The London School of Economics and Political Science

International and Innovation Activities of Firms

Claudia Steinwender

A thesis submitted to the Department of Economics of the London School of Economics for the degree of Doctor of Philosophy, London, April 2014.

Declaration

I certify that the thesis I have presented for examination for the PhD degree of the London School of Economics and Political Science is solely my own work other than where I have clearly indicated that it is the work of others (in which case the extent of any work carried out jointly by me and any other person is clearly identified in it).

The copyright of this thesis rests with the author. Quotation from it is permitted, provided that full acknowledgement is made. This thesis may not be reproduced without the prior written consent of the author.

I warrant that this authorization does not, to the best of my belief, infringe the rights of any third party.

I declare that my thesis consists of 43,310 words.

London, 17 April 2014

Claudic pinn

Claudia Steinwender

Statement of conjoint work

Chapter 2, "Survive another day: Using changes in the composition of investments to measure the cost of credit constraints", was jointly co-authored with Professor Luis Garicano.

This statement is to confirm that I contributed a minimum of 60% of this work. More specifically, I have carried out all of the empirical analyses. General decisions about the direction of the papers were made equally between the authors.

Claudic Brian

Claudia Steinwender

Luis Garicano

Acknowledgement

I am greatly indebted to my supervisor Steve Pischke for giving me all the support one could possibly wish for. Thanks for always pushing my identification strategy, replying to my emails in no time, and pointing out the flaws in my thinking. Much of what I have learned over the last couple of years is from trying to answer his critical questions.

I would also like to thank my advisors Luis Garicano and Gianmarco Ottaviano. Luis, thanks for being so enthusiastic about my work, for forcing me to think about the big picture, and for making research fun and exciting. I have thoroughly enjoyed working with him on the second chapter of my thesis. Gianmarco, thanks for your kind support and guidance, and for opening many doors.

A special thanks also to Steve Redding for his comments and advice, and for his generous invitation to Princeton, which provided an invaluable experience.

Many other researchers have provided feedback on several parts of the thesis: Walker Hanlon, Nikolaus Wolf, Thomas Sampson, Veronica Rappoport and Daniel Bernhofen on Chapter 1, Daron Acemoglu, Samuel Bentolila and especially Daniel Paravisini on Chapter 2, Emanuel Ornelas on Chapter 3; as well as other colleagues and seminar participants at the Centre for Economic Performance (CEP), at LSE and in various conferences. Financial support from the Royal Economic Society and the Toulouse Network for Information Technology is very gratefully acknowledged.

I will miss the inspiring atmosphere at the Centre for Economic Performance. Thanks to my fellow PhD students for making the PhD experience so much more enjoyable: Daniel Osorio, Anne Brockmeyer, Isabelle Roland, Rosa Sanchis-Guarner, Kati Szemeredi, Johannes Böhm, Giuseppe Berlingieri, Frank Pisch, Reka Juhasz, Esther Ann Boler, Jason Garred, Anna Sivropoulos-Valero, Michael Böhm, Christian Siegel, Hannes Schwandt, Georg Graetz, Michael Best, Alice Kügler, and many others.

I thank my parents for supporting the crazy idea to abandon my career and undertake a long, uncertain PhD instead - just to make my dream come true.

Finally, thanks to Jochen for coming with me to London, and sharing my exciting as well as my tough moments. Thanks for always believing in me.

Abstract

The economic environment in which a firm operates is constantly changing. This thesis contains three essays to examine how firms adapt their innovation and international activities to a variety of external changes.

The first paper, "Information Frictions and the Law of One Price: 'When the States and the Kingdom became United'", shows how information frictions affect the exporting behavior of merchants, exploiting a unique historical experiment: the transatlantic telegraph, established in 1866. Using a newly collected data set on cotton trade based on historical newspapers, I find that information frictions result in large and volatile deviations from the Law of One Price. There are also real effects, because exports respond to information about foreign demand shocks, and average exports increase after the telegraph and become more volatile. I provide a model in which exporters use the latest news about a foreign market to forecast expected selling prices when their exports arrive at the destination. Their forecast error is smaller and less volatile the more recent the available information. The welfare gains from the telegraph are estimated to be around 8% of annual export value.

The second paper, "Survive another day: Using changes in the composition of investments to measure the cost of credit constraints" is joint work with Luis Garicano. We introduce a novel empirical strategy to measure the credit shocks that were triggered by the recent financial crisis: Theoretically, we show that credit shocks affect long term investments by more than short term ones. Credit shocks can then be measured within firms by the relative drop of long run relative to short run investments; using firm-times-year fixed effects to absorb idiosyncratic demand shocks. Using data on Spanish manufacturing firms we find that credit constraints are equivalent to an additional tax rate of around 11% on the longest lived investment.

While the trade literature has established a positive impact of globalization on the productivity of firms, there is lacking consensus about the underlying mechanism at work: Trade theory focuses on a market access mechanism, but empirical papers point out that import competition matters as well. The third paper, *"The roles of import competition and export opportunities for technical change"*, conducts a "horse race" between the two mechanisms. Using Spanish firm level data, I find robust evidence that access to export markets leads to productivity increases, but only for firms that were already highly productive before. The evidence on import competition is weaker and very heterogeneous, pointing towards an omitted variable bias in earlier papers.

Contents

1	Info	ormation Frictions and the Law of One Price: 'When the States and the	
	Kin	gdom became United'	11
	1.1	Introduction	12
	1.2	Historical Setting	16
	1.3	Description of Data	20
	1.4	Reduced Form Evidence	22
	1.5	Model of Information Frictions in International Trade	30
	1.6	Welfare Gains from the Telegraph	41
	1.7	Conclusions	43
	1.A	Proofs	46
	1.B	Figures	51
	1.C	Tables	57
2	Sur	vive another day: Using changes in the composition of investments to	
	mea	sure the cost of credit constraints	64
	2.1	Introduction	65
	2.2	Theoretical Framework and Identification	68
		2.2.1 Theory: Investment Duration and Liquidity Risk	68
		2.2.2 Identification	70
	2.3	Data	71
		2.3.1 Identifying Long and Short Term Investments	71
		2.3.2 Description of Data Set	72
	2.4	Results	73
		2.4.1 Differential Effect across Investment Types	73
		2.4.2 Placebo Test	75
		2.4.3 Mechanism: Credit Crunch	75
	2.5	Conclusions	80
	2.A	Figures	82
	2.B	Tables	85
3	The	roles of import competition and export opportunities for technical change	96
	3.1	Introduction	97
	3.2	Description of Data	.00
	3.3	Trade Growth in Spain	01

Bibliography					
3.B	Tables	7			
3.A	Figures	3			
3.6	Conclusions	1			
3.5	Main Empirical Results	6			
3.4	Empirical Strategy	2			

List of Tables

1.C.1	Summary statistics	57
1.C.2	Average price difference	57
1.C.3	Variance of price difference	58
1.C.4	Variance of price difference using within period variation	58
1.C.5	Impact of telegraphed vs. steam shipped Liverpool price on New York price	59
1.C.6	Impact of known Liverpool price on exports	59
1.C.7	Average exports from New York to Liverpool	60
1.C.8	Variance of exports from New York to Liverpool	60
1.C.9	Estimation of the slope of the demand function	61
1.C.10	Estimation of the slope of the supply function	61
1.C.11	Parameters for calibration	62
1.C.12	Change from delayed to instantaneous information regime, in percent	62
1.C.13	Estimation of welfare gain from telegraph	63
2.B.1	Depreciation rates of different investment types	85
2.B.2	Summary statistics	86
2.B.3	Main results	87
2.B.4	Robustness checks: different measures for time-to-payoff	87
2.B.5	Placebo tests: treatment effect by year	88
2.B.6	Mechanism: Credit squeeze	89
2.B.7	Foreign versus domestic firms	90
2.B.8	Difference-in-differences	91
2.B.9	Robustness checks foreign versus domestic firms	92
2.B.10	Liabilities by maturity	93
2.B.11	Short term credit before crisis	94
2.B.12	Robustness checks: demand shock	95
3.B.1	Top 3 import origins by industry	117
3.B.2	Top 3 export destinations by industry	117
3.B.3	Import competition versus access to export markets	18
3.B.4	Heterogeneous effects	18
3.B.5	Estimated effect of import competition and export opportunities on pro-	
	ductivity growth of firms, by firms' initial productivity	19
3.B.6	Tariff changes as instrumental variables	19
3.B.7	First stages	20

3.B.8	Alternative measures for productivity
3.B.9	Comparing mean and quantile regressions
3.B.10	Comparing mean and quantile regressions, including exiting firms \ldots 122
3.B.11	Firm level exports
3.B.12	Imported inputs versus import competition
3.B.13	Ways to increase productivity
3.B.14	Implications for employment

List of Figures

1.B.1	Price series	51
1.B.2	Herfindahl index of market structure of cotton merchants	51
1.B.3	Information delay between New York and Liverpool over time	52
1.B.4	Distribution of information lags between New York and Liverpool (work	
	days), before and after the telegraph	52
1.B.5	Price difference and freight cost	53
1.B.6	Cotton receipts at the New York exchange	53
1.B.7	Reaction of New York prices to news from Liverpool	54
1.B.8	Impact of Liverpool price on New York price	54
1.B.9	Contemporaneous price difference and price difference to the known,	
	delayed Liverpool price	55
1.B.10	Surplus from export transaction, perfect foresight equilibrium	55
1.B.11	Surplus from export transaction when demand is overestimated	56
1.B.12	Price distortions: Difference between Liverpool price and lagged New	
	York price plus transport costs	56
2.A.1	Change in composition of investment towards short term investments	82
2.A.2	Investment change of long term investments by nationality of firm owner	83
2.A.3	Investment change by investment type, triple diff	84
2.A.4	Exit rates by domestic and foreign firms	84
3.A.1	Spain's world trade over time	13
3.A.2	Spain's trade by country over time	14
3.A.3	Spain's trade by industry over time	15
3.A.4	Estimated effect of import competition and export access on productivity	
	growth of firms, by firms' initial productivity	16
3.A.5	TFP distribution of all firms and exiting firms	16

Chapter 1

Information Frictions and the Law of One Price: 'When the States and the Kingdom became United'¹

How do information frictions distort international trade? This paper exploits a unique historical experiment to estimate the magnitude of these distortions: the establishment of the transatlantic telegraph connection in 1866. I use a newly collected data set based on historical newspaper records that provides daily data on information flows across the Atlantic together with detailed, daily information on prices and trade flows of cotton. Information frictions result in large and volatile deviations from the Law of One Price. What is more, the elimination of information frictions has real effects: Exports respond to information about foreign demand shocks. Average trade flows increase after the telegraph and become more volatile, providing a more efficient response to demand shocks. I build a model of international trade that can explain the empirical evidence. In the model, exporters use the latest news about a foreign market to forecast error is smaller and less volatile the more recent the available information. I estimate the welfare gains from information transmission through the telegraph to be roughly equivalent to those from abolishing a 6% ad valorem tariff.

¹The title of this paper is borrowed from a Citigroup advertisement in a campaign to celebrate its 200 year anniversary, as the bank provided capital for the laying of the telegraph cable (advertisement seen in the journal "City A.M." on 15 June 2012).

1.1 Introduction

The "Law of One Price" (LOP) states that if goods are efficiently allocated across markets, the price for identical goods in different locations should not differ by more than their transport costs. However, empirical studies document frequent and large deviations from the LOP (for example, Froot et al. 1995). Understanding the nature of the frictions that inhibit arbitrage across markets is one of the central objectives in international trade. Anderson and van Wincoop (2004) and Head and Mayer (2013) summarize the literature by observing that direct barriers to trade (for example transport costs and trade tariffs) have been found to be of minor importance. Therefore the recent emphasis of researchers has shifted to the search for more indirect barriers.

This paper focuses on information frictions as a potential explanation for "missing trade" (Trefler 1995) and deviations from the LOP. Information is essential for the efficient functioning of markets, but in reality often limited or costly (Jensen 2007; Stigler 1961). For example, exporting firms have to spend considerable time and money to learn about preferences of consumers in foreign countries and often fail while trying (Albornoz et al. 2010), especially if preferences are changing over time and production and export decisions have to be made in advance (Hummels and Schaur 2010; Evans and Harrigan 2005; Collard-Wexler 2013). The distortions from information frictions are hard to measure, as information flows are usually unobserved and also notoriously endogenous.

I use a historical experiment to circumvent these empirical issues: the construction of the transatlantic telegraph connection in the 19th century. First, the telegraph connection provides an exogenous and large reduction in information frictions. Before 28 July 1866, mail steam ships took between seven and 15 days to transmit information between the United States and Great Britain. The transatlantic cable reduced this information delay to a single day. The timing of the establishment of the connection was exogenous and not anticipated, because due to a series of technological setbacks over the course of ten years it remained unclear until the very end whether this new technology could ever work. It came as a big surprise when it not only worked, but also reliably and fast. Second, this paper is to my knowledge the first in the trade literature to observe information flows, which are based on news about foreign prices reported in historical newspapers.² The information flows are used to measure the impact of information on prices and exports, and to derive micro foundations for exporters' behavior under information frictions which I use to estimate welfare effects.

The empirical part of this paper focuses on cotton, the most important traded good between Great Britain and the United States in the mid-19th-century. The dominance of "King Cotton" in trade provides a unique setting to study information frictions, because historical newspapers published detailed and meticulous market reports on

²Previous papers using exogenous variation in information frictions used the presence of mobile phone coverage (Jensen 2007; Aker 2010), internet kiosks (Goyal 2010) or even the transatlantic telegraph connection (Ejrnaes and Persson 2010), but none of these papers observed information flows directly. Observing information ("news") and relating them to prices is much more common in the finance literature (for example Cutler et al. 1989, Koudijs 2013).

cotton. No other good was reported at a daily frequency and to such a degree of detail. Surprisingly, these rich data have never been systematically digitized. I use market reports from newspapers on both sides of the Atlantic – *The New York Times* and the *Liverpool Mercury* – to construct a new, daily data set that includes cotton prices in New York and Liverpool, export flows and freight cost between the two ports, stock levels in both markets and detailed information flows for the period of one year before and one year after the telegraph connection. The use of this data set has several advantages: First, using export as well as price data makes it possible to understand whether information has a real effect, as opposed to only distributing profits across agents. Second, it is possible to study the impact of information frictions on a durable good. Jensen (2007) provides evidence that information reduces spoilage of fish, a highly perishable good, but it is not clear whether the same is true for a storable commodity. Third, shipping time makes imports predetermined, which allows me to identify the supply and demand functions that are needed for the welfare estimation.

Using this detailed data, I am able to document six "Stylized Facts": (1) The telegraph caused a better adherence to the Law of One Price, as the mean and volatility of the price difference fell. (2) Within the pre-telegraph period, faster steam ships had a similar effect and reduced deviations from the Law of One Price. In contrast, within the post-telegraph period, temporary technical failures of the connection led to increased deviations from the Law of One Price. (3) New York prices respond to news from Liverpool. Before the telegraph, only Liverpool prices lagged by ten or more days are relevant in determining New York prices. After the telegraph, the transmission of shocks across prices is reduced to almost real-time.³ (4) Market participants base their search for arbitrage opportunities on the latest news from Liverpool. (5) Information frictions have real effects and are not just a reallocation of profits across market participants, because exports respond to news about Liverpool prices.⁴ (6) After the telegraph, exports are on average higher, and more volatile.

In order to establish a causal relationship between these findings and the telegraph, I use two complementary strategies. First, the findings are robust to a number of alternative explanations (for example transport cost variations, supply irregularities in the aftermath of the American Civil War, fluctuations within bounds given by trade cost in no trade periods, change in the market structure of merchants, futures or forward trading, and anticipation effects). Second, to rule out any confounding trends that happened over time, I use another source of exogenous variation in information flows *within* the period before the telegraph was established: the irregular passage times of steam ships across the Atlantic, which were driven by weather conditions.

I present a partial equilibrium model of trade under information frictions that

³A related paper by Ejrnaes and Persson (2010) estimates faster shock transmission after the telegraph using weekly wheat prices. However, their price series exhibits a gap of 15 years around the time of the establishment of the telegraph connection, which makes it difficult to distinguish the effect of the telegraph from other confounding trends such as the introduction of futures trading in the 1870s.

⁴To my knowledge, my paper is the first to provide evidence that the telegraph had real effects on exports. An earlier paper by Ejrnaes and Persson (2010) does not observe trade flows, takes real effects as given and estimates the welfare gains from the transatlantic telegraph using estimated demand and supply elasticities from other studies that are based over yearly rather than weekly time horizons.

provides a micro foundation for the empirical findings and can be used to study the welfare effects of information frictions. In the model, as in 19th century trade, intermediaries act as arbitrageurs across geographic markets. They buy cotton from suppliers at a centralized exchange in New York and ship it to Liverpool where they sell it to cotton millers, again at a centralized market place. Aggregate demand from cotton millers follows a stochastic, autocorrelated process. Shipping takes time, so merchants have to make their export decision before they know the realization of the demand shock,⁵ and will base it on the prices they *expect* to generate in Liverpool upon arrival of their shipment. Information frictions affect the information available to merchants when they build these conditional expectations: If frictions are low, market conditions in Liverpool are observed up to the current date. If frictions are high, lagged information about market conditions in Liverpool have to be used to predict future selling prices. In this model, merchants optimally choose exports taking their information as given.⁶

In the model I allow the commodity to be storable and study its consequences for information frictions. Storage softens the impact of information frictions. If exports are inefficiently high based on wrong expectations, cotton can be stored until demand is higher. However, the smoothing effect of storage is asymmetric, because negative amounts cannot be stored (Williams and Wright 1991; Deaton and Laroque 1996). If exports are inefficiently low based on wrong expectations, there might not be enough stock available to smooth prices. There is always a positive probability that long periods of especially high demand will run down inventories, and a finite stock can never fully insure against aggregate demand shocks.

I calibrate the model to match the historical data *after* the telegraph was introduced. Then I conduct a counterfactual analysis by increasing information frictions to simulate the effect of the telegraph. The resulting predictions are consistent with the reduced form evidence: The volatility of trade flows increases after the telegraph connection, because exports follow expected demand shocks in Liverpool.⁷ With better information, expected demand shocks are more volatile. Average exports are lower before the telegraph connection, because periods of high demand are systematically underestimated with information frictions. An asymmetry arises from restricting exports to be positive:⁸ While periods of low demand are also systematically overestimated with information frictions, in these periods it is never profitable to export. As a result, average exports increase after the telegraph, because in periods of high demand exports are higher. The distorted export flows are reflected in distorted price equalization: After the telegraph,

⁵Aggregate demand shocks cannot be fully insured away, as borrowing cotton from future harvests is impossible. Furthermore, since futures trading had not yet been established, the risk could also not be reallocated across market participants, for example from merchants to "speculators", so merchants had to bear the full market risk of their ex-ante export decision.

⁶This is different from Allen (2012), who models information frictions as a costly search across markets, and merchants optimally decide on how much information to acquire.

⁷To my knowledge, this is the first paper in this literature to model information directly as the way how conditional expectations are formed.

⁸In models without time lag due to shipping, negative exports can be interpreted as imports. However, with time-consuming shipping, negative exports are "imports from a future period", an unrealistic assumption.

the average and the volatility of the price difference falls.

The model provides an analytical solution for the lower bound of the deadweight loss arising from distorted trade flows under information frictions based on Harberger Triangles: The deadweight loss from information frictions is a function of the squared observed price difference between New York and Liverpool (taking into account the shipping lag) as well as the slopes of the demand and supply curves. The reduction in the absolute observed price difference after the construction of the telegraph connection correspond to the abolishment of an ad valorem tariff of around 6%. To see how this translates into welfare gains, the slopes of the supply and demand functions need to be estimated. This estimation is usually difficult due to the simultaneous determination of quantity and prices, and a valid instrument cannot always be found. I propose a novel identification strategy that exploits the fact that exports are predetermined once they arrive in Liverpool, since shipping takes time for transatlantic cotton trade. This breaks the simultaneity problem for the case of i.i.d. shocks. For the case of autocorrelated shocks and positive storage the model can be used to control appropriately for the endogenous part of the shocks, yielding identified regression equations. Combining the evidence I estimate the welfare gain from the telegraph to be equal to 8% of the annual export value of American cotton.

What are the implications of this paper for today's modern world, when optical glass fiber cables have long since replaced the copper wires of the telegraph? The historical example of the transatlantic telegraph provides a micro foundation for how exporters (or equivalently producers) use information about demand shocks to forecast demand and decide ex-ante on export (or production) quantities. Exporters and firms still face this problem today, and new emerging technologies such as the real-time analysis of "Big Data" have the potential to provide firms with immediate information about consumer behavior (McAfee and Brynjolfsson 2012). My model can be used to assess the welfare effects from these technologies.⁹

This paper contributes to an emerging literature on information frictions in trade. Information frictions can take different forms: One branch of the literature focuses on the information frictions in the search and matching process of buyers and sellers across international markets. Rauch and Trindade (2002) show that social networks help to overcome these frictions and increase trade.¹⁰ Other papers focus on the role of information technology to overcome these frictions. For example, Jensen (2007) and Aker (2010) use mobile phone coverage, while Goyal (2010) and Brown and Goolsbee (2002) use internet based price comparisons. This paper contributes to this strand of the literature in several ways: It observes also data on information flows and can relate this to outcomes; it demonstrates that information has real effects by observing exports and not only price differences; it provides a novel identification strategy to estimate the welfare effects; and, compared to Jensen (2007), it shows that information frictions are

⁹Technologies that reduce the time lag between the production decision and consumption, such as faster transport and supply chain management, have a similar effect, see also Hummels and Schaur (2010, 2012); Evans and Harrigan (2005); Aizenman (2004); Harrigan (2010).

¹⁰Similarly, Head and Ries (1998); Rauch (1999, 2001).

also relevant for the case of a storable good in the context of international trade. Instead of studying how technological innovations can overcome information frictions, Allen (2012) models the optimal behavior of agents when search is costly, and characterizes the resulting trade pattern. Another branch of the literature interprets information cost as fixed cost for entering an export market and shows how technology can reduce it (Freund and Weinhold 2004).

These interpretations differ from the mechanism in this paper which interprets information frictions as a source for making forecast errors when predicting foreign market conditions. This view is related to the finance literature's focus on the effect of news on capital prices (Cutler et al. 1989; Koudijs 2013) and on how information technologies can increase the efficient functioning of capital markets (Portes and Rey 2005, Garbade and Silber 1978¹¹, Field 1998).

My paper focuses exclusively on the information effects of the telegraph. However, in the long run there might be other, additional effects: For example, Lew and Cater (2006) argue that the telegraph reduced transport costs by increasing the capacity utilization of shipping, which increased trade flows (however, using only data from after 1870). Clark and Feenstra (2003) argue that the telegraph enabled international transfer of other production technologies. These and other additional, long-run effects would increase the welfare gains brought about by the telegraph beyond the ones estimated in this paper.

The structure of the paper is as follows: Section 1.2 describes the historical setting, and Section 1.3 describes the collected data set. Section 1.4 provides reduced form evidence on the effect of the telegraph. Section 1.5 develops a theoretical model of information frictions and intermediaries in international trade that is consistent with the empirical findings. Section 1.6 estimates the welfare effects of information frictions. I conclude in Section 1.7.

1.2 Historical Setting

Transatlantic cotton trade was the world's most important single trade linkage in mid 19th century. For the United States, half of exports to the world was in "King Cotton".¹² For Great Britain, a third of world imports (36% in 1866) was in cotton (The Economist 1866).

In the mid 19th century, cotton was grown primarily in the South of the United States (over 55% of world production, Ellison 1886). The second largest producing country was India (29%), followed by Egypt (9%) and Brazil (5%). The dominance of the United States in cotton production is mainly explained by the superior quality of "American cotton", whose longer and stronger fibers were preferred by spinners

¹¹Garbade and Silber (1978) show that the transatlantic telegraph connection reduced average spatial price difference and volatility of stock prices.

¹²Bruchey (1967). The term "King Cotton" was coined before the American Civil War and reflected its tremendous importance for 19th century economies (Surdam 1998). This dominance did not stop after the war: American cotton "was not toppled from his throne" and "resumed his former position of power", as Woodman (1968) phrases it.

(Irwin 2003). Other advantages of American cotton were lower production cost, lower transport costs and faster shipping time.

Cotton millers spun the raw cotton into yarn, which was then woven into fabrics, and sewn into a wide variety of apparel and accessories. The industrial revolution in Great Britain had led to several inventions in cotton manufacturing such as spinning machines, the spinning jenny, or the spinning frame, making the country the world's leading textile manufacturer: Great Britain produced 85% of worldwide cotton manufactures, and consumed half of the world's cotton production (Ellison 1886). Textile manufacturing was geographically highly concentrated and took mainly place in Lancashire, the hinterland of "Cottonopolis" Manchester.

Virtually all the cotton destined for Great Britain arrived at Liverpool, Lancashire's closest port. On the other side of the Atlantic, New York was the major port exporting to Great Britain: In 1866, 33% of cotton exported to Great Britain arrived from New York, followed by New Orleans (28%) and Mobile (18%).¹³

A thriving mercantile community was responsible for bringing cotton from source to destination.¹⁴ Most merchants were generalists: In the 1860s, only 11-13% specialized in a specific commodity, and 13-14% specialized in certain trade routes (Milne 2000). Merchants were early multinationals. They usually set up a subsidiary in important foreign port cities, mostly run by family members (Ellison 1886; Milne 2000; Chapman 1984). Merchant trade was associated with relatively low entry cost, leading to fierce competition.¹⁵ In fact, historical trade directories reveal that around 100–200 merchants were active in cotton trade from New York to Liverpool in 1866.¹⁶ Merchants were usually not credit constrained, as there was a well developed and functioning banking sector that provided trade financing (Chapman 1984; Brown 1909).

Organized exchanges for cotton existed in both New York and Liverpool. Merchants bought cotton at the New York exchange from cotton farmers, shipped it to Liverpool, and sold it at the Liverpool exchange to cotton millers. Due to the dominance of Great Britain in textile manufacturing, Liverpool essentially constituted the world price for cotton. Cotton futures trading had not yet been established.¹⁷ At each exchange there were also so called "speculators", who bought cotton when they thought prices would go up, stored it, and sold it at a later time. About 80% of the cotton stock was stored in warehouses near the ports by speculators, while spinners held only some widely

¹³Shipping and Commercial List, printed on 11 October 1866 in The New York Times

¹⁴Direct exporting by cotton farmers consisted of <1% of imports (own estimation based on a sample of the Bills of Entry generously provided by Graeme Milne and data from historical trade directories).

¹⁵Milne (2000) notes that in the 1850's and 1860's many people entered the trading profession and competition was so large that some traders were willing to work on a no-profit, no-commission basis.

¹⁶Own calculations based on the Bills of Entry.

¹⁷Futures trading involves the trading of a highly standardized contract based on a clearly defined quality of the underlying good, that can be enforced; and the possibility to short sell. Institutions are needed for objective assessment of the quality of the commodity, for drawing up standardized contracts, and for legal enforcement of the contract. These institutions were set in place only by the 1870s (Ellison 1886, Hammond 1897). There is some limited evidence of forward trading ("on arrival" business, the selling of a specific cotton lot that the seller possesses for delivery at a later date), but this was done only when a sample of the cotton in question could be inspected (again, because there was not yet a procedure for enforcement of a promised quality of cotton). In summary, no short selling of cotton was possible.

scattered stocks. Traders assumed the storage cost, while manufacturers stored only as much cotton as they needed to supply their mills in the short run (Milne 2000).

When merchants bought their cotton at the New York exchange, they had to forecast demand conditions in Liverpool upon arrival of their shipment. Demand for cotton at the exchange in Liverpool originated from cotton millers, whose customers were domestic but also foreign; mainly from Continental Europe. Market reports in historical newspapers describe how export demand for cotton textiles fluctuated frequently depending on the course of wars and peace negotiations on the continent, which could take quick surprising and unexpected turns. When a country on the continent was in war or in threat of war, its demand for cotton textiles dropped considerably, as the country shifted its funds towards war expenditures such as arms and munition. The Austro-Prussian and the threat of the Franco-Prussian war fall into my sample period, and historical newspapers frequently identified a change in the war conditions as source for increasing or falling demand from cotton millers.

Information was therefore important at the cotton exchanges. The 19th-centuryequivalent to computer screens with price tickers was a large billboard with the latest price information and news, together with circulars that summarized market developments, provided in the *Exchange Newsroom*. The news agency Reuters provided a subscription service with the most important news from all over the world. The compilation of news included the cotton prices from New York and Liverpool and was called *Reuter's Telegram* – even before the transatlantic cable was established, because the news traveled the overland part of the way via telegraph. Contemporaneous newspapers as well as the cotton exchanges were subscribers. Since these news were posted publicly at the exchange, the cost for individual merchants to obtain them was zero.

The first successful transatlantic telegraph connection between Great Britain and United States was established on 28 July 1866 and caused a dramatic reduction in the delay of information across the Atlantic. Before the telegraph connection, the only means of communication across the Atlantic was sending letters and messages (including a print of *Reuter's Telegram*) via steam ships. The so called "mail steam ships" were the fastest ships of those times, specialized in speedy transmission of information items such as letters, newspapers and other documents. There was fierce competition among mail steam ships to win the unofficial "Blue Riband" for record speeds, and by 1866 the fastest ship had crossed the Atlantic in little over eight days (Gibbs 1957). However, these speed records could only be achieved under the best possible weather conditions, which resulted in daily variation in communication times. If conditions were very bad, ships could take as long as three weeks to cross the Atlantic. Important commercial information was transmitted between the commercial hubs in the United States and Great Britain using a combination of existing land based telegraph cables and mail steam ships.¹⁸

¹⁸For example, the Liverpool price of cotton was telegraphed from Liverpool via the submarine Ireland/Great Britain connection on to a steam ship passing the coast of Ireland on its way to the United States. As soon as this ship reached the first telegraph post on the US coast, this information was further telegraphed to New York by land line, arriving faster at its destination than the steam ship.

The transatlantic telegraph connection changed communication flows dramatically and immediately. For the first time in history, information traveled faster than goods across the Atlantic (Lew and Cater 2006). From one day to the next, communication between the United States and Great Britain was possible within only one day.¹⁹ There were occasional technical break downs of the telegraph connection, but these were usually repaired within a couple of days and communication was restored.

The timing of the successful telegraph connection was unforeseen and exogenous to economic conditions, because the process of establishing a telegraph connection was characterized by a series of failures and setbacks over the course of around ten years, resulting in little confidence in the feasibility of a transatlantic telegraph connection. These technical difficulties arose because the transatlantic cable was the first undersea cable connecting two continents, which required to cover a greater distance (3,000 km) at a larger submarine depth (3,000 m) than any previous telegraph connection.²⁰ Consequently, it took 5 attempts over the course of almost 10 years until a lasting connection was established on 28 July 1866.²¹ The first attempt in 1857 had resulted in a snapped cable, whose ends were lost in the deep sea. The second attempt in 1858 produced a working connection; however with an extremely slow transmission speed that could not be used for commercial purposes,²² and the connection lasted only briefly. After three weeks the insulation of the cable was damaged, and the connection broke down permanently. After this failure the public lost faith in the telegraph project, and another attempt in the same year was delayed indefinitely. In fact, the faith in the technology had become so poor that the media suspected the working connection had been a "hoax" altogether. The Boston Courier asked: "Was the Atlantic Cable a Humbug?"

Although technical understanding of undersea electrical signal transmission had progressed, the fourth attempt in 1865 resulted again in a broken cable with ends that got lost in the ocean. By 1866 there was little confidence left. Even if people had expected this fifth attempt to work, the precise timing could not have been foreseen, as weather conditions determined the progress of the cable laying steam ship. Nonetheless, to everybody's surprise and excitement, on 28 July 1866 the first telegraph message, a congratulation message from the Queen of England to the President of the United States, was transmitted. From then on, the telegraph worked surprisingly reliably and fast. The newspapers of the next working day already reported cotton prices from the other side of the Atlantic in their commercial sections. By early September the 1865 cable was fished out of the sea and repaired. The two working transatlantic connections

¹⁹Messages sent from Great Britain to New York passed several telegraph posts along the route, and had to be retransmitted at each of the posts (Lew and Cater 2006). Therefore effective communication time between Liverpool and New York was around one day.

²⁰The previous submarine cables connected Great Britain to Ireland and France; they were much shorter and at a much shallower depth.

²¹Clarke (1992) provides a detailed and entertaining history of the cumbersome way towards a transatlantic connection.

²²The first message took 17 hours to transmit. Overall, the average transmission speed was 0.1 words per minute. The messages being sent were concerned with how to increase speed and trying to resolve misunderstandings (Clarke 1992).

provided reliable and fast transatlantic communication. The transatlantic cable was subsequently referred to as the "Eighth Wonder of the World", reflecting people's amazement about this technological milestone.

Once completed, the contemporary press had high hopes for the impact of the transatlantic telegraph: "The Atlantic Cable will tend to equalize prices and will eliminate from the transactions in bonds, in merchandise and in commodities, an element of uncertainty which has had the effect of [...] seriously damaging the commercial relations between this country and Europe."²³ This paper uses empirical and theoretical evidence to assess whether this prediction came true.

1.3 Description of Data

For establishing a causal relationship between delayed information, market integration and trade flows data requirements are substantial. First of all, I need price and export data on an *identical* good from at least two different market places. Many observed "violations" of the Law of One Price can be blamed on a lax interpretation of this requirement (Pippenger and Phillips 2008). For example, wheat grown in the United States and wheat grown in Great Britain are not identical, and even different varieties of wheat grown in the United States are not identical. This is a severe restriction on the data, as many local newspapers – the primary source of historical market information – report prices of the local variety and not foreign varieties. Sometimes, for example for wheat, they report prices on foreign varieties, but then not for the same variety over a consistent period of time.²⁴ Another pitfall when studying the Law of One Price is using retail instead of wholesale prices (Pippenger and Phillips 2008), so it would be ideal to have data on a good that is traded on organized exchanges rather than local farmers' markets.

Second, these prices and export flows should be reported at a daily frequency, to correspond to the actual adjustment horizon of prices to information in the real world. I can then relate price changes on a certain day to news arriving on that day. Using weekly data decreases the power of tests relating prices to news, and observed time periods for consistent varieties (e.g. for wheat) are not long enough to compensate for that (usually after 2–3 years the reported varieties change).

Third, I need data on information flows across the Atlantic. Newspapers report the arrival of some type of news, but often these reports consist of political news and information about stock prices and exchange rates rather than specific commodity markets.

²³New York Evening Post, 30 July 1866, as cited in Garbade and Silber (1978).

²⁴For example, the Aberdeen Journal reported weekly American winter and American spring prices for the London market, but stopped in July 1866 for no apparent reason, similarly the Economist and The Daily Courier for American and Canadian Red Wheat. In contrast, the Daily Courier started to report Chicago wheat prices only in August 1866. Some newspapers report weekly wheat prices for some American varieties over longer time series, but the prices do not have any variation, which means there was no underlying trade based on the commodity, and prices were just copied forward at the same level for months. Ejrnaes and Persson (2010), who also fail to find grain price data that cover the years around 1866, explain that these years were a period when the export of US grain ceased temporarily.

The importance of "King Cotton" in mid-19th-century allowed me to locate newspapers at important ports on either side of the Atlantic that provided detailed, daily information on cotton markets and trade flows. Furthermore, newspapers also reported news about foreign cotton prices, which makes it possible to reconstruct information flows. The richness of cotton data is extraordinary. No other good is consistently reported at such a high frequency in two different countries for the same variety around the mid 19th century.²⁵

The resulting data set combines four types of data: market information from the Liverpool exchange, market information from the New York exchange, trade flows between New York and Liverpool, and information flows between New York and Liverpool (and vice versa).

Market information from the Liverpool exchange was reported in the *Liverpool Mercury*, which had a daily section called "Commercial" that provided a detailed market report on cotton. The *Liverpool Mercury* published the daily price for "Middling American", where "middling" indicates a specific quality of American cotton (other qualities that existed, but were not reported consistently, were "fair", "middling fair", "ordinary"). In addition, the *Liverpool Mercury* provided weekly estimates of the stock of American cotton in the warehouses of Liverpool.

Market information from the New York was reported in *The New York Times*, which also published a daily commercial section with detailed information on cotton. Again, the prices reported there are for "middling" American cotton.²⁶ *The New York Times* reported also a weekly and later bi-weekly estimate of the stock of cotton in the warehouses, as well as the daily "receipts" of cotton from the hinterland that arrived at the exchange on that day. I convert the prices at the New York exchange from US dollars to Pound Sterling using daily exchange rates from the historical time series provided by *Global Financial Data*.Great Britain had adopted the gold standard in those times. Overall, the fluctuations in exchange rates were very small.²⁷ Figure 1.B.1 illustrates the resulting time series of daily New York and Liverpool cotton prices. The Liverpool price for cotton exceeded the New York price almost always, except for a short period in May 1866. The price series seem to follow each other by and large.

The New York Times also had a separate "Freights" section, which reported daily the bales of cotton that were shipped to Liverpool, as well as the freight cost paid for that shipment.²⁸

I can also reconstruct the data on information flows from the historical newspapers, as both newspapers reported the latest mail ship and telegraph arrivals on any given day and printed the main commercial indicators from the other country that these

²⁵Wheat data are often used for market integration studies. However, these are available at most at weekly frequency, and qualities of wheat are often not comparable as there exist so many local varieties. Furthermore wheat exports from the United States to Great Britain ceased for several years around 1866 (Ejrnaes and Persson 2010).

²⁶Several market reports pointed out that New York used the same classification scheme as Liverpool, as this was the most important destination for cotton.

²⁷Using the average exchange rate for the whole period as opposed to the daily exchange rates does not affect the results in this paper.

²⁸Very few shipments are reported to go to other ports in Europe.

shipped or telegraphed messages included. The relevant sections were headed "Latest and Telegraphic News" and "News from Europe", respectively. These indicators included certain bond and stock prices and the price of cotton. The newspapers also reported the origination date of these business indicators in the other market and the arrival date of the information. The difference in these dates yields the information transmission time across the Atlantic for any given day, which I call "information delay *l*". The measure of transmission times in my data corresponds to the fastest possible way of communicating between Liverpool and New York, and not to the corresponding steam ship travel times.²⁹ Sometimes steam ships were overtaken by other, faster steam ships, and its news were "old". In that case the newspapers reported "news were anticipated".

My final database comprises 605 observations, one for every work day between 29 July 1865 and 27 July 1867. The cotton exchange was open every week Monday through Saturday, except on holidays and a few other special occasions (for example, during a "visit of the Prince and Princess of Wales"). I discarded days which were holidays only in the UK or only in the US. The resulting time period comprises one calendar year before (301 work days) and one calendar year after the telegraph connection (304 work days).

The American Civil War between April 1861 and April 1865 severely disrupted cotton exports from the United States, restricting the period of analysis.³⁰ In addition, historical newspapers did not report cotton prices before that. While it is possible to extend the period of analysis to years after 1867, I kept symmetry between the before and after telegraph periods.

1.4 Reduced Form Evidence

The telegraph changed information frictions dramatically and suddenly: Figure 1.B.3 plots the time delay for information from Liverpool reaching New York for each day in the data set. This series shows a sharp drop on 28 July 1866, when the transatlantic telegraph was established. Before that, information from New York was around 10 days old when it reached Liverpool. After the telegraph, information from New York was usually from just the day before. Figure 1.B.4 shows the distribution of information lags before and after the telegraph: Before the telegraph, information lags varied between 7 and 15 days, caused by wind conditions that affected the speed of mail steam ships. After the telegraph information lags varied between 1 and 6 days, with lags over 1 day due to temporary technical breakdowns of the connection in the first few months of operation. However, these failures were usually quickly resolved. Table 1.C.1 confirms that the drop in average information transmission speed after the establishment of the telegraph connection was statistically significant.

²⁹The difference arises because steam ships often got the latest commercial news from England via telegraph while passing the last part of the Irish Coast, and upon arrival on the Newfoundland Coast the news were again transmitted via telegraph to New York, arriving faster than the steam ship.

³⁰I discuss potential implications of the American Civil War for my analysis in the empirical section.

How did this drop in information frictions affect the integration of the Liverpool and New York cotton markets? In this section I carefully develop six "Stylized Facts" that describe what happened to cotton prices and trade.

(1) The telegraph caused a better adherence to the Law of One Price, as the mean and volatility of the price difference fell.

Following the literature on the Law of One Price (LOP), I use the price difference between the two markets as a measure of market integration (Dybvig and Ross 1987; Froot and Rogoff 1995). When markets are perfectly integrated, the price difference should be zero, as any positive price difference has been arbitraged away. The telegraph reduced information frictions which are likely to have constituted a barrier to arbitrage, so we should expect to see the price difference go to zero.³¹

Figure 1.B.5 plots the difference between Liverpool and New York cotton prices. The vertical line indicates 28 July 1866, the date when the telegraph connection was established. The change in the behavior of the price difference due to the telegraph is striking: The volatility of the price difference falls sharply, and there are fewer very large and very small values. The average price difference falls as well. Table 1.C.1 shows that the average price difference was 2.56 pence/pound in the pre-telegraph period (16% of New York price), and fell to 1.65 pence/pound. This reduction is statistically significant, and corresponds to a fall of 35%. The variance of the price difference falls by even more, by 93%, and the coefficient of variation falls by more than half.³²

Are these drops causally related to the transatlantic telegraph? The troublesome history of the transatlantic telegraph connection is in favor of this interpretation: The timing of the successful establishment was driven by technical "luck" and the weather, and therefore exogenous to market conditions.³³ The date of the connection could not have been deliberately timed by market participants to coincide with other market events or developments, and anticipation effects can also safely be excluded.³⁴

In Table 1.C.2 I show that the observed deviations from the LOP are robust to a number of alternative explanations. For example, one alternative hypothesis is that the observed pattern in the price difference is caused by variation in transport costs rather than information frictions. In fact, the Law of One Price can only be expected to hold after taking into account transport costs. In empirical trade papers, transport cost is rarely observed and often derived from the price difference. However, in my case *The New York Times* listed the daily freight cost of cotton for shipment from New York to Liverpool. Cotton could be shipped either using the slow sailing ships (taking 1–2 months) or the faster steam ships (taking 2–4 weeks). Figure 1.B.5 plots the freight

³¹Garbade and Silber (1978) perform this check for stock prices

³²Note that usual explanations for a fall in volatility like exchange rates or sticky prices are not relevant in this setting (Froot et al. 1995).

³³Weather conditions affected the advancement of the cable laying steam ship, and its ability to locate and repair problems in the cable.

³⁴For example, by withholding cotton trade in the weeks before the telegraph until the telegraph gets established.

rates for both transport types. Freight costs are lower in the post-telegraph period (see Table 1.C.1), but the reduction is very small compared to the drop in the price difference. In columns (2) to (4) of Table 1.C.2 I subtract the freight cost – by sail ship, steam ship, or an average of both – from the price difference. However, the fall in freight cost is too small to explain the drop in the price difference after the telegraph connection.³⁵

While freight cost accounted for the major part of total transport costs, there were other transport costs such as fire and marine insurance, wharfage, handling at the port etc. Boyle (1934) provides a detailed account of all other transport costs, using historical bookkeeping figures of merchants.³⁶ The majority of transport costs, 83.1%, are charged based on weight, so they were unit cost.³⁷ Freight cost are the most important component of unit transport costs, comprising 65% of total transport costs. The remaining unit transport costs are paid for handling at the ports (including bagging, marking, wharfage, cartage, dock dues, weighing, storage at the port). Ad valorem transport costs constitute 16.9% of total transport costs and include fire and marine insurance, Liverpool town dues, and brokerage.³⁸ Based on these numbers I plot total transport costs in Figure 1.B.5. Column (5) in Table 1.C.2 shows that even after accounting for total transport costs we observe a large drop in the average price difference.

The Law of One Price does not hold in periods when there is no trade between two markets. In these periods, transport costs are too high, and the price difference will fluctuate freely between the bounds given by the transport costs (called commodity points): $|p_t^{LIV} - p_t^{NY}| < \tau$. If there were some periods before the telegraph when the price difference was not large enough to induce trade, this might explain why I observe high volatility before the telegraph. In fact, in my data exports occurred in every week in the sample except for a period of about four weeks during May 1866 (before the telegraph), when the threat of a war between Austria and Prussia depressed demand for cotton in Liverpool and lowered prices so much that exporting became unprofitable.³⁹ Column (6) of Table 1.C.2 excludes this period, but the results are again robust to this check.

Another concern might be that my observations begin in July 1865, three months after the end of the American Civil War. During the Civil War, the Northern states (the "Union") established a blockade of Southern ports ("Confederates") that stopped cotton exports almost completely. After the war, cotton production and trade were immediately taken up again: Woodman (1968) describes how the reopening of trade

³⁵Lew and Cater (2006) argue that the telegraph reduced freight rates, so the observed drop in freight cost could also be attributed to the telegraph. However, at least in the short run this is not the major contribution of the telegraph.

³⁶See online appendix for a detailed breakdown of total transport costs.

³⁷This is also why I do not use a log specification of the Law of One Price, this is only helpful with multiplicative transport costs.

³⁸No export tariff or import tariff was imposed during the period under consideration.

³⁹Only if the price difference becomes "negative enough" to cover transport costs, should we expect cotton re-exports to New York (and for those periods the LOP should hold again in absolute values). While the price difference was not large enough, there is no indication from the historical newspapers that this happened in that period.

with the South immediately induced a "scramble among cotton merchants". However, it took five to 10 years before the pre-war levels of cotton production were restored. Reasons for the slow recovery included the destruction of cotton during the war, the substitution of cotton production for food production, bankruptcy of many cotton planters, and the abolishment of slavery (Woodman 1968). Cotton production fell by three quarters from four to one million bales during the years of the Civil War, and reached 2 million bales, half of the pre-war production, again in the first harvest after the Civil War (cotton year 1865/1866).⁴⁰ The return to pre-war levels took until 1870.⁴¹ The first year of observations in my data coincides with the first year of cotton production after the Civil War. It is possible that cotton supply is still disrupted during that year. If there are no barriers to arbitrage, a larger volatility of production affects only price levels and not the price difference, as shocks are transmitted to the other country. To account for the possibility that there were some barriers to arbitrage, and to investigate whether supply irregularities therefore had an effect on the adherence to the LOP I use data on the quantity of cotton that arrived at the New York cotton exchange from the cotton farms on any given day, the so called "cotton receipts". Figure **1.B.6** illustrates the time pattern of cotton supply over the course of a harvest year. The cotton year starts in September, when the new harvest starts to come in. The winter months October to February are the months with the largest receipts of cotton, whereas the summer months June to August are the months with the smallest receipts. However, due to the time consuming cotton picking process and the long distances from the cotton fields in the interior to New York, the supply of cotton is positive on every single day in the sample. The visual evidence in Figure 1.B.6 does not suggest that the variation in cotton supplies differs very much before and after the telegraph.⁴² Column (7) of Table 1.C.2 controls for cotton supply, but again this is unable to explain the fall in the price difference after the telegraph.

Column (8) of Table 1.C.2 controls for shipping time by using the price in Liverpool at the time of the arrival of the shipment instead of the contemporaneous Liverpool price to construct the time difference. I use the steam ship travel times (around 10 days) for all shipments, because even if the cotton shipment was transported by the slower sail ship, samples of the cotton lot were usually shipped by a faster mail steam ship. The lot was then sold on the spot market in Liverpool upon inspection of the sample, while still on sea, which is called "forward trading" (Milne 2000). Again, correction for shipping time does not affect the results. Finally, column (9) shows the difference in log prices instead of price levels. Again, this does not affect the findings.

Contemporary observers describe that the transatlantic telegraph contributed to

⁴⁰The detailed time series is provided in the online appendix.

⁴¹During the Civil War American cotton was only partly substituted with Indian and Egyptian cotton. Irwin (2003) argues that the low supply elasticity of other countries was due to the fact that planters expected the war to be temporary and were therefore unwilling to make long-term investments in cotton cultivation. The advantages of American cotton (longer fibers, lower production and transport costs) still prevailed after the Civil War, which is why cotton millers returned to American cotton after the Civil War, as long as prices were not too high.

⁴²There is just one outlier in the pre telegraph area that is due to the closure of the New York cotton exchange over two Christmas holidays, when arrivals had piled up.

the development of futures trading in cotton across the Atlantic, as for the first time information traveled faster than goods (Hammond 1897; Ellison 1886). If the change in the pattern of the price difference is due to the introduction of futures trading, it is indirectly caused by the telegraph (rather than directly by changing information frictions). However, the development of futures trading was not immediate. The necessary institutions for futures trading were set in place only by the 1870s. Forward trading, as described earlier, was still limited and based on a sample of cotton made available for inspection. The telegraph did not change the speed at which cotton samples could be shipped, as they had to be transported physically, it only changed the speed of the transmission of *information*. An introduction of futures or forward trading due to the telegraph can therefore not explain the observed findings.

Finally, it is interesting to check whether the observed change in the price difference is the result of a change in the markups of merchants, maybe because of increased competition among merchants. The Bills of Trade record the shipment of every merchant arriving in Liverpool and allow for the computation of a Herfindahl Index. The Bills of Trade have been digitized for 3 months (February, June, October) of four years around the time of the introduction of the telegraph (1855, 1863, 1866, and 1870).⁴³ Figure 1.B.2 shows the development of the Herfindahl Index for cotton merchants, separately for shipments from the US and shipments from Egypt and the East Indies. In 1863 the American Civil War disrupted cotton trade, only few cotton shipments came from the US to Great Britain, and the Herfindahl Index shows high concentration. However, after the Civil War the Herfindahl Index immediately returned to a very small number (around 0.05), indicating a very competitive market structure, and stayed like this until the 1870s. The market structure of merchants shipping from the East has a similar Herfindahl Index, though without the disruption of the Civil War. Overall, the analysis shows that merchants were very competitive, and this was unchanged by the telegraph. Therefore it does not seem plausible that a change in markups could be responsible for the observed change in the price difference after the telegraph.

In Table 1.C.3 I conduct the same robustness checks for the variance of the price difference. I use the squared deviation of the price difference from the mean before and after the telegraph, respectively, as dependent variable and regress them on a dummy that indicates the period after the telegraph.⁴⁴ Column (1) shows that the variance of the price difference falls significantly after the telegraph; the drop is around 90%. Excluding no trade periods explains one third of the drop, but the remaining fall is large and robust to all other robustness checks.

⁴³The digitized sample of the Bills of Trade has been generously provided by Graeme Milne for the years 1855, 1863 and 1870. I digitized the three months for 1866 to check for any change around the introduction of the telegraph.

⁴⁴In the online appendix I normalize the dependent variable by the average price difference before and after the telegraph connection. The findings are unchanged.

(2) Faster steam ships had a similar effect to that of the telegraph in the pre-telegraph period: They also reduced deviations from the Law of One Price. Similarly, in the post-telegraph period, temporary technical failures of the connection led to deviations from the Law of One Price.

The analysis so far has only used the one-time change in information frictions brought about by the telegraph to explain the deviations from the Law of One Price. However, the data provide much richer exogenous daily variation in information delays. In the pre-telegraph period, weather and wind accelerated or delayed mail steam ships, and in the post-telegraph period a few occasional technical breakdowns stopped the transatlantic communication temporarily. Figure **1.B.3** illustrates this variation. It shows how old the latest information from Liverpool is on a given day in New York (or in other words, how many days the last passage across the Atlantic took). Table **1.C.4** relates this variation in information delay to the variance of the price difference. Column (1) shows that deviations from the Law of One Price dropped significantly after the telegraph was established. Column (2) uses the exogenous variation in information delays. For each additional day that information takes to get from Liverpool to New York, the deviation from the Law of One Price increases by 24%. Column (3) only uses the within period variation by conditioning on the telegraph dummy, with similar results.

(3) New York prices respond to news from Liverpool.

The response of New York to news from Liverpool is best illustrated by an example that explains the large upwards spike in the price difference in Figure 1.B.5. Figure 1.B.7 zooms into this period and explains what happened in detail: On 29 and 30 September 1865 the market in Liverpool experienced increased demand for cotton from cotton spinners and millers. The *Liverpool Mercury* from that day writes that the market was "stimulated by the increasing firmness of the Manchester [yarn] market". At the same time a mistake in the estimation of cotton stock in Lancashire was detected, leading to a downwards correction. As a result, the Liverpool cotton price jumped up by 20% within two days, from 20 to 24 pence/pound. However, due to the delayed information transmission by mail ships, market participants in New York were not aware of this demand shock. The next steam ship, arriving in New York on 2 October, still carried the outdated price information from 23 September, a week before the demand shock. Only on 9 October the news of the demand shock arrived, causing a jump in the New York cotton price, as export demand increased. *The New York Times* reports an "unusually large quantity" of exports "under the favorable advices from England" on that day.

To study more systematically whether news from Liverpool drive New York prices, column (1) in Table 1.C.5 starts with a parsimonious specification and regresses the New York price on a given day on the latest known price from Liverpool using only data from the pre-telegraph period. This latest known Liverpool price was transmitted by steam ship, on average 10 days old and is denoted as "steam shipped" Liverpool price in the table. In order to account for autocorrelation in prices, a maximum likelihood estimation including three lags of the dependent variable is implemented.

The coefficient on this latest known Liverpool price is positive, indicating a systematic reaction of the New York price to news from Liverpool. Since prices are serially correlated, it is possible that this coefficient picks up something other than the "news" about Liverpool. Therefore column (2) includes a "counterfactual" price, the Liverpool price from the previous day that was unknown to New Yorkers before the telegraph connection was established. Reassuringly, this unknown price has no impact on New York prices, while the coefficient on the steam-shipped price remains the same. Columns (3) and (4) perform the corresponding analysis for the period after the telegraph was established. The "telegraphed" Liverpool price, on average one day old, now is the major driving force of the New York price, and the outdated price information that the steam ship would have brought, had the telegraph not been in place, does not matter anymore.

This parsimonious specification is the most efficient regression to demonstrate the changing relevance of Liverpool's prices on the New York market, as it uses the timing of information arrivals explicitly. As an alternative specification I run a vector autoregression using both prices, separately before and after the telegraph. Figure 1.B.8 shows that before the telegraph, only lags on the Liverpool price larger than 10 days are relevant for the New York price. After the telegraph, lags between 1-5 days are most relevant, in line with the distribution of information lags in Figure 1.B.4. Interestingly, the lags around 14 days are significant after the telegraph, because steam ships were used to ship longer market reports such as circulars.⁴⁵

(4) Market participants base their search for arbitrage opportunities on the latest news from Liverpool.

Figure 1.B.9 plots the difference of the New York price and the latest known Liverpool price (with the light gray line repeating the contemporaneous price difference from Figure 1.B.5). Interestingly, most of the largest price deviations disappear (except for the period of no trade in July 1865, which is shaded in the figure). Column (10) of Table 1.C.2 shows that the average price difference to the latest known Liverpool price falls after the telegraph connection were established. In contrast, column (10) of Table 1.C.3 shows that the variance of the price difference using the latest known Liverpool price shows only a small drop. This evidence indicates that market participants seem to arbitrage away the price difference between the current New York and the latest known, delayed Liverpool price, probably using it as a proxy for the price they expect for their exports.

(5) Exports respond to news about Liverpool prices.

The analysis so far has only considered prices as outcomes. But does information have real effects? In order for prices to equalize across marketplaces, goods must be moved. The detailed daily data on export flows can be used to understand whether the

⁴⁵Full VAR estimation results are available in the online appendix.

observed changes in the price difference are driven by equivalent changes in exports.

Table 1.C.6 uses a similar specification as Table 1.C.5 with exports as outcome and tests whether news about Liverpool prices affects exports. Column (1) uses only data from the year before the telegraph was in place and shows that news about an increase in the Liverpool price leads to increased exports in the pre-telegraph period. This news was brought by steam ship and was on average 10 days old. Column (2) conducts a placebo test and includes the unknown Liverpool price from the previous day, called "telegraphed" price. This counterfactual "news" does not have a significant impact on exports, as we would expect. Column (3) adds a linear time trend to control for a potential build up of supply after the American Civil War. Columns (4) to (6) conduct a similar analysis for the period after the telegraph. The news via telegraph about the Liverpool market again has a positive effect on exports, but the coefficient is not significant due to large standard errors. Column (5) includes the Liverpool price that market participants would have known had there been no telegraph. The news from the steam ship does not have a positive impact on exports, but the results are only suggestive as standard errors are large.⁴⁶ Column (6) allows for a linear time trend; the results remain unchanged. While the coefficient on the known Liverpool price is smaller, equality of the coefficients before and after the telegraph cannot be rejected.

(6) After the telegraph, exports are on average higher, and more volatile.

Row (1) in Table 1.C.7 shows that average daily exports from New York to Liverpool increased substantially after the telegraph cable was established: Average daily exports amount to 460 bales before the telegraph and increase by 170 bales after the telegraph, which is an increase of 37%. Row (1) in Table 1.C.8 shows that the variance of exports increases by even more. The increase in the variance after the telegraph of 0.33 represents an increase of 114% compared to the variance of 0.29 before the telegraph.

The remaining columns in Table 1.C.7 perform similar robustness checks for average exports to the ones implemented for the price difference. The increase in average exports after the telegraph connection cannot be explained by a fall in transport costs or fewer no trade periods. Can it be explained by an expanding cotton production after the American Civil War? Column (7) includes the cotton receipts at the New York exchange from the fields as a control, which does not affect the result.⁴⁷ In case this

⁴⁷Cotton production can only be adjusted with a time lag, that is when a new harvest cycle starts.

⁴⁶One might be concerned that the regression in Table **1.C.6** is invalid because prices are endogenously determined. However, the regression uses lagged Liverpool prices which are predetermined as far as current exports are concerned. However, if the underlying supply shocks are autocorrelated, the coefficients on Liverpool prices might be biased downwards because current supply shocks – which both increase current exports and are negatively correlated (via lagged supply shocks) with past Liverpool prices – are omitted. This downward bias might be stronger for the more recent "telegraphed" Liverpool price, explaining why the coefficient on the "telegraphed" Liverpool price in column (2) is smaller than the coefficient on the "steam shipped" Liverpool price. In columns (3) and (6) I include a linear time trend to control for buildup of cotton supply after the American Civil War. Coefficients in column (3) are larger than in column (2), indicating that a small downward bias was corrected. In any case, a larger downward bias for the telegraphed Liverpool price cannot explain why the relationship between the coefficients after the telegraph switches around, as the downward bias should still be stronger for the telegraphed than for the steam shipped Liverpool price.

variable is not picking up the full increase in production column (8) adds harvest year dummies to control for a potential gradual increase in cotton production across years. Again, the increase in exports after the telegraph connection remains significant.

The remaining columns in Table 1.C.8 conduct the same robustness checks for the variance of exports. Again, the increase in variance after the successful telegraph connection cannot be explained by these alternative hypotheses.

The following section provides an intuitive model about how information influences the behavior of exporters which yields predictions that are consistent with the presented "Stylized Facts".

1.5 Model of Information Frictions in International Trade

I add information frictions to a basic two-country trade model with storage (based on Coleman 2009; Williams and Wright 1991) by changing the information set of market participants and studying the impact of information frictions on trade flows, prices and welfare. The model is a partial equilibrium model that focuses on the role of intermediaries, who arbitrage away price differences across markets, and take producer supply and consumer demand as given.

The model mimics cotton trade in the 19th century. There is a centralized market in the supplying country, market *NY*, where producers sell their homogeneous good to intermediaries (merchants). The merchants export the good to another country, where they sell it to consumers in another centralized market, market *LIV*. Shipping goods from the supplying to the consuming country takes one period,⁴⁸ and shipping cost τ per unit shipped are incurred.⁴⁹ Profit maximizing merchants are perfectly competitive, therefore price takers, and risk neutral.⁵⁰

In country *LIV* there is a competitive storage industry.⁵¹ Profit maximizing storers can buy a certain quantity of the good, store it for one period, and sell it in the next period.⁵² Storage cost is θ per unit stored.⁵³

The good is elastically supplied by producers in country NY, as given by the linear

The increase in cotton exports is instead reflected in a reduction of New York stock levels. Only in the following harvest season can production adjust to increased exports, and this is what can be observed in the data.

⁴⁸The predictions of the model are also true with immediate shipment, k = 0. However, it is not very realistic to add information frictions to such a model, as it would assume that goods can be shipped instantaneously, while information cannot.

⁴⁹Transport costs of cotton consisted predominantly of unit cost (based on weight not value), as mentioned earlier. The numerical predictions are robust to allowing for ad valorem instead of unit trade cost.

⁵⁰Assuming risk averse instead of risk neutral merchants would only reinforce the predictions of the model. With better information, the "risk" of exporting is reduced as the variance of expected profits falls, which will lead to an increase in average exports.

⁵¹In the estimation of the welfare effect I will also allow for a storage industry in NY.

⁵²Holding the good for more than one period is equivalent to storers selling their stored quantity and buying it immediately back, to store it for another period and so on.

⁵³Storage cost of cotton was also based on weight rather than value. However, the behaviour of storage does not depend on whether one assumes additive or proportional storage cost (Williams and Wright 1991).

aggregate inverse supply function $p^{S}(q_{t}) = \bar{a}_{S} + b_{S}q_{t}$.⁵⁴ Aggregate consumer demand in *LIV* is stochastic and given by a linear inverse demand function with a stochastic, autocorrelated intercept a_{Dt} : $p_{t}^{D}(q_{t}) = a_{Dt} - b_{D}q_{t}$.⁵⁵ Representative merchants act as arbitrageurs across the two markets and choose exports x_{t} in order to maximize their expected profits:

$$\max_{x_t \ge 0} E\left[\left(p_{t+1}^{LIV} - p_t^{NY} - \tau\right) x_t\right]$$

where p_t^{NY} and p_t^{LIV} are the market prices in period *t* in markets *NY* and *LIV*, respectively. Representative storers act as arbitrageurs across time and choose storage quantity s_t to maximize their expected profits:⁵⁶

$$\max_{s_t \ge 0} E\left[\left(p_{t+1}^{LIV} - p_t^{LIV} - \theta\right)s_t\right]$$

A particular feature of commodity storage models is that it is not possible for the market as a whole to store negative quantities. Each individual stock holder can in principle store a "negative amount" (that is, selling "short") by borrowing the commodity from other storers, selling it on the spot market, buying the same amount of stock in the next period and returning it to the lender. However, this is not true for the market as a whole (Williams and Wright 1991).

Equilibrium conditions

Equilibrium conditions are given by the market clearing conditions in both markets and the first order conditions of merchants as well as storers. In market *NY*, supply is given by the supply function, while demand is given by the export demand of merchants. The market clearing condition in *NY* equalizes supply and demand:

$$p_t^{NY} = \bar{a}_S + b_S x_t \tag{1.5.1}$$

In market *LIV* supply is given by imports (which are equal to the amount of goods exported from *NY* in the previous period) plus storage from the previous period, while demand consists of demand by consumers and storers. The market clearing condition in *LIV* is therefore:

$$p_t^{LIV} = a_{Dt} - b_D \left(x_{t-1} + s_{t-1} - s_t \right)$$
(1.5.2)

Since negative exports are not possible, the first order condition of merchants is

⁵⁴Later I allow for stochastic supply. This adds another layer of information frictions that affect the information storers in *LIV* have about supply shocks in *NY*, and another source of welfare gain from the telegraph.

⁵⁵The demand function is similar to Evans and Harrigan (2005), however, with an autocorrelated demand process. For welfare estimation it is not necessary to assume a specific time process for the demand shocks. In the numerical solution I assume demand shocks follow a AR(1) process around mean \bar{a}_D .

 $[\]bar{a}_D$. ⁵⁶A detailed discussion of the storer's maximization problem and solution is provided by Williams and Wright (1991).

a mixed complementarity problem. At least one the two following inequalities must hold exactly:⁵⁷

$$E\left[p_{t+1}^{LIV}\right] \leqslant p_t^{NY} + \tau \quad \perp \quad x_t \ge 0 \tag{1.5.3}$$

Merchants choose the export quantity that equalizes the difference between expected prices in *LIV* in the next period and current prices in *NY*, subject to transport cost, except if expected prices in *LIV* are smaller than current prices in *NY* plus transport cost. In this case it is not optimal to export at all. In either case, expected profits of merchants are zero.

Storers face a similar mixed complementarity problem with respect to expected prices in *LIV* in the next period and current prices in *LIV*. Their no arbitrage condition is:

$$E\left[p_{t+1}^{LIV}\right] \leqslant p_t^{LIV} + \theta \quad \perp \quad s_t \ge 0 \tag{1.5.4}$$

Storers increase storage until expected prices in Liverpool in the next period are equal to today's prices in Liverpool plus storage cost, except if expected prices are too low to make profits from storage. Also the expected profits of storers is zero.

The equilibrium storage and export quantities in the model are determined by the market clearing conditions 1.5.1 and 1.5.2 together with the mixed complementarity conditions of merchants 1.5.3 and storers 1.5.4.

Information frictions

Decisions about storage and exports are based on expected prices in *LIV*. Agents form their price expectations based on the information available to them at the time when they make their purchasing decision at the market.⁵⁸ There are three different information regimes:

• *Delayed information (DI)*. Assume merchants are based in *NY* where they make their exporting decision by buying from suppliers.⁵⁹ In the delayed information

⁵⁷That is, equation 1.5.3 is equivalent to the following two conditions: either $E\left[p_{t+1}^{LIV}\right] - p_t^{NY} = \tau$ and $x_t \ge 0$; or $E\left[p_{t+1}^{LIV}\right] - p_t^{NY} \le \tau$ and $x_t = 0$. These conditions result form the maximization problem of the representative merchants. Note that the profit of representative merchants is linear in exports, as they take prices as given. If $E\left[p_{t+1}^{LIV}\right] - p_t^{NY} > \tau$ merchants would like to export an infinite amount, which is not an equilibrium. If $E\left[p_{t+1}^{LIV}\right] - p_t^{NY} < \tau$ merchants would choose zero exports, which is one of the complementarity conditions. If $E\left[p_{t+1}^{LIV}\right] - p_t^{NY} = \tau$, individual merchants are indifferent about which quantity to export. However, in equilibrium aggregate exports are determined by plugging in the supply and demand functions into the first order condition.

⁵⁸Ex-ante decisions of exporters are also modeled in Hummels and Schaur (2010, 2012). A reduction of shipping time in these papers is equivalent to a reduction in the forecast horizon. A reduction of shipping time is complementary to a reduction in information frictions as modeled in this paper; it is equivalent to a reduction in information delay only if there are no supply shocks.

⁵⁹In practice, merchants had representatives (usually family members) in both New York and Liverpool. If a merchant would have been based in Liverpool, he would have had to travel to New York (or communicate with New York) in order to export cotton from Liverpool to New York, and therefore have

regime they possess information about all shocks in LIV up to the previous period t-1 and have to forecast LIV prices in the following period t+1, the time when their exports can be sold in LIV. Similarly, storers are based in LIV where they make their storage decision. They therefore have information about demand shocks in LIV up to the current period t when forecasting LIV prices for period t+1.

- *Instantaneous information (II)*. All market participants are informed about demand shocks in *LIV* up to the current period *t* when forecasting expected prices in the following period *t*+1.
- *Perfect foresight (PF)*. Merchants and storers can foresee demand shocks in future periods.

The *delayed information regime* mimics the information frictions that were in place before the telegraph was established. One period can be interpreted as around 10 days, the time it takes for a steam ship to ship cotton (or at least samples of a cotton shipment) and information. Merchants have only delayed information from Liverpool which is on average 10 days old at the time when they make their exporting decision. At the time of exporting they need to forecast Liverpool prices 10 days into the future, which is when their shipment can be sold in Liverpool. On the other hand, storers in Liverpool know current market conditions in Liverpool when they make their storage decision. How much they know about New York is only relevant if supply is stochastic, a case which will be addressed later.

The *instantaneous information regime* mimics the situation after the transatlantic telegraph was established. Merchants have roughly real-time information from Liverpool when making their exporting decision. Due to the time delay in shipping, they still have to forecast Liverpool prices 10 days in the future.

The *perfect foresight regime* is unrealistic, but serves as a useful benchmark for intuition and the welfare analysis, as it maximizes aggregate welfare.

In the following I refer to both *DI* and *II* regimes as having "information frictions" compared to the *PF* regime. However, information frictions in the *II* regime are smaller than those in the *DI* regime. Therefore the introduction of the telegraph can be interpreted as a reduction in information frictions.

Characterization of the solution

The first order condition for merchants, equation 1.5.3, together with market clearing conditions given by equations 1.5.1 and 1.5.2 yield the following expression for exports:

$$x_t = \max\left\{\frac{E\left[a_{D,t+1} + b_D s_{t+1} - b_D s_t\right] - \bar{a}_S - \tau}{b_S + b_D}, 0\right\}$$
(1.5.5)

the same information as a merchant already based in New York. Therefore we can assume that merchants are based in New York only.

This expression is not yet a solution, as it depends on endogenous storage. Exports depend on expected demand shocks and expected storage change in period t+1, when the shipment arrives in *LIV*. In order to understand how exports should change after the telegraph connection, we need to compare exports across different information regimes. In the *PF* case, expectations are equal to realizations. However, under information frictions, a difference between expected and realized values arises, which is equal to the forecast error of demand and storage. While we are able to say something about the forecast error of demand, it is difficult to understand what happens to storage. Optimal storage is endogenous and a different function across the information regimes. Due to the non-linearity constraints in the no-arbitrage conditions, it is not possible to derive an analytical solution for the storage function (Deaton and Laroque 1996; Williams and Wright 1991).

Before proceeding to the numerical solutions of the model, the special case without storage (for example, when storage is prohibitively costly, or when the good is perishable) provides some intuition about the impact of information frictions.

Lemma. Suppose there is no possibility of storage and demand follows a stationary AR(1) process around mean \bar{a}_D with innovations $\epsilon_t \sim N(0, \sigma_D^2)$. Suppose $\frac{\bar{a}_D - \bar{a}_S - \tau}{b_S + b_D} > 0$, which means that there are positive exports at the average demand shock. Then, when switching from delayed to instantaneous information:⁶⁰

1. Average exports increase:

$$E\left[x_t^{DI}\right] \leqslant E\left[x_t^{II}\right]$$

2. Assume exports are always positive. Then, the variance of exports increases:

$$Var\left[x_t^{DI}\right] \leqslant Var\left[x_t^{II}\right]$$

3. The average price difference falls:

$$E\left[\operatorname{pdiff}_{t}^{DI}\right] \ge E\left[\operatorname{pdiff}_{t}^{II}\right]$$

4. The variance of the price difference falls:

$$Var\left[\mathsf{pdiff}_{t}^{DI}\right] \ge Var\left[\mathsf{pdiff}_{t}^{II}\right]$$

Proof. In appendix.

The intuition for this result is as follows: From equation 1.5.5 (ignoring storage) we see that exports are a function of expected demand shocks. Depending on the information regime, expected demand shocks are based on more (in the *II* regime) or less (in the *DI* regime) information. The variance of any conditional expectation is larger, the more information it is based on. The extreme cases are the *PF* regime, in which the variance

⁶⁰In case of a white noise process, information is irrelevant, and the predictions hold with equality.

of exports is a function of the variance of the underlying demand shock; and "no information", when exports are a function of the unconditional mean of demand which has a zero variance, so exports are always the same. In between these two extreme cases, the variance of exports follows the variance of the expected demand shocks, which increases the more information is available. This gives the intuition for why the variance of exports increases when switching from the *DI* to the *II* regime, which is point 2 of the Lemma.

Given that the variance of expected demand shocks and therefore exports is larger the more information is known, this means that merchants underestimate high demand shocks and overestimate low demand shocks in the *DI* regime compared to the *II* regime. When they underestimate states of high demand, they export less in the *DI* compared to the *II* regime. When they overestimate low demand, there is no difference across information regimes as with low demand it is not profitable to export at all. This is why the first effect dominates and average exports are larger in the *II* compared to the *DI* regime. Average exports are higher in the *II* regime, because periods of high demand lead to appropriately high exports, point 1 of the Lemma.

Merchants equalize expected prices across the countries, and the resulting price difference equals the forecast error of merchants. If merchants were not making any forecast error, which would happen under the *PF* regime, the spatial price difference would be zero, and the no arbitrage condition would hold. With information frictions, the no arbitrage condition holds only in conditional expectations, and merchants make a forecast error depending on the information they have. The volatility of forecast errors falls when more information becomes available (a result well established in the finance literature), which explains why the price difference falls when switching from the *DI* to the *II* regime; which is point 4 of the Lemma.⁶¹

Point 3 of the Lemma states that the average price difference falls after switching to the *II* regime. This holds for the same reason that average exports increase: Under the *DI* regime, positive demand shocks are systematically underestimated, leading to high prices in Liverpool and a large price difference as exports are restricted. These positive price differences are eliminated under the *II* regime, as exports are high enough. Note that this does not mean that merchants were making profits under the *DI* regime, as high ex-post profits in cases when demand was higher than expected were offset by equivalently high losses in cases when demand was smaller than expected. Neither does this mean that merchants make losses under the *II* regime, as they avoid negative price differences by not exporting at all (when they would have exported under the *DI* regime).

As we shall see below, the predictions from the Lemma also hold for the case with storage. Storage is able to dampen the inefficiencies from information frictions: When merchants overestimate demand and ship too much, part of the imports can be stored for the future and consumed in cases when merchants underestimate demand and ship too little. As a result, in a model with storage, prices will fluctuate less than in a model

⁶¹Point 4 of the Lemma also holds if the demand shock is a random walk.

without storage. However, storage is not sufficient to fully eliminate the inefficiencies from information frictions. Perfect smoothing of prices over time cannot be achieved: First, storage is costly. Second, storage cannot be negative, and there is always a positive probability that long periods of particularly high demand will run down inventories. A finite stock can never fully insure against stock outs and its accompanying price spikes (Townsend 1977). As Williams and Wright (1991, p. 159) state: "Storage is asymmetric - able to support a glut but not alleviate every shortage".⁶²

Numerical solution

The commodity storage literature has provided numerical solution approaches for instantaneous information regimes (Coleman 2009; Williams and Wright 1991; Deaton and Laroque 1996). This paper adds information frictions in the form of delayed information to the model to obtain predictions for the effect of the telegraph.

The central task of the problem is the numerical solution of the control functions for storage and exports as a function of state variables. In the instantaneous information case, there are two types of state variables: "Stock on hand" m_t is the stock available at the beginning of each period, which consists of the sum of quantities stored in the previous period and arriving imports (equal to *NY*'s exports from the previous period), $m_t = s_{t-1} + x_{t-1}$; and the current realization of the demand shock, a_{Dt} . I approximate $x_t (a_{Dt}, m_t)$ and $s_t (a_{Dt}, m_t)$ simultaneously over a grid of state variables by checking the first order condition of merchants and storers for a guess for the control functions, and by updating the guess in every step. The approximation algorithm has converged when all the first order conditions in all grid points are satisfied up to a certain precision.

Then I add information frictions to the model. The delayed information regime requires a different set up of the problem, as storage and exports depend on different state variables because expectations are formed differently. In the delayed information regime merchants know only lagged demand shocks, so exports are a function of lagged demand and contemporaneous stock on hand (because this itself is a function of lagged storage and exports): $x_t (a_{D,t-1}, m_t)$. On the other hand, storage in Liverpool continues to be a function of contemporaneous demand, $s_t (a_{Dt}, m_t)$.

In the instantaneous information regime I approximate two control functions as a function of two state vectors: a_{Dt} and m_t . In the delayed information regime the number of state variables increases to three: a_{Dt} , $a_{D,t-1}$ and m_t . For a more realistic model that matches *NY* prices more closely, I extend the model to allow for stochastic supply shocks. Then the export decision also depends on the supply shocks: $x_t (a_{St}, a_{D,t-1}, m_t)$ in the *II* regime, and $x_t (a_{St}, a_{D,t-1}, m_t)$ in the *DI* regime. Storage in *LIV* depends on the

⁶²It is interesting to note that a problem with speculators engaging in futures trading leads to the same market-level equilibrium conditions as a model with explicit modeling of storage, as long as at least some speculators have rational expectations and are risk neutral. If there were a mix of stock holders and speculators, and stock holders were not risk neutral, as long as some speculators are risk neutral and have rational expectations, these equilibrium conditions for the market still hold. Expected prices in the equilibrium conditions 1.5.3 and 1.5.4 are equivalent to the prices of futures. While futures trading might lead to a different allocation of risk across market participants (from risk averse to less risk averse or risk neutral), the aggregate properties of a model with storage and one with futures trading are the same.
information about supply shocks: $s_t (a_{St}, a_{Dt}, m_t)$ in the *II* regime, and $s_t (a_{S,t-1}, a_{Dt}, m_t)$ in the *DI* regime. In total the number of state variables in the *DI* regime increases to five: a_{Dt} , $a_{D,t-1}$, a_{St} , $a_{S,t-1}$ and m_t .⁶³

What are the characteristics of the optimal storage functions? Because of the nonnegativity constraint, the storage function is zero until a "kink line" given by a combination of critical values for demand shocks and stock on hand. Beyond the kink line storage is increasing in stock on hand and decreasing in demand shocks. The slope of the storage function (the propensity to store) is everywhere less than one, and the behavior is non-linear.⁶⁴ The export function has a similar non-linear behavior, because it depends on storage: It is also zero up to a certain "kink line", beyond which it is increasing in demand shocks and decreasing in stock on hand. With supply shocks, the non-linear behavior extends to a third dimension: Exports and storage increase in the supply shock, again after a certain "kink plane". The behavior of the control functions with delayed information is qualitatively similar as a function of lagged instead of contemporaneous shocks. However, the slopes are different as they reflect different expectation forming about the future.

Estimation of parameters

The model depends on ten parameters that need to be estimated: mean \bar{a}_D , variance σ_D^2 and autocorrelation ρ_D of the stochastic demand process; mean \bar{a}_S , variance σ_S^2 and autocorrelation ρ_S of the stochastic supply process; the slope of the demand function b_D , the slope of the supply function b_S ; transport cost τ and storage cost θ .

First I focus on estimating the slopes of the demand and supply curves. These parameters will also be used for the estimation of welfare further below. Given these slopes, the underlying demand and supply shock processes can be backed out using data on prices, trade and storage. The mean, variance and autocorrelation coefficient of the shock processes are then estimates of the corresponding parameters. The remaining parameters of the model are transport cost, which is based on the data on freight cost as described earlier, and storage cost of cotton, which are obtained from the historical accounting statements of a merchant in Boyle (1934).

Estimation of the demand curve

Estimating the supply and demand functions is usually tricky, as quantities and prices are determined contemporaneously and finding a valid instrument is difficult. I propose a new identification strategy: Since shipping takes time, exports are predetermined once they arrive in Liverpool, breaking the simultaneity problem for the case of i.i.d. shocks. For the case of autocorrelated shocks and positive storage I use the model

⁶³As the number of calculations needed increase exponentially with the number of state variables (Williams and Wright 1991, p. 57), it is not possible to numerically solve a daily representation of the model. In this case the information lag would be around 10 days and I would need to keep track of 20 state variables (for a two-way trip; per market), which is computationally not feasible.

⁶⁴Further discussion is provided in Williams and Wright (1991).

to control appropriately for the endogenous part of the shocks, yielding identified regression equations.

The demand curve in Liverpool on a specific day t + k, where k indicates the time (in days) a shipment takes to get from New York to Liverpool, is determined by the realization of the demand shock on that day, $a_{D,t+k}$, the imports arriving in Liverpool on that day which are equivalent to exports from New York k days earlier, x_t , and net take-up or release of stock from storage on that day, Δs_{t+k} :

$$p_{t+k}^{LIV} = a_{D,t+k} - b_D \left(x_t - \triangle s_{t+k} \right)$$

Daily prices in Liverpool as well as daily imports can be observed. The main identification problem is the unobserved demand shock that is positively correlated with change in stock and exports. Note that exports are actually a function of lagged demand shocks, which are correlated with demand shocks at t+k only via the autocorrelation of the demand shock. My identification strategy will exploit this fact by modeling this dependence explicitly.

Assuming demand follows an AR(1) process around mean \bar{a}_D , $a_{Dt} - \bar{a}_D = \rho(a_{D,t-1} - \bar{a}_D) + \epsilon_t$ with $\epsilon_t \sim N(0, \sigma^2)$, we can express the demand shock in period t + k in terms of the demand shock in period t - l, where l denotes the information delay between Liverpool and New York, and the sum of demand innovations between t - l and t + k:

$$a_{D,t+k} = (1 - \rho^{k+l})\bar{a}_D + \rho^{k+k}a_{D,t-l} + \sum_{i=0}^{k+l-1} \rho^i \epsilon_{D,t+k-i}$$

We can use the lagged demand function to control for the lagged demand shock, as $p_{t-l}^{LIV} = a_{D,t-l} - b_D (x_{t-k-l} - \triangle s_{t-l}^{LIV})$. This results in an equation where all of the regressors except change in stock $\triangle s_{t+k}^{LIV}$ are uncorrelated with the unobserved demand shocks. However, current imports x_t can be used as an instrument for $x_t - \triangle s_{t+k}^{LIV}$. Data on stock in Liverpool is available only at weekly intervals, so I distribute the the weekly change equally across the day of the week, which introduces a measurement error, which is also addressed by the instrumental variables strategy. Table 1.C.9 shows the results of estimating the following equation:

$$p_{t+k}^{LIV} = \beta_0 + \beta_1 \left(x_t - \triangle s_{t+k}^{LIV} \right) + \beta_2 p_{t-l}^{LIV} + \beta_3 \left(x_{t-k-l} - \triangle s_{t-l}^{LIV} \right) + \sum_{i=0}^{k+l-1} \rho^i \epsilon_{D,t+k-i}$$

Column (1) shows the OLS results and column (2) shows the IV results. The first stage is strong, as indicated by the F-statistics of 125. The instrument addresses both the correlation of stock changes and demand shocks in the error as well as measurement error in the stock changes. The latter seems to dominate as the OLS estimate is biased towards zero. The sign of the lagged Liverpool price is positive and less than 1 as expected. According to the model $\beta_1\beta_2 - \beta_3 = 0$, which cannot be rejected. Column (3) uses a more efficient estimation method by imposing this restriction, applying nonlinear least squares estimation. Column (4) implements the instrumental variable estimation

using a control function approach, which again corrects for the measurement error in stock changes. In column (5) and (6) the non-linear IV specification are estimated separately for the period before and after telegraph, but this makes little difference to the estimates.

The last row in Table 1.C.9 computes the demand elasticity at mean values of prices and quantities. The resulting demand elasticities seem rather high when comparing them to estimates of demand elasticity of cotton in 19th century in the literature (Irwin 2003), which range between 1.7 and 2.3. Note however, that the estimates in the literature are based on yearly instead of daily data. Daily demand elasticities are much higher because they take into account the willingness of consumers (or cotton millers) to substitute consumption across time, which is easier across short periods compared to long periods. To empirically validate this argument, I run the demand estimation on different aggregation periods of my data. Aggregating the data reduces the demand elasticity strongly. For example, for 3-monthly data the demand elasticity is as low as -6 (see online appendix for details on aggregation patterns).

Estimation of the supply curve

The slope of the supply function is estimated in a similar way. In order to better match the data, the supply function given by equation 1.5.1 is extended by allowing both for supply shocks and for storage in New York:

$$p_t^{NY} = b_S \left(\triangle s_t^{NY} + x_t \right) + a_{St}$$

Again, the problems in estimating this equation are two-fold: First, New York stock data is only available at weekly intervals, so I distribute the weekly change equally across days, introducing measurement error. Second, exports and stock changes are correlated with current supply shocks. I add a dummy for the harvest year and include a quadratic in the day of the harvest year to model supply fluctuations, but as this cannot fully address endogeneity concerns, I pursue an instrumental variables approach for the estimation. In column (2) of Table 1.C.10 I use known prices from Liverpool p_{t-1}^{LIV} as instrument for the sum of export and stock changes, $\triangle s_t^{NY} + x_t$. The first stage is strong, as information about the latest prices from Liverpool influence exports and stock changes positively. If supply shocks are correlated, however, lagged Liverpool prices might reflect lagged supply shocks, and not be exogenous. Therefore I use implied demand shocks $a_{D,t-l} = p_{t-l}^{LIV} + b_D (x_{t-k-l} - \triangle s_{t-l}^{LIV})$ using the estimated slope of the demand function from the previous section as instrument for exports and stock changes in column (3). In column (4) lagged Liverpool prices are again used as instrument, but here $x_{t-k-l} - \triangle s_{t-l}^{LIV}$ is added as a control, leaving only demand variations in the instrument.

The estimates yield a estimate of around 1.7 in all specifications after eliminating measurement error in the OLS estimation, also when I estimate the equation separately for the periods before and after the telegraph. Again, the equivalent supply elasticities

are larger than the estimates based on yearly data mentioned in the literature which are between 1 and 2 (see Irwin 2003 for a review), as is expected when considering the substitution of supply across short time periods. When repeating the analysis of the supply estimation across data with an increasing aggregation horizon, the supply elasticity falls considerably and converges towards the estimates in the literature (see online appendix for details).

Delayed information as counterfactual

With the estimates of the slopes of the demand and supply functions I reconstruct the demand and supply processes from the data of the post-telegraph period and estimate the remaining supply and demand parameters. Table 1.C.11 gives an overview of all estimated parameters. The estimated AR(1) processes fit the data quite well, as it is not possible to reject white noise of the innovations in the supply and demand process (using a Portmanteau white noise test). The demand shocks are more autocorrelated than the supply shocks, while the supply shocks have a higher volatility than the demand shocks.

Together with an estimate for transport and storage cost I numerically solve for both the instantaneous (to which the parameters are calibrated) and delayed information regime (counterfactual analysis).

I use the counterfactual delayed information regime on the calibrated model to predict the effect of the telegraph on the data. Table 1.C.12 shows that the qualitative predictions of the model match the empirical section: The model predicts a fall in the average price difference between Liverpool and New York, and an increase in average exports. Similarly, the model predicts a fall in the variance of the price difference, and an increase in the variance of exports. This finding is robust to wide ranges of the parameter space. For example, the different columns vary storage cost from zero to prohibitively high storage cost. The panels below use the lower and upper 95% confidence interval for estimates of the slope of the supply and the demand functions.

Quantitatively, the model predicts the largest change for the fall in the standard deviation of the price difference, which is also the largest drop in the data. My preferred estimates with storage cost as given by Boyle (1934) are shaded in gray and predict a drop in the standard deviation of the price difference of around 40-60%, which is close to the drop in the data of 70%. The second largest change in the model is the increase in the standard deviation of exports, which again is also the second largest change in the data. Here, however, the model cannot fully predict the change. A possible explanation could be that storage cost are underestimated. Since storage and exports are complements in balancing out information frictions, higher storage costs imply that exports react more strongly. However, even with prohibitively high storage cost the model can only predict a increase in the standard deviation of exports deviation of exports of 6%.

The changes in the average price difference and in average exports are also larger in the data than in the model. A potential reason for this could be that in reality merchants were risk averse rather than risk neutral. An extension of the model with risk averse merchants is likely to predict larger changes in average exports and price difference, because the higher uncertainty in the delayed information regime should lead risk averse merchants to export less due to higher uncertainty.

1.6 Welfare Gains from the Telegraph

What are the welfare gains from reduced information frictions? This section shows that a lower bound of the deadweight loss (DWL) from information frictions is a function of only three parameters: the squared observed price difference across markets, and the slopes of the supply and demand functions presented in Section 1.5. This is an analytical result and does not rely on the numerical solutions obtained in the previous section, nor on assuming a specific time series process for the demand and supply shocks.

For intuition, consider a specific export transaction. The welfare arising from this transaction is the sum of immediate producer surplus, immediate consumer surplus, immediate merchant surplus, as well as the present value of future social surplus from the part of exports that is stored and not immediately consumed.⁶⁵ The red area in Figure 1.B.10 corresponds to the immediate producer surplus of exports. If a part of exports is stored, immediate consumer surplus corresponds to the blue area. The net present value of the social surplus from the quantity stored is given by the green area between the market demand curve and the price. The social surplus from storage is the sum of positive future consumer and negative future producer surplus, the expected surplus of storers is zero.

Note that Figure 1.B.10 ignores stock stored from the previous period (current stock on hand is equal to current imports only). This is because I measure welfare for each specific export transaction at the time when exporting occurs with the net present value (NPV) of social surplus. The stock in storage has been exported in previous periods and its welfare has already been accounted for by the net social surplus from storage then.

Figure 1.B.10 shows the perfect foresight equilibrium (*PF*), which I use as reference case to measure deadweight loss. In *PF*, merchants choose exports at the intersection of the lagged supply curve and the market demand curve, the price in Liverpool is equal to the lagged price in New York, and merchants make no profits.

In contrast, Figure 1.B.11 illustrates the case when there are information frictions⁶⁶

⁶⁵This discussion of welfare follows Williams and Wright (1991, p. 350). The current consumer surplus understates total consumer surplus, because positive stock raises current prices for consumers and reduces current consumption. However, eventually the stock is going to be consumed, and in that period prices for consumers will fall and consumption will increase. Similarly, current producer surplus overstates total producer surplus, because positive stock increases the current selling prices for producers and increases current production, whereas upon consumption it reduces prices and production.

⁶⁶The term "information frictions" in this paper is used both for delayed information (equivalent to before the telegraph) and instantaneous information (equivalent to after the telegraph), but in the former case information frictions are larger. In order to measure deadweight loss from reduced information frictions, I compare the deadweight loss of having delayed information (as opposed to having perfect foresight) to the deadweight loss of having instantaneous information (as opposed to having perfect foresight).

and merchants overestimate demand in Liverpool. In this case, exports are larger than *PF* exports, and prices are not equalized across markets. Merchants make a loss, and some of the inefficiently high export is stored for the future. However, optimal storage is not large enough to raise the Liverpool price to the level of the undistorted price in Figure 1.B.10. The size of the deadweight loss is given by the orange area. An equivalent deadweight loss triangle arises from an underestimation of demand.

Theorem. The deadweight loss from information frictions for a specific export transaction $x_{t-1} > 0$ is bounded from below by <u>DWL</u>:

$$DWL(x_{t-1}) \ge \frac{\left(p_t^{LIV} - p_{t-1}^{NY}\right)^2}{2(b_D + b_S)} =: \underline{DWL}$$

That is, the spatial price difference $p_t^{LIV} - p_{t-1}^{NY}$, the slope of the demand curve b_D and the slope of the supply curve b_S are sufficient statistics for the lower bound of the deadweight loss from information frictions.

Proof. In appendix.

It is not surprising that the welfare loss from information frictions is a function of the price difference. The Law of One Price states that any spatial price difference gets arbitraged away if agents are fully informed (due to the shipping lag, full information in this case would require foresight of market conditions upon arrival of shipments). The literature on the LOP therefore interprets observed price difference as a measure of the underlying market frictions and its associated deadweight loss. The theorem makes the relationship between deviations from the Law of One Price and welfare explicit.

For the actual welfare estimation I use the daily equivalent of the expression given in the theorem:

$$\underline{DWL}(x_{t-k}) = \frac{\left(p_t^{LIV} - p_{t-k}^{NY} - \tau_{t-k}\right)^2}{2(b_D + b_S)}$$

Figure 1.B.12 illustrates the observed price distortion $p_t^{LIV} - p_{t-k}^{NY} - \tau_{t-k}$, where *k* denotes the actual shipping time in days. The spatial price difference falls dramatically after the telegraph gets introduced. The fall in the price distortion after the telegraph is equivalent on average to a roughly 6% ad valorem tariff.⁶⁷ The largest price distortions during the pre-telegraph period were equivalent up to a 50% ad valorem tariff. For comparison, note that the average US tariff that was abolished during NAFTA in 1994 was 3%, while the highest abolished tariff was 12%, for textile trading with Mexico (Caliendo and Parro 2012).

In order to translate these price distortions into welfare effects, I use estimates of the slope of the supply and demand curves as described previously. Combining

⁶⁷The equivalent ad valorem tariff of the distortion is calculated for each day as the absolute price difference minus transport cost in percent of the lagged New York price p_{t-k}^{NY} . The average tariff equivalent of 6% is equal to the difference in the average of this measure between the pre- and post-telegraph period. Days with no trade are excluded from this calculation.

the estimates for the slope of the supply and the demand function with the observed price difference, Table 1.C.13 reports the welfare loss due to delayed or instantaneous information as compared to perfect foresight. The difference in welfare loss can be attributed to the telegraph, and corresponds to 8% of the annual export value of American cotton from New York to Great Britain in the data which is around 10 million pounds in 1866.⁶⁸ I construct a confidence interval for the welfare loss based on the confidence intervals for the estimates of the slopes of the supply and demand functions. The confidence interval of the welfare gains from the telegraph ranges from 5% to 22%.

The 8% welfare gain can be divided into a 6.7% efficiency gain from reducing the variance of the price difference (due to within year reallocation), and 1.7% of efficiency gain from reducing the average price difference (due to increased average trade). If I exclude the anecdotal episode with the especially large demand shock described earlier from the welfare calculation, the efficiency gain is 6.6%.

This is somewhat larger than previous estimates in Ejrnaes and Persson (2010) who estimate the welfare gains from the transatlantic telegraph to be around 0.5-3% of trade value. Their estimate is based on weekly data which averages out some of the variation in the data. Furthermore, they rely on demand and supply elasticities from the literature which are estimated over yearly and not weekly time horizons because they do not observe trade flows to estimate the reaction of exports directly.

How is the welfare gain distributed across producers, consumers and merchants? The surplus of merchants is zero under all information regimes, as unexpected gains and losses average out. The distribution of the additional surplus across consumers and producers depends on the slopes of the supply and demand functions. Since the estimated slope of the supply function is larger than the estimated slope of the demand function, producers gain more from the telegraph than consumers.

1.7 Conclusions

This paper exploits a clean historical experiment to understand the impact of information frictions on the Law Of Once Price and trade: the establishment of the transatlantic telegraph cable, connecting the United States and Great Britain in 1866. This episode provides a unique setting for studying information frictions. On one hand, it provides a dramatic and exogenous reduction in information frictions, as the information transmission times between these two countries fell unexpectedly from around ten days to only one day. On the other hand, a rich data set based on historical newspapers includes high-frequency data not only on prices, but also on trade and information flows.

This setting allows me to contribute to the literature in several ways. First of all, it is possible to *identify* and *measure* the impact of information, which is usually endogenous, complex and unobserved. This paper shows that a fall in information frictions causes

⁶⁸Note that the total exports from the United States to Europe is around three times as much. In reality, the welfare benefit of the transatlantic telegraph applies to all of transatlantic cotton trade (and potentially other goods as well).

better adherence to the Law of One Price. The average price difference between New York and Liverpool falls by 35%, and its standard deviation falls by 73%. This reduction in price distortions is equivalent to abolishing a roughly 6% ad valorem trade tariff.

Second, this paper shows that information frictions have real effects and are not just a reallocation of profits across market participants, because exports *respond* to information. After the telegraph, average trade flows increase and become more volatile. The model explains that this is the case because exports follow expected demand, conditional on information. More information makes expected demand more volatile, which explains why we observe more volatile exports after the telegraph. However, this effect is asymmetric, as there is a cutoff when it is not profitable to export at all. More information increases average exports because there are more incidents with high expected demand and therefore large exports.

The third contribution of the paper lies in estimating the welfare gains from reducing information frictions. Better information helps merchants to better forecast future demand, resulting in a more efficient alignment of supply and demand across countries. This is reflected in the better adherence to the Law of One Price. In order to translate the reduced price distortions into welfare, one needs to estimate supply and demand elasticities, which is usually difficult due to simultaneously determined prices and quantities. This paper uses a novel identification strategy that exploits the fact that exports are predetermined once they arrive in Liverpool (since shipping takes time) and controls adequately for the possibility of storage. Overall, the welfare gains from the telegraph are estimated to be around 8% of the annual export value.

The historical example of the transatlantic telegraph can be generalized to any setting in which exporters or producers have to make ex-ante decisions about production and/or exporting and face uncertainty about demand. In this setting, information technologies can improve the ability of firms to forecast demand. The forecast error of exporters becomes smaller and less volatile the better the available information is. This leads firms to decide on production or export quantities that are better matched with consumer demand, and therefore reduces deadweight loss. The model also points out that the benefits from information technology are larger, the larger the underlying volatility of demand, and the larger storage cost.

Identifying and reacting to demand changes is still critical in today's world. Demand fluctuates more rapidly and widely than it used to as new trends appear and disseminate via social media networks. Global supply chains and outsourced stages of production make it more difficult to communicate demand changes across the different firms involved in the production process. Newly emerging information technologies such as the real-time analysis of "Big Data" have the potential to impact trade in a similar, but probably even more drastic, way as the telegraph. The smart phone era has generated an enormous amount of real-time data on consumer behavior. The technologies for analyzing these large amounts of data are still being developed, but they have the potential to provide firms with much more accurate demand forecasts. The model in this paper can be used in this context. For example, it would predict that industries with a more volatile demand or higher storage cost should adopt and develop these technologies earlier than other industries. The model can also be used to assess the welfare effects of these technologies, and compare them with their cost.

1.A Proofs

Proof of Lemma

Exports are given by equation 1.5.5. Assume first that exports are always positive, so exports x_t are equal to uncensored exports $\tilde{x}_t = \frac{E[a_{D,t+1}] - \bar{a}_S - \tau}{b_S + b_D}$ (the assumption of point 2 of the Lemma). Exports in both regimes differ only in the way how expectations about future demand shocks are formed. In the instantaneous information regime the information set I_{II} includes everything up to period t, while the information set in the delayed information regime I_{DI} includes only information up to period t - 1: $I_{DI} \subset I_{II}$. By the Law of Iterated Expectations, average exports are the same in both information regimes:

$$E\left[x_{t}^{DI}\right] = E\left[\frac{E_{t-1}\left[a_{D,t+1}\right] - \bar{a}_{S} - \tau}{b_{S} + b_{D}}\right] = \frac{\bar{a}_{D} - \bar{a}_{S} - \tau}{b_{S} + b_{D}}$$
$$= E\left[\frac{E_{t}\left[a_{D,t+1}\right] - \bar{a}_{S} - \tau}{b_{S} + b_{D}}\right] = E\left[x_{t}^{II}\right]$$

The variance of exports is a function of the variance of expected demand shocks, conditional on the respective information set:

$$Var\left[x_{t}^{DI}\right] = \frac{Var\left[E_{t-1}\left[a_{D,t+1}\right]\right]}{\left(b_{S}+b_{D}\right)^{2}}$$
$$Var\left[x_{t}^{II}\right] = \frac{Var\left[E_{t}\left[a_{D,t+1}\right]\right]}{\left(b_{S}+b_{D}\right)^{2}}$$

Applying Jensen's inequality to the function $E_t[a_{D,t+1}]$, conditional on information in *t*-1:

$$(E_{t-1}[E_t[a_{D,t+1}]])^2 \leq E_{t-1}[(E_t[a_{D,t+1}])^2]$$

Taking the unconditional expectation:

$$E\left[\left(E_{t-1}\left[a_{D,t+1}\right]\right)^{2}\right] \leqslant E\left[\left(E_{t}\left[a_{D,t+1}\right]\right)^{2}\right]$$

The variance of the conditional expectations are:

$$Var [E_{t-1} [a_{D,t+1}]] = E [(E_{t-1} [a_{D,t+1}])^2] - (E [E_{t-1} [a_{D,t+1}]])^2$$

= $E [(E_{t-1} [a_{D,t+1}])^2] - (E [a_{D,t+1}])^2$

$$Var[E_t[a_{D,t+1}]] = E\left[(E_t[a_{D,t+1}])^2 \right] - (E[E_t[a_{D,t+1}]])^2$$
$$= E\left[(E_t[a_{D,t+1}])^2 \right] - (E[a_{D,t+1}])^2$$

Combining the last three equations shows that the variance of a conditional expected value increases if the information set increases: $Var[E_{t-1}[a_{D,t+1}]] \leq Var[E_t[a_{D,t+1}]]$. It follows directly that $Var[x_t^{DI}] \leq Var[x_t^{II}]$, which proves the second point of the Lemma. Note that this part of the proof is more general as it does not assume any specific process or distributional assumptions about the demand shocks.

Now assume that exports have to be positive. From the above analysis we know that uncensored exports have the same mean in both information regimes, but the variance of uncensored exports is smaller in the delayed information regime: $\tilde{x}_t^{II} \sim N\left(\frac{\tilde{a}_D - \tilde{a}_S - \tau}{b_S + b_D}, \frac{Var(E_t[a_{Dt}])}{(b_S + b_D)^2}\right)$ and $\tilde{x}_t^{DI} \sim N\left(\frac{\tilde{a}_D - \tilde{a}_S - \tau}{b_S + b_D}, \frac{Var(E_{t-1}[a_{Dt}])}{(b_S + b_D)^2}\right)$. Denoting $\tilde{\mu} := E[\tilde{x}_t]$ and $\tilde{\sigma}^2 := Var[\tilde{x}_t]$, average exports are given by $E[x_t] = \Phi\left(\frac{\tilde{\mu}}{\tilde{\sigma}}\right)\tilde{\mu} + \tilde{\sigma}\phi\left(\frac{\tilde{\mu}}{\tilde{\sigma}}\right)$ (Greene 2003). A change from *DI* to *II* increases the variance of censored exports $\tilde{\sigma}^2$, and this increases average exports, which proves the first point of the Lemma:

$$\begin{array}{ll} \frac{\partial E\left[x_{t}\right]}{\partial \tilde{\sigma}} & = & \phi\left(\frac{\tilde{\mu}}{\tilde{\sigma}}\right) \cdot \left(-\frac{\tilde{\mu}}{\tilde{\sigma}^{2}}\right) \cdot \tilde{\mu} + \phi\left(\frac{\tilde{\mu}}{\tilde{\sigma}}\right) + \tilde{\sigma}\phi'\left(\frac{\tilde{\mu}}{\tilde{\sigma}}\right) \left(-\frac{\tilde{\mu}}{\tilde{\sigma}^{2}}\right) \\ & = & \phi\left(\frac{\tilde{\mu}}{\tilde{\sigma}}\right) > 0 \end{array}$$

The average price difference $pdiff_{t+1} = p_{t+1}^{LIV} - p_t^{NY} - \tau$ is (plugging in the solution for exports):

$$E\left[\text{pdiff}_{t+1}\right] = E\left[a_{D,t+1} - E_{t-1}\left[a_{D,t+1}\right] | x_t > 0\right] Prob\left[x_t > 0\right] \\ + E\left[a_{D,t+1} - \bar{a}_S - \tau | x_t = 0\right] Prob\left[x_t = 0\right]$$

The first term is zero under both information regimes. For the second term consider that

$$E\left[a_{D,t+1} - \bar{a}_{S} - \tau \mid x_{t}^{II} = 0\right] = \bar{a}_{D} + \rho\sigma \frac{\phi\left(\frac{\bar{a}_{S} + \tau - \bar{a}_{D}}{\rho\sigma}\right)}{1 - \Phi\left(\frac{\bar{a}_{S} + \tau - \bar{a}_{D}}{\rho\sigma}\right)} - \bar{a}_{S} - \tau > \bar{a}_{D} + \rho^{2}\sigma \frac{\phi\left(\frac{\bar{a}_{S} + \tau - \bar{a}_{D}}{\rho^{2}\sigma}\right)}{1 - \Phi\left(\frac{\bar{a}_{S} + \tau - \bar{a}_{D}}{\rho^{2}\sigma}\right)} - \bar{a}_{S} - \tau = E\left[a_{D,t+1} - \bar{a}_{S} - \tau \mid x_{t}^{DI} = 0\right]$$

and

$$Prob\left[x_t^{II} = 0\right] = Prob\left[a_{Dt} < \frac{\bar{a}_S + \tau - \bar{a}_D}{\rho} + \bar{a}_D\right] > Prob\left[a_{D,t-1} < \frac{\bar{a}_S + \tau - \bar{a}_D}{\rho^2} + \bar{a}_D\right] = Prob\left[x_t^{DI} = 0\right]$$

From this, the third part of the proof follows:

$$E\left[\operatorname{pdiff}_{t+1}^{DI}\right] \ge E\left[\operatorname{pdiff}_{t+1}^{II}\right]$$

From the variance decomposition property it follows that the variance of the price difference is equal to the variance of the forecast error when exports are positive:

$$Var \left[pdiff_{t+1} \right] = E \left[V \left[a_{D,t+1} - E \left[a_{D,t+1} \right] | x_t > 0 \right] \right] + Var \left[E \left[a_{D,t+1} - E \left[a_{D,t+1} \right] | x_t > 0 \right] \right] = V \left[a_{D,t+1} - E \left[a_{D,t+1} \right] | x_t > 0 \right]$$

Under the II regime,

$$V\left[a_{D,t+1} - E_t\left[a_{D,t+1}\right] | x_t^{II} > 0\right] =$$

$$Var\left[\epsilon_{D,t+1} | a_{Dt} > \frac{\bar{a}_S + \tau - (1-\rho)\bar{a}_D}{\rho}\right] =$$

$$Var\left[\epsilon_{D,t+1}\right] = \sigma_D^2$$

while under the DI regime,

$$V\left[a_{D,t+1} - E_{t-1}\left[a_{D,t+1}\right] | x_t^{DI} > 0\right] =$$

$$Var\left[\epsilon_{D,t+1} + \rho_D\epsilon_{Dt} | a_{D,t-1} > \frac{\bar{a}_S + \tau - (1 - \rho^2)\bar{a}_D}{\rho^2}\right] =$$

$$Var\left[\epsilon_{D,t+1} + \rho_D\epsilon_{Dt}\right] = (1 + \rho_D^2)\sigma_D^2$$

which is larger than the variance under the *II* regime. This proves the last part of the Lemma:

$$Var\left[pdiff_{t+1}^{II} \middle| x_t^{II} > 0\right] < Var\left[pdiff_{t+1}^{DI} \middle| x_t^{DI} > 0\right]$$

Proof of Theorem

Current welfare is composed of immediate consumer surplus CS_t , social surplus from storage SS_t , immediate producer surplus PS_{t-1} (lagged because shipping takes one period time), and immediate merchant surplus MS_t .

The market demand curve p_t^M as a function of exports x_{t-1} is given by

$$p_t^M = a_{Dt} - b_D \left(x_{t-1} - s_t \left(a_{Dt}, x_{t-1} \right) \right)$$

Immediate consumer surplus and net future social surplus of storage is the area underneath the market demand curve, minus the price paid by consumers and storers (Williams and Wright 1991):

$$CS_t + SS_t = \int_0^{x_{t-1}} p_t^M(q) \, \mathrm{d}q - p_t^{LIV} x_{t-1}$$

Immediate producer surplus is the area between the price received by producers and the supply curve p_t^S :

$$PS_{t-1} = p_t^{NY} x_{t-1} - \int_0^{x_{t-1}} p_t^S(q) \, \mathrm{d}q$$

The surplus of merchants is given by their profits⁶⁹:

$$MS_t = \left(p_t^{LIV} - p_{t-1}^{NY}\right) x_{t-1}$$

Welfare of a specific export quantity is therefore

$$W_t(x_{t-1}) = \int_0^{x_{t-1}} p_t^M(q) \, \mathrm{d}q - \int_0^{x_{t-1}} p_t^S(q) \, \mathrm{d}q$$

In the perfect foresight equilibrium (PFE) export and storage are chosen such that prices are constant across markets and across time, denoting PFE outcomes with stars:

$$p_t^* := p_t^{LIV} = p_{t-1}^{LIV} = p_{t-1}^{NY} = p_t^{NY}$$

I define deadweight loss due to information frictions as the difference of welfare between a case with information frictions and the perfect foresight model.

$$DWL = W_t^* - W_t$$

= $\int_{x_{t-1}}^{x_{t-1}^*} \left[p_t^M(q) - p_t^S(q) \right] dq$

The storage function has a time-dependent slope (dependent on a_{Dt}) and can therefore not be estimated empirically, which precludes direct estimation of the deadweight loss (except in the numerical exercises). But I can estimate a lower bound for the deadweight loss as follows: Note that the value of the integrand equals $p_t^{LIV} - p_t^{NY}$ at the lower bound x_{t-1} , and 0 at the upper bound x_{t-1}^* . The integrand is monotonically decreasing, and its slope is smaller than $b_D + b_S$ in absolute values. To see this, note that the slope of the market demand function is in absolute value less than or equal to the slope of the consumption demand function, as the storage function is non-decreasing in exports (with slope between 0 and 1, Williams and Wright 1991, p. 101):

$$\frac{\partial p_t^M}{\partial x_{t-1}} = -b_D \left(1 - \frac{\partial s_t^B}{\partial x_{t-1}} \right) \ge -b_D$$

I denote \tilde{q} where $p_t^{LIV} - b_D(\tilde{q} - x_{t-1}) = p_t^S(\tilde{q})$ and define l(q) for $q \in [x_{t-1}, x_{t-1}^*]$ such that $l(q) \leq p_t^M(q) - p_t^S(q)$ in that intervall:

⁶⁹I ignore transport cost in this proof to avoid cluttered notation, but account for it in the empirical part.

$$l(q) := \begin{cases} p_t^{LIV} - b_D (q - x_{t-1}) - p_t^S(q) & \text{for } x_{t-1} \leq q \leq \tilde{q} \\ 0 & \text{for } \tilde{q} \leq q < x_{t-1}^* \end{cases}$$

Using integrand l(q) yields a lower bound for the deadweight loss from information frictions:

$$DWL = \int_{x_{t-1}}^{x_{t-1}^*} \left[p_t^M(q) - p_t^S(q) \right] dq \ge \int_{x_{t-1}}^{\tilde{q}} l(q) dq = \frac{\left(p_t^{LIV} - p_{t-1}^{NY} \right)^2}{2(b_D + b_S)}$$

1.B Figures



Figure 1.B.1: Price series



Figure 1.B.2: Herfindahl index of market structure of cotton merchants



Figure 1.B.3: Information delay between New York and Liverpool over time



Figure 1.B.4: Distribution of information lags between New York and Liverpool (work days), before and after the telegraph



Figure 1.B.5: Price difference and freight cost



Figure 1.B.6: Cotton receipts at the New York exchange



Figure 1.B.7: Reaction of New York prices to news from Liverpool



Figure 1.B.8: Impact of Liverpool price on New York price



Figure 1.B.9: Contemporaneous price difference and price difference to the known, delayed Liverpool price



Figure 1.B.10: Surplus from export transaction, perfect foresight equilibrium



Figure 1.B.11: Surplus from export transaction when demand is overestimated



Figure 1.B.12: Price distortions: Difference between Liverpool price and lagged New York price plus transport costs

1.C Tables

	Before telegraph	After telegraph	Difference
Information lag	10.03	1.31	-8.72***
-	(0.13)	(0.06)	(0.15)
Liverpool price	18.11	13.16	-4.95***
	(0.33)	(0.15)	(0.36)
New York price	15.55	11.51	-4.04***
	(0.21)	(0.13)	(0.25)
Price difference $(p_t^{LIV} - p_t^{NY})$	2.56	1.65	-0.91***
	(0.18)	(0.05)	(0.19)
Sail ship freight cost	0.28	0.24	-0.04***
	(0.01)	(0.01)	(0.01)
Steam ship freight cost	0.51	0.39	-0.12***
	(0.02)	(0.02)	(0.03)
Exports	459.88	631.80	171.90**
	(37.64)	(61.80)	(72.37)

Table 1.C.1: Summary statistics

Notes: Information lag is in work days. Prices and freight cost are in pence per pound of cotton. Exports from New York to Liverpool are given in bales. Newey West standard errors in parentheses. *** p<0.01, **p<0.05, *p<0.1.

Dependent variable: $p^{LIV} - p^{NY} - \tau$	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Constant	2.56***	2.27***	2.05***	2.17***	1.52***	1.84***	1.43***	1.27***	0.004	1.67***
	(0.18)	(0.18)	(0.17)	(0.17)	(0.17)	(0.15)	(0.16)	(0.21)	(0.040)	(0.13)
Telegraph dummy	-0.91***	-0.86***	-0.78***	-0.83***	-0.71***	-1.03***	-0.90***	-0.74***	-0.022**	-0.71***
• • •	(0.19)	(0.18)	(0.17)	(0.18)	(0.18)	(0.16)	(0.15)	(0.21)	(0.011)	(0.12)
Cotton supply							0.13***	0.13***	0.006**	0.06***
							(0.03)	(0.04)	(0.003)	(0.02)
Transport costs τ :										
Freight cost	none	sail	steam	avg						
Other transport costs					yes	yes	yes	yes	yes	yes
Excluding no trade periods						yes	yes	yes	yes	yes
Accounting for shipping time								yes	yes	yes
Observations	604	604	604	604	604	575	575	575	575	575

Table 1.C.2: Average price difference

Notes: Newey West standard errors in parentheses. *** p<0.01, **p<0.05, *p<0.1. As freight cost are not available for all the periods, they are interpolated when missing. The average freight cost is calculated as the average of sail and steam freight rate if both are available, or the freight rate that is available (assuming freight cost are not printed in newspapers if the freight type was not used). Total transport costs include the average freight cost, additional 0.17 pence/pound unit freight cost, and 3.1% ad valorem transport costs (on the New York price). In (6) the period of around four weeks during May 1866 (when exporters were inactive because the price in New York exceeded the price in Liverpool) is excluded. In (7) the price difference daily cotton supply (receipts, in thousand bales) is used to control for potential disruptions in cotton production after the American Civil War. In (8) the price in Liverpool at the time of the arrival of the shipment (around 10 days in the future) is used to control for transport time of shipment. In (9) I use the difference in log prices instead of levels, and log cotton supply. In (10) the last known Liverpool price is used to calculate the price difference.

Table 1.C.3: Variance of price difference

Dependent variable: \widehat{Var} (pdiff)	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Constant	3.36***	3.26***	3.00***	3.17***	3.09***	2.19***	2.07***	3.79***	0.007	1.37***
	(0.59)	(0.57)	(0.53)	(0.56)	(0.56)	(0.46)	(0.37)	(0.76)	(0.004)	(0.22)
Telegraph dummy	-3.11***	-3.03***	-2.82***	-2.96***	-2.89***	-1.99***	-1.95***	-3.91***	-0.010***	-0.43**
	(0.59)	(0.57)	(0.53)	(0.56)	(0.56)	(0.46)	(0.42)	(0.85)	(0.002)	(0.20)
Cotton supply							0.04	0.15	0.000	-0.10***
							(0.08)	(0.13)	(0.000)	(0.04)
Transport costs τ :										
Freight cost	none	sail	steam	avg	avg	avg	avg	avg	avg	avg
Other transport costs				, in the second s	yes	yes	yes	yes	yes	yes
Excluding no trade periods						yes	yes	yes	yes	yes
Accounting for shipping time								yes	yes	yes
Observations	604	604	604	604	604	575	575	575	575	575

Notes: Newey West standard errors in parentheses. *** p<0.01, **p<0.05, *p<0.1. The dependent variable is given by $\frac{N_{before}}{N_{before}-1} \left(\text{pdiff}_t - \overline{\text{pdiff}}_{before} \right)^2 \text{ and } \frac{N_{after}}{N_{after}-1} \left(\text{pdiff}_t - \overline{\text{pdiff}}_{after} \right)^2$, so that the coefficient on the constant yields the sample variance on the period before the telegraph, and the coefficient on the telegraph dummy yields the change in the sample variance before versus after the telegraph. The different columns repeat the different specifications from Table 1.C.2.

Table 1.C.4: Variance of price difference using within period variation

Dependent variable	(1)	(2)	(3)
ln Var (pdiff)			
Telegraph dummy	-2.21***		-0.97
	(0.24)		(0.78)
Information lag <i>l</i> , work days		0.24***	0.14*
		(0.03)	(0.08)
Supply	-0.05	-0.04	-0.05
	(0.06)	(0.06)	(0.06)
Observations	585	585	585

Notes: Newey West standard errors in parentheses. *** p<0.01, **p<0.05, *p<0.1. The dependent variable is given by log of $\frac{N_{before}}{N_{before}-1} \left(\text{pdiff}_t - \overline{\text{pdiff}}_{before} \right)^2$ and $\frac{N_{after}}{N_{after}-1} \left(\text{pdiff}_t - \overline{\text{pdiff}}_{after} \right)^2$. No trade periods are excluded as in Table 1.C.3.

Dependent variable:	(1)	(2)	(3)	(4)
ln(New York price)	Before to	elegraph	After te	legraph
ln("telegraphed"		0.00193	0.734***	0.710***
Liverpool price)		(0.0661)	(0.0612)	(0.0651)
ln("steam shipped"	0.434***	0.433***		0.0685
Liverpool price)	(0.0319)	(0.0321)		(0.0635)
Observations	301	300	303	303

Table 1.C.5: Impact of telegraphed vs. steam shipped Liverpool price on New York price

Notes: Counterfactual "telegraphed" price before telegraph is the Liverpool price in t - 1. Counterfactual "steam shipped" price after telegraph is Liverpool price in t - 10. Prices are measured in pence/pound. Estimation of an AR(3) model with maximum likelihood. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

Table 1.C.6: Impact of known Liverpool price on exports

Dependent variable:	(1)	(2)	(3)	(4)	(5)	(6)
ln(exports)	Bef	ore telegra	aph	Aft	er telegra	aph
ln("telegraphed"		-0.559	-0.313	1.497	1.608	2.285
Liverpool price)		(1.080)	(1.174)	(1.856)	(2.352)	(2.460)
ln("steam shipped"	2.482***	2.940***	3.137***		-0.164	0.827
Liverpool price)	(0.682)	(1.103)	(1.111)		(2.478)	(2.449)
Linear time trend Observations	216	215	yes 215	234	234	yes 234
eeservations	210	210	210	_ 01	_ 01	201

Notes: Columns (1) to (3) use data from the sample before the telegraph got established, columns (4) to (6) use data from the sample after the telegraph got established. Counterfactual "telegraphed" price before telegraph is the Liverpool price in t - 1. Counterfactual "steam shipped" price after telegraph is the Liverpool price in t - 10. Exports are measured in bales, Liverpool price is measured in pence/pound. Columns (3) and (6) include a linear time trend. Estimation of an AR(14) model with maximum likelihood. Standard errors in parentheses. *** p<0.01, ** p<0.05, * p<0.1

	(1)	(2)	(2)	(4)	(=)		(=)	(0)
Dependent variable:	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Exports								
Constant	0.46***	-0.39***	-0.26***	-0.34***	-1.10***	-1.08***	-1.06***	-1.25***
	(0.04)	(0.10)	(0.08)	(0.10)	(0.19)	(0.23)	(0.23)	(0.24)
Telegraph dummy	0.17**	0.30***	0.35***	0.33***	0.48***	0.47***	0.47***	0.27***
	(0.07)	(0.07)	(0.07)	(0.07)	(0.08)	(0.09)	(0.09)	(0.10)
Transport costs		3.02***	1.42***	2.10***	1.50***	1.49***	1.43***	1.43***
		(0.38)	(0.16)	(0.27)	(0.20)	(0.22)	(0.23)	(0.23)
Cotton supply							0.01	0.01
							(0.02)	(0.02)
Transport costs τ :								
Freight cost		sail	steam	avg	avg	avg	avg	avg
Other transport costs					yes	yes	yes	yes
Excluding no trade periods						yes	yes	yes
Harvest year dummy						-	-	yes
Observations	604	604	604	604	604	575	575	575

Table 1.C.7: Average exports from New York to Liverpool

Notes: Newey West standard errors in parentheses. *** p<0.01, **p<0.05, *p<0.1. Exports are in thousand bales.

Table 1.C.8: V	ariance of	exports	from New	York to	Liverpool
----------------	------------	---------	----------	---------	-----------

Dependent variable: \widehat{Var} (exports)	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Constant	0.29***	-0.51**	-0.20	-0.38*	-0.94**	-1.03**	-1.03**	-1.12**
Telegraph dummy	(0.04) 0.33^{***}	(0.24) 0.45^{***}	(0.18) 0.45^{***}	(0.22) 0.46^{***}	(0.39) 0.57*** (0.17)	(0.47) 0.61^{***}	(0.47) 0.61^{***}	(0.47) 0.21^{*}
Transport costs	(0.12)	(0.14) 2.85***	(0.15) 0.96*** (0.25)	(0.14) 1.77^{***}	(0.17) 1.18^{***}	(0.19) 1.25^{***}	(0.20) 1.24***	(0.12) 1.26***
Cotton supply		(0.87)	(0.35)	(0.58)	(0.39)	(0.44)	(0.48) 0.00 (0.03)	(0.48) -0.00 (0.03)
Transport costs τ: Freight cost		sail	steam	avg	avg	avg	avg	avg
Other transport costs Excluding no trade periods					yes	yes yes	yes yes	yes yes
Observations	604	604	604	604	604	575	575	575

Notes: The dependent variable is given by $\frac{N_{before}}{N_{before}-1} \left(\exp_t - \overline{\exp}_{before}\right)^2$ and $\frac{N_{after}}{N_{after}-1} \left(\exp_t - \overline{\exp}_{after}\right)^2$, so that the coefficient on the constant yields the sample variance on the period before the telegraph, and the coefficient on the telegraph dummy yields the change in the sample variance before versus after the telegraph. Exports are in thousand bales. Newey West standard errors in parentheses. *** p<0.01, **p<0.05, *p<0.1.

Dependent variable:	(1)	(2)	(3)	(4)	(5)	(6)
p_{t+k}^{LIV}	OLS	IV	NL	NL-IV	NL-IV	NL-IV
Constant	1.735*	2.043**	1.775***	2.010***	2.992**	1.415
	(0.903)	(0.794)	(0.642)	(0.537)	(1.256)	(1.134)
$x_t - \bigtriangleup s_{t+k}^{LIV}$	-0.045*	-0.073**	-0.036*	-0.049***	-0.050*	-0.033***
	(0.027)	(0.037)	(0.019)	(0.015)	(0.026)	(0.012)
p_{t-l}^{LIV}	0.884***	0.871***	0.881***	0.870***	0.822***	0.897***
	(0.056)	(0.051)	(0.032)	(0.027)	(0.049)	(0.070)
$x_{t-k-l} - \bigtriangleup s_{t-l}^{LIV}$	-0.020	-0.023	n/a	n/a	n/a	n/a
	(0.039)	(0.038)				
Observations	402	402	402	402	206	196
R squared	0.827	0.826	0.827	0.827	0.687	0.736
First stage F stat		125.8		125.8	65.3	171.3
First stage coefficient		0.647***		0.647***	0.676***	0.614***
0		(0.058)		(0.058)	(0.084)	(0.047)
Test: $\beta_1\beta_2 - \beta_3 = 0$	-0.020	-0.041				
	(0.046)	(0.058)				
Demand elasticity	-104.7	-64.8	-130.9	-97.7	-149.5	-90.7
Sample					before telegraph	after telegraph

Table 1.C.9: Estimation of the slope of the demand function

Notes: Prices are denoted in pence/pound. The quantities in the regressor are given in 1,000 bales (1 bale \approx 400 pounds). HAC standard errors in parentheses. *** p<0.01, **p<0.05, *p<0.1.

	$\begin{array}{c} (1) \\ p_t^{NY} \end{array}$	$\begin{array}{c} \textbf{(2)} \\ p_t^{NY} \end{array}$	$\begin{array}{c} \textbf{(3)} \\ p_t^{NY} \end{array}$	$\begin{pmatrix} 4 \\ p_t^{NY} \end{pmatrix}$	$(5) \\ p_t^{NY}$	$\begin{array}{c} \textbf{(6)} \\ p_t^{NY} \end{array}$
$\triangle s_t^{NY} + x_t$	0.071*** (0.026)	1.862*** (0.338)	1.715*** (0.286)	1.715*** (0.287)	1.574*** (0.300)	1.608*** (0.518)
Observations	554	554	469	469	227	242
Harvest year FE	yes	yes	yes	yes	yes	yes
Harvest cycle	yes	yes	yes	yes	yes	yes
Regression	OLS	ĪV	ĪV	ĪV	IV	ĪV
Instrument		known Liv price	known demand shock	known Liv price	known Liv price	known Liv price
First stage F-stat		46.95	57.81	57.54	46.85	10.49
First stage coefficient		0.27***	0.30***	0.30***	0.33***	0.49**
0		(0.04)	(0.04)	(0.04)	(0.05)	(0.15)
Supply elasticity	368.4	14.1	15.3	15.3	24.9	11.5
Sample					before telegraph	after telegraph

Table 1.C.10: Estimation of the slope of the supply function

Notes: Prices are denoted in pence/pound. The quantities in the regressor are given in 1,000 bales (1 bale \approx 400 pounds). Index *k* denotes shipping time from New York to Liverpool, and index *l* denotes information delay between Liverpool and New York. Harvest cycle controls for day of the harvest season, and the square of it. HAC standard errors in parentheses. *** p<0.01, **p<0.05, *p<0.1.

Parameter	Value	Method
Supply side (New York):		
b_S	1.608	Instrumental variables estimation
\bar{a}_S	13.03	Constant from estimation of AR(1) process on \bar{a}_S
ρ_S	0.24	AR(1) coefficient from estimation of AR(1) process on \bar{a}_S
σ_S^2	2.25	From estimation of AR(1) process on \bar{a}_S
Demand side (Liverpool):	-0.033	Instrumental variables estimation
0D ā-	-0.033	Constant from estimation of $AP(1)$ process on \bar{a}_{-}
ρ_D	0.91	AR(1) coefficient from estimation of AR(1) process on \bar{a}_D
σ_D^2	0.40	From estimation of AR(1) process on \bar{a}_D
Other parameters:		
Transport cost $ au$	0.81	Total transport cost as estimated in empirical section
Storage cost θ	0.004-0.01	From Boyle (1934)

Table 1.C.11: Parameters for calibration

Table 1.C.12: Change from delayed to instantaneous information regime, in percent

	Data	Model				
Storage cost:		0	0.004	0.1	1	∞
Main estimates for b_D and b_S :						
Mean						
LIV price - NY price	-35.9***	-1.36	-1.32	-1.04	-2.82	-2.79
Exports	35.7	0.33	0.32	0.25	0.63	0.63
Standard deviation						
LIV price - NY price	-72.3***	-41.84	-42.15	-57.40	-51.86	-52.62
Exports	65.7***	2.87	2.92	4.82	6.25	6.37
-						
Lower CI for b_D and b_S :						
Mean						
LIV price - NY price	-35.9***	-0.49	-0.48	-0.98	-1.14	-1.13
Exports	35.7	0.23	0.23	0.43	0.50	0.50
Standard deviation						
LIV price - NY price	-72.3***	-30.85	-31.31	-41.28	-42.02	-42.01
Exports	65.7***	1.50	1.58	3.07	3.36	3.35
-						
Upper CI for b_D and b_S :						
Mean						
LIV price - NY price	-35.9***	0	0	-0.68	-0.76	-0.49
Exports	35.7	0	0	0.08	0.09	0.06
Standard deviation						
LIV price - NY price	-72.3***	-1.10	-1.20	-7.14	-9.00	-9.21
Exports	65.7***	0.04	0.04	0.32	0.48	0.26

Notes: Change is from delayed (=before telegraph) to instantaneous (=after telegraph) information regime, in percent of the underlying variables. Model predictions are based on a simulation of the model over 10,000 periods. Summary statistics are based on weekly data. Storage cost of infinity mean prohibitively high storage cost. Preferred estimates are shaded in light gray.

	Annual welfare loss, thousand pounds	[95% Conf. Interval]		
Before telegraph	988	608	2,667	
After telegraph	125	77	342	
Change	-863	-531	-2,325	
Change in percent	-87%	-87%	-87%	
In % of annual export value	-8.4%	-5.2%	-22.7%	

Table 1.C.13: Estimation of welfare gain from telegraph

Notes: Confidence interval of welfare loss is based on confidence intervals for the slopes of the demand and supply functions. Annual export value in the data is 10.2 million pounds.

Chapter 2

Survive another day: Using changes in the composition of investments to measure the cost of credit constraints

We introduce a novel empirical strategy to measure the size of credit shocks. Theoretically, we show that credit shocks reduce the value of long term investments relative to short term ones. If demand shocks affect short and long run investments similarly, credit shocks can be measured *within firms* by the shift in the investment vector away from long run investments and towards short term ones. This within-firm strategy makes it possible to use firm-times-year fixed effects to capture unobserved between firm heterogeneity as well as idiosyncratic demand shocks. We implement this strategy using a rich panel data set of Spanish manufacturing firms before and after the credit crisis in 2008. This allows us to quantify the effect of the credit crunch: our theory suggests that credit constraints are equivalent to an additional tax rate of around 11% on the longest lived capital. To pin down credit constraints as the cause for this investment pattern we use two triple differences strategies where we show (i) that only Spanish owned firms became credit constrained during the financial crisis, and that the drop in long term investments after the crisis is indeed driven by credit constrained Spanish firms; and that (ii) the impact on long term investment is mostly noticeable in firms that started the crisis with more mature debt to roll over.

2.1 Introduction

Studying the impact of credit shocks on investment empirically requires solving an identification problem: separating the impact of the supply of credit from the impact of the aggregate demand shock that usually takes place concurrently. To do this, the more recent literature has proposed a within-firm estimator which holds the firm constant and compares the effect of different lenders on the same firm.¹ Here, we propose an alternative within-firm estimator, using variation in different investment duration classes within a firm.

Our strategy exploits the differential impact of demand shocks and liquidity constraints on the composition of investments. As we show formally in a simplified version of Aghion et al. (2010), absent liquidity constraints, firms equalize the value of the marginal dollar on short term and long term investments. However, under liquidity constraints, long term investments involve a risk, since the firm may have to liquidate before the payoff period. This creates a wedge between the value of short and long term investments: Firms are willing to give up some future expected payoffs in order to increase the probability of surviving another day.

Based on this result we propose an identification strategy that allows us to place a lower bound on the impact of credit shocks. Assuming that demand shocks affect short term and long term investments within firms similarly, the difference between the longer term and the shorter term investment, if positive, is a lower bound on a first order approximation of the impact of the credit shock. The assumption that demand shocks have no differential impact on the composition of investment by duration is likely to be a conservative assumption, since the recession itself, if anything, would shift investments towards the future. That the driving force for the changes is liquidity and not demand is also consistent with the differential impact on foreign versus domestic firms as well as on constrained versus unconstrained firms, as we explain below. A crucial advantage of our strategy is that it allows us to examine the shift in the composition of investment within firms before and after a financial shock, including firm-times-year fixed effects to make sure that neither demand shocks nor unobserved heterogeneity between firms (different firms react differently to the crisis, but this is absorbed by the year-firm fixed effects) bias the estimated impact of credit constraints. Our estimate thus has, as we shall show, a clear economic interpretation.

Our identification strategy requires formulating a taxonomy of investments by their time to payoff, or durability. We rely on an extensive existing literature (which we survey) to determine the relative durability of different investment categories. According to this literature the shortest lived investment is advertising, followed by IT, R&D, with fixed capital investment like equipment and machinery being, on average, the longest lived.

To conduct our empirical analysis we need two things: a credit crisis, and detailed data about different investment types. Luckily for the case of Spain both are available:

¹See for example Gan (2007), Khwaja and Mian (2008), Paravisini (2008), Iyer et al. (2014), Jiménez et al. (2010), Iyer and Peydró (2011), Schnabl (2012), Jiménez et al. (2012), Paravisini et al. (2011).

Spain suffered from a particular severe credit crisis in the wake of the financial crisis in 2008, and at the same time there exists detailed firm level data with investment information.

We use the financial crisis in 2008 as an exogenous shock to credit supply. This is possible because the 2008 crisis was at its core a banking crisis. Previous research has established that the reduced bank liquidity translated into a reduction of credit supply to firms (e.g. Iyer et al. 2014, Paravisini et al. 2011, Ivashina and Scharfstein 2010, Adrian et al. 2012, Santos 2011 for the financial crisis in 2008; and Chava and Purnanandam 2011 for the Russian crisis in 1998). This is particularly true for the case of Spain, where the liquidity crisis was exceptionally severe. Jiménez et al. (2012) show that weaker banks deny more loans, even when the loans are identical (which allows them to identify the supply rather than demand channel) and that firms can usually not substitute the weak bank with another bank. Bentolila et al. (2013) show that firms who borrowed more from weak financial institutions that were later bailed out (the old "Cajas") reduced employment by an additional 3.5 to 5 percentage points relative to firms who borrowed from healthier ones.

We use Spanish firm-level data from the Encuesta Sobre Estrategias Empresariales (ESEE, Survey of business strategies); a rich, high quality, long-term panel data set of Spanish manufacturing firms that breaks up investment information into a total of eight different investment categories: Advertising, IT, R&D, vehicles, machinery, furniture, buildings, and land. Since the Spanish financial crisis was based on a real estate bubble, we do not use land and buildings investment in our analysis, as it might bias our results towards finding the hypothesized effect.

Applying our estimation strategy to the Spanish data we find that after the financial crisis the longest term investments were reduced by 17 percentage points more than shortest term investments. This finding is robust to different classifications of short versus long term investment. For example, it does not matter whether we use depreciation rates from the literature directly as a measure for time-to-payoff, or if we just use the ranking of investment types based on their depreciation rates. Similarly, the finding holds for several ways of grouping some investment types with similar depreciation rates into one category. Also, we conduct placebo tests estimating the effects of the financial crisis in 2008 year by year and find reassuringly that the difference in investment behavior only appears in the crisis years. We conclude therefore that the 17 percentage point difference is our estimate of the impact of the financial crisis on investment. We show that, given our theory, this is equivalent to an 11% incremental tax rate on the longest term investment.

The second part of our empirical analysis aims to more precisely pin down credit constraints as the mechanism leading to the change in investment patterns (as in Bernanke and Gertler 1989), as opposed to other mechanisms which could lead to similar effects, such as an increase in uncertainty (Bloom 2009, Bernanke 1983), which could increase the option value of waiting. If credit constraints were indeed the cause of the change, we should see a stronger effect for firms that were more affected by credit constraints. We use two ways, suggested by the literature, to identify firms that were particularly affected by the financial crisis: domestic firms (as opposed to foreign ones), and firms with a lot of mature debt that needs to be rolled over at the beginning of the crisis.

Foreign firms are typically less affected by a credit squeeze since they have access to external finance via their parent companies (Desai et al. 2004, Kalemli-Ozcan et al. 2010). Indeed, this is the case in our data: the credit drop is only observed for Spanish firms. Thus under our hypothesis, the shift in the composition of investment (if linked to credit) should only occur in Spanish firms. Our analysis exploits that and amounts to a triple differences estimation: We compare long-term versus short-term investments before and after the financial crisis in 2008 for Spanish versus foreign firms. Our triple differences analysis confirms our initial analysis, and reassures us that credit constraints are indeed at play.

This robustness check may still fail to convince us, since domestic and foreign owned firms differ among a variety of other dimensions besides their access to external funding. For example, Spanish owned firms in our data are typically smaller and less likely to export, and might therefore show a different investment behavior. We address this concern by conducting a variety of robustness checks. First, we use only multinational firms for our comparison. These firms are all large, have subsidiaries in many countries and are heavily export oriented, the only difference between them being the nationality of their majority shareholder. Second, we use an inverse propensity score reweighting scheme based on the size, growth, export status and export development of firms before the financial crisis. This strategy basically matches foreign owned firms to comparable Spanish ones. Third, we make sure that firm size is not driving the results. Spanish firms are smaller, so this could be just a size effect. However, the magnitude of our estimated result does not vary at all when we control for firm size, so size is not driving the change in investment pattern. Overall, the results are very similar in magnitude across all alternative specifications, which gives us additional confidence that we are picking up the right effect. Fourth, the exit rates of Spanish and foreign firms are not statistically significantly different (consistent with our mechanism, as firms manage to avoid bankruptcy due to their changing investment behavior), so compositional effects are not driving our results either. A final alternative hypothesis could be that the liability side of Spanish firms' balance sheets could be responsible for our results: If firms cannot raise long term funding, then maturity matching could lead them to reduce the maturity of their asset side. However, we can rule out this explanation as well, as the data shows not difference between Spanish and foreign firms in the maturities of liabilities after the crisis.

A second way to investigate whether the mechanism we suggest — credit constraints — is at the root of the shift in the investment vector towards shorter time-topayoff relies on the observation that firms with a lot of mature debt at the beginning of a financial crisis also tend to be more affected by it because they experience difficulty in rolling their debt over under a credit crunch (Almeida et al. 2012). Therefore we use the ratio of short term debt over total debt to identify more affected firms in another triple difference estimation. Our estimates are again consistent with the evidence from before: these firms reduce long term investment relatively more than shorter term investments compared to other firms.

Another nice feature of both of our triple differences specifications is that they allow us to include category-year fixed effects to control for the cyclical behavior of different investment types. For example, it has been shown that R&D exhibits a pro-cyclical pattern (Barlevy 2007). However, this does not affect our results. Furthermore, the firms which are more credit constrained do not exhibit a difference in their output or exports compared to less credit constrained firms. This reassures us that we are not seeing a demand-driven shift towards short term investments caused by the recession, but rather a change in composition that only takes place for credit constrained firms.

Our finding that credit constraints are eating the long term future profits of firms in order to guarantee survival for another day complements a large literature that has established that financially constrained firms invest less,² and recent studies that use the world wide financial crisis in 2007/2008 as an exogenous shock to the credit supplied by banks.³ A smaller literature has studied how credit rationing affects the composition of firm investments. For example, see Eisfeldt and Rampini (2007) for the allocation of investment between new and used capital, as well as Campello et al. (2010), who point out that firms cut technology and marketing investment by more than capital investment, but do not offer an explanation why certain investment types might be more affected than others.

Beyond these findings we believe that our paper points a way forward in learning about credit shocks. We show how the rotation in the investment vector towards the present and away from the future informs us about the existence and the size of the credit crunch. Furthermore, we believe that this shift in the investment vector could have a macroeconomic impact: reducing long-term investment is likely to have a long term impact on the Spanish economy, impeding recovery after the financial crisis, and reducing long-term economic growth.

2.2 Theoretical Framework and Identification

2.2.1 Theory: Investment Duration and Liquidity Risk

Most theoretical analysis of liquidity constraints aggregates all investment into one single decision (e.g. Kiyotaki and Moore 1997). Instead, we assume that a profit maximizing firm can choose between two types of investment: **short-term investments** k_t yield an immediate payoff of $f(k_t)$, while **long-term investments** z_t yield a higher payoff $(1 + \rho)f(z_t)$ which is paid out at a later period. To capture this trade-off we rely on a model that is a simplified version of Aghion et al. (2010). The key difficulty of firms

²For example, Whited (1992), Carpenter et al. (1994), Hubbard et al. (1995), Bernanke et al. (1996), Bernanke et al. (1999), Kaplan and Zingales (1997), Lamont (1997), Cleary (1999), Klein et al. (2002), Amiti and Weinstein (2013), Fazzari et al. (1988).

³Campello et al. (2010), Duchin et al. (2010), Almeida et al. (2012), Kuppuswamy and Villalonga (2012).

is that with probability $1 - \lambda_{t+1}$ a liquidity crisis in the interim period before the payoff of the long term investment is realized, which may simply force the firm to liquidate. Thus the probability of survival λ_{t+1} measures the probability that the entrepreneur will have enough funds to cover the liquidity shock and is allowed to depend on the levels of short and long term investments. Specifically, reallocating investments from long to short term increases the probability of survival, $\left(\frac{\partial \lambda_{t+1}}{\partial k_t} - \frac{\partial \lambda_{t+1}}{\partial z_t}\right) > 0$. The choice of how much short run and long run investment to undertake is then given by:

$$\max_{k_t, z_t} E_t \left[f(k_t) + \beta \lambda_{t+1} (1+\rho) f(z_t) - q_t k_t - q_t z_t \right]$$
(2.2.1)

where λ_{t+1} measures the probability that the entrepreneur will have enough funds to cover the liquidity shock, ρ is the additional productivity of long term investment, and the rest of terms have their usual meanings.

The first order condition with respect to *k* is:

$$E_t \left[f'(k_t) \right] + \beta E_t \left[\frac{\partial \lambda_{t+1}}{\partial k_t} (1+\rho) f(z_t) \right] = q_t, \qquad (2.2.2)$$

and with respect to *z*:

$$\beta E_t \left[\lambda_{t+1} (1+\rho) f'(z_t) \right] + \beta E_t \left[\frac{\partial \lambda_{t+1}}{\partial z_t} (1+\rho) f(z_t) \right] = q_t.$$
(2.2.3)

Combining the two equations, we obtain the marginal condition:

$$E_t \left[f'(k_t) \right] = \beta E_t \left[(1 - \tau_{t+1}) \left(1 + \rho \right) f'(z_t) \right]$$
(2.2.4)

where

$$\tau_{t+1} = (1 - \lambda_{t+1}) + \left(\frac{\partial \lambda_{t+1}}{\partial k_t} - \frac{\partial \lambda_{t+1}}{\partial z_t}\right) \frac{f(z_t)}{f'(z_t)}$$

This contrasts with the first best, absent liquidity shocks, when it should be the case that the marginal value of a dollar is equalized across both types of investments:

$$\boldsymbol{E}_t\left[f'(k_t)\right] = \boldsymbol{\beta}\boldsymbol{E}_t\left[(1+\rho)f'(z_t)\right]. \tag{2.2.5}$$

Thus the risk that the firm will run out of cash in period t + 1 works exactly like a tax on investment τ_{t+1} and reduces the value of the (a priori more profitable) long term investments relative to the first best. The first term of this wedge, $(1 - \lambda_{t+1})$, captures the probability of failure. The second term captures the marginal change in this probability as we reallocate investment from long term to short term. Given that reallocating investments from long term to short term increases the probability of survival, the tax wedge $\tau_{t+1} > 0$. Hence the reallocation away from long term investment opportunities to short term ones is higher the higher the probability of avoiding bankruptcy by doing this, the higher the probability of not having enough liquidity next period, and the lower the marginal productivity of long run investments.

The model predicts that credit constrained firms will reduce long term investment

by more than short term investment in order to secure survival. In the next section we discuss how we implement this idea empirically.

2.2.2 Identification

Our theoretical framework suggests a new empirical strategy, closely linked to the theory, that can help us to identify credit shocks. To get exact expressions, we assume that the function *f* is Cobb Douglas, that is $y = k^{\alpha}$ (as usual everything goes through as a log linear approximation otherwise).

Suppose that there are good ex-ante reasons to expect liquidity to be plentiful before the shock to credit supply (denoted by subscript b), and to expect liquidity to be scarce after the credit shock (denoted by subscript a). Then we have from equation (2.2.5) that, for a given firm i,

$$f'(k_b^i) = \beta(1+\rho)f'(z_b^i)\varepsilon_b^i$$
(2.2.6)

where we assume ε_i is an i.i.d. log normal error term with mean 1. Thus, in logs, and using the Cobb-Douglas specification

$$(\alpha - 1)\ln k_b^i = \ln \left(\beta(1+\rho)\right) + (\alpha - 1)\ln z_b^i + \ln \varepsilon_b^i.$$

While after the liquidity crunch we have, from equation (2.2.4):

$$(\alpha - 1) \ln k_a^i = \ln (\beta (1 + \rho)) + \ln \left(1 - \tau_{t+1}^i\right) + (\alpha - 1) \ln z_a^i + \ln \varepsilon_a^i.$$

This immediately suggests a difference in differences estimator as the way to identify the wedge introduced by the liquidity shock in firm *i*. Specifically, the difference in difference estimator is:

$$(1-\alpha)\left(\left(\ln z_a^i - \ln z_b^i\right) - \left(\ln k_a^i - \ln k_b^i\right)\right) = \ln\left(1 - \tau_{t+1}^i\right) + \ln \varepsilon_a^i - \ln \varepsilon_b^i$$

where $E\left(\ln \varepsilon_{a}^{i} - \ln \varepsilon_{b}^{i}\right) = 0.$

Now consider the following difference-in-differences specification using investment *I* in investment category c = k, z as dependent variable:

$$\ln I_{ict} = \beta_0 + \beta_1 * \operatorname{crisis}_t * \operatorname{longterm}_c + \operatorname{crisis}_t + \operatorname{longterm}_c + \nu_{ict}$$

where crisis_t is a dummy variable that turns 1 in the years of a financial crisis, and $\operatorname{longterm}_c$ is a dummy variable indicating a long term investment (i.e. it equals 1 if investment type c = z). In this specification the coefficient on the interaction term equals:

$$\beta_1 = \boldsymbol{E} \left(\left(\ln I_{iza} - \ln I_{izb} \right) - \left(\ln I_{ika} - \ln I_{ikb} \right) \right)$$

However, this last expression equals, up to a factor, the wedge between long term

and short term investments, which has a clear economic interpretation in the theory:

$$\beta_1 = \frac{E\left(\ln\left(1 - \tau_{t+1}\right)\right)}{(1 - \alpha)}$$

In reality and in our data we have more than two investment categories, thus we generalize our formula above to multiple investment types. Furthermore we can include firm times year fixed effects as well as investment category fixed effects to make sure that the structural equation above is identified. This leads to our estimated regression equation:

$$\ln I_{ict} = \beta_0 + \beta_1 * \text{crisis}_t * \text{duration-of-investment}_c$$

+firm * year FE_{it} + category FE_c + ν_{ict} (2.2.7)

2.3 Data

2.3.1 Identifying Long and Short Term Investments

The theory allows us to make predictions about the behavior of different investment categories depending on the horizon over which they pay off. For the model to guide our empirical work, we need a taxonomy of tangible and intangible investments by their durability. While accountants and growth accountants have produced a large body of work aiming to estimate the durability of tangible investments, the literature on intangible investment lifespan is somewhat less extensive.

The shortest lived investment category is brand equity and advertising. Landes and Rosenfield (1994) estimate the annual rates of decay of advertising to be more than 50% for most industries, using 20 two-digit SIC manufacturing and service industries. For a number of industries they even find that the effect of advertising does not persist until the following year. A more recent literature review by Corrado et al. (2009) concludes that the depreciation rate for advertising is 60%. They also note that 40% of advertising expenditure is spent on advertisements that last less than a year, e.g. on "this week's sale", which partly explains the short-lived impact of advertising.

The literature reports a depreciation rate of around 30% for software investments. The Bureau of Economic Analysis (1994) estimated a depreciation rate of 33% for a 5 year service life, according to Corrado et al. (2009). Tamai and Torimitsu (1992) report a 9 years average life span for software (between 2 and 20 years), relying on industry estimates based on survey evidence. The Spanish accounting rules give a depreciation rate of 26% for IT equipment and software, so we use a value of around 30% as summarizing the evidence in our main specification.

The evidence on the average depreciation rates and average lifespans of R&D capital is extensive, and estimates range from 10-30%. Pakes and Schankerman (1984) and Schankerman and Pakes (1986) propose 25% based on 5 European countries, and 11-26% in a later study for Germany, UK and France. Nadiri and Prucha (1996) estimated

a rate of 12% for R&D, while Bernstein and Mamuneas (2006) estimate the depreciation rate at 18-29%. Corrado et al. (2009) review the literature and settle on a value of 20% for R&D, which is the value we are using for our main analysis.

Longer lived investments include fixed tangible assets like machinery, vehicles and other equipment. Both the Spanish⁴ and the BEA's⁵ accounting rules yield similar values for these types of investment, with vehicles having a depreciation rate of around 16%, machinery around 12% and furniture and office equipment around 10%.

The longest lived investments are investments into real estate, i.e. land and buildings. According to the BEA's estimates, industrial and office buildings have a depreciation rate between 2-3%. Spanish accounting rules specify a very similar depreciation rate for buildings of 3%. It's harder to make a general statement about the depreciation of land. While land clearly is long-lived, many factors determine the price of land and therefore the implicit depreciation rate. In any case, since the Spanish financial crisis affected real estate prices strongly, we exclude this category from our analysis.

Our summary of the literature to classify investment types into time-to-payoff is presented in Table 2.B.1, ranked from the shortest to the longest time-to-payoff. We use our summarized depreciation rates as well as the ranks for our estimation and regroup some categories when there is some ambiguity as robustness checks. However, our results are robust to these checks.

2.3.2 Description of Data Set

We rely on the Encuesta Sobre Estrategias Empresariales (ESEE), a panel of Spanish manufacturing firms. This data has been collected by the Spanish government and the SEPI foundation every year since 1990.⁶ The survey covers around 1,800 Spanish manufacturing firms per year, which include all firms with more than 200 employees and a stratified sample of smaller firms. The coverage is about 50-60% of large firms, and 5-25% of small firms. The sample started out as a representative sample of the population of Spanish manufacturing firms. In order to reduce the deterioration of representativeness due to non-responding firms, every year new companies are resampled in order to replace exiting ones.

In contrast to balance sheet firm level data bases, which usually report only a single investment number for each firm, the Spanish data covers a number of different investment choices made by firms. A number of variables can be linked to our investment categories based on time-to-payoff: advertising expenditure, IT expenses, R&D expenses, investment into vehicles, machinery, and furniture & office equipment, as well as investment into land and buildings.

Besides these main investment variables, we also have data on the credit ratio of firms, and other complementary variables such as sales, exports and foreign ownership, which we will use for several types of robustness checks.

⁴Please see http://www.individual.efl.es/ActumPublic/ActumG/MementoDoc/MF2012_Coeficientes %20anuales%20de%20amortizacion_Anexos.pdf

⁵Please see http://www.bea.gov/scb/account_articles/national/wlth2594/tableC.htm

⁶For more information see http://www.fundacionsepi.es/esee/sp/spresentacion.asp
Table 2.B.2 presents summary statistics for the main variables that are the object of our analysis, before and after the crisis. The data shows that investment in all categories fell after the financial crisis in 2008. However, ex ante it is not clear whether this investment drop is triggered by the credit squeeze or the adverse demand shock. Our empirical strategy aims to disentangle these effects.

It is notable that investment into buildings show the largest statistically significant drop in investment. Land and buildings are also the longest lived investment categories. However, since the financial crisis in Spain was based on a real estate bubble which led to falling real estate prices, it seems safer to exclude land and building from our analysis as it would bias our results towards finding our hypothesized effect: The fall in building and land investment might reflect a fall in prices due to the burst of the real estate bubble rather than be caused by the credit squeeze.

The credit crunch triggered by the financial crisis is also reflected in the Spanish credit data: total credit as a percentage of total assets (the credit ratio) fell by 3 percentage points after the crisis, from 57% to 54%. At the same time, observed average credit cost increased by 0.22 percentage points, from 4.06% to 4.28%. This is obviously a lower bound on the increased cost, as firms often simply could not get access to credit. Together with the observed immediate drop in the credit ratio this suggests that we observe a credit supply rather than a credit demand shock immediately after the financial crisis hit.

2.4 Results

2.4.1 Differential Effect across Investment Types

Table 2.B.3 presents our main results from estimating regression equation 2.2.7. The dependent variable is the log of investment of firm *i* in year *t* in investment category *c*, where investment categories include the six investment types specified above: advertising, IT, R&D, vehicles, machinery, and furniture & office equipment. The main regressor is an interaction term of the inverse of the depreciation rate of an investment category as a measure of the time-to-payoff of an investment type as given by Table 2.B.1 and a time dummy variable that indicates the financial crisis (=1 in and after 2008). Standard errors are clustered at the firm level, allowing for autocorrelation across time and across investment categories within the firm.⁷ Column (1) implements the regression equation with just category and firm fixed effects. The coefficient on the interaction term is negative, implying that investments with a longer time-to-payoff fell more after the financial crisis than investments with a shorter time-to-payoff. The coefficient on the crisis dummy is also significant and strong, showing that most short

⁷Bertrand et al. (2004) point out that serially correlated outcomes in differences-in-differences estimations produce serially correlated residuals, and standard errors need to be adjusted accordingly. They recommend clustering errors at the group level, if the number of groups is large enough (e.g. 50). For our regression this would mean clustering at the investment category level, unfortunately our number of categories is too small for clustering (6 categories). We decided to correct for autocorrelation of the residuals by clustering at the firm level instead, which allows for arbitrary serial correlation over time as well as across investment categories within a firm.

run investment fell by 17 percentage points after the crisis. In column (2) we replace the crisis dummy with year fixed effects, which doesn't change the coefficient on the interaction term.

It is possible that the demand shock rather than the credit squeeze drives our result. In column (3) we control for firm times year fixed effects. If demand shocks don't have a differential effect across investment types, we manage to control for them in column (3). The magnitude of the effect increases somewhat and stays significant after introducing the firm year fixed effects. Note that in contrast to other papers on the effect of credit squeezes on investment this is likely to be a lower bound of the true estimate, because if demand shocks affect investment types differentially, they are likely to affect investments with a shorter time-to-payoff by more than investments with a longer time-to-payoff (the recession will, after all, finish at some point in the future). It is common that investment observations are often 0 and thus excluded from the analysis (in logs). Column (4) codes the 0's as 1 euro and thus includes all those observations. The results are substantially stronger, suggesting our baseline analysis is very conservative.

Table 2.B.4 uses alternative measures for time-to-payoff instead of the inverse of the depreciation rate. For example, column (2) uses the rank of investment types according to depreciation rates, using the highest rank for investment with the longest time-to-payoff. Since the estimated depreciation rates in the literature vary within investment types and sometimes overlap, we regroup the investment categories in columns (2) to (5) of the table. For example, the depreciation rates of R&D and IT are not that different in the literature, so we group them together. Also machinery and furniture have quite similar depreciation rates, justifying a similar treatment. However, our finding is robust across all these alternative measures for time-to-payoff. The magnitudes of the effect differ, but this is because the different measures use different units. In column (6) we check whether our result is sensitive to the non-linearity implicit in measuring time-to-payoff as the inverse of the depreciation rate, and use the depreciation rate instead. The sign reverses, as a higher depreciation rate now means *less* time to payoff, but the result remains again robust to this specification.

How can we interpret the economic significance of our effect? Our preferred specification in column (3) in Table 2.B.3 tells us that investment falls by 2 percentage points for a unit increase in the inverse depreciation rate. When we compare advertising, the investment category with the highest depreciation rate, to furniture and office equipment, the category with the lowest depreciation rate, the inverse depreciation rate increases by 8.3, so we need to multiply the regression coefficient by this value. This leads to our main result: Investment in office equipment gets reduced by 17 percentage points more than advertising expenditure. This is quite a sizable difference across investment categories.

Our theory suggests another way to interpret this result: as a tax on capital. Recall that

$$\beta_1 \approx \frac{\ln\left(1 - \tau_{t+1}\right)}{(1 - \alpha)}$$

Given that the investment gap between capital with the shortest and longest time-topayoff (β_1 in the theory) is 17% and using $\alpha = 1/3$ (the capital share), this means that the credit crunch is equivalent to an 11% tax on the long run investments relative to the shortest run one ($\tau_{t+1} = 1 - \exp(-0.17 * 2/3) = 10.6\%$).

2.4.2 Placebo Test

So far we have pooled the estimated effect across all years before and after the crisis, respectively. In order to make sure we are capturing the effect of the credit squeeze instead of something else, we check whether the timing of the effect really coincides with the credit squeeze. Therefore we allow the interaction term to vary with each year of the sample, the results are given in Table 2.B.5. The change in coefficients over time supports our story: In column (1) the coefficient becomes negative (but still insignificant) in the year 2008, and becomes even more negative and highly significant thereafter. This timing is consistent with the development of the credit squeeze: After the failure of Lehman Brothers in September 2008, conditions tightened severely. 2009 was the first full year in which the effects of the credit crunch were fully spread.

The effect is even more visible in Figure 2.A.1, where we plot the coefficients of the regression estimated in column (1) in Table 2.B.5 over time. While there was not much going on before the crisis, from 2008 on the reduction in long term investment became apparent.

2.4.3 Mechanism: Credit Crunch

In this section we aim to further pin down credit constraints as cause for the observed change in the investment behavior, as opposed to, for example, the effect of an increase in uncertainty. If our hypothesis is true, we should expect to see a differential effect on firms that are more affected by the credit crunch compared to firms that are less affected.

The literature suggests two types of firms that are typically affected more by credit constraints than others. First, domestic firms are typically more affected by a credit squeeze on domestic banks, since foreign firms have access to external finance in other countries that are less affected via their parent companies (Desai et al. 2004, Kalemli-Ozcan et al. 2010).

Second, firms that happen to have a lot of mature debt at the beginning of a financial crisis also tend to be more affected by it because they experience difficulty in rolling their debt over under a credit crunch (Almeida et al. 2012).

In the following we test whether we see a differential effect of the credit crisis across these two types of firms.

Foreign versus Domestic Firms

We start our analysis by looking at foreign versus domestic firms. If it is true that foreign firms are less affected by a credit squeeze, then we should observe a fall in the credit ratio only for domestic firms. Table 2.B.6 tests this. In column (1) we find that the credit ratio, defined by total credit divided by assets, on average fell after the crisis, a result that was already visible from the summary statistics in Table 2.B.2. Column (2) controls for industry specific demand conditions using the industry's exports and size as a time varying control. Also, firm level fixed effects allow us to control for any other time invariant unobserved firm heterogeneity.

Column (3) compares the drop between Spanish and foreign owned firms and answers the question: comparing two firms of the same size that are facing the same demand conditions, does the firm that happens to be Spanish suffer a significant drop in credit after the crisis? The answer is unambiguous and highly significant: Spanish firms suffer a drop in credit of around 2.5 percentage points after the crisis compared to non-Spanish firms. In column (4) we add time fixed effects to capture any common, time varying aspects of the crisis that are not yet captured by industry exports or size, and the effect remains the same. Column (5) is our most demanding specification, which allows for industry specific time effects (and thus absorbs our previous industry specific controls), and the result is again stronger, with Spanish firms facing a credit drop of 3.5 percentage points. This is equivalent to a 6.1% drop in credit relative to the 2007 baseline of 57.8% credit to assets for Spanish firms before the crisis.

Thus under our hypothesis, the shift in the composition of investment (if linked to credit) should only occur in Spanish firms. In Table 2.B.7 we start the analysis by running the main regression separately for domestic and foreign firms. Only the effect in column (1) is negative and statistically significant, suggesting that only domestic firms cut their long term investment relatively more than their short term investment. There is no significant difference across investment types for foreign owned firms.

In columns (3) and (4) we again conduct the placebo tests by allowing the interaction term to vary by year. Again we see that the effect is driven by domestic firms, in line with our hypothesis. Figure 2.A.2 shows this visually, we see the strong drop in long term investments only for domestic firms after 2008.

We can test this more formally by extending our analysis therefore to a triple difference estimation, comparing long-term versus short-term investments before and after the financial crisis in 2008 for Spanish versus foreign firms. This allows us to further challenge our results by including *category times year* fixed effects to control for the possibility that firms might reduce or increase investment in certain categories during recessions. The resulting estimating equation is:

- $\ln I_{ict} = \beta_0 + \beta_1 * \text{crisis}_t * \text{duration-of-investment}_c * \text{domesticfirm}_i$
 - + $\beta_2 * \text{duration-of-investment}_c * \text{domesticfirm}_i$
 - + firm * year FE_{it} + category * year FE_{ct} + ν_{ict} . (2.4.1)

Table 8 shows the results of the triple difference specification. It shows a significant differential negative effect for long term investments after the crisis undertaken by Spanish firms. As in our baseline case, column (4) codes the 0's as 1 euro and thus includes all the 0 investments. Again the results are substantially stronger, suggesting our baseline analysis is very conservative.

The differential effects for domestic firms by investment category and over time are visualized in Figure 2.A.3. Darker lines depict investment types with a longer time-to-payoff, i.e. for which we would expect a larger drop. The visual evidence is broadly in line with our hypothesis, as lighter lines show a smaller, and darker lines show a larger drop after 2008. It is also notable that until 2007 there is no differential effect by investment types, the lines are all parallel and very close. The differential effect only starts to come in after 2007, when the credit crunch hits.

A worry is that domestic and foreign owned firms differ among a variety of other dimensions besides access to external funding. For example, Spanish owned firms in our data are typically smaller and less likely to export and might therefore show a different investment behavior. To address this concern, Table 2.B.9 conducts a variety of robustness checks. One dimension of time-varying, unobserved heterogeneity might be differences between companies that operate across countries and those that operate in a single country. Companies that operate in many countries belong to a corporate group, and this could provide companies with advantages that go beyond their access to capital. For example they might face a more diversified demand. Column (2) conducts our analysis only for companies that belong to a corporate group, presumably most of them are multinationals. The results are pretty remarkable. Even though the sample size drops very substantially (by more than half), the effect remains remarkably similar and highly significant. Column (3) restricts the sample to firms that have non-industrial plants in foreign countries. The drop in the sample size is now enormous, yet the effect remains. The finding is similar in column (4), in which we restrict the sample to firms that have share holdings in foreign countries.

Column (5) uses another way to make the control group of foreign firms a more suitable counterfactual for the treatment group of domestic firms by applying inverse propensity score weights. This type of matching estimator reweights each observation by its (inverse) propensity score (the "likelihood" that a firm belongs to the treatment group, i.e. is under Spanish ownership) in order to generate the same distribution of (observed) characteristics of treatment and control group, and therefore hopefully also match the unobserved time varying heterogeneity better. We construct propensity scores based on sales and export status (as these observables seem to be the major differences between Spanish and foreign owned firms) of all pre-treatment years based on probit regression of the treatment (i.e. Spanish ownership) on sales and export status in all years between 2003 and 2007. The predicted values of these regressions, *treat*, are then used to calculate inverse propensity score weights $psw = \frac{treat}{1-treat}$, which we use as weights for all firms in the control group in our regression (for more details on the method, see DiNardo et al. 1996 and Nichols 2007 and Nichols 2008).

Our results from the inverse propensity score reweighing in column (5) are also robust to this test. Most of the results are numerically very close to the baseline specification, suggesting that selection is not a major concern in our analysis.

The last two columns in Table 2.B.9 analyze whether firm size or productivity differences are driving the results by including interaction terms with ln(sales) and ln(TFP) besides the interaction term with domestic firms.⁸ However, both size and productivity fail to explain the differential drop in investment, the ownership interaction remains significant and its magnitude is unchanged in spite of including this competing explanation.

A separate concern is the extent to which differential exit rates of Spanish and foreign owned firms could explain these results. Suppose simply that "worse" firms are exiting. If "worse" firms are those that feature more long term investments, then we shall see more short term investment and less long term ones in the surviving data. This seems unlikely a priori, as we tend to think of better firms as the ones doing more long term investment. In any case, the exit rates among Spanish versus foreign firms behave very similarly, as Figure 2.A.4 shows: Exit rates are not statistically significantly different, which in fact suggests that our mechanism is operating: Firms reduce their long run investments to generate liquidity, and manage to survive another day.

A final concern is the mechanism through which this process takes place. Specifically, while we postulate in the theory that it takes place through the asset side of the balance sheet (firms have less access to credit in general and decide to cut long-term investments), an alternative hypothesis is that it takes place through the liability side: firms have less access to long-term credit, and therefore cut long-term investment because otherwise they cannot match the liabilities and investments by debt maturity. To test this, in Table 2.B.10 we check whether domestic firms suffered a differential drop in long term credit (as a ratio of total credit) compared to foreign firms, using the same specification as in Table 2.B.6. However, while Spanish firms suffer from access to credit in general as shown in Table 2.B.6, there is no differential effect with respect to long-term credit as opposed to short term credit. So a differential liability matching does not explain our results.

Firms with Maturing Debt just before the Crisis

An alternative approach to studying the mechanism that does not rely on using nationality of ownership as the driver of credit constraints is to use firms whose debt is maturing just before the crisis as a treatment group. These firms are likely to be more severely affected by the credit squeeze as they have to roll over their debt when the crisis starts. We use short term credit with financial institutions divided by total credit in 2007, the year before the crisis, as measure for more credit constrained firms in Table 2.B.11. Column (1) repeats our main specification from before using domestic

⁸TFP is estimated by the Levinsohn-Petrin method with separate regressions by industry and furthermore includes a correction for both changes in input or output prices (input and output price changes are asked in the survey).

firms as treatment. Column (2) uses a dummy variable if this short term credit ratio is larger than average, and column (3) uses the ratio itself as a continuous measure. Both columns show a very similar effect than our comparisons of domestic to foreign firms, and the magnitude is also similar: More credit constrained firms cut long-term investment relatively more.

Credit Constraints or Demand Shock?

A concern for identification might be that the credit squeeze in Spain went hand in hand with a recession. Are we picking up the effects of credit constraints as opposed to a pure demand shock without financial frictions? We do not think so, for the following reasons. First of all, theoretically we should expect the effects of a demand shock to go into the *opposite* direction, which means that our estimates are a lower rather than an upper bound for estimating credit shocks. Second, our empirical evidence is not consistent with a purely recession-driven explanation.

To see why in theory demand shocks should decrease short-term rather than longterm investments, consider that a demand shock means that consumers want to consume less now and save for the future instead. The increased demand for savings reduces interest rates. Firms will also invest less because demand is lower, but falling interest rates also reduce the opportunity cost of investment in general. The presence of the (temporary) demand shock provides a differential effect depending on the time-to-payoff of the investment: It is optimal to invest more in long-term investment (which increase output later when the demand shock is over) rather than short-term investment (which increase output now when demand is low). Overall long-term investment should fall *less* than short term investment (or even increase), which is the opposite of what we find.

We can adjust our model to illustrate this differential effect of a demand shock. Consider that a demand shock lowers the output in the current period, i.e. the output of short-term investments, by factor $\gamma < 1$. The maximization problem becomes:

$$\max_{k_t, z_t} E_t \left[\gamma f(k_t) + \beta \lambda_{t+1} (1+\rho) f(z_t) - q_t k_t - q_t z_t \right]$$
(2.4.2)

and the first order condition can be rearranged to yield:

$$E_t \left[f'(k_t) \right] = \beta E_t \left[(1 - \tau_{t+1}) \left(1 + s \right) (1 + \rho) f'(z_t) \right]$$
(2.4.3)

with τ_{t+1} as before and $s = \frac{1}{\gamma} - 1 > 0$ acting as a subsidy rather than a tax on long-term investment. The recession might be longer lived than just one period, but as long as we expect the recession to end at some point, it will differentially reduce short-term investment by more.

While this is reassuring, Bernanke (1983) suggests a different way how a recessions can influence investment behavior: He argues uncertainty increases during a recession, which decreases investment. If uncertainty about expected returns on long-run investments increases by more than uncertainty about expected returns on short-run investments, we might expect to see a differential effect in line with our results, i.e. a fall in long-term investments during the recent recession. Also, some empirical evidence in the literature points out that certain investment types exhibit a cyclical behavior. For example, R&D has been found to be pro-cyclical, see Barlevy (2007). However, our triple differences specifications allow us to include category-year fixed effects (which is impossible in the simple difference in differences analysis as it is collinear with the interaction term) to control for recession driven changes in the composition of long and short term investments, and our results are robust to this inclusion, so we do not believe a differential impact of the recession on different investment types are driving the results.

Furthermore, we would not expect to see a differential effect of the crisis on longterm investments only for more credit constrained firms, unless our measures for credit constraints (i.e. foreign ownership and share of short-term debt just before the crisis) are correlated with a larger exposure to the demand shocks. This concern might be true for our measure of foreign ownership (which might be correlated with export behavior and therefore differential exposure to other markets), but our matching estimates control for exporting and size and therefore compare companies with similar exposure to demand shocks. Our measure of short term credit is driven by debt maturity which seems unlikely to be correlated with different exposure to demand shocks. In any case, a negative demand shock should be reflected in the output of the firms. So in Table 2.B.12 we compare the output of credit constrained firms to the output of unconstrained firms. The results are in line with our claim. Column (1) shows that more credit constrained firms (panel A measures credit constraints by ownership, panel B by the share of short term credit) reduce total investment by more than unconstrained firms. Column (2) reflects our earlier results and shows that credit constrained firms reduce long term investments by even more (using the share of the two investment types with the highest durability, machinery and furniture) in total investment as measure of long term investment - a cruder measure as the one we used before). However, there is no differential impact on sales: Column (3) uses the log of total sales as dependent variable, which has a very small and statistically insignificant coefficient. There is also no differential effect on exports in column (4) or the propensity to export in column (5). There is also no sign of a differential creditworthiness or quality of the firms, as both pay the same credit cost, as shown in column (6).⁹ There is also no significant difference in the underlying quality of firms, as measured by TFP in column (7).

All this together we read as evidence that our empirical specifications are indeed measuring the effect of credit constraints, as opposed to demand shocks.

2.5 Conclusions

We have shown how to measure the extent of a credit crunch by analyzing changes in the composition of investment *within firms*. Intuitively, the extent to which firms are

⁹Note that the magnitude of the coefficients are very small; e.g. the coefficient of 0.096 would mean that Spanish firms after the crisis pay 0.09% higher credit cost, but it is insignificant.

altering the composition of investment away from longer time-to-payoff towards more immediate payoff is a measure of the risk that the firms perceive of facing liquidation due to lack of access to cash over the relevant period. In this sense, our measure of the credit crunch yields a clearly identified economic parameter which is readily interpretable: the credit shock is equivalent to a 11% additional tax on the investment with the longest payoff horizon.

We have tested the hypothesis underlying our methodology by conducting a wide range of robustness and alternative specification tests. Our results have proven remarkably resilient to quite demanding alternatives, such as including firm size times investment duration, category year fixed effects, etc. in addition to the firm times year fixed effects which we include already in the baseline specification. We have also studied the linkage we proposed by analyzing whether the effects are particularly strong for firms that are a priori expected to suffer stronger from the credit crunch: domestic firms, and firms with more maturing debt. Indeed, the effects are stronger for these sets of firms.

Our results suggest that the breakdown of the single European capital market is likely to have long term effects on Spanish firms. Spanish firms which are affected by the credit squeeze cut investments with a medium to long term payoff, such as R&D, innovation and capital investment, by more than investment with a short-term payoff such as advertising. Credit constraints force Spanish firms to eat up their future and act as if only the immediate future, tomorrow, mattered. This is likely to have a long term impact on the Spanish economy, impeding recovery after the financial crisis, and reducing long-term economic growth.

Methodologically, our analysis yields estimates of the impact of the crunch that can serve as input for other models. The analysis can be easily extended to other locations, crises and other capital choices, for example by comparing changes in the ratio of used versus new capital equipment, which are induced by the financial crisis to measure the cost of the crunch.

2.A Figures



Notes: This figure shows the coefficients of the regression in column (1) in Table 2.B.5 over time.

Figure 2.A.1: Change in composition of investment towards short term investments



Notes: This figure shows the coefficients of the regressions in columns (3) and (4) in Table 2.B.7 over time.

Figure 2.A.2: Investment change of long term investments by nationality of firm owner



Notes: Darker lines depict investment types with a longer time-to-payoff, for which we would expect a larger drop.





Note: The solid line shows the exit rate (firms that went bankrupt over active firms) separately for foreign and domestic firms. The dashed lines show the 95% confidence interval of the rates.

Figure 2.A.4: Exit rates by domestic and foreign firms

2.B Tables

-		Consolidated depreciation	
Investment type	Estimates in literature	rate	Rank
Advertising and	 Landes/Rosenfield (1994): >50% for most 	60%	1
brand equity	industries; up to 100% for some industries		
	Corrado/Hulten/Sichel (2009) conclude on 60%		
	from literature review, with some studies having		
	larger and smaller depreciation rates (lower		
	bound: Ayanian (1938) with 7 years)		
Software/IT	• Corrado/Hulten/Sichel (2009): 33% for own-	30%	2
	account software based on BEA (1994)		
	 Tamai/Torimitsu (1992): 9 year average 		
	lifespan, ranging from 2 to 20 years		
	• Spain accounting rules: 26% (IT equipment and		
	software)		
R&D	Corrado/Hulten/Sichel (2009): 20% based on	20%	3
	literature review		
	• Pakes/Schankerman (1984): 25% based on 5		
	European countries		
	 Pakes/Schankerman (1986): 11-12% for 		
	Germany, 17-26% for UK, 11% for France		
	 Nadiri/Prucha (1996): 12% 		
	• Bernstein/Mamuneas (2006): 18%-29% for		
	different US industries		
Vehicles	Spain accounting rules: 16%	16%	4
Machinery	Spain accounting rules: 12%	12%	5
	• BEA accounting rules: 10.31%-12.25%		
Furniture &	• Spain accounting rules: 10%	10%	6
office equipment	• BEA accounting rules: 11.79%		
Buildings	Spain accounting rules: 3%	n/a*	n/a*
	• BEA: 2-3% (industrial and office buildings)		
Land	• Spain accounting rules: depends on land prices	n/a*	n/a*
	• BEA: depends on land prices		

Table 2.B.1: Depreciation rates of different investment types

Notes: Spanish accounting rules are given in Table 2, "Tabla simplificada" of http://www.individual.efl.es/ActumPublic/ActumG/MementoDoc/MF2012_Coeficientes% 20an-uales%20de%20amortizacion_Anexos.pdf

* As the real estate crisis coincided with the credit crunch and resulted in large price drops in rest estate

(e.g. of up to 90% in land), we have chosen conservatively not to include this in our analysis, see the text.

	Me	an January)		
	(Standar Before crisis (2003-2007)	After crisis (2008-2010)	Change	Change in %
Investment categories, mn EUR (ordered by depreciation rate)				
Advertising	150.99	118.78	-32.21**	21.3%**
C C	(9.86)	(12.79)		
IT	6.20	3.86	-2.34**	-37.7%**
	(0.52)	(0.53)		
R&D	1.12	1.05	-0.07	-6.3%
	(0.13)	(0.16)		
Vehicles	4.20	6.10	1.90	45.2%
	(0.60)	(2.33)		
Machinery	198.57	141.78	-56.79***	-28.6%***
	(13.86)	(13.50)		
Furniture and office equipment	37.73	33.98	-3.75	-9.9%
	(4.70)	(5.35)		
Buildings	40.18	23.17	-17.01***	-42.3%***
	(4.71)	(2.42)		
Land	7.47	5.57	-1.90	-25.4%
	(0.75)	(1.29)		
Credit				
Credit ratio (total credit/	0.57	0.54	-0.03***	-4.4%***
total assets)	(0.00)	(0.00)		
Credit cost*, %	4.06	4.28	0.22***	5.4%***
	(0.02)	(0.03)		

Table 2.B.2: Summary statistics

* Total cost of a credit (including interest rates, but also other fees) as a percentage of obtained credit.

DEPENDENT VARIABLE: ln(investment value)	(1)	(2)	(3)	(4)
(1/depreciation rate)* after 2008 dummy After 2008 dummy	-0.012** (0.005) -0.170*** (0.031)	-0.012** (0.005)	-0.020*** (0.006)	-0.072*** (0.016)
Observations Partial R-squared Category FEs Firm FEs Year FEs	43,900 0.710 YES YES	43,900 0.711 YES YES YES	43,900 0.582 YES	88,331 0.153 YES
Firm*year FEs Including 0's		120	YES	YES YES

Table 2.B.3: Main results

Notes: The dependent variable is log of investment of firm i in year t in investment category c, where investment categories include 6 investment types: advertising, IT, R&D, vehicles, machinery, and furniture & office equipment. The main regressor is an interaction term of the inverse of the depreciation rate (a measure for time-to-payoff) of a investment category as a measure of the time-to-payoff of an investment type as given by Table 1 and a time dummy variable that indicates the financial crisis (=1 in and after 2008). All standard errors are clustered at the firm level, allowing for autocorrelation across time and across investment categories within the firm. *** p<0.01, ** p<0.05, * p<0.1. Observations are between 2003 and 2010. Column (4) codes the 0 euro observations as 1 euro in order to use the 0s as well, and shows our estimates are conservative without the 0s.

Table 2.B.4: Robustness checks: different measures for time-to-payoff

DEPENDENT VARIABLE: ln(investment value)	(1)	(2)	(3)	(4)	(5)	(6)
(Time-to-payoff measure)*	-0.020***	-0.028***	-0.062***	-0.089***	-0.097***	0.264***
after 2008 dummy	(0.006)	(0.008)	(0.014)	(0.022)	(0.022)	(0.089)
Observations	43,900	43,900	43,900	43,900	43,900	43,900
R-squared	0.582	0.582	0.582	0.582	0.582	0.582

Notes: This table replicates the final specification in the previous table, but uses different measures for time-to-payoff (larger the more long-term an investment) across columns:

(1) 1/depreciation rate; as before

(2) Rank of investment type according to depreciation rate (higher rank for more long term investments)

(3) 4 categories: short term (advertising; value 1), short/mid term (R&D, IT; value 2), long/mid term (vehicles; value 3), long term (machinery, furniture; value 4)

(4) 3 categories: short term (advertising; value 1), mid term (R&D, IT; value 2), long term (vehicles, machinery, furniture; value 3)

(5) 3 categories: short term (advertising; value 1), mid term (R&D, IT, vehicles; value 2), long term (machinery, furniture; value 3)

(6) Depreciation rate (note that the sign reverses, as a higher depreciation rate means a shorter time-to-payoff)

All columns include a full set of firm*year fixed effects to capture any demand specific effects driven by the crisis, as well as category fixed effects. All standard errors are clustered at the firm level, allowing for autocorrelation across time and across investment categories within the firm. *** p<0.01, ** p<0.05, * p<0.1. Observations are between 2003 and 2010.

DEPENDENT VARIABLE:	(1)	(2)	(3)	(4)	(5)
ln(investment value)					
(year==2004)*	-0.003	-0.005	-0.016	-0.021	-0.035
(Time-to-payoff measure)	(0.008)	(0.010)	(0.019)	(0.028)	(0.030)
(year==2005)*	-0.003	-0.006	-0.019	-0.013	-0.047
(Time-to-payoff measure)	(0.009)	(0.011)	(0.021)	(0.032)	(0.034)
(year==2006)*	0.005	0.006	0.004	0.010	-0.007
(Time-to-payoff measure)	(0.009)	(0.012)	(0.023)	(0.034)	(0.036)
(year==2007)*	0.004	0.005	-0.003	-0.009	-0.025
(Time-to-payoff measure)	(0.009)	(0.012)	(0.023)	(0.035)	(0.036)
(year==2008)*	-0.003	-0.004	-0.021	-0.034	-0.060
(Time-to-payoff measure)	(0.010)	(0.013)	(0.024)	(0.036)	(0.038)
(year==2009)*	-0.026**	-0.037***	-0.088***	-0.123***	-0.149***
(Time-to-payoff measure)	(0.010)	(0.013)	(0.025)	(0.038)	(0.040)
(year==2010)*	-0.032***	-0.046***	-0.101***	-0.134***	-0.161***
(Time-to-payoff measure)	(0.010)	(0.014)	(0.025)	(0.039)	(0.040)
Observations	43,900	43,900	43,900	43,900	43,900
R-squared	0.583	0.583	0.583	0.583	0.583
F tests on equality					
of coefficients:					
F-stat 2007=2008	0.752	0.751	0.888	0.752	1.476
p-val	0.386	0.386	0.346	0.386	0.224
F-stat 2007=2009	11.46	12.85	15.10	11.72	13.10
p-val	0.001	0	0	0.001	0
F-stat 2007=2010	15.86	17.86	18.99	13.34	14.80
p-val	0	0	0	0	0

Table 2.B.5: Placebo tests: treatment effect by year

Notes: This table replicates the specifications of Table 4, i.e. each column uses a different measure for time-to-payoff, with the exception that it allows the interaction term to vary in each year. All columns include a full set of firm*year fixed effects to capture any demand specific effects driven by the crisis, as well as category fixed effects. All standard errors are clustered at the firm level, allowing for autocorrelation across time and across investment categories within the firm. *** p<0.01, ** p<0.05, * p<0.1. Observations are between 2003 and 2010.

DEPENDENT VARIABLE: Credit ratio (between 0 and 1)	(1)	(2)	(3)	(4)	(5)
Dummy if after crisis	-0.015***	-0.016***	0.006		
(year>=2008)	(0.004)	(0.004)	(0.010)		
Interaction term (Spanish firms) *			-0.025***	-0.025***	-0.035***
(after 2008)			(0.010)	(0.010)	(0.010)
ln(industry exports to EU)		0.021	0.024	0.058	
		(0.037)	(0.037)	(0.040)	
ln(industry exports to World)		-0.009	-0.013	-0.056	
		(0.037)	(0.037)	(0.041)	
ln(industry output)		0.013	0.014	0.005	
		(0.013)	(0.013)	(0.015)	
Observations	13 915	13 915	13 897	13 897	13 897
Partial R-squared	0.004	0.005	0.007	0.002	0.003
Number of firms	2 650	2 650	2 650	2 650	2 650
Firm FEs	YES	YES	YES	YES	YES
Year FEs	1 LO	TLC	120	YES	TLU
Ind*year FEs				120	YES

Table 2.B.6: Mechanism: Credit squeeze

Notes: This table shows that the financial crisis in 2008 triggered a credit squeeze, which especially affected Spanish owned firms. The dependent variable is credit ratio (total credit divided by total assets, ratio between 0 and 1). The main regressors are a dummy variable that indicates the financial crisis (=1 in and after 2008), and an interaction term of a Spanish ownership dummy (defined by <=50% foreign ownership in same year) and the crisis dummy variable. All columns include firm fixed effects. Columns (2)-(4) also control for industry specific demand variables (Spanish exports to EU, Spanish exports to the World, domestic value added per industry) to capture industry specific demand shocks of the recession driven by the financial crisis. Export data is from the WITS database provided by the Worldbank, and Spanish value added per industry is from National Accounts data provided by the Spanish National Institute of Statistics (INE). Column (4) includes year fixed effects to capture common time effects (the crisis dummy is absorbed by these fixed effects driven by the crisis (our industry controls are absorbed by these fixed effects and therefore omitted). Column (5) includes a full set of industry*year specific fixed effects and therefore omitted). All standard errors are two-way clustered at the firm and industry*year level. *** p<0.01, ** p<0.05, * p<0.1. Observations are between 2003 and 2010.

DEPENDENT VARIABLE:	(1)	(2)	(3)	(4)
ln(investment value)	(domestic only)	(foreign only)	(domestic only)	(foreign only)
Long term investment*	-0.029***	0.018		
after 2008 dummy	(0.006)	(0.014)		
(year==2004)*			-0.004	-0.002
Long term investment			(0.009)	(0.015)
(year==2005)*			-0.009	0.009
Long term investment			(0.010)	(0.018)
(year==2006)*			0.002	0.005
Long term investment			(0.010)	(0.020)
(year==2007)*			0.003	-0.003
Long term investment			(0.010)	(0.021)
(year==2008)*			-0.016	0.045**
Long term investment			(0.011)	(0.022)
(year==2009)*			-0.035***	0.006
Long term investment			(0.011)	(0.023)
(year==2010)*			-0.043***	0.005
Long term investment			(0.011)	(0.024)
Observations	35,346	8,479	35,346	8,479
R-squared	0.566	0.661	0.566	0.661
Sample	domestic firms	foreign firms	domestic firms	foreign firms
F tests on equality				
of coefficients:				
F-stat 2007=2008			4.719	7.697
p-val 2008			0.030	0.006
F-stat 2007=2009			15.25	0.213
p-val 2009			0	0.644
F-stat 2007=2010			20.51	0.153
p-val 2010			0	0.696

Table 2.B.7: Foreign versus domestic firms

Notes: This table conducts our preferred specification separately for domestic firms (<50% foreign owned) and foreign firms (>50% foreign owned). Columns (1) and (2) pool the effect across post crisis years, columns (3) and (4) show the effect evolving over time. All columns include a full set of firm*year fixed effects to capture any demand specific effects driven by the crisis, as well as category fixed effects. All standard errors are clustered at the firm level, allowing for autocorrelation across time and across investment categories within the firm. *** p<0.01, ** p<0.05, * p<0.1. Observations are between 2003 and 2010.

DEPENDENT VARIABLE: ln(investment value)	(1)	(2)	(3)	(4)
(Time-to-payoff measure)*	-0.020*	0.020		
(after 2008 dummy)	(0.012)	(0.019)		
(Time-to-payoff measure)*		-0.049***	-0.052***	-0.152***
(after 2008 dummy)*(domestic firm dummy)		(0.019)	(0.020)	(0.058)
(Time-to-payoff measure)*		0.033**	0.034**	-0.314***
(domestic firm dummy)		(0.017)	(0.017)	(0.048)
Observations	41,550	41,475	41,475	88,223
R-squared	0.582	0.583	0.585	0.005
Firm [*] year FEs	YES	YES	YES	YES
Category FEs	YES	YES		
Category*year FEs			YES	YES
Including 0's				YES

Table 2.B.8: Difference-in-difference-in-differences

Notes: This table implements a triple difference estimation using the interaction of time-to-payoff (here measured by the inverse of the depreciation rate), after crisis dummy (=1 in 2008 and after) and a domestic firm dummy (defined by <=50% foreign ownership in the same year). Note that the domestic firm dummy variable is not time invariant, as it changes with ownership changes. However, there are very few of those in the data, and they are not driving our results. All standard errors are two-way clustered at the firm and industry*year level. *** p<0.01, ** p<0.05, * p<0.1. Observations are between 2003 and 2010. Column (4) codes the 0 euro observations as 1 euro in order to use the 0s as well, and shows our estimates are conservative without including the 0s.

DEPENDENT VARIABLE:	(1)	(2)	(3)	(4)	(5)	(6)	(7)
ln(investment value)							
(Time-to-payoff measure)*	0.034**	0.031*	0.005	0.003	0.063	0.007	0.008
(domestic firm dummy)	(0.017)	(0.017)	(0.041)	(0.023)	(0.038)	(0.015)	(0.016)
(Time-to-payoff measure)* (after 2008	-0.052***	-0.049***	-0.066*	-0.061**	-0.074*	-0.045***	-0.042**
dummy)*(domestic firm dummy)	(0.020)	(0.019)	(0.039)	(0.024)	(0.039)	(0.018)	(0.018)
(Time-to-payoff measure)*	. ,		. ,	. ,	. ,	-0.013***	. ,
ln(sales)						(0.004)	
(Time-to-payoff measure) * (after 2008						0.002	
dummy)*ln(sales)						(0.006)	
(Time-to-payoff measure)*							-0.014***
ln(TFP)							(0.005)
(Time-to-payoff measure) * (after 2008							0.001
dummy)*ln(TFP)							(0.007)
5,							· · ·
Observations	41,475	22,909	3,181	11,318	23,965	41,475	38,791
Partial R-squared	0.001	0.001	0.001	0.001	0.004	0.003	0.002
Number of firmyr	11,028	5,731	710	2,584	6,302	11,028	10,300
Firm*year FEs	YES	YES	YES	YES	YES	YES	YES
Category*year FEs	YES	YES	YES	YES	YES	YES	YES

Table 2.B.9: Robustness checks foreign versus domestic firms

Notes: This table implements a triple difference estimation using the interaction of time-to-payoff (here measured by the inverse of the depreciation rate), after crisis dummy (=1 in 2008 and after) and a domestic firm dummy (defined by <=50% foreign ownership in the same year). The regressions in the columns are as follows:

(1) Baseline, all companies

(2) Only firms that belong to a corporate group

(3) Only firms that have foreign non-industrial plants

(4) Only firms that have foreign shareholdings

(5) Inverse propensity score reweighting based on sales and exports in all the years between 2003 and 2007 (before crisis years)

(6) Uses the baseline specification, but adds another interaction term with size as measured by ln(sales)

(7) Uses the baseline specification, but adds another interaction term with size as measured by ln(TFP); with productivity estimation by industry with Levinsohn-Petrin method, corrected for input and output price changes

All standard errors are two-way clustered at the firm and industry*year level. *** p<0.01, ** p<0.05, * p<0.1. Observations are between 2003 and 2010.

DEPENDENT VARIABLE: Long term credit/total credit (between 0 and 1)	(1)	(2)	(3)	(4)	(5)
-					
Dummy if after crisis (year>=2008)	0.051***	0.050***	0.050***		
	(0.005)	(0.005)	(0.011)		
Interaction term (Spanish firms) *			0.000	0.001	0.000
(after 2008)			(0.012)	(0.012)	(0.013)
In(industry exports to EU)		-0.071	-0.073	-0.042	
		(0.046)	(0.046)	(0.044)	
ln(industry exports to World)		0.060	0.061	0.037	
		(0.045)	(0.045)	(0.047)	
ln(industry output)		-0.050***	-0.050***	-0.029	
		(0.018)	(0.018)	(0.020)	
Observations	14,410	14,410	14,392	14,392	14,392
Partial R-squared	0.029	0.032	0.033	0.002	0.001
Number of firmid	2,707	2,707	2,707	2,707	2,707
Firm FE	YES	YES	YES	YES	YES
Year FE				YES	
Ind*Year FE					YES

Table 2.B.10: Liabilities by maturity

Notes: The dependent variable is long term credit divided by total credit (ratio between 0 and 1). The main regressors are a dummy variable that indicates the financial crisis (=1 in and after 2008), and an interaction term of a Spanish ownership dummy (defined by <=50% foreign ownership in same year) and the crisis dummy variable. All columns include firm fixed effects. Columns (2)-(4) also control for industry specific demand variables (Spanish exports to EU, Spanish exports to the World, domestic value added per industry) to capture industry specific demand shocks of the recession driven by the financial crisis. Export data is from the WITS database provided by the Worldbank, and Spanish value added per industry is from National Accounts data provided by the Spanish National Institute of Statistics (INE). Column (4) includes year fixed effects to capture common time effects (the crisis dummy is absorbed by these fixed effects and therefore omitted). Column (5) includes a full set of industry*year specific fixed effects and therefore omitted). All standard errors are two-way clustered at the firm and industry*year level. *** p<0.01, ** p<0.05, * p<0.1. Observations are between 2003 and 2010.

DEPENDENT VARIABLE: ln(investment value)	(1)	(2)	(3)
(Time-to-payoff measure)*	-0.052***	-0.042***	-0.086***
(after 2008 dummy)*(treatment)	(0.020)	(0.011)	(0.027)
(Time-to-payoff measure)*	0.034**	0.020*	0.019
(treatment)	(0.017)	(0.011)	(0.029)
Observations	41,475	36,135	36,135
R-squared		0.585	0.584
Firm*year FEs	YES	YES	YES
Category*year FEs	YES	YES	YES

Table 2.B.11: Short term credit before crisis

Notes: This table implements a triple difference estimation using the interaction of time-to-payoff (here measured by the inverse of the depreciation rate), after crisis dummy (=1 in 2008 and after) and a treatment variable. This specification is equivalent to the triple difference estimation conducted in column (3) in Table 8. Column (1) uses domestic firms as treatment variable as in our baseline specification. Column (2) uses a dummy variable if short term credit with financial institutions/total credit is larger than average in 2007. Column (3) uses the ratio of short term credit with financial institutions/total credit in 2007 as treatment variable. All standard errors are two-way clustered at the firm and industry*year level. *** p<0.01, ** p<0.05, * p<0.1. Observations are between 2003 and 2010.

	(1) , , , ,	, , (2) ,	(3)	(4)	(5)	(9)	(2)
	In(total inv)	long term inv/ total inv	ln(sales)	ln(exports)	export dummy	Credit cost, %	ln(TFP)
ANEL A. SPANISH FIRMS							
nteraction term (Spanish	-0.147**	-0.065***	-0.032	-0.027	0.001	0.096	-0.031
irms) * (after 2008)	(0.068)	(0.013)	(0.029)	(0.059)	(0.011)	(0.188)	(0.026)
bservations	12,990	12,990	14,414	9,064	14,414	3,584	13,001
t-squared	0.001	0.002	0.000	0.000	0.001	0.004	0.000
Jumber of firmid	2,549	2,549	2,710	1,763	2,710	1,002	2,340
irm FEs	YES	YES	YES	YES	YES	YES	YES
nd*year FEs	YES	YES	YES	YES	YES	YES	YES
ANEL B. SHORT TERM CRI	EDIT BEF	ORE CRISIS					
nteraction term (Short term	-0.108*	-0.038***	-0.011	0.022	-0.014	0.037	-0.018
redit dummy) * (after 2008)	(0.056)	(0.013)	(0.018)	(0.050)	(0.010)	(0.098)	(0.016)
bservations	11,429	11,429	12,608	7,964	12,851	3,131	12,050
artial R-squared	0.001	0.001	0.000	0.000	0.000	0.000	0.000
Jumber of firmid	1,925	1,925	2,005	1,331	2,005	805	1,995
irm FEs	YES	YES	YES	YES	YES	YES	YES
nd*vear FEs	χ_{FS}	$\chi_{\rm ES}$	YFS	γES	$\gamma_{\rm ES}$	$\chi_{\rm ES}$	$\chi_{\rm FS}$

Table 2.B.12: Robustness checks: demand shock

equipment. (2) share of long term investment (machinery and furniture/office equipment) in total investment (ratio between 0 and 1), (3) log of total sales, (4) log of exports (only for exporting firms), (5) exporter dummy, (6) credit cost of firms (in %, i.e. 1 means 1%), (7) log productivity of firms (productivity estimation by industry with Levinsohn-Petrin method, corrected for input and output Notes: The dependent variables are different for each column: (1) log of total investment (sum of investment over the six investment types advertising, IT, R&D, vehicles, machinery and furniture/office price changes).

Panel A and panel B show the results of two separate regressions. The main regressors are the interaction term of either a Spanish ownership dummy in panel A (defined by <=50% foreign ownership in same year) or a dummy variable if short term credit with financial institutions/total credit is larger than average in 2007 in panel B and a dummy variable that indicates the financial crisis (=1 in and after 2008). All columns include firm fixed effects and a full set of industry*year specific fixed effects to capture any demand specific effects driven by the crisis. All standard errors are two-way clustered at the firm and industry*year level. *** p<0.01, ** p<0.05, * p<0.1. Observations are between 2003 and 2010.

Chapter 3

The roles of import competition and export opportunities for technical change

A variety of empirical and theoretical trade papers have suggested and documented a positive impact of trade on the productivity of firms. However, there is less consensus about the underlying mechanism at work. While trade theory focuses on a market access mechanism, other papers point out that import competition also matters. This paper conducts a "horse race" between export opportunities and import competition, focusing on the heterogeneity in firm responses. Using Spanish firm level data, instrumenting for exports and imports with tariff changes, and applying quantile regressions to account for selection, I find robust evidence that access to export markets leads to productivity increases, but only for firms that were already highly productive before. The evidence on import competition is weaker, with possibly some initially low-tech firms managing to increase their productivity in response to increased competition from abroad, but responses are very heterogeneous. Productivity upgrades are driven by increased R&D, patenting, product innovation, and the adaptation of certain technologies like CAD. There is no evidence that either mechanism leads to increased full time employment, instead full time workers seem to be replaced by part-time or temporary workers.

3.1 Introduction

A variety of empirical and theoretical trade papers have established a positive impact of trade on the productivity of firms (e.g. Bustos 2011; Pavcnik 2002; Bloom et al. 2011). However, there is less clarity about the underlying mechanism at work. Trade theory focuses almost exclusively on a market access mechanism, i.e. the opportunity to access new export markets increases firms' revenues and makes productivity enhancing investments more profitable. On the other hand, empirical and theoretical papers in industrial organization point out that increased (import) competition induces firms to increase their productivity. This paper wants to disentangle the potential productivity enhancing effects of import competition and access to export markets.

There are several reasons why it is important to distinguish between the "import competition" and the "export opportunities" mechanism: First, it is plausible to think that each mechanism affects a different type of firm. For example, the market access mechanism is only relevant for firms that are eventually able to export, or at least think they might have the capacity to do so. It is widely known in the trade literature that these are typically only the most productive firms in an industry. There is less consensus about which firms might be induced by import competition to upgrade their productivity. On one hand, import competition could be more relevant for firms with low profits, which are close to the border of bankruptcy and want to avoid exit. On the other hand, import competition could discourage already unproductive firms from trying harder, whereas already productive firms might be induced to innovate in order to escape competition. Second, each mechanism might have a different implication on employment. Larger export markets give firms the opportunity to expand, while increasing competition might force firms to contract.

This paper picks up the focus on firm heterogeneity prevalent in the trade literature since Melitz (2003) by highlighting the heterogeneous response of firms to trade shocks, more specifically the response by initial level of productivity. Furthermore, this paper also sheds some light on the way how firms are increasing their productivity, by exploiting a rich Spanish data set on a number of productivity enhancing measures. Summarizing, this paper attempts to answer the following questions: Do import competition and/or access to export opportunities induce firms to increase their productivity? Is there a heterogeneity in the productivity response? And finally, how do firms achieve productivity increases?

Several trade theory papers show that firms could be induced by export opportunities to update their productivity (e.g. Melitz and Costantini 2007; Bustos 2011; Atkeson and Burstein 2010). In standard trade models with CES preferences markups are constant and unaffected by import competition. With variable markups as in Melitz and Ottaviano (2008), import competition decreases the incentive to innovate as profits fall. However, there are a variety of (not necessarily trade) models that predict an increase in innovation resulting from competition, for example relying on trapped factors (Bloom et al. 2013), non-profit maximizing managers and X-inefficiency (Aghion et al. 1999; Horn et al. 1995), a differential impact of competition on post- and pre-innovation rents (Aghion et al. 2005), or imitation based on a search model (Waugh et al. 2014). So, in theory both export opportunities and import competition could lead to productivity upgrades of firms.

The vast majority of empirical papers in this area focus on either the (import) competition mechanism or the access to export market mechanism. Several papers in empirical IO have shown that increased competition (from either domestic or foreign entrants) increase the productivity of incumbent firms (e.g. Aghion et al. 2009; Blundell et al. 1995; Bloom et al. 2011; Tybout 2004; also Aghion et al. 2014 in a lab experiment). On the other hand, papers in empirical trade have shown that exporting leads to increased productivity (e.g. Lileeva and Trefler 2010; De Loecker 2007).

However, increased export opportunities often coincide with increased competition from abroad, as many trade liberalization episodes are bilateral and increase trade in both directions (even within narrowly defined industries). Regressions focusing on either the import or the export side omit an important variable, and might be picking up the productivity effect coming from the other mechanism. This paper avoids this problem and distinguishes empirically between the import competition and export access effect. Only a few papers focus on the response of productivity to both increased imports and exports. Pavcnik (2002) finds productivity increases among incumbents in import-competing industries but not in exporting industries, but does not allow for both forces to affect different firms within the same industry. In contrast, Trefler (2004) incorporates both import and export tariffs of the US and Canada and finds in firm level regressions that only the export effect is a significant driver of productivity increases, but he is not able to control for a selection effect.

The literature also suggests heterogeneity in productivity responses for both export access and import competition. For example, Aghion et al. (2005) provide evidence that competition discourages firms with initially lower productivity from innovating but encourages already productive firms to innovate more in order to escape competition. In contrast, Waugh et al. (2014) show in a theory model that firms at the lower end of the productivity distribution are more incentivized by import competition to update productivity because they can gain more, but they do not provide empirical evidence for this mechanism. Bustos (2011) and Lileeva and Trefler (2010) show that mostly large firms increase productivity as response to export access. This paper allows for heterogeneous responses to both the export access and import competition mechanism.

The main empirical specification in this paper is a regression of productivity on both imports and exports of an industry as measures for import competition and export opportunities to address the omitted variables bias, and interactions of those measures with a firm's initial productivity to allow for heterogeneous responses. I estimate this regression in first differences to take out firm specific, time-invariant factors, and add firm level fixed effect and year fixed effects to consider only deviations from firm specific growth rates. Both import and export changes are then instrumented with changes in domestic and foreign trade tariffs to address any remaining endogeneity issue. The Spanish data set allows me to use firm specific input and output price changes to obtain a measure of total factor productivity that is not driven by changes in markups. In robustness checks I show that import competition is not picking up the effect of better access to imported inputs, and that firms which increase their productivity because they are exposed to export opportunities increase their exports at the extensive and intensive margin, but not their output prices.

Empirical papers focusing on import competition face yet another problem, as they might falsely pick up an effect on productivity, because firms that are hit very hard by a negative productivity shock exit the sample. Without correcting for this selection effect, the effect of trade on productivity might be overestimated. I propose a simple fix to this problem using quantile instead of mean regressions. If exiting firms are assumed to have been hit by very low productivity shocks, and there are not too many of them, the median effect on productivity will be unaffected by exits, while the mean effect is biased.

A final contribution of this paper lies in shedding light on the way how firms increase their productivity. Most empirical papers focus on labor productivity or total factor productivity, exceptions are Bloom et al. (2011) who include patents, IT spending and R&D spending for subsets of their data, and Bustos (2011) who distinguishes between process and product innovation. The rich Spanish firm level data used in this paper includes most of these variables such as R&D investment, patenting, product innovation, but also other variables such as adaptation of foreign technologies or implementation of other technologies (e.g. computer-aided design CAD).

The results suggest that empirical papers focusing on import competition pick up the effect of access to export markets by omitting this variable. There is robust evidence that access to export markets leads to productivity increases, but only for firms that were already highly productive to begin with. The evidence on import competition is weaker and heterogeneous, with possibly some initially low-tech firms managing to upgrade their productivity in response to increased competition from abroad.

Productivity upgrades are driven by increased R&D, patenting, and product innovation, and the adaptation of certain technologies like CAD. There is no evidence that either mechanism leads to increased full time employment, instead full time workers seem to be replaced by part-time or temporary workers.

The remainder of the paper is organized as follows: Section 3.2 contains a description of the used Spanish firm level data of manufacturing firms that allows for TFP estimation and the analysis of firm exits, and provides a rich set of outcome variables. Section 3.3 provides an overview of Spain's trade flows, which have grown strongly over the observed 15 years. Section 3.4 describes the empirical strategy and section 3.5 discusses and interprets the empirical results. Section 3.6 concludes.

3.2 Description of Data

This paper uses panel data from a Spanish survey of manufacturing firms (ESEE; Encuesta Sobre Estrategias Empresariales), that is collected by the Fundación SEPI, a foundation affiliated with the Spanish Ministry of Finance and Public Administration.¹ The survey is designed to cover a representative sample of Spanish Manufacturing firms and includes around 1,800 firms per year. Participation of firms with more than 200 employees is required, while firms with more than 10 but less than 200 employees are sampled via a stratified sampling approach. SEPI makes a great effort to replace non-responding and exiting firms to ensure the continuing representativeness of the sample, leading to a total number of around 4,000 observed firms between 1993 and 2007.

The most distinctive feature of this data set is the very rich information it provides on several important dimensions: Detailed capital stock and investment needed for TFP estimation; input and output price changes to distinguish TFP changes from markup changes; information on exits (distinct from non-response) and entry to deal with selection; and a wide variety of productivity related activities such as R&D, patenting, and the adaptation of certain technologies (e.g. use of robots, computer aided design, flexible manufacturing systems).

Total factor productivity. We need detailed data on capital stock, output, employment and intermediate inputs to estimate TFP at the firm level. In many firm level data sets capital stock is not available and must be reconstructed using investment data (often using only average depreciation rates). The problem of a missing initial capital stock is only negligible if data over a long period of time is available and initial capital stock is depreciated for much of the observed sample period. Fortunately, the Spanish data set provides both gross and net capital stock together with firm level depreciation and investment, which allows a precise construction of the capital stock at any point in time.

Estimation of total factor productivity with OLS suffers from several problems: Employment and capital choices are endogenous, and TFP cannot easily be distinguished from markup changes (Beveren 2012). I use the well-established Levinsohn-Petrin procedure to deal with the endogeneity problem, which uses intermediate inputs to control for unobserved expected productivity changes. This is preferable to the Olley-Pakes method which uses investment as control, because investment is often reported as zero, which casts doubt on the validity of the assumed monotonicity condition requiring that investment is strictly increasing in unobserved productivity. The monotonicity condition is more likely to be satisfied for intermediate inputs, as firms usually report positive numbers.

Beveren (2012) points out that policy evaluations are usually robust to the TFP estimation method, with one exception: It is necessary to control for input and output prices (Loecker 2011). Luckily, the Spanish firm level survey provides a remedy to this omitted price bias as it also asks for input and output prices. Firms are asked by how

¹For more information see http://www.fundacionsepi.es/esee/sp/spresentacion.asp

much % the sales price of its products and the purchasing price of its intermediate inputs and services has changed compared to the previous year. The price changes are a weighted average across final products and markets (for output prices), and a weighted average across intermediate inputs, energy consumption and purchased services (for input prices), which I use to deflate output and intermediate inputs at the firm level (instead of usually used industry-wide deflators).

The results in this paper are robust to using different productivity measures such as labor productivity and simple TFP fixed effects estimation.

Other productivity related activities. In order to understand what firms do in order to increase productivity, I use a variety of productivity related activities such as R&D expenditure, number of filed patents, and product innovation. Every four years the survey contains additional questions about the use of specific technologies. These technologies are the use of robots, computer aided design (CAD), flexible manufacturing systems, and whether firms make an effort to assimilate foreign technologies.

Trade and tariff data. This paper exploits the variation in industry-specific imports and exports over time. I merge the firm-level data with industry level trade data from COMTRADE using the NACECLIO industry classification of firms (20 NACECLIO categories²). Section 3.3 provides an overview of Spain's imports and exports over time and by industry. I use the tariffs that the EU imposes on imports from the rest of the world ("import tariffs") and tariffs that other countries impose on imports from the EU ("export tariffs") as instrumental variable for trade flows. All trade and tariff data used in the analysis is from COMTRADE (provided by UNSD) and TRAINS (provided by UNCTAD); all data sets are accessed via the WITS software provided by the World Bank.³

3.3 Trade Growth in Spain

This paper uses import and export growth in Spain between 1993 and 2007 as source of variation in access to export markets and import competition. Trade has grown substantially over this time period. Figure 3.A.1 shows that Spain's world exports in 2007 (171 bn EUR) were almost four times larger than in 1993 (45 bn EUR). Spain's world imports grew even stronger, from 60 bn EUR in 1993 to 265 bn EUR in 2007. The trade growth was not entirely linear (neither in levels nor in logs): It increased strongly in the 90's, slowed down in the early 2000's, and picked up again around 2003 (especially imports). Spain incurred a trade deficit from goods trade in every single year over the observed time period. The trade deficit increased from around 2% of GDP to a staggering 10% of GDP at the end of the sample period.

Most of Spain's trade is with the European Union: In 2007, 73% of Spain's exports

²The 20 industries are: Meat related products; Food and tobacco; Beverage; Textiles and clothing; Leather, fur and footwear; Timber; Paper; Printing and publishing; Chemicals; Plastic and rubber products; Nonmetal mineral products; Basic metal products; Fabricated metal products; Industrial and agricultural equipment; Office machinery, data processing, precision instruments and similar; Electric materials and accessories; Vehicles and accessories; Other transportation materials; Furniture; Miscellaneous.

³http://wits.worldbank.org/wits/

went to EU25 countries, compared to 60% of imports. Figure 3.A.2 graphs Spain's trade with its most important trading partners over time. A large share of exports are destined for Spain's neighboring countries in Europe: France, Germany, Italy, Portugal and Great Britain. Among those, exports to France have increased the most over the sample period. Most of Spain's imports are also coming from EU countries: Germany has the largest import share, followed by France and Italy. However, imports from China have been skyrocketing: China's share in Spain's imports has increased from 2% in 1993 to 6% in 2007. Table 3.B.1 shows that the rise of China is most prevalent in certain industries like leather/fur/footwear, textiles/clothing, and nonmetal mineral products.

Table 3.B.2 lists the top 3 export destinations by industry. Portugal has become a prominent destination for computer products and electronics; printing and publishing; leather, fur and footwear; and nonmetal mineral products. But also United Kingdom and the United States have increasingly become the destination for Spain's exports.

Overall, the distribution of trade across industries has remained fairly stable over time. Figure 3.A.3 shows that the most important export and import industry is vehicles and accessories, covering a quarter of exports and one fifth of imports. Second is the Chemicals sector, followed by industrial and agricultural equipment. The food industry ranks fourth among Spanish exports. The dominant industries have only become more dominant over time.

Researchers have attributed the increased trade with EU countries to European integration such as the introduction of the euro, the European Single Market, and the European Monetary Union (e.g. Berger and Nitsch 2008; Bergin and Lin 2012; Brouwer et al. 2008). Another important trade liberalization episode that occurred during the sample period was China's accession to the WTO in 2001, which was accompanied by a fall of tariff and non-tariff barriers between China and the European Union, among others, and led to increased trade with China (e.g. Bloom et al. 2011).

3.4 Empirical Strategy

Import competition versus export opportunities. How do access to export markets ("export opportunities") and competition from foreign firms ("import competition") affect firm productivity? In order to test this, I estimate the following firm-level equation that relates a firm's productivity (*TFP*) to measures of export access (*EXP*) and import competition (*IMP*):

$$TFP_{ist} = \beta_0 + \beta_1 IMP_{st} + \beta_3 EXP_{st} + \text{yearFE} + \text{firmFE} + \varepsilon_{ist}$$
(3.4.1)

where *i* indicates the firm, *s* indicates one of 20 industries (NACECLIO classification), and *t* is year (between 1993 and 2007). The main empirical measure for TFP_{ist} is obtained via Levinsohn-Petrin estimation (using material inputs to control for unobserved productivity) with an extra adjustment for changes in input and output prices, as described above.

Spain's world imports at the industry level are used as proxy for competition from foreign firms to domestic firms IMP_{st} (instead of a firm's actual imports, because firm imports are inputs and not outputs, but I want to measure import competition at the level of a firm's end product). Similarly, access to export markets EXP_{st} is proxied by Spain's world exports in industry *s* and year *t* (instead of actual firm level exports, to proxy for *potential* export opportunities). Year fixed effects control for unobserved common time trends, and firm level fixed effects control for unobserved, time-invariant heterogeneity at the firm level. All variables are in logs.

I proceed by estimating this equation in first differences, which controls for unobserved firm heterogeneity, as the firm level fixed effects cancel out. However, industries might have both higher productivity growth and higher trade growth for reasons other than export access or import competition. I add industry fixed effects to avoid using the cross-sectional variation and exploit the time-variation in trade within industries instead. In the main specifications I also add firm level fixed effects (which absorb the industry fixed effects) to control for firm characteristics that affect the productivity growth rate:

$$\Delta TFP_{ist} = \beta_1 \Delta IMP_{st} + \beta_3 \Delta EXP_{st} + \text{yearFE} + \text{firmFE} + \nu_{ist}$$
(3.4.2)

All standard errors are clustered at the industry level, in the spirit of Bertrand et al. (2004).

Heterogeneous effects. Since Melitz (2003) the trade literature has focused on firm level heterogeneity, usually with respect to the initial productivity of firms. In models with heterogeneous firms we do not expect all firms to be affected in the same way by import competition and export opportunities. For example, import competition might affect firms with initially lower productivity by more, because the threat of bankruptcy is stronger for them. On the other hand, firms with already low productivity might be discouraged by import competition, and only firms with a high enough productivity level to start with might even try to push the productivity frontier further out to become productivity leader. Similarly, we should expect a heterogeneous response to a better access to export markets. For example, the trade literature finds that usually only the most productive firms export.

In order to test for a heterogeneous response of firms depending on their initial productivity level I interact changes in import competition and export access with *TFP* in the first year of the analysis (in 1993, denoted as *TFP*93_{*i*}):

$$\Delta TFP_{ist} = \beta_1 \Delta IMP_{st} + \beta_2 (TFP93_i \cdot \Delta IMP_{st}) + \beta_3 \Delta EXP_{st} + \beta_4 (TFP93_i \cdot \Delta EXP_{st}) + \text{yearFE} + \text{firmFE} + \eta_{it} (3.4.3)$$

As robustness check I add industry*year specific fixed effects, which are collinear with the main effects of import competition and export access, but still allow me to estimate the sign on the interaction terms, β_2 and β_4 .

Endogenous trade flows. A potential threat to identification is the endogeneity of industry level imports and exports. For example, there might be reverse causality: If firm productivity in a sector is high, this might lead to fewer imports (e.g. due to Ricardian comparative advantage of Spain) and more exports in that industry. Industry (or firm) level fixed effects in regression 3.4.1 avoid such cross-sectional comparisons across industries with different productivity levels and focus on within industry comparisons. However, even within industries reverse causality might hold over time: Imports are higher and exports lower when sectoral TFP is lower. Regression 3.4.2 addresses some of this endogeneity concern by using year fixed effects to absorb common time trends and industry (or firm) level fixed effects to look only at deviations from the industry (firm) specific TFP growth rate.

If these sets of fixed effects cannot address the endogeneity problem fully, this would lead us to underestimate the effect of import competition and overestimate the effect of access to export markets. Many empirical trade papers use exogenous trade liberalization events and/or tariff reductions to instrument for trade. Very often trade liberalization episodes lead to a bilateral reduction of tariffs, therefore own tariff changes are usually highly correlated with tariff changes abroad. I will use the word "import tariff" to refer to Spain's tariffs, and "export tariff" to refer to import tariffs that foreign countries impose on imports from Spain. The high correlation between import and export tariffs means that using tariff changes as instruments will not solve the omitted variable bias problem in papers studying the effect of only import competition or export access, as import tariff changes are not uncorrelated with changes in exports and vice versa. Any study focusing on the effects of export or imports on firms needs two instruments, one for exports and one for imports that are not too highly correlated. Econometrically, with multiple endogenous regressors it is not sufficient for identification to have a significant first stage for each endogenous regressor, because each first stage might use the same level of exogenous variation. Instead, the matrix of first stages needs to be of full rank to ensure identification. The Kleibergen-Paap statistics (Kleibergen and Paap 2006, implementation in Stata provided by Baum et al. 2010) provides a rank test to check this.

I use the tariffs that the EU imposes on imports from the rest of the world and the tariffs that other countries impose on imports from the EU as instruments for trade. More specifically, I use the maximum tariff⁴ for each product category (ISIC Rev. 3; 244 product categories) and aggregate them to NACECLIO industries using the import shares of each product within an industry and across countries. Empirically, some tariff changes are more relevant for trade patterns than others, i.e. some are "binding" and inhibit trade, whereas other tariff changes do not seem to be as relevant for trade. In order to use only binding tariff changes which have an impact on trade flows (and therefore a strong first stage), I multiply the weighted tariffs with an "importance weight" if tariff changes led to trade changes in the previous period. The lagged importance weights are given by $w = -\Delta \tan i f * \Delta \ln(\operatorname{trade})$ if w > 0, and 0 otherwise.

⁴Changes in the maximum tariff were empirically the changes with the most relevant impact on trade.

It is not always clear whether tariff changes can be interpreted as exogenous to firms and industries, as large companies often try to influence policy makers to negotiate favorable tariffs. However, in the case of Spain tariffs are negotiated at the European level, and it is less likely that Spanish firms are able to influence European decision making. Furthermore, many tariff changes are a part of a larger political process, for example, EU enlargement, or China's WTO accession, and therefore likely out of control for specific Spanish firms.

Selection. There is a potential selection bias in the estimation above because some firms went bankrupt during the sample period. It is plausible to assume that these firms exit because they have been hit hard by a negative productivity shock. Omitting these firms from the sample and carrying out the estimation on the surviving firms which experienced relatively more positive productivity shocks would lead to an overestimation of the effect on productivity. This problem might be most severe for the productivity estimation of the firms that have a very low productivity to begin with, because negative productivity shocks are even more likely to cause bankruptcy for them.

The threat from selection bias to estimation can be eliminated by estimating the median instead of the mean effect. Identification of the median will then depend on an assumption that seems reasonable: Exiting firms are assumed to have been hit by a very negative productivity shock, resulting in a censored dependent variable.⁵ Censoring the dependent variable from below a certain point has no effect on the estimating quantile regressions that are above this point. Therefore, as long as there are not too many exiting firms, the median effect and other quantiles are still identified. In our sample there are only between 0.1% and 2.5% exits per year. I pursue a censored median regression approach as in Powell (1986) by assigning exiting firms the lowest observed productivity change in the exit year. The median is only identified if censoring is only on one side of the conditional median. Buchinsky (1994) proposes a estimation algorithm to ensure this condition. The proposed estimation algorithm (which amounts to checking that the predicted values for productivity changes for exiting firms are above the censored values) converges already in the first round in all of my estimations.⁶

As in the mean regressions I would like to allow for firm and year fixed effects in the quantile regressions. However, estimating quantile regressions with large number of fixed effects is tricky. For example, it is not valid to transform variables to deviations from means, as the conditional quantile function is not linear. Estimating a large number of fixed effects is computationally very intensive. Furthermore, a large number of fixed effects increases the variability of other estimates. In order to solve this problem,

⁵Note that while the data set allows me to distinguish between non responding and exiting firms, exiting firms include closed firms, firms in liquidation, but also firms that are taken over by other firms. The last category of firms might not necessarily have been hit by a negative productivity shock, if firms like to take over only targets that are very productive ("cherry picking"). However, there is also evidence that take over targets are very unproductive ("lemon grabbing"), because they have a potential of high returns after a successful turnaround (Gelübcke 2012). In the cherry picking case, I might be underestimating the true effect on productivity. In any case, note that the number of exiting firms is very small (between 0.1% and 2.5% per year), and take overs are even rarer, so this should not affect the results very much.

⁶This approach is also proposed and discussed in Angrist and Pischke (2009), chapter 7.1.1.

I follow an approach suggested by Koenker (2004) using penalized quantile regression and sparse linear algebra.⁷

Besides being robust to selection, quantile regressions are also insightful as to the heterogeneity in productivity responses to import competition and export access.

Robustness checks. In the remaining part of the paper I perform a variety of robustness checks. For example, I show that the results are not sensitive to the way of measuring productivity. I also provide evidence that the import competition mechanism is not capturing an increased access to imported inputs, and that increased export opportunities indeed lead to larger exports and a higher propensity to exports at the firm level, without affecting output prices.

Furthermore, I show what type of activities Spanish firms undertake to increase productivity, by looking at a large number of other variables such as R&D investment, patenting, product innovation, adaptation of foreign technologies or implementation of other technologies. Finally, I use employment as a dependent variable to check whether import competition and export opportunities have any or differing effects on firm employment.

3.5 Main Empirical Results

Import competition versus export opportunities. Table 3.B.3 conducts a "horse race" between import competition and access to export markets by estimating equation 3.4.2 to see whether either one affects firm level productivity. Columns (1) and (2) regress productivity on import competition and export access separately, and both regressions show a similar sized, significant effect of trade on productivity. However, if including both measures in column (3), the coefficient on import competition falls and becomes insignificant, showing that the regression in column (1) suffered from severe omitted variable bias and captured the joint effect of both imports and exports. As imports and exports are usually highly correlated, there is a substantial omitted variable bias in empirical studies focusing only on import competition or export opportunities. A similar argument holds when comparing columns (2) and (3), but the effect of export opportunities on productivity remains significant.

Adding industry fixed effects in column (4) and firm fixed effects in column (5) makes the difference between import competition and export access even more pronounced. Access to export opportunities has a strong and significant positive effect on firm level productivity. If exports increase by 10%, average firm level productivity increases by 1.1%. On the other hand, import competition has a much smaller and insignificant average effect. A regression focusing only on import competition is actually

⁷Ideally I would like to combine an IV approach with the quantile regression approach described before, but it is not easy to implement IV quantile regression with more than one endogenous regressors and a large number of fixed effects. Since the OLS and IV mean estimations showed qualitatively similar results, I conduct simple quantile regressions without controlling for potential endogeneity of the regressors at this point. In the future, a possible way to deal with this endogeneity could be to implement a control function approach for censored quantile regressions, for example as described in Blundell and Powell (2007).

picking up the effect of export opportunities. However, this does not necessarily mean that there is no economic role of import competition, as the average effect washes out heterogeneity in the response of firms.

Heterogeneous effects. Not all firms might react in the same way to increased import competition or increased export opportunities. It might even be that the insignificant coefficient on import competition hides a heterogeneous reaction that averages out. Table 3.B.4 therefore checks for heterogeneous effects by adding interaction terms with the firm's initial productivity level (in year 1993) as in equation 3.4.3. An interesting pattern emerges in column (1): While the overall average effect of import competition is positive but insignificant, the firms with the lowest initial productivity levels actually do increase their productivity. This effect fades out as firms' initial productivity increases. At the same time, while the overall average effect of export opportunities is positive and significant, it turns out that it is really only firms with an initially already high level of productivity that are driving these results. This should not come as a surprise, as the trade literature on heterogeneous firms (e.g. Melitz 2003) predicts that only firms with high productivity are exporting, and the empirical literature (e.g. Bernard et al. 2007) find that exporting firms are usually the most productive firms in an industry.

This pattern is even more pronounced in column (2) which adds firm fixed effects. Column (3) adds industry*year fixed effects that absorb imports and exports as well as any difference across industries and over time. However, the interaction effects can still be estimated and their coefficients are robust to this inclusion.

In order to interpret the coefficients in Table 3.B.4, I plot the predicted change in productivity from import competition and export opportunities by initial productivity for the observed sample of firms. Figure 3.A.4 uses my preferred specification in column (2) of Table 3.B.4 and shows that the average annual increase in import competition (11.1%) leads firms with the lowest productivity level to increase their productivity by 2.4%. The effect from export opportunities is even stronger, but only for the firms with the highest productivity levels: The average annual increase in export opportunities (10.9%) makes highly productive firms increase their productivity by 3.5%. Scaled up to the overall observed increase in trade over the 15 years in the sample, import competition could have been responsible for a 35% productivity increase for low-TFP firms, and export opportunities could have been responsible for a 53% productivity increase for high-TFP firms (Table 3.B.5).

Endogenous trade flows. If trade flows depend on the productivity growth in an industry, our estimates might be biased. Therefore I use weighted import and export tariff changes (constructed as described in section 3.4) as instruments for import competition and export opportunities in Table 3.8.6, and the interaction of tariff changes with initial TFP as instruments for the interaction terms. Column (1) repeats my preferred OLS specification in column (2) in Table 3.8.4, including both firm and year fixed effects. Column (2) in Table 3.8.6 shows the instrumented version of that regression. The Kleibergen-Paap statistics reported in the last row is 22.10, confirming

a strong joint first stage. Table 3.B.7 reports the four first stages associated with column (2). Import tariffs have a negative impact on imports, and the same is true for exports, as expected. This relationship also holds for the interaction terms.

The magnitudes of the IV estimates is larger compared to the OLS results, pointing to measurement error in imports and exports in the OLS regression.⁸ The effect of import competition is no longer significant, but the effect of access to export markets remains significant. According to the IV estimates in column (2), the average annual increase in export opportunities increased the productivity of high-TFP firms by 9.8% (compared to 3.5% in the OLS estimation). Columns (3) and (4) add industry*year fixed effects, omitting the main effects, but the findings in terms of the interaction terms are very similar.

Measurement of productivity. The results are not sensitive to the way of measuring productivity. For example, in Table 3.B.8 I use labor productivity or TFP estimates obtained via simple fixed effects regression (of sales on employment and capital including firm and year fixed effects) as dependent variable, and the results are very similar. If anything, the Levinsohn-Petrin measure of TFP is more conservative than the other measures, probably because it addresses the endogeneity of input use best, and also corrects for input and output price changes.

Selection. As previously explained, we might still be overestimating the effect of trade on productivity because there is a selection effect coming from the exits of unproductive firms. As a first simple step I assign the lowest observed productivity change to an exiting firm in a given industry and year, and include these firms in the OLS and IV regressions in columns (5) and (6) of Table 3.B.6. As there are not very many exiting firms, the sample increases by less than 2%. Figure 3.A.5 shows that exiting firms tend to have a lower initial productivity as continuing firms, but there is still quite a wide dispersion. Compared to the regressions without exiting firms, the estimates remain largely unchanged. In the IV regressions, the import competition effect becomes significant when including exiting firms. But of course this might still overestimate the true effect of import competition or export access on productivity, if exiting firms have a much lower productivity change than the lowest observed in the sample.

Quantile regressions are robust to the precise value that is assigned to exiting firms, as long as the share of exiting firms is smaller than the estimated quantile. But if the dependent variable is skewed, as in the case with productivity changes, the mean and the median will be different, and I start out with simple comparisons of mean and quantile regressions on the sample without exiting firms in Table 3.B.9. Panel A looks only at the main effects of import competition and export access. While the mean effect of import competition on productivity is small and insignificant, the effect at the median is positive and significant. Other quantiles show that there is large heterogeneity in the observed responses. A similar picture emerges for the effect of

⁸It is not that surprising to find measurement error in export and imports. Other studies have found that trade flows reported by the importing and exporting country often differ, and attribute this to misreporting.
export opportunities, but the effect is stronger at lower percentiles, and flattens out for higher percentiles. The median effect of export access is smaller than the mean effect.

Panel B include all the interaction terms, and again it becomes apparent that the effect of import competition on productivity is very dispersed. Only few low-TFP firms manage to upgrade their productivity as a result of import competition, while others end up with decreases in TFP. In fact, the effect of import competition at the median is negative for low-TFP firms, and larger for high-TFP firms. The effect of export access on productivity is also dispersed, but the direction of the effect is more similar across the quantiles: The high-TFP firms always show the largest TFP gains resulting from an increased access to export markets.

The quantile regressions reveal a large degree of heterogeneity in the productivity response of firms to trade liberalization. I now show the results of censored quantile regressions including exiting firms to check whether a selection effect drives these results. Table 3.B.10 replicates the regressions performed in Table 3.B.9 but includes exiting firms. Comparing the two tables, there is little change in the estimated effects, so selection does not seem to be a major explanation for observed effects on productivity.

Overall, the regressions conducted so far have shown a very robust effect of access to export markets on the productivity of firms, but only for firms that initially already had a high productivity. The effect of import competition on productivity is very dispersed and less robust, with some initially low productive firms increasing their TFP by a lot.

Robustness checks. Are those firms who are induced by export opportunities to update their productivity also the ones who actually increase their exports or start to export? Otherwise I might be not be capturing the right mechanism. To check this, column (1) in Table 3.B.11 uses firm level exports as dependent variable in the IV regressions from before (IV specification as in column (2) of Table 3.B.6). It is reassuring to see that the result is consistent with the export access mechanism and not picking up any spurious correlation, as firm level exports increase in response to increased export opportunities (driven by tariff changes), but only for the most productive firms. Column (2) uses the change in exporter status as a dependent variable, and the pattern is similar. Access to export markets increase exports at both the extensive and intensive margin. Column (3) checks whether the increase in exports is related to prices or quantities by using the change in output prices related to access to export markets. Note that the combined first stage is again strong enough in all three regressions, as indicated by the large Kleibergen-Paap statistics.

Another concern is that increased imports might not be only a measure of import competition, but also provide firms with the opportunity to use (potentially cheaper) imported goods as intermediate goods which might be reflected in TFP. The used instrumental variable captures tariff changes at the output level of an industry and should therefore not pick up variation in imports driven by inputs. However, tariff

⁹This variable refers to all products produced by a firm, not only for exported goods, which is not available in the survey.

changes for inputs and outputs of an industry might be correlated, and therefore the regression might still pick up a change in imported inputs. Table **3.B.12** uses firm level imports as dependent variable in the IV regressions to check this. Column (1) shows no significant change in firm level imports as response to import competition. Column (2) uses the change in the probability to import as dependent variable, and these results go into the opposite direction as the productivity results. Column (3) uses the change in intermediate input prices as dependent variable.¹⁰ Again, the results go into the opposite direction as the productivity results. Interestingly, these changes are also observed for firms who get access to export markets. The most productive firms which upgrade their productivity as response to export opportunities reduce imports and pay higher input prices which is the opposite of what we would expect to happen if our measures are picking up a better access to imported inputs.

Ways to increase productivity. What do firms actually do to increase their productivity? Table 3.B.13 checks whether firms that are exposed to increased import competition or get access to export markets engage in some activities that have the potential to increase TFP. Columns (1) and (2) show that firms upgrade their productivity by increasing R&D expenses and starting to engage in R&D. Firms also increase the number of patents and start to engage in patenting (columns (3) and (4)).¹¹ Column (5) shows that the firms also engage in product innovation, as they increase their number of products. For firms with access to export markets this could mean that they adapt their products to foreign tastes or standards. Firms under competition from abroad might be forced to develop new products in niches where they can still be competitive, or to adapt their products to compete with foreign firms.

Every four years the survey contains additional questions about the use of specific technologies. These technologies are the use of robots, computer aided design (CAD), flexible manufacturing systems, and whether they make an effort to assimilate foreign technologies in their production process. As these variables are only surveyed every four years, the sample size drops significantly, and the first stage becomes too weak to use IV regressions. In columns (6) and (7) I report therefore the OLS results of the adaptation of those technologies that had significant coefficients in the regressions: the use of CAD and the assimilation of foreign technologies. Interestingly, both of those technologies are only used by productive firms that are exposed to export opportunities.

Implications for employment. There is strong evidence for the productivity enhancing effect of export opportunities, and weak and heterogeneous evidence of the productivity enhancing effect of import competition. Both channels might have different implications for employment: Import competition might lead to labor-saving productivity increases, while access to export markets might lead to employment growth. Table 3.B.14 uses employment as the dependent variable in IV regressions. Interestingly, neither import competition nor access to export opportunities lead to significant employment changes. Note that although exports and sales increase for

¹⁰This variable refers to all inputs of a firm, not only imported goods, which is not available in the survey.

¹¹The results are unchanged if I normalize the number of filed patents by employment.

high-TFP firms that are exposed to export opportunities, labor productivity increases as well, leading to an insignificant net change in employment. However, when taking a closer look by employment type, e.g. full-time and part-time employment as well as temporary workers (who are not included in total employment), there is some evidence that low-TFP firms induced by import competition to upgrade their productivity have increased their temporary staff (potentially replacing full time workers, but the coefficient is not significant), whereas high-TFP firms induced by export opportunities to increase their productivity increased their part time staff (again, potentially replacing full time workers, but the coefficient is not significant). In any case, there is no evidence that employment increased as a result of increased trade.

3.6 Conclusions

Trade liberalization affects firms in several ways: One the one hand, firms get access to new export markets, providing them with an opportunity for growth. On the other hand, foreign firms enter the home market and create more competition for domestic firms. Both of these two "faces" of globalization might induce firms to upgrade their productivity. Increased export opportunities can make it worthwhile for firms to invest in new technology, while increased competitive pressure from foreign companies might force firms to engage in innovation in order to avoid bankruptcy or retain monopoly rents.

Existing papers have mostly focused on the productivity inducing effect of either export access or import competition. However, usually imports and exports are highly correlated (even within narrowly defined industries), and it is not clear whether empirical papers are picking up the right mechanism, or whether they suffer from omitted variable bias.

This paper disentangles the the two channels and finds strong productivity-enhancing effects from export opportunities. However, the effect is heterogeneous and depends on the initial productivity of firms: Only already very productive firms update their productivity when subject to new export opportunities. The evidence on import competition is weaker, with possibly some initially low-tech firms managing to increase their productivity in response to increased competition from abroad, but responses are very heterogeneous.

I use tariff changes as instruments for import and export changes to deal with the potential endogeneity of trade flows, and conduct censored quantile regressions which are robust to firm exits to control for selection. The result on export access is robust across all specifications, but the evidence on import competition is more mixed across different specifications. Most notably is the heterogeneous response of firms in quantile regressions: It seems that only few initially unproductive firms respond with large productivity increases to competition from abroad. Overall, productivity upgrades are driven by increased R&D, patenting, product innovation, and the adaptation of certain technologies like CAD.

My results suggest that papers studying the effects of either import competition or access to export markets need to control for the other channel in order to eliminate omitted variable bias. Furthermore, in order to deal with the endogeneity of imports and exports, it is necessary to find two instrumental variables, one for each trade direction. Trade liberalization episodes usually affect tariffs of both the importing and exporting country, making them highly correlated. It is therefore not sufficient to instrument for only exports or imports. The two instruments need to be sufficiently uncorrelated such that the matrix of first stages is full rank, which can be checked using the Kleibergen-Paap statistics on top of checking the F-statistics on each first stage separately.

Furthermore, the paper shows that median regressions can be a useful tool to deal with the selection problem of firms coming from firm exits, if one is willing to assume that the exiting firms are hit by the most unfavorable shock, an assumption that will often be sensible.

Finally, the paper supports the focus on the heterogeneity of firms that has been prevalent in the trade literature since Melitz (2003). The heterogeneity in firm responses to trade is astonishing, and the heterogeneous firm focus should also be reflected in the empirical trade literature.

There is no evidence that either mechanism leads to increased full time employment, instead full time workers seem to be replaced by part-time or temporary workers, which is probably disappointing for policy makers. Growth in firm size seems to be offset by increase in (labor saving) productivity, leveling out the effect on employment.

3.A Figures



Figure 3.A.1: Spain's world trade over time

Source: United Nations COMTRADE database, accessed by World Integrated Trade Solution (WITS), wits.worldbank.org



Figure 3.A.2: Spain's trade by country over time

Source: United Nations COMTRADE database, accessed by World Integrated Trade Solution (WITS), wits.worldbank.org



Figure 3.A.3: Spain's trade by industry over time

Figure 3.A.4: Estimated effect of import competition and export access on productivity growth of firms, by firms' initial productivity



Notes: The average annual increase in trade is 11.1% for imports and 10.9% for exports over the observed sample period between 1993 and 2007.



Figure 3.A.5: TFP distribution of all firms and exiting firms

Notes: The graph shows the kernel density of ln(TFP) in 1993 for all firms that exist in 1993 and separately for the subset of firms that exit at some later point in time. TFP is estimated by Levinsohn-Petrin method, adjusted for changes in input and output prices.

3.B Tables

		import share		import share			import share		import share
Industry	Country	1993	Country	2007	Industry	Country	1993	Country	2007
Total trade	France	17%	Germany	16%					
	Germany	15%	France	13%					
	Italy	8%	Italy	8%					
Meat related products	France	24%	France	21%	Nonmetal mineral products	France	20%	China	21%
	Netherlands	15%	Germany	15%		Germany	13%	Portugal	15%
	United Kingdom	10%	Netherlands	15%		Italy	13%	Italy	13%
Food and tobacco	France	21%	France	16%	Basic metal products	France	27%	France	17%
	Netherlands	9%	Germany	11%		Germany	17%	Germany	10%
	Germany	7%	Argentina	8%		United Kingdom	14%	China	9%
Beverage	United Kingdom	53%	United Kingdom	32%	Fabricated metal products	Germany	25%	Germany	22%
-	Netherlands	11%	France	12%		Italy	20%	Italy	17%
	France	10%	Italy	8%		France	19%	France	14%
Textiles and clothing	Italy	19%	China	22%	Industrial and agricultural equipment	Germany	24%	Germany	22%
-	France	11%	Italy	13%		Italy	21%	Italy	19%
	China	9%	Turkey	8%		France	14%	France	11%
Leather, fur and footwear	Italy	22%	China	35%	Computer products, electronics and op	United States	17%	Netherlands	16%
	China	15%	Italy	15%		Germany	14%	Germany	16%
	Korea, Rep.	7%	Vietnam	8%		France	13%	China	12%
Timber	United States	17%	Portugal	13%	Electric materials and accessories	Germany	19%	Germany	19%
	Portugal	16%	France	10%		France	14%	China	15%
	France	12%	China	10%		Japan	12%	France	9%
Paper	Finland	20%	France	20%	Vehicles and accessories	France	31%	Germany	34%
	France	16%	Finland	14%		Germany	30%	France	24%
	Germany	13%	Germany	14%		United Kingdom	8%	Italy	7%
Printing and publishing	United Kingdom	19%	United Kingdom	20%	Other transportation materials	United States	47%	France	21%
	Germany	17%	Germany	16%		Italy	10%	United States	20%
	Italy	12%	China	11%		Japan	8%	United Kingdom	12%
Chemicals	Germany	20%	Germany	17%	Furniture	Italy	22%	China	20%
	France	17%	France	14%		France	20%	Italy	14%
	United Kingdom	9%	United States	8%		Germany	18%	Germany	14%
Plastic and rubber products	France	27%	Germany	20%	Games & toys, sports instr	China	20%	China	22%
	Germany	23%	Italy	16%		Italy	14%	Germany	16%
	Italy	14%	France	16%		Japan	12%	United Kingdom	15%

Table 3.B.1: Top 3 import origins by industry

Note: Countries with the largest change are shaded.

Table 3.B.2:	Top 3 ex	port destina	ations by	industry
		1		

		export share		export share			export share		export share
Industry	Country	1993	Country	2007	Industry	Country	1993	Country	2007
Total trade	France	19%	France	20%					
	Germany	15%	Germany	11%					
	Italy	9%	Italy	9%					
Meat related products	France	32%	France	29%	Nonmetal mineral products	France	16%	France	19%
incut related products	Portugal	19%	Portugal	20%	nonnetar millerar products	Germany	10%	Portugal	10%
	Germany	12%	Germany	10%		United States	11%	United Kingdom	8%
Food and tobacco	Italy	16%	France	19%	Basic metal products	France	12%	France	15%
	France	10%	Italy	18%	busic metal products	Germany	10%	Italy	15%
	Portugal	10%	Portugal	15%		China	8%	Germany	15%
Beverage	Germany	10%	United Kingdom	14%	Fabricated metal products	France	17%	France	21%
beveluge	United Kingdom	14%	Germany	14%	rabilitatea inetai producio	Germany	13%	Germany	13%
	France	10%	France	9%		Portugal	1376	Portugal	12%
Textiles and clothing	France	15%	Portugal	15%	Industrial and agricultural equipment	France	15%	France	12%
rextiles and clothing	Portugal	13%	France	13%	industrial and agricultural equipment	Germany	12%	Germany	11%
	Italy	11%	Italy	9%		Portugal	7%	Portugal	8%
Leather, fur and footwear	Germany	19%	France	20%	Computer products, electronics and op	Germany	23%	Portugal	20%
	United States	17%	Portugal	10%		France	12%	France	10%
	France	15%	Italy	9%		Italy	9%	Germany	9%
Timber	France	21%	France	20%	Electric materials and accessories	Germany	19%	France	14%
	Portugal	15%	Portugal	20%		France	15%	Germany	13%
	United Kingdom	10%	United States	9%		Portugal	6%	Italy	9%
Paper	France	21%	France	22%	Vehicles and accessories	France	31%	France	32%
	Portugal	16%	Portugal	18%		Germany	21%	Germany	14%
	Germany	12%	Italy	9%		Italv	14%	United Kingdom	11%
Printing and publishing	Argentina	14%	France	23%	Other transportation materials	France	12%	France	14%
	Mexico	13%	Portugal	12%		Liberia	11%	United Kingdom	10%
	France	12%	Mexico	11%		Norway	11%	United States	9%
Chemicals	France	13%	Italy	13%	Furniture	France	26%	France	30%
	Germany	12%	France	12%		Germany	15%	Portugal	13%
	Italy	10%	Germany	11%		Portugal	11%	United Kingdom	6%
Plastic and rubber products	France	23%	France	23%	Games & toys, sports instr	France	20%	France	19%
	Germany	15%	Germany	13%		Germany	9%	Portugal	17%
	Portugal	9%	Portugal	11%		Portugal	9%	United States	8%

Note: Countries with the largest change are shaded.

	(1)	(2)	(3)	(4)	(5)
VARIABLES	$\Delta \ln (TFP)$				
$\Delta \ln (IMP)$	0.094*		0.065	0.035	0.015
	(0.049)		(0.047)	(0.052)	(0.053)
$\Delta \ln (EXP)$		0.096***	0.065**	0.085***	0.106***
		(0.034)	(0.025)	(0.024)	(0.028)
Observations	14 027	14 027	14 027	14.027	12 802
Observations	14,027	14,027	14,027	14,027	13,092
R-squared	0.001	0.001	0.002	0.002	0.002
Year fixed effects	YES	YES	YES	YES	YES
Industry fixed effects				YES	
Firm fixed effects					YES

Table 3.B.3: Import competition versus access to export markets

Notes: The dependent variable $\Delta \ln (TFP)$ denotes the change in log TFP (estimated by Levinsohn-Petrin method, adjusted for changes in input and output prices). The main regressors are $\Delta \ln (IMP)$ measuring the change in log of Spain's world imports, and $\Delta \ln (EXP)$ measuring the change in log of Spain's world exports, both at the NACECLIO industry level. * p<0.05, ** p<0.01, *** p<0.001. Standard errors in parentheses are clustered by NACECLIO industries.

	(1)	(2)	(3)
VARIABLES	$\Delta \ln (TFP)$	$\Delta \ln (TFP)$	$\Delta \ln (TFP)$
$\Delta \ln IMP$	0.385**	0.572**	
	(0.195)	(0.238)	
$(\Delta \ln IMP) * \ln (TFP_{93})$	-0.026	-0.040**	-0.026*
	(0.016)	(0.017)	(0.014)
$\Delta \ln EXP$	-0.521***	-0.493***	
	(0.198)	(0.188)	
$(\Delta \ln EXP) * \ln (TFP_{93})$	0.043***	0.043***	0.026*
	(0.014)	(0.014)	(0.015)
Observations	14,027	13,892	13,892
Partial R-squared	0.003	0.003	0.001
Year fixed effects	YES	YES	
Industry fixed effects	YES		
Firm fixed effects		YES	YES
Industry*year fixed effects			YES

Table 3.B.4: Heterogeneous effects

Notes: The dependent variable $\Delta \ln (TFP)$ denotes the change in log TFP (estimated by Levinsohn-Petrin method, adjusted for changes in input and output prices). The main regressors are $\Delta \ln IMP$ measuring the change in log of Spain's world imports, and $\Delta \ln EXP$ measuring the change in log of Spain's world exports, both at the NACECLIO industry level. Both main effects are interacted with the log of a firm's initial productivity in year 1993, $\ln (TFP_{93})$. * p<0.05, ** p<0.01, *** p<0.001. Standard errors in parentheses are clustered by NACECLIO industries.

Table 3.B.5: Estimated effect of import competition and export opportunities on productivity growth of firms, by firms' initial productivity

		% increase in trade				
% change in TFP	1%	11% (average annual)	160% (1993-2007)			
Import competition:						
Firms with lowest observed TFP	0.21%	2.36%	35.46%			
Firms with highest observed TFP	-0.19%	-2.10%	-31.4%			
Export opportunities:						
Firms with lowest observed TFP	-0.11%	-1.16%	-17.36%			
Firms with highest observed TFP	0.32%	3.54%	53.05%			

Notes: The predicted changes are calculated based on column (2) of Table 3.B.4.

DEPENDENT	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLE: $\Delta \ln (TFP)$	OLS	IV	OLS	IV	OLS	IV
$\Delta \ln IMP$	0.572**	1.448			0.564**	1.634*
	(0.238)	(1.007)			(0.253)	(0.846)
$(\Delta \ln IMP) * \ln (TFP_{93})$	-0.040**	-0.106	-0.026*	-0.0947	-0.0399**	-0.130**
	(0.017)	(0.068)	(0.014)	(0.0616)	(0.0184)	(0.0602)
$\Delta \ln EXP$	-0.493***	-2.711***			-0.376	-2.973***
	(0.188)	(0.753)			(0.247)	(0.699)
$(\Delta \ln EXP) * \ln (TFP_{93})$	0.043***	0.190***	0.026*	0.179***	0.0367**	0.215***
	(0.014)	(0.049)	(0.015)	(0.0512)	(0.0176)	(0.0436)
Observations	13.892	13.892	13.892	13.892	14.178	14.178
Firm fixed effects	YES	YES	YES	YES	YES	YES
Year fixed effects	YES	YES			YES	YES
Industry*year fixed effects			YES	YES		
Including exiting firms					YES	YES
First stage Kleibergen-Paap statistics		22.10		22.24		20.45

Table 3.B.6: Tariff changes as instrumental variables

Notes: The dependent variable $\Delta \ln (TFP)$ denotes the change in log TFP (estimated by Levinsohn-Petrin method, adjusted for changes in input and output prices). The main regressors are $\Delta \ln IMP$ measuring the change in log of Spain's world imports, and $\Delta \ln EXP$ measuring the change in log of Spain's world exports, both at the NACECLIO industry level. Both main effects are interacted with the log of a firm's initial productivity in year 1993, $\ln (TFP_{93})$. Columns (2), (4) and (6) use weighted tariff changes (multiplied by importance weights) as described in the text as instrumental variables for import competition and export access, and their interactions with initial productivity as instrumental variables for the interaction terms. Columns (5) and (6) include exiting firms and assigns them the lowest observed productivity change in their exiting year in their industry. * p<0.05, ** p<0.01, *** p<0.001. Standard errors in parentheses are clustered by NACECLIO industries.

	(1)	(2) $(\Lambda \ln IMP) *$	(3)	(4) $(\Delta \ln EXP) *$
VARIABLES	$\Delta \ln IMP$	$\ln (TFP_{93})$	$\Delta \ln EXP$	$\ln\left(TFP_{93}\right)$
$\Delta IMPTAR$	-2.733	4.563	-1.538	18.457
	(1.693)	(32.088)	(1.188)	(24.868)
$(\Delta IMPTAR) * \ln (TFP_{93})$	0.044	-2.386	-0.010	-2.948
	(0.098)	(2.155)	(0.072)	(1.899)
$\Delta EXPTAR$	-0.000	0.013	-0.006*	0.025
	(0.003)	(0.033)	(0.003)	(0.038)
$(\Delta EXPTAR) * \ln (TFP_{93})$	-0.000	-0.003*	-0.000	-0.009***
	(0.000)	(0.002)	(0.000)	(0.002)
Observations	13,892	13,892	13,892	13,892
Firm fixed effects	YES	YES	YES	YES
Year fixed effects	YES	YES	YES	YES

Table 3.B.7: First stages

Notes: This Table shows the four first stages of the IV regression in column (2) in Table 3.B.6. $\Delta IMPTAR$ and $\Delta EXPTAR$ are weighted tariff changes (multiplied by importance weights) as described in the text and used as instrumental variables for import competition and export access, and their interactions with initial productivity are used as instrumental variables for the interaction terms. * p<0.05, ** p<0.01, *** p<0.001. Standard errors in parentheses are clustered by NACECLIO industries.

	(1)	(2)	(3)	(4)	(5)	(6)
	Levinsoł	Levinsohn-Petrin		Labor productivity		effects
VARIABLES	OLS	IV	OLS	IV	OLS	IV
$\Delta \ln IMP$	0.572**	1.448	-0.0374	4.937**	0.283	6.683**
	(0.238)	(1.007)	(0.661)	(2.342)	(0.482)	(2.855)
$(\Delta \ln IMP) * \ln (TFP_{93})$	-0.040**	-0.106	0.00495	-0.448**	-0.0205	-0.579**
	(0.017)	(0.068)	(0.0566)	(0.210)	(0.0400)	(0.249)
$\Delta \ln EXP$	-0.493***	-2.711***	-0.781**	-3.652*	-1.006***	-6.922***
	(0.188)	(0.753)	(0.376)	(2.023)	(0.263)	(2.432)
$(\Delta \ln EXP) * \ln (TFP_{93})$	0.043***	0.190***	0.0830**	0.334*	0.0989***	0.595***
	(0.014)	(0.049)	(0.0339)	(0.172)	(0.0239)	(0.203)
Observations	13,892	13,892	15,100	15,100	14,134	14,134
Firm fixed effects	YES	YES	YES	YES	YES	YES
Year fixed effects	YES	YES	YES	YES	YES	YES
First stage Kleibergen-Paap statistics		22.10		23.11		26.26

Table 3.B.8: Alternative measures for productivity

Notes: Columns (1) and (2) replicate the specifications in columns (1) and (2) in Table 3.B.6, using the change in log TFP estimated by Levinsohn-Petrin method, adjusted for changes in input and output prices as dependent variable. Columns (3) and (4) use changes in log of labor productivity as dependent variable, where labor productivity is defined as total sales divided by total employment. Columns (5) and (6) use changes in log of TFP estimated by an OLS regression of sales on employment and capital, including firm and year fixed effects.

DEPENDENT	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLE:	Quantile regressions					
$\Delta \ln (TFP)$	Mean	0.1	0.25	Median	0.75	0.9
PANEL A. No interaction effects						
$\Delta \ln IMP$	0.015	-0.116***	0.031***	0.082***	0.139***	0.231***
	(0.053)	(0.058)	(0.031)	(0.018)	(0.031)	(0.055)
$\Delta \ln EXP$	0.106***	0.196***	0.099***	0.043***	0.018***	0.0002***
	(0.028)	(0.066)	(0.023)	(0.018)	(0.029)	(0.057)
PANEL B. Interaction effects						
$\Delta \ln IMP$	0.572**	-1.081***	-0.419***	-0.194***	0.318***	0.942***
	(0.238)	(0.296)	(0.196)	(0.144)	(0.191)	(0.236)
$(\Delta \ln IMP) * \ln (TFP_{93})$	-0.040**	0.066***	0.031***	0.020***	-0.012***	-0.053***
	(0.017)	(0.019)	(0.014)	(0.010)	(0.013)	(0.018)
$\Delta \ln EXP$	-0.493***	-0.773***	-0.571***	-0.229***	-0.023***	-0.052***
	(0.188)	(0.319)	(0.163)	(0.111)	(0.158)	(0.317)
$(\Delta \ln EXP) * \ln (TFP_{93})$	0.043***	0.067***	0.045***	0.017***	0.003***	0.004***
	(0.014)	(0.021)	(0.011)	(0.008)	(0.011)	(0.022)
Observations	13,892	13,892	13,892	13,892	13,892	13,892

Table 3.B.9: Comparing mean and quantile regressions

Notes: The dependent variable $\Delta \ln (TFP)$ denotes the change in log TFP (estimated by Levinsohn-Petrin method, adjusted for changes in input and output prices). The main regressors are $\Delta \ln IMP$ measuring the change in log of Spain's world imports, and $\Delta \ln EXP$ measuring the change in log of Spain's world exports, both at the NACECLIO industry level. Both main effects are interacted with the log of a firm's initial productivity in year 1993, ln (*TFP*₉₃). Panel A and Panel B show separate sets of regressions, with and without interaction terms. Exiting firms are not included in the sample in their exiting year, as productivity in the exiting year is unobserved. All regressions include full sets of firm level and year fixed effects. * p<0.05, ** p<0.01, *** p<0.001. Bootstrapped standard errors by firm in parentheses.

DFPFNDFNT	(1)	(2)	(3)	(4)	(5)	(6)
VADIARIE.	(1)	(2)	(0)	(1)	()	(0)
			Qual	ittle regress	510115	
$\Delta \ln (TFP)$	Mean	0.1	0.25	Median	0.75	0.9
PANEL A. No interaction effects						
$\Delta \ln IMP$	0.006	-0.196***	0.023***	0.079***	0.146***	0.226***
	(0.043)	(0.060)	(0.024)	(0.019)	(0.032)	(0.046)
$\Delta \ln EXP$	0.135***	0.282***	0.110***	0.050***	0.016***	0.006***
	(0.025)	(0.058)	(0.029)	(0.022)	(0.030)	(0.055)
PANEL B Interaction offects						
TAINED D. Interaction effects						
$\Delta \ln IMP$	0.564**	-1.208***	-0.458***	-0.214***	0.329***	0.985***
	(0.253)	(0.390)	(0.217)	(0.140)	(0.222)	(0.411)
$(\Delta \ln IMP) * \ln (TFP_{93})$	-0.040**	0.074***	0.033***	0.021***	-0.013***	-0.056***
	(0.018)	(0.029)	(0.015)	(0.010)	(0.015)	(0.028)
$\Delta \ln EXP$	-0.376	1.319***	-0.752***	-0.307***	-0.101***	-0.170***
	(0.247)	(0.303)	(0.162)	(0.122)	(0.173)	(0.344)
$(\Delta \ln EXP) * \ln (TFP_{93})$	0.037**	0.109***	0.058***	0.024***	0.008***	0.013***
	(0.018)	(0.022)	(0.011)	(0.008)	(0.013)	(0.024)
Observations	14,178	14,178	14,178	14,178	14,178	14,178

Table 3.B.10: Comparing mean and quantile regressions, including exiting firms

Notes: The dependent variable $\Delta \ln (TFP)$ denotes the change in log TFP (estimated by Levinsohn-Petrin method, adjusted for changes in input and output prices). The main regressors are $\Delta \ln IMP$ measuring the change in log of Spain's world imports, and $\Delta \ln EXP$ measuring the change in log of Spain's world exports, both at the NACECLIO industry level. Both main effects are interacted with the log of a firm's initial productivity in year 1993, $\ln (TFP_{93})$. Panel A and Panel B show separate sets of regressions, with and without interaction terms. Exiting firms are included in the sample; they are assigned the lowest observed productivity change in their exiting year in their industry. All regressions include full sets of firm level and year fixed effects. * p<0.05, ** p<0.01, *** p<0.001. Bootstrapped standard errors by firm in parentheses.

	(1)	(2)	(3)
	Change in	Change in	Change in
VARIABLES	ln(firm level exports)	exporter status	output price
$\Delta \ln IMP$	13.76	0.815	180.2
	(12.76)	(1.440)	(163.8)
$(\Delta \ln IMP) * \ln (TFP_{93})$	-0.915	-0.0209	-9.775
	(0.778)	(0.113)	(10.93)
$\Delta \ln EXP$	-6.220	-2.864***	-30.35
	(3.790)	(0.642)	(48.71)
$(\Delta \ln EXP) * \ln (TFP_{93})$	0.463**	0.163***	1.380
	(0.216)	(0.0400)	(3.124)
Observations	7 800	10 578	22 220
Firm fixed effects	VES	12,576 VES	23,329 VES
View fine die ffe etc	I E5 VEC	I E5 VEC	I ES
rear fixed effects	IES	165	1 ES
First stage Kleibergen-Paap statistics	12.12	34.44	19.16

Table 3.B.11: Firm level exports

Notes: The dependent variables are the change in log of firm level exports in column (1), change in exporter dummy variable in column (2) and percentage change in output prices (weighted average across all outputs) in column (3). The main regressors are $\Delta \ln IMP$ measuring the change in log of Spain's world imports, and $\Delta \ln EXP$ measuring the change in log of Spain's world exports, both at the NACECLIO industry level. Both main effects are interacted with the log of a firm's initial productivity in year 1993, ln (*TFP*₉₃). All regressions use tariff changes (multiplied by importance weights) as described in the text as instrumental variables for import competition and export access, and their interactions with initial productivity as instrumental variables for the interaction terms. * p<0.05, ** p<0.01, *** p<0.001. Standard errors in parentheses are clustered by NACECLIO industries.

	(1)	(2)	(3)
	Change in	Change in	Change in
VARIABLES	ln(firm level imports)	importer status	input price
$\Delta \ln IMP$	-31.61	-4.353**	199.6*
	(27.99)	(1.692)	(109.4)
$(\Delta \ln IMP) * \ln (TFP_{93})$	2.255	0.282**	-13.30*
	(1.936)	(0.120)	(7.236)
$\Delta \ln EXP$	34.48***	6.101***	-41.27
	(9.857)	(0.982)	(28.22)
$(\Delta \ln EXP) * \ln (TFP_{93})$	-2.349***	-0.389***	3.135*
	(0.678)	(0.0641)	(1.801)
Observations	7 646	12 502	23 291
Firm fixed effects	VES	VFS	VFS
Year fixed effects	YES	YES	YES
First stage Kleibergen-Paap statistics	36.38	33.17	19.01
i not stage i delle el gent i dap statistics	20.00	22.17	12.01

Table 3.B.12: Imported inputs versus import competition

Notes: The dependent variables are the change in log of firm level imports in column (1), change in importer dummy variable in column (2) and percentage change in input prices (weighted average across all inputs) in column (3). The main regressors are $\Delta \ln IMP$ measuring the change in log of Spain's world imports, and $\Delta \ln EXP$ measuring the change in log of Spain's world exports, both at the NACECLIO industry level. Both main effects are interacted with the log of a firm's initial productivity in year 1993, $\ln (TFP_{93})$. All regressions use tariff changes (multiplied by importance weights) as described in the text as instrumental variables for import competition and export access, and their interactions with initial productivity as instrumental variables for the interaction terms. * p<0.05, ** p<0.01, *** p<0.001. Standard errors in parentheses are clustered by NACECLIO industries.

VARIABLES	(1) $A \ln R \& D$	(2) Change in R&D dummy	(3) A#patents	(4) Change in patenting dummy	(5)	(6) Change in CAD dummy	(7) Assimilate imported technology
	Annae	dunniny	Anpaterito	dullility	Anpiouucio	dunnity	teennology
$\Delta \ln IMP$	20.08** (10.14)	3.822** (1.847)	233.1* (134.7)	-2.675** (1.162)	5.918* (3.434)	0.444 (0.355)	0.680 (0.497)
$(\Delta \ln IMP) * \ln (TFP_{93})$	-1.167*	-0.207**	-17.78*	2.035***	-0.397*	-0.0298	-0.0550
Alm FYD	(0.685) -9.625***	(0.0995) -4 812***	(9.916) -347 4***	(0.290) -0.127***	(0.233)	(0.0250)	(0.0371)
	(3.528)	(1.005)	(57.45)	(0.0210)	(1 412)	-0.000	(0.439)
$(\Delta \ln EXP) * \ln (TFP_{93})$	(0.235) 0.468** (0.235)	(1.000) 0.284^{***} (0.0631)	(37.43) 26.06*** (4.242)	(0.0210) 0.174^{**} (0.0844)	(1.412) 0.486^{***} (0.0909)	(0.050) 0.0635** (0.0248)	(0.439) 0.0594* (0.0340)
Observations Firm fixed effects Year fixed effects First stage Kleibergen- Paap statistics	4,292 YES YES 17.58	12,491 YES YES 33.34	12,582 YES YES 34.92	21,797 YES YES 21.05	12,523 YES YES 33.07	2,516 YES YES	2,562 YES YES

Table 3.B.13: Ways to increase productivity

Notes: The dependent variables are the change in log of total R&D expenditures (internal and external) in column (1), change in R&D dummy variable in column (2), change in number of patents in column (3), change in patenting dummy in column (4), change in number of products in column (5), change in dummy variable indicating whether the firm used computer aided design (CAD) in column (6), and change in a dummy variable to indicate whether a firm spent effort to assimilate imported technologies in column (7). The latter two variables are only asked every 4 years in the survey. The main regressors are $\Delta \ln IMP$ measuring the change in log of Spain's world imports, and $\Delta \ln EXP$ measuring the change in log of Spain's world imports, and $\Delta \ln EXP$ measuring the change in log of spain's world imports, and $\Delta \ln EXP$ measuring the change in log of spain's world imports, and $\Delta \ln EXP$ measuring the change in log of spain's world imports, and $\Delta \ln EXP$ measuring the change in log of spain's world imports, and $\Delta \ln EXP$ measuring the change in log of spain's world imports, and $\Delta \ln EXP$ measuring the change in log of spain's world imports, and $\Delta \ln EXP$ measuring the change in log of spain's world imports, and $\Delta \ln EXP$ measuring the change in log of spain's world exports, both at the NACECLIO industry level. Both main effects are interacted with the log of a firm's initial productivity in year 1993, $\ln (TFP_{93})$. Regressions (1) to (5) use tariff changes (multiplied by importance weights) as described in the text as instrumental variables for import competition and export access, and their interactions with initial productivity as instrumental variables for the interaction terms. Columns (6) and (7) are estimated via OLS (the first stage was too weak for IV estimation because of the reduced number of observations). * p < 0.05, ** p < 0.01, *** p < 0.001. Standard errors in parentheses are clustered by NACECLIO industries.

	(1)	(2)	(3)	(4)
VARIABLES	$\Delta \ln EMPL$	$\Delta \ln FULLTIME$	$\Delta \ln PARTTIME$	$\Delta \ln TEMP$
$\Delta \ln IMP$	-1.789	-4.416	0.229	11.19*
	(2.514)	(2.728)	(8.186)	(6.309)
$(\Delta \ln IMP) * \ln (TFP_{93})$	0.144	0.302	0.0377	-0.678
	(0.189)	(0.194)	(0.609)	(0.458)
$\Delta \ln EXP$	-1.372	1.641	-11.09**	-13.28
	(1.070)	(1.215)	(4.619)	(10.83)
$(\Delta \ln EXP) * \ln (TFP_{93})$	0.0891	-0.106	0.768**	0.814
	(0.0755)	(0.0824)	(0.312)	(0.754)
Observations	12,635	12,555	2,548	8,931
Firm fixed effects	YES	YES	YES	YES
Year fixed effects	YES	YES	YES	YES
First stage Kleibergen-Paap statistics	34.98	34.95	28.73	36.86

Table 3.B.14: Implications for employment

Notes: The dependent variables are the change in log of firm level employment (excluding temporary workers) in column (1), change in log of firm level full time employment (excluding temporary workers) in column (2), change in log of firm level part time employment (excluding temporary workers) in column (3), and change in log of firm level temporary workers in column (4). The main regressors are $\Delta \ln IMP$ measuring the change in log of Spain's world imports, and $\Delta \ln EXP$ measuring the change in log of Spain's world imports, and $\Delta \ln EXP$ measuring the change in log of spain's world imports, and $\Delta \ln EXP$ measuring the change in log of spain's world imports, and $\Delta \ln EXP$ measuring the change in log of spain's world imports, and $\Delta \ln EXP$ measuring the change in log of spain's world imports, and $\Delta \ln EXP$ measuring the change in log of spain's world imports, and $\Delta \ln EXP$ measuring the change in log of spain's world imports, and $\Delta \ln EXP$ measuring the change in log of spain's world imports, and $\Delta \ln EXP$ measuring the change in log of spain's world exports, both at the NACECLIO industry level. Both main effects are interacted with the log of a firm's initial productivity in year 1993, $\ln (TFP_{93})$. All regressions use tariff changes (multiplied by importance weights) as described in the text as instrumental variables for import competition and export access, and their interactions with initial productivity as instrumental variables for the interaction terms. * p < 0.05, ** p < 0.01, *** p < 0.001. Standard errors in parentheses are clustered by NACECLIO industries.

Bibliography

- Adrian, T., P. Colla, and H. S. Shin: 2012, 'Which Financial Frictions? Parsing the Evidence from the Financial Crisis of 2007 to 2009'. In: *NBER Macroeconomics Annual* 2012, Volume 27, NBER Chapters. National Bureau of Economic Research, Inc, pp. 159–214.
- Aghion, P., G.-M. Angeletos, A. Banerjee, and K. Manova: 2010, 'Volatility and growth: Credit constraints and the composition of investment'. *Journal of Monetary Economics* 57(3), 246–265.
- Aghion, P., S. Bechtold, L. Cassar, and H. Herz: 2014, 'The Causal Effects of Competition on Innovation: Experimental Evidence'. Working Paper 19987, National Bureau of Economic Research.
- Aghion, P., N. Bloom, R. Blundell, R. Griffith, and P. Howitt: 2005, 'Competition and Innovation: An Inverted-U Relationship'. *The Quarterly Journal of Economics* 120(2), 701–728.
- Aghion, P., R. Blundell, R. Griffith, P. Howitt, and S. Prantl: 2009, 'The Effects of Entry on Incumbent Innovation and Productivity'. *The Review of Economics and Statistics* 91(1), 20–32.
- Aghion, P., M. Dewatripont, and P. Rey: 1999, 'Competition, Financial Discipline and Growth'. *Review of Economic Studies* **66**(4), 825–52.
- Aizenman, J.: 2004, 'Endogenous Pricing to Market and Financing Costs'. Journal of Monetary Economics 51(4), 691–712.
- Aker, J. C.: 2010, 'Information from Markets Near and Far: Mobile Phones and Agricultural Markets in Niger'. *American Economic Journal: Applied Economics* **2**(3), 46–59.
- Albornoz, F., H. Calvo-Pardo, G. Corcos, and E. Ornelas: 2010, 'Sequential Exporting'. CEP Discussion Paper, Centre for Economic Performance, LSE.
- Allen, T.: 2012, 'Information Frictions in Trade'. Working Paper.
- Almeida, H., M. Campello, B. Laranjeira, and S. Weisbenner: 2012, 'Corporate Debt Maturity and the Real Effects of the 2007 Credit Crisis'. *Critical Finance Review* 1(1), 3–58.

- Amiti, M. and D. E. Weinstein: 2013, 'How Much do Bank Shocks Affect Investment? Evidence from Matched Bank-Firm Loan Data'. NBER Working Papers 18890, National Bureau of Economic Research, Inc.
- Anderson, J. E. and E. van Wincoop: 2004, 'Trade Costs'. *Journal of Economic Literature* **42**(3), 691–751.
- Angrist, J. and J. Pischke: 2009, *Mostly Harmless Econometrics: An Empiricist's Companion*. Princeton University Press.
- Atkeson, A. and A. T. Burstein: 2010, 'Innovation, Firm Dynamics, and International Trade'. *Journal of Political Economy* **118**(3), 433–484.
- Barlevy, G.: 2007, 'On the Cyclicality of Research and Development'. American Economic Review 97(4), 1131–1164.
- Baum, C., M. E. Schaffer, and S. Stillman: 2010, 'ivreg2: Stata Module for Extended Instrumental Variables/2SLS, GMM and AC/HAC, LIML and k-class Regression'. Technical report.
- Bentolila, S., M. Jansen, G. Jiménez, and S. Ruano: 2013, 'When Credit Dries Up: Job Losses in the Great Recession'. IZA Discussion Papers 7807, Institute for the Study of Labor (IZA).
- Berger, H. and V. Nitsch: 2008, 'Zooming Out: The Trade Effect of the Euro in Historical Perspective'. *Journal of International Money and Finance* **27**(8), 1244 1260.
- Bergin, P. R. and C.-Y. Lin: 2012, 'The Dynamic Effects of a Currency Union on Trade'. *Journal of International Economics* 87(2), 191 – 204.
- Bernanke, B. and M. Gertler: 1989, 'Agency Costs, Net Worth, and Business Fluctuations'. American Economic Review 79(1), 14–31.
- Bernanke, B., M. Gertler, and S. Gilchrist: 1996, 'The Financial Accelerator and the Flight to Quality'. *The Review of Economics and Statistics* **78**(1), 1–15.
- Bernanke, B. S.: 1983, 'Irreversibility, Uncertainty, and Cyclical Investment'. *The Quarterly Journal of Economics* **98**(1), 85–106.
- Bernanke, B. S., M. Gertler, and S. Gilchrist: 1999, 'The financial accelerator in a quantitative business cycle framework'. In: J. B. Taylor and M. Woodford (eds.): *Handbook of Macroeconomics*, Vol. 1 of *Handbook of Macroeconomics*. Elsevier, Chapt. 21, pp. 1341–1393.
- Bernard, A. B., S. J. Redding, and P. K. Schott: 2007, 'Comparative Advantage and Heterogeneous Firms'. *Review of Economic Studies* **74**(1), 31–66.
- Bernstein, J. I. and T. P. Mamuneas: 2006, 'R&D depreciation, stocks, user costs and productivity growth for US R&D intensive industries'. *Structural Change and Economic Dynamics* 17(1), 70–98.

- Bertrand, M., E. Duflo, and S. Mullainathan: 2004, 'How Much Should We Trust Differences-in-Differences Estimates?'. *The Quarterly Journal of Economics* **119**(1), 249–275.
- Beveren, I. V.: 2012, 'Total Factor Productivity Estimation: A Practical Review'. Journal of Economic Surveys 26(1), 98–128.
- Bloom, N.: 2009, 'The Impact of Uncertainty Shocks'. Econometrica 77(3), 623-685.
- Bloom, N., M. Draca, and J. V. Reenen: 2011, 'Trade Induced Technical Change? The Impact of Chinese Imports on Innovation, IT and Productivity'. NBER Working Papers 16717, National Bureau of Economic Research, Inc.
- Bloom, N., P. M. Romer, S. J. Terry, and J. V. Reenen: 2013, 'A Trapped-Factors Model of Innovation'. *American Economic Review* 103(3), 208–13.
- Blundell, R., R. Griffith, and J. Van Reenen: 1995, 'Dynamic Count Data Models of Technological Innovation'. *Economic Journal* **105**(429), 333–44.
- Blundell, R. and J. L. Powell: 2007, 'Censored Regression Quantiles with Endogenous Regressors'. *Journal of Econometrics* **141**(1), 65–83.
- Boyle, J. E.: 1934, *Cotton and the New Orleans Cotton Exchange. A Century of Commercial Evolution.* The Country Life Press.
- Brouwer, J., R. Paap, and J.-M. Viaene: 2008, 'The Trade and FDI Effects of EMU Enlargement'. *Journal of International Money and Finance* 27(2), 188 – 208. Euro Area Expansion: Current State and Future Prospects.
- Brown, J.: 1909, A Hundred Years of Merchant Banking. New York: Privately Printed.
- Brown, J. R. and A. Goolsbee: 2002, 'Does the Internet Make Markets More Competitive? Evidence from the Life Insurance Industry'. *Journal of Political Economy* **110**(3), 481– 507.
- Bruchey, S. W.: 1967, Cotton and the Growth of the American Economy, 1790-1860; Sources and Readings. New York, NY, USA: Harcourt, Brace & World.
- Buchinsky, M.: 1994, 'Changes in the U.S. Wage Structure 1963-1987: Application of Quantile Regression'. *Econometrica* **62**(2), 405–58.
- Bureau of Economic Analysis: 1994, 'A Satellite Account for Research and Development'. Survey of current business.
- Bustos, P.: 2011, 'Trade Liberalization, Exports, and Technology Upgrading: Evidence on the Impact of MERCOSUR on Argentinian Firms'. *American Economic Review* 101(1), 304–40.
- Caliendo, L. and F. Parro: 2012, 'Estimates of the Trade and Welfare Effects of NAFTA'. NBER Working Papers 18508, National Bureau of Economic Research, Inc.

- Campello, M., J. R. Graham, and C. R. Harvey: 2010, 'The real effects of financial constraints: Evidence from a financial crisis'. *Journal of Financial Economics* 97(3), 470–487.
- Carpenter, R. E., S. M. Fazzari, and B. C. Petersen: 1994, 'Inventory Investment, Internal-Finance Fluctuation, and the Business Cycle'. *Brookings Papers on Economic Activity* **25**(2), 75–138.
- Chapman, S.: 1984, *The Rise of Merchant Banking*, Economic history series. Taylor & Francis.
- Chava, S. and A. Purnanandam: 2011, 'The effect of banking crisis on bank-dependent borrowers'. *Journal of Financial Economics* **99**(1), 116–135.
- Clark, G. and R. C. Feenstra: 2003, *Technology in the Great Divergence*, pp. 277–322. University of Chicago Press.
- Clarke, A. C.: 1992, *How the World Was One: Towards the Tele-Family of Man.* New York, NY, USA: Bantam Books, Inc.
- Cleary, S.: 1999, 'The Relationship between Firm Investment and Financial Status'. *Journal of Finance* **54**(2), 673–692.
- Coleman, A.: 2009, 'Storage, Slow Transport, and the Law of One Price: Theory with Evidence from Nineteenth-Century U.S. Corn Markets'. *The Review of Economics and Statistics* **91**(2), 332–350.
- Collard-Wexler, A.: 2013, 'Demand Fluctuations in the Ready-Mix Concrete Industry'. *Econometrica* **81**(3), 1003–1037.
- Corrado, C., C. Hulten, and D. Sichel: 2009, 'Intangible Capital And U.S. Economic Growth'. *Review of Income and Wealth* **55**(3), 661–685.
- Cutler, D. M., J. M. Poterba, and L. H. Summers: 1989, 'What Moves Stock Prices?'. NBER Working Papers 2538, National Bureau of Economic Research, Inc.
- De Loecker, J.: 2007, 'Do Exports Generate Higher Productivity? Evidence from Slovenia'. *Journal of International Economics* **73**(1), 69–98.
- Deaton, A. and G. Laroque: 1996, 'Competitive Storage and Commodity Price Dynamics'. *Journal of Political Economy* **104**(5), 896–923.
- Desai, M. A., C. F. Foley, and J. R. Hines: 2004, 'A Multinational Perspective on Capital Structure Choice and Internal Capital Markets'. *Journal of Finance* **59**(6), 2451–2487.
- DiNardo, J., N. M. Fortin, and T. Lemieux: 1996, 'Labor Market Institutions and the Distribution of Wages, 1973-1992: A Semiparametric Approach'. *Econometrica* 64(5), 1001–44.

- Duchin, R., O. Ozbas, and B. A. Sensoy: 2010, 'Costly external finance, corporate investment, and the subprime mortgage credit crisis'. *Journal of Financial Economics* 97(3), 418–435.
- Dybvig, P. H. and S. A. Ross: 1987, 'Arbitrage'. In: J. Eatwell, M. Milgate, and P. Newman (eds.): *The New Palgrave: A Dictionary of Economics*. Basingstoke: Palgrave Macmillan.
- Eisfeldt, A. L. and A. A. Rampini: 2007, 'New or used? Investment with credit constraints'. *Journal of Monetary Economics* 54(8), 2656–2681.
- Ejrnaes, M. and K. G. Persson: 2010, 'The Gains from Improved Market Efficiency: Trade Before and After the Transatlantic Telegraph'. *European Review of Economic History* **14**(03), 361–381.
- Ellison, T.: 1886, The Cotton Trade of Great Britain. Effingham Wilson.
- Evans, C. L. and J. Harrigan: 2005, 'Distance, Time, and Specialization: Lean Retailing in General Equilibrium'. *American Economic Review* **95**(1), 292–313.
- Fazzari, S. M., R. G. Hubbard, and B. C. Petersen: 1988, 'Financing Constraints and Corporate Investment'. *Brookings Papers on Economic Activity* 19(1), 141–206.
- Field, A. J.: 1998, 'The Telegraphic Transmission of Financial Asset Prices and Orders to Trade: Implications for Economic Growth, Trading Volume, and Securities Market Regulation'. *Research in Economic History* 18, 145–184.
- Freund, C. L. and D. Weinhold: 2004, 'The Effect of the Internet on International Trade'. *Journal of International Economics* **62**(1), 171–189.
- Froot, K. A., M. Kim, and K. Rogoff: 1995, 'The Law of One Price Over 700 Years'. NBER Working Papers 5132, National Bureau of Economic Research, Inc.
- Froot, K. A. and K. Rogoff: 1995, 'Perspectives on PPP and Long-Run Real Exchange Rates'. In: G. M. Grossman and K. Rogoff (eds.): *Handbook of International Economics*, Vol. 3 of *Handbook of International Economics*. Elsevier, Chapt. 32, pp. 1647–1688.
- Gan, J.: 2007, 'The Real Effects of Asset Market Bubbles: Loan- and Firm-Level Evidence of a Lending Channel'. *Review of Financial Studies* **20**(6), 1941–1973.
- Garbade, K. D. and W. L. Silber: 1978, 'Technology, Communication and the Performance of Financial Markets: 1840-1975'. *Journal of Finance* **33**(3), 819–832.
- Gelübcke, J. P. W.: 2012, 'Foreign and Domestic Takeovers in Germany: Cherry-picking and Lemon-grabbing'. Working Paper Series in Economics 240, University of Lüneburg, Institute of Economics.
- Gibbs, C. R. V.: 1957, Passenger Liners of the Western Ocean: A Record of Atlantic Steam and Motor Passenger Vessels from 1838 to the Present Day. Staples Press.

- Goyal, A.: 2010, 'Information, Direct Access to Farmers, and Rural Market Performance in Central India'. *American Economic Journal: Applied Economics* **2**(3), 22–45.
- Greene, W. H.: 2003, *Econometric Analysis*. Upper Saddle River, NJ: Prentice Hall, 5th edition.
- Hammond, M. B.: 1897, *The Cotton Industry: An Essay in American Economic History. Part 1, The Cotton Culture and the Cotton Trade,* Publications of the American economic association. American Economic Association.
- Harrigan, J.: 2010, 'Airplanes and Comparative Advantage'. *Journal of International Economics* **82**(2), 181–194.
- Head, K. and T. Mayer: 2013, 'What Separates Us? Sources of Resistance to Globalization'. CEPR Discussion Papers 9614.
- Head, K. and J. Ries: 1998, 'Immigration and Trade Creation: Econometric Evidence from Canada'. *Canadian Journal of Economics* **31**(1), 47–62.
- Horn, H., H. Lang, and S. Lundgren: 1995, 'Managerial Effort Incentives, X-inefficiency and International Trade'. *European Economic Review* **39**(1), 117–138.
- Hubbard, R. G., A. K. Kashyap, and T. M. Whited: 1995, 'Internal Finance and Firm Investment'. *Journal of Money, Credit and Banking* **27**(3), 683–701.
- Hummels, D. and G. Schaur: 2012, 'Time as a Trade Barrier'. NBER Working Papers 17758, National Bureau of Economic Research, Inc.
- Hummels, D. L. and G. Schaur: 2010, 'Hedging Price Volatility Using Fast Transport'. *Journal of International Economics* 82(1), 15–25.
- Irwin, D. A.: 2003, 'The Optimal Tax on Antebellum US Cotton Exports'. *Journal of International Economics* **60**(2), 275–291.
- Ivashina, V. and D. Scharfstein: 2010, 'Bank lending during the financial crisis of 2008'. *Journal of Financial Economics* 97(3), 319–338.
- Iyer, R. and J. L. Peydró: 2011, 'Interbank Contagion at Work: Evidence from a Natural Experiment'. *Review of Financial Studies* **24**(4), 1337–1377.
- Iyer, R., J. L. Peydró, S. da Rocha-Lopes, and A. Schoar: 2014, 'Interbank Liquidity Crunch and the Firm Credit Crunch: Evidence from the 2007–2009 Crisis'. *Review of Financial Studies* 27(1), 347–372.
- Jensen, R.: 2007, 'The Digital Provide: Information (Technology), Market Performance, and Welfare in the South Indian Fisheries Sector'. *The Quarterly Journal of Economics* 122(3), 879–924.

- Jiménez, G., A. R. Mian, J. L. Peydró, and J. Saurina: 2010, 'Local Versus Aggregate Lending Channels: The Effects Of Securitization On Corporate Credit Supply In Spain'. NBER Working Papers 16595, National Bureau of Economic Research, Inc.
- Jiménez, G., S. Ongena, J.-L. Peydró, and J. Saurina: 2012, 'Credit Supply and Monetary Policy: Identifying the Bank Balance-Sheet Channel with Loan Applications'. *American Economic Review* 102(5), 2301–26.
- Kalemli-Ozcan, S., H. Kamil, and C. Villegas-Sanchez: 2010, 'What Hinders Investment in the Aftermath of Financial Crises: Insolvent Firms or Illiquid Banks?'. NBER Working Papers 16528, National Bureau of Economic Research, Inc.
- Kaplan, S. N. and L. Zingales: 1997, 'Do Investment-Cash Flow Sensitivities Provide Useful Measures of Financing Constraints'. *The Quarterly Journal of Economics* **112**(1), 169–215.
- Khwaja, A. I. and A. Mian: 2008, 'Tracing the Impact of Bank Liquidity Shocks: Evidence from an Emerging Market'. *American Economic Review* **98**(4), 1413–42.
- Kiyotaki, N. and J. Moore: 1997, 'Credit Cycles'. *Journal of Political Economy* **105**(2), 211–48.
- Kleibergen, F. and R. Paap: 2006, 'Generalized Reduced Rank Tests using the Singular Value Decomposition'. *Journal of Econometrics* **133**(1), 97–126.
- Klein, M. W., J. Peek, and E. S. Rosengren: 2002, 'Troubled Banks, Impaired Foreign Direct Investment: The Role of Relative Access to Credit'. *American Economic Review* 92(3), 664–682.
- Koenker, R.: 2004, 'Quantile Regression for Longitudinal Data'. *Journal of Multivariate Analysis* **91**(1), 74–89.
- Koudijs, P.: 2013, 'The Boats That Did Not Sail: News and Asset Price Volatiliy in a Natural Experiment'. Working Paper.
- Kuppuswamy, V. and B. Villalonga: 2012, 'Does Diversification Create Value in the Presence Of External Financing Constraints? Evidence from the 2008-2009 Financial Crisis'. Working paper.
- Lamont, O.: 1997, 'Cash Flow and Investment: Evidence from Internal Capital Markets'. *Journal of Finance* **52**(1), 83–109.
- Landes, E. M. and A. M. Rosenfield: 1994, 'The Durability of Advertising Revisited'. *Journal of Industrial Economics* **42**(3), 263–76.
- Lew, B. and B. Cater: 2006, 'The Telegraph, Co-ordination of Tramp Shipping, and Growth in World Trade, 1870-1910'. *European Review of Economic History* **10**, 147–173.

- Lileeva, A. and D. Trefler: 2010, 'Improved Access to Foreign Markets Raises Plant-Level Productivity... for Some Plants'. *The Quarterly Journal of Economics* 125(3), 1051–1099.
- Loecker, J. D.: 2011, 'Product Differentiation, Multiproduct Firms, and Estimating the Impact of Trade Liberalization on Productivity'. *Econometrica* **79**(5), 1407–1451.
- McAfee, A. and E. Brynjolfsson: 2012, 'Big Data: The Management Revolution'. *Harvard Business Review* **90**(10), 60–68.
- Melitz, M. J.: 2003, 'The Impact of Trade on Intra-Industry Reallocations and Aggregate Industry Productivity'. *Econometrica* **71**(6), 1695–1725.
- Melitz, M. J. and J. A. Costantini: 2007, *The Dynamics of Firm-Level Adjustment to Trade Liberalization*. Harvard University Press.
- Melitz, M. J. and G. I. P. Ottaviano: 2008, 'Market Size, Trade, and Productivity'. *Review* of *Economic Studies* **75**(1), 295–316.
- Milne, G. J.: 2000, *Trade and Traders in Mid-Victorian Liverpool. Mercantile Business and the Making of a World Port.* Liverpool University Press.
- Nadiri, M. I. and I. R. Prucha: 1996, 'Estimation of the Depreciation Rate of Physical and R&D Capital in the U.S. Total Manufacturing Sector'. *Economic Inquiry* **34**(1), 43–56.
- Nichols, A.: 2007, 'Causal inference with observational data'. Stata Journal 7(4), 507-541.
- Nichols, A.: 2008, 'Erratum and discussion of propensity-score reweighting'. *Stata Journal* **8**(4), 532–539.
- Pakes, A. and M. Schankerman: 1984, 'The Rate of Obsolescence of Patents, Research Gestation Lags, and the Private Rate of Return to Research Resources'. In: *R & D, Patents, and Productivity*, NBER Chapters. National Bureau of Economic Research, Inc, pp. 73–88.
- Paravisini, D.: 2008, 'Local Bank Financial Constraints and Firm Access to External Finance'. *Journal of Finance* **63**(5), 2161–2193.
- Paravisini, D., V. Rappoport, P. Schnabl, and D. Wolfenzon: 2011, 'Dissecting the Effect of Credit Supply on Trade: Evidence from Matched Credit-Export Data'. NBER Working Papers 16975, National Bureau of Economic Research, Inc.
- Pavcnik, N.: 2002, 'Trade Liberalization, Exit, and Productivity Improvements: Evidence from Chilean Plants'. *Review of Economic Studies* **69**(1), 245–276.
- Pippenger, J. and L. Phillips: 2008, 'Some Pitfalls in Testing the Law of One Price in Commodity Markets'. *Journal of International Money and Finance* **27**(6), 915–925.

- Portes, R. and H. Rey: 2005, 'The Determinants of Cross-Border Equity Flows'. *Journal* of *International Economics* **65**(2), 269–296.
- Powell, J. L.: 1986, 'Censored Regression Quantiles'. *Journal of Econometrics* **32**(1), 143–155.
- Rauch, J. E.: 1999, 'Networks Versus Markets in International Trade'. Journal of International Economics 48(1), 7–35.
- Rauch, J. E.: 2001, 'Business and Social Networks in International Trade'. *Journal of Economic Literature* 39(4), 1177–1203.
- Rauch, J. E. and V. Trindade: 2002, 'Ethnic Chinese Networks In International Trade'. *The Review of Economics and Statistics* **84**(1), 116–130.
- Santos, J. A. C.: 2011, 'Bank Corporate Loan Pricing Following the Subprime Crisis'. *Review of Financial Studies* **24**(6), 1916–1943.
- Schankerman, M. and A. Pakes: 1986, 'Estimates of the Value of Patent Rights in European Countries during the Post-1950 Period'. *Economic Journal* **96**(384), 1052–76.
- Schnabl, P.: 2012, 'The International Transmission of Bank Liquidity Shocks: Evidence from an Emerging Market'. *Journal of Finance* **67**(3), 897–932.
- Stigler, G. J.: 1961, 'The Economics of Information'. Journal of Political Economy 69, 213.
- Surdam, D. G.: 1998, 'King Cotton: Monarch or Pretender? The State of the Market for Raw Cotton on the Eve of the American Civil War'. *The Economic History Review* 51(1), 113–132.
- Tamai, T. and Y. Torimitsu: 1992, 'Software Lifetime and its Evolution Process over Generations'. In: *Proceedings of 1992 Conference on Software Maintenance*. pp. 63–69.
- The Economist: 1866, 'Accounts Relating to Trade And Navigation in the Eleven Months Ended November 30, 1866'.
- Townsend, R. M.: 1977, 'The Eventual Failure of Price Fixing Schemes'. *Journal of Economic Theory* 14(1), 190–199.
- Trefler, D.: 1995, 'The Case of the Missing Trade and Other Mysteries'. *American Economic Review* **85**(5), 1029–1046.
- Trefler, D.: 2004, 'The Long and Short of the Canada-U. S. Free Trade Agreement'. *American Economic Review* **94**(4), 870–895.
- Tybout, J. R.: 2004, *Plant- and Firm-Level Evidence on "New" Trade Theories*. Wiley-Blackwell.
- Waugh, M., C. Tonetti, and J. Perla: 2014, 'Equilibrium Technology Diffusion, Trade, and Growth'. Working paper.

- Whited, T. M.: 1992, 'Debt, Liquidity Constraints, and Corporate Investment: Evidence from Panel Data'. *Journal of Finance* **47**(4), 1425–60.
- Williams, J. C. and B. D. Wright: 1991, *Storage and Commodity Markets*. Cambridge University Press.
- Woodman, H. D.: 1968, *King Cotton and His Retainers. Financing and Marketing the Cotton Crop of the South, 1800-1925.* Lexington.