}}$.

Proof. See appendix □

Proposition 3 stipulates that the treatment effects (2.9) are increasing in anticipation error differences $\delta_{iA} - \delta_{i-A}$ as well as in the share of market in the total sales of firm i , ρ_{iAt^*} . In practice, 3 predicts that if there is a negative demand shock in A relatively larger in absolute value than a shock in $-A$ (i.e. $\delta_{iA} - \delta_{i-A} > 0$, as firms over-estimate their demand in A more than in $-A$), then the treatment effect is positive on sales and negative on prices. The more so the larger the share of market A in firms total sales. If now shocks are the same in both countries, the treatment effect is zero. Likewise, in traditional trade models, the treatment effect is zero as there is no demand uncertainty so that prediction errors $\delta_{ij}, j = A, -A$ are both equal to zero for all firms. In other words, using this DD approach on a shock which is equal across markets does not allow us to test the reallocation effect because our model's prediction aligns with those of traditional models.

We now present the data we use to implement our DD approach.

4 Data

4.1 Why Champagne during the 2000-2001 US crisis?

In the empirical section below, we apply our DD approach to the Champagne wine industry during the 2001 US economic crisis. Formally, the date of the shock is $t^* = \{2000, 2001\}$ as it is usually what is considered as the period of the crisis.⁹ The market where the shock originates from, A , is the US. As to B , the market over which we test the presence of a potential reallocation, it will be alternatively the home market of firms in the sample, namely France, or the export market excluding the US.

On top of the anecdote described at the beginning this article, there are at least five reasons why Champagne wine is the appropriate industry to test our theory. First, the production of Champagne takes time. Second, Champagne exhibits high storage cost. Both these features fit the assumptions of our model. Third, US is an important destination market for Champagne which implies that producers are greatly affected with the shock in the US. Forth, Champagne is also widely exported outside the US which makes it possible to test for reallocation towards other destination markets. Fifth, due to legal constrained, Champagne can not be produced outside France. This rules out relocation, a competing mechanism to the one we want to test, as a response to the shock. We will now detail these points.

First, the production of Champagne takes time which makes demand uncertain at the time production decisions are made, consistently with the assumption of the model. This idea is supported by Declerck & Cloutier (2002):

“Because it takes about three years to sell Champagne wine after the vintage of grapes, industry operators have to cope with the uncertainty associated to long time delays in the formation of price expectations, in the adjustment of supply,

9. While the crash of the dot-com bubble occurred in 2000 (see DeLong & Magin 2006), the NBER considered the the real economy peaked in early 2001 Hall et al. (2001).

seasonal production and demand. [...] The consumption of non-basic goods such as Champagne wine is more sensitive to the fluctuations of economic cycles.”

Second, the storage of Champagne is costly. Unlike many other wines, most Champagne bottles do not improve with age. On the contrary, Champagne rapidly deteriorates if strict storage conditions regarding light, humidity and temperature are not met.¹⁰

Third, the US market is a large market for Champagne. Over the 1995-2007 period covered by our data, Champagne producers make close to 9% of their sales in the US and only the UK imports more Champagne than the US. As a consequence, the 2001 US economic crisis is likely to have a large impact on Champagne exporting firms which is necessary to detect a potential reallocation effect. Fourth, Champagne is a highly exported product, not only to the US. In our data and not counting exports to the US, Champagne exporting firms serve 201 foreign markets where they make 33% of their total sales. In the case of the US crisis that we investigate, this large share of foreign sales makes it possible to test for the existence of reallocation strategies not only from the US to the home market but also to other destination markets. Fifth, a sparkling wine must be produced in the Champagne region of France to be sold under the “Champagne” label. This rules out the possibility that firms spatially re-allocate their production (rather than their exports) as a response to a local shock. As a consequence, customs data is the correct type of data to detect the presence of spatial reallocation.

The next question regards our motivation for choosing the 2001 US crisis as a shock. This choice mainly rests upon the unexpected nature of the crisis. The events are usually seen considered at the origin of the 2001 US recession: the crash of the financial bubble on information technology (the “dot com” bubble) and the september 11th terrorist attack. Had the financial crisis been expected it would have been reflected into asset prices. Had the attack been expected, it would have been prevented. Another interest of this shock is that although it has diffused to other economies, it originates in the US. Hence, we may expect the shock to be relatively weaker in the rest of the world. As explained in section 3, we need a relative shock for our test to be valid.

4.2 Data Sources

Customs dataset Our empirical analysis relies on two databases. The first is the French customs database reporting the almost entire universe of exports by firms located in France. Only very small export flows are missing from the data due to reporting thresholds. More specifically, Firms which total exports to a non-EU country are below 1000 euros (or 1000 kilograms) do not need to complete an export declaration. For similar reasons, trade flows by firms that exported less than 100,000 euros to the whole EU in the previous year do not enter our data set.

A trade flow is a unique combination of a year, a firm identifier, a destination country and a product code. For each trade flow, the free-on-board value in euro and the quantity in kilogram are reported. From both these variables, we construct FOB prices (often referred to as “unit values”) as the ratio of the value of a flow to its quantity. As documented in Crozet

10. See <http://en.wikipedia.org/wiki/Storageofwine> for details regarding conservation conditions.

et al. (2012), quantity declaration are subject to errors which result into price outliers. The data contains two alternative measures of quantity which Crozet et al. (2012) combine to limit the number of observations with aberrant prices. We follow this procedure. The resulting prices had a first percentile of 5.45 Euro/kg, a 99th percentile of 86.94 Euro/kg as well as 284 missing values .

Table 2.1: Unit Values of non-US Exports

# observations	mean	percentile			min	max
		1	50	99		
2264	16.91	5.45	13.66	86.94	4.13	948.95

The product classification is the 8-digit product classification (hereafter “cn8”) which is obtained by adding a 2 digit suffix to the codes from the 6-digit harmonized system. In the cn8, Champagne sparkling wine is attributed is single code. This is fortunate as it guarantees that export flows utilized in the regression do not correspond to different types of wines which mitigates potential aggregation biases.

BRN dataset The second source is the balance sheet dataset BRN (for Benefice Reels Normaux). The BRN database is constructed from mandatory reports of French firms to the tax administration). This dataset reports information including firms main activity, total sales, export sales, number of employees, capital stock, value added, the industry, year, and balance-sheet variables. The dataset does not break down firms sales by product which may be an issue when it comes to measuring domestic sales of Champagne. We come back to this in the next paragraph as we present the data on quality.

Our data covers the period 1993-2008 but only are used years from 1995 to 2007 as we merge the BRN data to customs data. The BRN covers all firms subject to the normal tax regime, i.e. firms which annual sales are at least 763 K euros (230 K euros for services) or which chose the normal tax regime. This implies that small firms may be in the sample. The BRN contains between 650,000 and 750,000 firms per year over the period, which is around 60% of the total number of French firms. Out of these, we only keep firms which are present in the customs data and the BRN before 2000.

Quality Data Our last source is the expert assessed measure of Champagne established by Juhlin (2008). This measure is a rating ranging from 1 to 5 stars which aims at assessing the historical quality of main producers. In order to analyze the impact of quality of firms’ export performance, Crozet et al. (2012) have created the correspondence from producer names used by Juhlin (2008) to SIREN codes which identify firms in both the BRN and the French customs dataset. We use this correspondence, which was kindly made available to us by the authors, to incorporate quality ratings in our dataset. We are able to match 196 rated firms to either the BRN and customs dataset.

We limit our sample to the set of firms for which quality is observed. This limitation serves two purposes. First, it leaves us with a sample a of “the most important active producers and their wines.” for which selling Champagne is arguably the largest source of revenues. This

is an important point as we do not have information on firms domestic sales by product. As documented in table 2.2, a large majority of domestic sales in our final sample (70%) are made by Champagne wine-makers. This is comforting as to the fact that domestic sales mostly measure Champagne sales.

Table 2.2: Champagne producers' domestic sales by main activity

activity	code	Share home sales
Champagne Wine-maker	159F/1102A	71%
Wholesalers	511N/513J/1101Z 4617B/4634Z	21.6%
Others	various	7.4%

Another way to build confidence as to the product structure of home sales is to look at the product composition of exports. As reported in table 2.3, the median share of Champagne in the exports of the 196 firms in the sample is 99.7% while the fifth percentile is 71.4%.

Table 2.3: Share of Champagne in total exports

# Observations	Percentile					
	1	5	10	25	50	75
196	0	71.4%	95%	98.8%	99.7%	100%

Assuming that the share of Champagne in exports is a strong predictor of the share in domestic sales, table 2.3 suggests that domestic revenues in the sample should essentially result from Champagne sales.¹¹

The second reason for incorporating quality in our analysis is to control for unobservables which might endanger our identification. As Crozet et al. (2012) show, quality is a strong predictor of Champagne producers' export performance, even more so than physical productivity which is usually considered as the main source of firm heterogeneity in trade literature.

Our final sample has 196 firms over 13 years (1995-2007). Out of the $196 \times 13 = 4936$ firm-year combinations, 2467 (97%) have non-missing home sales values. As to observations on total non-US exports, 2292 (90%) are non-missing values. Total exports by firms in the sample make (88.3%) of the total Champagne exports over the period.

Treated and control groups A firm is defined as treated if it has positive Champagne exports to the US over the pre-shock period 1995-1999. It falls in the control group otherwise. table 2.4 provides the distribution of exposure to the US crisis where exposure is the share of exports to the US in the total Champagne exports of a firm over 1995-1999. Around 56% of firms are treated.

Table 2.5 compares firms from the control and treatment group along a series of observable characteristics. Treated firms make significantly larger value added and employ significantly

11. Since it is still possible that domestic sales poorly measure Champagne sales in our data in spite of these descriptive statistics, as a robustness check we run our main statistical tests over a restrictive sample of firms which main activity is Champagne and which only exported good is Champagne.

Table 2.4: Distribution of Exposure to the 2001 US crisis

firms	#	mean exposure	percentile exposure			
			25	50	75	99
All	196	8.2%	0%	0.4%	9.6%	43.0%
Treated	110	14.7%	1.6%	6.4%	16.2 %	91.4%

Notes: Variable “exposure” is the share of exports to the US in the total sales of a firms over 1995-1999

Table 2.5: Comparison of treated and untreated firms

<i>variable</i>	<i>untreated</i>	<i>treated</i>	<i>difference</i>
employees	10.34	43.57	33.23**
value added (KEuros)	1.08e+03	5.93e+03	4.84e+03*
unit value (Euros)	17.25	17.17	-0.08
quality (Stars)	1.83	2.86	1.02***

more workers. Note, however, that to be valid, our DD approach only requires treated, untreated firms to follow the same dynamics. We show in next section that this is verified.

5 Stylized Facts

This section documents patterns of Champagne sales at home and abroad during the 2001 crisis. Proposition 3 tells us that the crisis (i) should be larger in the US and (ii) the US should be a large importer of Champagne for the treatment effect to be large, and hence detectable. We now show that the US crisis meets both these requirements.

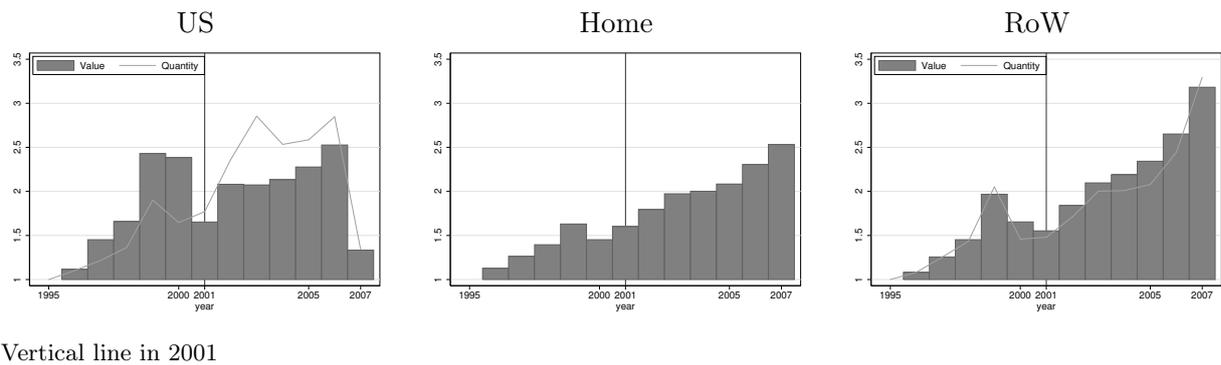
figure 2.1 reports the evolution of total sales in value and quantity at home, in the US and to the rest of the world (RoW). US imports of Champagne decline slightly in 2000, yet the biggest share of the 2000-2001 drop occurs in 2001 with a 40% drop in imports between 2000 and 2001. If the drop in 2000 is small (around 4%) it should be compared to the 40% growth in 1999 to appreciate the magnitude of the shock.

Sales at home and to the rest of the world exhibit similar patterns, although of a lesser magnitude: home sales are almost back to their 1999 level in 2001 while RoW sales drop by 15% between 1999 and 2001. Therefore, for the Champagne industry, the 2001 recession is global, although more marked in the US.

To get a better sense of differences in the magnitude of the shock between the US and other markets, figure 2.2 shows the evolution of the sales to the US respectively relative to home sales and to non-US exports. Both lines rise at the beginning of the period. The increases slows down between 1999 and 2000 and both variables fall back to their 1995 level in 2001, which clearly shows that the shock was larger in the US.

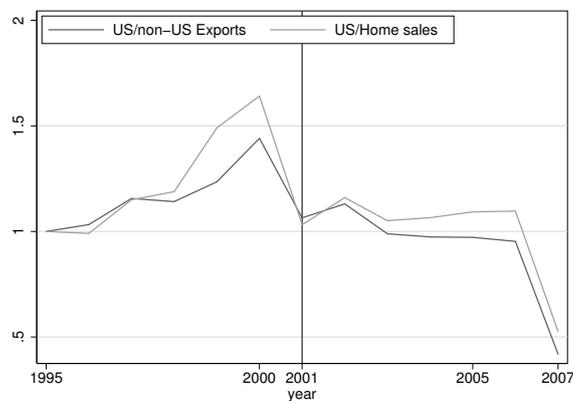
Let us now document the importance of the US as a market for Champagne. Figure 2.3 plots the value of sales to the main Champagne importing countries for years 1999 and 2001, i.e. before and after the shock. The lion’s share of world imports of Champagne is done by two countries, namely Great Britain and the US. This is interesting since, to observe the reallocation from a country to another, we need the country hit by a shock to be relatively

Figure 2.1: Champagne sales and the crisis - Home sales versus US and non-US exports



Vertical line in 2001

Figure 2.2: The 2001 Crisis was stronger on the US market

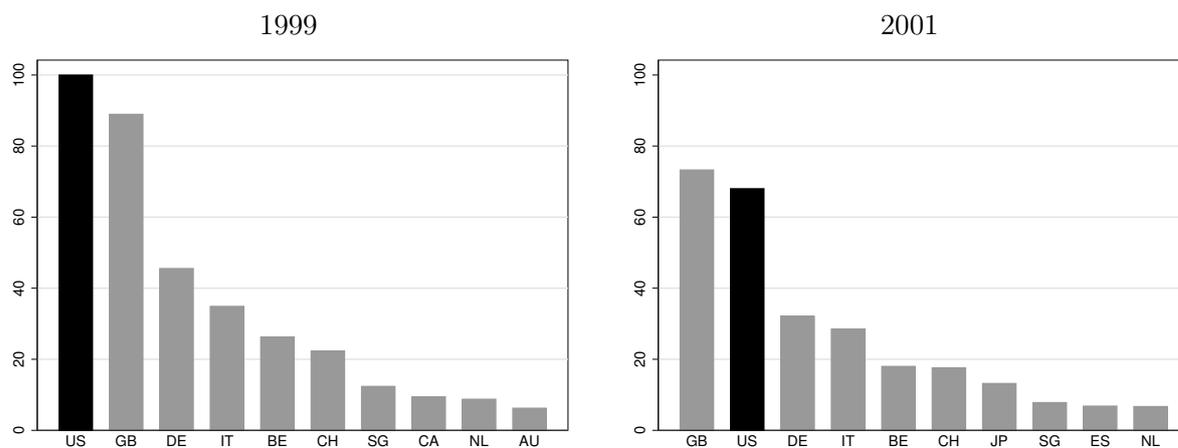


large. The second striking fact from the comparison of the two graphs is that from 1999 to 2001, the US left its place of first importer at the expense of Great Britain.

We conclude this section by graphically looking at the differential impact of the crisis on the treated group and the control group. Figure 2.4 displays the total sales and the average unit value of non-US export, as well as the domestic sales for both the treated group and the control group. In the left-hand side panel, it appears that the pre-trend of total exports are very similar across. According to the model, this pre-shock common trend is interpretable as the occurrence of demand shocks on non-US exports as well cost shocks equal across groups of average. This is very supportive of the validity of our DD approach. At the time of the shock now, both lines drop, yet the drop is smaller for the firms from the treated group. This is consistent with our theory: firms which expected to serve the US in 2000 reallocate their sales to the RoW which partly mitigates the sales drop in the RoW. A nice feature of this graph is that lines converge back to each other after the shock: once the surprise of the shock is over, producer adjust their total output downward and do not have a surplus of sales to reallocate outside the US anymore.

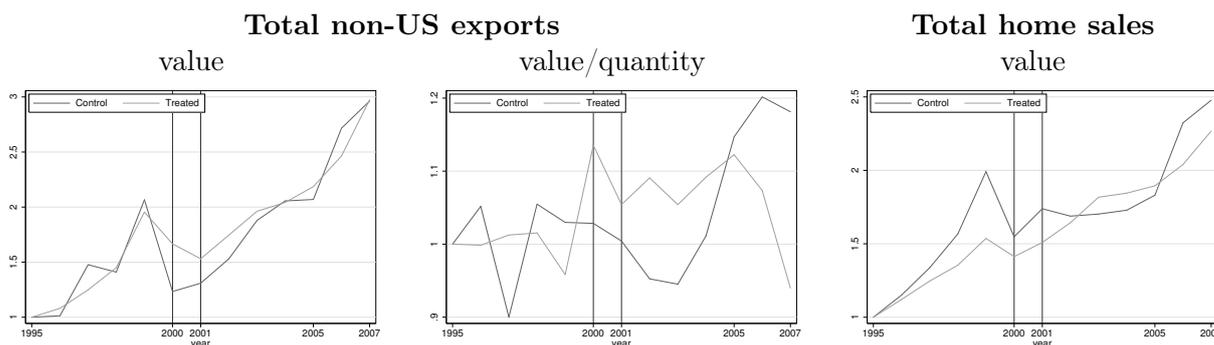
The message from the middle and right hand side panels in figure 2.4 is less clear. Both prices and domestic sales seem to violate the common trend assumption. Domestic sales grow

Figure 2.3: The 2001 crisis and the US among top Champagne importers



Notes: Total value of Champagne exports to top Champagne importing countries in 1999 and 2001. Normalized to 100 for the US in 1999.

Figure 2.4: Exposure to the Crisis and sales at home and abroad



Vertical lines in 2000 and 2001

faster for the control group before the shock while prices simply do not exhibit a clear trend. This calls for a cautious interpretation of the results from our DD regressions that we present in next section.

6 Results

6.1 Baseline Results

In this section, we estimate model (2.9) with firm fixed effects in place of controls X_{it} . Another detail of our specification is that we include a full set of year dummies after 2000 interacted with the treated dummy variable T_i . In that way we allow the treatment effect to evolve non-parametrically over time after year 2000:

$$\begin{cases} r_{iBt} &= \exp\left(\sum_{t'=2000}^{2007} \alpha_{rt} T_i \times S_t + \gamma_r T_i + \lambda_{rt} + FE_i + \eta_r\right) u_{rit} \\ p_{iB\tau} &= \exp\left(\sum_{t'=2000}^{2007} \alpha_{pt} T_i \times S_t + \gamma_p T_i + \lambda_{pt} + FE_i + \eta_p\right) u_{pit} \end{cases}, B = \{\text{home, non-US exports}\}$$
(2.11)

With S_t a vector of dummy variable for years 2000 to 2007 and α_{vt} , $v = r, p$ the average effect of the crisis on treated firms at date t . As pointed out by Silva & Tenreyro (2006), an OLS estimation of a log-linearized version of model (2.11) would deliver inconsistent coefficient estimates unless strong distributional assumptions, such as the homoskedasticity of the residual, are made. We rely on the popular solution, proposed by Silva & Tenreyro (2006), which consists in estimating constant-elasticity models via Poisson pseudo-maximum likelihood. Results are reported in table 2.1.

Table 2.1: A reallocation effect of the US Crisis?

	(1)	(2)	(3)
Dependent variable:	<i>non-US Export</i> _{it}	<i>non-US unit value</i> _{it}	<i>home sales</i> _{it}
Method:	Poisson	Poisson	Poisson
{ year = 2000 } × Treated _{<i>i</i>}	0.334** (0.027)	0.106* (0.051)	0.027 (0.799)
{ year = 2001 } × Treated _{<i>i</i>}	0.190 (0.231)	0.094* (0.082)	-0.023 (0.767)
{ year = 2002 } × Treated _{<i>i</i>}	0.166 (0.339)	0.118* (0.052)	0.092 (0.411)
{ year = 2003 } × Treated _{<i>i</i>}	0.077 (0.743)	0.062 (0.369)	0.184 (0.194)
{ year = 2004 } × Treated _{<i>i</i>}	0.027 (0.943)	0.091 (0.130)	0.184 (0.344)
{ year = 2005 } × Treated _{<i>i</i>}	0.089 (0.798)	0.114* (0.056)	0.153 (0.458)
{ year = 2006 } × Treated _{<i>i</i>}	-0.063 (0.862)	0.086 (0.142)	-0.012 (0.955)
{ year = 2007 } × Treated _{<i>i</i>}	0.038 (0.910)	0.067 (0.278)	0.030 (0.912)
Firm FE, Year FE	Yes	Yes	Yes
Observations	2548	2264	2548

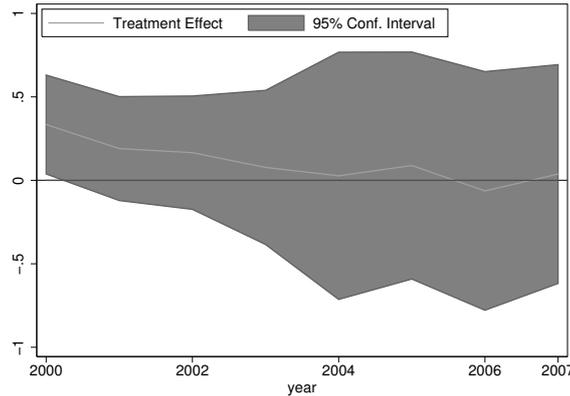
Heteroskedasticity robust p -values in parentheses

* $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Column (1) of table 2.1 reports the results for non-US export. As a way to visualize the evolution of the treatment effect over time, captured by variable {**year = t** } × **Treated**_{*i*}, $t =$

2000, ..., 2007, we plot the point estimates as well as their 95% confidence interval in figure 2.1. The point estimates are decreasing over time. As a result, the estimated treatment effect is statistically significant at 5% in 2000 and is not significant in following years. This pattern is consistent with treated firms reallocating their sales surplus from the US at the beginning of the crisis and then adjusting their output over time to avoid surplus.

Figure 2.1: Treatment Effect for non-US exports over time



6.2 Discussion and Robustness

Is the estimated treatment effect too large? The point estimate in 2000 is $\widehat{\exp(\alpha_r)}_{DD} = \exp(0.34) = 1.40$. The subscript “DD” signals that the estimate was obtained from the DD regression and aims at differentiating from the alternative estimate we will obtain via an alternative strategy detailed below. The 1.40 coefficient means that treated firms have decreased their sales 40% less between 1999 and 2000 than their untreated counterpart. How big is that coefficient? We address that question by first calibrating the structural expression of parameter $\exp(\alpha_r)$ given in proposition 3. Then we compare the calibrated value obtained, $\widehat{\exp(\alpha_r)}_{\text{calib.}}$, to the estimated one, $\widehat{\exp(\alpha_r)}_{DD}$. To achieve this, we first re-write $\exp(\alpha_r)$ as follows:

$$\exp(\alpha_r) = E \left[\left(\frac{\rho_{i,non-US,2000}}{\hat{\rho}_{i,non-US,2000}} \right)^{\frac{\varepsilon-1}{\varepsilon}} \mid i \in T, X_{it} \right] \quad (2.12)$$

, with $\hat{\rho}_{i,non-US,t} \equiv \frac{\sum_{j \in \mathcal{D}_i, j \neq US} (1 + \delta_{ij}) \beta_{ijt}}{\sum_{j' \in \mathcal{D}_i} (1 + \delta_{ij'}) \beta_{ij't}}$ the anticipated share of sales to all market (including home sales) but the US at date t .¹² The ratio $\frac{\rho_{i,non-US,2000}}{\hat{\rho}_{i,non-US,2000}}$ measures by how much firms had under-estimated the share of their sales outside the US in 2000. The larger this ratio, the larger the amount of output reallocated from the US to the RoW, the larger the treatment effect. The nice feature of expression (2.12) is that it involves $\rho_{i,non-US,2000}$ which is easy to compute in the data as the ratio in 2000 of home sales plus non-US export of Champagne to home sales plus total exports of Champagne.

To calibrate expression (2.12), we also need to put a value on anticipations $\hat{\rho}_{i,non-US,2000}$. To this end, we assume that firms make following extrapolative anticipation:

$$\hat{\rho}_{i,non-US,2000} = \rho_{i,non-US,1999} \times \left(\frac{\rho_{i,non-US,1999}}{\rho_{i,non-US,1998}} \right) \quad (2.13)$$

. Namely, at the beginning of 2000, firms expect $\hat{\rho}_{i,non-US}$ to grow at the same rate between 1999 and 2000 than between 1998 and 1999. From (2.13), we can easily compute $\hat{\rho}_{i,non-US,2000}$ for each firm in the data. In table 2.2, we report the distribution of $\frac{\hat{\rho}_{i,non-US,2000}}{\rho_{i,non-US,2000}}$ that we obtain for the treated group.

Table 2.2: Distribution of $\frac{\hat{\rho}_{i,non-US,2000}}{\rho_{i,non-US,2000}}$

firms	#	mean	min	max	percentile		
					1	50	99
treated	110	1.68	1.81e-5	51.31	9.62e-4	0.98	33.39
treated, 1st & 99th percentiles excluded	106	0.94	0.22	1.69	0.24	0.94	1.36

Notes: This table reports the distribution of the ratio of the anticipated to realized share of non-US sales in 2000, $\frac{\hat{\rho}_{i,non-US,2000}}{\rho_{i,non-US,2000}}$, for firms from the treated group. The realized share is computed as the ratio of home sales+non-US exports to home sales+total exports. The anticipated share is computed following formula (2.13) in the main text.

A simple way to check whether extrapolative assumptions are reasonable is to verify whether they imply that firms under-estimate their share of non-US sales in 2000. Table 2.2 provides the distribution of the anticipated to realized ratio of non-US sales in 2000 computed for the sample of non treated firms under the assumption of extrapolative anticipations. We can verify that once we exclude outliers, as reported in the last row of table 2.2, firms indeed

12. Let us re-write the argument of the expectation in the expression of $\exp(\alpha_r)$ in proposition 3:

$$\begin{aligned} \left(1 + \rho_{iAt^*} \frac{\delta_{iA} - \delta_{i-A}}{1 + \delta_{i-A}} \right) &= \left(\frac{1 + \rho_{iAt^*} \delta_{iA} + \rho_{i-A} \delta_{i-A}}{1 + \delta_{i-A}} \right) \\ &= \left(\frac{1 + \delta_i}{1 + \delta_{i-A}} \right) \\ &= \frac{\sum_{j \in \mathcal{D}_i} (1 + \delta_{ij}) \beta_{ijt^*} / \sum_{j' \in \mathcal{D}_i} \beta_{ij't^*}}{\sum_{j \in \mathcal{D}_i, j \neq A} (1 + \delta_{ij}) \beta_{ijt^*} / \sum_{j' \in \mathcal{D}_i, j' \neq A} \beta_{ij't^*}} \\ &= \frac{\rho_{i-At^*}}{\hat{\rho}_{i-At^*}} \end{aligned}$$

with $\delta_i \equiv \frac{\sum_{j \in i} \beta_{ijt^*} (1 + \delta_i)}{\sum_{j \in i} \beta_{ijt^*}}$. Applying expectancy to both side of last equation, substituting US for A and 2000 for t^* gives equation (2.12).

under estimate on average. Specifically, the mean ratio of 0.94 implies that firms underestimated $\rho_{i,non-US,2000}$ by 6 percents on average .

With firm level estimates of $\frac{\hat{\rho}_{i,non-US,2000}}{\rho_{i,non-US,2000}}$ in hand, we obtain an alternative estimator of $\exp(\alpha_r)$,

$$\widehat{\exp(\alpha_r)}_{\text{calib.}} = \frac{1}{\#T} \sum_{i \in T} \left\{ \left(\frac{\rho_{i,non-US,1999}}{\rho_{i,non-US,2000}} \right) \left(\frac{\rho_{i,non-US,1999}}{\rho_{i,non-US,1998}} \right) \right\}^{-\frac{\varepsilon-1}{\varepsilon}} \quad (2.14)$$

$\widehat{\exp(\alpha_r)}_{\text{calib.}}$ is the empirical counterpart to (2.12) under the assumption of extrapolative anticipations. As a last step, computing $\widehat{\exp(\alpha_r)}_{\text{calib.}}$ requires to calibrate the price elasticity of demand ε . Here we follow Crozet et al. (2012) and use $\varepsilon = 5$ and $\varepsilon = 10$. In third column of table 2.3 we report the value of the treatment effect $\widehat{\exp(\alpha_r)}_{\text{calib.}}$ for both values of ε and compare it to the treatment effect estimated from DD, $\widehat{\exp(\alpha_r)}_{\text{DD}}$. It appears that $\widehat{\exp(\alpha_r)}_{\text{DD}}$ is significantly larger than $\widehat{\exp(\alpha_r)}_{\text{calib.}}$. In other words, under extrapolative anticipation, our model predicts a reallocation effect significantly smaller than the one implied by $\widehat{\exp(\alpha_r)}_{\text{DD}}$.

Table 2.3: Treatment effect and anticipation error: comparing the calibration and the diff-in-diff

	treatment effect			Mean $\frac{\hat{\rho}_{i,non-US,2000}}{\rho_{i,non-US,2000}}$
	$\widehat{\exp(\alpha_r)}_{\text{DD}}$	$\widehat{\exp(\alpha_r)}_{\text{calib.}}$	data	implied by $\widehat{\exp(\alpha_r)}_{\text{DD}}$
$\varepsilon = 5$		1.11		0.78
$\varepsilon = 10$	1.40	1.13	0.94	0.81

Notes: This table reports the value of the treatment effect as well as the average expected to realized share of non-US sales in 2000. For the treatment effect, we report the value estimated from the DD regression as well as the values obtained from calibrating expression with different values of ε . For the average expected to realized ratio, we report the value computed from the data and the value which is necessary to replicate $\widehat{\exp(\alpha_r)}_{\text{DD}}$ which is obtained by solving (2.15).

A way to measure how excessively large $\widehat{\exp(\alpha_r)}_{\text{DD}}$ is, is in term of anticipation error. We can compute how larger the anticipation error on $\rho_{i,non-US,2000}$ should be for our model to generate $\exp(\alpha_r) = \widehat{\exp(\alpha_r)}_{\text{DD}}$. Formally, we look for the constant λ by which one should shift the anticipated/realized ratio $\frac{\hat{\rho}_{i,non-US,2000}}{\rho_{i,non-US,2000}}$ to equalize $\widehat{\exp(\alpha_r)}_{\text{calib.}}$ to $\widehat{\exp(\alpha_r)}_{\text{DD}}$:

$$\lambda : \frac{1}{\#T} \sum_{i \in T} \left\{ \left(\frac{\rho_{i,non-US,1999}}{\rho_{i,non-US,2000}} \right) \left(\frac{\rho_{i,non-US,1999}}{\rho_{i,non-US,1998}} \right) + \lambda \right\}^{-\frac{\varepsilon-1}{\varepsilon}} = 1.40 \quad (2.15)$$

We compute λ as a solution to (2.15) for $\varepsilon = 5$ and $\varepsilon = 10$. The average value of the anticipated to realized non-US share, $\frac{\hat{\rho}_{i,non-US,2000}}{\rho_{i,non-US,2000}} + \lambda$, is reported in the last column of table 2.3. While extrapolative assumption imply that firms under anticipate $\rho_{i,non-US,2000}$ by 6 percents, firms should under anticipate by 22 percents (respectively 19 percents) for the model to rationalize $\widehat{\exp(\alpha_r)}_{\text{DD}}$ if $\varepsilon = 5$ (respectively $\varepsilon = 10$). In other words, our DD estimate of the treatment effect implies firms' average anticipation error to be $22/6=3.6$ (respectively

19/6=3.2) times larger than under extrapolative anticipations.

This is clearly a large number which suggests that other phenomena are at play to generate a large estimated coefficient. In next paragraph we investigate two potential sources explaining this excessive size.

Why is the estimated treatment effect too large? One possibility is that the 2001 financial crisis has tightened firms' financial constraints and that this has more impacted untreated firms which are smaller and financially more fragile. However, this explanation is not consistent with the common trend observed before the shock and at the end of the sample between groups. If untreated firms' sales were more sensitive to global financial conditions, then they should have grown faster before the crisis as the IT bubble was inflating.

Another competing story is that treated and untreated firms are located at different levels of the quality ladder and as such sell to different customers and face different demand shocks. The positive treatment effect could therefore be explained by treated firms facing a smaller demand shock in 2000.

We test both previous competing stories by including controls for the potentially specific financial and demand shocks faced by treated firms in 2000. Formally, we estimate following specification:

$$r_{i,nonUS,t} = \exp \left(\sum_{t'=2000}^{2007} \alpha_{rt} T_i \times S_t + \gamma_r T_i + \lambda_{rt} + \chi_{rt} \times X_i + FE_i + \eta_r \right) u_{rit} \quad (2.16)$$

Where X_i are variables which are suspected to affect the magnitude of the financial and demand shock faced by firms at each date and in particular at the time of the shock. χ_{rt} is a full set of year dummies. The inclusion of the interaction between X_i and the time dummies implies that the treatment effect α_{rt} identifies the difference in sales dynamics between treated and untreated, *conditional on* X_i . Thus, the value of the treatment effect is not driven by systematic differences in X_i between groups.

To control for the sensitivity to the financial shock, we construct the working capital requirement ratio (WCR) at the firm level. This variable is the difference between current assets and short-term debts normalized by the stable resources of the firms. Intuitively, it measures the difficulty with which a firm can access liquidity to finance its working capital. As for differences in demand shocks, we control them by including the quality rating by Juhlin (2008) as an X variable.

Remark that if we are interested in the value of the treatment coefficient, the value of χ_{r2000} is also of interest as it is revealing of the direction of the bias on the treatment effect resulting from omitting variable X . More specifically, if the positive treatment effect found in table 2.1 was driven by the fact that firms in the case where $X = WCR$, if one should expect the coefficient associated to χ_{r2000}

Results from the estimation of model (2.16) are reported in table 2.4. Although only coefficients of interest are reported, they were obtained by including full sets of year dummies as well as year dummies interact with control variables, WCR and quality. The first finding

Table 2.4: Treatment effect: controlling for group differences in financial and demand shocks

	(1)	(2)
	non-US Export _{it}	non-US Export _{it}
Method:	Poisson	Poisson
$\{t = 2000\} \times \text{Treated}_i$	0.351** (0.029)	0.384 (0.130)
$\{t = 2000\} \times \text{WCR}_i$	0.468 (0.112)	
$\{t = 2000\} \times \text{Quality}_i$		-0.067 (0.603)
Firm FE	Yes	Yes
Year FE	Yes	Yes
Observations	2535	2548

Notes: Quality is the number of stars from Juhlin (2008). WCR is the working capital requirement ratio which construction is described in the text. Larger WCR proxies more stringent liquidity constraints. Both specifications, including interactions of Treated_i with year dummies for years 2000-2007. Specification (1) includes interactions of WCR_i with a full set of year dummies. Specification (2) includes interactions of Quality_i with a full set of year dummies. Heteroskedasticity robust p -values in parentheses. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

from the analysis of table 2.4 is that the treatment effect is barely affected with the inclusion of the time-control interactions. This suggests that the treatment effect that we find is not biased by group differences in financial and demand shocks in 2000. This interpretation is confirmed by the fact that the interaction between the control and year 2000 is not significant. In other words, quality differences and differences in WCR do not generate different shocks during the crisis which implies that omitting these variables does not bias the estimation of the treatment effect. Overall, this exercise re-inforces the reallocation interpretation of the treatment effect.

7 Conclusion

In this paper, we provide for the first time causal evidence that firms substitute sales between foreign markets in order to mitigate the impact of unexpected relative demand shocks. We see two implications of this result that future research should investigate. First, our results suggest that demand volatility is costly (decreases the expected profit of a firm) because it generates gaps between production and demand. This suggests a new gain for firms to serve many markets, namely, it allows them to hedge against the volatility of demand. By symmetry, this new gain might change the exporting costs required to replicate the observed exporting activity of firms.

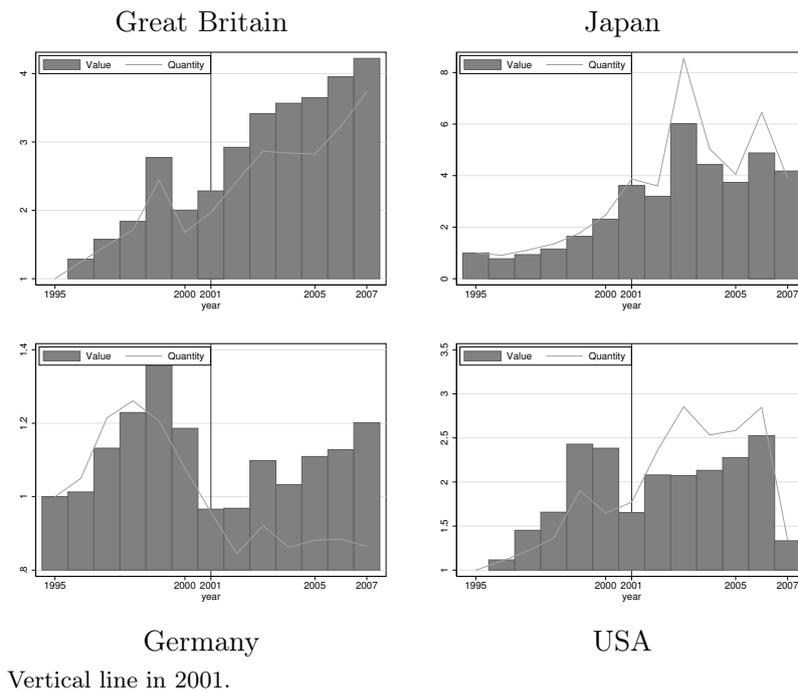
Another implication of our finding regards the relationship between international trade and international business cycle co-movement. Since in our model exporting is a way for a firm to smooth its revenue in the presence of demand volatility, it might be that in a general equilibrium model embedding our model, trade liberalization actually asynchronizes (or weakly synchronizes) international business cycles.

Appendix

A Additional Graphs

Figure A.1 provides an information similar to figure 2.1 but for four of the largest destination markets for Champagne, namely Great Britain, Japan and Germany. Information for the US is reported to ease the comparison. The main take away point from A.1 is that sales to main foreign markets all exhibit, to different extents, a dip in 2000-2001.

Figure A.1: Value and Volume of Exports to Great Britain, Japan and the US



B Deriving the Structural Expression of the Treatment Effect

In this section, we prove proposition 3. Under assumption 1, and outside the shock period, the expression of Q in (2.8) simplifies to:

$$Q_{it} = \left(\frac{\varepsilon}{\varepsilon - 1} c_{it} \right)^{-\varepsilon} \left(\sum_{j \in \mathcal{D}_i} \tau_j \beta_{ijt} \right), \quad \forall t \neq t^*$$

and so sales and prices to market B verify

$$\begin{cases} r_{iBt} &= \beta_{iBt} \tau_B^{1-\varepsilon} \left(\frac{\varepsilon}{\varepsilon-1} c_{it} \right)^{1-\varepsilon}, \quad \forall t \neq t^* \\ p_{iBt} &= \frac{\varepsilon}{\varepsilon-1} c_{it} \tau_B \end{cases}$$

During the shock period now, firms make a δ_{ij} percent anticipation error on their demand shifter in market j so their production is

$$Q_{it^*} = \left(\frac{\varepsilon}{\varepsilon - 1} c_{it^*} \right)^{-\varepsilon} \left(\sum_{j \in \mathcal{D}_i} \tau_j \beta_{ijt^*} (1 + \delta_{ij}) \right)$$

and sales and prices verify

$$\begin{cases} r_{iBt^*} &= \beta_{iBt^*} \tau_B^{1-\varepsilon} \left(\frac{\sum_{j \in \mathcal{D}_i} \tau_j \beta_{ijt^*} (1 + \delta_{ij})}{\sum_{j \in \mathcal{D}_i} \tau_j \beta_{ijt^*}} \right)^{\frac{\varepsilon-1}{\varepsilon}} \left(\frac{\varepsilon}{\varepsilon-1} c_{it^*} \right)^{1-\varepsilon} \\ p_{iBt^*} &= \left(\frac{\sum_{j \in \mathcal{D}_i} \tau_j \beta_{ijt^*}}{\sum_{j \in \mathcal{D}_i} \tau_j \beta_{ijt^*} (1 + \delta_{ij})} \right)^{\frac{1}{\varepsilon}} \left(\frac{\varepsilon}{\varepsilon-1} c_{it^*} \right) \tau_B \end{cases}$$

We can re-write last expression

$$\begin{cases} r_{iBt^*} &= \beta_{iBt^*} \tau_B^{1-\varepsilon} \left(\sum_{j \in \mathcal{D}_i, j \neq A} \rho_{ijt^*} (1 + \delta_{ij}) + \rho_{iAt^*} (1 + \delta_{iA}) \right)^{\frac{\varepsilon-1}{\varepsilon}} \left(\frac{\varepsilon}{\varepsilon-1} c_{it^*} \right)^{1-\varepsilon} \\ p_{iBt^*} &= \left(\sum_{j \in \mathcal{D}_i, j \neq A} \rho_{ijt^*} (1 + \delta_{ij}) + \rho_{iAt^*} (1 + \delta_{iA}) \right)^{\frac{1}{\varepsilon}} \left(\frac{\varepsilon}{\varepsilon-1} c_{it^*} \right) \tau_B \end{cases}$$

, which can further be re-written as

$$\begin{cases} r_{iBt^*} &= \beta_{iBt^*} \tau_B^{1-\varepsilon} (\rho_{i-A} (1 + \delta_{i-A}) + \rho_{iAt^*} (1 + \delta_{iA}))^{\frac{\varepsilon-1}{\varepsilon}} \left(\frac{\varepsilon}{\varepsilon-1} c_{it^*} \right)^{1-\varepsilon} \\ p_{iBt^*} &= (\rho_{i-A} (1 + \delta_{i-A}) + \rho_{iAt^*} (1 + \delta_{iA}))^{\frac{1}{\varepsilon}} \left(\frac{\varepsilon}{\varepsilon-1} c_{it^*} \right) \tau_B \\ \delta_{i-A} &\equiv \frac{\sum_{j \in \mathcal{D}_i, j \neq A} \beta_{ijt^*} \tau_j^{1-\varepsilon} \delta_{ij}}{\sum_{j' \in \mathcal{D}_i, j' \neq A} \tau_{j'}^{1-\varepsilon} \beta_{ij't^*}} \\ \rho_{i-A} &\equiv \frac{\sum_{j \in \mathcal{D}_i, j \neq A} \beta_{ij't^*} \tau_j^{1-\varepsilon}}{\sum_{j' \in \mathcal{D}_i} \beta_{ij't^*} \tau_j^{1-\varepsilon}} \end{cases}$$

, with δ_{i-A} the anticipation error average over all markets but A and ρ_{i-A} the share of total sales to all markets but A .

Re-arranging this last expression gives

$$\begin{cases} r_{iBt^*} &= \beta_{iBt^*} \tau_B^{1-\varepsilon} \left\{ (1 + \delta_{i-A}) \left(1 + \rho_{iAt^*} \frac{\delta_{iA} - \delta_{i-A}}{1 + \delta_{i-A}} \right) \right\}^{\frac{\varepsilon-1}{\varepsilon}} \left(\frac{\varepsilon}{\varepsilon-1} c_{it^*} \right)^{1-\varepsilon} \\ p_{iBt^*} &= \left\{ (1 + \delta_{i-A}) \left(1 + \rho_{iAt^*} \frac{\delta_{iA} - \delta_{i-A}}{1 + \delta_{i-A}} \right) \right\}^{\frac{1}{\varepsilon}} \left(\frac{\varepsilon}{\varepsilon-1} c_{it^*} \right) \tau_B \end{cases}$$

The expression of sales (respectively prices) can be re-written in a single expression for both the period of the shock and outside the shock

$$\begin{cases} r_{iBt^*} &= \beta_{iBt^*} \tau_B^{1-\varepsilon} \exp \left\{ \frac{\varepsilon-1}{\varepsilon} \ln(1 + \delta_{i-A}) S_t + \frac{\varepsilon-1}{\varepsilon} \ln \left(1 + \rho_{iAt^*} \frac{\delta_{iA} - \delta_{i-A}}{1 + \delta_{i-A}} \right) S_t \times T_i \right\} \left(\frac{\varepsilon}{\varepsilon-1} c_{it^*} \right)^{1-\varepsilon} \\ p_{iBt^*} &= \exp \left\{ \frac{1}{\varepsilon} \ln(1 + \delta_{i-A}) S_t + \frac{1}{\varepsilon} \ln \left(1 + \rho_{iAt^*} \frac{\delta_{iA} - \delta_{i-A}}{1 + \delta_{i-A}} \right) S_t \times T_i \right\} \left(\frac{\varepsilon}{\varepsilon-1} c_{it^*} \right) \tau_B \end{cases}$$

Making use of assumption 2 we get

$$\begin{cases} r_{iBt} &= \exp(\gamma_r T_i + \lambda_{rt} + \pi_r X_{it} + \eta_r) \exp \left\{ \frac{\varepsilon-1}{\varepsilon} \ln \left(1 + \rho_{iAt^*} \frac{\delta_{iA} - \delta_{i-A}}{1 + \delta_{i-A}} \right) S_t \times T_i \right\} \times \tilde{u}_{rit} \\ p_{iBt} &= \exp(\gamma_p T_i + \lambda_{pt} + \pi_p X_{it} + \eta_p) \exp \left\{ \frac{1}{\varepsilon} \ln \left(1 + \rho_{iAt^*} \frac{\delta_{iA} - \delta_{i-A}}{1 + \delta_{i-A}} \right) S_t \times T_i \right\} \times \tilde{u}_{pit} \\ \lambda_{rt} &= \tilde{\lambda}_{rt} + \frac{\varepsilon-1}{\varepsilon} \ln(1 + \delta_{i-A}) S_t \\ \lambda_{pt} &= \tilde{\lambda}_{pt} - \frac{1}{\varepsilon} \ln(1 + \delta_{i-A}) S_t \end{cases}$$

The last step consists in multiplying and dividing the model for sales and prices respectively by $E \left[\left(1 + \rho_{iAt^*} \frac{\delta_{iA} - \delta_{i-A}}{1 + \delta_{i-A}} \right)^{\frac{\varepsilon-1}{\varepsilon}} \mid i \in T, X_{it} \right]$

and $E \left[\left(1 + \rho_{iAt^*} \frac{\delta_{iA} - \delta_{i-A}}{1 + \delta_{i-A}} \right)^{\frac{-1}{\varepsilon}} \mid i \in T, X_{it} \right]$. This gives:

$$\begin{cases} r_{iBt} &= \exp(\gamma_r T_i + \lambda_{rt} + \pi_r X_{it} + \eta_r + \alpha_r S_t \times T_i) \times u_{rit} \\ p_{iBt} &= \exp(\gamma_p T_i + \lambda_{pt} + \pi_p X_{it} + \eta_p + \alpha_p S_t \times T_i) \times u_{pit} \end{cases}$$

where (u_{rit}, u_{pit}) are residuals with mean 1 and orthogonal to (t, T_i, X_{it}) and (α_r, α_s) are the treatment effects. The expression of these different variables is

$$\left\{ \begin{array}{l}
\exp(\alpha_r) = E \left[\left(1 + \rho_{iAt^*} \frac{\delta_{iA} - \delta_{i-A}}{1 + \delta_{i-A}} \right)^{\frac{\varepsilon-1}{\varepsilon}} \mid i \in T, X_{it} \right] \\
\exp(\alpha_p) = E \left[\left(1 + \rho_{iAt^*} \frac{\delta_{iA} - \delta_{i-A}}{1 + \delta_{i-A}} \right)^{\frac{-1}{\varepsilon}} \mid i \in T, X_{it} \right] \\
u_{rit} = \left(\frac{\left(1 + \rho_{iAt^*} \frac{\delta_{iA} - \delta_{i-A}}{1 + \delta_{i-A}} \right)^{\frac{\varepsilon-1}{\varepsilon}}}{E \left[\left(1 + \rho_{iAt^*} \frac{\delta_{iA} - \delta_{i-A}}{1 + \delta_{i-A}} \right)^{\frac{\varepsilon-1}{\varepsilon}} \mid i \in T, X_{it} \right]} \right)^{S_t \times T_i} \tilde{u}_{rit} \\
u_{pit} = \left(\frac{\left(1 + \rho_{iAt^*} \frac{\delta_{iA} - \delta_{i-A}}{1 + \delta_{i-A}} \right)^{\frac{-1}{\varepsilon}}}{E \left[\left(1 + \rho_{iAt^*} \frac{\delta_{iA} - \delta_{i-A}}{1 + \delta_{i-A}} \right)^{\frac{-1}{\varepsilon}} \mid i \in T, X_{it} \right]} \right)^{S_t \times T_i} \tilde{u}_{pit}
\end{array} \right.$$

□

Chapter 3

Size-dependent Regulation and Factor Income Distribution

Abstract

In many countries, labor regulation is more binding for large firms than for small firms. By inducing a reallocation of productive resources away from large firms, such policy designs are likely to generate a loss of aggregate output. Moreover, as small firms are less capital intensive than large firms, size-dependent regulation might also induce a reallocation of income from capital to labor. The objective of this paper is to quantify these macroeconomic implications in the case of the 50 employees threshold in France. In order to achieve this, I first estimate the distortion induced by the regulation at the micro-level using French firm-level data. Then I use my estimates to calibrate policy experiments in a general equilibrium framework. Estimation is based on a structural model in which firms produce by combining capital and labor through a CES technology. Capital productivity and labor productivity both differ across firms. At the threshold, the model predicts (i) a break in the distribution of firm sizes, (ii) a break in capital intensity and (iii) the presence of an excessive mass of firms. These predictions are empirically verified and identify my parameters. I find that regulation raises the unit cost of labor by 7.5% and total costs of firms at the threshold by 3%. Firms above 50 raise their capital-labor ratio by 5% as a response to the regulation, which corresponds to a capital-labor elasticity of substitution of 0.6. Full deregulation raises aggregate welfare by 0.2% and decreases labor share of income by 0.14%.

1 Introduction

In many countries, labor regulation is more stringent for large firms than for small firms. An illustration is the obama care which constrains employers with more than 50 employees to purchase health insurance for their workers. Such regulation designs as suspected of impacting aggregate outcomes because (i) they induce a reallocation of market shares toward small firms and (ii) small firms are different from large firms. Most of the literature so far has explained the impact of size dependent regulations through the lens of productivity differences. By contrast, I focus on firm-level heterogeneity in factor intensity.

In particular, I consider the implications for the aggregate productivity and factor income distribution. Micro differences in factor intensity matter for two reasons. Firstly, the extend of misallocation resulting from regulating the use a factor depends on the share of this factor in the total costs of *regulated firms*. Ignoring factor heterogeneity therefore leads to mis-estimate the extend of misallocation by inferring it from the aggregate (wrong) factor share. Secondly, because large firms tend to be more capital intensive, size dependent regulations induce a reallocation of market shares towards labor intensive firms and thence potentially distort the aggregate distribution of income across factors.¹

I carry out my analysis in the context of the 50 employees threshold in the employees rights legislation which applies to French firms. Drawing upon a census of French manufacturing firms above 20 employees over the period 1995-2007, I structurally estimate the implicit cost of the regulation above 50 employees using observed distortions in firms factor demands in the neighborhood of the threshold. Using my estimates to calibrate a general equilibrium model, I get preliminary results suggesting that the regulation shifts the distribution of income towards labor, even if regulated firms substitute capital for labor. The rationale for this result being that the regulation generates a reallocation of market shares from large capital intensive firms to small labor intensive firms.

There are three main contributions to this paper. First I document new facts on firms behavior around a workforce threshold in the labor regulation. Namely, I show that regulated firms significantly increase there capital labor ratio in response to the regulation. Second, I propose a simple way to structurally estimate the implicit cost of the regulation. Incidentally, my approach delivers an estimate of the elasticity of substitution at the firm level. Since my structural model is stationary, I interpret it as long-run elasticity of substitution. Third, I extend the closed version of Melitz (2003) to allow firms to produce out of both capital and labor and to have non-hick neutral productivity differences. The calibration of this model allows me to analyze the distributional implication of the regulation.

The analysis developed in this paper proceed in several steps. First, I develop a structural model from which I identify the parameters of interest. It is a partial equilibrium model with heterogeneous firms producing a differentiated good through a production function which is CES in capital and labor. Firms face a regulation discontinuity which is modeled as an increase in the payroll tax rate above a given number of employees. Productivity differences across firms are driven by a single parameter and can be factor-biased. Based on this parameter, firms

1. See Raval (2011) and Forslid & Okubo (2011) for examples of papers documenting a positive relationship between size and capital intensity.

self-select into size so that high productivity firms are above the threshold, low productivity firms are below the threshold and there is a range of intermediate productivity firms which choose to bunch at the threshold. The model predicts a break in the capital-labor ratio and in the firm size distribution at the threshold respectively because (i) firms at the left and the right of the threshold have different technologies and (ii) because labor is more costly above the threshold.

Then I describe the French regulation along with some descriptive facts on firms behavior around the 50 threshold. The French case offers a perfect set-up to empirically assess the effect of size-dependent labor regulations. In fact, when a French firm exceeds the 50 threshold, it has to provide its employees with a handful of important additional rights (e.g. set-up a works council, agree on a profit sharing rule with its employees, etc.). I verify that main qualitative predictions of the model are empirically verified. In particular, I get reduced form evidence suggesting that firms substitute capital for labor at the threshold as well as above the threshold.

Using the model, I estimate the tax equivalent of the regulation and the capital-labor elasticity of substitution. My approach consists in estimating the individual factor demands for a given technology parameter under two counterfactual situations: one where regulation binds all firms irrespective of size and another one where there is no regulation. The estimated parameters are those which equate estimated counterfactual factor demands to their structural counterparts. I estimate the vector of factor demands for a firm which would locate exactly at the 50 threshold in the absence of regulation. Its counterfactual labor demand when regulation applies to all firms is identified from the excessive mass of firms at the threshold. The insight for why the excess mass is related to the size of a firm in the counterfactual situation when regulation binds all firms is simply that both are a function of the cost of the regulation: the model predicts that a larger cost of the regulation leads more firms to bunch and would also lead bunching firms to be smaller, should they be regulated. Counterfactual capital demands are identified from the relationship between the number of employees and the average relative factor expenditures in a neighborhood of the threshold.

I find that the regulation raises the unit cost of labor by 7.5% and total costs by 3%. Moreover, firms above 50 raise their capital-labor ratio by 5%. Estimated capital-labor elasticity of substitution is 0.6 and consistent with the literature².

Finally, I use estimated parameters to implement counterfactual exercises. I am especially interested in the response of the aggregate production and of the share of labor in national income to a variation of the micro cost of the regulation. To this end, I superimpose a general equilibrium layer to the problem of the firm. The general structure builds upon essentially follow Melitz (2003). I assume that upon entry, firms pay a sunk cost and randomly draw a capital and labor productivity. Once on the market, firms pay a fixed operating cost. The joint entry distribution of capital and labor productivity is calibrated to replicate the empirical size distribution and the empirical size-factor cost ratio below the threshold. I am still working on the results from this section.

My work relates to several strands of literature. It contributes to the literature on the

2. See Chirinko et al. (2004) for a review.

impact of (workforce) size-dependent regulations on firm's policy. These papers differ in the source of identification. Some have an exogenous source of variation in the regulation. This is the case of Bauernschuster (2009) and Kugler & Pica (2008) who use time-variations in the size threshold for dismissal protection legislation to identify an impact on the employment dynamics at the firm-level. Given their source of identification, they estimate a short-run to medium-run effect of the change in regulation. Some other papers like Boeri & Jimeno (2005) document the impact of the threshold by regressing firm level variables on a dummy variable for whether the firm is above the threshold or not. However, the obtained coefficient can not have a causal interpretation as firms are susceptible of choosing their position to the threshold based on unobservables. Ceci-Renaud & Chevalier (2010) and Schivardi & Torrini (2004) do not have an instrument variable either and so identify threshold effects from the distortion in firms' size dynamics around the discontinuity under the assumption that firms' dynamics is a smooth function of size in the absence of a threshold. I depart from these different contributions by adopting a structural approach while they rely on reduced forms.

Closer to what I do, Gourio & Roys (2012) and Garicano et al. (2013) estimate a structural model. Both are able to estimate the micro and macro cost of the regulation. However, since labor is the only production factor in both papers, no analysis of the substitution effects across factors is possible, whether at the micro or macro level. The present paper, by contrast, considers both capital and labor as production factors.

Finally, size-dependent labor regulation belongs to the broader family of heterogeneous firm-level regulation. As pointed out in a series of recent paper (e.g. Guner et al. 2008, Restuccia & Rogerson 2008, Hsieh & Klenow 2009), the heterogeneity of policy distortions at the firm-level can generate large aggregate inefficiencies by inducing a reallocation of resources across firms. I add by letting firms capital intensity vary within industries. This makes it possible to investigate the macro-level substitution effect induced by heterogeneous policies.

The remainder of the paper is organized as follows. Next section introduces the dataset, the content of the French regulation as well as some evidences of its effect on firms' behavior. Section 3 describes the structural model. Section 4 presents the estimation strategy. Section 5 reports the results. section 6 provides the results from preliminary policy experiments.

2 Regulatory Background and Motivating Evidence

In this section, I first describe the French Regulation and my dataset. Then I provide reduced form evidences that labor and capital demand respond to the regulation at the firm-level. These evidences motivate the structural model.

2.1 The French Regulation

French regulation has important thresholds at 10,20 and 50 employees. In the present study, I will focus on the 50 threshold. There are two reasons why the 50 threshold is arguably the most interesting one. First, the 50 threshold is generally recognized as the most binding to firms and as such, identifying its effects is easier from a statistical point of view and has a greater interest for economic policy. Second, existing works on the thresholds in

France have mostly focused on the 50 threshold so that my results will be comparable to the largest possible set of existing results.³

When a firm exceeds 50 employees in France, it has to deal with a long list of additional rules on labor and accounting (see Ceci-Renaud & Chevalier (2010) for the exhaustive list). For instance, above 50, a firm must:

- organize the election of a works council upon the request of at least one employee. In France, a works council is in charge of providing services to employees and to represent them with respect to executive management.
- set-up a committee for health, safety and working conditions.
- produce a *plan social* when more than 9 employees are fired simultaneously. A plan social is a legal process through which, among other things, firms must prove to some legal institutions that they have tried to find a new job to dismissed employees.
- share profits with their employees. The sharing rule must be jointly established by the managers and workers' representatives at the firm level.

When, exactly, does the regulation binds to a firm? The definition of the number of employees in the regulation is the monthly full-time equivalent workforce. This measure takes into account part-time workers and temporary workers but not trainees nor subsidized employment. The 50 threshold triggers a set of laws and not all of them have the exact same criterion as to how long a firm should stay above 50 before it applies. Most of the rules, however, become binding as soon a firm has spent more than 12 months, over the three past years, above 50. As I will explain at the time of presenting the data, my measure of workforce is the monthly full-time equivalent averaged over the 12 months of an accounting year. As such it is a good, although imperfect, measure of whether a firm is regulated or not. Moreover, some elements of the regulation are based on the firm exceeding a given number of employees *or* a given turnover (around 4 million Euros).⁴ This is another factor which makes my measure of workforce a slightly noisy proxy for whether a firm is regulated or not⁵

2.2 Data

My empirical analysis is ran on the EAE (Annual Business Survey) dataset. This dataset is a panel of the universe of French manufacturing firms above 20 employees over 1995-2007. The data is collected by INSEE (French National Statistical institute) out of about 20000 firms each year. The dataset also contains a randomly drawn sample of firms below 20 which I never use in my analysis. I also limit my analysis to firms which main activity is in the manufacturing sector. The main elements of firms' balance sheet are present in the database (e.g., turnover, value added, investment, materials, etc.). There is no variable at the worker level (e.g., individual workers' wage, occupation or education). The dataset has no information either on whether the main elements of the regulation are implemented by a firm (works council, committee for health and safety, etc.). Only is there information on profit

3. Note that even if I focus on the 50 employees threshold, my estimation strategy virtually applies to any threshold of the regulation.

4. The turnover criterion is only on accounting rules, not on labor rules.

5. note, however, that following the insight from the ordinary least squares literature, this noise should play against me by biasing estimated parameters toward zero.

shared by the firm with its employees.

I essentially use three variables: the workforce size, the wage bill and the gross stock of fixed assets. I have separate information on tangible and intangible fixed asset, at the beginning of the year and at the end of the year. Hereafter, the variable I name “capital” is the sum of these four measures of fixed assets divided by two (I average the beginning and the end of year stocks). For now I have not implemented the usual treatments to the stock of capital (i.e., deflating each vintage of capital by the price index at the time of purchase and implementing perpetual inventory methods). Since capital is one of the most important variable in my analysis, it is clear that these treatments will be present in the final version of the paper.

At the core of my analysis is the factor cost ratio variable which I define as the ratio of the user cost of capital and wage bill. I obtain the user cost of capital very simply by multiplying the stock of capital, defined above, by a unit user cost of capital $r = 0.1$. Here I follow the value proposed by Hsieh & Klenow (2009). It corresponds to a rental rate of 5% and a depreciation rate of 5%. The purpose of this calibration is just to present factor cost ratios with sensible levels. My estimation, however, is independent of this calibration as only do relative factor cost ratios across firms serve my identification.

In order to compute TFP, I also construct the value added and materials variables. I define materials as the sum of the cost of goods and the cost of materials. I compute TFP using Levinsohn & Petrin (2003) semi-parametric method. Once again, value added and materials are not deflated although they will be in future versions of the paper.

Table 3.1: Correlation Matrix for Main Variables

	Value Added	Fixed Assets	Employees	Factor Cost Ratio
Value Added	1.00			
Fixed Assets	0.88 (0.00)	1.00		
Employees	0.88 (0.00)	0.81 (0.00)	1.00	
Factor Cost Ratio	0.03 (0.00)	0.04 (0.00)	0.01 (0.00)	1.00

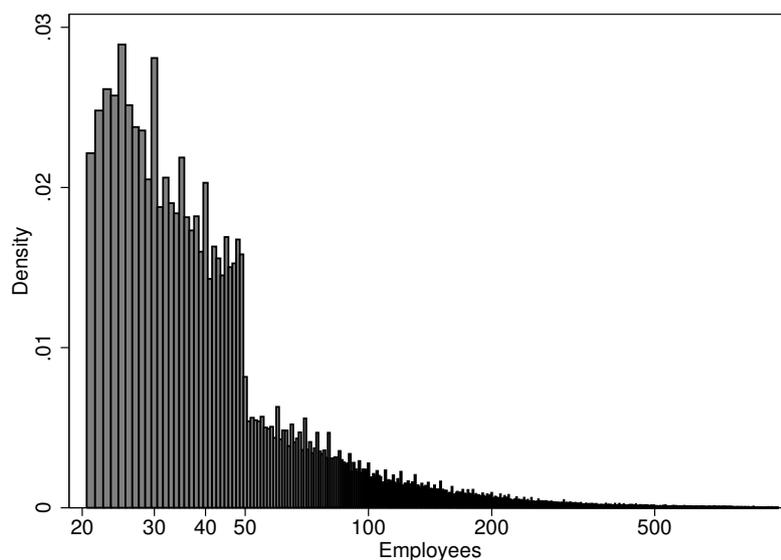
Note: P-values in parentheses. Correlations are computed over the whole sample.

Table 3.1 reports correlations for main variables over the whole sample. Not surprisingly, value added, capital and employees are positively correlated. More interestingly, these different measures of size are positively and significantly correlated with the factor cost ratio. This is a finding consistent with previous evidences in the literature (see Raval 2011). It is suggestive of a model where productivity differences across firms are factor biased. In appendix A, I present additional descriptive statistics on the main variable of interest.

2.3 Motivating evidence

The present paper proposes to identify the implicit cost of the regulation along with the capital-labor elasticity of substitution. To this end, I will structurally estimate firm-level response of capital and labor demands. This sub-section gives suggestive evidence that there are such responses.

Figure 3.1: Firm Size Distribution between 20 and 1000 Employees



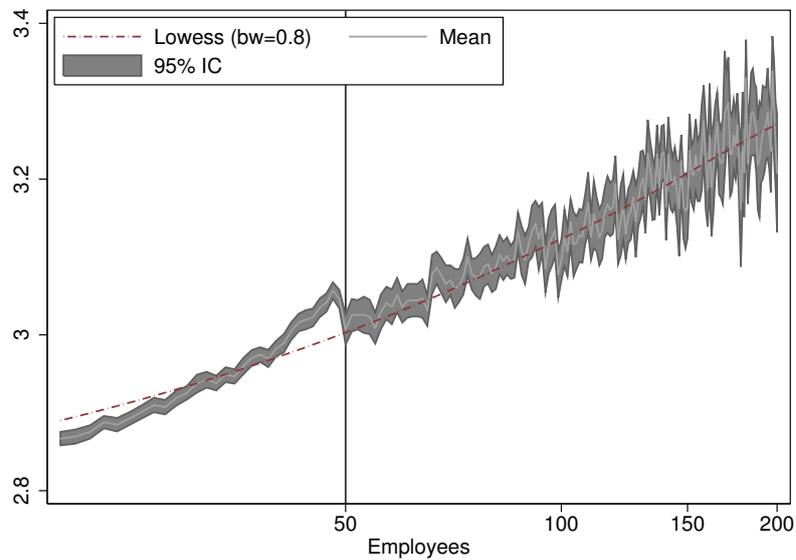
In figure 3.1, I plot the workforce size distribution of French firms on a log x-axis. This distribution exhibits features which are common in the literature: it is highly skewed and its upper tail is akin to a pareto distribution.⁶ Yet, visual inspection reveals two departures from the traditional pattern of firm size distribution. First, there is a break at 50 employees: the number of firms is approximately twice larger at 49 than at 50. Second, while the distribution follows a negative trend everywhere, it seems to be increasing just below the threshold. This excessive mass supports the existence of firms bunching at the threshold so as not to get regulated. In my estimation, this excessive mass identifies the labor demand response to the regulation.

As an additional evidence of bunching, in figure 3.2 I plot workforce against log TFP. The x-axis is log-scaled. Figure 3.2 features a positive relationship between TFP and size: it seems that more productive firms self-select into larger sizes. However, one can observe that the average productivity below the threshold, i.e. between 40 and 49, significantly differs from the value predicted from non-parametrically regressing productivity on size. Put differently, some firms below the threshold are too small given their productivity, which indicates that some firms shrink their labor demand to stay below 50.

Let us now look at the response of the factor cost ratio as it informs on the substitution behavior of firms. I first describe the relationship between employees and the factor cost ratio.

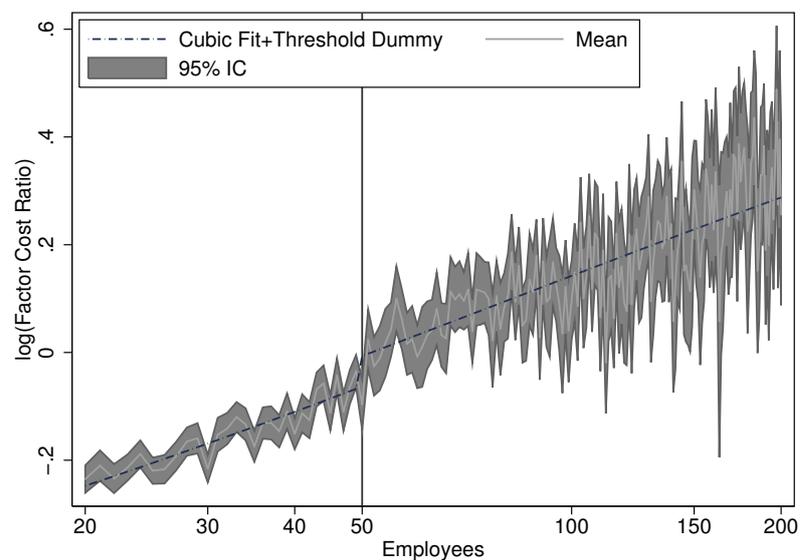
6. Influential papers documenting the pareto shape of firm size distribution include Simon & Bonini (1958) and Luttmer (2007)

Figure 3.2: log TFP Versus Size



Note: TFP is estimated using Levinsohn and Petrin (2003) methodology with intermediate input as a control for unobserved productivity.

Figure 3.3: Factor Cost Ratio Versus Size



Note: 'Cubic Fit+Threshold Dummy' is the prediction of a regression of log Factor Cost Ratio over a dummy ($Employees \geq 50$), $Employees$, $Employees^2$ and $Employees^3$, a set of year-2-digit industry dummies and a constant term.

In figure 3.3, I report the firm-level relationship between employees and the average factor cost ratio. Prior to computing the average, I demean the factor cost ratio at the year-2-digit industry level (36 industries) so that relationship in figure 3.3 is within year-industry. This relationship is upward sloping consistently with the correlation matrix presented in previous subsection. More interestingly, there is a positive break in the relationship at the threshold. This means that firms tend to be more capital intensive above 50. In order to check the statistical magnitude of this break, I regress the factor cost ratio on a third order polynomial and a dummy equal to one if a firm has more than 50 employees. The results are in table 3.1.

Table 3.2: Factor Cost Ratio VS. Size

	log Factor Cost Ratio		
	(1)	(2)	(3)
“Above 50 Employees” Dummy	0.082 ^a (0.000)	0.057 ^a (0.000)	0.010 (0.104)
Employees	-0.286 ^a (0.000)	-0.139 ^b (0.036)	0.302 ^a (0.000)
Employees ²	0.113 ^a (0.000)	0.083 ^a (0.000)	-0.084 ^a (0.000)
Employees ³	-0.008 ^a (0.000)	-0.006 ^a (0.000)	0.005 ^a (0.000)
Fixed Effect Model	No OLS	2-digits Industry×year OLS	Firm OLS
Observations	242437	242437	242437
R^2	0.047	0.043	0.005
rho		0.150	0.859

p-values in parentheses

^c $p < 0.1$, ^b $p < 0.05$, ^a $p < 0.01$

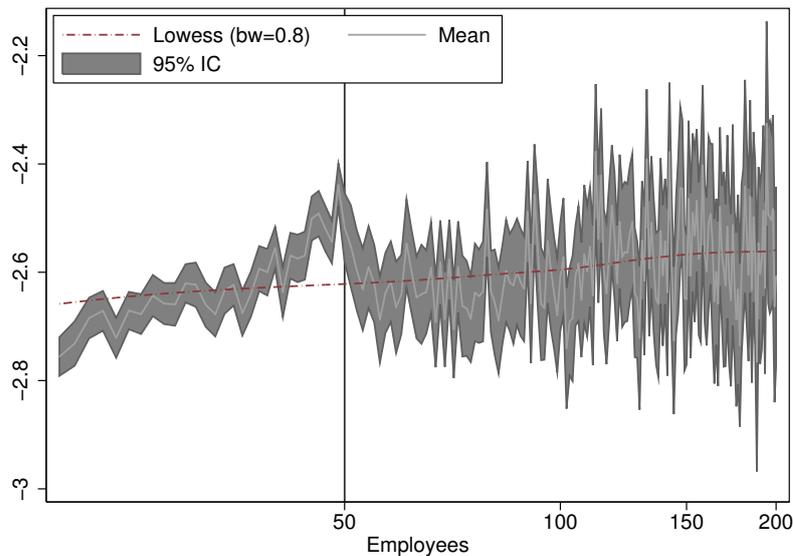
Coefficients from column (1) were obtained without controls. The coefficient associated to the dummy variable is 0.08 which means that a firm at 50 employees is on average 8% more capital intensive. Adding industry fixed effects (Column (2)) lowers the coefficient to 5% but significance remains. The effect of adding industry fixed effects on the coefficient suggests that more capitalistic industries also have a larger proportion of firms above 50, on average. Specification in column (3) has firm fixed effects and the “above 50 employees” dummy is close to be significant at 10%.

How to interpret these results? Clearly, the dummy in both first specifications has no causal interpretation. This is because, as the analysis of the TFP and the firm distribution make clear, firms self-select with respect to the threshold. Therefore, firms at 49 *are not* comparable to firms at 50, a priori. There could be a positive break in capital intensity at 50 employees even if firms do not substitute capital for labor as a response to the regulation. It will for instance be the case if (i) more productive firms self-select above 50 and (ii) more-productive firms are also more capital intensive (for reasons exogenous to the regulation). For now, all one can conclude from the positive break is that it is consistent with firms substituting

capital for labor because of the regulation.

I see three possible reasons why the dummy coefficient is not significant in the within-firm specification. First, firms do not substitute capital for labor at all. Second, within-firm capital-labor substitution takes time so that it can not be observed in the time horizon of the panel. My structural estimation of the capital-labor substitution will be in the cross-section of firms. It will have a long-run interpretation so that it will be consistent with the absence of substitution obtained in the time series. Third, firms might anticipate that they will exceed 50 in the future and therefore increase their capital-labor ratio before to get regulated.

Figure 3.4: Investment over Fixed Assets Versus Size



As a way to test these rationales for the low explanatory power of the “above 50” dummy variable, I look at the investment behavior of firms. In figure 3.4 is plotted the mean log ratio of investment to fixed assets (hereafter the investment rate), conditional on workforce. The investment rate exhibits a clear bulge below the threshold. Interestingly, there is no equivalent distortion visible on the factor cost ratio below 50, in figure 3.3. Combining these features suggests that firms do not bunch long enough so their investment effort translates into an excess factor cost ratio below the threshold. This analysis supports the idea that firms invest more intensively below the threshold in order to have a higher capital labor ratio once above the threshold.

With these empirical facts in hand, I turn to the description of the structural model. The structure will guide the rest of our empirical analysis and will allow me to identify a causal effect of the regulation on factor demands.

3 Theory: Heterogeneous Firms and Legal Threshold

This section first lays down the structure of the model. At this stage, I limit my presentation to the problem of the firm in a partial equilibrium set up as it suffices for my estimation

(I postpone the description of the rest of the model to the counterfactual section). Then, I derive the empirical predictions that will be used to identify the parameters.

3.1 The Set-up

Technology and Demand I consider a closed and static economy populated with heterogeneous firms producing a differentiated good out of capital and labor. The technology of a firm is CES and is summarized by a scalar parameter α which determines both the efficiency of capital $A(\alpha)$ and the efficiency of labor $B(\alpha)$. The expression of the production function is

$$y(k, l; \alpha) = \left[(A(\alpha_i)k_i)^{\frac{\sigma-1}{\sigma}} + (B(\alpha_i)n_i)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1}}, \quad (3.1)$$

with n_i the number of employees, k_i the stock of capital and $\sigma \geq 0$ the capital-labor elasticity of substitution. I do not postulate the pattern of complementarity across factors: σ can be smaller or larger than unity. Through following assumption, I ensure that the production cost of the firm are decreasing with α

Assumption 1. *Functions $A(\bullet)$ and $B(\bullet)$ are strictly increasing*

As indicated in assumption 1, I assume both labor efficiency and capital efficiency to increase with α . Parameter α is therefore a positive measure of a firm's TFP and I refer to it as "efficiency" in the rest of the paper. α is distributed across firms according to a PDF function $g(\alpha)$. This function is subject to following assumption:

Assumption 2. *$g(\alpha)$ is smooth.*

Assumption 2 will be useful in the estimation presented in section 4. Intuitively, under assumption 2 and in the absence of a regulation threshold, the model delivers a smooth distribution of firm size. So the impact of the regulation can be identified from deviations to a smooth distribution in the neighborhood of the threshold .

All firms face an identical isoelastic demand function

$$y_i = p_i^{-\varepsilon} Z, \quad (3.2)$$

with ε the price elasticity of demand, p_i the price charged by firm i and Z a demand shifter common to all firms. As is standard, we impose $\varepsilon > 1$ to ensure that firms charge finite prices.

The Regulation I model regulation as a tax rate τ on labor which applies to firms above \bar{n} ($\bar{n} = 50$ in my application to the French regulation).⁷

The expression of total labor costs of a firm with n employees is:

$$(1 + \tau \mathbb{1}\{l \geq \bar{l}\})wn.$$

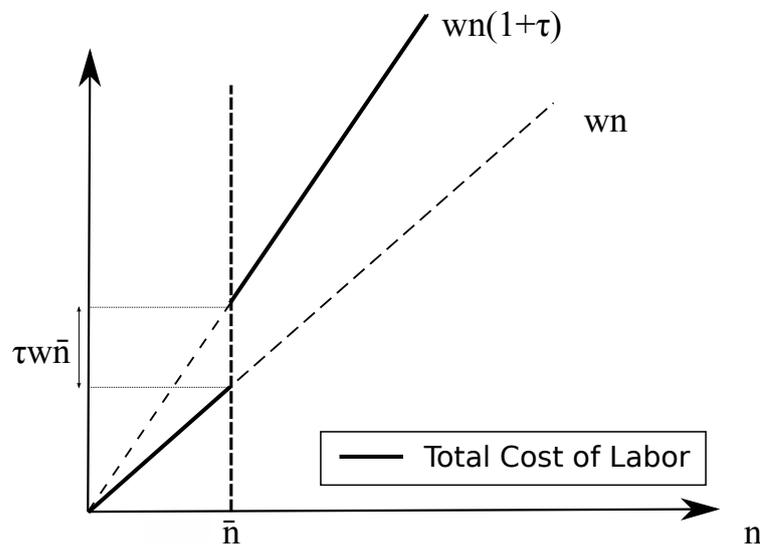
Here, I implicitly assume that the regulation does not impact the cost of capital. This is obviously a simplifying assumption which will not be verified in the general case of a size-

7. I do not include a fixed cost component to the regulation. This choice is motivated by Garicano et al. (2013) who estimate a small and insignificant fixed component of the 50 employees threshold in France.

dependent regulation. However, most rules binding above 50 employees in France are relative to EPL. In the French case, it seems therefore reasonable to assume that the productivity of other production factors is not affected by the regulation.

The regulation could also have dynamic components that we omit here. For instance, the French regulation contains elements relative to redundancy rules which potentially impact the firing costs of the firm. However, as proved by Bentolila & Bertola (1990), when financial markets are complete and firms can fully hedge, firing costs are equivalent to a per period cost on labor. The regulation could also factor in firms' dynamics by taking the form of a sunk cost. Gourio & Roys (2012) investigate that possibility by assuming that firms incur an adjustment cost to activate the regulation and that once activated, the regulation can be maintained at no cost.⁸

Figure 3.1: Total Cost of Labor



In figure 3.1, I plot the cost function of a firm. Observe that at \bar{n} , the cost of labor features both a positive discontinuity of magnitude $\tau w \bar{n}$ and a change of slope from w to $w(1 + \tau)$. The discontinuity results from the fact that when regulation binds, it applies to every workers and not only to the \bar{n} -th worker and beyond. This discontinuity in firms' production set distorts firms' decisions. Next subsection solves the problem of the firm and analyzes these distortions.

3.2 The problem of the firm

In this subsection, I derive the expression of firms optimal output and factor demand. I start by investigating the counterfactual cases where regulation binds all firms ($\bar{n} = 0$) and where regulation binds no firms ($\bar{n} = +\infty$). The fact that in these cases firms do not have to choose to be regulated or not facilitates the analysis. Then I present the actual case:

8. Although Gourio & Roys (2012) get results suggestive of the presence of a sunk cost, they are unable to estimate this jointly with a per period cost of the regulation. Moreover, the break in the size distribution at the threshold found in the data is strongly supportive of a per period-per unit cost of the regulation.

$0 < \bar{n} < +\infty$.

Firm's Decisions in Absence of a Threshold

Let $c_0(\alpha)$ and $c_1(\alpha)$ be the unit production cost function of a firm with efficiency α respectively in the absence of regulation and when regulation applies to all firms. The expression of these functions is standard under CES technology

$$\begin{cases} c_0(\alpha) &= \left[\left(\frac{r}{A(\alpha)} \right)^{1-\sigma} + \left(\frac{w}{B(\alpha)} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}} \\ c_1(\alpha) &= \left[\left(\frac{r}{A(\alpha)} \right)^{1-\sigma} + \left(\frac{(1+\tau)w}{B(\alpha)} \right)^{1-\sigma} \right]^{\frac{1}{1-\sigma}}. \end{cases} \quad (3.3)$$

The production cost is a CES index of the cost per efficiency unit of labor and capital. Given assumption 1, $c_i(\alpha)$ is strictly decreasing with α . Due to the regulation cost, $c_1(\alpha)$ is larger than $c_0(\alpha)$. The optimal price and output are

$$\begin{aligned} p_i(\alpha) &= \frac{\varepsilon}{\varepsilon - 1} c_i \\ y_i(\alpha) &= \left(\frac{\varepsilon}{\varepsilon - 1} c_i \right)^{-\varepsilon} Z, \end{aligned}$$

Because a firm's production costs are larger in presence of regulation, $p_1(\alpha) > p_0(\alpha)$ and, as a result, output is $y_1(\alpha) > y_0(\alpha)$.

Labor demand $n_i(\alpha)$ involved in the production of output $y_i(\alpha, \tau)$, $i = 0, 1$ has usual CES expression:

$$\begin{cases} n_0(\alpha) &= \left(\frac{w}{c_0(\alpha)} \right)^{-\sigma} B(\alpha)^{\sigma-1} y_0(\alpha) \\ n_1(\alpha) &= \left(\frac{(1+\tau)w}{c_1(\alpha)} \right)^{-\sigma} B(\alpha)^{\sigma-1} y_1(\alpha). \end{cases} \quad (3.4)$$

Regulation reduces labor demand through two effects. First, it leads firms to substitute capital for labor as it shifts up the relative cost of labor. Second, it causes a scale effect as the production costs go up which shrinks the output of the firm. It follows that $n_1(\alpha)$ is smaller than $n_0(\alpha)$.

The way labor demand varies with α is more ambiguous. To see this, let us write down the elasticity of labor demand to α

$$\frac{\partial \log n_i(\alpha)}{\partial \alpha} = \underbrace{-\frac{\partial \log B(\alpha)}{\partial \alpha}}_{\text{Labor Saving Effect} < 0} + \underbrace{\sigma \frac{\partial \log \frac{B(\alpha)}{c_i(\alpha)}}{\partial \alpha}}_{\text{Substitution Effect}} + \underbrace{\varepsilon \frac{\partial \log c_i(\alpha)}{\partial \alpha}}_{\text{Scale Effect} > 0}. \quad (3.5)$$

Efficiency α has three effects on labor demand. First effect is what we call a "Labor saving effect". When α goes up, so does labor productivity so that a firm needs fewer physical units of workers for a given number of efficiency units of labor. This effect is unambiguously negative on labor demand. The second effect is a scale effect: production costs decrease with α which allows a firm to charge lower prices and sell more. This effect is unambiguously positive. The last effect is a substitution effect. It comes from the fact that I did not impose α to be a Hicks

neutral technology parameter. Therefore, a change in α may induce a substitution effect by shifting the relative efficiency $\frac{A}{B}(\alpha)$.

Assumption 3. $\frac{\partial \log B}{\partial \alpha} < \sigma \left(\frac{\partial \log B}{\partial \alpha} - \frac{\partial \log A}{\partial \alpha} \right) + \varepsilon \frac{\partial \log A}{\partial \alpha}$

3 guarantees that labor demand is strictly increasing with α which will prove useful in the estimation. Intuitively, it constraints the labor saving effect to be more than compensated by both other effect. If $\varepsilon > 1 > \sigma$, as we find in the data, this amounts to assuming the labor productivity B do not increase “too much” with α relative to capital productivity A .

I now define $\rho_i(\alpha) \equiv \frac{rk_i}{wn_i}(\alpha)$ as the factor cost ratio. This function will show useful in my estimation as it will allow me to use the observed substitution behavior of firms to back out structural parameters. $\rho_i(\alpha)$ verifies

$$\begin{cases} \rho_0(\alpha) &= \left(\frac{r}{w}\right)^{1-\sigma} \left(\frac{A}{B}(\alpha)\right)^{\sigma-1} \\ \rho_1(\alpha) &= (1+\tau)^\sigma \left(\frac{r}{w}\right)^{1-\sigma} \left(\frac{A}{B}(\alpha)\right)^{\sigma-1}. \end{cases} \quad (3.6)$$

ρ_1 is larger than ρ_0 and $\rho_1(\alpha)$ is increasing with τ .⁹

The factor cost ratio depends on α through $\frac{A}{B}(\alpha)$. Actually, relative factor efficiency is the only source of variation in the factor cost ratio across firms (conditional on the regulation status). We could have made the structure of the model more complex in order to incorporate alternative sources of variation like adjustment costs, heterogeneous factors costs, heterogeneous credit constraints could come as complements. However, as will appear clearly below, the parsimony of my model pays off through the simplicity of the estimation. In now describe firms’ decisions in the presence of a regulation threshold.

The General Case

I now consider the case $\infty > \bar{n} > 0$. Unlike counterfactual cases studied above, firms now have to choose whether to be regulated or not, consistently with the French context. Let α_c be such that

$$n_0(\alpha_c) = \bar{n}.$$

Firms below α_c are unaffected with the regulation since their optimal demand, $n_0(\alpha)$, is inferior to \bar{n} . Firms above α_c trade-off between (i) bunching at the threshold and thence giving up revenues generated by a larger workforce or (ii) expanding beyond the threshold and paying a higher unit cost per worker. The profit from bunching is

$$\pi_b(\alpha) = \max_k \{R(k, \bar{n}, \alpha) - w\bar{n} - r.k\},$$

9. The ranking of ρ_1 and ρ_0 would be more ambiguous if I had included the cost of the regulation in the cost of labor used to compute the factor cost ratio (i.e. if I had considered an alternative factor cost ratio $\tilde{\rho}_1(\alpha) \equiv \frac{rk_1}{(1+\tau)wn_1}(\alpha)$). In that case, there would be a negative price effect of the regulation on the factor cost ratio. This effect would even dominate the substitution effect when $\sigma < 1$. In terms of the identification, this point implies that we should make sure that what we observe in the data above 50 employees is actually ρ_1 and not $\tilde{\rho}_1$ because it matters for the way my structural model rationalizes the discontinuity in the factor cost ratio at the threshold. As a matter of fact, in the French case, regulation triggers monetary transfers which are not recorded as wages. Therefore, I am confident that the factor cost ratio plotted in figure 3.3 is indeed the empirical counterpart of ρ_i .

with $R(k, n; \alpha)$ the revenue production function of a firm α obtained by combining equations (3.1) and (3.2) :

$$R(k, l; \alpha) = Z^{\frac{1}{\varepsilon}} \left[(A(\alpha)k)^{\frac{\sigma-1}{\sigma}} + (B(\alpha)l)^{\frac{\sigma-1}{\sigma}} \right]^{\frac{\sigma}{\sigma-1} \frac{\varepsilon-1}{\varepsilon}}. \quad (3.7)$$

It results that the profit loss from bunching instead of being regulated is

$$\Delta(\alpha) \equiv \pi_1(\alpha) - \pi_b(\alpha),$$

, with $\pi_1(\alpha)$ the optimal profit of a regulated firm:

$$\pi_i(\alpha) = Z^{\frac{1}{\varepsilon}} y_i(\alpha)^{\frac{\varepsilon-1}{\varepsilon}}. \quad (3.8)$$

A firm bunches at the threshold if it is affected by the regulation ($\alpha \geq \alpha_u$) and its profit loss from bunching is negative ($\Delta \leq 0$). Respectively, a firm chooses to exceed \bar{n} if $\alpha \geq \alpha_u$ and $\Delta > 0$. Simple analysis of $\Delta(\alpha)$ delivers following proposition:

Proposition 4. *Let α_u be such that $\Delta(\alpha_u) = 0$ and $\alpha_u > \alpha_c$ then:*

1. α_u is unique
2. $\Delta'(\alpha_u) > 0$
3. $n_1(\alpha_u) > \bar{n}$

Proof. See Appendix B □

Point 1 of proposition 4 imply that there is a cutoff firms α_c indifferent between bunching or exceeding the threshold:

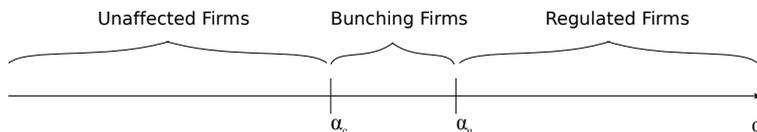
$$\Delta(\alpha_c) = 0. \quad (3.9)$$

From point 2 of proposition 4 it follows that firms between α_c and α_u bunch at \bar{n} and firms above α_u exceed \bar{n} . I represent the sorting of efficiencies in figure 3.2. The rationale for this sorting is that returns to capital are decreasing and optimal output is increasing with α . Since the only way of increasing its output for a bunching firm is to consume more capital, the profit loss from bunching, $\Delta(\alpha)$, is increasing. As established in third bullet of proposition 4, no firm finds it optimal to have a workforce size in a range $[\bar{n}, n_1(\alpha_u)]$. This is because regulation shifts up the average cost of labor: by crossing \bar{n} , a firm is hit by a finite cost $\tau w \bar{n}$ which can not be compensated for by a marginal increase of n .

To summarize, labor demand $n(\alpha)$ and factor cost ratio $\rho(\alpha)$ verify:

$$\{n(\alpha), \rho(\alpha)\} = \begin{cases} \{n_0(\alpha), \rho_0(\alpha)\} & \text{if } \alpha < \alpha_c \\ \{\bar{n}, \frac{rk_b(\alpha)}{w\bar{n}}\} & \text{if } \alpha_c \leq \alpha < \alpha_u \\ \{n_1(\alpha), \rho_1(\alpha)\} & \text{if } \alpha \geq \alpha_u, \end{cases} \quad (3.10)$$

Figure 3.2: Sorting of Efficiencies



With $k_b(\alpha)$ the optimal capital of a bunching firm. $k_b(\alpha)$ equalizes the marginal revenue from capital at \bar{n} with the rental rate:

$$\frac{\partial R}{\partial k}(k_b(\alpha), \bar{n}; \alpha) = r. \quad (3.11)$$

In figure 3.3 and 3.4, I give a sketch of $n(\alpha)$ and $\rho(\alpha)$. Labor demand grows with α below α_c . It is flat over $[\alpha_c, \alpha_u]$ as firms bunch at \bar{n} . There is a positive discontinuity in $n(\alpha)$ at α_u , then $n(\alpha)$ increases again over $[\alpha_u, \alpha_{max}]$, although more slowly than below α_c , due to the per unit cost of the regulation. As to $\rho(\alpha)$, the model does not determine whether it is increasing, decreasing or even monotonic. Consistently with figure 3.3 where factor cost ratio is increasing, I plotted figure 3.4 assuming $\rho(\alpha)$ is increasing. Over $[\alpha_u, \alpha_{max}]$, bunching firms substitute capital for labor in order to mitigate the cost of bunching so that $\rho(\alpha)$ diverges positively from $\rho(\alpha, 0)$ beyond α_c . So far I have expressed relationships between observables and α , which is unobservable. Next subsection derives empirical predictions from the model.

3.3 Empirical Predictions

The purpose of this subsection is twofold. First, I show that the simple model presented above is capable of replicating the features observed in the data around the threshold. Second, I lay the ground for the estimation section below. In order to achieve these ends, I derive the relationship between the factor cost ratio and workforce as well as the firm size distribution. I verify that, consistently with the data, the model predicts (i) an excessive mass of firms below the threshold, (ii) a break in the size distribution and (iii) a break in the factor cost ratio at the threshold.

Figure 3.3: Labor Demand and Efficiency

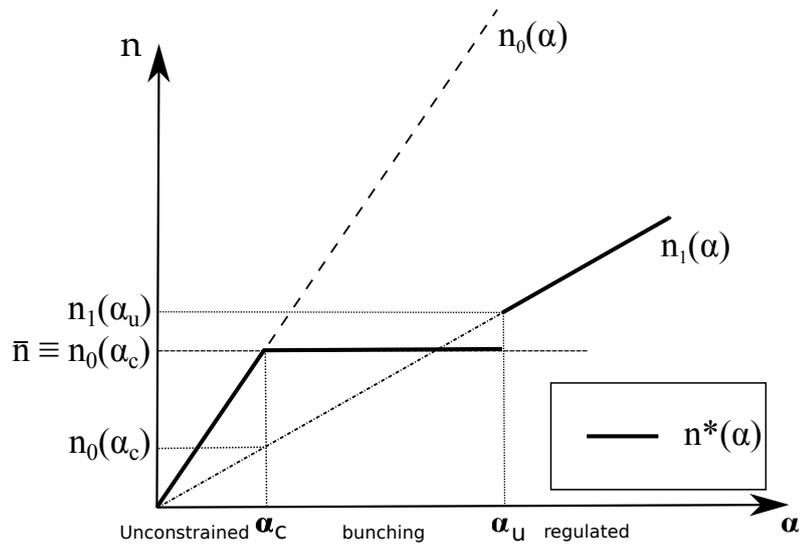
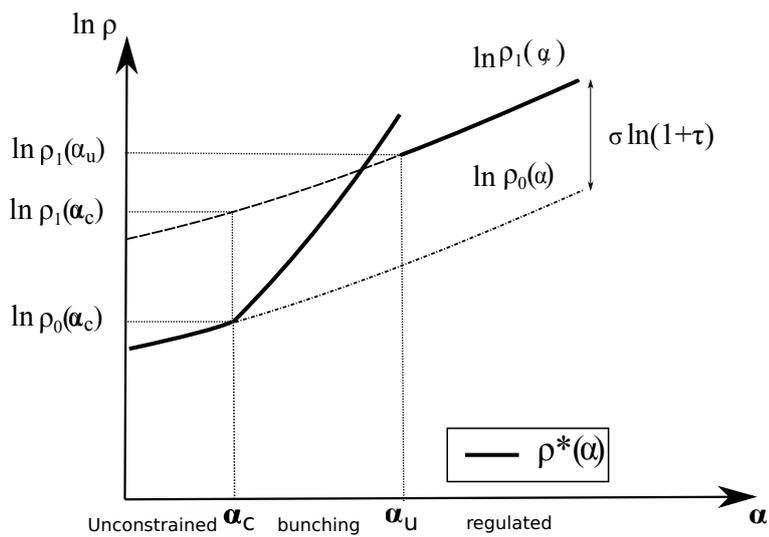


Figure 3.4: Factor Cost Ratio and Efficiency



Size Distribution While I assume workforce size is a continuous decision variable of the firm, I model *observed* labor as a discrete variable. In other words, firms pick their workforce size n in \mathbb{R}^+ but the econometrician only observes the integer part $\lfloor n \rfloor \in \mathbb{N}$. Let $\alpha_i(n)$ be the inverse of function $n_i(\alpha)$. For $n \notin [\bar{n}, n_1(\alpha_u)]$, i.e for n not in the region above the threshold ‘deserted’ by firms, observed size distribution $g^n(n)$ verifies:

$$g^n(n) = \begin{cases} g_0^n(n) & \text{if } n < \bar{n} \\ g_1^n(n) & \text{if } n \geq \lfloor n(\alpha_u, \tau) \rfloor + 1, \end{cases} \quad (3.12)$$

with

$$g_i^n(n) = \int_{\alpha_i(n+1)}^{\alpha_i(n)} g(\alpha) d\alpha$$

$g_1^n(n)$ (respectively $g_0^n(n)$) can be interpreted as the size distribution in a counterfactual economy where the regulation applies to all firms (respectively applies to no firm). $g_1^n(n)$ is parametrized by τ because τ drives the relationship between efficiency and size $\alpha_1(n)$: the larger τ , the larger the efficiency α associated to a given size. Consequently, and as shown in equation (3.12), regulation distorts the size distribution of firms not only in a neighborhood of the threshold but for any size above \bar{n} .

True size of bunching firms is \bar{n} minus some epsilon so that their observed size is $\bar{n} - 1$. This mass of firms comes on top of firms which are unaffected by the regulation and which optimal size is in $[\bar{n} - 1, \bar{n}]$:

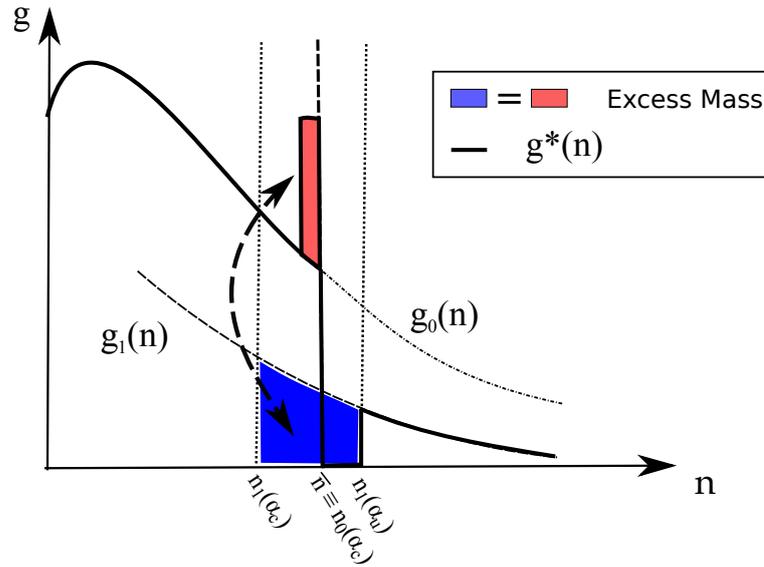
$$g^n(\bar{n} - 1) = g_0^n(\bar{n} - 1) + \int_{\alpha_c}^{\alpha_u} g(\alpha) d\alpha. \quad (3.13)$$

So the model predicts an excessive mass of firms $\delta^g \equiv \int_{\alpha_c}^{\alpha_u} g(\alpha) d\alpha$ at $\bar{n} - 1$. All else equal, this mass is increasing with τ as α_u , the efficiency level making a firm indifferent to bunch or the exceed \bar{n} , is obviously increasing with τ itself. To conclude on the size distribution, no firm is observed over the segment $[\bar{n}, \lfloor n_1(\alpha_u, \tau) \rfloor - 1]$. To sum up, the general expression of firm distribution is:

$$g^n(n) = \begin{cases} g_0^n(n) & \text{if } n < \bar{n} \\ g_0^n(\bar{n}) + \delta^g & \text{if } n = \bar{n} \\ 0 & \text{if } n \in [\bar{n}, \lfloor n_1(\alpha_u) \rfloor - 1] \\ \int_{n_1(\alpha_u)}^{\lfloor n_1(\alpha_u) \rfloor + 1} g(\alpha) d\alpha & \text{if } n = \lfloor n_1(\alpha_u) \rfloor \\ g_1^n(n) & \text{if } n \geq \lfloor n_1(\alpha_u) \rfloor + 1, \end{cases}$$

Figure 3.5 gives a representation of the firm size density. Consistently with the data, it features a break at 50 due to the fact that the labor demand function is distorted above the threshold, and hence the size distribution. The model also predicts an excessive mass below 50 and the absence of firms over a size interval above 50. This is not verified in the data. There are, at least, two possible explanations to that. First, the size distribution is measured with error. As explained when presenting the data, this might be due to the fact that my

Figure 3.5: Firm Size Density



measure of workforce does not exactly match the one from the regulation. Second, unlike what I assume in the model, the regulation tax might induce a fixed cost on top of a marginal cost. If this fixed cost is negative, this might compensate the finite cost $w\bar{n}\tau$ which hits firms at the threshold so that firms do not need to exceed the threshold by a finite amount of workers to find it profitable to be regulated. In order to test the latter explanation, in figure 3.6 I plot the expected annual growth of the number of employees conditional on growing. This variable peaks at the threshold and decreases quickly above it. It seems that firms “jump” when they cross the threshold. This is consistent with the idea that regulation induces a positive discontinuity in the cost function of firms, and that the discontinuity in the support of the distribution above the threshold is masked by some measurement error.

Capital-Labor ratio VS. Size I now derive the model’s prediction regarding the relationship between the factor cost ratio and size, which the second relationship I use in my estimation. Mean factor cost ratio given a size $n \notin [\bar{n}, n_1(\alpha_u)]$ equals

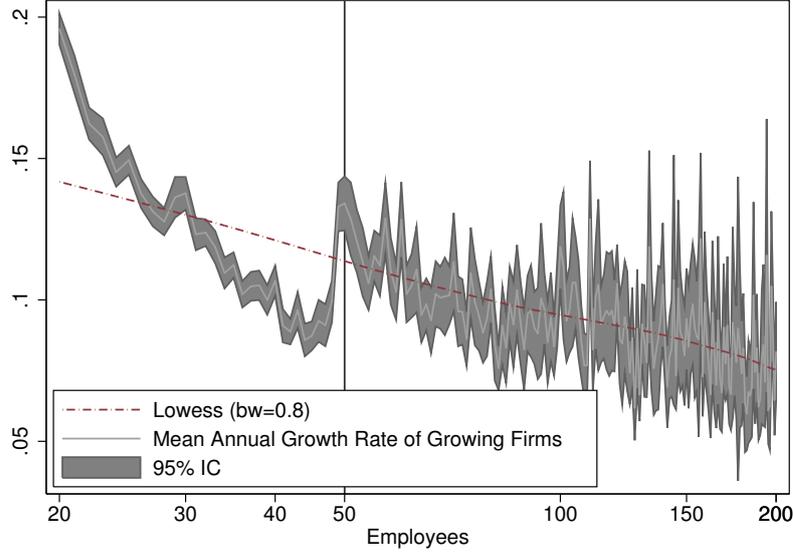
$$\rho^n(n) \equiv \mathbb{E} \left[\frac{rk}{wn} \middle| n \right] = \begin{cases} \rho_0^n(n) & \text{if } n < \bar{n} \\ \rho_1^n(n) & \text{if } n \geq \lfloor n_1(\alpha_u) \rfloor + 1, \end{cases}$$

with

$$\rho_i^n(n) = \frac{\int_{\alpha_i(n)}^{\alpha_i(n+1)} \rho_i(\alpha) g(\alpha) d\alpha}{\int_{\alpha_i(n)}^{\alpha_i(n)} g(\alpha) d\alpha}, \quad (3.14)$$

the average factor cost ratio of firms with size n when regulation applies to all firms ($i=1$) and to no firms ($i=0$). Hereafter, I make the approximation

Figure 3.6: Do firms “jump” when crossing the 50 threshold ?



Notes: This figure displays the average growth of the number of employees conditional on (i) the current number of employees and on (ii) the current number of employees being strictly larger than the number of employees in the following year. Lowess line is a locally weighted scatterplot smoothing.

$$\rho_i^n(n) = \rho_i(\alpha_i(n)), \quad (3.15)$$

i.e. I do as if firm' size was not observed with a rounding error. Similarly to what happens on the size distribution, regulation distorts the size-capital intensity relationship above \bar{n} . To better understand this distortion, one can write the expression of the distortion $\delta^\rho(n) \equiv \rho_1^n(n) - \rho_0^n(n)$ for τ small:

$$\delta^\rho(n) = \overbrace{\frac{\partial \rho_1(\alpha_1(n))}{\partial \alpha} \times \frac{\partial \alpha_1(n)}{\partial \tau}}^{\text{Selection Effect}} + \overbrace{\frac{\partial \rho_1(\alpha_1(n))}{\partial \tau}}^{\text{Substitution Effect}}.$$

$\delta^\rho(n)$ has two parts: a selection effect and a substitution effect. Selection effect comes from the fact that regulated firms are more efficient for a given size and that higher efficiency is potentially associated with a different capital intensity. Selection effect differs from 0 insofar as α causes factor biased technological differences ($\frac{\partial \rho_1(\alpha_1(n))}{\partial \alpha} \neq 0$). Substitution effect reflects the within-firm capital-labor substitution that occurs as a response to the upward shift in the relative cost of labor to capital created by the regulation. It is of the same sign as σ . Remark that if σ tends to 0 (and so substitution effect is nil) and technological differences are biased, regulation still creates a distortion in the factor cost ratio-size relationship, just because of the selection effect. A regression discontinuity approach would take $\delta^\rho(\bar{n})$ as an estimate of the substitution effect induced by the regulation. It would thus be biased upward if more productive firms are also more capitalistic (as suggested by my data). This analysis points out the importance of allowing for factor biased technological differences in my structural

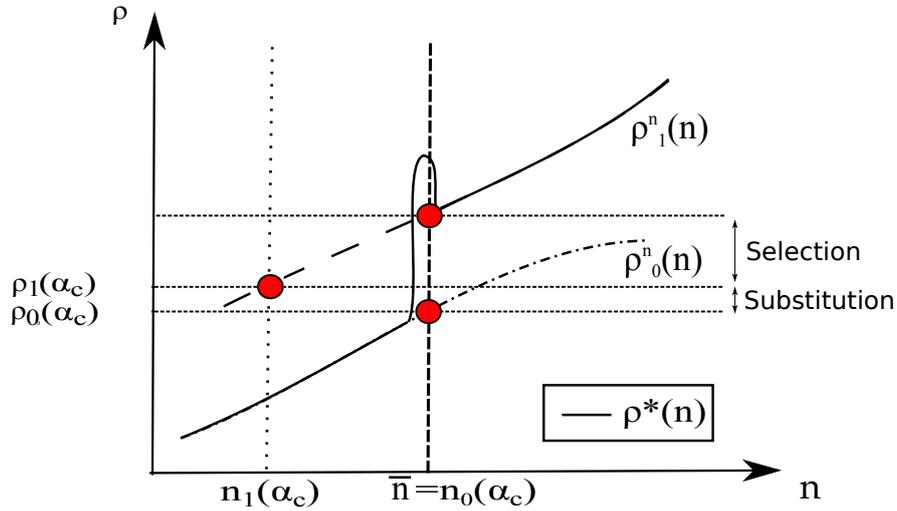
estimation.

I conclude with describing the behavior of the factor cost ratio at the threshold. There is a range of efficiencies $[\alpha_0(\bar{n} - 1), \alpha_u]$ such that firms bunch at $\bar{n} - 1$. Firms with $\alpha \in [\alpha_c, \alpha_u]$ demand $k_b(\alpha)$ units of capital. So capital intensity at $\bar{n} - 1$ is

$$\rho^n(\bar{n} - 1) = \frac{\int_{\alpha_0(\bar{n}-1)}^{\alpha_c} \rho_0(\alpha)g(\alpha)d\alpha + \int_{\alpha_c}^{\alpha_u} \frac{rk_b(\alpha)}{wl}g(\alpha)d\alpha}{\int_{\alpha_0(\bar{n}-1)}^{\alpha_u} g(\alpha)d\alpha}.$$

Bunching firms mitigate the cost of being small by imbalancing their capital-labor ratio, this suggests that the factor cost ratio should be abnormally high in a neighborhood below the threshold.

Figure 3.7: Theoretical Relationship Between Factor Cost Ratio and Size



The main patterns of the relationship between factor cost ratio and size are illustrated in figure 3.7. As in the data, this relationship is broken at the threshold. However, the excessive factor cost ratio below the threshold is not visible in the data. As explained when presenting the motivating evidence, it might be that firms only transiently bunch at the threshold so that they do not have time to adjust their capital stock. In other words, my model fails at replicating the investment behavior of firms below the threshold because it is a zone where firms transit from being not regulated to being regulated. And as a static model, my model is obviously not equipped to predict a transition. Nevertheless, as firms move away from the threshold, they converge back to a long run investment behavior for which my static model is more adapted. In the estimation, we address that limitation of the model by not using the information on capital around the threshold as a source of identification.

Next section describes how I use these predictions to identify the tax equivalent τ and the elasticity of substitution σ .

4 Estimation Strategy

My estimation strategy consists in two steps. In a first step, I estimate the factor demands of firm α_c (i.e. the lowest efficiency firm bunching at 50) in two counterfactual situations: a situation where regulation binds all firms (situation 1) and a situation where regulation binds no firm (situation 0). At this stage, I do not make use of the CES assumption on technology. Then, in a second stage, I use the structural expression of the factor demands (equations (3.4) and (3.6)). The estimated tax equivalent and elasticity of substitution equalize the structural factor demands to factor demands estimated in the first step.

4.1 Estimation of Counterfactual Factor Demands

The first step of the estimation has two substeps which I explain now. First I identify the counterfactual labor demands from the firm size distribution. Intuitively, the excess mass of firms below the threshold will identify by how much firms would shrink their workforce by switching from full regulation to no regulation. Then, I use the counterfactual labor demands and the relationship between the factor cost ratio and size to estimate the counterfactual factor cost ratios.

Estimation of Labour Response In this paragraph, we propose an estimation strategy for $\{n_0(\alpha_c), n_1(\alpha_c)\}$. One of these two labor demands is known by definition since $n_0(\alpha_c) \equiv \bar{n}$, with $\bar{n} = 50$ in the French case. As for the counterpart $n_1(\alpha_c)$, it is identified from the distortion in the size distribution at the threshold. I will first give a graphical intuition of the estimation.

Let us consider figure 3.5 again. In red is the mass of bunching firms. How would this mass of firms distribute itself in a counterfactual situation where regulation applies to all firms? First of all, in that situation there would be no incentive to bunch at 50 since regulation applies anyway, whether a firm is at 49 or at 51. Therefore bunching firms would spread over an interval of sizes. The boundaries of this interval would correspond to the counterfactual labor demand $n_1(\alpha)$ of, respectively, the least efficient and the most efficient bunching firms, namely α_c and α_u . This is precisely what the blue area is, namely the density of firms located between $n_1(\alpha_c)$ and $n_1(\alpha_u)$ in the full regulation counterfactual. It results that *the blue area is equal to the red area*. I can now write this mathematically:

$$\delta^g = \int_{n_1(\alpha_c)}^{n_1(\alpha_u)} g_1^n(n') dn'. \quad (3.16)$$

The left hand side of equation(3.16) is the excess mass at 49, it corresponds to the red area in figure 3.5. The right hand side corresponds to the blue area. Based on this insight and up to a few technical details which I cover below, the way I estimate $n_1(\alpha_c)$ is the following:

1. I estimate the excessive mass at the threshold (the red area)
2. I look for the value of $n_1(\alpha_c)$ which equates the blue area to the excessive mass.

In order to implement this estimation, I must be able to evaluate the blue area. Put differently, I must know the value of function $h(x) \equiv \int_x^{n_1(\alpha_u)} g_1^n(n') dn'$ for any x in a neighborhood

of 50. I identify function $g^n(n)$ by regressing firm size distribution above 50 over a polynomial of n . In order to account for the fact that the counterfactual distribution, $g_1^n(n)$, departs from the actual distribution over $[50, n_1(\alpha_u)]$, I add to the regression size dummies over the interval $[50, n_B]$ where n_B is a size calibrated to be arguably larger than $n_1(\alpha_u)$. Another issue with evaluating $h(x)$ is that the upper bound of the integral, $n_1(\alpha_u)$, is unknown. However, after a few lines of algebra, it is possible to get rid of $n_1(\alpha_u)$ in $h(x)$:

$$\begin{aligned} h(x) &= \int_x^{50} g_1^n(n') dn' + \int_{50}^{n_1(\alpha_u)} g_1^n(n') dn' \\ &= \int_x^{50} g_1^n(n') dn' + \int_{50}^{n_1(\alpha_u)} \underbrace{g_1^n(n') - g^n(n')}_{=0} dn' + \int_{n_1(\alpha_u)}^{n_B} \underbrace{g_1^n(n') - g^n(n')}_{=0} dn' \\ &= \int_x^{50} g_1^n(n') dn' + \int_{50}^{n_B} g_1^n(n') - g^n(n') dn' \end{aligned}$$

$\int_x^{50} g_1^n(n')$ is picked up by extrapolating the estimate of function $g_1^n(n)$ (obtained from size distribution above n_B) to the interval $[x, 50]$. $\int_{50}^{n_B} g_1^n(n', \tau) - g^n(n) dn'$ is estimated as the sum of the dummies over $[50, n_B]$, in the regression fitting the size distribution.

Since I estimate $n_1(\alpha_c)$ as a solution to equation (3.16), I also need to estimate the excessive mass δ^g . The way I proceed is similar to the literature on bunching (e.g. Saez (2010), Chetty et al. (2009)). I regress the observed firm size distribution below the threshold over a polynomial of n and a set of size dummies over an interval $[n_A, 50]$, with n_A calibrated to fall in a region in which firm size distribution is arguably not distorted by the presence of bunching firms at the threshold. These dummies capture the excessive mass.

To summarize and formalize, to estimate $n_1(\alpha_c)$, I first fit the firm size distribution above the threshold and below the threshold by running following regression:

$$g^n(n) = \sum_{p=0}^{P_0} b_{p,0} n^p + \sum_{p=0}^{P_1} b_{p,1} n^p \times \mathbb{1}\{n \geq 50\} + \sum_{i=n_A}^{n_B} d_i \mathbb{1}\{n = i\} + u_g \quad (3.17)$$

With $g^n(n)$ the firm size density, observed in the data, P_0 and P_1 the order of the polynomial used to fit the size distribution respectively below and above 50, and u_g the error term. Coefficients $\{b_{p,0}\}$ are identified from the size distribution below n_A . Coefficients $\{b_{p,1}\}$ are identified from the size distribution above n_B . Dummies $\{d_i\}$ are identified from the distortion in the distribution between n_A and n_B .

From equation (3.17), excess mass is identified by the sum of dummies below the threshold

$$\hat{\delta}^g = \sum_{i=n_A}^{49} \hat{d}_i$$

while $h(x)$ is obtained as the sum over $[x, 49]$ of the distribution fitted above 50, minus the size dummies $\{\hat{d}_i\}_{i=50, \dots, n_B}$:

$$\hat{h}(x) = \sum_{n=x}^{49} \sum_{p=0}^{P_1} \hat{b}_{p,1} n^p - \sum_{i=50}^{n_B} \hat{d}_i$$

Then, the estimator of $\hat{n}_1(\alpha_c)$ is a solution to $\hat{\delta}^g = \hat{h}(x)$:

$$\hat{n}_1(\alpha_c) : \hat{\delta}^g = \hat{h}(\hat{n}_1(\alpha_c)). \quad (3.18)$$

4.2 Estimation of Capital Intensity Response

By definition of ρ_i^n (see equation (3.15)), $\{\rho_0(\alpha), \rho_\tau(\alpha)\}$ verifies

$$\begin{aligned} \rho_0(\alpha_c) &= \rho_0^n(\bar{n}) \\ \rho_1(\alpha_c) &= \rho_1^n(n_1(\alpha_c)) \end{aligned}$$

Since \bar{n} is known and $n_1(\alpha_c)$ has been estimated in a first step, the estimation of counterfactual capital intensities amounts to estimating functions $\rho_i^n(n)$. Analogously to the labour response estimation, I assume that functions ρ_0^n and ρ_1^n can be approximated by finite order polynomials which I estimate by regressing factor cost ratio over size through following regression:

$$\rho^n(n) = \sum_{p=0}^{P'_0} b'_{p,0} n^p + \sum_{p=0}^{P'_1} b'_{p,1} n^p \times \mathbb{1}\{n \geq \bar{n}\} + \sum_{i=n_A}^{n_B} d'_i \mathbb{1}\{n = i\} + u'_n \quad (3.19)$$

With $\rho^n(n)$ the factor cost ratio conditional on size observed in the data and u_ρ the error term. In this equation, the set of dummies only controls for the wedge between actual capital intensity and counterfactual capital intensity in a neighborhood of \bar{n} . It is of no further use in the estimation. $\{\rho_0(\alpha_c), \rho_1(\alpha_c)\}$ is estimated by using the value of counterfactual capital intensities predicted by the polynomials at the sizes of interest, \bar{n} and $\widehat{n_1(\alpha_c)}$:

$$\begin{aligned} \widehat{\rho_0(\alpha_c)} &= \sum_{p=0}^{P'_0} \hat{b}'_{p,0} \bar{n}^p \\ \widehat{\rho_1(\alpha_c)} &= \sum_{p=0}^{P'_1} \hat{b}'_{p,1} \widehat{n_1(\alpha_c)}^p \end{aligned}$$

4.3 Mapping Counterfactual Factor Demands to Structural Parameters

The system of input demands by α_c in the “no regulation” and the “full regulation” situations is:

$$\begin{cases} \bar{n} &= \left(\frac{w}{c_0(\alpha_c)}\right)^{-\sigma} B(\alpha_c)^{\sigma-1} y_0(\alpha_c) \\ n_1(\alpha_c) &= \left(\frac{(1+\tau)w}{c_1(\alpha_c)}\right)^{-\sigma} B(\alpha_c)^{\sigma-1} y_1(\alpha_c) \\ \rho_0(\alpha_c) &= \left(\frac{r}{w}\right)^{1-\sigma} \left(\frac{A}{B}(\alpha_c)\right)^{\sigma-1} \\ \rho_1(\alpha_c) &= (1+\tau)^\sigma \left(\frac{r}{w}\right)^{1-\sigma} \left(\frac{A}{B}(\alpha_c)\right)^{\sigma-1}. \end{cases}$$

This system can be concentrated over \bar{n} , $n_1(\alpha_c)$, $\rho_0(\alpha_c)$, $\rho_1(\alpha_c)$, τ , σ and ε :

$$\begin{cases} \frac{n_1(\alpha_c)}{\bar{n}} &= (1+\tau)^{-\sigma} \left(\frac{(1+\tau)^{1-\sigma} + \rho_0(\alpha_c)}{1+\rho_0(\alpha_c)}\right)^{\varepsilon-\sigma} \\ \frac{\rho_1(\alpha_c)}{\rho_0(\alpha_c)} &= (1+\tau)^\sigma. \end{cases} \quad (3.20)$$

Factor demands were estimated in previous steps. So I can plug their estimates in system (3.20), calibrate ε and then estimate $\{\hat{\tau}, \hat{\sigma}\}$ as a solution to the system. I estimate standard errors via bootstrap. Namely, I draw with replacements 500 samples of firms of the size of my dataset and I run the estimation for each of them.

Next section presents the results.

5 Results

There is a handful a parameters that I need to calibrate to run the estimation. First parameters are P_0 and P_1 . I fit the distribution below and above the threshold with polynomials of order one and two respectively. In practice, choosing too large numbers for P_0 and P_1 generates an over fitting of the distribution and results into a bad estimation of the distortion over $[n_A, n_B]$. Another parameter to calibrate is n_A , namely the lower bound to the interval over which distortions around the threshold are estimated. I choose $n_A = 40$ because it is the size where the bulge of productivities starts in figure 3.2. I set $n_B = 56$ because it the expected destination size of firms which grow at the threshold (see figure 3.6). The idea is that the distortion in the size distribution just above 50 comes from firms not wanting to be there. I use the size dynamics to get a sense of where firms ‘land’ when they jump over this unprofitable range of sizes.

In order to visually assess the quality of the fit, I plot the prediction of the regression of size density (model (3.17)) in figure 3.1. The sum of size dummies from 40 to 50, which identifies the mass of bunching firms, represents 3% of firms in my sample. This is a small number, especially if one considers the fact that I do not observe firms below 20 employees. However, this small mass of bunching firms does not imply a small aggregate impact as all firms above 50 are also constrained by the regulation.

I estimate $n_1(\alpha_c)$ by equalizing the sum of size dummies with the mass of firms between $n_1(\alpha_c)$ and n_B . I illustrate the solution to this problem in figure 3.1. I find $n_1(\hat{\alpha}_c) = 46$ meaning that the bunching firm with the lowest productivity would be 4 employees smaller if regulation were to bind all firms, all else equal.

In order to estimate counterfactual factor cost ratio, I estimate model (3.19) and then take the prediction at $n = 50$ and $n = n_1(\hat{\alpha}_c)$. This procedure is illustrated in figure 3.2.

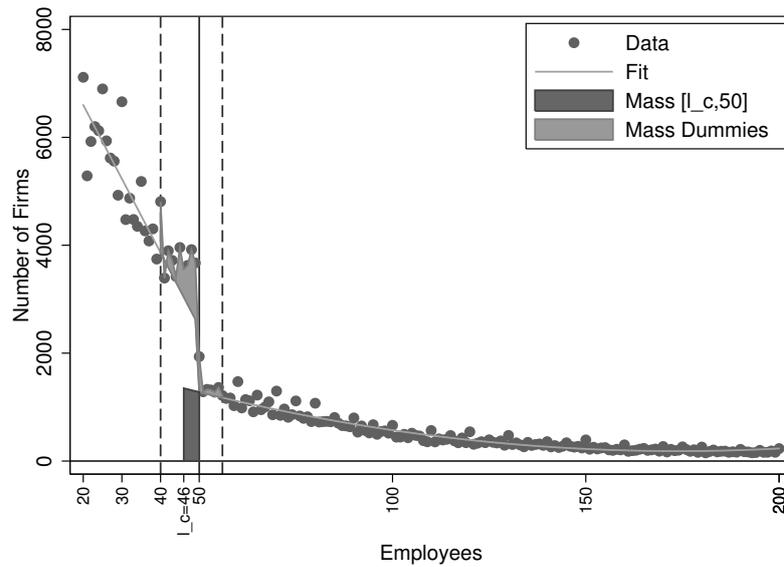
Table 3.1: Parameter Estimates

<i>Parameter</i>	<i>Estimate</i>	<i>95% Confidence Interval</i>	
		Lower Bound	Upper Bound
response # employees ($50 - n_1(\alpha_c)$):	4.1	3.4	4.8
response factor cost ratio ($\ln \frac{\rho_1(\alpha_c)}{\rho_0(\alpha_c)}$):	0.05	0.01	0.08
excessive mass at 50 (# Firms)	6.67e+03	5.75e+03	7.59e+03
excessive mass at 50 (% of firms)	3.05	2.63	3.47
τ	0.077	-0.001	0.135
σ	0.58	0.07	4.80

Notes: Estimates of confidence interval are obtained with 500 bootstrap replications.

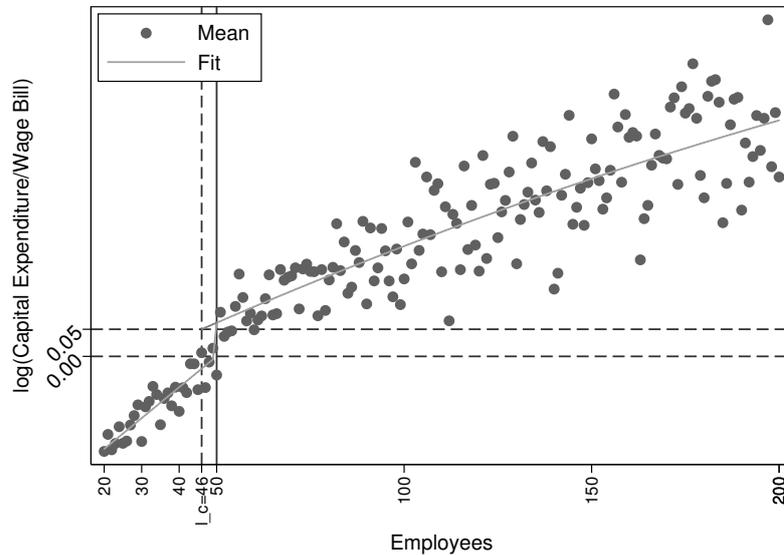
The confidence interval is at 95% and is estimated from the basic percentile method.

Figure 3.1: Fitting Firm Size Distribution I



Note: Area 'Mass over $[l_c; 50]$ ' is the area between the x-axis and the extrapolation of the distribution fit above 50 over the range $[46, 50]$. Area 'Mass over $[l_c; 50]$ ' and area 'Mass Dummies' are equal.

Figure 3.2: Fitting Capital Intensity Versus Size



Note: Capital Expenditure=0.1*Fixed Assets. Fit is the prediction of a regression of log Capital Expenditure on Employees, a dummy ($size \geq 50$), a collection of interaction terms $Employees \times dummy(size \geq 50)$ and $Employees^2 \times dummy(size \geq 50)$, a set of dummies for employees $\in [40, 56]$, a set of 2-digit industry dummies and a constant term.

With the counterfactual factor demands estimated, I now invert system (3.20) over σ and τ . To this end, I need to calibrate ε . The way I proceed is I implement Levinsohn & Petrin (2003) semi-parametric estimation of a cobb-douglas production function with capital and

labor. Put simply, the estimation consists in regressing value added on capital and labor and using materials as a control for (unobserved) total factor productivity. I use the sum of obtained coefficients on labor and capital as an estimate of the curvature of the revenue function defined in (3.7). I get $\frac{\varepsilon-1}{\varepsilon} = 0.837$.

Finally, I find a tax equivalent of 8% and an elasticity of substitution of 0.6. Table 3.1 summarizes the results. The value of τ is slightly larger than in comparable studies.¹⁰ However, in present paper, the cost of labor is only one part of the total cost of production. In order to make our result more comparable to the relevant benchmark, I compute the increase in percents in the total unit cost of production of a firm at the threshold:

$$\ln \frac{\widehat{C_1(\alpha_c)}}{C_0(\alpha_c)} = \frac{1}{1 + \hat{\rho}_0(\alpha_c)} \tau = 0.035$$

I get that regulation increases the total cost of a firm at the threshold by 3.5% which is smaller (although not significantly) than existing estimates. This is consistent with the idea that controlling for the substitution behavior of firms leads to infer a smaller cost of the regulation from a given dataset. The reason for that is simply that substitution allows firms to bunch at a lower cost so that a smaller regulation cost is necessary to rationalize the excessive mass at the threshold.

As to the elasticity of substitution, I get a value close to recent studies (Chirinko et al. (2004) and Raval (2011) estimate long run capital-labor elasticity of substitution around 0.5). I obtain a standard error too large, however, to reject the null hypothesis that σ is larger or equal to one.

In order to put these results in perspective, next section makes us of our estimates to implements policy experiments.

6 Counterfactual Analysis

In this section, I investigate the impact of variations in τ over the equilibrium. Since regulation potentially affects relative price of labor to capital, I will be especially interested in the factor income redistribution effects of modifying the regulation. To quantify these general equilibrium effects, I first need to specify the structure of the economy.

6.1 The Structure of the General Equilibrium

My general equilibrium model builds upon the closed version of Melitz (2003) model and embed the partial equilibrium model of the firm developed in section 3. More specifically, I generalize the assumption on technology made by Melitz (2003) by allowing firms to produce out of both capital and labor through a CES function. The economy is at a stationary equilibrium and is populated with a unitary mass of individuals. Each individual is endowed with N_i units of labour and K_i units of capital, with the sum of individual endowments summing to N and K respectively Labor is the numeraire ($w = 1$).

10. Garicano et al. (2013) and Gourio & Roys (2012) get estimates of τ between 4% and 7%.

Demand Consumers have CES preferences over a continuum of varieties of a differentiated good. Given the homotheticity of these preference, I rely on a representative consumer approach even if individual consumers have different budget constraints. The demand of the representative consumer for a variety ω is

$$y(\omega) = p(\omega)^{-\varepsilon} Z.$$

Residual demand functions therefore have the same expression as in section 3 although Z is now made endogenous and verifies

$$Z = P^{1-\varepsilon} R,$$

with R the total expenditure of the representative consumer over the differentiated good and P is the CES price index over the varieties. Letting M be the mass of incumbent firms in the economy, P verifies:

$$P = \left(M \int p(\alpha)^{1-\varepsilon} g(\alpha) d\alpha \right)^{\frac{1}{1-\varepsilon}}. \quad (3.21)$$

Production, entry and exit The technology and the regulatory context are those assumed in section 3. As a consequence the optimal price, output, factor demand and operating profit of incumbent firms extend to present general equilibrium setting.

There is a large pool of ex-ante identical prospective entrants in the industry. To enter, firms pay a sunk cost f_e , labeled in capital and draw an efficiency α in a distribution $\lambda(\alpha)$. Drawn productivity remains constant over time. Upon entry, firms may decide to exit immediately. Given the fact that α is constant, that the equilibrium is stationary, and that the value of exiting the market is zero, endogenous exit will happen if and only if the current profit of a firm is negative. To allow for this possibility, I assume that firms need to pay a fixed cost per period f , labeled in capital, to stay in the market. Since operating profits are strictly decreasing with α , there exists a cutoff firm α_e , indifferent to exit or stay in the market:

$$\pi(\alpha_e) = rf, \quad (3.22)$$

and the probability of staying in the market for a new entrant is $1 - \Lambda(\alpha)$, with Λ the CDF associated to λ . The stationary distribution of firms, $g(\alpha)$ is therefore linked to entry distribution as follows

$$g(\alpha) = \begin{cases} \frac{\lambda(\alpha)}{1-\Lambda(\alpha)} & \text{if } \alpha \geq \alpha_e \\ 0 & \text{otherwise.} \end{cases}$$

The value of drawing a productivity to a prospective entrant, v_e , is equal to the probability of staying in the market multiplied by the expectation of the infinite sum of future profits, discounted by the probability of the death shock: $v_e = (1 - \Lambda(\alpha)) \int \sum_{t=0}^{\infty} (1 - \delta)^t [\pi(\alpha) - rf] g(\alpha) d\alpha$. To decide whether to enter the market or not, a prospective entrant compares v_e to the cost of entry rf_e . Free entry in the economy equalizes v_e to rf_e .

$$\frac{(1 - \Lambda(\alpha))}{\delta} \int [\pi(\alpha) - rf] g(\alpha) d\alpha = rf_e. \quad (3.23)$$

Equilibrium At the equilibrium, factor markets clear:

$$\begin{aligned} K &= M \int_{\alpha_c} [k(\alpha) + rf] g(\alpha) d\alpha + M_e rf_e \\ N &= M \int_{\alpha_c} n(\alpha) g(\alpha) d\alpha, \end{aligned} \quad (3.24)$$

with M_e the masse of new entrant firms at every period.

On top of its revenues from factors, the representative consumer owns firms and so gets dividends, and she also gets redistributed the tax on labor. The representative consumer exhaust its income by consuming the differentiated good and creating new firms. Therefore her budget constraint is:

$$rK + L + M \int_{\alpha_c} [\pi(\alpha) - rf] g(\alpha) + \tau \int_{\alpha_u} n(\alpha) g(\alpha) d\alpha = M_e rf_e + R. \quad (3.25)$$

Finally, stationarity imposes the flow of exiting firms to equal the flow of new firms:

$$(1 - \Lambda(\alpha_c)) M_e = \delta M, \quad (3.26)$$

with M_e the mass of entering firms at each period.

An equilibrium is characterized:

- by a set of cutoffs $\{\alpha_e, \alpha_u, \}$ such that the zero-profit condition (3.22) and indifference condition on bunching (3.9) are verified,
- by a set of individual input demands $\{k(\alpha), n(\alpha)\}$ such that profit is maximized, i.e verifying (3.10) and (3.11),
- by an aggregate expenditure consistent with the budget constraint (3.25),
- by a price index P consistent with its definition (3.21)
- by a mass of incumbent M such that the free entry condition (3.23) is verified,

- by a flow of new entrants M_e such that the flow balance condition (3.26) applies
- by an interest rate r insure that capital market clears (equation (3.24)),

6.2 Results

This subsection is preliminary. In table 3.1 I describe the calibration which I intend to use.

Table 3.1: Calibration

Parameter	method	source/matched moment
$A(\alpha) = \alpha$	<i>normalized</i>	
$B(\alpha) = \alpha^b$	<i>calibrated</i>	capital ratio difference below/above 50
$\lambda(\alpha) = \theta\alpha^{-\theta}$	<i>calibrated</i>	share workers below 50
f	<i>normalized</i>	
f_e	<i>calibrated</i>	number of firms in the economy
δ	<i>normalized</i>	Burstein & Melitz (2011)
σ	<i>estimated</i>	
ε	<i>estimated</i>	
τ	<i>estimated</i>	
K	<i>calibrated</i>	aggregate labor share=0.6
L	<i>calibrated</i>	number of workers in the economy

My intention is to compute the effect of **de-regulation** on the real wage $1/P$, the real interest rate r/P , the relative share of capital rK/wL and aggregate welfare R/P . In order to isolate the role of capital intensity heterogeneity, I intend to run this experiment in two cases: one where all firms have the same factor share and one where the relative capital intensity between and above the threshold matches the data.

7 Conclusion

In this paper, I have analyzed the impact of size dependent regulation on factor income distribution. Drawing upon a census of French firms, I find suggestive evidence of the opposite effects. First, the regulation is costly - as revealed by the excess mass of firms at the threshold - and regulated firms are relatively more capital intensive. This suggests that the regulation, by imposing a cost on large (capital-intensive) firms tend to make small firms expand so that the whole economy is less capital intensive. Second, the pattern on capital labor ratio around the threshold suggest that firms substitute capital for ratio in order to mitigate the cost of the regulation. This second effect tends to make the economy more capital intensive.

To find out which of the between and within firm substitution effect dominates, I write down a general equilibrium model which I intend to calibrate. Among the parameters to be calibrates, two important ones - the implicit cost of the regulation and the capital labor elasticity of substitution - have been structurally estimated.

Appendix

A Appendix Tables

Table A.1: Descriptive Statistics

Year	# Firms	# Employees	Wage Bill	Fixed sets	As-	Value Added
1995	2.14e+04	2.81e+06	6.69e+10	1.95e+11		1.60e+11
1996	2.10e+04	2.77e+06	7.09e+10	2.28e+11		1.59e+11
1997	2.12e+04	2.75e+06	7.25e+10	2.42e+11		1.67e+11
1998	2.08e+04	2.73e+06	7.37e+10	2.42e+11		1.71e+11
1999	2.07e+04	2.74e+06	7.55e+10	2.54e+11		1.79e+11
2000	2.05e+04	2.77e+06	7.83e+10	2.66e+11		1.88e+11
2001	2.06e+04	2.80e+06	8.08e+10	2.81e+11		1.81e+11
2002	2.03e+04	2.74e+06	8.17e+10	2.90e+11		1.81e+11
2003	1.98e+04	2.66e+06	8.16e+10	3.01e+11		1.81e+11
2004	1.92e+04	2.58e+06	8.09e+10	3.05e+11		1.84e+11
2005	1.88e+04	2.52e+06	8.16e+10	3.17e+11		1.87e+11
2006	1.80e+04	2.47e+06	8.31e+10	3.25e+11		1.87e+11
2007	1.70e+04	2.32e+06	8.07e+10	3.22e+11		1.84e+11

Table A.2: Descriptive Statistics

Variable	Mean	q1	q25	q50	q75	q99
Above 50 Dummy	0.42	0	0	0	1	1
Fixed Assets	1.4e+07	0.00	5.2e+05	1.4e+06	4.3e+06	1.8e+08
Employees	133.7	20	30	44	91	1416
Factor Cost Ratio	0.21	0.00	0.06	0.13	0.24	1.15
Wage Bill	3.9e+06	3.2e+05	6.8e+05	1.1e+06	2.2e+06	4.7e+07
Value Added	8.9e+06	3.3e+05	1.2e+06	2.0e+06	4.3e+06	1.1e+08

B Proof Proposition 4

We first demonstrate that $\Delta(\alpha)$ has a unique root over $[\alpha_c, +\infty)$. Let $\tilde{\alpha}$ be such that

$$n_1(\tilde{\alpha}) = \bar{n}$$

$\frac{\partial n_1(\alpha)}{\partial \alpha} > 0$ guarantees the uniqueness of $\tilde{\alpha}$. We will proceed by showing that:

- (1) $\Delta(\alpha_c) < 0$
- (2) $\Delta'(\alpha) \leq 0$ over $[\alpha_c, \tilde{\alpha}]$
- (3) $\Delta'(\alpha) > 0$ over $[\tilde{\alpha}, +\infty]$

(1) and (2) guarantee that $\Delta(\alpha)$ has no root over $[\alpha_c, \tilde{\alpha}]$. (1), (2) and (3) guarantee that $\Delta(\alpha)$ has a unique root over $[\tilde{\alpha}, +\infty]$.

Regulation does not bind to firms α_c so it is clear that $\Delta(\alpha_c) < 0$. From the envelope theorem we get the general expression of the derivative of $\Delta(\alpha)$:

$$\Delta'(\alpha) = \frac{\partial R}{\partial \alpha}(k_1(\alpha), n_1(\alpha), \alpha) - \frac{\partial R}{\partial \alpha}(k_b(\alpha), \bar{n}, \alpha) \quad (3.27)$$

With $k_b(\alpha)$ the optimal capital demand of a bunching firm and $k_1(\alpha)$ the optimal capital of a regulated firm:

$$\begin{aligned} k_b(\alpha) : \frac{\partial R}{\partial k}(k_b(\alpha), \bar{n}; \alpha) &= r \\ k_1(\alpha) : \frac{\partial R}{\partial k}(k_1(\alpha), n_1(\alpha); \alpha) &= r \end{aligned}$$

The definition of $\tilde{\alpha}$ and the fact that $\frac{\partial n_1(\alpha)}{\partial \alpha} > 0$ imply that $n_1(\alpha) \leq \bar{n}$ over $[\alpha_c, \tilde{\alpha}]$. $\sigma > 0$ implies that inputs are (imperfect) complements and so the same ranking holds for capital: $\forall \alpha \in [\alpha_c, \tilde{\alpha}], k_1(\alpha) \leq k_b(\alpha)$. Since $\frac{\partial R}{\partial \alpha}(k, n; \alpha)$ is increasing in (k, n) , it follows from equation (3.27) that $\Delta'(\alpha) \leq 0$ over $[\alpha_c, \tilde{\alpha}]$. Applying similar reasoning over $[\tilde{\alpha}, \infty]$, one establishes that $\forall \alpha > \tilde{\alpha}, \Delta'(\alpha) > 0$. $\Delta(\alpha)$ has therefore a unique root $\alpha_u > \tilde{\alpha}$.

Moreover, α_u verifies $\Delta'(\alpha_u) > 0$. Finally, since $n_1(\tilde{\alpha}) = \bar{n}$ by definition, and $n_1(\alpha)$ is increasing in α , $n_1(\alpha_u) > \bar{n}$.

□

Chapter 4

General Conclusion

Through this dissertation, I wanted to shed new light on three questions: the interdependence of firms' sales across markets, the firm-level response of the quality of exported goods to low-cost competition and the macroeconomic consequences of size-dependent labor regulation. Throughout my work, I have made use of firm-level data as well as of simple theory to guide my empirical strategy.

I am optimistic that my work will contribute to three topics in international economics and macroeconomics. First, chapter 1 adds to the literature on quality and trade as we propose a new solution to estimate product quality at the firm-level. This measure is appealing as it requires widely and increasingly available information of firms. Moreover, our measure of quality may vary in time and space within the firm. Thus, it can be used to tackle a vast class of questions regarding firm-level quality and can be applied to any firm-level customs dataset.

Second, chapter 1, 2 and 3 enhance our knowledge of the way firms respond to external constraints. In chapter 1, we show that firms shift up the quality of their exports when low-cost competition gets fiercer. To the extent that this quality upgrading spills over domestic cost of living and wages, our result may change the way we understand the impact of low-cost competition.

In chapter 2, we find that Champagne exporters spatially reallocate their sales when demand falls in a destination market. This is of importance to understand the way exporting firms shape the international diffusion of shocks.

In chapter 3, my analysis of firms' response to size-dependent regulation reveals that some firms bunch at regulation thresholds in order not to face additional legal constraints. This behavior is consistent with the fact that regulation is costly. I also find evidence that regulated firms substitute capital for labor in order to mitigate the cost of the regulation. These strategies implemented by firms shape the macroeconomic implications of the regulation.

Third, chapter 3 shows that firm heterogeneity has implications for the labor share of income. The aggregate labor share is an important object in macroeconomics because it is often used to discipline theoretical models. I show that since firms are heterogeneous in labor intensity within industries, any policy reallocating market shares across firms has potential implications for the labor share. An example of such policies is size-dependent regulations. By constraining more large (capital intensive) large firms, such policies are likely to redistribute national income toward labor.

As I am concluding my thesis, I am left with more questions than answers. Part of these questions fall into macroeconomics, others are more related to international trade, but all of them regard the way firm behavior shapes aggregate outcomes.

One question I would like to investigate is the impact of trade liberalization on factor income distribution. Since the seminal contribution of Melitz (2003), we know that trade liberalization increases the weight of large firms in the economy. Since large firms are more capital intensive, as I document in chapter 3, trade liberalization could have contributed to the fall in the labor share observed in the world during three past decades. In a future project, I would like to evaluate the share of this fall which can be attributed to the reallocation induced by trade liberalization.

Another project I intend to carry out is on the pricing behavior of exporting firms. Workhorse trade models have predictions on mark-ups charged by firms. These predictions have different implications for the gains from trade and it is therefore of importance to empirically discriminate between them. For instance, Melitz (2003) predicts that firms charge identical mark-ups at home and abroad while Melitz & Ottaviano (2008) predicts that firms charge lower mark-ups in distant markets. In the meantime, Melitz & Ottaviano (2008) that a gain from trade liberalization is to increase the intensity of competition while this gain is absent from Melitz (2003). By estimating the price elasticity of demand through the demand estimation developed in chapter 1, I will be able to look at the way price elasticity varies within firms across destination markets. I will be for instance able to test if firms face more elastic demand on distant markets as predicted by Melitz & Ottaviano (2008).

Bibliography

- Ahn, J. & McQuoid, A. F. (2012), ‘Capacity constrained exporters: micro evidence and macro implications’.
- Amiti, M. & Cameron, L. (2012), ‘Trade liberalization and the wage skill premium: Evidence from indonesia’, *Journal of International Economics* **87**(2), 277–287.
- Amiti, M., Itskhoki, O. & Konings, J. (forthcoming), ‘Importers, exporters, and exchange rate disconnect’, *American Economic Review* .
- Anderson, S. P., De Palma, A. & Thisse, J.-F. (1987), ‘The ces is a discrete choice model?’, *Economics Letters* **24**(2), 139–140.
- Autor, D. H., Dorn, D. & Hanson, G. H. (2013), ‘The china syndrome: Local labor market effects of import competition in the united states.’, *American Economic Review* .
- Baldwin, R. & Harrigan, J. (2011), ‘Zeros, quality, and space: Trade theory and trade evidence’, *American Economic Journal: Microeconomics* **3**(2), 60–88.
- Bastos, P., Silva, J. & Verhoogen, E. (2012), ‘Export destinations and input prices: Evidence from portugal’.
- Bauernschuster, S. (2009), ‘Relaxed dismissal protection: Effects on the hiring and firing behaviour of small firms’, *Jena Economic Research Papers* **2009**, 015.
- Bentolila, S. & Bertola, G. (1990), ‘Firing costs and labour demand: how bad is eurosclerosis?’, *The Review of Economic Studies* **57**(3), 381–402.
- Berman, N., Berthou, A. & Héricourt, J. (2011), Export dynamics and sales at home, Technical report, Graduate Institute of International and Development Studies Working Paper.
- Berman, N., Martin, P. & Mayer, T. (2012), ‘How do different exporters react to exchange rate changes?’, *The Quarterly Journal of Economics* **127**(1), 437–492.
- Bernard, A. B. & Jensen, J. B. (2004), ‘Why some firms export’, *Review of Economics and Statistics* **86**(2), 561–569.
- Bernard, A. B., Jensen, J. B. & Schott, P. K. (2009), Importers, exporters and multinationals: a portrait of firms in the us that trade goods, *in* ‘Producer Dynamics: New Evidence from Micro Data’, University of Chicago Press, pp. 513–552.
- Berry, S., Levinsohn, J. & Pakes, A. (1995), ‘Automobile prices in market equilibrium’, *Econometrica: Journal of the Econometric Society* pp. 841–890.
- Berthou, A. & Fontagne, L. (2008), The euro and the intensive and extensive margins of trade: Evidence from french firm level data, Working Papers 2008-06, CEPII research center.

- Bloom, N., Draca, M. & Van Reenen, J. (2011), Trade induced technical change? the impact of chinese imports on innovation, it and productivity, Technical report, National Bureau of Economic Research.
- Boeri, T. & Jimeno, J. (2005), 'The effects of employment protection: Learning from variable enforcement', *European Economic Review* **49**(8), 2057–2077.
- Brambilla, I., Lederman, D. & Porto, G. (2012), 'Exports, export destinations, and skills', *The American Economic Review* **102**(7), 3406–38.
- Broda, C. & Weinstein, D. E. (2006), 'Globalization and the gains from variety', *The Quarterly Journal of Economics* pp. 541–585.
- Broda, C. & Weinstein, D. E. (2010), 'Product creation and destruction: Evidence and price implications', *The American economic review* **100**(3), 691–723.
- Burstein, A. & Melitz, M. (2011), Trade liberalization and firm dynamics, Technical report, National Bureau of Economic Research.
- Cancé, R. (2009), 'L'Appareil exportateur français: une réalité plurielle', *Trésor-Eco, Direction générale du Trésor et de la Politique Economique*.
- Cebeci, T., Fernandes, A. M., Freund, C. L. & Pierola, M. D. (2012), 'Exporter dynamics database', *World Bank Policy Research Working Paper* (6229).
- Ceci-Renaud, N. & Chevalier, P. (2010), 'How much do legal size-contingent thresholds impact the size distribution of french firms?', *Documents de Travail de la DESE-Working Papers of the DESE*.
- Chetty, R. (2012), 'Bounds on elasticities with optimization frictions: A synthesis of micro and macro evidence on labor supply', *Econometrica* **80**(3), 969–1018.
- Chetty, R., Friedman, J., Olsen, T. & Pistaferri, L. (2009), Adjustment costs, firm responses, and micro vs. macro labor supply elasticities: Evidence from danish tax records, Technical report, National Bureau of Economic Research.
- Chirinko, R., Fazzari, S. & Meyer, A. (2004), 'That elusive elasticity: a long-panel approach to estimating the capital-labor substitution elasticity'.
- Crozet, M., Head, K. & Mayer, T. (2012), 'Quality sorting and trade: Firm-level evidence for french wine', *The Review of Economic Studies* **79**(2), 609–644.
- Declerck, F. & Cloutier, M. L. (2002), *The Champagne Wine Industry: An Economic Dynamic Model of Production and Consumption*, Citeseer.
- DeLong, J. B. & Magin, K. (2006), A short note on the size of the dot-com bubble, Technical report, National Bureau of Economic Research.
- Dixit, A. K. & Pindyck, R. S. (1996), 'Investment under uncertainty : Avinash k dixit and robert s pindyck princeton university press, princeton, nj, 1994, xiv + 468 pp (hardcover), isbn 0-691-03410-9', *Resources Policy* **22**(3), 218–218.
- Dubé, J.-P. (2004), 'Multiple discreteness and product differentiation: Demand for carbonated soft drinks', *Marketing Science* **23**(1), 66–81.
- Eaton, J. & Kortum, S. (2002), 'Technology, geography, and trade', *Econometrica* **70**(5), 1741–1779.

- Eaton, J., Kortum, S. & Kramarz, F. (2011), ‘An anatomy of international trade: Evidence from french firms’, *Econometrica* **79**(5), 1453–1498.
- Eaton, J., Kortum, S., Kramarz, F. & Sampognaro, R. (2011), ‘Dissecting the french export wage premium’, *Penn State University mimeo* .
- Evans, C. & Harrigan, J. (2005), ‘Distance, time, and specialization: lean retailing in general equilibrium’, *American Economic Review* pp. 292–313.
- Faber, B. (2012), ‘Trade liberalization, the price of quality, and inequality: Evidence from mexican store prices’, *London School of Economics* .
- Faljaoui, A. (2009), ‘Les 3 raisons d’un champagne à prix bradé’.
- Feenstra, R. C. (1994), ‘New product varieties and the measurement of international prices’, *The American Economic Review* pp. 157–177.
- Feenstra, R. C. & Romalis, J. (2012), International prices and endogenous quality, Technical report, National Bureau of Economic Research.
- Fontagné, L., Gaulier, G. et al. (2008), *Performances à l’exportation de la France et de l’Allemagne*, la Documentation française.
- Forslid, R. & Okubo, T. (2011), *Are capital intensive firms the biggest exporters?*, Centre for Economic Policy Research.
- Foster, L., Haltiwanger, J. & Syverson, C. (2008), ‘Reallocation, firm turnover, and efficiency: Selection on productivity or profitability?’, *American Economic Review* **98**(1), 394–425.
- Gabaix, X. (2011), ‘The granular origins of aggregate fluctuations’, *Econometrica* **79**(3), 733–772.
- Garicano, L., Lelarge, C. & Van Reenen, J. (2013), Firm size distortions and the productivity distribution: Evidence from france, Technical report, National Bureau of Economic Research.
- Gaulier, G. & Zignago, S. (2010), ‘Baci: International trade database at the product-level (the 1994-2007 version)’.
- Gervais, A. (2011), ‘Product quality and firm heterogeneity in international trade’, *University of Notre Dame Unpublished Manuscript* .
- Gourio, F. & Roys, N. (2012), Size-dependent regulations, firm size distribution, and reallocation, Technical report, National Bureau of Economic Research.
- Grossman, G. M. & Helpman, E. (1991), ‘Quality ladders in the theory of growth’, *The Review of Economic Studies* **58**(1), 43–61.
- Gruben, W. C., Koo, J. & Millis, E. (2002), *How much does international trade affect business cycle synchronization?*, Vol. 2, Federal Reserve Bank of Dallas.
- Guimaraes, P. & Portugal, P. (2010), ‘A simple feasible procedure to fit models with high-dimensional fixed effects’, *Stata Journal* **10**(4), 628.
- Guner, N., Ventura, G. & Xu, Y. (2008), ‘Macroeconomic implications of size-dependent policies’, *Review of Economic Dynamics* **11**(4), 721–744.

- Hall, R., Feldstein, M., Bernanke, B., Frankel, J., Gordon, R. & Zarnowitz, V. (2001), 'The business-cycle peak of march 2001', *Business Cycle Dating Committee, National Bureau of Economic Research* .
- Hallak, J. C. (2006), 'Product quality and the direction of trade', *Journal of International Economics* **68**(1), 238–265.
- Hallak, J. C. & Schott, P. K. (2011), 'Estimating cross-country differences in product quality*', *The Quarterly journal of economics* **126**(1), 417–474.
- Hallak, J. C. & Sivadasan, J. (2011), 'Firms' exporting behavior under quality constraints'.
- Handbury, J. (2012), 'Are poor cities cheap for everyone? non-homotheticity and the cost of living across us cities'.
- Hausman, J. A. (1996), Valuation of new goods under perfect and imperfect competition, in 'The economics of new goods', University of Chicago Press, pp. 207–248.
- Helpman, E. (1981), 'International trade in the presence of product differentiation, economies of scale and monopolistic competition: a chamberlin-heckscher-ohlin approach', *Journal of international economics* **11**(3), 305–340.
- Helpman, E., Itskhoki, O. & Redding, S. (2010), 'Inequality and unemployment in a global economy', *Econometrica* **78**(4), 1239–1283.
- Hsieh, C. & Klenow, P. (2009), 'Misallocation and manufacturing tfp in china and india', *The Quarterly Journal of Economics* **124**(4), 1403.
- Imbs, J. & Méjean, I. (2008), 'Elasticity optimism'.
- Johnson, R. C. (2012), 'Trade and prices with heterogeneous firms', *Journal of International Economics* **86**(1), 43–56.
- Juhlin, R. (2008), *Champagne Guide*, Richard Juhlin Publishing AB.
- Khandelwal, A. (2010), 'The long and short (of) quality ladders', *The Review of Economic Studies* **77**(4), 1450–1476.
- Khandelwal, A., Schott, P. & Wei, S. (forthcoming), 'Trade liberalization and embedded institutional reform: Evidence from chinese exporters', *American Economic Review* .
- Krugman, P. (1980), 'Scale economies, product differentiation, and the pattern of trade', *American Economic Review* **70**(5), 950–959.
- Kugler, A. & Pica, G. (2008), 'Effects of employment protection on worker and job flows: Evidence from the 1990 italian reform', *Labour Economics* **15**(1), 78–95.
- Kugler, M. & Verhoogen, E. (2012), 'Prices, plant size, and product quality', *The Review of Economic Studies* **79**(1), 307–339.
- Leland, H. E. (1972), 'Theory of the firm facing uncertain demand', *The American Economic Review* **62**(3), 278–291.
- Levinsohn, J. & Petrin, A. (2003), 'Estimating production functions using inputs to control for unobservables', *The Review of Economic Studies* **70**(2), 317–341.
- Liu, Y. (2012), Capital adjustment costs: Implications for domestic and export sales dynamics, Technical report, mimeo.

- Luttmer, E. G. (2007), 'Selection, growth, and the size distribution of firms', *The Quarterly Journal of Economics* **122**(3), 1103–1144.
- Manova, K. & Zhang, Z. (2012), 'Export prices across firms and destinations', *The Quarterly Journal of Economics* **127**(1), 379–436.
- Martin, J. & Méjean, I. (2011), 'Low-wage countries' competition, reallocation across firms and the quality content of exports'.
- Mayer, T. & Ottaviano, G. I. (2008), 'The happy few: The internationalisation of european firms', *Intereconomics* **43**(3), 135–148.
- Melitz, M. J. (2003), 'The impact of trade on intra-industry reallocations and aggregate industry productivity', *Econometrica* **71**(6), 1695–1725.
- Melitz, M. J. & Ottaviano, G. I. (2008), 'Market size, trade, and productivity', *The review of economic studies* **75**(1), 295–316.
- Mula, J., Poler, R., Garcia-Sabater, J. & Lario, F. (2006), 'Models for production planning under uncertainty: A review', *International Journal of Production Economics* **103**(1), 271–285.
- Nevo, A. (2000), 'Mergers with differentiated products: The case of the ready-to-eat cereal industry', *The RAND Journal of Economics* pp. 395–421.
- Nguyen, D. X. & Schaur, G. (2010), Cost linkages transmit volatility across markets, Technical report, EPRU Working Paper Series.
- Pavcnik, N. (2002), 'Trade liberalization, exit, and productivity improvements: Evidence from chilean plants', *The Review of Economic Studies* **69**(1), 245–276.
- Pierce, J. R. & Schott, P. K. (2012), 'Concording us harmonized system codes over time', *Journal of Official Statistics* **28**(1), 53–68.
- Raval, D. (2011), Estimation of a ces production function with factor augmenting technology, Technical report, mimeo.
- Restuccia, D. & Rogerson, R. (2008), 'Policy distortions and aggregate productivity with heterogeneous establishments', *Review of Economic Dynamics* **11**(4), 707–720.
- Roberts, M., Xu, D., Fan, X. & Zhang, S. (2012), 'A structural model of demand, cost, and export market selection for chinese footwear producers', *NBER* .
- Saez, E. (2010), 'Do taxpayers bunch at kink points?', *American Economic Journal: Economic Policy* **2**(3), 180–212.
- Schivardi, F. & Torrini, R. (2004), 'Threshold effects and firm size: the case of firing costs'.
- Schott, P. K. (2004), 'Across-product versus within-product specialization in international trade', *The Quarterly Journal of Economics* **119**(2), 647–678.
- Silva, J. S. & Tenreyro, S. (2006), 'The log of gravity', *The Review of Economics and statistics* **88**(4), 641–658.
- Simon, H. A. & Bonini, C. P. (1958), 'The size distribution of business firms', *The American Economic Review* **48**(4), 607–617.
- Sutton, J. (2001), *Technology and market structure: theory and history*, The MIT Press.

- Trefler, D. (2004), 'The long and short of the canada-us free trade agreement', *The American Economic Review* **94**(4), 870–895.
- Vannoorenberghe, G. (2012), 'Firm-level volatility and exports', *Journal of International Economics* **86**(1), 57–67.
- Verhoogen, E. A. (2008), 'Trade, quality upgrading, and wage inequality in the mexican manufacturing sector', *The Quarterly Journal of Economics* **123**(2), 489–530.
- Westphal, L. E. (2002), 'Technology strategies for economic development in a fast changing global economy', *Economics of Innovation and New Technology* **11**(4-5), 275–320.