The London School of Economics and Political Science

Essays in Applied Microeconomics

Fabio Pinna

June 30, 2014

Thesis submitted to the Department of Economics of the London School of Economics for the degree in Doctor of Philosophy. London, May 2014

To my family.

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 acknowledgment 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 47,741 words.

Statement of conjoint work

I confirm that Chapter 2 was jointly co-authored with Stephan Seiler. My co-author Stephan Seiler and I collaborated closely in determining paper's objectives. My contribution to the implementation was concentrated especially on the estimation. In sum, I believe that I contributed with 50% of this work.

I confirm that Chapter 4 was jointly co-authored with Marco Gallo, Giacinto Micucci and Francesco Palazzo at Bank of Italy. I contribute together with Francesco Palazzo with the main idea. The general directions of the paper were taken jointly with Francesco Palazzo. I also contributed to the identification strategy, the estimation and the writing of the paper. In total, I believe I contribute with 40% of this work. The views expressed in this work are those of the authors and do not necessarily reflect the opinions of the Bank of Italy.

Acknowledgments

I would like to thank my advisors Doctor Guy Michaels and Professor John Van Reenen for the support throughout the PhD.

Financial support during my PhD from the Government of Sardinia (program Master and Back), LSE, LP Group Asset Management, Massimo Gentile, Prince Salman Bin Sultan Al Saud and Deutsche Bank is gratefully acknowledged.

I would like to thank my co-authors Stephan Seiler and Francesco Palazzo for proving an essential contribution for this thesis. I also benefited from discussions with Francisco Costa, Daniel Silva Junior, Mathew Gentry, Joao Paulo Pessoa, Pasquale Schiraldi, Mark Schankerman, Ronny Razin, Louis Garicano, John Sutton, Alessandro Gavazza, Daniel Osorio, Erika Deserranno, Erik Eyster, Steve Pischke, Oriana Bandiera, Martin Pesendorfer, Alan Manning, Nitika Bagaria, Oliver Pardo, Min Zhang, Qi Zang, Michael Best, Dimitri Szerman, Patrick Andreoli and in general with the discussions had at seminars in the Labour Economics group, Industrial Economics group and the STICERD research centre at the LSE.

I am indebted to Professor Michele Polo, Professor Pierpaolo Battigalli and Professor Donato Michele Cifarelli for having supported my application for the PhD. I am also indebted to Gerardo De Luca, Annamaria Chessa and Antonello Cossu to have supported and believed in me during high school studies.

I am indebted to a number of friends that have made my PhD a more pleasant experience: Maristella Cotelessa, Matteo Manzoni, Francesco Garbuglia, Elisa Russo, Viviana Mucci, Valentina Giuliani, Stephanie Denamps Gomez, Emilio Grisolia, Matteo Malesani, Raffaella Tenconi, Elisa Ferrari, Zozan Kaya, Ana McDowal, Francesca Polo, Cecilia Nardini, Edward Jefferies, Emanuele Degortes, James O'Neil, Lis Meyers, Elena Della Rosa, Can Celiktemur, Nicola Peat, Tal Shapsa, the friends at Make Sense and at Paragon Gym, as well as all the friends I have had the chance to spent good times with in London.

Finally, I would like to thank my family, specially my father Sebastiano, my mother Mariadomenica, my brother Antonello, my sister Cristina and my nieces Alice, Adele Jolanda and Ginevra, for all the support given to me during all these years.

Contents

1.	. Introduction			
2.	Consumer Search: Evidence from Path-Tracking Data			
	2.1.	Introduction	4	
	2.2.	Data	7	
		2.2.1. Purchase data	7	
		2.2.2. Path data	7	
	2.3. Descriptive Statistics		9	
		2.3.1. Search Time	9	
		2.3.2. Speed (and its relationship to search-time)	10	
		2.3.3. Price Dispersion and Possible Savings from Search	12	
	2.4.	A simple model of sequential search	13	
	2.5.	Identification Strategy	15	
		2.5.1. Category-level Price Variation over Time	16	
		2.5.2. Measurement Error	18	
		2.5.3. Chance and Search Spell Duration	19	
	2.6.	Main Results	20	
	2.7.	Robustness checks	21	
		2.7.1. Search over other product attributes	22	
		2.7.2. Measurement Error and Alternative Instruments	25	
		2.7.3. Price Distribution and Expectations	26	
		2.7.4. Differences in search-efficiency	29	
	2.8.	Effect Magnitude and Cross-Category Heterogeneity	30	
		2.8.1. Interpreting the Effect Magnitude	31	

		2.8.2.	Cross-Category Heterogeneity in Search	32
	2.9.	Conclu	usions	35
	Tabl	es		40
3.	Bon	us sche	emes under price-delegation: Evidence from furniture sales	55
	3.1.	Introd	uction \ldots	56
	3.2.	Contex	xt and theory	59
		3.2.1.	Context	59
		3.2.2.	Theory	61
		3.2.3.	Numerical analysis and model predictions	65
	3.3.	Data a	and descriptive statistics	70
	3.4.	Empir	ical model and estimates	74
	3.5.	Robus	tness checks	76
		3.5.1.	Individual characteristics	77
		3.5.2.	Target on revenues and salary	77
		3.5.3.	Heterogeneity across pricing groups	78
		3.5.4.	Outliers	79
		3.5.5.	Within-year dynamics	79
	3.6.	Conclu	isions	79
	Tabl	es		84
4.	Crea	lit Cha	nnel and Non-Performing Loans: Evidence from the Italian Sovereign	
	Deb	t Crisis		97
	4.1.	Introd	uction	99
	4.2.	Literat	ture Review	.03
		4.2.1.	NPLs and firm-level characteristics	.03
		4.2.2.	NPLs and bank characteristics	.04
		4.2.3.	NPLs and macro shocks 1	.06
		4.2.4.	Panel VAR models	.07
	4.3.	Data a	and sample selection	.07
		4.3.1.	The database	.07
		4.3.2.	Sample selection	.09

	4.4. Descriptive statistics	111
	4.4.1. Distribution of non-performing loans (NPLs)	112
	4.4.2. Banks' financial distress and firms' exposure	113
	4.4.3. Product market demand	114
	4.4.4. Firm-level characteristics	114
	4.5. Identification strategy	115
	4.6. Econometric model and results	119
	4.7. Robustness checks	121
	4.8. Conclusions	123
	Tables	130
E	Conduciona	1/1
э.	Conclusions	141
Ap	ppendices	145
Α.	. Appendix to chapter 2	147
	A.1. Linking Sales and Path Data	147
	A.2. Tables	149
в.	. Appendix to chapter 3	151
	B.1. Compensation schemes	151
	B.2. Theoretical model: numerical example	156
	B.3. Data appendix	164
С.	. Appendix to chapter 4	179
	C.1. Geographic markets	180
	C.2. Sample selection	182
	C.3. Share of risky bonds	183
	C.4. Validity of the instrumental variable	184
	C.5. Robustness checks	188

List of Figures

2.1.	Data Structure	48
2.2.	Search-Time Histogram	49
2.3.	Variation in Search-Time Across and Within Trips	50
2.4.	Speed Histogram	51
2.5.	Variation in Speed Across and Within Trips	52
2.6.	Estimated Local Average Treatment Effect	53
2.7.	Treatment Effects Across Categories with varying Degree of Price Dispersion	54
3.1.	Compensation scheme	84
4.1.	Time series BTP-BUND spread	130
4.2.	Exposure to funding gap	138
B.1.	Compensation Scheme I	151
B.2.	Compensation Scheme II	152
В.З.	Compensation Scheme III	153
B.4.	Compensation Scheme IV	154
B.5.	Compensation Scheme V	155
C.1.	Share of risky bonds	183

List of Tables

Correlation between Search-Time and Speed	40
Descriptive Statistics: Prices	41
The Effect of Category-level Pricing on Search	42
Baseline OLS and IV regressions	43
Robustness Check: Fixed Effect Regressions and Promotional Dummy as De-	
pendent Variable	44
Robustness Check: Lagged and Trip-Level Speed Instruments	45
Cross-Category Heterogeneity: Price Dispersion	46
Cross-Category Heterogeneity: Other Category Characteristics	47
Equilibrium outcomes	85
Frequency of employment contracts prescribing a bonus ex-ante - by year and	
pricing rights	86
Employees' history with the firm	87
Frequency of bonuses and pricing rights	88
Summary statistics of outcomes	89
Summary statistics of <i>annual</i> outcomes	90
Mean and std. dev. by bonus	91
Mean and std. dev. by pricing rights	92
Mean of outcomes by bonus and pricing rights	93
The effect of incentives and price delegation	94
Revenues-based vs mixed compensation	95
Distribution of NPLs by geography and by sector	131
	Contention between beaten finite and speed Descriptive Statistics: Prices The Effect of Category-level Pricing on Search Baseline OLS and IV regressions Robustness Check: Fixed Effect Regressions and Promotional Dummy as Dependent Variable Robustness Check: Lagged and Trip-Level Speed Instruments Cross-Category Heterogeneity: Price Dispersion Cross-Category Heterogeneity: Other Category Characteristics Equilibrium outcomes Frequency of employment contracts prescribing a bonus ex-ante - by year and pricing rights Employees' history with the firm Frequency of bonuses and pricing rights Summary statistics of outcomes Summary statistics of annual outcomes Mean and std. dev. by bonus Mean of outcomes by bonus and pricing rights Mean of outcomes by bonus and pricing rights The effect of incentives and price delegation Revenues-based vs mixed compensation

4.2.	Time-series of NPLs - semi-annual data	132
4.3.	Distribution of NPLs before and after the crises	133
4.4.	Distribution of NPLs before and after the Italian crisis by geography and by	
	sector	134
4.5.	Banks' characteristics	135
4.6.	Change in banks' NPLs and total loans before and after the Italian crisis	
	grouped by sign of funding gap at June 2010 \ldots \ldots \ldots \ldots \ldots \ldots	136
4.7.	Firms' exposure to banks' financial distress	137
4.8.	Firm characteristics	139
4.9.	Main regressions	140
A.1.	The Effect of Trip-Duration and Within-Trip Decile on Search-Time and Speed	149
B.1.	Theoretical model: solution for $d = 0$	156
B.2.	Theoretical model: solution for $d = 5$	157
В.З.	Theoretical model: solution for $d = 10$	158
B.4.	Theoretical model: solution for $d = 15$	159
B.5.	Theoretical model: solution for $d = 20$	160
B.6.	Theoretical model: solution for $d = 30$	161
B.7.	Theoretical model: solution for $d = 40$	162
B.8.	Theoretical model: solution for $d = 50$	163
B.9.	ANCOVA	164
B.10	Percentiles of annual Gross Profits-SC	165
B.11	. The effect of incentives and price delegation controlling for employee FEs	166
B.12	The effect of incentives and price delegation controlling for target on revenues.	167
B.13	. The effect of incentives and price delegation controlling for salary	168
B.14	.Gender distribution by pricing group	169
B.15	.Differences in ex-ante characteristics of employment contracts	170
B.16	Distribution of bonus scheme across pricing groups	171
B.17	.Size of expected amount of bonus by pricing groups	171
B.18	Price-elasticity of demand	172
B.19	Price-elasticity for a static monopolist by pricing group	173
B.20	.Impact on outcomes using robust regression with bi-weighting	174

75
76
77
80
81
82
82
84
85
86
87
89
90

Chapter 1. Introduction

This thesis is composed by three essays and applies econometric methods to analyze different economic research questions using microeconomic data. The first essay (chapter 2) analyzes consumer searching behavior in a grocery context. The second essay (chapter 3) studies the implications of the introduction of a bonus scheme in a principal-agent context using data from furniture sales. The third essay (chapter 4) proposes an empirical strategy to estimate the impact that a worsening in banks' wholesale funding opportunities (such as the Italian sovereign debt crisis of 2011) has on borrowers' ability to repay their loans. Chapter 5 concludes the thesis and provides some directions for future work.

The first essay (chapter 2), written jointly with Stephan Seiler, estimates the effect of time spent searching in a supermarket on consumers' expenditure. The analysis is implemented using a unique data-set obtained from radio frequency identification tags which are attached to supermarket shopping carts. This allows us to record consumers' purchases as well as the time they spent in front of the shelf when contemplating which product to buy, giving us a direct measure of search effort. We estimate the effect of extending search on the price consumers pay within a category while controlling for a host of confounding factors such as category-level price variation over time and measurement error. Our results show that an additional minute spent searching lowers category-level expenditure by \$1.40. Extending search-time by one standard deviation allows consumers to appropriate 8 percent of the possible category-level price savings.

The second essay (chapter 3) uses data on the staff of a furniture firm to show that, when a fixed bonus scheme conditional on revenues was introduced, it increased the revenues generated by all sales employees, but I find no significant heterogeneous effect of the bonus scheme depending on whether the employee is given control over price or not. The essay also shows that giving the sales staff control over price does not significantly increase revenues. The effect of the bonus scheme and of price delegation on gross profits minus paid bonuses, commissions, and wages were similar. These results are robust to a number of checks, and are consistent with a model of moral hazard and price delegation.

The third essay (chapter 4), written jointly with Marco Gallo, Giacinto Micucci and Francesco Palazzo, proposes an empirical strategy to estimate the impact that a worsening in banks' wholesale funding opportunities has on the ability of firms to repay their loans. We exploit the Italian sovereign debt crisis of July 2011 as a significant funding shock to Italian banks. This chapter investigates whether this severe shock to credit supply hampered borrowers' ability to repay their loans. We find that, following the funding shock of July 2011, one standard-deviation increase in our measure of firms' exposure to banks' financial distress increases their probability not to honour their loans by about 0.4% (i.e., the size of the credit channel). Our results also suggest that the aggregate demand channel led to a 2.4% increase in the share of non-performing loans.

Chapter 5 concludes and discusses the limitations of the current work and provides some directions for future research.

Chapter 2. Consumer Search: Evidence from Path-Tracking Data

Abstract

We estimate the effect of consumer search on the price of the purchased product in a physical store environment. The analysis is implemented using a unique data-set obtained from radio frequency identification tags which are attached to supermarket shopping carts. This allows us to record consumers' purchases as well as the time they spent in front of the shelf when contemplating which product to buy, giving us a direct measure of search effort. Controlling for a host of confounding factors such as category-level price variation over time and measurement error, we estimate that an additional minute spent searching lowers price paid by \$1.70. We also find that search intensity varies substantially across product to the potential benefits.

2.1. Introduction

When consumers make a purchase decision they might often not be aware of prices for all products due to informational and cognitive constraints. In many categories a large number of products are available and obtaining relevant information can be a costly process. In a grocery shopping context, consumers can search across stores, time their purchase in order to benefit from temporary price reductions and search across various products within a particular store when standing in front of the shelf. In this chapter we focus on the final part of this decision process: the consumer's search-effort when processing information and comparing products and prices immediately before putting the chosen product into her shopping cart. Specifically, our goal is to estimate the effect of the extent of consumers' search activity within a particular product category on the price they pay.

A key challenge in analyzing consumer search behavior in a physical store environment lies in the fact that it is very hard to observe and record which products the consumer was considering before picking one particular product from the shelf. This differs from studies that use online data such as De Los Santos et al. [2012], Chen and Yao [2012] or Koulayev [2014] where one typically observes the sequence of searches. An alternative in a brick and mortar environment would be to provide consumers with eye-tracking equipment as in Stüttgen et al. [2012]. This provides a great level of detail but has the disadvantage of disrupting the "natural" shopping experience of the consumer. In this chapter we therefore propose an approach to understanding search behavior without such an intervention. To this end we use "path-tracking" data obtained from shopping karts that are equipped with radio-frequency identification (RFID) tags combined with store-level data on purchases and product prices.¹. The data allows us to measure the time a consumer spends in front of a particular category before deciding to purchase a specific product. This gives us a direct measure of the extent of the consumer's search activity.²

The central contribution of this chapter is to demonstrate how the benefits from search (per unit of time) can be estimated using data on the total duration of search as well as the

¹A further source of data on consumer search behavior / considerations is represented by survey information directly levied from consumers. This kind of data is used in Draganska and Klapper [2011] and Honka [2014]

²Apart from RFID other technology such as video capture (see Hui et al. [2013a]) or smart-phone wi-fi signals might also be used to measure search-time in a similar fashion. Speaking with experts in the industry we were told that most systems are not precise enough, too costly or difficult to implement due to privacy concerns. Also RFID is not frequently used in practice because implementation costs are high.

price of the chosen product. To the best of our knowledge this is the first work to gather data on search effort and to estimate search benefits in a physical store environment. We find that an additional minute spent searching lowers expenditure by \$1.70. The magnitude is economically significant: Extending search-time by one standard deviation in each product category lowers total trip-level expenditure by \$1.80 or 7 percent for the average shopping trip. Given that we are analyzing very frequently purchased products, the potential unrealized savings are large and suggest that consumers engage in a limited amount of search activity. Secondly, we leverage the fact that we observe consumers searching across a large set of product categories, a unique feature of our data. Consumers face very different price distributions across categories and the potential search benefits therefore differ substantially. We would expect a rational consumer to extend her search-time in such a way that the marginal benefits from search are equated across categories. However, we find this not to be the case with search intensity being lower relative to the benefits from search in categories with high price dispersion as well as categories with short inter-purchase spells. The number of products as well as the average price-level of the category instead are not correlated with marginal search benefits. Taken together our results suggest that given the limited amount of search activity, marketing tools that influence the search process such as displays and feature advertising can have a strong impact. Moreover, such marketing tools are likely to be more effective in categories which are characterized by low search intensity.

In order to guide our empirical analysis, we rely on the canonical sequential model of consumer search. Using the model we show that consumers' search-time is affected by the total number of products being promoted within a category on a given day. When consumers are faced with a distribution that contains a larger number of low prices they will find a sufficiently low price faster. We test this empirically and find that *within categories* weeks with more promotions are characterized by shorter search-spells. Furthermore, consumers also pay lower prices when more promotions are available, thereby creating an endogeneity concern for the relationship of search-time and price paid. Secondly, our data records the time spent in the vicinity of the product category, which is a noisy measure of actual categorylevel search activity. The presence of this measurement error will lead to attenuation bias in an OLS regression setup. In order to deal with both concerns we instrument search-time with the consumer's walking speed before reaching the product. Speed is highly negatively correlated with search-time. We interpret this correlation as being driven by underlying variation in search costs: Consumer with higher search costs will walk faster and search less at each pickup. Speed is a valid instrument as it is determined before the actual search process and therefore before the consumer can learn about prices. Furthermore measurement error in search-time is unlikely to be correlated with speed which allows us to eliminate the attenuation bias.³

Our work is closely related to a series of papers by Hui, Bradlow and Fader (Hui et al. [2009a], Hui et al. [2009c], and Hui et al. [2009b]) which introduced path-tracking data to the academic literature. Relative to their work, which jointly describes the path as well as purchase decisions of consumers, we make little use of the actual path the consumer takes. Instead, we focus more narrowly on the consumer's search process when standing in front of the shelf containing a particular product category.⁴ In addition to the path-data, we also make use of detailed product-level price and purchase data that we are able to link to the path-tracking data-set. The combination of the two data sources allows us to analyze how consumers spend time in the store (recorded by the path-data) impacts the purchases they make (measured in the sales data). In this way we are able to link the novel information we can get out of the path-tracking data to the literature on consumer search and consideration set formation. To the best of our knowledge when analyzing consideration sets in a physical store context (see for example Roberts and Lattin [1991], Andrews and Srinivasan [1995], Bronnenberg and Vanhonacker [1996], Mehta et al. [2003] and Seiler [2013]), the search process was usually unobserved. In this chapter we instead have a direct measure of the extent of search activity.⁵

The remainder of the chapter is organized as follows. Section 2.2 provides a detailed explanation of the data used in our analysis followed by descriptive statistics in section 2.3. In section 2.4, we provide a theoretical framework to guide our empirical strategy, which is presented in section 2.5. In section 2.6, we present the main results, followed by robustness checks in section 2.7. In section 2.8 we provide some interpretation for the magnitude of the

³Note that the measurement error in search-time might be due to the consumer looking at other nearby categories, leaving her cart behind or spending time doing something unrelated to search. There is a possibility that an event that led slows the consumer down will influence her speed prior to pick-up as well as the search-time recorded at the pick-up. We run a set of robustness check to deal with this particular issue.

⁴Another application of path-data is the analysis of unplanned purchases in Hui et al. [2013b].

⁵A small number of studies on consumer search in a physical store environment such as Cobb and Hoyer [1985] or Dickson and Sawyer [1990] and Hoyer [1984] employed teams of trained investigators, who observed consumers in the store and recorded their search-time manually. This allows them to record search-duration albeit only for a relatively small sample of consumers.

estimated effect and explore heterogeneity in search intensity across categories. Finally we make some concluding remarks.

2.2. Data

We use data from a large store in Northern California that belongs to a major supermarket chain.⁶ The complete dataset comprises three pieces: (1) sales data from the supermarket, (2) a store-map with information on product-locations, (3) data on the path a consumer took through the store for a subset of trips over a period of 26 non-consecutive days.⁷ Importantly, we are able to link the path-data to the corresponding purchase baskets from the sales data with the help of the store map. In Section (A.1) of the appendix we provide details on how the two pieces of data are combined.

2.2.1. Purchase data

We have complete purchase data for all consumers that visited the store during the 26 days for which we also observe the path-data. This part of the data is a standard supermarket scanner data-set similar to the IRI dataset (see Bronnenberg et al. [2008]) for instance. At the consumer-level we observe the full basket of products as well as the price paid for each item. Unfortunately, prices for items that do not come in specific pack-sizes (e.g. fresh fruit, vegetables, meat etc.) are not reported in meaningful units (i.e. per kilogram for instance). We are therefore unable to use those products in our analysis. Apart from these problematic products we are going to use data across about 7,500 unique products belonging to roughly 150 different product-categories which are stocked in the store. Over our sample period we observe a total of about 220,000 shopping baskets. However, the path-data is only available for a subset of those.

2.2.2. Path data

In addition to the sales data we also have data on the path that consumers took when walking through the store. The paths are obtained using RFID tags that are attached

⁶We are not able to disclose the identity of the supermarket. The store has a fairly typical format with a trading area of about 45,000 square-feet and a product range of 30,000 UPCs.

⁷The days in the path data are 8/24/2006 - 8/29/2006 and 9/7/2006 - 9/26/2006.

to consumers' shopping karts and baskets (see Sorensen [2003]). Each RFID tag emits a signal about every 4 seconds that is received by a set of antennas throughout the store. Based on the signal, triangulation from multiple antennas is used to pin-point the precise location of the consumer. The consumer's location is then assigned to a particular point on a grid of so-called "traffic-points" which is overlaid onto the store-map. The points used to assign consumers' locations are four feet apart from each other, allowing for a fairly granular tracking of the consumer. For every path we observe a sequence of consecutive traffic points with a time stamp associated to each point.⁸

Not all shopping karts and baskets in the store are equipped with RFID tags however. We therefore only observe path-data for a subset of about 7 percent of all store visits. This is somewhat limiting as we rarely observe multiple trips for the same consumer despite the fact that we have more of a panel dimension in the purchase data. We will discuss how this affects our analysis later when we present the empirical strategy. Second, even if a shopping basket is matched to the path-data, it is possible that not all purchased items in the basket have a match in the path-data. This can happen for instance if the consumer leaves her cart or basket behind and the item pick-up can therefore not be captured in the data.

The primary variable of interest derived from the path-data is the time a consumer spends stationary at a certain point in the store when picking up a product. An individual item purchase, or more precisely the "pick-up" of the item from the shelf, constitutes the unit of observation in our regressions and we observe a total of around 29,000 pick-ups in the data. Using the store map we match the grid of traffic points to product locations that are within reach of the consumer from a given traffic point.⁹ For a given path and set of products in the basket at the check-out we can then use the store map to determine when the product was picked up by the consumer as well as how long she spent in front of the shelf. In other words, the item pickup is defined as the moment in time the consumer walked past a specific product that we later see in her purchase basket. In order to compute search-time, we measure the time elapsed between (1) the moment the consumer is first located on a traffic point assigned to the product and (2) the point in time when she moves on to a traffic

⁸If a consumer moves further than to an adjacent traffic point between signals, the movement over traffic points in between the signals is interpolated. As the signal is emitted at a high frequency little interpolation is necessary for most trips.

⁹The linkage between traffic- and product-points is provided in the data. Mostly any product location is associated with two or three traffic points. However, at a few special locations such as the end of an aisle more traffic points can be associated with a given product location.

point outside of the assigned area. Figure (2.1) illustrates graphically how search-time is assigned to a product pick-up. This metric gives us a measure of time spent in the vicinity of the product which was ultimately purchased. For convenience of exposition we will refer to this metric as search-time. However, we recognize that it is a noisy measure of actual search activity and the consumer might have been doing other things at the same time. The presence of such measurement error will inform our empirical strategy later.¹⁰

Secondly, we compute the speed at which the consumer moves during various parts of her trip using time-stamps and distances between consecutive traffic points. Speed, although not the primary focus of this chapter will play an important role in our empirical strategy.

2.3. Descriptive Statistics

All of our analysis is going to be conducted *within* product categories. In other words we model how a consumer's search activity within a category affects the particular product she buys from that category. In total we have around 150 categories which are defined as groups of products that are naturally substitutes for each other but not with other products outside of the category. Examples for categories defined in this way are Bacon, Beer and Bird Food. Due to a relatively small amount of observations per category, we pool data across categories for most of our analysis. In all pooled specifications we control for a set of category fixed effects.

2.3.1. Search Time

The main novelty of the work lies in the introduction of a direct measure of in-store searchtime. For each item picked up during a shopping trip we can compute the extent of the consumer's search activity as described in the previous section. Figure (2.2) shows the histogram for our search metric. The variable is roughly log-normally distributed with a mean of 10.3 seconds and a standard deviation of 8.5 seconds.

In order to explore where the variation in search time originates from we take the data from 29,000 item pick-ups and regress time spent searching on a set of category fixed effects. These

¹⁰Furthermore, we only observe the movement and stationarity of the cart, but not the consumer herself. To the extent that karts or baskets are left behind, this might also contribute to measurement error in the duration of stationary periods. This is something we explicitly deal with in the empirical analysis.

would play an important role if specific product locations and/or category characteristics such as the price dispersion or the number of available products were the source of systematic differences in search behavior. With a full set of category fixed effects we find an r-square of 0.064. When adding a set of trip fixed effects into the regression as well the r-square goes up to 0.57. However, this is to a large extent due to the fact that we have many trips with only a small number of matched item pick-ups in the data. When we constrain the sample to trips with over 5 items in the basket, the r-square of a regression with category and trip fixed effects decreases to 0.38. In other words there is substantial variation, about two-thirds, in search behavior for the same consumer within a given shopping trip. The rich within-trip variation is going to be helpful for conducting some important sensitivity checks later on.

Figure (2.3) analyzes further how search-time varies both within and across trips. Specifically, we plot for trips of different duration the average search-time within each decile of the trip. Unsurprisingly, we find that longer trips are characterized by substantially longer periods of search. A change of 15 minutes in total trip-length leads to about a 1 second increase in the average search-time per item pick-up.¹¹ Interestingly, the extent of search activity evolves non-monotonically over the duration of each trip. Regardless of the trip's total duration, we find an inverted u-shape of search-time across deciles. The within-trip differences are of a similar magnitude as the across trip ones. For example for trips of up to 15 minutes, average search-time varies between 6.5 second at the first decile to a maximum of 9.5 seconds at the sixth decile. Generally, most of the variation within shopping trips comes from shorter search spells in the first two deciles as well as the last two deciles. There is less variation in search-time in the middle part of the trip.

2.3.2. Speed (and its relationship to search-time)

The speed at which consumers walk during the shopping trips plays an important role in our identification strategy as it will serve as an instrument for search-time. Figure (2.4) shows a histogram of speed, the unit of observation is a minute-long interval within a shopping trip. Speed is approximately normally distributed with a mean of 2.12 feet per second and a standard deviation of 0.75 feet per second. This corresponds to roughly half the average

¹¹Note that most of the trip duration variation is due to consumers walking a longer distance through the store. Only 6 percent of total time in the store is spent searching on the average trip. The positive relationship between category-level search-times and trip duration is therefore not one that occurs by construction.

walking speed of an adult which is about 4.3 feet per seconds. Speed varies substantially both across and within trips. Analogous to search-time we plot speed across deciles of the trip for shopping trips of different length in Figure (2.5). The graph shows the mirror image of what we saw in the case of search-time: the longer the total trip duration the slower do consumers walk at any given point in the trip. Second, speed is fastest at the very beginning and the end of the trip. The relationship between the length of search spells and speed is quite intuitive and presumably reflects that consumers which are in more of a rush will both walk faster and also spend less time contemplating which product to pick. Our interpretation is that both variables reflect underlying variation in search costs / opportunity cost of time. We explicitly exploit this relationship for our identification strategy.

In order to investigate the relationship between speed and search-time in more detail we regress search-time at each pick-up on the consumer's walking speed over the 60 seconds preceding the pick-up. Specifically, we calculate our speed-instrument as feet per second, i.e. by dividing the distance the consumer walked within the minute leading up to the product pick-up by 60 seconds.¹² This yields a highly significant coefficient of -3.60 with a standard deviation of 0.10 and an F-stat of 1,267 and constitutes the first stage regression of the baseline IV-specification we estimate later. We confirm, as figures (2.3) and (2.5) suggest, that the relationship holds both within and across trips: In column (2) of Table (2.1) we regress pick-up specific search-time on speed over the 60 seconds prior to the pick-up and include a set of trip fixed effects. In column (3) we regress average trip-level search-time on average trip-level speed. For both the across- and the within-regression we find a significant correlation with F-stats of 572 and 618 respectively. In later sensitivity checks we will make use of these relationships and isolate only within (across) trip variation in order to estimate the benefits from search.

Note that we interpret speed variation within the trip as representing changes in consumer search costs over a very short period of time. This contrasts with a notion of search costs as being a consumer-specific trait that varies little over time. We think however that the observed variation of speed and search-time over the course of the shopping trip does indicate that consumers' capacity and willingness to process information varies considerably within the trip. Alternatively search costs could vary for each consumer across different product categories and the within-trip patterns in speed and search-time reflects the sequence in

 $^{^{12}}$ In the case of the first pick-up happening less than 60 seconds into the trip, we use speed between the beginning of the trip and the first pick-up.

which those categories are visited. While the sequence of categories might affect search-time it is less clear why speed would be affected by it. We also directly test whether the pattern are driven by the store layout by regressing speed and search-time on location dummies. Results are reported in Table (A.1) in the appendix and we find that within-trip search and speed patterns are similar even after controlling for product location.¹³ Furthermore, our main analysis controls for category fixed effects which will take out any category-specific variation that might be caused by certain categories being visited earlier or later during the typical trip. Finally, our results are also robust to using speed at the trip-level as an instrument at which point any within-trip dynamics have no influence on the estimated coefficient.

2.3.3. Price Dispersion and Possible Savings from Search

In order to quantify the possible benefits of search, we report the category-specific differences between the highest and the lowest price in the category. Because prices for the same product vary substantially over time, we compute the difference between the minimum and maximum price for each day/category combination. We then compute the average of this variable across days for each category. The first row of Table (2.2) reports the distribution of the min-max price difference across categories. On average there is a price difference of \$4.21, but this varies across the set of about 150 categories. At the 25th percentile the price difference is equal to \$1.78 and it rises to \$5.26 at the 75th percentile. We also report the *percentage* difference of the lowest daily price relative to the highest daily price in the category in the second row of the same Table.

Because there is substantial variation in prices due to promotional activity we also report some descriptive statistics on the time series variation in prices. For the purpose of this exercise we define a promotion as a daily price which lies at least 15 percent below the maximum price of that product over our sample period. Similar to the calculation for the price difference, we compute the share of promoted products for each day/category pair and then take the average across days for each category. The distribution across categories is reported in the third row. On average about 30 percent of UPCs within a category are on promotion. Furthermore, even within our short time window many different products go on promotion. In order to capture this, we compute the percentage of UPCs that went on

¹³The tables reports regressions of speed and search-time on dummies for each decile of the trip as well as dummies for trips of different duration (0 to 15 minutes, 15 to 30 minutes, etc.). When we add location fixed effect to the regression the coefficients on the dummies are very similar.

promotion *at some point* during our sample period for each category. The average across categories is almost 60 percent which is substantially higher than the daily share of promoted products indicating that the identity of the set of promoted products changed frequently.

Taken together, the large within-category price dispersion as well as the substantial degree of promotional activity suggest that there are gains from search. The average category-level saving of \$4.21 might seem relatively small compared to other (non-CPG) product categories, however relative to the amount of total shopping expenditure it is not trivial. Consumers buy on average 8 products on a shopping trip, which would allow for maximum savings of roughly \$34. Furthermore, these gains can be realized by consumers on each shopping trip, i.e. on a very regular basis and are therefore of a large overall magnitude.

2.4. A simple model of sequential search

In this section we outline the predictions of the canonical sequential search model described in McCall [1970] and describe how the model maps onto our specific context and data. We make some minor modifications to the model in order to adapt it to our setting and data. Note that the model is not estimated structurally but we use it to guide the estimation strategy.

In the sequential search model consumers receive draws from a distribution of utilities and optimally decide when to stop searching. In our context, consumers search across products within a category. For simplicity of exposition we outline a pure price search model, i.e. consumers care only about price but not about other product characteristics. We will discuss the implications of relaxing this assumption later on.

Assume a consumer gets gross utility v if she consumes any product within the category. Further, assume the consumer incurs a search cost $c_{product}$ when evaluating an additional product and receives a draw from the price distribution F(p) with support $[\underline{p}, \overline{p}]$ for each search attempt. The optimal stopping rule is a time invariant threshold-rule λ (i.e. the consumer will accept any price below λ) which maximizes the consumer's value function¹⁴

¹⁴We ignore discounting due to the short amount of time that consumers spent searching in a given category in our data. We also assume that v is high enough such that the consumer searches at least one option. We interpret this as the consumer being committed to buying in the category but deciding which specific product to pick from within the category.

$$EV = -c_{product} + \int_{\underline{p}}^{\lambda} (v-p)dF(p) + (1-F(\lambda))EV$$
(2.1)

Alternatively one can interpret the optimal stopping rule as the value of λ which equates the marginal benefit with the marginal cost of searching

$$\int_{\underline{p}}^{\lambda} (\lambda - p) dF(p) = c_{product}$$

One can easily see that the optimal threshold λ is increasing in search costs $c_{product}$. Intuitively, a higher search costs will make the consumer less picky and therefore willing to accept a higher price.

In the standard search model we can think of $c_{product}$ as representing the cost of resolving uncertainty about one more option. In our data however, we are not able to measure the number of options evaluated, instead we only know the extent of search activity measured in real time. In order to adapt the model to our setting we model the search cost of evaluating one more alternative as $c_{product} = TimePerSearch * c_{time}$, the product of time needed to search one option (TimePerSearch) and the opportunity cost of time (c_{time}). TimePerSearch represents the efficiency of the search process. It might (as any of the other model primitives) vary across consumers. For simplicity of exposition we ignore any consumer *i* subscripts. The conversion into real-time leads to a slightly modified optimality condition

$$\int_{\underline{p}}^{\lambda} (\lambda - p) dF(p) = TimePerSearch * c_{time}$$
(2.2)

Using this condition it is easy to show that the expected price paid is equal to

$$E(p) = \frac{1}{F(\lambda)} \int_{\underline{p}}^{\lambda} p dF(p)$$
(2.3)

and the expected number of searches is determined by

$$E(N) = \frac{1}{F(\lambda)}$$

therefore the expected time spent searching is given by

$$E(SearchTime) = TimePerSearch * E(N) = \frac{TimePerSearch}{F(\lambda)}$$
(2.4)

Note that λ is the optimal stopping rule defined by equation (2.2) and is therefore a function of *TimePerSearch*. A larger amount of time needed to make an additional search will increase search costs and therefore increase the stopping threshold λ , making the consumer willing to accept higher prices.

2.5. Identification Strategy

One way to think about our empirical strategy is the following: Our lives would be easy if all the observed differences in search-time were caused by differences in consumers' search costs c_{time} . In this case we would have low search cost consumers choosing a lower stopping threshold which would lead them to search for a longer amount of time and pay a lower price on average. In this scenario we could simply estimate the relationship between searchtime and price paid by OLS. Put differently, the object of interest that we want to estimate is the effect that an increase in search-time caused by a decrease in search costs has on price. As we outline in more detail below, there are likely to be factors other than search cost variation which affect the joint-distribution of search time and price paid. The OLS estimate is therefore unlikely to allow us to recover the desired causal estimate.

In the absence of direct information on search costs any variable that is correlated with search costs can be used as an instrumental variable to shift search-time.¹⁵ As long as the instrument is uncorrelated with other factors that affect the price paid, the IV will allow us to estimate the causal effect of search time on price. Moreover, the IV is useful because it translates the search cost movement into more meaningful units: dollars saved per unit of time spent searching.

In our baseline specification we use the speed at which the consumer is walking before picking up the product as an instrument. As we have shown in Section (2.3.2) speed is highly correlated with search-time both within and across trips. Our interpretation is that

¹⁵No paper that we are aware of has direct data on consumers' search costs. A typical approach is to use a structural model in order to back out search costs under some set of assumptions. We do not take this approach here.

variation in search costs is driving both speed and search-time. For instance on a trip on which the consumer has a higher opportunity cost of time, she will both walk faster and spend less time searching for each product. In other words speed and search costs are correlated because they are both affected by a latent third variable: search costs. However, searchtime is also influenced by other factors such as the prices the consumer samples during the search process. Speed is arguably not affected by these factors and can therefore serve as an instrument.

In the following sections we lay out evidence in support of our exclusion restriction in more detail. In particular, two prominent reasons that would introduce bias into an OLS regression are variation in category-level promotional activity over time as well as measurement error in search-time. We believe that the speed instrument does a good job in overcoming these issues. In robustness checks and extensions we consider further sources of bias such as heterogeneity in preferences over product characteristics other than price and incorrect price expectations.

2.5.1. Category-level Price Variation over Time

Consumers form expectations knowing that prices vary both across products and over time. The latter dimension is particularly important in the grocery shopping context due to the presence of high frequency price movements. As was shown in Section (2.3) price reductions due to promotions are very common in our data. Both dimensions are embodied in the price distribution governing the expectation process F(p). On any given day t there exists a price distribution $F_t(p)$ across products that is (in most cases) not known to the consumer¹⁶, but that will influence the length of the search process as well as the expected price. Formally this situation corresponds to the threshold value of the stopping rule being determined by F(p) whereas the expected number of searches and the expected price are a function of $F_t(p)^{17}$

¹⁶It could be known to the consumer in some circumstances such as information about promotions being available through feature advertising. We will address this issue later.

¹⁷Strictly speaking both E(SearchTime) and E(p) are also still a function of F(p) which determines the optimal stopping threshold λ .

$$E(p) = \frac{1}{F_t(\lambda)} \int_{\underline{p}}^{\lambda} p dF_t(p)$$
$$E(SearchTime) = \frac{TimePerSearch}{F_t(\lambda)}$$

Days with more promotional activity are characterized by a price CDF with more weight in the left part of the distribution. This leads to a lower expected search duration as can be easily seen from the equation above. The impact on E(p) is in principle ambiguous and depends on how the mass of the probability density function moves with respect to the threshold. When more products are promoted this will lead to more prices lying below λ . However, depending on where those prices lie within the *truncated* distribution will determine whether the expected price paid increases or decreases. With our data we are able to directly test whether changes in $F_t(p)$ have any impact on search-time and price. We do this by regressing time spent searching (and price paid) by consumer *i* in category *c* on day *t* on the fraction of products promoted within the category and a set of category fixed effects as well as day fixed effects

$$SearchTime_{ict} = \alpha * FractionPromotedProducts_{ct} + \xi_c + \delta_t + \varepsilon_{ict}$$
(2.5)

$$Price_{ict} = \widetilde{\alpha} * FractionPromotedProducts_{ct} + \widetilde{\xi}_c + \widetilde{\delta}_t + \widetilde{\varepsilon}_{ict}$$
(2.6)

where $\xi_c(\tilde{\xi}_c)$ denotes the category fixed effect and $\delta_t(\tilde{\delta}_t)$ the day fixed effect. The predictions outlined above correspond to a negative coefficient α in the first regression. The prediction for $\tilde{\alpha}$ instead is ambiguous. Note that controlling for category fixed effects is important here as promotional activity and search might vary across categories for a host of other reasons.

Table (2.3) shows that we indeed find a negative and significant coefficient when regressing search-time on our measure of promotional activity, confirming the prediction of the search model. The regression also provides some first evidence that our search metric varies in an intuitively plausible way.¹⁸ When regressing search-time on the share of promoted UPCs we get a negative and significant coefficient. In terms of magnitude, a movement from no

¹⁸Note that we are less concerned with measurement error in search-time in this regression as search is used as the dependent variable. Later search-time will appear as an explanatory variable and measurement error will play a more important role.

promotions in the category to all products being promoted would lower search-time by one second which corresponds to a 10 percent decrease. Although the theoretical predictions regarding the effect on price paid are ambiguous we do find a negative and significant effect as reported in columns (2) and (3). Quantitatively, price paid is about 19 percent lower in full promotion weeks versus no promotion weeks.¹⁹

The analysis above shows that the variation in promotional activity potentially leads to a correlation of search-time and price which is unrelated to the consumer's search cost as a driver of the extent of search activity. An OLS regression will therefore not allow us to estimate the causal effect of search activity on price paid. In order to deal with the endogeneity problem we need an instrument that shifts search-time by affecting search costs but that is unrelated to the extent of category-level promotional activity at any given point in time. Speed does qualify as an instrument as long consumers do not have any price information before arriving at the shelf. If consumers do have information about pricing before they arrive at the shelf, from promotional flyers and/or in-store displays, this might influence their expectation and potentially violates the exclusion restriction. However, even in this case, the IV is only invalid in the case where consumers adjust their walking speed to the price information, by hurrying to the shelf with the promoted product for instance. We don't think that this scenario is very likely but it is hard to rule out entirely. We deal with this issue in greater detail in the robustness check section later in the chapter.

2.5.2. Measurement Error

In our data we are able to measure time spent in the vicinity of the product category, which presumably is a noisy measure of actual category-level search activity. In particular, measurement error in search-time might arise for a variety of reasons: the consumer might

¹⁹Note, that apart from promotional activity having an impact on the *set* of prices being available, it could also affect the probability with which a particular price is drawn. This is an issue specific to our setup as all product prices are visually "accessible" on the shelf immediately. Promotions might therefore provide visual cues that draw the consumer's attention to the promoted product. This would be captured by a shift in the CDF from which prices are drawn which would now assign more probability weight to products which are promoted on the particular day. In a pure price search model most typically the probabilities of drawing a particular price are uniformly distributed across products. In our case there might be reason to think of promotions as shifting the probability of drawing a certain product. This type of effect would also lead to a negative correlation of promotional activity with search-time. Most likely both mechanism contribute to the negative coefficient α estimated in the above regression. For the purpose of the empirical exercise in this chapter we do not attempt (or need) to disentangle the two channels.

be looking at other categories nearby, leave her cart behind or simply spend part of the time engaging in search-unrelated activity. In other words we are not dealing with measurement error that arises simply from imperfections in the data recording process. Instead, searchtime as recorded in the data can be seen as a proxy for actual search effort. As usual, the presence of this measurement error will lead to attenuation bias in an OLS regression setup. Given the nature of our data this issue could potentially be quite severe.

As long as we think of measurement error as arising from local and isolated occurrences such as the consumer leaving the cart behind or contemplating a purchase in another nearby category, using speed as an instrument will alleviate the problem. Potentially the issue is more complicated in our context as things that slow the consumer down during her search process might also affect her speed leading up to the search incidence. This could occur for instance if the consumer slows down because her attention is grabbed by a display in a particular aisle. This might impact her behavior over a longer time window during the trip. In particular she might walk more slowly as well as spend more time near the product due to the distraction created by the nearby display. In other words the same measurement error might also affect the instrumental variable. This kind of issue is most likely to contaminate speed over the time period right before the pick-up. In order to alleviate concerns we therefore use lagged versions of our speed instrument in a robustness check. In particular, we implement a set of regressions using different time-spans over which the lagged measure is computed. Second, we also run a sensitivity check in which we use trip-level average speed as an instrument. In other words we are using only across trip variation in speed in order to estimate the impact of search. We consider this a very conservative approach to dealing with measurement error as events affecting individual pick-ups are unlikely to be correlated with speed over the entire shopping trip.

2.5.3. Chance and Search Spell Duration

There is another issue, specific to our context, which might cause attenuation bias in a similar way as measurement error. A sequential searcher can be more or less lucky in how quickly she comes across a price draw which lies below her stopping threshold. However, the expected price conditional on having already searched a certain number of times remains unchanged. In other words whether the consumer searched only once or 10-times, conditional on not having stopped yet, the expected price is always equal to the unconditional price expectation at the beginning of the search process:

$$E(p|p_1 > \lambda, ...p_k > \lambda) = E(p) = \frac{1}{F(\lambda)} \int_{\underline{p}}^{\lambda} p dF(p)$$

where p denotes the price of the actually purchased product, p_1 to p_k denote the price draws for the k options searched so far (without having stopped). The intuition for this result can be easily obtained from the basic dynamic optimization problem in equation (2.1). As long as prices above the threshold are drawn, the consumer always finds herself back in the same situation with an unchanged value function when making the decision to continue searching.

To fix ideas, assume that there is a set of consumer with identical search costs (in terms of both c_{time} and TimePerSearch) and therefore identical threshold value λ . The actual duration of their respective search spells will in general be different, although the expected duration is the same, and this difference depends entirely on the sequence of price draws they receive. Furthermore consumers with longer spells will not pay different prices on average because the expected price conditional on the number of unsuccessful searches is the same as the unconditional expected price. We therefore have variation in the duration of search spells which is uncorrelated with price.

Remember that we want to find the effect of search-time on price caused by a change in search costs. In other words we want to know how much less a consumer pays who searches more *on average* because she is pickier. We therefore want to get rid of the variation in search duration which is caused by similarly picky consumer being more or less lucky with their price draws. In a similar vein as measurement error the chance-induced variation in search spell duration would lead to an underestimated effect of search-time on price. It seems safe to assume that the speed instrument is not correlated with chance during the search process and the IV should therefore deal with this issue.

2.6. Main Results

In order to analyze the impact of search time on the price paid within a category we run the following regression
$$p_{ijt} = \beta * SearchTime_{ijt} + \zeta_c + \varepsilon_{ijt}$$

$$(2.7)$$

Where p_{ijt} denotes the price consumer *i* pays for product *j* which she purchased on day *t*. ζ_c denotes a category fixed effect, the subscript *c* denotes the category which product *j* belongs to. ε_{ijt} denotes the error term. A full set of category fixed effects is used across all our specifications as we want know whether *within* a given category longer search leads to a consumer picking a lower priced product. We cluster standard errors at the customer-level to allow for an arbitrary within-customer correlation of the error terms. Results are reported in Table (2.4)

We start by running the above regression by OLS. Doing so, we find a negative and significant effect of search time on price. The coefficient is equal to -0.0057, in other words an additional second of search time leads to a half a cent lower price. An additional minute spent searching would therefore lower the price paid by about 30 cents. However, as described in the previous section, the coefficient on search-time might be biased for various reasons. We therefore implement an IV-strategy using the consumer's walking prior to reaching the product as an instrument. We outlined in the previous section why speed should deal with both the endogeneity of search-time as well as measurement error.

As reported before in Section (2.3.2), the first stage regression of search time on speed reported in column (2) is highly significant with an F-stat of 1,267. Column (3) reports the coefficient of the effect of our (instrumented) measure of search time on price. We find a negative and significant effect of -0.0275 which is over 4-times larger than the OLS estimate of -0.0057 showing that the issues described above had a substantial impact on the magnitude of the OLS coefficient. Quantitatively the point estimate of the IV corresponds to about a \$1.7 drop in price for an additional minute searched. We will return to an interpretation of the effect magnitude later, after probing the robustness of our result with a set of sensitivity checks.

2.7. Robustness checks

We use the sequential search model in order to systematically run through a battery of robustness checks. Despite the fact that we do not structurally estimate the search model, it nevertheless provides a natural starting point to guide the sensitivity analysis. In particular we consider how variation in each of the model primitives influences search-time and price paid as well as how it relates to the consumer's walking speed, our instrument. The search model is quite parsimonious, therefore the set of model primitives we have to consider is small and comprises the price distribution F_p and search efficiency (*TimePerSearch*). We further investigate several extensions of the simple model: (1) a model where consumers have preferences over non-price characteristics and therefore search not only for a lower price, (2) deviations from rational expectations which influence the consumer's perceived benefit from searching and (3) the scenario where consumer have information about prices before arriving at the product location in the store. Finally, we also provide a more in-depth discussion of issues related to measurement error in search-time.

2.7.1. Search over other product attributes

One threat to the validity of our estimation lies in the fact that consumers are likely to not only consider price but rather search over a whole set of product characteristics. As products in most categories are quite differentiated and consumers presumably have heterogeneous tastes over product attributes it is natural to ask how this interferes with our analysis. In the search model this would be captured by the product valuation term v becoming consumer-product-specific

$$EV = -c_{product} + \int_{\lambda}^{\overline{u}} (u_{ij}) dG(u) + G(\lambda)EV$$

where $u_{ij} = (v_{ij} - \alpha_i p)$ denotes utility which is a function of both price and brand preferences. α_i denotes the individual-specific price coefficient and v_{ij} represent the consumer specific valuation of product j. G(u) is the cumulative density function that describes the distribution of utilities across products.²⁰ In this framework consumers will find higher utility products as they search longer. A higher utility could be achieved either by a lower price or by finding a product which is preferable along other product dimensions, i.e. that has a higher realization of v_{ij} .

 $^{^{20}}$ Note that the threshold now denotes the *minimum* utility level at which the consumer will stop searching. In the price search model the threshold denoted the *maximum* price at which to stop.

First note that the presence of preferences over other product characteristics does not per se invalidate our analysis. Consider for instance the situation where consumers have preferences over brand characteristics beside price and that the former are randomly distributed across consumers. The higher the weight on non-price characteristics the lower will be the effect of search-time on price, but it does not introduce bias into our analysis. If instead product tastes are not randomly distributed, this could potentially pose a problem, in particular if product preferences are correlated with search costs across consumers. For instance one could imagine that lower income consumers have a stronger preference for lower prices relative to quality and also have lower search costs. These consumers would be searching longer and also pay lower prices due to their preferences. This would lead to an upward bias (in absolute terms) in the effect of search-time on price.²¹

We tackle this issue in two ways. First, we run a robustness check which controls for individual- and trip-specific differences in search and purchase behavior by including a set of trip fixed effects. This sensitivity check leverages the fact that there is substantial variation in both speed and search-time over the course of a consumer's shopping trip (see Section (2.3)). In this way, we are only identifying the effect of search from *within trip* variation across categories. In other words, we identify our main coefficient of interest from consumers paying lower than average prices in categories in which they search more relative to their average search-time across categories on the particular trip. Note that this approach is more conservative than using consumer fixed effects. However, because most consumers do not appear multiple times in the path-data, the two approaches are very similar. The results from this regression are reported in Column (2) of Table (2.5). We replicate our baseline specification without fixed effects is -0.0206 (standard error of 0.0088), which is similar to the results of our baseline specification.²²

Note that the robustness check deals with preference heterogeneity only as long as a consumer's taste for quality relative to price is common across categories. If instead consumers have a strong preference for quality over price only in some categories, but not in others,

²¹This issue is not unique to our setting. We are not aware of a structurally estimated search model that allows for a flexible joint distribution of search costs and price sensitivity which would capture the dynamics described above.

²²Note that the number of observations for this robustness checks varies slightly relative to the baseline IV regression. This is due to the fact that we drop trips for which only one item pickup is recorded when we include trip fixed effects. We re-estimated the baseline model using without the single-item trips (not reported) and find that the change in the sample size does not affect our results.

then trip fixed effects do not fully address the issue. We see no reason why search costs would in general be category-specific, however it could be the case that consumers are more likely to buy certain categories towards the beginning, end or middle of their trip. However, even if preferences were category specific for whatever reason, this would only be an issue if search-costs were also category specific in a way that would create a spurious correlation. I.e. categories in which consumers have stronger preferences over quality would have to be categories for which search costs are higher in order to overestimate the effect. Also, we later show that results are very similar when using a specification which only uses *across* trip variation in search-time. Any within-trip dynamics could not possibly contaminate the results in this alternative specification.

Second, we re-run our main specification, but change the dependent variable: Instead of price paid we use an indicator variable that is equal to one if the consumer picked a product that was on promotion. Note that the number of observations is smaller as we need to observe regular purchases of a particular product in order define when it went on promotion. This is only possibly if the product is purchased relative frequently. As before, the instrument is strongly correlated with search-time with an F-stat of 529. The results slightly differs from our baseline first stage only due to the difference in the number of observations used.²³ In the second stage the magnitude of the coefficient (standard error) on search-time is 0.0041 (0.0016), i.e. an additional minute spent searching increases the likelihood of finding a promotion by 25 percentage points (0.0041 * 60 = 0.246). As in our baseline case, we find a much larger effect when instrumenting search time relative to the OLS case. Finally, we also estimate the effect on the promotion dummy when including a set of trip fixed effects. When doing so we obtain a coefficient (standard error) of 0.0044 (0.0024).

Using only time variation in product-specific prices due to promotions does mitigate some of the concerns raised above. In particular, if quality differences are only reflected in different baseline prices, then this approach deals with concerns arising from preferences over quality relative to prices. The specification using a promotional dummy shows that our effect is not estimated purely from consumers with longer search spells buying products with lower base prices which are presumably of lower quality. Instead it is the case that longer search spells make it more likely for a consumer to buy a promoted product. It could of course be the

²³We replicated the baseline regression using only the observations for which the promotion dummy is defined and find results that are not significantly different from the ones using the full sample. This reassures as those issues of sample selection are unlikely to contaminate the analysis.

case that lower quality products go on promotion more often and that is the reason why we find consumers with longer search spells purchasing on promotion more often. In our data we find however no relationship between price and promotional frequency. To test this we regress the fraction of days a product is promoted on the baseline price and a set of category dummies. The regression is run at the product-level for the set of 5,848 UPCs for which we are able to define the promotion dummy. The coefficient on the baseline price is very small and insignificant with a coefficient (standard error) of 0.0016 (0.0018).

2.7.2. Measurement Error and Alternative Instruments

In our data we are able to measure how much time the consumer spends in the vicinity of a product she purchased before picking it up from the shelf. This is however only a proxy for "true" search-time as the consumer might spend only part of her time in the product's vicinity on search. She might leave the cart or basket behind, some of the recorded time in the category's vicinity might be spent looking at other nearby products or simply not engaging in any search related activity at all. The assumption for our identification strategy to work is that whatever factor affects our search-time proxy is of a very immediate nature and therefore only affects search-time, but not the consumer's shopping behavior prior to arriving at the category. If this is true, then speed leading up the pick-up is correlated with "true" search-time because both reflect variation in the consumer's search costs. On the other hand, speed would in this case not be correlated with time spent on non-search related activity that we capture as part of our proxy variable.

In our context it is conceivable that some part of the measurement error in search-time is correlated with speed leading up to the pick-up. This could occur for instance if the consumer slows down because her attention is grabbed by a large display in a particular aisle. This might impact her behavior over a longer time window during the trip. In particular she might walk more slowly as well as spent more time near the product due to the distraction created by the nearby display. Although we do not think that such a scenario is very likely to occur, we do run a set of robustness checks to further explore the issue. Specifically, we re-run our baseline estimation with a slightly different instrument: we use speed lagged by 10 second, i.e. speed from 70 second up to 10 seconds before the pick-up. In this way we are allowing measurement error to affect speed directly before the pick-up as we are cutting out the part of the consumer's trip that is closest to the actual pick-up. We repeat the same exercise for longer lags of 20 and 30 seconds as well. The results are reported in Columns (2) to (4) of Table (2.6). In comparison to our baseline case, reported in the first column of the same table, the effect is similar. When increasing the lag, the coefficient on search-time increase. This is consistent with the idea that the lags are able to get rid of issues of measurement error. However, none of the coefficients using the lagged instruments is significantly different from our baseline specification. We also lose some precision when employing lagged variables due to the fact that the instrument is slightly weaker the more seconds we exclude. This is unsurprising as speed directly before the pick-up presumably has the highest correlation with pick-up time.

Second, we implement an IV-regression that uses only trip-level variation for identification and constitutes probably the most conservative way to deal with the type of measurement error present in our data. We try to explicitly capture the idea that longer trips with more purchases tend to be trips on which the consumer is less in rush, i.e. has lower search costs. She therefore walks more slowly and search-spells are longer as we have shown in the descriptive statistics earlier. In order to use this variation we simply use average speed at the trip-level as an instrument, rather than speed immediately prior to the pick-up. The type of measurement error that is most likely to occur is relatively localized likely to affect only a small part of the trip. Measurement error from individual instances during the trip is therefore likely to be averaged out at the trip-level. Even if there is measurement error at the trip-level, we see little reason why it would be correlated with average speed over the whole duration of the trip. The results from this regression are reported in the last column of Table (2.6). Our instrument is somewhat weaker as we are not using within-trip variation, but still strong in absolute terms with an F-stat of 580. Our second stage coefficient on search-time is significant and of similar magnitude as our baseline specification: The point estimate is equal to -0.0299 with a standard error of 0.0142.

2.7.3. Price Distribution and Expectations

A model primitive that has a key influence on search behavior is the price distribution F(p). We already discussed endogeneity concerns which arise from the fact that categoryspecific price distributions vary over time due to the fact that different products go on promotion at different points in time. We now turn to other two issues related to the price distribution. Firstly, we consider the effect of consumers having biased expectations about the price distribution. Secondly, we investigate the consequences of consumers having information about daily prices, and in particular promotions, before engaging in search. The latter is likely to arise in our setting due to the presence of feature advertising and in-store displays which provide price information to the consumer before she arrives at the shelf and starts searching.

Incorrect Consumer Expectations

A dimension in which consumers' behavior might differ from the stylized model is in the way they form expectations about prices. As in any search model, expectations play a crucial role because they determine the marginal benefit of searching and therefore the optimal amount of search activity.²⁴ In our search model a deviation from rational expectations can be captured by the fact that the optimal stopping rule would be based on an incorrect price distribution. In other words the optimal price threshold λ would solve

$$\int_{\underline{p}}^{\lambda} (\lambda - p) d\widetilde{F}(p) = c_{product}$$
(2.8)

where $\widetilde{F}(p)$ represent the price distribution used to form expectations. In the case of nonrational expectation $\widetilde{F}(p)$ will be different from the actual price distribution F(p). Note, that when the consumer engages in search, prices are still drawn from the true price distribution F(p), however the stopping threshold might differ from the one of a rational consumer. $\widetilde{F}(p)$ therefore only affects search-time and price via its impact on λ :

$$\begin{split} E(p) &= \frac{1}{F(\lambda(\widetilde{F}(p)))} \int_{\underline{p}}^{\lambda(F(p))} p dF(p) \\ E(SearchTime) &= \frac{TimePerSearch}{F(\lambda(\widetilde{F}(p)))} \end{split}$$

It is easy to see that more pessimistic expectations will lead to shorter search spells as well as a higher expected price paid. The negative correlation between search-time and price that our estimation captures could therefore be in part due to heterogeneity in expectations across

²⁴In virtually all structural models of search, consumers are assumed to know the true price distribution. Indeed, imposing the expectation process is usually necessary for identification in any dynamic model, including models of search.

consumers. However, this is actually "good" variation in the data rather than a confounding factor that might interfere with a causal interpretation of our estimates. To see this, note that expectations only influence search-time and price paid through their influence on the stopping threshold λ . Overly optimistic consumers do overestimate the marginal benefit from search and therefore search longer and pay a lower price on average. Moreover, it is always possible to mimic the behavior of an overly optimistic consumer with a rational consumer that has a lower opportunity cost of time, i.e. one could lower the marginal benefit but increase the marginal cost in order to keep the stopping threshold unchanged. In other words it matters little to our estimation whether optimism or low search costs make the consumer less picky.²⁵

Information obtained before searching

Prices at the daily level are likely to be, at least partially, observed by some set of consumers due to feature advertising and in-store displays. This affects behavior in two ways. Consumers with prior knowledge about daily prices will base their expectations on this information whereas other consumers form expectations based on the distribution of prices over time and across products. This issue is very similar to the case of consumers having biased expectations. As discussed above, any type of variation in expectation formation does not cause any problems in terms of causal inference.

Apart from promotional activity having an impact on the *set* of prices being available and on consumers' expectations, it could also affect the probability with which a particular price is drawn. This is an issue specific to our setup because all product prices are visually "accessible" on the shelf immediately. Promotions might therefore provide visual cues that draw the consumer's attention to the promoted product. This could happen either because the consumer knows about the promotion and specifically tries to find to particular product or because promotional signs on the shelf capture her attention. Formally such an effect would be captured by a shift in the CDF from which prices are drawn which would now assign more probability weight to products which are promoted on the particular day.²⁶ This type

²⁵Note that if there is any such variation in expectations in the data, our IV-strategy (in particular in conjunction with trip fixed effects) will most likely not make use of it. It does seem unlikely that consumers' category-specific price expectations do influence their walking speed leading up to the pick-up within the particular category.

²⁶In a pure price search model most typically the probabilities of drawing a particular price are uniformly distributed across products.

of effect would lead to a negative correlation of promotional activity with search-time similar to the effect of variation in F(p) over time discussed earlier.

Our instrument is valid as long as prior knowledge of prices does not alter the speed at which consumers walk leading up to the item pick-up, by hurrying to the shelf with the promoted product for instance. We don't think that this scenario is very likely. Even if it did occur, our sensitivity check using trip-level speed as an instrument is unlikely to be affected by price knowledge in one particular category. The fact that our results are robust to this particular test indicates that prior price knowledge is unlikely to pose a threat to identification.

2.7.4. Differences in search-efficiency

The final model primitive whose influence on our analysis we need to look at is *TimePerSearch*, the efficiency of the consumer's search process. Most likely there is variation across consumers in how much time they need in order to resolve uncertainty regarding a specific number of options. The first order effect of a decrease in *TimePerSearch* is that it lowers the consumer's search cost and therefore leads to a lower stopping threshold λ . In other words, consumers which search more efficiently are willing to wait for a lower price draw as it is less costly for them to evaluate additional options in the search process. Search efficiency only affects price via this channel. The impact on search-time is however more complicated. On the one hand search-time will be longer due to the fact that a more efficient consumer is pickier, i.e. has a lower λ . At the same time however search-time is lower simply because it takes less time to evaluate an additional option. This is easy to see from equation (2.4), where TimePerSearch enters in the numerator and λ (which is a function of TimePerSearch) in the denominator. The consequences of variation in search efficiency for our estimation are similar in nature to a measurement error problem. Ideally we would like to measure variation in the extent of search activity in terms of the number of options evaluated, but we only observe search effort in real-time. The total search duration can be decomposed into two components: the number of options evaluated and the time it takes to evaluate each option. The former has an impact on price paid, but the latter does not. Variation in search-time due to differences in search efficiency therefore causes attenuation bias in our estimate.

Because search efficiency is a latent concept, it is very hard to assess how much this issue could affect estimation. We are less sure in this case that our speed instrument is able to purge out the problematic variation in search efficiency. It is conceivable that consumers which are less efficient when searching also generally walk at a lower speed. However, most likely TimePerSearch is a consumer-specific trait and does not vary across categories within a trip. Our specification with trip fixed effects should therefore be able to deal with the issue.²⁷

2.8. Effect Magnitude and Cross-Category Heterogeneity

We find returns from searching that are fairly large with roughly \$1.70 per minute. However, because our measure of search-time is distributed with a mean of 10 seconds and a standard deviation of 8 seconds, a minute constitutes a strong linear extrapolation relative to the typical search time. In this section we therefore provide some guidance on how to interpret the magnitude of the effect.

By the nature of the search problem, the relationship between search-time and price paid is necessarily a non-linear one. Figure (2.6) illustrates this relationship when varying consumers' search costs. In terms of our search model, we can trace out how lowering search costs leads to a lower stopping thresholds (λ) which in turn increases expected search-time and decreases expected price (see equations (2.2 to (2.4)). The relationship is non-linear with extensions in search-time from a lower level being associated with larger gains in terms of finding lower prices. Moreover the potential gains within a category are bounded by the lower bound of the price distribution. Given the shape of this relationship, the magnitude of our estimate will depend on whether consumers in our data search relatively little (represented by the red scatter-plot) or a lot (the blue scatter-plot). In the latter case the average consumer realizes more of the potential gains from search and the incremental benefit at the margin is smaller. We can therefore interpret our estimates as the average consumer's marginal benefit from search or the unrealized potential gain from extending search by another second. Due to the local nature of the effect and non-linear shape of the relationship a linear extrapolation is a good approximation only for small changes in search-time.

In order to study the non-linearity of the relationship more directly we also re-estimate our baseline model with an additional quadratic search-time term. When doing so we find a coefficient (standard error) of -0.092 (0.041) for the linear term and 0.00166 (0.00097) for

 $^{^{27}}$ Even if there was variation in TimePerSearch within a trip this would if anything lead to an underestimation of the effect of search-time on price.

the quadratic term. The sign on both terms is as expected and consistent with the shape in Figure (2.6).

2.8.1. Interpreting the Effect Magnitude

Due to the reasoning above we have to be careful not to extrapolate out linearly too far. In this vain we use some back-of-the-envelope calculation in order to compute how large the gains from search can be within a given trip using the non-linear relationship estimated above. Extending search-time by one standard deviation from its mean, i.e. from 10 to 18 seconds, lowers price by 36 cents²⁸. The average consumer purchases 8 products on a typical trip and could therefore save about \$2.90 in total when extending search-time by one standard deviation in each product category. This constitutes roughly 11 percent of the typical total shopping basket size of \$27. Another way to quantify potential savings from search is to put them into the broader context of the total time budget allocated to the shopping trip rather than just the time spent searching. Consumers spend on average 23 minutes in the store and spend only about 80 second, i.e. 6 percent of their trip, searching. Extending search time by one standard deviation (from the mean-level) in each category, i.e. by 64 seconds, corresponds to a 4.5 percent increase in total shopping time and lowers expenditure by \$2.90. Relative to the average trip-level expenditure of \$27 this translates into an elasticity of expenditure with respect to shopping time of -2.3 at the trip-level.

Second, the magnitude of the benefits from search we find are in line with search cost estimates of papers that estimate search costs structurally. Although our approach is different in nature, it is fairly straightforward to compare our estimates to structural estimates of search costs. In the typical empirical search model, search costs are identified as the monetary value that is equal to the marginal benefit from searching another option.²⁹ In our case we directly estimate the marginal benefit from search. Therefore the only missing element to compare estimates is an assumption about how much time searching another option takes. De Los Santos et al. [2012] find search costs of \$1.35 in the internet book market, Honka [2014] estimates a cost of \$80 for acquiring an additional car insurance quote, Koulayev [2014] reports a search cost of around \$6 that a consumer needs to incur to flip to another

 $^{{}^{28}(-0.0921*18+0.0017*18^2)-(-0.0921*10+0.0017*10^2)=-0.3638}$

²⁹Strictly speaking the search cost magnitude has to be equal or higher than the marginal benefit at the point where the consumer stops, but lower at all previously searched options. This identifies search cost bounds. Point identification usually comes from functional form.

page when using an online meta-searcher for hotel bookings.³⁰ If we assume that each of the search activities in those papers takes about a minute, our estimate of \$1.70 saved per minute (or \$2.90 saved per trip) is roughly of a similar magnitude. The one estimate that is considerably higher is Honka [2014], possibly due to the fact that procuring a car insurance quote might take longer. Of course, the markets for which search costs are estimated differ in many ways and one would therefore not expect to necessarily find search costs of exactly the same magnitude.

Note that we cannot easily translate monetary savings into welfare gains as we do not model preferences over product characteristics other than price. We therefore do not know how much price gains weigh in the consumer's utility function relative to brand preferences along other dimensions. This is a shortcoming of this study due to the nature of our data. Pooling data across many categories and products makes it difficult to model preferences over other product characteristics for almost 30,000 UPCs. This also affects the comparison with the structural search cost estimates above. If the monetary gain is only part of the full utility gain then search costs might actually be larger than the purely monetary benefit we estimate.

2.8.2. Cross-Category Heterogeneity in Search

Contrary to most studies that we are aware of in the search literature, we are able to use data from a large set of about 150 categories of grocery shopping products. This allows us to analyze whether consumers' search efforts vary across categories and if so whether category characteristics are correlated with variation in search activity.

The starting point to an investigation of search patterns across categories is to ask why we would expect to see any difference in search behavior at all. In terms of the model presented earlier, the two drivers of search duration are search costs and search benefits, where the latter depend on the degree of price dispersion in the category. If, for ease of exposition, we assume for the moment that search costs do not vary by category, price dispersion is the one key category characteristic which should influence how much consumers engage in search activity. The top graph of Figure (2.7) illustrates the search-time / price relationship for two categories, one with high and another with a low degree of price dispersion. In the high price

³⁰All papers allow for some form of heterogeneity in search costs, the reported values roughly correspond to the average value of search costs in the respective paper.

dispersion category the gain from extending search by a small amount are large for any given level of search-time. This translates into the slope of the curve being steeper at any level of search activity.³¹ This implies that if a consumer were to search the same amount of time in both categories, then she would forgo higher potential benefits from search in the high dispersion category relative to the low dispersion one. In other words she could reallocate her search-time across the two categories in order to achieve lower total expenditure. For a rational consumer our model would therefore predict that she will search for longer in the high dispersion category and that the benefits from search at the margin are equated. In this section we set out to test this prediction: do consumers adjust their search-time across categories in reaction to differences in the benefits from search and do they do it in such a way that marginal benefits are equalized?

In order to implement such a test, we compute price dispersion at the category-level and split our sample into pick-ups of products from categories with above and below median price dispersion. We start by testing whether search-time differs systematically for purchases in categories with different degrees of price dispersion. To this end we regress search-time on a constant as well as a dummy for whether the purchased item belongs to a high (above median) price dispersion category. Results are reported in column (2) of Table (2.7). Consistent with the prediction outlined above we find a positive and significant effect of 0.907. Relative to an estimated constant of 9.770 this represents a roughly 10 percent increase in search-time. In order to rule out that this difference originates from different types of consumers buying in different categories, we report results from a regression with trip fixed effects in column (3) and find a very similar coefficient magnitude. While we think that consumer pool differences across categories are unlikely to be quantitatively important in the grocery shopping context, it is re-assuring to see that we find a very similar coefficient magnitude in this specification. Next, we run our baseline IV-regression of price paid on search-time but also include an interaction of the high price dispersion dummy with search-time.³² Relative to our baseline regression, which is reported in column (1) of the table, we estimate an insignificant negative coefficient on search-time and a negative and significant coefficient on the interaction term.³³

³¹For ease of exposition and without loss of generality we assume that the two categories have the same minimum price.

 $^{^{32}}$ We use speed and speed interacted with the high price dispersion dummy as instruments.

³³Note that the interaction term on its own is significant only at the 10 percent level. However, the effect of search for high dispersion categories, i.e. the search-time coefficient plus the interaction term, is significantly different from zero at the 1 percent level (coefficient (standard error) of -0.0360 (0.0065))

This implies that the marginal returns from search are significantly higher in categories with high price dispersion. The difference is also economically significant with estimated gains of \$2.16 versus \$0.9 per minute for the two types of categories. Together with the results from columns (2) and (3) our analysis shows that consumers do extent their search time when benefits are higher, but not enough to equate marginal benefits between categories. Consumers therefore have higher unrealized search benefits in high price dispersion categories and in relative terms "leave more money on the table" in those categories. At the end of this section we explore in more detail how those differences can be explained.

We next investigate whether other category characteristics apart from price dispersion have an influence on search behavior. For this purpose we picked a set of another three characteristics for which we considered it to be possible to see differences in search effort. Specifically we analyze (1) average inter-purchase spell duration to see whether the frequency at which a product is purchased influences search behavior, (2) average price-level to test whether consumers search more in more expensive categories, (3) number of products to analyze whether product proliferation might make search more difficult. We run the same set of regressions as we did for price dispersion difference by defining the category-level median for each characteristic in turn. Results are reported in Table (2.8). We find that longer inter-purchase spells, higher average-price and a larger number of products are all associated with an increase in average search-time. However, only for the case of purchase frequency do we find any difference in marginal search benefits. For the other two characteristics the interaction effect is insignificant and very small in magnitude. Interestingly, in terms of purchase frequency we find that consumers search significantly more in categories that they purchase less frequently. Moreover, they extend their search-time in such a way that the marginal benefits from search are lower. In other words, they exhaust the benefits from search more in categories (such as laundry detergent) with longer inter-purchase spells relative to categories with shorter spells (like milk).

Taken together we conclude from the above regressions that there is a substantial amount of variation in the amount of search activity as well as the extent to which search benefits are realized across categories. Most importantly, we find that higher price dispersion categories are characterized by higher marginal search benefits. This could be explained either by consumers' search cost being higher in categories with more dispersion or by consumers under-estimating the potential benefits.³⁴ Regardless of the source, the more limited amount of search indicates that marketing activity which interferes with the search process such as features and displays might be more important and effective in categories with high dispersion. Other category characteristics like the number of UPCs in the category or purchase frequency do not enter the search model outlined in Section (2.4) directly. Instead, they only influence search behavior to the extent that they correlate with price dispersion or if they affect search costs or perceived search benefits. We find significant differences in search benefits only for categories with higher inter-purchase spells which are characterized by more search benefits being realized and therefore lower benefits at the margin. This pattern is consistent with the idea that consumers make a more conscious effort to search in categories they buy from less often. Similarly to high dispersion categories it might therefore be more effective to influence the search process through marketing activity in frequently purchased categories more so than in less frequently purchased ones.

2.9. Conclusions

We estimate the effect of search intensity on the price a consumer pays within a particular category using data from RFID tags on supermarket shopping carts. Recording search in a physical store environment is generally challenging and even our detailed data is only able to capture total search-time, but not which options the consumer evaluated. The technology does however have the advantage of not interfering in any way with the consumer's natural shopping experience and might be the best possible way to gain insights into consumer search in a brick-and-mortar store. To the best of our knowledge this work is the first to use direct data on search effort to analyze consumer search within a brick-and-mortar environment.

We employ a reduced-form approach to estimate the effect of search intensity on the price a consumer pays within a particular category. We find that an additional minute of search lowers expenditure by about \$1.4. The gains from search are substantial, increasing category-level search-time by one standard deviation in each purchased category leads to a 6 percent reduction in total shopping basket expenditure. This result is robust to a host of

³⁴Finally, it could also be the case that the pool of consumers differs across categories. We think this is generally not likely to be a major issue in consumer packaged goods. Furthermore the fact that search-time differences are very similar when including trip fixed effects also suggests that consumer pool differences are not driving the observed differences in search behavior.

sensitivity checks which deal with possible confounds such as variation in prices over time and preferences over product characteristics other than price. Due to the limited amount of observations per category in the data our evidence comes from regressions which are pooled across categories. Going forward, with path-data over a longer time-horizon for only one category it should be possible to model the search process in more detail (possibly by means of a structural model). In particular, our approach only looks at the effect on price paid and does not directly analyze the role of other product characteristics. We are therefore not able to make any statements about the effect of search on consumer utility. However, we believe that the effect of search-time on price is a dimension of the search process which is particularly relevant for informing optimal supply-side behavior. Our findings imply that, due to the limited amount of search, the use of marketing tools such as feature advertising and in-store displays can be very effective. Furthermore, firm behavior that influences consumer search interacts in an interesting way with pricing decisions. Because more search makes finding a lower price or promoted product more likely, firms have an incentive to encourage search when running a promotion.³⁵

Finally, our setting allows us to explore heterogeneity in search behavior across categories. This is of particular interest from a managerial perspective because it speaks to the fact that marketing tools which influence search might be more effective in some categories than others. We find that consumers' unrealized search benefits are higher in categories with higher price dispersion as well as categories with shorter inter-purchase spells suggesting that the role of feature advertising and in-store displays might be particularly important for this set of categories.

³⁵The data and empirical approach could also be used to study seasonal variation in search behavior which (as posited by Haviv [2013]) might be a source of counter-cyclical pricing.

Bibliography

- ANDREWS, R. L., AND T. C. SRINIVASAN (1995): "Studying Consideration Effects in Empirical Choice Models Using Scanner Panel Data," *Journal of Marketing Research*, 32(1), 30-41.
- BRONNENBERG, B. J., M. W. KRUGER, AND C. F. MELA (2008): "The IRI Marketing Data Set," *Marketing Science*, 27(4), 745-748.
- BRONNENBERG, B. J., AND W. R. VANHONACKER (1996): "Limited Choice Sets, Local Price Response and Implied Measures of Price Competition," *Journal of Marketing Research*, 33(2), 163-173.
- CHEN, Y., AND S. YAO (2012): "Search with Refinement," unpublished manuscript.
- COBB, C. J., AND W. D. HOYER (1985): "Direct Observation of Search Behavior in the Purchase of Two Nondurable Products," *Psychology and Marketing*, 2(3), 161-179.
- DE LOS SANTOS, B. I., A. HORTACSU, AND M. WILDENBEEST (2012): "Testing Models of Consumer Search using Data on Web Browsing and Purchasing Behavior," American Economic Review, 102(6), 2955-2980.
- DICKSON, P. R., AND A. G. SAWYER (1990): "The Price Knowledge and Search of Supermarket Shoppers," *The Journal of Marketing*, 54(3), 42-53.
- DRAGANSKA, M., AND D. KLAPPER (2011): "Choice Set Heterogeneity and the Role of Advertising: An Analysis with Micro and Macro Data," *Journal of Marketing Research*, 48, 653-669.

- HAVIV, A. (2013): "Does Purchase Without Search Explain Counter-Cyclic Pricing?,"unpublished manuscript.
- HONKA, E. (2014): "Quantifying Search and Switching Costs in the U.S. Auto Insurance Industry,"unpublished manuscript.
- HOYER, W. D. (1984): "An Examination of Consumer Decision Making for a Common Repeat Purchase Product," *The Journal of Consumer Research*, 11(3), 822-829.
- HUI, S. K., P. S. FADER, AND E. T. BRADLOW (2009a): "An Integrative Framework and Prospectus for Model Building," *Marketing Science*, 28(2), 320-335.
- HUI, S. K., P. S. FADER, AND E. T. BRADLOW (2009b): "Testing Behavioral Hypotheses Using an Integrated Model of Grocery Store Shopping Path and Purchase Behavior," *Journal of Consumer Research*, 36(3), 478-493.
- HUI, S. K., P. S. FADER, AND E. T. BRADLOW (2009c): "The Traveling Salesman Goes Shopping: The Systematic Deviations of Grocery Paths from TSP Optimality," *Marketing Science*, 28(3), 566-572.
- HUI, S. K., Y. HUANG, J. SUHER, AND J. INMAN (2013a): "Deconstructing the First Moment of Truth: Understanding Unplanned Consideration and Purchase Conversion Using In-Store Video Tracking," *Journal of Marketing Research*, 50(4), 445-462.
- HUI, S. K., Y. HUANG, J. SUHER, AND J. INMAN (2013b): "Estimating the Effect of Travel Distance on Unplanned Spending: Applications to Mobile Promotion Strategies," *Journal of Marketing*, 77(2), 1-16.
- KOULAYEV, S. (2014): "Estimating Demand in Search Markets: The Case of Online Hotel Bookings?, *RAND Journal of Economics*, forthcoming
- McCALL, J. J. (1970): "Economics of Information and Job Search, *The Quarterly Journal* of *Economics*, 84(1), 221-242.
- MEHTA, N., S. RAJIV, AND K. SRINIVASAN (2003): "Price Uncertainty and Consumer Search: A Structural Model of Consideration Set Formation," *Marketing Science*, 22(1), 58-84.

- ROBERTS, J. H., AND J. M. LATTIN (1991): "Development and Testing of a Model of Consideration Set Composition," *Journal of Marketing Research*, 28(4), 429-440.
- SEILER, S. (2013): "The Impact of Search Costs on Consumer Behavior: A Dynamic Approach," Quantitative Marketing and Economics, 11(2), 155-203.
- SEILER, S. (2013): "The Science of Shopping," Marketing Research, 15(3), 31-35.
- STÜTTGEN, P., BOATWRIGHT, AND R. T. MONROE (2012): "A Satisficing Choice Model," Marketing Science, 31, 878-899.

	(1)	(2)	(3)
Sample	All Item Pick-ups	All Item Pick-ups	Trips
Dependent Variable	Pickup-specific Search-Time	Pickup-specific Search-Time	Trip-level Search-Time
Speed 60 Second Prior to Pick-up	-3.600^{***} (0.101)	-3.260^{***} (0.136)	
Average Trip-level Speed			$\begin{array}{c} -4.766^{***} \\ (0.171) \end{array}$
F-Stat Trip FEs	1,267 No	572 Yes	779n/a
Observations	28,603	23,446	12,286

Table 2.1.: Correlation between Search-Time and Speed

Notes. Reports coefficients of OLS regressions of search-time on speed. Columns (1) and (2) use specific item pick-ups, column (3) uses trips as the unit of observation. Trips with only one pick-up are dropped in column (2). No further controls are added except trip fixed effects for column (2). Standard error are clustered at the consumer-level.

	Mean	S.D.	P25	Median	P75
Absolute Difference between Daily Min and Max Price	4.21	3.52	1.78	3.21	5.26
Percentage Difference between Daily Min and Max Price	0.66	0.23	0.53	0.71	0.84
Fraction of UPCs promoted on a specific day	0.3	0.17	0.17	0.29	0.42
Fraction of UPCs promoted during the sample period	0.58	0.32	0.41	0.62	0.83

Table 2.2.: Descriptive Statistics: Prices

Notes. The unit of observation for all distributions of price differences / fractions of promoted items is a category. There are 150 categories in our data.

	(1)	(2)	(3)
Dependent Variable	Search-Time	Price Paid	Log(Price Paid)
Share of Promoted UPCs within the Category	-0.965^{*} (0.542)	-0.563^{***} (0.147)	-0.188^{***} (0.036)
Average of the Dependent Variable	10.322	3.268	0.894
Category FEs Day FEs	Yes Yes	Yes Yes	Yes Yes
Observations	28,603	28,603	28,603

Table 2.3.: The Effect of Category-level Pricing on Search

Notes. Search-time at the item pick-up level is regressed on the number of promoted item within the category of the purchased product. No further control variables (other than the indicated fixed effects) are used.

	(1)	(2)	(3)
Type of Regression	OLS	IV: 1st Stage	IV: 2nd Stage
Dependent Variable	Price	Search Time	Price
Search-Time	-0.0057*** (0.0016)		-0.0275^{***} (0.0064)
Speed	(0.0010)	-3.600^{***} (0.101)	(0,000 1)
First-stage F-stats		1,267	
Category FEs	Yes	Yes	Yes
Observations	$28,\!603$	28,603	28,603
Trips Consumers	$12,286 \\ 7,882$	$12,286 \\ 7,882$	$12,286 \\ 7,882$

Table 2.4.: Baseline OLS and IV regressions

Notes. Standard errors are clustered at the consumer-level.

	(1)	(2)	(3)	(4)
Type of Regression	IV: 2nd Stage	IV: 2nd Stage	IV: 2nd Stage	IV: 2nd Stage
Dependent Variable	Price	Price	Promotion Dummy	Promotion Dummy
Search-Time	-0.0275^{***} (0.0064)	-0.0206^{**} (0.0088)	0.0041^{***} (0.0016)	0.0044^{*} (0.0024)
First Stage F-stat Category FEs Trip FEs	1,267 Yes No	572 Yes Yes	475 Yes No	336 Yes Yes
Observations	28,603	23,446	19,718	14,498

 Table 2.5.: Robustness Check: Fixed Effect Regressions and Promotional Dummy as Dependent Variable

Notes. Standard errors are clustered at the consumer-level in columns (1) and (3) and at the trip-level in columns (2) and (4). Sample size changes due to the fact that we exclude trips with only one pickup when including trip fixed effects and the promotion dummy is only is defined only for products for which we see regular purchases.

	(1)	(2)	(3)	(4)	(5)
Type of Regression	IV: 2nd Stage	IV: 2nd Stage	IV: 2nd Stage	IV: 2nd Stage	IV: 2nd Stage
Dependent Variable	Price	Price	Price	Price	Price
Instrument	Speed 60 Sec. Before Pick-up	Speed Lagged by 10 Sec.	Speed Lagged by 20 Sec.	Speed Lagged by 30 Sec.	Trip-level Av. Speed
Search- Time	-0.0275^{***} (0.0064)	-0.0392*** (0.0141)	-0.0446*** (0.0144)	-0.0550^{***} (0.0171)	-0.0299** (0.0142)
First Stage F-stat Category FEs	1,267 Yes	406 Yes	222 Yes	144 Yes	580 Yes
Observations	28,603	26,159	24,892	23,840	28,603

Table 2.6.: Robustness Check: Lagged and Trip-Level Speed Instruments

Notes. Standard errors are clustered at the consumer-level.

Dependent Variable	(1)	(2)	(3)	(4)
	Price	Search-Time	Search-Time	Price
Search-Time Search-Time * Above Median Price Dispersion Dummy Above Median Price Dispersion Dummy Constant	-0.0275*** (0.0064)	$\begin{array}{c} 0.907^{***} \\ (0.098) \\ 9.770^{***} \\ (0.069) \end{array}$	0.849*** (0.129) n/a	-0.015 (0.009) -0.021* (0.011)
Category FEs	Yes	No	No	Yes
Trip FEs	No	No	Yes	No

Table 2.7.: Cross-Category Heterogeneity: Price Dispersion

Notes. Standard errors are clustered at the consumer-level in columns

(1) and (4) and at category-level in columns (2) and (3).

		(1)	(2)	(3)
	Dependent Variable	Search-Time	Search-Time	Price
Inter- Purchase	Search-Time			-0.037^{***} (0.010)
Duration	Search-Time * Above Median Dummy	1 409***	1 000***	0.023^{**} (0.011)
	Above Median Dummy	$(0.097) \\ 9.502^{***}$	(0.126) n/a	
	Constant	(0.068)		
Average Price	Search-Time			-0.024^{***}
Level	Search-Time * Above Median Dummy		0 00 04444	-0.003 (0.010)
	Above Median Dummy	$\begin{array}{c} 0.577^{***} \\ (0.098) \\ 9.944^{***} \end{array}$	0.886^{***} (0.127) n/a	
	Constant	(0.068)	,	
Number	Search-Time			-0.022^{***}
UPCs	Search-Time * Above Median Dummy			-0.006 (0.011)
	Above Median Dummy	$\begin{array}{c} 0.288^{***} \\ (0.097) \\ 10.049^{***} \end{array}$	0.262** (0.127) n/a	
	Constant	(0.068)	/ ~	
	Category FEs	No	No	Yes
	тир г ця	INO	168	INO

Table 2.8.: Cross-Category Heterogeneity: Other Category Characteristics

Notes. Standard errors are clustered at category-level in columns (1) and (2) and at the consumer-level in column (3).



The picture illustrates a consumer traversing an aisle. Consumer location within the aisle is recorded on a grid of traffic points. Products are located at specific locations on the shelf, which are coded up as a grid of product points. Product points are matched to nearby traffic points. This allows to measure how long a consumer remained near the product when picking it up. The dashed black line denotes the consumer's path when traversing the aisle.



Figure 2.2.: Search-Time Histogram



Figure 2.3.: Variation in Search-Time Across and Within Trips



Figure 2.4.: Speed Histogram



Figure 2.5.: Variation in Speed Across and Within Trips

Figure 2.6.: Estimated Local Average Treatment Effect



The picture illustrates the local nature of the estimated search benefit. The relationship between expected price paid and search-time for varying levels of search costs is represented by the solid line. The relationship is non-linear with extensions in search-time from a lower level being associated with larger gains. Moreover the potential gains within a category are bounded by the lower bound of the price distribution. The magnitude of our estimate depends on whether consumers in our data search relatively little (red scatter-plot) or a lot (blue scatter-plot). In the latter case the average consumer realizes more of the potential gains from search and the incremental benefit at the margin is therefore smaller.

Figure 2.7.: Treatment Effects Across Categories with varying Degree of Price Dispersion



The source of differences in search benefits is variation in the degree in price dispersion. The graph shows the relationship between search-time and price paid for two categories with high (low) price dispersion respectively. For ease of exposition in the lower bound of the price distribution is assumed to be the same. The high price dispersion category is characterized by higher incremental benefits from search which lead to a steeper slope of the curve at every level of search-time (illustrated in the top graph). If search costs are identical across categories, then consumers should equate the benefits from search across categories by extending their search-time in the high dispersion category (illustrated in the bottom graph).

Chapter 3. Bonus schemes under price-delegation: Evidence from furniture sales

Abstract

What is the effect of a performance-related pay scheme when it is used in combination with price delegation? Using data on the staff of a furniture firm, I show that when a fixed bonus scheme conditional on revenues was introduced, it increased the revenues generated by all sales employees, but I find no significant heterogeneous effect of the bonus scheme depending on whether the employee is given control over price or not. Moreover, I show that giving the sales staff control over price does not significantly increase revenues. The effects of the bonus scheme and of price delegation on gross profits minus paid bonuses, commissions, and wages were similar. These results are robust to a number of checks, and are consistent with a model of moral hazard and price delegation.

3.1. Introduction

This is one of the first works that provides empirical evidence on the incentive problem between a firm and the staff of its sales force when the latter are given control over the price¹. When the firm cannot perfectly observe the efforts of its sales staff, misalignment of objectives between the firm and the employees can lead to under-provision of services to the end-customers. Similarly, if the firm gives the sales employees control over the price, interest conflicts between the two parties might induce the sales staff to set a sub-optimal level of price, with consequences for the firm's profitability. This is especially the case when a firm incentivizes its sales staff with a pay that is conditional on revenues, like the one I analyze in this chapter. Consider the situation where a firm compensates its sales staff with a fixed salary, commissions as a share of revenues, and a fixed bonus conditional on revenues (see Figure 3.1). This compensation scheme creates conflicts between the interests of the firm, which maximizes profits, and those of the employees, who maximize utility and do not take account of the firm's costs. In addition to the classic moral hazard problem (see Bolton and Dewatripont [2004]), a revenues-based pay gives the employee an incentive to set the price below the level that the firm would optimally set when it sets the price centrally. On the one hand, Weinberg [1975] shows this conflict does not arise when salesmen are paid a commission based on gross margin and are given control over price². On the other hand, gross profits are a noisier measure than revenues for detecting the employee's effort, as the former also include costs. This is the setting where the firm studied in this chapter operates. The firm prefers bearing the cost associated with the price distortion generated by revenues-based pay because it believes profits are too noisy a measure of the employee's effort and, thus, will not be a good performance-measure for tackling moral hazard issues. To eliminate the price distortion created by revenues-based pay, the firm could set the price centrally. However, by setting the price centrally, the firm prevents the employee from exploiting some private information she might have about the customers that can be valuable for the firm's pricing³.

¹The only paper I found that explicitly studies the interaction between bonuses and price delegation is the laboratory experiment designed by Ham and Lim [2013], though the aim of their paper is to understand reciprocity motives in the principal-agent relationship and how this relates to price delegation. Moreover, Ham and Lim's work studies the impact of the bonus scheme only when the principal gives the agent control over price but not when the principal sets the price centrally.

 $^{^{2}}$ In this chapter I study bonuses but the logic is the same as for the case of commissions.

³Examples that highlight the importance of the employee's private information can be found in the finance literature. Baron [1982], for instance, studies investment banking advising and distribution services when the investment banker is better informed about the capital market than is the issuer and the issuer cannot
In summary, the firm faces a trade-off between attenuating moral hazard and extracting more consumer surplus by allowing the employee to price discriminate the customers based on her private information.

I present a stylized principal-agent model between a firm and its sales employee. The model accommodates moral hazard, as the firm cannot perfectly observe the efforts exerted by its employee. The model further allows for the possibility that the employee has private information about the market that can be valuable to the firm. In this setting, I study the implications of augmenting the power of incentives in the employee's compensation scheme via the introduction of a bonus scheme which is conditional on revenues. In particular, I study how the employee responds to the bonus scheme under two scenarios. Firstly, I study the equilibrium of the model when the firm gives the employee control over price. Secondly, I study the equilibrium when the firm sets the price centrally. The model predicts that the bonus scheme increases revenues and gross profits when the customers are responsive to services provision (represented by the agent's effort). If the employee has private information that is valuable to the firm (what I call *innovative news*), the model predicts that revenues and gross profits are larger when the firm gives the employee control over price. The model also predicts that, the greater the disparity between the employee's private information and what the firm normally expects about the customers' willingness to pay, the stronger the impact of the bonus scheme on revenues and gross profits is when the firm gives the employee control over price.

I then take the theory to the data. This work exploits a quasi-natural experiment that took place at a furniture manufacturing company between 2003 and 2011. In 2006, the furniture firm introduced a bonus system to *almost* all its sales employees. In 2010, the firm started to *gradually* phasing out the bonus system, eliminating it completely at the end of 2012. Another interesting feature of the data is the variation in price delegation. The firm sells office furniture to dealers or directly to companies. The sales-staff who directly serve the end-customers are given control over price, while the firm sets the price centrally in the dealers channel. I exploit this additional variation in the decision rights delegated to the sales staff for assessing the implications of price delegation on employees' performance. The empirical identification comes from the quasi-experimental setting where the firm operated. Talking to the firm, the sales director admitted the firm tested the impact of bonuses

observe the distribution effort expended by the banker.

on employees' performance. Thus, bonuses should be independent of employees' individual characteristics that are observable to the firm. Although we do not observe individual characteristics (a part from gender), the panel structure of the data allows to controlling for unobservable individual characteristics that may be correlated with the bonus scheme or with price delegation.⁴

Using a reduced-form approach, this work uses ordinary least squares for estimating the impact of the bonus scheme on a set of performance measures, such as revenues and gross profits. I also study the implications of price delegation and the possibility that the bonus scheme has a heterogeneous effect on the outcome variables depending on whether the firm sets the price centrally or not. I find that the bonus scheme has a positive and significant impact on revenues, gross profits, gross profits minus commissions bonuses and salary, and on quantity. For example, offering the bonus scheme increases quarterly revenues by about 70% relative to the mean value. I also find the bonus scheme has a negative impact on the price. In the setting of the theoretical model discussed in this chapter, these results suggest that furniture sales are responsive to non-price sales efforts. The estimates also show that there is no significant heterogeneous effect of the bonus scheme depending on whether the employee has control over price or not. I also find that price-delegation has no significant impact on any of the outcome variables studied here. In the context of the theoretical model, these results are a sign of the employee having private information that is not sufficiently different from what the firm usually expects (that is, the employee's private signal does not reveal innovative news). These results are robust to a number of checks.

I believe the context where the furniture firm operates is suitable for the study of the incentive problem and of price delegation because of the presence of the firm's corporateclients segment. The firm generates 90% of its revenues from corporate clients such as British Airways or The London School of Economics. Dealing with these customers involves a certain degree of uncertainty regarding their needs because of the complexity of the design projects they request. Thus, employees can gather valuable information about the needs of the corporate clients on a case-by-case basis. This makes the decision to give the employee control over price interesting.

This paper is related to the literature in human resource management as summarized in

⁴The use of individual fixed effects does not fully control for the selection issues that might exist in the assignation of the price delegation as only one employee in our data changes pricing regime passing from the direct channel to the dealers channel.

Bloom and Van Reenen [2010]; to a series of papers in the multi-tasking literature such as Holmstrom and Milgrom [1991], Baker [1992] and Baker [2002], Griffith and Neely [2009]; to the literature in insider econometrics such as Lazear [2000], Bandiera et al. [2007], Misra and Nair [2011], and summarized in Ichniowski and Shaw [2012], which explores how econometric methods can be used to assess productivity of managerial practices within firms by combining observational or experimental data with non-quantitative information provided by the management of the firm; and to Paul Oyer's literature on fiscal year ends and nonlinear incentive contracts Oyer, P. [1998] and Oyer, P. [2000] as the bonus scheme studied here is indeed a non-linear incentive scheme. Finally, this work is related to the industrial economics literature in resale-price maintenance and vertical contracting such as Romano [1994] and Mortimer [2008]; to the marketing literature on sale-force compensation such as Coughlan and Sen [1989], Frenzen et al. [2010], Ghosh et al. [2013], and John et al. [2013].

The remainder of the chapter is organized as follows. Section 3.2 describes the context where the furniture manufacturing firm operates and provides a theoretical framework to guide the empirical strategy. Section 3.3 provides a detailed explanation of the data used in our analysis followed by descriptive statistics. In section 3.4, I explain the empirical strategy and show the main results of the chapter. Section 3.5 discusses a number of empirical issues that might challenge the findings of this chapter. Finally, section 4.8 concludes.

3.2. Context and theory

3.2.1. Context

The firm in our data is one of the national HQ's of a large furniture manufacturing company (henceforth, the firm) that sells home and office furniture. The firm sells furniture either through authorized dealers or directly to the end-customers. The dealers channel and the direct channel generate approximately the same amount of revenues. The customer base is divided by retail customers and corporations, accounting respectively for about 10% and 90% of total revenues generated by the firm over the entire sample. The dealers channel generates 10% of the firm's total revenues from retail customers and 40% from corporate clients, while the direct channel generates 50% of the firm's total revenues from corporate clients only. Clients in the corporations segment usually hire external consultants for managing the complexity and risks involved in their interior design projects. I find this information

useful for justifying the potential role of the employees' private information regarding the preferences of the corporate clients (i.e., their willingness to pay).

For each channel, the firm recruits employees and provides them with compensation plans aimed to achieving its end-goals. Talking to the sales director, it emerged that the firm follows very simple contracting rules when they hire potential employees. I can summarize the hiring process as follows. At the beginning of each year, the firm and each employee (individually) sign the employment contract. At this stage, the firm offers the full compensation plan to the employee and tells her the tasks she is expected to do under the employment contract. The compensation plan includes a fixed salary (determined at this stage) and the commissions the employee will get as a share of the revenues she generates. The compensation plan might also include the possibility for the employee to get a bonus conditional on achieving a certain level of revenues (see the different bonus schemes used by the firm in section B.1 of the Appendix). The threshold on revenues is revealed to the employee at the contracting stage and there is no possibility to renegotiate it for the entire duration of the contract, unless exceptional events such as the recent financial crisis occur. The possibility of renegotiating the targets on revenues in case of exceptional events does not affect the predictions of the theoretical model presented below, as these rare events occur with very low (i.e., zero) probability and do not affect the employee's expected payoff⁵. Finally, the firm tells the employee what tasks she is expected to do. Examples of such tasks include managing existing customers as well as finding new customers.⁶ At this stage, the firm also tells the employee whether she is given control over price or not.

Every year the process is repeated for both the new employees and those who already work for the firm. Some people may argue whether the firm and the employee can negotiate some characteristics of the compensation plan. Talking to the firm, the sales director consults with the employees who already work in the firm regarding what they expect about future trends in the market. Hence, the employee's beliefs are somehow incorporated in the employment contracts. However, the sales director also responds to other guidelines imposed by the

⁵Given the information provided by the furniture firm, there were no renegotiation of targets over the entire sample. However, one can argue that the firm could have had a strong incentive to reduce the targets at the end of 2008 as the shock induced by the collapse of Lehman Brothers might have induced corporate clients to cancel some of their orders. For 2009 things look more reliable as the firm set the targets for 2009 at the end of 2008 taking into consideration the new landscape. This leaves 2008 as the only year where the furniture firm might have renegotiated the targets and should be taken into consideration as a possible issue in the data.

⁶To this end, the firm trains the employees regarding products and sales practices

top management of the firm's world HQ. Thus, employment contracts can be seen as a compromise between the goals of the whole group and the view of the national HQ which is primarily determined by the views of the sales director, the chief financial officer, the managing director and, in part, by the employees' expectations about future market trends.

3.2.2. Theory

I develop a stylized model of the firm to analyse the impact of the bonus scheme in equation 3.2 on the equilibrium mean of a set of performance measures both under the case where the firm sets the price centrally and under the case where the firm gives the employee control over price. The model is tailored to fit the context where the furniture firm operates. The firm's hierarchy has two layers: the top executives (finance, managing and sales directors) who represent the firm, and the employees. For expositional ease, I assume there is one top executive (i.e., the firm) and one employee. The division of tasks can be summarized as follows. The firm designs the employment contract deciding whether to include a bonus scheme or not and whether to delegate the pricing to the employee or to centralize it. The employee chooses how much time to dedicate to existing customers or for finding new customers, as well as how to spend the allocated time (i.e., sales efforts). Moreover, the employee might negotiate the price with the customer if she has the right to do so. To account for the possibility that the employee might have private information that is valuable to the firm, I assume that the employee receives a perfect signal that shifts the mean of the distribution of the market demand.

Formally, consider a situation where a risk-neutral firm hires a risk-neutral person (i.e., the employee) for selling its product. The market demand for the product is represented by the function

$$\begin{aligned} \tilde{q} &= q + \epsilon \\ &= a - bp + dE + \epsilon \end{aligned} \tag{3.1}$$

where p is the price of the product, E represents the non-price efforts exerted by the employee (unobservable to the firm), and a, b and d are non-negative scalar parameters which represent the intercept and the marginal effects of the demand function. The variable ϵ represents a random shock that is unobservable to both the firm and the employee. The firm and the employee form beliefs about the range of values that the shock ϵ can take. In particular, I assume that the firm believes ϵ is a zero-mean random shock with cumulative distribution function F. The employee, on her side, receives a perfect signal x such that the distribution of ϵ conditional on x is identical to F except for the mean which now equals μ (that is, a horizontal shift in F). In other words, the employee has an informational advantage relative to the firm because she receives a signal x that tells her (with certainty) the *true* mean of the distribution of ϵ . I motivate this assumption by the fact that the employee can collect detailed information about the preferences of the customers when spending time with them and, thus, incorporates this information in her decisions. In particular, when the employee has control over price, she can use the information regarding customers for setting the price. When the firm sets the price centrally it prevents itself from incorporating the employee's private information in the price. Note that if the employee has no informational advantage, that is $x = \emptyset$, then the mean of the distribution F conditional on x will be zero, $\mu = 0$, same as the firm's beliefs.

The firm compensates the employee with a fixed salary, W, commissions paid as a share β of the revenues $\tilde{R} = \tilde{q}p$, and an additional fixed bonus, B, if the revenues exceed a given threshold R^* . Formally, the employee is compensated according to the function

$$W + \beta \widetilde{R}, \quad \widetilde{R} < R^*$$

$$W + \beta \widetilde{R} + B, \quad \widetilde{R} \ge R^*$$
(3.2)

For simplicity, I assume that the probability of revenues being larger or equal than the threshold R^* is a linear function of the deterministic part of revenues R. Formally

$$\mathbb{P}(R \ge R^*) = \alpha R \tag{3.3}$$

where R is the deterministic part of \widetilde{R} , and α is a positive scalar parameter that ensures the probability is smaller than one.⁷

I represent this situation as a two-stage game. In the first stage, the risk-neutral firm

⁷You can think of α as a deterministic and decreasing function of the threshold R^* which represents the beliefs of both the firm and the employee. I am assuming symmetry of the beliefs, conditional on the employee's private signal being empty, for simplicity. Further, this model assumes revenues will never be negative, though quarterly revenues happen to be negative for two observations in our data. For the purpose of this work, I believe this model provides a simple framework in support of the empirical findings.

chooses the employment contract, which includes salary, commissions, the bonus and the pricing schemes in order to maximize its expected profits. For simplicity, I consider the salary W, the commission rate β , and the bonus-scheme pair (R^*, B) as fixed, that is, I do not model how the firm would optimally choose the salary, the commission rate and the parameters of the bonus scheme. This leaves the firm with two policies. Firstly, the firm can choose whether to include a bonus scheme or no. Secondly, the firm can decide whether to set the price, p, centrally or give the employee control over price. In summary, the firm has four possible strategies in the first stage: offer no bonus scheme and set the price centrally, offer the bonus scheme and give the employee control over price, do offer the bonus scheme and set the price centrally and, finally, offer the bonus scheme and give the employee control over price.

In the second stage, the risk-neutral employee chooses the level of sales effort, E, and potentially the price p, to maximize her expected utility conditional on the employment contract (set by the firm in the first stage) and the signal x she receives.

The firm's payoff is

$$\widetilde{\pi} = \widetilde{R} - \widetilde{C} - W - \beta \widetilde{R} - 1(Bonus \quad Yes) * 1(\widetilde{R} \ge R^*) * B$$
(3.4)

where $\tilde{C} = c\tilde{q}$ are the firm's variable costs for producing \tilde{q} units of the product and c is the unit cost. 1(Bonus Yes) is an indicator variable that equals one if the firm offers the bonus scheme in the employment contract and zero otherwise⁸. $1(\tilde{R} \ge R^*)$ is an indicator variable that equals one if the revenues generated by the employee will be greater or equal to the threshold R^* and zero otherwise (this captures the possibility that the firm's costs might change depending on the employee's performance). Note the tilde on top of π to indicate the randomness of the firm's payoff in equation 3.4. The employee's payoff is

$$\widetilde{U} = W + \beta \widetilde{R} + 1(Bonus \quad Yes) * 1(\widetilde{R} \ge R^*) * B - \frac{E^2}{2}$$
(3.5)

where $\frac{E^2}{2}$ is the employee's cost of exerting effort level E.⁹ Equation 3.5 makes no dis-

 $^{^{8}1(}Bonus Yes)$ is a choice variable for the firm and it is known to the employee in stage two.

⁹Note that $\frac{E^2}{2}$ does not depend on any employee-specific parameter. This does not affect the interpretation of the coefficient of effort, d, in the demand function. There, d represents the responsiveness of demand to non-price efforts and solely depends on the preferences of the consumers. I simply do not model heterogeneity in the employee's ability as in Bandiera et al. [2007] or John et al. [2013] because employee's heterogeneity does not significantly alter the estimates when I control for employees' time-invariant

tinction between the asymmetry that arises in the distribution of the optimal level of effort exerted by the employee before and after having achieved the threshold R^* as documented by John et al. [2013]. This simplification of the incentive problem does not change the implication of the model as this work focuses on the interaction between incentives and price delegation, instead of looking at gaming across time periods as in Oyer, P. [1998], Misra and Nair [2011] and John et al. [2013]. In brief, the presence of the bonus scheme in the context of the furniture firm analysed in this chapter can be interpreted as an increase of the power of the incentives offered to the employee. This affects the trade-offs involved in the principal-agent relationship.

I now solve the model. Conditional on its beliefs, the firm maximizes the expected value of equation (3.4)

$$E(\tilde{\pi}) = R - cq - W - \beta R - 1(Bonus \quad Yes) * \alpha B \tag{3.6}$$

Conditional on her beliefs, the signal x, and the employment contract, the employee maximizes her expected utility

$$E(\widetilde{U}|x) = W + \beta E(\widetilde{R}|x) + 1(Bonus \quad Yes) * \alpha B - \frac{E^2}{2}$$
(3.7)

I now show how the mean of the equilibrium outcomes is affected by the use of the bonus scheme and its interaction with price delegation, and how the employee's private information affects the equilibrium. In Table 3.1, I report the equilibrium mean price and effort under the different scenarios. Comparing the prices in Table 3.1, one can see how the bonus and pricing schemes affect the equilibrium of the principal-agent problem. In particular, the optimal price under delegation incorporates the true mean of the demand shock, μ , in contrast to the case where the firm sets the price centrally. Thus, the optimal price under centralized pricing is the same regardless of whether the employee has informational advantage ($\mu \neq 0$) or not ($\mu = 0$). Also, the bonus scheme affects equilibrium prices via the factor αB , which disappears from the solution when the firm does not offer the bonus scheme as 1(Bonus Yes) = 0.

I now show how the bonus scheme and price delegation affect other performance measures

characteristics (see section 3.5). The model can be easily adapted to incorporate employee's heterogeneity in the level of productivity as in John et al. [2013] by setting $\tilde{q} = \theta q + \epsilon$, where θ is the productivity parameter.

such as revenues, gross profits and gross profits minus commissions, bonuses and salary. Since the equilibrium outcomes are non-linear in the parameters, I prefer to calibrate the model using the information available in the data.

3.2.3. Numerical analysis and model predictions

Using annual data for the firm's sales-force, I replace some of the parameters that in this model are considered as exogenous with their sample median.¹⁰ I replace the salary W with its sample median of £35,000. The unit cost, c, is also replaced with its sample median of $\pounds 120$. The commission rate is replaced with the sample median of the firm's expected commission rate for each employee, which equals 0.75%. For each employee, the expected commission rate is computed by plugging the firm's expected revenues into the compensation scheme to get expected commissions and, then, by dividing the latter by the firm's expected revenues. The firm's expected revenues were provided to me by the firm's sales director. In a similar way, I replace αB with the sample median of the firm's expected commissions and bonuses divided by the firm's expected revenues minus the expected commission rate, which equals 0.15%. The parameters of the demand function a and b are estimated from a linear least-squares regression of (observed) quantity on price. Since price is determined by the equilibrium between demand and supply, the estimate of the price coefficient of the demand function, b, can be subject to the simultaneous-equation bias (for instance, see Hamilton [1994]). I use unit cost as an instrumental variable for copying with the endogeneity of price in the regression for the demand function. The estimates for the demand coefficients are a = 14,000 and b = 32. The remaining parameters, d and μ , cannot be directly estimated from the data. I choose to calibrate d with the values 0, 5, 10, 15, 20, 30, 40, and 50 because they cover a wide range of admissible cases where prices and quantities are nonnegative in equilibrium (such that d < 60). In particular, cases where d is larger or smaller than b are covered. I also calibrate the parameter μ in such a way to reflect different information structures between the employee and the firm. For each given value of d, I solve the model using seven values for μ . When the employee receives an empty signal (i.e. $x = \emptyset$) $\mu = 0$. When the employee receives a non-empty signal, the information revealed can be either customary or innovative. The information is customary if μ is within a reasonable

 $^{^{10}}$ In this section, the symbol of the British pound sterling \pounds is used for expositional purposes only as the real currency cannot be revealed for confidentiality issues.

range of values, which I set to be between minus and plus one standard error of the leastsquare estimate of the intercept, a, of the demand function: in our data, $\mu = -2,355$ and $\mu = 2,355$.¹¹ The information is *innovative* if μ is outside the rage of values I consider as *customary*. I propose two ranges of *innovative* values for μ . The first range outside the *customary* range is between minus and plus twice the standard error of the least-squares estimate of a (i.e., $\mu = -4,711$ and $\mu = 4,711$). I consider the values $\mu = -4711$ and $\mu = 4,711$ to be *innovative* as they shift the intercept of the demand quite far from the a = 14,000 the firm expects. The last range of *innovative* values is plus ten-times the standard error of the least-square estimate of a, $\mu = 23,554$, and a large negative values which I set to $\mu = -5,061$.¹² I consider $\mu = 23,554$ and $\mu = -5,061$ as exceptionally innovative news as the firm might expect the intercept a to be so large/small with very low probability. Below, I will show what role these different information structures play in rationalizing the firm's choice to give the employee control over price.

The Tables in Appendix B.2 provide the equilibrium outcomes of the model under different scenarios. Each Table provides sets of equilibrium outcomes for a given value of d. Within each Table, there are twenty-eight sets of equilibrium outcomes which depend on the information structure between the employee and the firm (seven cases), on whether the firm offers the bonus scheme or not, and on whether the firms delegates the pricing or not. The equilibrium outcomes I compute are the price p, the non-price effort E, and a set of expected outcomes conditional on the different realization of the signal x such as the quantity q, the revenues R = pq, the gross profits GP = R - C, and the gross profits minus commissions, bonuses and salary $GP - SC = GP - \beta R - \alpha RB - W$.

The main results I discuss in this chapter are (i) the impact of the bonus scheme on the equilibrium, (ii) the impact of price delegation on the equilibrium, and (iii) the impact of the interaction between the bonus scheme and price delegation on the equilibrium. These results will be tested in section 3.3.

The impact of the bonus scheme on the equilibrium outcomes mostly depends on the responsiveness of the firm's customers to the non-price effort, which is represented by the

¹¹As μ shifts the intercept of the demand function, it comes natural to think of the standard error of the least-squares estimate of a as the delimiter of the range of most likely values μ can take.

¹²The asymmetry between $\mu = 23,554$ and $\mu = -5,061$ comes from the fact that, in this numerical example, I want to exclude the possibility that quantity is negative. Therefore, the size of the negative values μ can take is bounded by the non-negativity of quantity, and minus ten-times the standard error of the estimate of *a* violates this constraint.

demand parameter d. The higher the value of d, the stronger the impact of E in shifting the demand is. The model suggests that the bonus scheme has a positive impact on all equilibrium outcomes, including prices, only if d is sufficiently large. In our example, when d > 5 offering the bonus scheme is always optimal for the firm, in GP - SC terms. On the contrary, when $d \leq 5$, offering the bonus scheme is never optimal for the firm, in GP - SCterms. This is intuitive because, when the unobservable non-price effort E is ineffective (i.e., small d), there is no point for firm to incentivize the employee in this direction.

When d is sufficiently large, the revenues-based bonus scheme has a positive impact on price as well. This result might seem surprising but it is a consequence of the fact that the optimal level of effort E is increasing in the power of the bonus scheme, B, and the equilibrium price increases with the effort level E.

Another interesting result is that the returns to offering the bonus scheme increase as d increases. This result is also intuitive: when the unobservable non-price effort E is highly effective (d is large), the potential benefits to incentivizing the employee in increasing E are large. Note that these results are only slightly altered by the presence of price delegation and by the information structure. The main implications still remain. Hence, the parameter d, which in our model reflects the importance of moral hazard, determines the impact of the bonus scheme on the equilibrium.

These results provide the first set of hypotheses to be tested in the data.

Hypothesis 1 Does the bonus scheme have a positive impact on the equilibrium level of R, GP and GP - SC?

The impact of price delegation on the equilibrium outcomes is more complicated. The impact of price delegation on revenues, R, is positive unless d is very large (in our example, d = 50). When the firm gives the employee control over price, the employee will set p and E to maximize her share of the revenues. When the price is set centrally, the firm maximizes its profits which are a function of the firms' variable costs. Thus, the maximum revenues under delegation should be higher other things equal. However, when d is very large, the impact of moral hazard is more severe. Looking at the Tables in Appendix B.2, we can see that the (positive) difference between the equilibrium price set by the firm centrally and that set by the employee under delegation increases significantly as the parameter d increases. Since the effort E is an increasing function of p in equilibrium (both under delegation and when the firm sets the price centrally), for very large value of d the level of effort E will not be

high enough to shift the quantity demanded significantly to compensate the impact of a high price. As a consequence, very low price and low effort will drive revenues under delegation below the level the firm would achieve if the price is set centrally.¹³

The second set of outcomes includes the price p and the effort E. The impact of price delegation on p and E mostly depends on the information structure. In particular, price delegation reduces p and E relative to centralized pricing if the employee receives an empty signal (i.e., $\mu = 0$) or if the employee receives a non-empty and negative signal (i.e., $\mu < 0$). When the signal is empty, price delegation reduces equilibrium price as the employee has incentive to under-cut because it does not internalize the costs of the firm. Effort moves in the same direction as price as E is in equilibrium increasing in price. When μ is negative, the employee optimally reduces the price to accommodate the news that the consumer's willingness to pay is lower than previously expected, though her pricing is still distorted because she does not internalize the costs of the firm. On the contrary, equilibrium price is higher under delegation if the employee receives a non-empty signal with a positive and innovative news (in our example, $\mu \ge 4,711$), unless d is very large. The first part of this result is intuitive. As the signal reveals that the consumer is willing to pay more than previously expected, the employee will, under delegation, incorporate this information into the pricing. Moreover, the incentives to raise the price will be sufficiently large to off-set the employee's incentive to undercut only if the news is innovative (μ sufficiently large). What seems less intuitive is the fact that equilibrium price is higher under delegation if the employee receives news unless dis very large. Again, this is a consequence of the fact that the (positive) difference between the equilibrium price set by the firm centrally and that set by the employee under delegation increases significantly as the parameter d increases (see Appendix B.2).¹⁴

I do not discuss the impact of price delegation on the quantity q because this is not intuitive. The last set of outcomes I discuss in this paragraph includes the gross profits GPand the gross profits minus commissions, bonuses and salary GP - SC. The impact of price delegation on this set of outcomes again depends on the information structure. In particular, price delegation always reduces GP and GP - SC if the employee receives an empty signal

¹³This result should not be taken literally as it might simply be the consequence of the simplistic functional form chosen for the demand function.

 $^{^{14}}$ Again, this result should not be taken literally as it might simply be the consequence of the simplistic functional form chosen for the demand function.

(i.e., $\mu = 0$). Moreover, delegation has a positive impact on GP and GP - SC only if the employee receives an non-empty signal and the news contained in the signal is sufficiently innovative (in our numerical example, this corresponds to $\mu \ge 4,711$ or $\mu \le -4,711$). This is intuitive. Delegating the price to the employee generates a negative externality on the firm because of moral hazard. Therefore, delegation is profitable for the firm (in terms of GP or GP - SC) only if the employee has some private information ($\mu \ne 0$) that is valuable to the firm (i.e., μ is sufficiently large in absolute value). In brief, giving the employee has informational advantage over the firm (in GP or GP - SC terms) only if the employee has informational advantage over the firm and the information contained in the signal is innovative. If the employee has private information that is valuable to the firm, price delegation allows the employee to incorporate this information into the pricing and, thus, to extract plenty of surplus from the customer to the benefit of the firm.

These results provide the second set of hypotheses to be tested in the data.

Hypothesis 2 Does price delegation have a positive impact on the equilibrium level of R, GP and GP - SC?

I now show when the bonus scheme has heterogeneous effects on equilibrium outcomes depending on whether the firm sets the price centrally or gives the employee control over price. From our numerical example, the interaction between the bonus scheme and price delegation has a positive impact on R, GP and GP - SC when μ is negative and the employee receives a non-empty signal that reveals innovative news (i.e., $\mu \leq -4,711$ or $\mu \geq -4,711$).¹⁵

These results provide the third and last set of hypotheses to be tested in the data.

Hypothesis 3 Does the bonus scheme have heterogeneous effects on the equilibrium level of R, GP and GP - SC depending on whether the firm sets the price centrally or gives the employee control over price?

¹⁵The employee increases effort *E* only if she receives a non-empty signal that reveals an extremely innovative news. In our example, this corresponds to $\mu = 23,554$, ten-times the standard error of the least-squares coefficient of the demand intercept *a*. I do not see this when $\mu < 0$ has the size of negative shocks are bounded by non-negativity constraints in equilibrium quantities.

3.3. Data and descriptive statistics

This chapter makes use of detailed individual-level panel data from one of the national HQ's of a multinational furniture manufacturing company between January 2003 and December 2011. The data used in this chapter comes from two sources of information.

The first dataset contains annual data on the employment contracts signed between the firm and the people employed in its sales force, with extremely detailed information regarding compensations and tasks. The second dataset contains quarterly data on different measures of employees' performance such as revenues, gross profits, transaction prices and quantities, as well as bonuses and commissions accumulated (and paid) at the end of each quarter. I now discuss the annual data on employment contracts more in detail.

Before January 2006, the firm did not use bonuses for compensating its sales employees. In January 2006, the firm introduced a bonus scheme based on (employee-level) revenues to almost all of the sales employees. In January 2010, the firm started phasing the bonus scheme out. In January 2013, the firm abolished the bonus scheme completely. This is out-of-sample information as the dataset stops on December 2011. Talking to the sales director, the firm abolished the bonus scheme because it was perceived as ineffective. Table 3.2 clearly shows how the bonus scheme was implemented over the sample period. The Table is divided in two macro columns. The left-hand column (called Centralized pricing) shows how the bonus was introduced for those employees who had no control over price. The right-hand column (called Delegated pricing) shows how the bonus was introduced for those employees who did have control over price. For each of these two macro columns, an observation is a year-quarteremployee combination, while a contract is a year-employee pair. The number of observations with B are year-quarter-employee combinations where the employment contract includes a bonus scheme. In other words, B is an indicator variable, the bonus dummy, that equals one if the contract includes a (conditional) bonus and zero otherwise. At the moment the contract is signed, any employees' performance is random. Thus, B represents an ex-ante characteristic of the employment contract.

For the more curious readers, Table 3.3 provides each employee's history with the firm and clearly shows how the bonus scheme was implemented over the years. Each column represents the history of a specific employee. The rows represent the years. When the number in the cell equals four (for quarters), the employee is employed by the firm in that year. If this number equals zero the employee is not employed by the firm.¹⁶ If the number four is circled, then the contract includes a bonus scheme. If the cell is shaded, the contract gives the employee control over price. Combining these different sources of information, one can easily identify the sources of variation present in the data. Table 3.3 also shows how the bonus regime changes over time for the same employee. To see this, consider employee M01 who stays in the firm from 2004 to 2006. M01 has control over price for the three years he works at the firm and he is given a bonus scheme only in year 2006. An important question to be considered is whether the bonus scheme was introduced randomly or not. Talking to the firm, the sales director admitted the firm tried to assess the impact of bonuses on employees' performance. Thus, the bonus scheme should be independent of employee's characteristics that are observable to the firm. Unfortunately, we cannot test whether this is true in the data as gender is the only individual characteristic available. Table 3.4 summarizes the four sources of variation in the data.

Another important question concerns whether the employees were randomly assigned the control over the price. The measure of price delegation used in this chapter coincides with the definition of sales channel. Some employees serve *dealers* while others serve end-customers *directly.* Pricing in the dealers channel is set at a central level through dealership contracts signed at the beginning of each year, which prescribe quantity discounts and other commercial terms and conditions. The employees in the direct channel have control over price as the firm believes there are significant gains from using employees' private information for better price-discriminate end-customers. Nonetheless, the nature of the corporations segment managed via the dealers channel (40% of total revenues) does not preclude the possibility that employees in this channel might have valuable private information regarding the preferences of their customers, as the dealers they serve may be in turn serving some corporate clients. Thus, the sales skills required to operate in the two channels are in principle the same. Talking to the sales director, he recognizes that employees in the direct channel are recruited in a similar way of those employed in the dealers channel. Unfortunately, our data lacks of individual-level characteristics, such as education, age and work experience that could be used to test any systematic allocation of the employees across the sales channels. In section 3.5, I try to assess the selection issues that can be embodied with both the assignation of the bonus scheme as well as price delegation by controlling for employee-level annual targets

¹⁶This number is set equal to zero also for employee M05 in 2010 who was employed by the firm but is dropped because of excessively different compensation scheme.

on revenues as well as for employee fixed-effects.

The second dataset contains quarterly data on different measures of employees' performance such as revenues, gross profits, transaction prices and quantities, as well as bonuses and commissions accumulated (and paid) at the end of each quarter. Table 3.5 reports basic summary statistics of the outcome variables used in our empirical study. The unit of observation is a year-quarter-employee combination. For each employee, revenues are computed by multiplying the quantity by the price per transaction, and then aggregating at the quarterly level. Gross Profits are revenues minus variable costs for producing the products sold in each quarter. Gross Profits-SC equals to gross profits minus bonuses, commissions, and salary paid at the end of each quarter.¹⁷ All outcomes other than quantity are deflated using country-specific CPI with 2005 as base year. For each employee, the unit price at the quarterly level is computed as the ratio between total revenues over total quantity. Price is missing for those quarters where an employee recorded nil revenues.¹⁸ I dropped two observations with negative revenues as these represent an exceptional event in the firm history: the revenues were negative due to products returned by a large customer. Negative revenues can regularly occur for individual transactions but (almost) never occur as the total over the quarter. In section B.3 of the appendix, I show the main findings of the chapter when the two observations with negative revenues are also included in the estimation. Table 3.6 reports the same basic summary statistics of the outcome variables aggregated at the annual level (i.e., the unit of observation is a year-employee pair). I show also annual statistics for providing a more informative view of the firm. In particular, Table 3.6 shows that Gross Profits-SC were observed to be negative only once. This means that in a given year the firm lost money on a specific employee. I believe this to be reasonable as it only represents about 2% of the observations. Appendix B.3 also provides the percentiles of annual Gross Profits-SC for further reference.

I conclude this section by providing the sample means of the outcome variables grouping the employees according to the different treatments. Table 3.7 reports the mean and the standard deviation of the outcome variables depending on whether the employment contract includes the bonus scheme or not, using year-quarter-employee as a unit of observation. Table 3.7 shows that the bonus scheme has a positive impact on the average value of the

¹⁷These are ex-post bonuses and commissions and can differ from the expected values that emerge ex-ante from the compensation plan.

¹⁸In principle, I could use the firm's list price though this is only available for a small subset of transactions.

outcomes, with exception for price. In the context of the theoretical model of section 3.2.2, this result suggests that the demand parameter d is sufficiently large, that is, the nonprice effort E effectively shifts the firm's demand.¹⁹ Table 3.8 reports the mean and the standard deviation of the outcome variables depending on whether the employee has control over price or not, using year-quarter-employee as a unit of observation. Table 3.8 show the average value of all outcomes. With exception for price, this is higher for those employees who cannot independently set the price. In our theoretical model, this is consistent with the product demand being extremely responsive to non-price effort E (large d) and with the employee having not sufficiently innovative informational advantage over the firm.²⁰ The last set of descriptive statistics shown in this section appears in Table 3.9. There, I report the mean of each outcome variable for the four treatment groups determined by the combinations between the bonus scheme and price delegation. The Table is divided into two macro columns. The left-hand column called Centralized pricing and the right-hand column called Delegated pricing. Each macro column is divided into two sub-columns, one for those observations where the firm offers the bonus scheme and another for those observations where the firm does not offer the bonus scheme. Table 3.9 suggests that the introduction of the bonus scheme has a heterogeneous effect across the two pricing groups. In particular, the bonus scheme seems to have a higher average effect on Gross Profits and on Gross Profits-SC for those employees who have control over price. This might be consistent with the presence of significant informational advantage on the side of the employee, though the average impact of the interaction between the bonus scheme and price delegation on revenues points to the other direction. At this stage, it is important to recognize that there is the possibility of a heterogeneous effect of the bonus scheme across the two pricing groups. Other factors such as business-cycle fluctuations and within-year seasonality can affect the average effects of the bonus scheme and of price delegation.

¹⁹Another interesting result is that the presence of the bonus scheme seems to reduce the dispersion of the employees' performance, something in contrast with the empirical findings in Bandiera et al. [2007]. I do not investigate the impact of incentives on the dispersion of performance.

²⁰Another interesting results is that price delegation seems to reduce the dispersion of the employees' performance, something apparently new in the literature. Again, I do not investigate the impact of price delegation on the dispersion of performance.

3.4. Empirical model and estimates

In order to analyze the impact of the bonus scheme and price delegation on the outcome variables presented in the previous section, I use ordinary least squares to run the following regression

$$Y_{it} = \alpha + \beta * B_{it} + \delta * PD_{it} + \gamma * PD_{it} * B_{it} + \eta_t + \varepsilon_{it}$$
(3.8)

 Y_{it} is a quarterly, individual-level outcome such as revenues or gross profits. B_{it} is an indicator variable that equals 1 if employee *i*'s compensation at time *t* prescribes the *possibility* of getting a fixed bonus if cumulated annual revenues up to time *t* exceed a given threshold, and equals 0 otherwise.²¹ PD_{it} is an indicator variable that equals 1 if employee *i*'s contract at time *t* gives her control over price, and equals 0 otherwise. η_t are time fixed effects computed as the sum of year and quarter dummies. Formally, $\eta_t = \sum_{s=2003}^{2011} year_{s,t} + \sum_{j=1}^{4} quarter_{j,t}$, where $year_{s,t}$ is an indicator variable that equals 1 if period *t* occurs in year *s* and equals 0 otherwise, and *quarter_{j,t}* is an indicator variable that equals 1 if period *t* occurs in the *j*-th quarter and equals 0 otherwise. η_t is always included in all specifications because of the presence of significant business-cycle fluctuations across years and of quarterly seasonality in furniture sales. Appendix B.3 provides empirical evidence in support of this choice. I cluster standard errors at the employee-level to allow for an arbitrary within-employee correlation of the error terms. Results are reported in Table 3.10.

I find that the bonus scheme has a positive and significant²² impact on revenues, gross profits, gross profits minus commissions bonuses and salary, and on quantity. For example, offering the bonus scheme increases quarterly revenues by 252,729, about 70% increase relative the mean of 350,229. This increase looks quite large. There could be other factors that may be driving this result, and the large variation in the outcome variables definitely plays an important role: for example, standard deviation of quarterly revenues in our sample is 269,244, about 80% of the sample mean. Despite the fact that the bonus scheme has a negative impact on the price, these results suggest that the demand coefficient d in the setting of the model of section 3.2.2 might be sufficiently large: that is, the firms' customers seem to be sufficiently responsive to the non-price sales effort, E. Thus, our findings provide

 $^{^{21}}B_{it}$ is an *ex-ante* characteristic of the employment contract and will equal one even if (*ex-post*) cumulated annual revenues up to time t do not exceed the given threshold.

 $^{^{22}\}mathrm{At}$ least at the 10%.

some evidence in support of hypothesis 1.

There seems to be no significant evidence for the presence of a heterogeneous effect of the bonus scheme depending on whether the employee has control over price or not. The coefficient of the interaction Bonus * PD is not statistically different from zero, except for price. Thus, our findings do not provide evidence in support of hypothesis 3.

Table 3.10 also shows that price-delegation has a negative impact on revenues and on the two measures of gross profits, though the coefficient of the price-delegation dummy, PD, is not significantly different from zero. In the context of our theoretical model, I can interpret these results as a situation where the employee might have some private information about the customers' willingness to pay, but this information is not sufficiently innovative to off-set the distortion created by revenue-based incentives on pricing. In particular, looking at the impact of price delegation on gross profits minus bonuses commissions and salary (column (3) Table 3.10), I find some evidence that price delegation does not increase the profitability of the firm. This result does not contradict other results in the literature. For example, Stephenson et al. [1979] find that price delegation reduced profitability among a sample of 108 firms. On the other hand, Weinberg [1975] suggest that price delegation can increase profits if the employees are compensated with commissions based on gross profits, instead of revenues-based pay. In the context of the theoretical model of section 3.2.2, compensating the employee with bonuses and commissions based on gross profits would reduce the distortion on pricing created by revenues-based pay. For price delegation to have a positive impact on gross profits, the employee's private information when commissions are based on gross profits would need to be less innovative than the information required under revenues-based pay. In our data, there are 56 observations where the employees are paid with commissions based on gross profits. Unfortunately, when the employee is compensated based on gross profits she is never offered the bonus scheme and she always has control over price. To overcome this shortcoming of the data, I restrict the sample to only those observations where the employees are not offered the bonus scheme and they are given control over price. Of these 111 observations, 55 refers to employee-time pairs where the employee is compensated with revenues-based pay, while the remaining 56 observations refer to cases where the employee is compensated with pay based on gross profits or a mix between gross profits and revenues. Table 3.11 reports the results from the ordinary least squares regression where the dependent variable is one of the five outcomes discussed in this chapter and the independent variable is the indicator variable *Revenues* – *based* which equals 1 if, in a given year, the employee is compensated with a revenues-based pay and zero if the employee is compensated with a pay based on gross profits or a mix between gross profits and revenues. I find that compensating the employee with revenues-based pay has no significant impact on the outcome variables. This reinforces the findings in Table 3.10. Price delegation does not significantly increase revenues or gross profits in the context of furniture sales even if the firm compensates its sales-force with pays based on gross profits. Thus, our findings do not provide evidence in support of hypothesis 2.

In summary, the findings of this chapter suggest that sales employees in the furniture industry are responsive to revenue-based bonus schemes. These findings also suggest that price delegation has not significant impact on employees' performance.

3.5. Robustness checks

There are a number of empirical issues that might challenge the findings of this chapter. For example, if the firm assigned the bonus scheme to those employees conditional on certain individual characteristics, such as age education or work experience that are observable to the firm, then the estimates presented in section 3.4 may be biased. Further, if the firm correlated the assignment of the bonus scheme to other characteristics of the compensation scheme such as target on revenues or salary, then the estimates may be biased.

The presence of systematic heterogeneity between the employees in direct channel (who have the right to set price) and those in the dealers channel imposes additional challenges to the identification strategy proposed in this chapter. One can think of a number of differences between the direct channel and dealers channel. For example, the firm might sort employees across the sales channels according to some individual characteristics which is observable to the firm but not to the econometrician. The compensation plan may systematically differ across the two sales channels: for example, the salary can be significantly lower/higher for those employees in the direct channel; or, the price-elasticity of demand may differ across the sales channels.

The use of quarterly data in our empirical investigation might impose additional challenges to the findings of this chapter. As employment contract are set at the beginning of each year and are valid for the entire year, by using quarterly data in the estimation I disregard the possibility of within-year dynamics in the behavior of the employees as a response to the incentives provided by annual target on revenues.

Finally, the sample size and the presence of outliers can make the estimates less stable.

In this section, I propose a number of robustness checks that attempt to assess the potential biases the aforementioned empirical issues can impose on the estimates found in this chapter. The results I discuss in this section refer to the Tables in appendix B.3.

3.5.1. Individual characteristics

I include employee fixed-effects in the baseline model of equation 3.8. Employee fixed-effects control for time-invariant observed and unobserved individual-level heterogeneity, such as gender or ability. Thus, if the firm assigned the bonus scheme, or the control over price, according to some time-invariant individual characteristics, the inclusion of employee fixed-effects should significantly affect the estimates of the model 3.8. I find employee fixed-effects do not significantly affect the coefficient of the bonus scheme nor they affect that of the price delegation.²³

3.5.2. Target on revenues and salary

The furniture manufacturing firm designs complex compensation plans. These include salary and commissions, as well as bonuses the employees can get conditional on achieving predetermined targets on revenues (see section B.1 of the appendix). The bonus scheme may be correlated with some other dimension(s) of the compensation plan: for example, target on revenues. Controlling for target on revenues is important because this measure takes account of possible non-random allocation of customers or projects to some employees and, hence, may contain important information regarding the firm's beliefs about future market trends. I find that the inclusion of target on revenues does not significantly affect the coefficient of the bonus scheme nor does it affect the coefficient of the price delegation.

Salaries are another important part of employees' compensation. It seems natural to think the firm sets the salary and the bonus scheme altogether when it designs the employment contract. Salaries may also contain information regarding employees' characteristics such as

²³The use of individual fixed effects does not fully control for the selection issues that might exist in the assignation of the price delegation as only one employee in our data changes pricing regime passing from the direct channel to the dealers channel.

work experience and skills, or at least the firm's beliefs about these. The inclusion of salaries does seem to alter the magnitude of our estimates significantly, although the effect of the bonus scheme on the outcomes is now stronger.

3.5.3. Heterogeneity across pricing groups

Individual characteristics can systematically differ across the two pricing groups. In appendix B.3, I show that the distribution of gender, the only observable individual characteristic, is sufficiently even across the two pricing groups.

The characteristics of the compensation plans can also differ systematically across the two pricing groups. In appendix B.3, I show the mean of target on revenues, salaries, expected commissions and other features of the compensation plans across the two pricing groups. With exception for target on revenues and salaries, the other dimensions of the compensation plans do not show marked differences across the two pricing groups. I also show the distribution of the bonus scheme and the distribution of the size of the bonus for the employee who have control over price are 10% less likely to be offered the bonus scheme, the size of the bonus for those who are offered the scheme looks quite similar across the two pricing groups.

Another source of heterogeneity across the two pricing groups is the types of customers served. The employees who cannot set the price independently serve dealers. These dealers serve both retail customers (10% of the total revenues generated by the furniture firm) and business customers (40% of the firm's total revenues). The employees who have control over price solely serve corporations. A simple statistics that summarizes differences in the customer base across the two pricing groups is the price-elasticity of demand. In appendix B.3 I show that the price elasticity of demand differs across the two sales channels. Unfortunately, the nature of the data does not allow to control for observed heterogeneity across the sales channels because the price delegation indicator variable, PD, is perfectly collinear with the sales channel fixed-effects. The different price elasticity may signal different customer bases across the sales channels. This issue probably represents the biggest limitation to the findings of this chapter.

3.5.4. Outliers

Given the small sample size, extreme values of the outcome variables can alter our estimates. Moreover, the inclusion of dummy variables on the right-hand side of the regression tends to weaken the robustness of linear regression estimators when the sample contains anomalous observations (see Blankmeyer [2006]). I run different specifications using weighting methods that control for outliers and find these do not significantly affect our findings (see appendix B.3)

3.5.5. Within-year dynamics

Employment contract are set at the beginning of each year and last for the entire year: for example, targets on revenues are set on an annual basis. The problem is that I use quarterly data in the estimation. Although the quarter dummy variables partially control for withinyear dynamics, I need to run a stronger check to test whether the results are contaminated by within-year dynamics created by the annual incentive scheme. Using *annual* outcomes as dependent variable should control for within-year dynamics. Annual data should also mitigate excessive within-year volatility in outcomes such as negative or large realizations of revenues or gross profits. In appendix B.3, I show that the estimates of the main specification using annual data look similar to those of the baseline model reported in Table 3.10. I believe the annual model specification provides a very important robustness check because averaging outcome variables for the entire year eliminates plenty of noise and excessive volatility that is present when considering these variables at the quarterly level.

Finally, employees may "game" the compensation scheme by shifting output across years (see Oyer, P. [1998]). Unfortunately, I do not have the variation between annual and fiscal calendars as in Oyer, P. [1998] for testing the presence of "gaming" across years.

3.6. Conclusions

Using data on the staff of a furniture firm, I show that when a bonus scheme conditional on revenues was introduced, it increased the revenues generated by all of the sales employees, but I find no significant heterogeneous effect of the bonus scheme depending on whether the employee is given control over price or not. Moreover, I show that giving the sales staff control over price does not significantly increase revenues. The effects of the bonus scheme and of price delegation on gross profits minus paid bonuses, commissions, and wages were similar. These results are robust to a number of checks, and are consistent with a model of moral hazard and price delegation. We can interpret these results in light of the theoretical model presented in this chapter. The results suggest that the agent might have some private information but this is not enough to fully off-set the negative impact of moral hazard on pricing.

Although these results are robust to alternative model specifications, empirical challenges might be important. For example, the small sample size does not allow to perfectly controlling for all observed heterogeneity. Future works should provide more transparent evidence by using larger data-sets, variation across firms and industries or even engineering ad-hoc randomized-control trials that clearly generate the variation in the data necessary for identifying the implications of the bonus scheme and of price delegation.

Bibliography

- BAKER G. P. (1992): "Incentive Contracts and Performance Measurement," Journal of Political Economy, 100(3), 598-614.
- BAKER G. P. (2002): "Distortion and Risk in Optimal Incentive Contracts," The Journal of Human Resources, 37(4), 728-751.
- BANDIERA, O., I. BARANKAY, AND I. RASUL (2007): "Incentives for Managers and Inequality among Workers: Evidence from a Firm-Level Experiment," *The Quarterly Journal* of Economics, 122, 729-773.
- BARON, D., P. (1982): "A Model of the Demand for Investment Banking Advising and Distribution Services for New Issues," *The Journal of Finance*, 37(4), 955-976.
- BLANKMEYER, E. (2006): "How robust is linear regression with dummy variables?,"Faculty Publications Department of Finance and Economics Texas State University.
- BLOOM, N., AND J. V. Reenen (2010): "Human Resource Management and Productivity,"in Handbook of Labor Economics, ed. by O. Ashenfelter, and D. Card, vol. 4b, pp.1039-1089. Elsevier Science.
- BOLTON, P., AND M. DEWATRIPONT (2004): *Contract Theory* The MIT Press, Cambridge, MA.
- COUGHLAN, A. T., AND S. K. SEN (1989): "Salesforce Compensation: Theory and Managerial Implications," *Marketing Science*, 8(4), 324-342.
- FRENZEN, H., A. K. HANSEN, M. KRAFFT, M. K. MANTRALA, AND S. SCHMIDT (2010): "Delegation of pricing authority to the sales force: An agency-theoretic perspective

of its determinants and impact on performance,"*Intern. J. of Research in Marketing*, 24, 333-355.

- GHOSH, M., F. LAFONTAINE, AND D. HO-FU-LO (2013): "Delegation and Performance Pay: Evidence from Industrial Sales Forces," Ross School of Business Working Paper, Working Paper No. 1190.
- GRIFFITH, R., AND A. NEELY (2009): "Performance Pay and Managerial Experience in Multitask Teams: Evidence from within a Firm," *Journal of Labor Economics*, 27(1), 49-82.
- HAM, S. H., AND N. LIM (2013): "Relationship Organization and Price Delegation: An Experimental Study," *Management Science*, Articles in Advance, 1-20.
- HAMILTON, J. D. (1994): Times Series Analysis Princeton University Press, Princeton.
- HOLMSTROM, B., and P. MILGROM (1991): "Multitask Principal-Agent Analyses: Incentive Contracts, Asset Ownership, and Job Design," *Journal of Law Economics and Orga*nization, 7, 24-52.
- ICHNIOWSKI, C., AND K. L. SHAW (2012): "Insider Econometrics: Empirical Studies of How Management Matters," in The Handbook of Organizational Economics, ed. by R. Gibbons, and J. Roberts, pp. 263-312. Princeton University Press.
- JOHN, G., S. KISHORE, O. NARASIMHAN, AND R. S. RAO (2013): "Bonuses Versus Commissions: A Field Study," *Journal of Marketing Research*, 50(3), 317-333.
- LAZEAR, E. P. (2000): "Performance Pay and Productivity," *American Economic Review*, 90(5), 1346-1361.
- MISRA, S., AND H. S. NAIR (2011): "A Structural Model of Sales-Force Compensation Dynamics: Estimation and Field Implementation," *Quantitative Marketing and Economics*, 9, 211-257.
- MOTIMER, J. H. (2008): "Vertical Contracts in the Video Rental Industry," *Review of Economic Studies*, 75, 165-199.
- OYER, P. (1998): "Fiscal Year Ends and Nonlinear Incentive Contracts: The Effect on Business Seasonality," *The Quarterly Journal of Economics*, 113(1), 149-185.

- OYER, P. (2000): "A Theory of Sales Quotas with Limited Liability and Rent Sharing," *Journal of Labor Economics*, 18(3), 405-426.
- ROMANO, R. E. (1994): "Double Moral Hazard and Resale Price Maintenance," *The RAND Journal of Economics*, 25(3), 455-466.
- STEPHENSON, P. R., W. L. CRON, AND G. L. FRAZIER (1979): "Delegating Pricing Authority to the Sales Force: The Effects on Sales and Profit Performance," *Journal of Marketing*, 43(2), 21-28.
- WEINBERG, C. B. (1975): "An Optimal Commission Plan For Salesmen's Control Over Price," Management Science, 21(8), 937-943.



Figure 3.1.: Compensation scheme

Bonus scheme	Information st	ructure	Pricing scheme				
			Decentralized		Centralized		
		μ	р	E	р	Ε	
Bonus Yes	No advantage	$\mu = 0$	$\frac{a}{2b - (\beta + \alpha B)d^2}$	$(\beta+\alpha B)dp$	$\frac{(1-\beta-\alpha B)a+(b-[\beta+\alpha B]d^2)c}{2(1-\beta-\alpha B)(b-[\beta+\alpha B]d^2)}$	$(\beta+\alpha B)dp$	
	Advantage	$\mu \neq 0$	$\frac{a+\mu}{2b-(\beta+\alpha B)d^2}$	$(\beta + \alpha B)dp$	$\frac{(1-\beta-\alpha B)a+(b-[\beta+\alpha B]d^2)c}{2(1-\beta-\alpha B)(b-[\beta+\alpha B]d^2)}$	$(\beta + \alpha B)dp$	
Bonus No	No advantage	$\mu = 0$	$\frac{a}{2b-\beta d^2}$	eta dp	$\frac{(1-\beta)a+(b-\beta d^2)c}{2(1-\beta)(b-\beta d^2)}$	eta dp	
	Advantage	$\mu \neq 0$	$\frac{a+\mu}{2b-\beta d^2}$	βdp	$\frac{(1-\beta)a+(b-\beta d^2)c}{2(1-\beta)(b-\beta d^2)}$	βdp	

Table 3.1.: Equilibrium outcomes

		<u>a</u>			D 1 . 1 .		
		Centralized pric	ing	Delegated pricing			
Year	Tot $\#$ obs	# obs with B	% obs with B	Tot $\#$ obs	# obs with B	% obs with B	
2003	16	0	0%	8	0	0%	
2004	16	0	0%	12	0	0%	
2005	16	0	0%	20	0	0%	
2006	20	20	100%	20	16	80%	
2007	16	16	100%	20	12	60%	
2008	16	16	100%	20	16	80%	
2009	16	16	100%	8	8	100%	
2010	20	12	60%	0	0	N/A	
2011	12	8	67%	0	0	N/A	
Total	148	88	59%	108	52	48%	

Table 3.2.: Frequency of employment contracts prescribing a bonus ex-ante - by year and pricing rights

Notes. An observation is a year-quarter-employee combination. A contract is a year-employee pair. The bonus dummy B describes the *possibility* for the employee to get extra money conditional on performance. Thus, B represents an *ex-ante* characteristic of the employment contract. Not all people employed in the firm's sales force are included in this data. Reasons for this are missing data and excessive heterogeneity in the contracts with lack of variation in the use of bonuses (namely, these employees are compensated with a threshold on gross profits and no bonus is ever prescribed to them).

Year]	Empl	oyee											Total
	F01	F02	F03	F04	F05	M01	M02	M03	M04	M05	M06	M07	M08	M09	M10	M11	M12	M13	M14	M15	M16	M17	M18	
2003	0	0	4	4	0	0	0	4	0	0	4	0	0	0	0	0	0	4	0	4	0	0	0	24
2004	0	0	0	4	0	4	0	4	0	0	0	0	4	0	4	0	0	4	0	4	0	0	0	28
2005	0	0	0	4	0	4	0	0	4	0	0	0	4	0	4	4	0	4	4	0	0	0	4	36
2006	4	0	0	0	4	4	0	0	4	4	0	0	4	0	4	0	0	4	4	0	0	④	0	40
2007	0	0	0	0	(4)	0	0	0	0	(4)	0	0	(4)	4	4	0	0	(4)	(4)	0	(4)	(4)	0	36
2008	0	0	0	0	Ā	0	0	0	0	$\overline{\mathbb{A}}$	0	0	Ā	(\mathcal{P})	4	0	0	Ā	$\overline{(4)}$	0	$\overline{(4)}$	Ā	0	36
2009	0	0	0	0	Ť	0	4	0	0	Ē	0	0	ō	Ē	0	0	0	Ť	0	0	0	Ť	0	24
2010	0	4	0	0	0	0	4	0	0	0	0	4	0	0	0	0	4	4	0	0	0	0	0	20
2011	0	0	0	0	0	0	Ō	0	0	4	0	4	0	0	0	0	Ō	4	0	0	0	0	0	12
Total	4	4	4	12	16	12	8	8	8	20	4	8	20	12	20	4	4	36	16	8	8	16	4	256

Table 3.3.: Employees' history with the firm

Notes. F refers to female and M to male. Circles refer to contracts that prescribe a bonus. Shaded numbers refer to employees who have control over price. In 2005, Ms F04 switched from not being able to set price to have the right to. In 2010 Mr M5 had a mixed scheme and he was revoked the right to set prices. In 2011 Mr M5 had a scheme based on revenues and he still could not set prices. I had to drop 9 employees because compensated on gross profits and a couple more because of missing info.

Price delegation	Bonus No	Bonus Yes	Total
No	60	88	148
	(23.4%)	(34.4%)	(57.8%)
Yes	56	52	108
	(21.9%)	(20.3%)	(42.2%)
Total	116	140	256
	(45.3%)	(54.7%)	(100.%)

Table 3.4.: Frequency of bonuses and pricing rights

Notes. An observation is a year-quarteremployee combination.

Variable	Obs	Mean	Std. Dev.	Min	Max
Qrt Revenues	254	350,229	269,244	0	1,209,114
Qrt Gross Profits	254	$139,\!682$	$126,\!956$	-47,975	$990,\!857$
Qrt Gross Profits-SC	254	126,724	124,779	-60,539	$979,\!480$
Qrt Quantity	254	$1,\!838$	1,771	0	15,507
Qrt Price	241	236	231	6	$3,\!295$

Table 3.5.: Summary statistics of outcomes

Notes. All outcomes are quarterly, individual-level data. Revenues are total quantity times price. Gross Profits are revenues minus variable costs for producing the product. Gross Profits-SC equals to gross profits minus bonuses, commissions, and salary. All outcomes but quantity are deflated using country-specific CPI with 2005 as base year. I dropped two observations with negative sales. Price is missing for nil revenues.

Variable	Obs	Mean	Std. Dev.	Min	Max
Annual Revenues	64	$1,\!384,\!100$	$883,\!283$	$172,\!470$	3,781,046
Annual Gross Profits	64	$552,\!265$	409,998	50,784	$2,\!602,\!684$
Annual Gross Profits-SC	64	$500,\!355$	$403,\!359$	$-39,\!697$	$2,\!554,\!629$
Annual Quantity	64	$7,\!280$	$5,\!550$	648	$27,\!442$
Annual Price	64	211	72	87	498

Table 3.6.: Summary statistics of annual outcomes

Notes. All outcomes are annual, individual-level data. Revenues are total quantity times price. Gross Profits are revenues minus variable costs for producing the product. Gross Profits-SC equals to gross profits minus bonuses, commissions, and salary, which occur to be negative once. All outcomes but quantity are deflated using country-specific CPI with 2005 as base year.

	Bon	us No†	Bonus Yes‡		
Variable	Mean	Std. Dev.	Mean	Std. Dev.	
Qrt Revenues	280,746	262,742	407,714	261,769	
Qrt Gross Profits	$123,\!418$	$151,\!047$	$153,\!138$	$101,\!486$	
Qrt Gross Profits-SC	$110,\!976$	$150,\!623$	139,752	$97,\!084$	
Qrt Quantity	1,599	1,761	2,036	1,761	
Qrt Price	241	326	232	110	

Table 3.7.: Mean and std. dev. by bonus

Notes. All outcomes are quarterly, individual-level data. All outcomes but quantity are deflated using country-specific CPI with 2005 as base year. I dropped two observations with negative sales. Price missing for nil revenues. † 115 observations except price with 106. ‡ 139 observations with except price with 135.

	Centraliz	ed Pricing [†]	Delega	ted Pricing [‡]
Variable	Mean	Std. Dev.	Mean	Std. Dev.
Qrt Revenues	383,190	280,158	304,206	247,221
Qrt Gross Profits	$148,\!413$	138,771	$127,\!491$	107,790
Qrt Gross Profits-SC	$136,\!336$	136,323	$113,\!302$	105,765
Qrt Quantity	2,079	$1,\!818$	1,502	$1,\!654$
Qrt Price	227	281	249	127

Table 3.8.: Mean and std. dev. by pricing rights

Notes. All outcomes are quarterly, individual-level data. All outcomes but quantity are deflated using country-specific CPI with 2005 as base year. I dropped two observations with negative sales. Price missing for nil revenues. † 148 observations except price with 142. ‡ 106 observations with except price with 99.
	Centraliz	ed Pricing	Delegated Pricing	
Variable	Bonus No Mean†	Bonus Yes Mean‡	Bonus No Mean††	Bonus Yes Mean‡‡
Qrt Revenues	$302,\!277$	438,359	257,258	354,837
Qrt Gross Profits	142,791	$152,\!247$	$102,\!285$	$154,\!675$
Qrt Gross Profits-SC	$132,\!373$	139,039	$87,\!634$	$140,\!984$
Qrt Quantity	$1,\!631$	$2,\!384$	1,563	$1,\!437$
Qrt Price	275	196	204	295

Table 3.9.: Mean of outcomes by bonus and pricing rights

Notes. All outcomes are quarterly, individual-level data. All outcomes but quantity are deflated using country-specific CPI with 2005 as base year. I dropped two observations with negative sales. Price missing for nil revenues. † 60 observations with exception of price with 56. ‡ 88 observations with exception of price with 86. †† 55 observations with exception of price with 50. ‡‡ 51 observations with exception of price with 49.

Sample Type of Regression Dependent Variable	(1) Pooled OLS Qrt Revenues	(2) Pooled OLS Qrt Gross Profits	(3) Pooled OLS Qrt Gross Profits-SC	(4) Pooled OLS Qrt Quantity	(5) Pooled OLS Qrt Unit Price
Bonus	252,729**	71,689*	73,227*	1,871**	-78*
	(111, 561)	(39, 824)	(38,688)	(687)	(44)
PD	-18,543	-21,468	-25,071	360	-74
	(84, 403)	(39, 292)	(38,571)	(617)	(64)
Bonus*PD	-61,656	20,512	23,970	-1,081	162^{**}
	(119, 947)	(47,051)	(45, 400)	(765)	(65)
Year dummy	Yes	Yes	Yes	Yes	Yes
Quarter dummy	Yes	Yes	Yes	Yes	Yes
Employee dummy	No	No	No	No	No
Other controls	No	No	No	No	No
Clustering	Employee	Employee	Employee	Employee	Employee
Observations	254	254	254	254	241
	1.0 1.1		NRX 1 1 00 00 1		

Table 3.10.: The effect of incentives and price delegation

Notes. All outcomes are deflated using country-specific CPI with 2005 as base year. All specifica-

tions are clustered by employee to allow for serial correlation in individual error terms.

	(1)	(2)	(3)	(4)	(5)
Sample	Pooled	Pooled	Pooled	Pooled	Pooled
Type of Regression	OLS	OLS	OLS	OLS	OLS
Dependent Variable	Qrt Revenues	Qrt Gross Profits	Qrt Gross Profits-SC	Qrt Quantity	Qrt Unit Price
Revenues-based	-180,046	-70,267	-67,940	-307	-49
	(192, 597)	(88, 441)	(86, 998)	(1,402)	(48)
Year dummy	Yes	Yes	Yes	Yes	Yes
Quarter dummy	Yes	Yes	Yes	Yes	Yes
Employee dummy	No	Yes	No	Yes	No
Other controls	Yes	Yes	Yes	Yes	Yes
Clustering	Employee	Employee	Employee	Employee	Employee
Observations	111	111	111	111	111

Table 3.11.: Revenues-based vs mixed compensation

Notes. Other controls include the annual target on revenues. 56 observations refer to employees compensated by a mix pay, and the remaining 55 refer to employees compensated with revenuesbased pay. All outcomes are deflated using country-specific CPI with 2005 as base year. All specifications are clustered by employee to allow for serial correlation in individual error terms.

Chapter 4. Credit Channel and Non-Performing Loans: Evidence from the Italian Sovereign Debt Crisis

1

Abstract

We propose an empirical strategy to estimate the impact that a worsening in banks' wholesale funding opportunities has on firms' ability to repay their loans. We exploit the Italian sovereign debt crisis of July 2011 as a significant funding shock to Italian banks. In summer 2011 the yields on Italian bonds suddenly increased causing an immediate decline in the price of government securities. Italian banks suffered significant losses as they hold a large share of their assets in domestic government bonds. Moreover, as Italian bonds are used as collateral in repo markets, the decline in bonds prices restricted wholesale funding opportunities and it has likely led to a contraction in the credit supply. This paper investigates whether this severe shock to credit supply hampered borrowers' ability to repay their loans. To capture how each bank has been affected by the sudden increase on Italian yields, we use

¹This chapter was jointly co-authored with Marco Gallo, Giacinto Micucci and Francesco Palazzo at Bank of Italy. The views expressed in this work are those of the authors and do not necessarily reflect the opinions of the Bank of Italy.

the difference between each bank's lending and retail funding over its total assets. This measure captures each bank's dependence on the wholesale funding market and, thus, its exposure to the sovereign-debt shock. This credit crunch episode has likely involved all kinds of borrowers (firms, households, government) and it has contributed to an overall reduction in aggregate demand. To distinguish how much firms suffered from more restrictive credit conditions rather than from a significant reduction in firms' product demand, we consider firms that borrow from *stressed* banks (i.e., banks that were more vulnerable to a wholesale funding shock) to *otherwise-similar* firms that borrow from *non-stressed* banks and compare their loan-repayment performance. To compare firms that are otherwise-similar, we control for several firm-level characteristics and we create several geographic-product markets where firms are likely to face the *same* demand shocks. We find that, following the funding shock of July 2011, one standard-deviation increase in our measure of firms' exposure to banks' financial distress increases their probability not to honour their loans by about 0.4% (i.e., the size of the credit channel). Our results also suggest that the aggregate demand channel led to a 2.4% increase in the share of non-performing loans.

4.1. Introduction

In the last two decades the interdependence between financial markets and economic activity has been one of the most prolific research areas. Bernanke and Gertler [1989] and Kiyotaki and Moore [1997] initiated a vast theoretical literature on the credit channel. They stress the role of financial frictions between lenders and non-financial borrowers in amplifying business cycles. Holmstrom and Tirole [1997], Stein [1998], Diamond and Rajan [2001] and Freixas and Rochet [2008] extend the analysis to include financial friction between financial intermediaries and their providers of funds (wholesale interbank funding, retail depositors).

Although the empirical literature has extensively documented co-movements between credit aggregates and real economic variables, there are only a few empirical works that succeed in disentangling credit demand and supply shocks. The seminal work of Khwaja and Mian [2008] proposes an effective identification strategy that exploits asymmetric funding shocks to banks in order to measures credit supply decisions with respect to identical firms. Di Patti and Sette [2012], Gobbi and Sette [2013] and Bofondi et al. [2013] use a similar methodology to measure heterogeneous credit supply restrictions across banks operating in Italy. The first two consider the 2007-08 global credit crunch event and compare banks with different balance-sheet characteristics; the latter considers the Italian sovereign debt crisis of July 2011 and compares Italian and foreign banks.

In this work we propose an empirical strategy to identify the impact that a deterioration in banks' funding opportunities has on firms' ability to honour their credit contracts. Specifically, we study whether a funding shock to the wholesale interbank market translates into a *deterioration* of the borrowing status of enterprises in closer credit relationships with more liquidity constrained banks. In a nutshell, our goal is to identify and quantify how shocks to banks' wholesale liquidity provision may be transmitted to the real sector, leading to additional disruption in firms' economic activity. Previous works on the transmission of credit supply shock to the non-financial sector focus on overall credit supply. Our choice to study non-performing loans is partially motivated by the direct implications that non-performing loans have for banking regulation as their sharp increase poses a serious threat to financial stability. This "jump" to final outcomes makes the credit shock identification problematic. Khwaja and Mian [2008] identification strategy when the dependent variable is firms' borrowing status with respect to the whole banking sector cannot be directly applied because it is not possible to exploit variation of the dependent variable across different lenders. To measure how the credit shock affected the borrowing status we ideally want to consider firms that differ only in their lenders. Therefore, we devote extra care to select our sample and we introduce additional controls on firm observed and unobserved characteristics.²

We consider the Italian sovereign debt crisis of July 2011 because it is an ideal case-study of unanticipated shock to the wholesale funding opportunities of Italian banks. In July 2011, Italian government bonds experienced a drastic price drop because of a very volatile political and economic environment (see Figure 4.1). Importantly, Italian government bonds were extensively used by domestic banks as collateral in secured repo transactions in the wholesale funding market. The decline in Italian bond prices and the contemporaneous increase in haircuts and margin requirements severely affected banks costs to raise money on international capital markets. This shock was particularly severe because collateralized borrowing represented at the time of the shock the main wholesale venue to raise additional funds. In fact, other venues were *de facto* shut down for Italian banks: unsecured short-term liquidity had not recovered from the 2007-08 interbank market freeze, while longer term financing on capital markets was prohibitively expensive. Although an increase in retail deposits might have partially provided an alternative, the contemporaneous bank competition for deposits makes it an expensive as well as slow alternative to wholesale funding. We think it is a reasonable approximation to restrict attention to the banks' different ability to access secured wholesale funding. A priori we expect that the sovereign debt crisis is likely to have hit harder financial institutions that heavily relied on wholesale funding. We therefore focus our identification strategy on the banks' reliance on wholesale funding which we measure with the difference between each bank's lending and retail funding over its total assets and call it the funding gap.

The focus of our analysis is to empirically identify how the shock is transmitted to the real economy. We encounter two main issues to empirically identify how the shock is transmitted to the real economy. Firstly, we need to construct a measure of exposure to the shock at the firm level. To this purpose, we compute a weighted average of the lenders' funding gap, weighting each lender with the credit granted by the bank to the firm over the firm's total credit granted in a given period. Secondly, it is difficult to disentangle the relative contributions of weak product demand or heightened credit constraints. Periods of aggregate credit contraction are positively correlated with economic downturns. In turn, more expen-

²Although our identification strategy involves a higher risk of model misspecification, we believe that the final trade-off is worth taking this risk.

sive financing costs and weaker product demand increase firms' financial vulnerability. We propose an econometric strategy to distinguish the *product market* channel from the *direct* transmission of banks funding shocks to the borrowing firms (i.e. the *credit channel*). In order to measure how much the wholesale funding shock affected borrowers' ability to repay their loans, we compare firms that predominantly borrow from stressed banks (i.e. banks more exposed to a wholesale funding shock) to **otherwise-similar** firms that borrow from non-stressed banks. To identify whether firms are otherwise-similar we control as much as we can on the product demand they face. To this purpose, we consider several geographicproduct markets where firms are exposed to similar demand shocks and compare their loan-repayment performance within the same market. Thanks to this identification strategy, we try to disentangle the direct credit supply channel from the product demand channel. We present a simple example to better explain our identification strategy. In an ideal empirical setting, we would compare the performance, before and after July 2011, of two otherwiseidentical firms that differ only for the bank they borrow the money from. Consider two FIAT car dealers, dealer A and dealer B. They operate in the town of Olbia, in the north-east coast of Sardinia. Assume both dealers have similar total assets and non-financial cost structures, and that before July 2011 the two dealers share equally the same consumer market, that is, the market of people that would purchase new FIAT cars in Olbia. This implies that dealer A and dealer B are both subject to the same macro-demand shocks. The two dealers only differ for the banks they borrow the money from. Assume dealer A borrows money only from bank C while dealer B only borrows the money from bank D. For simplicity, we further assume that bank C and D only differ along one dimension. Bank C can only fund itself on the wholesale market through collateralized repo transactions with other private banks or the central bank. Bank D, on the contrary, is able to additionally fund all it loans with a combination of deposits and equity. If capital markets are efficient before the crisis, the sources of funding (equity and deposits or wholesale funding) should be equivalent. As a further simplification, we assume that before July 2011 both banks charged an interest rate of 4% for a credit line of \in 500,000 (granted); after July 2011 bank D charges 5% while bank C charges 8% and unilaterally reduces its granted credit line to $\in 200,000$. The two dealers struggle more to repay their pre-existing debts because the shock on July 2011 impacted the cost of borrowing and the demand of cars in Olbia. However, we expect dealer B to struggle more because of the marked increased in bank C's interest rate and the restricted access to

credit.

We implement this strategy using ordinary least squares on semi-annual data at the firm level. The dependent variable is a binary variable that equals 1 if, in a given period, the firm is recorded as *non-performing* in Italy's Credit Register and equals 0 if the firm is recorded as *performing*. As a treatment variable for assessing the impact of the funding shock on loans repayment we use a measure for the firms' exposure to lenders' financial distress (henceforth referred as *firms' exposure*). Moreover, we also use the same measure of the firms' exposure interacted with a binary variable that equals 0 before July 2011 and equals 1 afterwards. This should capture the impact of the Italian sovereign debt crisis. To disentangle the credit channel from the aggregate-demand channel, we control for market-time fixed effects in order to focus our analysis on the residual variation left each period within the same geographicproduct market.

This econometric specification suffers from three main issues. Firstly, there could be feedback effects in our treatment variable in response to the funding shock, as banks can adjust their funding practices. To control for this possibility, we instrument our treatment variable by its equivalent in June 2010.

The second issue concerns the empirical identification of the channels underlying the funding shock to the banks. As mentioned above, the sovereign debt crisis is likely to have hit harder financial institutions that heavily relied on wholesale funding. We measure firms' exposure to the funding shock by using the weighted-average of the funding gap of the firms' lenders. Banks' funding gap captures their reliance on wholesale markets whose functioning was severely affected during the sovereign debt crisis. Moreover, our data show substantial variation in this variable across Italian banks so it is likely to represent a valid treatment variable for our purposes. However, by limiting the treatment of the shock to the funding gap only we neglect other potential transmission channels.³ To control for other dimensions of banks' heterogeneity, we use other observable characteristics of banks, such as Tier 1 capital ratio, ROA and total assets.

³The sovereign debt crisis is likely to have hit harder financial institutions that non only heavily relied on wholesale funding but that also (i) held a considerable proportion of eligible collateral in Italian government bonds, and (ii) did not have a comfortable buffer of eligible securities to post as additional collateral to meet higher haircuts and margin calls. Ideally, we would measure the extent to which each Italian bank suffered from this funding shock using an index that considers all three dimensions of dependence on wholesale funding, exposure to Italian government bonds and spare capacity to post collateral. However, the lack of variation in bond holdings across Italian banks prevents us to fully exploit these three dimensions.

Finally, our identification strategy is valid only if firms that operate in the same market were trending similarly before the crisis. To account for this possibility, we include the most significant firms' characteristics that we observe from balance sheet data such as total assets, ROA, EBIT and leverage.

Holding all included variables constant, the results suggest that one standard-deviation increase (around 8%) in our measure of firms' exposure to financial distress marginally increases firms' probability to be non-performing by about 0.4%. This is our estimate for the size of the credit channel. Interestingly, the aggregate demand channel lead to a 2.4% increase in the share of non-performing loans.

The rest of the chapter is organized as follows. Section 4.2 provides a summary of the related literature. Section 4.3 explains our data and section 4.4 provides summary statistics. Section 4.5 explains the identification strategy. Section 4.6 presents the econometric model and the results, followed by robustness checks in section 4.7. Finally we make some concluding remarks.

4.2. Literature Review

Non-performing loans (NPLs) have recently gained substantial attention both in policy and academic circles. During the last five years the banking sector has been greatly exposed to the financial and economic turmoil. Indeed, banks suffered significant financial stress because asset prices substantially declined and an increasing number of outstanding loans defaulted. Although the financial crisis provides new reasons to study NPLs, in the last decade the literature has already focused on several issues that were central to policymakers' discussions. For the sake of exposition and without any pretence of strict methodological classification, we prefer to divide the literature in four main strands. Although there are some inevitable overlaps, each strand has different research questions, and uses different datasets as well as different econometric techniques. Moreover, we only consider the more recent works as they relate more closely to our research questions.

4.2.1. NPLs and firm-level characteristics

This strand of literature aims to estimate the probability to default (PD) of small and medium sized enterprises (SMEs) as a function of some observable firm characteristics. The famous

Z-score in Altman [1968] initiates the conventional use of financial ratios as a predictor for bankruptcy.

Behr et al. [2004] estimate PDs for German SMEs using a logit scoring model on a loan level dataset provided by a SME financier. They find that indebtedness profitability, liquidity and regional location are all significant factors where estimates have all intuitive signs. Recently, Fidmuc and Hainz [2010] obtain similar results using a probit model on loan level data from a Slovakian commercial bank for the period 2000-2005. While the former paper highlights how relationship lending may lead to monopoly rents for banks, the latter stresses how Slovakian default rates converged to the levels found in more developed economies.

Dietsch and Petey [2004] and Jacobson et al. [2005] use different PD's models to test the validity of some implicit assumptions in the Basel II framework. The former work uses default data for German and French SMEs from two European financial information providers to estimate a one-factor credit risk model and assess default correlations. The latter adopts a non-parametric Monte-Carlo technique to illustrate how SMEs are significantly riskier than other corporate credit.

Recently, McCann and McIndoe-Calder [2012] build a cross-sectional dataset of loan-level exposures and borrowers' balance sheet characteristics for approximately 6,000 Irish SMEs. They run several probit regressions controlling for financial and economic characteristics (leverage, liquidity, profitability, exposure, firm size, economic sector) and they get results qualitatively consistent with the previous literature.

4.2.2. NPLs and bank characteristics

Another strand of literature looks at NPLs differences across banks, rather than firm-level (borrower) heterogeneity. The main focus is on the internal organization, collateral requirements, and sector specialization of banks. Although macroeconomic variables are included in the analysis, they are used as exogenous covariates with respect to NPLs. In Section 4.2.3 we discuss why this partial equilibrium framework has been criticized in several works that prefer to model NPLs dynamics through a Vector Autoregressive (VAR) approach.

Berger and Young [1997] use Granger-causality techniques to consider the relationships among loan quality, cost efficiency and bank capital for US commercial banks in the 1985-1994 period. As they consider aggregate variables, they do not control for individual loan characteristics or other bank related specificities. A more comprehensive analysis has been carried out by Salas and Saurina [2002] thanks to a loan-level panel data of Spanish banks in the 1985-1997 period. They find that macroeconomic variables have strong explanatory power for NPLs but they also highlight the importance of some microeconomic factors such as the banks' rapid credit and branch expansion, portfolio composition, size, net interest margin, capital ratio and market power. In particular, they stress the difference between commercial and savings banks.

Jimenez and Saurina [2004] use Spanish Credit Register data over the period 1988-2000 to assess how borrowers' PD depend on collateral used, the type of lender and previous bankborrower relationships. They report that saving banks have riskier loan portfolios, that closer bank-borrower relationship seem to increase lender's willingness to take more risk and that collateralized loans have a higher default probability. Jimenez et al. [2006] conduct a more thorough study of the effects of collateral on loan PD. They analyse a sample of Spanish bank loans from 1984 to 2002 and they claim that collateral usage is higher among riskier borrowers, i.e. those who previously defaulted or new borrowers who default after receipt of the loan. As a result, the signalling theory of collateral posting seems to have some validity only for younger cohorts of borrowers.

Quagliariello [2007] builds a large panel dataset of Italian financial intermediaries to study the cyclical behaviour of NPLs over the business cycle. His analysis confirms that macroeconomic conditions have a significant and long-standing impact on NPLs. In the same spirit, Louzis et al. [2012] consider a panel dataset of loans originated in the Greek banking sector from 2003 to 2009. The authors estimate NPLs responsiveness to macroeconomic variables for consumer loans, business loans and mortgages (the latter being the less responsive) and they conclude that NPLs are mainly explained by macroeconomic factors and management quality.

Beyond macroeconomic factors, Jimenez and Saurina [2006] stress how expansionary credit policies have a statistically significant effect on future NPLs. Indeed, they report a strong positive relationship between rapid credit growth and subsequent NPLs. Finally, Jimenez et al. [2007] focus on the interplay between banking competition and NPLs. They provide some empirical evidence in support of the "charter value" hypothesis (see Keeley [1990], Repullo [2004]), i.e. a positive relationship between market competition in banking and their credit risk profile.

4.2.3. NPLs and macro shocks

All works in Section 4.2.2 assume macroeconomic variables to be exogenous with respect to bank losses. However, an extensive literature in macroeconomics highlights the importance of the credit channel as a powerful amplifying mechanism (see Bernanke and Gertler [1989], Kiyotaki and Moore [1997]). In a nutshell, as NPLs directly erode banking capital, credit supply may significantly decline and have a negative effect on employment and output.

A first strand of literature estimates VAR models to capture the simultaneous determination of NPLs and macroeconomic variables. Data include aggregate variables, with no reference to microeconomic characteristics at the loan level for borrowers and lenders.

Gambera [2000] use US quarterly data from 1987 to 1999 to estimate a bivariate VAR and assess the impact of regional and national macroeconomic variables on different types of loans (agricultural, commercial, industrial and residential). He reports how a limited number of macroeconomic variables well predict NPLs behaviour. Hoggarth et al. [2005] consider the interaction between banks' write-offs to loan ratio and several macroeconomic variables over the period 1988-2004. They found a negative relationship between changes in output and the write-offs ratio, while the feedback effect from NPLs to output is much weaker. Filosa [2007] and Marcucci and Quagliariello [2008] study credit cycles in the Italian economy. The former finds that the behaviour of NPLs is weakly pro-cyclical and it explains only a modest fraction of the historical variability in bad loans dynamics. Moreover, he finds statistically significant feedback effects from bad loans to real economic activity. The latter work highlights the importance of macroeconomic shocks on the banking sector while they find a somewhat weaker feedback effect.

Marcucci and Quagliariello [2009] use a threshold regression methodology on a panel dataset for Italian bank borrowers' default rates to detect possible asymmetries in the relationship between credit risk and macroeconomic variables. They report that negative effects of the business cycle on credit risk are more pronounced during downturns and for riskier banks. Another strand of literature uses single-equation time series approach to assess the macroe-

conomic determinants of NPLs. Arpa et al. [2001] focus on a sample of Austrian banks over the 1990-199 period and they find that output growth, interest rates and real estate price dynamics have a good explanatory power. Recently Bofondi and Ropele [2011] use a similar econometric technique to study how different macroeconomic shocks impact households and firms. In particular, they find significant differences between the two groups in terms of endogenous persistence and time-response to shocks. Finally Caporale et al. [2013] adopt a Structural VAR approach to study whether "excessive" loans during expansionary phases can explain the more than proportional increase in NPLs during recessions.

4.2.4. Panel VAR models

Panel Vector Auto Regressive (PVAR) models⁴ aim to combine the salient features of the strands of literature discussed in Section 4.2.2 and 4.2.3. Indeed, this approach allows to consider the endogenous relationship between macroeconomic variables and NPLs while, at the same time, it allows to control for some idiosyncratic characteristics. As a result, it is possible to assess the relative contribution of common and idiosyncratic factors to NPLs' determination.

Espinoza and Prasad [2010] estimates a PVAR model on a dataset for Gulf Cooperation Council countries over the period 1995 to 2008. He finds evidence of a strong but shortlived feedback effect of NPLs on non-oil growth. Similarly, Nkusu [2011] considers a panel dataset of developed countries over the period 1998 to 2009 and he reports strong linkages between macroeconomic and financial variables. Finally, Inessa and Rima [2013] control for bank specific characteristics in a dataset of Egyptian banks for the period 1993-2010. They stress the importance of the credit channel in the transmission of macroeconomic shocks to the banking sector. Moreover, they point out that capital inflows have significant effects on banks' loan portfolio quality.

4.3. Data and sample selection

4.3.1. The database

Our analysis uses three main pieces of information: data on individual firms, data on bankfirm relationships, and data on banks' balance sheets. We collect this data from January 2006 to June 2012.

As for data on individual firms, we use the Company Accounts Data Service (henceforth CERVED), a commercial dataset which virtually includes all Italian corporations. Our sample only takes into account non-financial companies where information on financial liabilities

 $^{{}^{4}}$ A recent introduction to this methodology is Canova and Ciccarelli [2013].

such as financial leverage is available.⁵ The CERVED dataset provides detailed information on firms' geographic location and sector of economic activity as well as annual accounting information on revenues, total assets, liquidity and financial liabilities. In the CERVED dataset there are about 280,000 firms over the time period considered.

Data on bank-firm relationships for our sample of CERVED corporations comes from the Italian Credit Register (henceforth CR). This database lists all performing loans above \in 30,000 and the universe of non-performing loans from banks operating in Italy, including branches and subsidiaries of foreign banks. This dataset is updated monthly and every financial intermediary reports granted and drawn amounts for each single borrower. Moreover, it assigns to each borrower one out of five possible statuses: *performing, past-due*,⁶ *restructured*,⁷ *sub-standard*,⁸ and *bad debt*.⁹ The *global* position of a borrower with respect to the whole Italian banking system follows the Basel 2 classification:

- Bad debt: 10% or more out of total cash credit used is reported as bad debt.
- *Sub-standard:* 20% or more out of total cash credit used is reported as sub-standard or bad debt.
- *Restructured:* 20% or more out of total cash credit used is reported as restructured, sub-standard or bad debt.
- *Past-due:* 50% or more out of total cash credit used is reported as past-due, restructured, sub-standard or bad debt.
- *Performing:* None of the above condition applies.

At this exploratory stage of the work, we simplify the analysis by studying the variation of the binary outcome variable NPL_{it} which equals 0 if the global position of borrower *i* at time *t* with respect to the whole Italian banking system is classified as *performing* and equals 1 otherwise (whence the acronym NPL for non-performing loan). For every firm in the CERVED sample we consider its global borrowing status in each semester (ending in June and in December).

⁵It is a necessary requirement to meaningfully analyze non-performing loans. The vast majority of corporations provide this information.

 $^{^6\}mathrm{Debt}$ repayment at least 90 days late.

 $^{^7\}mathrm{Renegotiated}$ contractual terms.

⁸Temporary problem to reimburse debt that might be solved within an "appropriate" time lapse.

⁹Insolvency status or situations *de facto* comparable.

Lastly, we collect balance-sheet information on Italian banks from the Supervisory Reports submitted by financial intermediaries to the Bank of Italy. If a financial intermediary is part of a banking group in a given period we only consider aggregated data referred to the whole group because the parent bank controls funding decision for all subsidiaries. In the remainder of the work we generically use the term "bank" to refer to an economic unit that controls its own funding and lending decisions, irrespective of its specific legal status. There are about 1,050 credit institutions in our sample. Some of them are foreign banks such as BNP-Paribas or Deutsche Bank. Branches and subsidiaries of foreign banks are treated separately because of their peculiar status. To measure how the sovereign debt crisis affected each Italian bank, we construct several standard accounting ratios at the bank level such as ROA, ROE, Tier 1 and Tier 2 Capital ratios¹⁰ and we introduce a variable more closely related to the specific "distress" episode caused by the Italian sovereign debt crisis:

• *Funding gap:* It is the ratio to total assets of the difference between lending and retail funding. Data is available monthly for the whole sample period.¹¹

To have semi-annual observations all monthly data is aggregated by semester with a simple average.

4.3.2. Sample selection

Our dataset does not provide enough information to directly observe firms' product demand. Therefore, our identification strategy relies on the ability to construct groups of non-financial corporations that face similar demand shocks. Using information on the geographic location and sector of economic activity, we control for unobserved market-demand factors in those sectors where firms predominantly serve a *local* geographic market and sell sufficiently *homogeneous* products. As a consequence, we restrict attention to only those firms that operate in local markets. Examples of such markets can be found in constructions and services to

¹⁰For Tier 1 and Tier 2 Capital ratios we have information from June 2008 to June 2013 at semi-annual frequency.

¹¹It follows the definition in the Bank of Italy Financial stability Report (November 2012). Lending includes bank loans net of provisions and securitized loans repurchased in the form of liabilities issued by the securitization vehicle. Retail funding includes residents' deposits and other forms of retail fund-raising (such as bank bonds subscribed by households) but it excludes liabilities related to securitization. Repos with central counter-parties are excluded both from lending and funding.

mention a few. Bresnahan and Reiss [1988] and Bresnahan and Reiss [1991] use a similar approach to identify the "relevant market" for the sake of competition law issues.¹²

In the CERVED dataset there are about 280,000 firms over the time period considered, which represent about 2.9 million firm - half-year observations. We drop firms who work in the agricultural and manufacturing sectors as these are traditionally export-intensive sectors in Italy and we have information on exports for none of the firms present in the sample. This leaves only firms that work in the construction and services sectors which are traditionally (more) local businesses. These firms represent roughly 70% of the sample and constitute about 1.9 million firm - half-year observations.

We also drop all firms that in a given half-year borrowed a positive amount from a foreign bank as we have detailed information on the balance sheet of the Italian subsidiary but not equivalent information for the foreign group as a whole. We also drop all firms that in a given half-year borrowed a positive amount from a credit institution for which we have no information on the funding gap (about 925 small financial institutions such as leasing or factoring companies). Our sample only includes firms that borrow from the 92 Italian banks for which we have complete data.¹³ Further, we drop those firms for which total assets is smaller than 10,000 euros as we believe the accounting data reported by these firms is not fully reliable. There are about 200 such observations (less than 1% of the sample). This leaves us with about 1.2 million observations.

Finally, our identification strategy requires to comparing firms that in a given period operate in the same market. Given the information available in CERVED, we define a market as a combination of a geographic market and a product market. The geographic market is defined as a province, such as Milan or Rome. There are 103 provinces considered in the sample (see Appendix C.1). The product market is defined using the ATECO-2007 Classification of Economic Activity provided by Italy's National Institute of Statistics (ISTAT). ATECO-2007 classifies each economic activity as a 6-digit code. We use the full 6-digit code to define a product market as this represents the more homogeneous unit of geographicproduct market we are able to identify given the information available in CERVED. For example, the 6-digit code for the services sector can tell whether a firm operates in the retail

¹²For example, Bresnahan and Reiss [1988] consider thirteen retail and professional service industries: farm equipment dealers, movie theatres, new or used tire dealers, beauty shops, barbers, plumbers, electricians, new auto dealers, physicians, veterinarians, dentists, drugstores, and optometrists and opticians.

¹³The 92 Italian banks we consider represent a big share of the Italian market covering roughly 77% of the amount of credit granted annually to the firms in CERVED between 2005 and 2012.

sale of beverages or in the retail sale of packaged fruit and vegetables. There are about 800 product markets in our sample which in combination with the 103 provinces form about 25,000 geographic-product market pairs.¹⁴ For our identification strategy to work, we only keep those markets that in a given half year have at least two active firms present in the sample. In this way, we are able to compare firms that share the same (or potentially the same) demand shocks. This leaves us with about 130,000 firms and 1 million firm - half-year observations. Appendix C.2 shows how the selected sample differs from the original CERVED sample. Median firms do not differ significantly across the two samples. However, our selection removes very large firms as one can see by comparing the maximum statistics of the firm-level characteristics. This is mainly due to the fact that we remove firms that borrow from foreign banks as well as very large manufacturing firms which represent the bulk of the top-exporters of the "made-in-Italy".

4.4. Descriptive statistics

Our identification strategy aims to assess whether the sovereign debt crisis contributed to worsen firms' borrowing status. To achieve this goal we consider four main factors:

- Distribution of the share of firms recorded as non-performing (i.e., NPL) across geographic areas, sectors and over time
- Individual bank exposure to the sovereign debt crisis and firm-level credit linkages to the banking sector
- Product market demand
- Firm financial and economic characteristics

The first factor is the outcome we aim to explain. The second factor describes the extent to which the sovereign shock affected banks' wholesale funding opportunities and, in turn, how this shock got transmitted to firms as a function of their credit links to the banking system. Together, they determine the *direct* credit supply shock to the firms that we hope to identify. The last two factors play a fundamental role to determine borrowers' ability to

 $^{^{14}\}mathrm{However},$ not all product markets are present in every province.

repay their loans but they are not directly related to the wholesale funding shock. Therefore, we need to control for these last two dimensions if we hope to identify the shock.

In this section, we present some descriptive statistics.

4.4.1. Distribution of non-performing loans (NPLs)

The outcome variable we study in this chapter is the binary variable NPL_{it} which equals 1 if firm *i* is recorded as non-performing in Italy's Credit Register in period *t* and equals 0 otherwise.

Table 4.1 reports the distribution of the share of non-performing loans in our selected sample of firms by geography and by sector. The share of non-performing loans is markedly larger in Central and Southern Italy and for the construction sector.

Table 4.2 reports the semi-annual time series of the share of non-performing loans in our selected sample. There is a positive trend for the whole series, with some marked jumps starting from the second half of 2008 when Lehman Brothers collapsed. In the second half of 2011 the share of non-performing loans jumped by 1.1% relative the previous half year. At the end of 2012 the share of non-performing loans in our selected sample reached 7%.

On the top of Table 4.3 we report the share of non-performing loans in our selected sample of firms for three time periods: the period before the collapse of Lehman Brothers, the period between the collapse of Lehman and the Italian sovereign debt crisis, and the period post Italian crisis. The share of non-performing loans increased markedly after the collapse of Lehman Brothers and it increased again after the Italian crisis. On the bottom of Table 4.3 we report the same statistics for the periods before and after the Italian sovereign debt crisis, as the focus of this chapter is on the Italian sovereign debt crisis. The share of non-performing loans almost doubled following the Italian crisis compared to the previous period.¹⁵

Table 4.4 completes the analysis of the distribution of non-performing loans before and after the Italian crisis by filtering the series with respect to geographic areas and sectors. It shows that the impact of the Italian crisis, as measured by the change of the share of non-performing loans after the crisis, is larger for Central and Southern Italy as well as for the construction sector.

¹⁵At this preliminary stage, we do not model the dynamics of the share of non-performing loans. However, one can imagine that there is strong persistence in the series: firms that become non-performing stay in this state for a long time over the sample. In our estimation we allow arbitrary correlation between disturbance terms by clustering at the market level.

4.4.2. Banks' financial distress and firms' exposure

The literature uses different variables to capture banks' vulnerability to different types of shocks. Typical measures of banks' distress include bank solvency and liquidity ratios as well as profitability and credit quality indicators (see Jimenez et al. [2012]).¹⁶ The sovereign debt crisis is likely to have hit harder financial institutions that heavily relied on wholesale funding. We therefore focus our identification strategy on the banks' reliance on wholesale funding which we measure with the difference between each bank's lending and retail funding over its total assets and call it the funding gap.

Table 4.5 reports a summary of balance-sheet characteristics for the 92 Italian banks for which we have compete information, using bank - half-year as unit of observation. At the bottom of Table 4.5 one can see summary statistics for the funding gap and the share of risky bonds both expressed as a percentage. The standard deviation of the funding gap is more than eight-times its mean, while that for the share of risky bonds is about one-eighth of the mean. This provides a rough indication of the difference in the cross-bank variation in the two measures of distress.

Table 4.6 compares two groups of banks before and after the Italian crisis. The first group of banks are those that as of June 2010 had a negative funding gap. A negative gap can be interpreted as banks with stable sources of funding as they do not crucially rely on wholesale funding. The second group of banks are those that as of June 2010 had a positive funding gap. This separation is only indicative of different banks' exposure to the wholesale funding shock. In the top part of Table 4.6 we report the change in the ratio between the total value of non-performing loans over the bank's total assets. The share of non-performing loans increased more (about 2.2% increase) for the group of banks that in June 2010 had a positive funding gap compared to the other group (which experienced an increase of about 1.2%). At the bottom of Table 4.6 we report the change in the difference between a bank's total loans (in euros) in the current half year and the previous one over the bank's total assets. The change in total loans has become negative for the group of banks that in June 2010 had a positive funding gap compared to the other group where this change has only decreased but remained positive. Although the bank's amount of total loans is determined by demand and supply, these numbers provide an overview of the dynamics going on during the Italian crisis.

¹⁶For our same purposes, Bofondi et al. [2013] use a dummy for foreign banks.

Let $DISTRESS_{j,t}$ be the index of banks' financial distress (i.e., the funding gap).

 $DISTRESS_{j,t}$ summarizes the funding vulnerability of bank j at time t. In our baseline model we use data where the unit of observations is a firm-time pair. As the firms in our data normally borrow from more than one lender, we aggregate our index $DISTRESS_{j,t}$ at the firm level. For each non-financial corporation i and each period t, we construct $EXP_{i,t}$ as the weighted average of the index $DISTRESS_{j,t}$ for those banks with credit relationship with firm i in period t. Bank j's weight in $EXP_{i,t}$ depends on the share of credit granted by bank j to firm i out of the total credit granted by the banking sector to firm i at time t. Table 4.7 reports summary statistics for the firm's exposure to its lenders funding gap (i.e., $EXP_{i,t}$). The first variable in the Table (Exposure to Funding gap) is a time-dependent variable, while the second variable (Exposure to Funding gap at June 2010) is the corresponding value of the firm's exposure in June 2010. This last variable is used as an instrumental variable to control for possible feedbacks effects in the treatment variable $EXP_{i,t}$ in response to the funding shock of July 2011. Figure 4.2 provides the histogram of firms' exposure to their lenders' distress which looks about normally distributed with mean 6.5% and standard deviation 8.4%, i.e. in our sample firms were exposed on average to banks with a funding gap of 6.5%of total assets.

4.4.3. Product market demand

In the selected sample of about 130,000 firms and about 1 million semi-annual firm-level observations, there are 103 provinces and about 600 product markets at 6 digits of the ATECO code. Over the entire sample, these make about 140,000 market-time combinations that we use as fixed effects. In each half year, the number of firms operating in a given geographic-product market ranges from a minimum of 2 to a maximum of about 2,000, with a median number of 14, an average number of 83 and a standard deviation of about 220 firms.

4.4.4. Firm-level characteristics

Table 4.8 reports the main firm characteristics used in our analysis. The median firm in our sample has annual assets of about 1.2 million euros, a return on assets of about 4%, and a leverage ratio, measured by the annual share of financial debt over the sum of financial

debt plus equity, of about 74%. ¹⁷ The median firms in our selected sample borrow from two different banks. We also report HHI, the Herfindahl-Hirschman Index of a firm's loans portfolio across its lenders in a given half year. HHI can take values between 10,000/No banks when the firm borrows an equal amount from each of its lenders, and 10,000 when the firm only borrows from one bank. The median value of HHI in our data is about 5,700 which is very close to 5,000 (= 10,000/2). This suggests that the median firms in our sample distribute their borrowed funds evenly across lenders.

4.5. Identification strategy

The goal of this work is to assess whether the Italian crisis significantly aggravated the performance of borrowers linked to banks that suffered a more severe restriction on their funding opportunities. To identify this question empirically, one needs to consider the following issues.

Firstly, we need to construct a measure of the firms' exposure to their lenders' financial distress, as measured by the funding gap, and map this into a *deterioration* of the borrowing status of the same firms in credit relationships with more liquidity constrained banks (i.e. our treatment variable). For each firm i we compute a weighted average of its lenders' funding gap in every half year t, $EXP_{i,t}$.¹⁸ In the analysis, we lag the treatment $EXP_{i,t}$ one period (i.e., EXP_{it-1}) because we believe that bank i's credit supply as well as loan-portfolio management decisions at time t are based on the information available to the bank at the end of period t - 1, which also includes the bank's reliance on the wholesale funding market in t - 1 (i.e., its funding gap). Thus, EXP_{it-1} is the relevant metrics for approximating firm i's exposure to the credit-supply shock in period t and, thus, for assessing its impact on the performance of borrower i.¹⁹ In addition to the firm's exposure, EXP_{it-1} , our treatment also includes the interaction between EXP_{it-1} and the binary variable $Crisis_t$ which equals 0 before July 2011 and equals 1 afterwards. The interaction is intended to capture the regime change in response to the Italian crisis. We consider the Italian sovereign debt crisis of July

¹⁷The negative values of the leverage ratio are puzzling as well as values of the leverage that are greater than 100%. These cases represent about 7% of the observations.

¹⁸The weight for each lender b in period t is computed as the share of the credit granted by lender b to firm i over firm i's total credit granted in period t.

¹⁹In general, longer lags of EXP_{it-1} might be relevant to explain banks' decisions but at this exploratory stage of the work we only focus on the first lag.

2011 because it is an ideal case-study of unanticipated shock to the wholesale funding opportunities of Italian banks. The rapid increase in the interest rate spread between Italian and German bonds pushed the market value of Italian government bonds down (see Figure 4.1). In turn, this event worsened banks' conditions to obtain liquidity on the interbank market or through central bank refinancing operations because it led to additional margin calls and higher haircuts. Therefore, to get the same cash amount banks had to post a higher nominal amount of Italian bonds. This shock was particularly severe because collateralized borrowing represented at the time of the shock the main wholesale venue to raise additional funds. In fact, other venues were *de facto* shut down for Italian banks: unsecured short-term liquidity had not recovered from the 2007-08 interbank market freeze, while longer term financing on capital markets was prohibitively expensive.

Secondly, there is a fundamental economic issue if one intends to disentangle the impact of the shock on the firms' access to credit from the impact of the shock on the firms' product demand. Periods of aggregate credit contraction are positively correlated with economic downturns. During these periods, more expensive financing costs and weaker product demand increase firms' financial vulnerability. Therefore, it is difficult to disentangle the relative contributions of weak product demand or heightened credit constraints. We distinguish the *product market channel* from the *direct* transmission of funding shocks on banks to the borrowing firms (i.e., the *credit channel*) by comparing firms that predominantly borrow from *stressed* banks (i.e. banks more exposed to a wholesale funding shock) to **otherwise-similar** firms that borrow from *non-stressed* banks. To identify whether firms are otherwise-similar we control as much as we can on the product demand they face. To this purpose, we consider several geographic-product markets where firms are exposed to **similar** demand shocks and compare their loan-repayment performance.

This identification strategy works if three conditions hold. Firstly, the banks should not adjust their reliance on the wholesale market in response to the shock of the Italian crisis. In other words, there should be no feedback effect in our treatment variable EXP_{it-1} in response to the funding shock of July 2011. This assumption is obviously too restrictive as non-performing loans directly erode banking capital and in turn might induce banks to reduce credit supply and adjust their funding strategies (see section 4.2.3).²⁰ We control for this possibility by instrumenting our treatment EXP_{it-1} by its equivalent computed on

 $^{^{20}}$ Hoggarth et al. [2005] document that feedbacks affects from non-performing loans to output is much weaker than the reverse effect.

June 2010, $EXP_{i,6-2010}$.²¹ For each firm *i*, $EXP_{i,6-2010}$ is a time-invariant variable which summarizes firm *i*'s exposure to its lenders distress at June 2010. We choose June 2010 because, on the one side, it represents a moment in time that is close enough to the sovereign debt crisis shock of July 2011 to provide a good picture of the firm's exposure before the crisis and, on the other hand, it is far enough from July 2011 to believe that on June 2010 the shock of July 2011 was unexpected by banks and companies. The intuition behind this instrumental variable is the following: had the crisis not occurred, two similar firms would have been trending similarly.

The second issue concerns the empirical identification of the channels underlying the funding shocks to the banks. As mentioned previously, the sovereign debt crisis started in the wholesale money market and is likely to have hit harder financial institutions that heavily relied on wholesale funding. We therefore measure the firms' exposure to the funding shock in the wholesale money market of July 2011 by using the weighted-average of the funding gap of the firms' lenders, $EXP_{i,t}$, as we believe that the funding gap is a necessary channel of transmission of the shock from the wholesale market to the individual banks. By limiting the treatment of the shock to the funding gap, we neglect other channels through which the shock might have hit the Italian banks. This leaves plenty of cross-banks heterogeneity out of our analysis. In our dataset, we observe many time-variant characteristics of the banks such as total assets, return on assets (ROA), capital requirements such as Tier 1 and Tier 2 capital ratios, the amount of non-performing loans as a share of total assets, as well as the share of the "risky" bonds over the total bonds used by the banks as collateral in the wholesale money market.²² Therefore, we use the firms' exposure on June 2010 to these other dimensions of banks' heterogeneity as a simple way to control for other possible trans-

²¹We instrument the interaction $EXP_{it-1} * Crisis_t$ with $EXP_{i,6-2010} * Crisis_t$.

²²Securities holdings by country of counter-party and pledge status is a crucial dimension for assessing the impact of the Italian sovereign debt crisis on the amount of credit supplied by each individual bank. However, this measure does not exhibit the cross-sectional variation that would be necessary to identify its contribution to the credit supply shock in our data. For example, the median share of bank's holding in "risky bonds" over the total bonds is roughly 95%. Also 90% of the times the share of risky securities exceeds 80% (see Appendix C.3). The national focus of our data could explain the lack of cross-sectional variation in bonds holdings across banks. In principle, we could explain the lack of cross-sectional variation between Italian and foreign groups but the balance-sheet information we observe for foreign banks operating in Italy is incomplete as it only refers to the balance sheet of the Italian subsidiary of the foreign group. As documented in Battistini et al. [2013], banks tend to hold a large share of national securities relative to securities of foreign governments. To have the sufficient cross-sectional variation in banks' securities holdings, one would ideally have cross-country data where banks of different nationalities are likely to have different holdings of sovereign securities.

mission channels of the shock. Although very rudimental, we believe this approach provides valid insights to understand how the liquidity shock can propagate into the economy using detailed micro-economic data.

Finally, the identification strategy summarized by the ordinary least-squares specification with time-dependent treatment EXP_{it-1} and market-time fixed effects is valid only if firms that operate in the same market were trending similarly before the crisis. If the firms did not show similar patterns before the Italian crisis, then our identification strategy might not work. For example, if firms that borrowed from banks with heavy reliance on the wholesale funding market had been systematically highly or slightly leveraged our estimates will be biased. In Appendix C.4, we report the results of ordinary-least squares regressions of annual firm-level characteristics, such as the leverage ratio or the total assets, on our treatment variables. The regressions show that even within the same market, firms with different exposure to their lenders' distress were trending differently before the crisis. For example, firms with higher exposure to the banks' funding gap at June 2010 were significantly less leveraged than firms less exposed to banks' distress.²³ Although this evidence might reinforce the impact of our treatment variable as firms that borrowed from highly distressed banks were likely to be the "healthiest" before the crisis (for instance, less leveraged), the possible presence of different firm-level trends before the crisis can bias our results. To account for this possibility, one would ideally control for firm fixed effects (α_i) as well as firm-specific time trends (for instance, $\alpha_i * time$). Given the large dimensionality of our panel dataset, one would have to compute more than 100,000 firm-specific fixed effects and as many linear trends. Due to time and technological constraints, we are not able to estimate fixed effects and time trends for each individual firm. Therefore, we opt for a simple solution and so include the most significant firms' characteristics that we observe in the data.²⁴

²³This result might seem puzzling but this is what the data say. We computed the same regressions dropping the doubtful observations where firms have negative leverage or leverage larger than 100% and we still get the same result.

²⁴Although the inclusion of firms' characteristics do not control for the full individual trend, we can get an idea of the direction of the bias induced by not controlling for firm-specific trends.

4.6. Econometric model and results

In order to identify the impact of the funding shock to Italian banks following the Italian crisis on the ability of the borrowing firms to repay their loans, we run the following regression

$$\mathbb{P}(NPL_{it}=1) = NPL_{it} = \beta_0 EXP_{it-1} + \beta_1 EXP_{it-1} * Crisis_t + \mu_{mt} + X_{it}\gamma + \epsilon_{it} \qquad (4.1)$$

where EXP_{it-1} is our measure of firm *i*'s exposure to banks' financial distress in period t-1. The binary variable $Crisis_t$ equals 1 after July 2011 and 0 otherwise. The market-time fixed effect μ_{mt} equals 1 if firm *i* operates in market *m* at time *t* and equals 0 otherwise. μ_{mt} captures unobserved demand conditions and it helps to separate the pure credit channel of the shock from the general aggregate demand channel.²⁵ X_{it} is a vector that includes firm characteristics such as liquidity and solvency ratios, as well as the firm's exposure to several dimensions of its lenders' financial health such as Tier 1 capital ratios, total assets and ROA. ϵ_{it} is an idiosyncratic error term which is allowed to exhibit arbitrary correlations within the same province and product market (defined using the first three digits of the ATECO code), but it is assumed to be uncorrelated across markets. Parameters β_0 and β_1 are the main parameters of interests as they measure the marginal effects of the firms' exposure to their lenders' financial distress before and after the funding shock of July 2011.

Table 4.9 reports the results of our estimation. Column (1) reports the estimates of the ordinary least squares regression of NPL_{it} on the treatment variables EXP_{it-1} and $EXP_{it-1} * Crisis_t$ with no additional controls. The results suggest that, before July 2011, the firms' exposure to the banks' funding gap in the previous period is negatively correlated with the probability that the firm enters the *non-performing* status. This relation reverts after the start of the Italian crisis as the correlation becomes positive.

Column (2) reports the estimates of the ordinary least squares regression of NPL_{it} on the treatment variables EXP_{it-1} and $EXP_{it-1} * Crisis_t$ with the addition of market-time fixed effects μ_{mt} to control for unobserved shocks to product demand. Comparing the results in column (2) to those in column (1), we can see that the measure of exposure has a weaker impact on the firm's credit performance once we look at the residual variation within markets. Moreover, the impact of the firm exposure remains negative even following the start of the Italian crisis. The presence of feedbacks effects from the dependent variable to our

 $^{^{25}\}mathrm{Section}$ 4.4.3 provides some statistics about the markets.

treatment variable can explain part of this result. It is easy to believe that current and past values of the lagged dependent variable, NPL_{it} , are correlated with EXP_{it-1} . Simple pairwise correlation shows a coefficient of -4% between EXP_{it} and NPL_{it} , which is consistent with a downward bias in our model specification in column (2). This links to the results in column (3).

Column (3) reports the estimates of the two-stage least squares regression of NPL_{it} on the treatment variables EXP_{it-1} and $EXP_{it-1} * Crisis_t$, with market-time fixed effects μ_{mt} , where the treatment variables are instrumented with their respective values at June 2010, $EXP_{i,6-2010}$ and $EXP_{i,6-2010} * Crisis_t$. The F-stats of excluded instrument are significantly larger than 10 for both the plain treatment EXP_{it-1} and its interaction term $EXP_{it-1} * Crisis_t$. In particular, the estimates in column (3) look quite reasonable. The firms' exposure is negatively related to the firms' financial performance before the crisis and it is positively correlated after the Italian crisis. One problem with the results in column (3) is represented by the net sign of the marginal effect of EXP_{it-1} post-crisis which is about -.03, a negative number. This is against what we expect from the understanding of the channels underlying the funding shock as more distressed banks should have tightened money supply in response to shock more than less distressed banks. The model in column (3), however, suffers from two additional issues: unaccounted complexity of the mechanisms of transmission of the shock, and the fact that some firms within the same market could have been trending differently already before the Italian crisis started.

Column (4) reports the estimates of the second stage of the two-stage least squares regressions of NPL_{it} on the treatment variables with market-time fixed effects, when other dimensions of the firms' exposure to banks' heterogeneity (all computed at June 2010) are included. Looking at the estimates, the measure of exposure has no significant impact on the firms' credit performance before the start of the Italian crisis but it has a significant and positive impact after the Italian crisis. This result is what we expect to see from our understanding of the transmission mechanisms.

Finally, column (5) reports the same specification as in column (4) with the addition of annual firm characteristics. The estimates in column (5) looks very similar to those in column (4), though the coefficient of the interaction EXP_{it-1} is larger. This last result might suggest that the exclusion of firm-specific time trends creates a downward bias in our estimates. To make sense of this, we need to consider the evidence provided in Appendix C.4 where the results of several least squares regressions of firms' annual characteristics on our treatment variables are reported. These Tables suggest that the healthiest firms, for example those with low leverage, are likely to have borrowed from the most distressed banks before the Italian crisis started. To understand how this can affect the size of the estimates in column (5), we have to compare the setting we have in our data with a hypothetical situation where firms are randomly allocated to banks. In this second scenario the impact of the shock should on average be stronger than what we observe in the data as the financial health of the firms that borrow from distressed banks is lower in the hypothetical case compared to what we observe in the data.

To make sense of the estimates, we consider the specification in column (5) as our baseline model. The coefficient of our treatment variable EXP_{it-1} suggests that before the start of the Italian crisis of July 2011 the financial health of the lenders have no significant impact on the financial performance of the borrowing firms. Following the Italian crisis, this relation becomes positive and significant. We find that, holding all control variables constant, one standard-deviation increase (about 8%) in our measure of the firms' exposure to their lenders' financial distress increases the probability that the firm becomes non-performing by about 0.4%. From Table 4.3, we see that the rate of non-performing loans in Italy increased by about 2.8% (6.2%-3.4%) points following the start of the Italian crisis in July 2011. If we consider this change as the total effect of the funding shock on the financial performance of the borrowing firms in Italy, and we accept the estimates of column (5) in Table 4.9 as reasonable, we can say that the size of the credit channel of the funding shock of July 2011 is about 0.4% points while the size of aggregate demand shock is about 2.4% points.²⁶

4.7. Robustness checks

In this section, we discuss some issues that might be left unsolved with our identification strategy as well as some downsize in the data.

We do not observe the statistical units or plants of production for the firms we have in

²⁶The estimates of the other control variables look reasonable apart two of them: the exposition to the share of risky bonds at June 2010 and to the firms' total assets. We still do not have a clear explanation for why this is the case, but we believe this might be the consequence of unobserved heterogeneity and non-linearity in the mechanisms of transmission of the shock that are not controlled for in our model specification.

the data. This can make the definition of the markets using provinces and the full 6-digit ATECO code problematic as some firms might have production plants in several markets that we do not observe. To partially control for this, we run an additional specification of the baseline specification of column (5) in Table 4.9 where only small and medium sized firms are included. The firms are classified as small or medium sized if either their total assets or their total revenues for a particular year are below the median we observe in the data. The results of this regression are reported in Appendix C.5 and are very similar to the estimates in our baseline model. The intuition behind this robustness check is the following: smaller firms are more likely to operate in a local market which could be roughly included within the border of an Italian province, making our market-time fixed effects more effective in controlling shocks to product demand.

In the estimations we use data from 2006 to 2012. It is natural to question whether feedbacks effects between the firms' exposure to the banks' distress and the firms' financial status (i.e., performing vs non-performing) already existed before the start of the Italian crisis in July 2011. In Appendix C.5 we show the results of our baseline model specification where we use only data from December 2010 to December 2012 as this makes our instrument at June 2010 exempt of any feedback effect from the dependent variable. The results of this regression are very similar to those of our baseline model specification of Table 4.9.

One could argue why we do not analyse the impact of the collapse of Lehman Brothers instead. The collapse of Lehman Brothers on September 2008 seems to have hit Italian banks to a lesser extent compared to banks based in the UK and in the US as the former were much less exposed to collateralized-debt obligations such as sub-prime mortgage-backed securities. Moreover, Italian banks continued to fund part of their operations in the repo market where Italian government bonds are pledged as collateral. In the context of our analysis, this can be interpreted as follows: Italian banks believed that the wholesale money market in Europe, where government bonds are heavily used as collateral, was sufficiently liquid. After July 2011 the banks drastically changed these beliefs as the spread between Italian bonds and German bonds jumped to values never seen in recent history, making Italian bonds a riskier asset than what was previously believed. Appendix C.5 shows three regressions of the baseline model specification for three different periods: before Lehman collapse, between Lehman and the Italian crisis, and after the start of the Italian crisis. The regressions show that the sign of the coefficient of the treatment variable exposure changes over time and becomes significantly positive only after the start of the Italian crisis.

4.8. Conclusions

In this work we study how wholesale funding shocks may affect the credit performance of banks' borrowers. We construct a novel dataset of bank-firm relationships that provides detailed information on bank and firm characteristics. We also propose an empirical strategy to estimate the impact that a worsening in banks' wholesale funding opportunities has on firms' ability to repay their loans. We exploit the Italian sovereign debt crisis of July 2011 as a significant funding shock to Italian banks. To capture how each bank has been affected by the sudden increase on Italian yields, we use the monthly difference between each bank's lending and retail funding over its total assets. This measure captures each bank's dependence on the wholesale funding market and, thus, its exposure to the sovereign-debt shock. We distinguish how much firms suffered from more restrictive credit conditions rather than from a significant reduction in firms' product demand by considering firms that borrow from stressed banks (i.e., banks that were more vulnerable to a wholesale funding shock) to otherwise-similar firms that borrow from *non-stressed* banks and compare their loan-repayment performance. To compare firms that are otherwise-similar, we control for several firm-level characteristics and we create several geographic-product markets where firms are likely to face the same demand shocks. We find that, following the funding shock of July 2011, a standard-deviation increase in our measure of the firms' exposure to the banks' financial distress increases the probability that the firm is recorded as non-performing in Italy's Credit Register by about 0.4% (i.e., the size of the credit channel). The results also suggest that the aggregate demand channel lead to a 2.4% increase in the share of non-performing loans.

To conclude the chapter, we highlight some other limitation of our methodology and provide suggestions for future research.

Another limitation of our identification strategy is the focus on the constructions and services sectors to control for the lack of information on firms' exports. A more comprehensive dataset that includes exports would provide a more complete view of the impact of the Italian crisis on the performance of the borrowing firms as this would allow the econometrician to include sectors other than constructions and services such as agricultural and manufacturing. Further, we do not observe balance sheet information of foreign bank groups nor we have

cross-country data. A dataset which includes information on foreign groups and information on the performance of firms in other countries would provide more variation in several dimensions of banks' heterogeneity, such as the share of risky bonds that could help the researcher to better identify the mechanisms of transmission of the funding shock.

In the analysis of this chapter, we never used the average interest rate charged by the banks to their clients. Bank of Italy has this info in the TAXIA database which can be merged to the information already used in the work and see how the results would change if we control for the interest rates.

Also, we know that some firms within the same product market were already trending differently before the crisis and that some of them were already borrowing from distressed banks. Although we included annual firm characteristics, future research should develop methods for dealing with multi-way fixed effects and instrumental variables in high-dimensional panel data models in order to fully control for the unobserved heterogeneity present in the data. Finally, we only discussed the impact of the funding shock in July 2011 in a static framework. In some of our investigation, we find that the outcome variable NPL_{it} exhibits inertia as firms that enter the status non-performing stay in that status for a very long time. Future research should study the dynamics and persistence of non-performing loans, or better study the full dynamics of the five borrowing statuses enlisted in the Basel 2 classification with the use of dynamic ordered models.

Bibliography

- ALTMAN E.I (1968): "Financial Ratios, Discriminant Analysis and the Prediction of Corporate Bankruptcy," *The Journal of Finance*, 23, 589-609.
- ARPA M., I. GIULINI, A. ITTNER, AND F. PAUER (2001): "The influence of macroeconomic developments on Austrian banks: implications for banking supervision." *BIS Paper*, 1, 91-116.
- BATTISTINI N., PAGANO, M., LOPEZ, AND S. SIMONELLI (2013): "Systemic Risk, Sovereign Yields and Bank Exposure in the Euro Crisis" Working Paper, 347, Naples.
- BEHR, P., A. GÜTTLER, AND D. PLATTNER (2004): "Credit Scoring and Relationship Lending: The Case Of German SME,"Working Paper, University of Frankfurt.
- BERGER, A. N., AND R. DE YOUNG (1997): "Problem loans and cost efficiency in commercial banks," *Journal of Banking and Finance*, 21, 849-870.
- BERNANKE, B. AND M. GERTLER (1989): "Agency Costs, Net Worth and Business Fluctuations," *American Economic Review*, 79, 14-31.
- BIKKER, J. A., AND H. HU (2002): "Cyclical patterns in profits, provisioning and lending of banks and procyclicality of the new Basel capital requirements," *BNL Quarterly Review*, 55, 143-175.
- BOFONDI, M., CARPINELLI, L. AND E. SETTE (2013): "Credit supply during a sovereign debt crisis,"Temi di discussione, Working Papers, Banca d'Italia.
- BOFONDI, M. AND T. ROPELE (2011): "Macroeconomic determinants of bad loans: Evidence from Italian banks," *Questioni di Economia e Finanza*, Bank of Italy Occasional Papers 89.

- BRESNAHAN, T. F. AND P. C. REISS (1988): "Do Entry Conditions Vary across Markets?," *Brookings Papers on Economic Activity: Microeconomic Annual*, 3, 833-882.
- BRESNAHAN, T. F. AND P. C. REISS (1991): "Entry and Competition in Concentrated Markets," *Journal of Political Economy*, 99, 977-1009.
- CANOVA, F. AND M. CICCARELLI (2013): "Panel Vector Autoregressive Models: a Survey,"Working Paper, 1507, European Central Bank.
- CAPORALE, G. M., DI COLLI, S. AND J.S. LOPEZ (2013): "Bank Lending Procyclicality and Credit Quality During Financial Crises," *DIW Berlin Discussion Paper*, 1309.
- DIETSCH, M., AND J. PETEY (2004): "Should SME exposures be treated as retail or corporate exposures? A comparative analysis of default probabilities and asset correlations in French and German SME, "Journal of Banking and Finance, 28, 773-788.
- DI PATTI, E. B. AND E. SETTE (2012): "Bank balance sheets and the transmission of financial shocks to borrowers: evidence from the 2007-2008 crisis,"Temi di discussione, Working Papers, Banca d'Italia.
- DIAMOND, D. W., AND R. G. RAJAN (2001): "Liquidity Risk, Liquidity Creation, and Financial Fragility: A Theory of Banking," *Journal of Political Economy*, 109, 287-327.
- ESPINOZA, R., AND A. PRASAD (2010): "Nonperforming Loans in the GCC Banking Systems and their Macroeconomic Effects,"Working Paper, 10/224, International Monetary Fund.
- FIDRMUC, J., AND C. HAINZ (2010): "Default rates in the loan market for SMEs: Evidence from Slovakia," *Economic Systems*, 34, 133-147.
- FILOSA R. (2007): "Stress Testing of the Stability of the Italian Banking System: a VAR Approach," *Heterogeneity and Monetary Policy*, 703, 1-46.
- FREIXAS, X., AND J. C. ROCHET (2001): *Microeconomics of Banking*. The MIT Press, second edition.
- GAMBERA M. (2000): "Simple Forecasts of Bank Loan Quality in the Business Cycle,"*Emerging Issues Series*, Federal Reserve Bank of Chicago.

- GOBBI, G., AND E. SETTE(2013): "Do Firms Benefit from Concentrating their Borroing? Evidence from the Great Recession," *Review of Finance*, 18, 2, 1-34.
- HOGGARTH G., SORENSEN, S., AND L. ZICCHINO (2005): "Stress Tests of UK Banks Using a VAR Approach,"Working Paper, 282, Bank of England.
- HOLMSTROM, B., AND J. TIROLE(1997): "Financial Intermediation, Loanable Funds, and The Real Sector," *The Quarterly Journal of Economics*, 112, 663-691.
- JACOBSON, T., J. LINDÈ, AND K. F. ROSZBACH "Credit Risk versus Capital Requirements under Basel II: Are SME Loans and Retail Credit Really Different?," Journal of Financial Services Research 28 (2005), 43-75.
- JIMENEZ, G., LOPEZ, J.A., AND S. SAURINA (2007): "How Does Competition Impact Bank Risk Taking?"Working Paper, 23, FRB San Francisco.
- JIMENEZ, G., ONGENA, S., PEYDRO, J. AND J. SAURINA (2012): "Credit Supply and Monetary Policy: Identifying the Bank-Balance Sheet Channel with Loan Applications," *American Economic Review*, 102, 2301-2326.
- JIMENEZ,G., AND J. SAURINA (2004): "Collateral, type of lender and relationship banking as determinants of credit risk, "*Journal of Banking and Finance*, 28, 2191-2212.
- JIMENEZ, G., AND J. SAURINA (2006): "Credit Cycles, Credit Risk, and Financial Regulation," *International Journal of Central Banking*, 2, 65-98.
- JIMENEZ, G., SALAS FUMÁS, V., AND J. SAURINA (2006): "Determinants of collateral," *Journal of Financial Economics*, 81, 255-281.
- KALIRAI, H., AND M. SCHEICHER (2002): "Macroeconomic Stress Testing: Preliminary Evidence for Austria,"Austrian National Bank Financial Stability Report, May, 3.
- KEELEY M. (1990): "Deposit Insurance, Risk and Market Power in Banking," American Economic Review, 80, 1183-1200.
- KHWAJA, A. I., AND A. MIAN (2008): "Tracing the Impact of Bank Liquidity Shock: Evidence from an Emerging Country," *American Economic Review*, 98, 1413-1442.

- KIYOTAKI, N., AND J.H. MOORE (1997): "Credit Cycles," Journal of Political Economy, 105, 211-248.
- INESSA, L., AND T.A. Rima (2013): "Macro-Financial Linkages in Egypt: A Panel Analysis of Economic Shocks and Loan Portfolio Quality,"Working Paper, University of Hawai'i at Mānoa.
- LOUZIS D.P., VOULDIS, A.T., AND V.L. METAXAS (2012): "Macroeconomic and bankspecific determinants of non-performing loans in Greece: A comparative study of mortgage, business and consumer loan portfolios," *Journal of Banking and Finance*, 36, 1012-1027.
- MARCUCCI, J., AND M. QUAGLIARIELLO (2008): "Is Bank Portfolio Risk Procyclical? Evidence from Italy Using a Vector Autoregression," Journal of International Financial Markets, Institutions and Money, 18, 46-63
- MARCUCCI, J., AND M. QUAGLIARIELLO (2009): "Asymmetric Effects of the Business Cycle on Bank Credit Risk," *Journal of Banking and Finance*, 33, 1624-1635.
- McCANN, F., AND T. McINDOE-CALDER (2012): "Determinants of SME Loan Default: The Importance of Borrower-Level Heterogeneity,"Working Paper, Central Bank of Ireland.
- NKUSU M. (2011): "Nonperforming Loans and Macrofinancial Vulnerabilities in Advanced Economies," Working Paper, 11/161, International Monetary Fund.
- QUAGLIARIELLO M. (2007): "Banks' Riskiness over the Business Cycle: A Panel Analysis on Italian Intermediaries," *Applied Financial Economics*, 17, 119-138.
- PESOLA J. (2001): "The Role of Macroeconomic Shocks in Banking Crises," Bank of Finland Discussion Papers, 6.
- REPULLO R. (2004): "Capital Requirements, Market Power and Risk-Taking in Banking," *Journal of Financial Intermediation*, 13, 156-182.
- SALAS, V., AND J. SAURINA (2002): "Credit Risk in Two Institutional Regimes: Spanish Commercial and Savings Banks," *Journal of Financial Services Research* 22, 203-224.
- STEIN J. (1998): "An Adverse Selection Model of Bank Asset and Liability Management with Implications for the Transmission of Monetary Policy," *The RAND Journal of Economics*, 29, 466-486.
ZOLI E. (2013): "Italian Sovereign Spreads: Their Determinants and Pass-through to Bank Funding Costs and Lending Conditions,"Working Paper, 13/84, International Monetary Fund.



Figure 4.1.: Time series BTP-BUND spread

This chart is taken from Zoli [2013].

Macro areas	Percent	NPLs Obs.	Tot. obs.
Central North East North West Southern-Insular Total	$ \begin{array}{r} 4.4 \\ 3.5 \\ 3.1 \\ 6.0 \\ 4.1 \end{array} $	$10,090 \\ 8,274 \\ 9,487 \\ 11,832 \\ 39,683$	$\begin{array}{c} 227,327\\ 234,854\\ 310,693\\ 198,398\\ 971,272\end{array}$
Sector	Percent	NPLs Obs.	Tot. obs.
Constructions Services Total	$5.0 \\ 3.7 \\ 4.1$	14,031 25,652 39,683	279,522 691,750 971,272

Table 4.1.: Distribution of NPLs by geography and by sector

Notes. The number of observations refers to those firm - half-year combinations where the firms are recorded as non-performing in a given geographic area or sector.

			NPLs	
Year	Half	Percent	Obs.	Tot. obs.
2006	1 st	1.5	630	41,231
2006	2nd	1.5	657	42,756
2007	1 st	1.6	810	$51,\!471$
2007	2nd	1.6	865	$54,\!125$
2008	1 st	1.7	$1,\!136$	$65,\!070$
2008	2nd	2.3	1,565	$68,\!527$
2009	1 st	3.1	2,058	66,293
2009	2nd	3.8	2,531	67,007
2010	1 st	5.1	4,840	$94,\!395$
2010	2nd	5.9	6,084	103,336
2011	1st	4.9	4,189	86,083
2011	2nd	6.0	5,083	84,474
2012	1st	5.5	4,010	$72,\!695$
2012	2nd	7.1	$5,\!225$	$73,\!809$
Total		4.1	$39,\!683$	971,272

Table 4.2.: Time-series of NPLs - semi-annual data

Notes. The number of observations refers to the number of firms that in a given half year are recorded as non-performing.

	Percent	NPLs Obs.	Tot. obs.
Before Lehman	1.6	4,098	254653
After Italian crisis Total	$4.4 \\ 6.2 \\ 4.1$	14,318 39.683	230,978 971 272
rouur	T • T	00,000	01110
		NPLs	,
	Percent	NPLs Obs.	Tot. obs.

	Table $4.3.$:	Distribution	of NPLs	before an	nd after t	the crises
--	----------------	--------------	---------	-----------	------------	------------

Notes. The number of observations refers to those firm - half-year combinations where the firms are recorded as non-performing in a given geographic area or sector.

			NPLs					
		Before			After			
Macro areas	Percent	Obs.	Tot. obs.	Percent	Obs.	Tot. obs.		
Central	3.7	$6,\!867$	184,033	7.0	4,111	58,733		
North East	2.9	$5,\!557$	188,515	5.5	$3,\!533$	$64,\!070$		
North West	2.5	6,300	$247,\!996$	4.9	$3,\!987$	$81,\!953$		
Southern-Insular	5.1	8,440	164,128	9.3	4,666	50,088		
			NPLs					
		Before			After			
Sector	Percent	Obs.	Tot. obs.	Percent	Obs.	Tot. obs.		
Constructions	4.0	8,512	213,777	8.4	5,519	65,745		
Services	3.2	$16,\!853$	$526,\!517$	5.3	8,799	165,233		

Table 4.4.: Distribution of NPLs before and after the Italian crisis by geography and by sector

Notes. The number of observations refers to those firm - half-year combinations where the firms are recorded as non-performing in a given geographic area or sector before or after the Italian crisis.

Variable	Obs	Mean	Std. Dev.	Min	Median	Max
Total assets	$1,\!204$	$31,\!636$	106,424	69	$3,\!866$	849,759
Total loans	$1,\!204$	19,771	$63,\!127$	1	$2,\!454$	488,078
% NPLs/Assets	$1,\!085$	3.2	3.2	0	2.5	40.1
% Tier 1 ratio	868	11.4	5.7	1.6	10.0	44.6
% Tier 2 ratio	868	13.4	5.1	3	12.1	44.6
$\% \mathrm{ROA}$	$1,\!204$	1.7	1.7	-5.7	1.3	14.9
$\% \mathrm{ROE}$	$1,\!204$	3.9	10.2	-71.4	3.5	65
% Funding gap	$1,\!204$	2.8	23.6	-84.1	3	111.5
% Risky share	333	90.8	11.1	43.3	94.8	100

Table 4.5.: Banks' characteristics

Notes. The unit of observation is a bank - half-year pair. Total assets and total loans are expressed in millions of euros. NPLs/Assets is the share of a bank's non-performing loans over its total assets. Tier 1 ratio and Tier 2 ratio are capital requirements imposed by Basel II. ROA is the return on assets. ROE is the return on equity. Funding gap is the ratio to total assets of the difference between lending and retail funding. Risky share is the ratio to total Bond holding of the bonds that are considered "risky" by the regulator. All ratios are expressed as percentages.

NPLs/Assets				
	Before		After	
Funding gap June 2010	Percent	Observations	Percent	Observations
Negative	2.4	315	3.6	112
Positive	2.6	459	5.8	199
Total	2.6	774	5.0	311
Delta loans/Assets				
Delta loans/Assets	Before		After	
Delta loans/Assets Funding gap June 2010	Before Percent	Observations	After Percent	Observations
Delta loans/Assets Funding gap June 2010 Negative	Before Percent 3.7	Observations 324	After Percent 3.2	Observations 123
Delta loans/Assets Funding gap June 2010 Negative Positive	Before Percent 3.7 4.5	Observations 324 460	After Percent 3.2 -1.1	Observations 123 204

Table 4.6.: Change in banks' NPLs and total loans before and after the Italian crisis grouped by sign of funding gap at June 2010

Notes. The unit of observation is a bank - half-year pair. NPLs/Assets is the share of a bank's non-performing loans over its total assets. Delta loans/Assets is the change of a bank loans between two half-year over its total assets in the first half-year. All ratios are expressed as percentages. The table reports how NPLs/Assets and Delta loans/Assets changed before and after the Italian crisis for two groups of banks: those banks with a liquidity surplus at June 2010 (i.e. negative funding gap) and those with a liquidity deficit at June 2010 (i.e. positive funding gap).

Table 4.7.: Firms' exposure to banks' financial distress

Variable	Obs	Mean	Std. Dev.	Min	Median	Max
Exposure to % Funding gap	971,272	6.5	8.4	-65.5	6.7	108.1
Exposure to $\%$ Funding gap at June 2010	$971,\!272$	5.1	7.5	-65.8	6.3	98.4

Notes. The unit of observation is a bank - half-year pair. Exposure to funding gap is a firm-level average of its lenders funding gap in a given half-year. The weights used for averaging are the share of the credit granted by a bank to the firm over the firm's total credit granted in a given half year. Exposure to funding gap at June 2010 is the value of Exposure to funding gap at June 2010. All ratios are expressed as percentages.



Figure 4.2.: Exposure to funding gap

The histogram is computed using firm - half-year as unit of observation, for the selected sample of firms.

Variable	Obs	Mean	Std. Dev.	Min	Median	Max
EBIT	971,272	127	3,143	-1,478,188	39	$559,\!607$
% EBITDA/Assets	$971,\!272$	6	24.4	-5,240	6.3	848
HHI	$971,\!272$	$6,\!417$	$2,\!950$	0	$5,\!669$	10,000
% Leverage	$971,\!272$	67.7	395.8	-69,300	74.1	$138,\!900$
% Liquidity	$971,\!272$	5.8	9.6	0	1.8	100
No banks	$971,\!272$	2.6	1.8	1	2	30
$\% \mathrm{ROA}$	$971,\!272$	2.7	34	-9,203	3.7	$10,\!851$
Total assets	971,272	4,080	$55,\!419$	10	1,166	$20,\!900,\!000$

Table 4.8.: Firm characteristics

Notes. EBIT is annual earnings before interests and taxes. EBIT and Total assets are annual firm-level variables expressed in thousands of euros. EBITDA is annual earnings before interests, taxes, depreciation and amortization. EBITDA/Assets is EBITDA over total assets. Leverage is the annual share of financial debt over the sum of financial debt plus equity. Liquidity is the annual fraction of a firm's cash holdings over its total assets. ROA is returns on assets. EBITDA/Assets, Leverage, Liquidity and ROA are annual firmlevel indicators expressed as percentages. The number of banks is the number of lenders a firm borrow from in a given half year. HHI is the Herfindahl-Hirschman Index of a firm's loans portfolio across its lenders in a given half year. HHI can take values between 10,000/No banks when the firm borrows an equal amount from each of its lenders, and 10,000 when the firm only borrows from one bank.

		(1)	(2)	(3)	(4)	(5)
	Sample	Selected	Selected	Selected	Selected	Selected
	Type of Regression	OLS	OLS	2SLS-IV	2SLS-IV	2SLS-IV
	Dependent Variable	NPL	NPL	NPL	NPL	NPL
Treatment	Expo. Funding Gap	-0.1293***	-0.0606***	-0.1142***	-0.002	0.0023
		(0.0059)	(0.0066)	(0.0122)	(0.0202)	(0.02)
	Expo. Funding Gap*Crisis	0.1631^{***}	-0.0922***	0.0765^{***}	0.0461^{**}	0.0507^{***}
		(0.0078)	(0.0136)	(0.0135)	(0.0182)	(0.0181)
Banks characteristics	Expo. ROA June 2010				-0.0423***	-0.0428***
	-				(0.0031)	(0.0034)
	Expo. ROE June 2010				-0.0008*	-0.0007
					(0.0005)	(0.0005)
	Expo. Tier 2 Ratio June 2010				-1.1056^{***}	-1.1508***
					(0.1277)	(0.1352)
	Expo. Risky Share June 2010				-0.0005***	-0.0004***
	1 U				(0.00004)	(0.00004)
	Expo. NPLs/Assets June 2010				-0.0039***	-0.0039***
	- ,				(0.0009)	(0.001)
	Expo. Tier 1 Ratio June 2010				0.1761	0.2229*
	1				(0.1237)	(0.1334)
	Expo. Total Assets June 2010				-0.0000***	-0.0000***
	-				(0.0000)	(0.0000)
Dimon al ana stariation	EDIT				()	0.00001***
Firms characteristics	EDII					-0.000001
						(0.0000003)
	EBI1DA/Assets					-0.0007
						(0.0003)
	HHI					-0.000006***
	T					(0.00001)
	Leverage					(0.000000)
	T 11,					(0.00002)
	Liquidity					-0.0007***
	NT 1 1					(0.0001)
	No banks					-0.0082***
	DOL					(0.0008)
	ROA					-0.0005
						(0.0003)
	Total assets					0.00000003*
						(0.00000001)
	F-stat 1			8,531	3,149	4,140
	F-stat 2 (interaction)			6,128	6,499	6,828
	Firm dummy	No	No	No	No	No
	Market6-Time dummy	No	Yes	Yes	Yes	Yes
	No of Market6-Time dummies	0	132,708	132,708	122,689	107,266
	Other controls	No	No	No	No	No
	Clustering	Market3-Time	Market3-Time	Market3-Time	Market3-Time	Market3-Time
	Observations	971,272	971,272	971,272	864,272	751,930

Table 4.9.: Main regressions

Notes. The unit of observation is a firm - half-year pair. NPL is a binary variable that equals 1 if the firm is recorded as non-performing in a given period and equals 0 otherwise. Expo. Funding Gap is the time-variant one-period lagged firm's exposure to its lenders funding gap, and Expo. Funding Gap*Crisis its interaction with the binary variable Crisis which equals 1 for after July 2011 and equals 0 before. Expo. Funding Gap is expressed as a fraction of the total assets (not as a percentage). In columns (3)-(5) Expo. Funding Gap and Expo. Funding Gap*Crisis are instrumented using the respective measures computed at June 2010. The variables Expo. Var June 2010, for Var equal to ROA, Tier 1 and Tier 2 ratios, Risky Share, NPLs/Assets, and Total Assets is the time-invariant firm's exposure to its lenders Var at June 2010. EBIT is annual earnings before interests and taxes. EBIT and Total assets are annual firm-level variables expressed in thousands of euros. EBITDA is annual earnings before interests, taxes, depreciation and amortization. EBITDA/Assets is EBITDA over total assets. Leverage is the annual share of financial debt over the sum of financial debt plus equity. Liquidity is the annual fraction of a firm's cash holdings over its total assets. ROA is returns on assets. EBITDA/Assets, Leverage, Liquidity and ROA are annual firm-level indicators expressed as number (not percentages). The number of banks is the number of lenders a firm borrow from in a given half year. HHI is the Herfindahl-Hirschman Index of a firm's loans portfolio across its lenders in a given half year. HHI can take values between 10,000/No banks when the firm borrows an equal amount from each of its lenders, and 10,000 when the firm only borrows from one bank. Market6-Time dummy represent market-time fixed effects where the time is represented by an half year and a market is defined as a geography-product pair where the geography is a province while the product market is defined using the full 6-digit ATECO 2007 code. Market3 is a geography-product pair where the geography is a province while the product market is defined using the first 3 digits of the ATECO 2007 code. We always cluster standard errors with respect to Market3 to allow for arbitrary correlations between disturbances within a broad definition of market. The number of observations in columns (4) and (5) differs because the included variables are missing for some firms and we need to drop some more markets that are left with only one firm.

Chapter 5. Conclusions

This thesis applies econometric methods to analyze different economic research questions using microeconomic data. The first essay (chapter 2) analyzes consumer searching behavior in a grocery context. The second essay (chapter 3) studies the implications of the introduction of a bonus scheme in a principal-agent context using data from furniture sales. The third essay (chapter 4) proposes an empirical strategy to estimate the impact that a worsening in banks' wholesale funding opportunities (such as the Italian sovereign debt crisis of 2011) has on borrowers' ability to repay their loans. Chapter 5 concludes the thesis and provides some directions for future work.

Chapter 2 contributes to the understanding of the consumer decision journey. We estimate the effect of search intensity on the price a consumer pays within a particular category using data from RFID tags on supermarket shopping carts. Recording search in a physical store environment is generally challenging and even our detailed data is only able to capture total search-time, but not which options the consumer evaluated. The technology does however have the advantage of not interfering in any way with the consumer's natural shopping experience and might be the best possible way to gain insights into consumer search in a brickand-mortar store. To the best of our knowledge this work is the first to use direct data on search effort to analyze consumer search within a brick-and-mortar environment. Due to the limited amount of observations per category in the data our evidence comes from regressions which are pooled across categories. Going forward, with path-data over a longer time-horizon for only one category it should be possible to model the search process in more detail (possibly by means of a structural model). In particular, our approach only looks at the effect on price paid and does not directly analyze the role of other product characteristics. We are therefore not able to make any statements about the effect of search on consumer utility. However, we believe that the effect of search-time on price is a dimension of the search process which is particularly relevant for informing optimal supply-side behavior. Our findings imply that, due to the limited amount of search, the use of marketing tools such as feature advertising and in-store displays can be very effective. Furthermore, firm behavior that influences consumer search interacts in an interesting way with pricing decisions. Because more searching makes finding a lower price or promoted product more likely, firms have an incentive to encourage search when running a promotion.

Chapter 3 contributes to the literature in incentives when the agent is given control over price. To the best of my knowledge, this is one of the first empirical papers that exploits variation in both the nature of the incentive scheme and price delegation. Using data on the staff of a furniture firm, I show that when a bonus scheme conditional on revenues was introduced, it increased the revenues generated by all of the sales employees, but I find no significant heterogeneous effect of the bonus scheme depending on whether the employee is given control over price or not. Moreover, I show that giving the sales staff control over price does not significantly increase revenues. The effects of the bonus scheme and of price delegation on gross profits minus paid bonuses, commissions, and wages were similar. These results are robust to a number of checks, and are consistent with a model of moral hazard and price delegation. We can interpret these results in light of the theoretical model presented in this chapter. The results suggest that the agent might have some private information but this is not enough to fully off-set the negative impact of moral hazard on pricing. Although these results are robust to alternative model specifications, empirical challenges might be important. For example, the small sample size does not allow to perfectly controlling for all observed heterogeneity. Future works should provide more transparent evidence by using larger data-sets, variation across firms and industries or even engineering ad-hoc randomizedcontrol trials that clearly generate the variation in the data necessary for identifying the implications of the bonus scheme and of price delegation.

Chapter 4 explores the impact of the Italian sovereign debt crisis of summer 2011 on the real economy using detailed microeconomic data on credit linkages between firms and banks. We find that, following the funding shock of July 2011, a standard-deviation increase in our measure of the firms' exposure to the banks' financial distress increases the probability that the firm is recorded as non-performing in Italy's Credit Register by about 0.4% (i.e., the size of the credit channel). The results also suggest that the aggregate demand channel lead to a 2.4% increase in the share of non-performing loans. At this exploratory stage of

the research we recognize some limitations of our analysis that can be improved in future works. Examples of these limitations include the lack of information on firms' exports, on balance sheet information of foreign bank groups, on other countries, and on interest rates. Our analysis is also affected by limitations in available techniques for estimating multi-way fixed effects estimators in highly-dimensional panel data models with endogenous variables as well as dynamic panel data modes for discrete outcome variables. Although our results represent the outcome of an exploratory study, we believe they provide a solid basis for the understanding of the transmission mechanisms of shocks from the financial sector to the real economy.

Appendices

Appendix A. Appendix to chapter 2

A.1. Linking Sales and Path Data

One of the interesting features of our dataset is the linkage of sales to trip records. As part of the RFID tracking process, the data reports when the consumer arrives at the checkout. Independently, the sales data also has a time-stamp for each shopper's transaction at the checkout. Comparing the timestamp of a particular path with the sales data allows to define a set of "candidate" checkout product baskets that occured at a similar point in time.¹ Matching which trip goes with which specific transaction involves considering the physical location (i.e., longitude = x and latitude = y relative to the store map) of all the UPCs in each candidate basket. Based on how many of those locations lay on the path we are trying to match, a score is created for the baskets and the highest scoring one is matched to the path.² The matches do not necessarily yield a perfect score as consumer might occasionally leave the cart and pick up an item. Because of this we might not see the path of the consumer going past a specific item, even if the item part of her matched purchase basket. In this case no information on search-time will be available for the particular item.

Finally, when recording the data the location of products within the store is established once at the beginning of the sample period. As it is to costly to continuously track product placement at a daily level, there is a (small) level of noise in the data. The big majority of products in the store do not move within the short time window of our data. However, some movement does occur, primarily due to special promotional displays (end of aisle displays

¹The path-data timestamp that record the arrival at the checkout can be noisy as the consumer will be stationary when standing in line at the cashier. Therefore checkout baskets within a certain time-window after the consumer became stationary in the check-out area qualify as possible matches.

 $^{^2\}mathrm{The}$ data provider did not disclose the precise algorithm to us.

for instance). Overall we (and the data provider) believe that this is a relatively minor issue regarding the quality of our data.

A.2. Tables

	(-)	(2)	(9)	(4)	(G)	(0)
Dependent Variable S	Speed	Speed	Speed	Search-Time	Search-Time	Search-Time
1st Decile 0.	758***	0.707***	0.601***	-2.250***	-1.737***	-1.324***
ist Deene 0.	0.020)	(0.021)	(0.021)	(0.239)	(0.237)	(0.238)
2nd Decile 0.	128***	0.109***	0.098***	-0.230	-0.001	-0.068
	(0.020)	(0.020)	(0.020)	(0.239)	(0.235)	(0.228)
3rd Decile	0.008	0.001	0.000	0.067	0.178	0.153
(0.020)	(0.020)	(0.020)	(0.232)	(0.227)	(0.220)
4th Decile	/	()	()	omitted	()	()
				category		
5th Decile	0.013	0.019	0.028	-0.204	-0.218	-0.383*
((0.020)	(0.020)	(0.019)	(0.230)	(0.225)	(0.218)
6th Decile).033*́	0.040**	0.043**	-0.293	-0.364	-0.445**
(0.020)	(0.019)	(0.019)	(0.229)	(0.224)	(0.216)
7th Decile	0.029	0.027	0.033*	-0.435*	-0.460**	-0.603***
(0.019)	(0.019)	(0.019)	(0.226)	(0.221)	(0.214)
8th Decile 0.	063***	0.063***	0.051***	-0.491**	-0.430**	-0.447**
(0.019)	(0.019)	(0.019)	(0.222)	(0.217)	(0.210)
9th Decile 0.	109***	0.110***	0.105***	-0.586***	-0.504**	-0.525**
(0.019)	(0.019)	(0.019)	(0.221)	(0.217)	(0.210)
10th Decile 0.	113***	0.118^{***}	0.119^{***}	-1.651^{***}	-1.511^{***}	-1.665***
(0.018)	(0.018)	(0.018)	(0.208)	(0.209)	(0.208)
Below 15 Minutes 0.	416***	0.391^{***}	0.381***	-3.830***	-3.532^{***}	-3.281***
(0.016)	(0.016)	(0.016)	(0.190)	(0.188)	(0.183)
15 to 30 Minutes 0.	179^{***}	0.170^{***}	0.158^{***}	-2.629^{***}	-2.490^{***}	-2.295^{***}
(0.016)	(0.016)	(0.015)	(0.182)	(0.179)	(0.174)
30 to 45 Minutes 0.	059^{***}	0.056^{***}	0.049^{***}	-1.530^{***}	-1.550^{***}	-1.420^{***}
(0.017)	(0.017)	(0.016)	(0.195)	(0.191)	(0.185)
45 to 60 Minutes	0.029	0.027	0.028	-1.017^{***}	-1.037^{***}	-0.941^{***}
(0.019)	(0.019)	(0.019)	(0.226)	(0.221)	(0.214)
Above 60 Minutes				omitted		
				category		
Constant 2.	025^{***}	n/a	n/a	13.268^{***}	n/a	n/a
(0.019)			(0.226)		
Category FEs	No	Yes	No	No	Yes	No
Location coordinate FEs	No	No	Yes	No	No	Yes
Observations 2	28,603	$28,\!603$	$28,\!603$	$28,\!603$	$28,\!603$	$28,\!603$

Table A.1.: The Effect of Trip-Duration and Within-Trip Decile on Search-Time and Speed

We regress speed and search-time and dummies for trips of different length and deciles within the trip. The omitted categories are pegged to the trip / decile with the slowest speed: above 60 minutes trips and the 4th decile. We add category fixed effects (149) in columns (2) and (5) and a fixed effect for every unique product location coordinate (974) in columns (3) and (6).

Appendix B. Appendix to chapter 3

B.1. Compensation schemes

Figure B.1.: Compensation Scheme I

Scheme 1 (no bonus) - 28 observations



Figure B.2.: Compensation Scheme II



Scheme 2 (no bonus) - 88 observations

Figure B.3.: Compensation Scheme III



Scheme 3 (with bonus) - 20 observations











B.2. Theoretical model: numerical example

Bonus scheme	Information st	ructure	Pricing scheme											
				Decentralized							Centralized			
		μ	р	Е	\mathbf{q}	R	GP	GP-SC	р	Е	q	R	GP	$\operatorname{GP-SC}$
Bonus Yes	No advantage	0	£219	0	7,000	$\pounds 1,531,250$	$\pounds 691,250$	£642,469	$\pounds 279.3$	0	5,063	$\pounds 1,413,948.07$	£806,440	£758,715
	Advantage	23,554	$\pounds 587$	0	18,777	$\pounds 11,018,038$	$\pounds 8,764,794$	$\pounds 8,\!630,\!631$	$\pounds 279.3$	0	28,617	$\pounds7,992,483$	$\pounds 4,558,485$	$\pounds 4,451,553$
		4,711	$\pounds 292$	0	9,355	$\pounds 2,735,114$	$\pounds 1,\!612,\!465$	$\pounds 1,552,849$	$\pounds 279.3$	0	9,773	$\pounds 2,729,655$	$\pounds 1,556,849$	$\pounds 1,497,283$
		2,355	$\pounds 256$	0	8,178	$\pounds 2,089,839$	$\pounds 1,108,514$	$\pounds 1,054,706$	$\pounds 279.3$	0	7,418	$\pounds 2,071,802$	$\pounds 1,181,645$	$\pounds 1,127,999$
		-2,355	$\pounds 182$	0	5,822	$\pounds 1,059,348$	$\pounds 360,672$	$\pounds 316,138$	$\pounds 279.3$	0	2,707	$\pounds756,094.62$	$\pounds 431,236$	£389,431
		-4,711	$\pounds 145$	0	$4,\!645$	$\pounds 674,132$	$\pounds 116,781$	$\pounds75,714$	$\pounds 279.3$	0	352	£98,241.16	$\pounds 56,032$	£20,147
		-5,061	$\pounds 140$	0	$4,\!470$	$\pounds 624,263$	$\pounds 87,923$	$\pounds 47,305$	$\pounds 279.3$	0	2	$\pounds 436.56$	$\pounds 249$	$-\pounds34,755$
Bonus No	No advantage	0	$\pounds 219$	0	7,000	$\pounds 1,531,250$	$\pounds 691,250$	£644,766	$\pounds 279.2$	0	5,065	$\pounds 1,414,302$	£806,443	£760,836
	Advantage	23,554	$\pounds 587$	0	18,777	$\pounds 11,018,038$	$\pounds 8,764,794$	$\pounds 8,647,158$	$\pounds 279.2$	0	28,620	$\pounds7,990,682$	$\pounds 4,556,333$	$\pounds 4,461,403$
		4,711	$\pounds 292$	0	9,355	$\pounds 2,735,114$	$\pounds 1,612,465$	$\pounds 1,556,952$	$\pounds 279.2$	0	9,776	$\pounds 2,729,578$	$\pounds 1,556,421$	$\pounds 1,500,950$
		2,355	$\pounds 256$	0	8,178	$\pounds 2,089,839$	$\pounds 1,108,514$	$\pounds1,057,841$	$\pounds 279.2$	0	7,421	$\pounds 2,071,940$	$\pounds 1,181,432$	$\pounds 1,130,893$
		-2,355	$\pounds 182$	0	5,822	$\pounds 1,059,348$	$\pounds 360,672$	$\pounds 317,727$	$\pounds 279.2$	0	2,710	$\pounds756,664$	$\pounds 431,454$	£390,779
		-4,711	$\pounds 145$	0	$4,\!645$	$\pounds 674,132$	$\pounds 116,781$	$\pounds76,725$	$\pounds 279.2$	0	355	£99,027	$\pounds 56,465$	£20,723
		-5,061	$\pounds 140$	0	4,470	$\pounds 624,263$	$\pounds 87,923$	$\pounds 48,241$	$\pounds 279.2$	0	4	£1,254	£715	$-\pounds34,294$

Table B.1.: Theoretical model: solution for d = 0

Bonus scheme	Information st	ructure						Pricing scl	neme					
						Decentralized	l					Centralize	ed	
		μ	р	Е	q	R	GP	GP-SC	р	Е	q	R	GP	GP-SC
Bonus Yes	No advantage	0	£220	10	7,025	$\pounds 1,\!542,\!073.65$	$\pounds 699,110$	$\pounds 650,231$	$\pounds 281$	13	5,076	$\pounds 1,\!425,\!616$	$\pounds 816,473$	$\pounds768,\!643$
	Advantage	23,554	$\pounds 588.9$	26	18,843	£11,095,920	$\pounds 8,834,725$	$\pounds 8,699,862$	$\pounds 281$	13	28,630	$\pounds 8,040,635$	$\pounds 4,605,003$	$\pounds 4,497,637$
		4,711	$\pounds 293.4$	13	9,388	$\pounds 2,754,448$	$\pounds 1,627,838$	$\pounds 1,568,048$	$\pounds 281$	13	9,787	$\pounds 2,748,620$	$\pounds 1,574,179$	$\pounds 1,514,442$
		2,355	$\pounds 256.5$	12	8,207	$\pounds 2,104,611$	$\pounds 1,\!119,\!824$	$\pounds 1,065,883$	$\pounds 281$	13	7,432	$\pounds 2,\!087,\!118$	$\pounds 1,\!195,\!326$	$\pounds 1,141,542$
		-2,355	$\pounds 182.6$	8	5,843	$\pounds 1,066,836$	$\pounds 365,695$	£321,094	$\pounds 281$	13	2,721	$\pounds764,114$	£437,620	$\pounds 395,743$
		-4,711	$\pounds 145.7$	7	4,661	$\pounds 678,897$	£119,580	£78,470	$\pounds 281$	13	365	£102,612	£58,767	£22,844
		-5,061	$\pounds 140.2$	6	4,485	$\pounds 628,676$	$\pounds90,444$	$\pounds 49,786$	$\pounds 281$	13	15	$\pounds 4,265$	$\pounds 2,443$	-£32,596
Bonus No	No advantage	0	£219	8	7,021	$\pounds 1,540,262$	£697,794	£651,242	$\pounds 280$	11	5,077	$\pounds 1,424,013$	£814,794	$\pounds769,113$
	Advantage	23,554	$\pounds 588.5$	22	18,832	$\pounds 11,082,882$	$\pounds 8,823,017$	$\pounds 8,704,895$	$\pounds 280$	11	$28,\!631$	$\pounds 8,030,760$	$\pounds 4,595,051$	$\pounds 4,499,821$
		4,711	$\pounds 293.2$	11	9,383	$\pounds 2,751,211$	$\pounds 1,\!625,\!263$	$\pounds 1,569,629$	$\pounds 280$	11	9,788	$\pounds 2,745,362$	$\pounds 1,570,845$	$\pounds1,515,255$
		2,355	$\pounds 256.3$	10	8,202	$\pounds 2,102,138$	$\pounds 1,117,930$	$\pounds 1,067,164$	$\pounds 280$	11	7,432	$\pounds 2,084,687$	$\pounds 1,192,819$	$\pounds 1,142,184$
		-2,355	$\pounds 182.5$	7	5,839	$\pounds 1,065,582$	$\pounds 364,854$	£321,862	$\pounds 280$	11	2,721	$\pounds763,338$	$\pounds 436,768$	$\pounds 396,043$
		-4,711	$\pounds 145.6$	5.5	$4,\!658$	£678,100	$\pounds 119,111$	$\pounds79,025$	$\pounds 280$	11	366	$\pounds 102,663$	£58,742	$\pounds 22,972$
		-5,061	$\pounds 140.1$	5.3	$4,\!483$	$\pounds 627,937$	£90,021	$\pounds 50,312$	$\pounds 280$	11	16	$\pounds 4,439$	£2,540	-£32,493

Table B.2.: Theoretical model: solution for d = 5

Domino och omo	Information of					Prioing advance								
bonus scheme	information st	ructure						Friding sc	neme			<i>a</i>		
						Decentralized						Centralize	d	
		μ	р	Е	q	R	GP	GP-SC	р	Е	q	R	GP	GP-SC
Bonus Yes	No advantage	0	£222	20	7,100	$\pounds 1,575,242.18$	£723,261	$\pounds674,084$	$\pounds 286$	26	5,117	$\pounds 1,461,560$	$\pounds 847,513$	$\pounds799,359$
	Advantage	23,554	$\pounds 595$	54	19,045	£11,334,582	$\pounds 9,049,199$	$\pounds 8,912,188$	£286	26	28,671	£8,189,201	$\pounds 4,748,665$	$\pounds 4,639,962$
		4,711	$\pounds 297$	27	9,489	$\pounds 2,813,693$	$\pounds 1,\!675,\!032$	$\pounds 1,\!614,\!708$	$\pounds 286$	26	9,828	$\pounds 2,807,088$	$\pounds 1,\!627,\!744$	$\pounds 1,567,480$
		2,355	$\pounds 259.2$	23	8,294	$\pounds 2,149,879$	$\pounds 1,154,558$	$\pounds 1,100,209$	$\pounds 286$	26	7,472	$\pounds 2,134,324$	$\pounds 1,237,629$	$\pounds 1,183,420$
		-2,355	$\pounds 185$	17	5,905	$\pounds 1,089,782.46$	$\pounds 381,142$	$\pounds 336,334$	$\pounds 286$	26	2,762	$\pounds788,795.79$	$\pounds 457,398$	$\pounds 415,299$
		-4,711	$\pounds 147.2$	13.2	4,711	$\pounds 693, 499.88$	$\pounds 128,199$	$\pounds 86,958$	$\pounds 286$	26	406	$\pounds 116,031.70$	$\pounds 67,283$	£31,239
		-5,061	$\pounds 142$	12.7	4,533	$\pounds 642, 198.27$	$\pounds 98,208$	$\pounds 57,429$	$\pounds 286$	26	56	$\pounds 16,010.29$	£9,284	$-\pounds 25,860$
Bonus No	No advantage	0	£221	17	7,083	£1,567,780	£717,819	£671,061	£284	21	5,111	£1,453,793	£840,494	£794,590
	Advantage	23,554	$\pounds 594$	45	19,000	$\pounds 11,280,885$	$\pounds9,000,922$	$\pounds 8,\!881,\!315$	$\pounds 284$	21	$28,\!665$	$\pounds 8,153,831$	$\pounds4,714,042$	$\pounds 4,\!617,\!888$
		4,711	$\pounds 296$	22	9,466	$\pounds 2,800,363$	$\pounds1,\!664,\!402$	$\pounds1,\!608,\!400$	$\pounds 284$	21	9,822	$\pounds 2,793,801$	$\pounds1,\!615,\!203$	$\pounds 1,559,250$
		2,355	$\pounds 258.6$	19	8,275	$\pounds 2,139,694$	$\pounds 1,146,733$	$\pounds 1,095,686$	$\pounds 284$	21	7,466	$\pounds 2,123,797$	$\pounds 1,227,848$	$\pounds 1,176,920$
		-2,355	$\pounds 184$	14	5,891	$\pounds 1,\!084,\!619.66$	£377,659	$\pounds 334,525$	$\pounds 284$	21	2,755	$\pounds783,789.51$	$\pounds 453,139$	£412,260
		-4,711	$\pounds 146.9$	11.0	4,700	$\pounds 690,214.45$	$\pounds 126,255$	$\pounds 86,078$	$\pounds 284$	21	400	$\pounds 113,785.69$	$\pounds 65,784$	£29,930
		-5,061	$\pounds 141$	10.6	4,522	$\pounds 639,155.89$	$\pounds 96,456$	$\pounds 56,662$	$\pounds 284$	21	50	$\pounds 14,\!174.66$	$\pounds 8,\!195$	$-\pounds 26,911$

Table B.3.: Theoretical model: solution for d = 10

	Information of							Deleinere	oh om o					
bonus scheme	information st	ructure				December	4	Friding s	cheme			C		
						Decentralize	ea		Centralized					
		μ	р	Ε	q	R	GP	GP-SC	р	Ε	q	R	GP	GP-SC
Bonus Yes	No advantage	0	$\pounds 226$	30	7,229	$\pounds 1,\!632,\!950.34$	$\pounds765,504$	£715,807	$\pounds 294$	40	5,185	$\pounds 1,\!524,\!817$	£902,597	$\pounds 853,873$
	Advantage	23,554	$\pounds 606$	82	19,391	$\pounds 11,749,819$	$\pounds 9,422,950$	$\pounds 9,282,202$	$\pounds 294$	40	28,739	$\pounds 8,451,432$	$\pounds 5,002,722$	$\pounds 4,891,659$
		4,711	$\pounds 302$	41	9,661	$\pounds 2,916,771$	$\pounds 1,757,440$	$\pounds 1,\!696,\!189$	$\pounds 294$	40	9,896	$\pounds 2,910,140$	$\pounds 1,722,622$	$\pounds1,\!661,\!431$
		2,355	$\pounds 264$	36	8,445	$\pounds 2,228,639$	$\pounds 1,215,250$	$\pounds 1,160,192$	$\pounds 294$	40	7,541	$\pounds 2,217,478$	$\pounds 1,312,609$	$\pounds 1,257,652$
		-2,355	$\pounds 188$	25	6,013	$\pounds 1,129,706$	$\pounds408,202$	$\pounds 363,034$	$\pounds 294$	40	2,830	$\pounds 832,155$	$\pounds 492,584$	$\pounds 450,095$
		-4,711	$\pounds 150$	20	4,796	£718,906	$\pounds 143,344$	£101,874	$\pounds 294$	40	474	$\pounds 139,494$	£82,572	$\pounds 46,316$
		-5,061	$\pounds 144$	19	4,616	$\pounds 665,725$	$\pounds 111,860$	$\pounds70,869$	$\pounds 294$	40	124	$\pounds 36,514$	$\pounds 21,\!614$	$-\pounds13,715$
Bonus No	No advantage	0	$\pounds 225$	25	7,190	£1,615,309	£752,561	£705,446	$\pounds 291$	33	5,168	$\pounds 1,505,714$	£885,614	£839,321
	Advantage	$23,\!554$	$\pounds 603$	68	19,286	$\pounds 11,\!622,\!883$	$\pounds 9,308,618$	$\pounds 9,\!186,\!446$	$\pounds 291$	33	28,722	$\pounds 8,368,931$	$\pounds 4,922,341$	$\pounds 4,\!824,\!574$
		4,711	$\pounds 300$	34	9,609	$\pounds 2,885,261$	$\pounds 1,732,209$	$\pounds 1,675,570$	$\pounds 291$	33	9,878	$\pounds 2,878,358$	$\pounds 1,692,959$	$\pounds 1,636,371$
		2,355	$\pounds 262$	30	8,399	$\pounds 2,204,562$	$\pounds 1,\!196,\!662$	$\pounds 1,145,128$	$\pounds 291$	33	7,523	$\pounds 2,\!192,\!036$	$\pounds 1,289,286$	$\pounds 1,237,846$
		-2,355	$\pounds 187$	21	5,980	$\pounds 1,117,502$	$\pounds 399,905$	£356,524	$\pounds 291$	33	2,812	$\pounds 819,393$	£481,941	$\pounds 440,795$
		-4,711	$\pounds 149$	17	4,770	£711,139	$\pounds 138,695$	$\pounds 98,361$	$\pounds 291$	33	457	£133,071	£78,268	£42,270
		-5,061	$\pounds 143$	16	4,591	$\pounds 658,533$	$\pounds 107,\!668$	$\pounds 67,729$	$\pounds 291$	33	107	$\pounds 31,034$	$\pounds 18,253$	$-\pounds 16,980$

Table B.4.: Theoretical model: solution for d = 15

Bonus scheme	Information st	ructure						Pricing s	cheme					
Bonus scheme	mormation st	ructure				Decentraliza	d	1 fields c	cheme			Centralize	d	
		μ	р	Е	q	R	GP-SC	р	Е	q	R	GP	GP-SC	
Bonus Yes	No advantage	0	£232	42	7,417	£1,719,223	£829,157	£778,684	£307	55	5,281	£1,621,247	£987,584	£937,992
	Advantage	23,554	$\pounds 622$	112	19,896	$\pounds 12,370,588$	£9,983,044	$\pounds 9,836,709$	£307	55	28,835	$\pounds 8,852,909$	$\pounds 5,392,757$	$\pounds 5,278,081$
		4,711	$\pounds 310$	56	9,913	$\pounds 3,070,871$	£1,881,309	$\pounds 1,818,671$	£307	55	9,991	£3,067,579	$\pounds 1,868,618$	£1,806,010
		2,355	$\pounds 271$	49	8,665	$\pounds 2,346,383$	$\pounds 1,306,569$	$\pounds 1,250,451$	$\pounds 307$	55	7,636	$\pounds 2,344,413$	$\pounds 1,428,101$	$\pounds 1,372,001$
		-2,355	$\pounds 193$	35	6,169	$\pounds 1,189,391$	$\pounds 449,073$	$\pounds403,368$	£307	55	2,925	£898,080	£547,066	£503,984
		-4,711	$\pounds 154$	28	4,921	$\pounds756,887$	$\pounds 166,317$	$\pounds 124,505$	£307	55	570	£174,914	$\pounds106,549$	$\pounds 69,975$
		-5,061	$\pounds 148$	27	4,736	$\pounds700,897$	$\pounds 132,\!589$	$\pounds 91,281$	$\pounds 307$	55	220	$\pounds 67,399$	$\pounds41,\!056$	$\pounds 5,\!450$
Bonus No	No advantage	0	$\pounds 230$	34	7,344	$\pounds 1,685,568$	£804,257	$\pounds756,615$	$\pounds 302$	45	5,247	£1,583,671	£954,049	£907,172
	Advantage	23,554	$\pounds 616$	92	19,701	$\pounds 12,128,429$	$\pounds 9,764,369$	$\pounds9,638,406$	$\pounds 302$	45	28,801	$\pounds 8,693,063$	$\pounds 5,236,951$	$\pounds 5,136,753$
		4,711	$\pounds 307$	46	9,816	$\pounds 3,010,757$	$\pounds 1,832,896$	$\pounds 1,775,316$	$\pounds 302$	45	9,958	$\pounds 3,005,550$	$\pounds 1,810,630$	$\pounds 1,753,088$
		2,355	$\pounds 268$	40	8,580	$\pounds 2,300,451$	$\pounds 1,270,865$	$\pounds 1,218,612$	$\pounds 302$	45	7,602	$\pounds 2,294,611$	$\pounds 1,382,339$	$\pounds 1,330,130$
		-2,355	$\pounds 191$	29	6,109	$\pounds 1,166,108$	£433,072	£389,326	$\pounds 302$	45	2,891	$\pounds 872,732$	$\pounds 525,759$	£484,213
		-4,711	$\pounds 152$	23	4,873	£742,071	$\pounds 157,309$	$\pounds 116,744$	$\pounds 302$	45	536	$\pounds 161,793$	$\pounds 97,469$	$\pounds 61,255$
		-5,061	$\pounds 147$	22	$4,\!689$	$\pounds 687,176$	$\pounds 124,459$	$\pounds 84,305$	$\pounds 302$	45	186	$\pounds 56,096$	£33,794	$-\pounds1,\!627$

Table B.5.: Theoretical model: solution for d = 20

Bonus scheme	Information st	ructure						Pricing s	cheme					
						Decentraliz	zed					Centralize	d	
		μ	р	Е	q	R	GP	GP-SC	р	Е	q	R	GP	GP-SC
Bonus Yes	No advantage	0	$\pounds 250$	68	8,014	$\pounds 2,007,162$	$\pounds 1,045,445$	£992,380	$\pounds 353$	95	5,553	$\pounds 1,962,599$	$\pounds 1,296,242$	$\pounds 1,243,579$
	Advantage	23,554	$\pounds 672$	181	$21,\!498$	$\pounds 14,\!442,\!441$	$\pounds 11,862,698$	$\pounds 11,697,716$	$\pounds 353$	95	29,107	$\pounds 10,\!287,\!363$	$\pounds 6,794,517$	$\pounds 6,666,930$
		4,711	£335	90	10,711	$\pounds 3,585,187$	$\pounds 2,299,864$	$\pounds 2,232,598$	$\pounds 353$	95	10,264	$\pounds 3,\!627,\!552$	$\pounds 2,395,897$	$\pounds 2,328,249$
		2,355	$\pounds 293$	79	9,363	$\pounds 2,739,360$	$\pounds 1,615,840$	$\pounds 1,556,186$	$\pounds 353$	95	7,908	$\pounds 2,795,076$	$\pounds 1,846,070$	$\pounds 1,785,914$
		-2,355	$\pounds 208$	56	6,666	$\pounds 1,388,593$	$\pounds 588,678$	$\pounds 541,181$	$\pounds 353$	95	3,198	$\pounds 1,130,123$	$\pounds746,415$	$\pounds701,244$
		-4,711	$\pounds 166$	45	5,318	£883,652	$\pounds 245,540$	£202,587	$\pounds 353$	95	842	$\pounds 297,647$	$\pounds 196,587$	$\pounds 158,908$
		-5,061	$\pounds 160$	43	5,117	$\pounds 818,284$	$\pounds 204,228$	$\pounds 161,863$	$\pounds 353$	95	492	$\pounds 173,880$	$\pounds 114,\!843$	$\pounds 78,278$
Bonus No	No advantage	0	$\pounds 245$	55	7,825	$\pounds 1,913,617$	$\pounds974,578$	£925,226	£338	76	5,474	$\pounds 1,848,315$	$\pounds 1,\!191,\!489$	$\pounds 1,142,627$
	Advantage	23,554	$\pounds 656$	148	20,991	$\pounds 13,769,344$	$\pounds 11,250,433$	$\pounds 11,112,163$	£338	76	29,028	$\pounds 9,802,083$	$\pounds 6,318,767$	$\pounds 6,210,252$
		4,711	$\pounds 327$	74	$10,\!458$	$\pounds 3,418,098$	$\pounds 2,163,084$	$\pounds 2,102,448$	£338	76	10,184	$\pounds 3,439,069$	$\pounds 2,216,945$	$\pounds 2,156,152$
		2,355	$\pounds 286$	64	9,142	$\pounds 2,\!611,\!691$	$\pounds 1,514,664$	$\pounds 1,460,077$	£338	76	7,829	$\pounds 2,\!643,\!692$	$\pounds 1,704,217$	$\pounds 1,649,389$
		-2,355	$\pounds 203$	46	6,509	$\pounds 1,323,877$	$\pounds 542,825$	£497,896	£338	76	3,118	$\pounds 1,052,938$	$\pounds 678,761$	$\pounds 635,864$
		-4,711	$\pounds 162$	37	5,192	$\pounds 842,469$	$\pounds 219,404$	£178,086	£338	76	763	$\pounds 257,561$	$\pounds 166,033$	$\pounds 129,101$
		-5,061	$\pounds 156$	35	4,996	$\pounds780,148$	$\pounds 180,571$	$\pounds 139,720$	$\pounds 338$	76	413	$\pounds 139,311$	$\pounds 89,805$	$\pounds 53,760$

Table B.6.: Theoretical model: solution for d = 30

D	sahama Information structure Prising sahama														
Bonus scheme	Information st	ructure				D		Pricing	scheme	9		<i>a</i>			
						Decentraliz	Decentralized				Centralized				
		μ	р	Е	q	R	GP	GP-SC	р	Е	q	R	GP	GP-SC	
Bonus Yes	No advantage	0	$\pounds 282$	102	9,032	$\pounds 2,549,428$	$\pounds 1,465,557$	$\pounds1,407,612$	$\pounds 458$	165	5,934	$\pounds 2,719,574.85$	£2,007,446	$\pounds 1,947,970$	
	Advantage	$23,\!554$	£757	273	24,228	$\pounds 18,344,289$	$\pounds 15,\!436,\!876$	$\pounds 15,236,778$	$\pounds 458$	165	29,488	$\pounds 13,513,754.36$	$\pounds 9,975,136$	$\pounds 9,818,512$	
		4,711	£377	136	12,071	$\pounds 4,553,780$	£3,105,201	$\pounds 3,029,217$	$\pounds 458$	165	$10,\!645$	$\pounds 4,878,410.75$	$\pounds 3,600,984$	$\pounds 3,522,078$	
		2,355	£330	119	10,552	$\pounds 3,479,440$	$\pounds 2,213,215$	$\pounds 2,146,900$	$\pounds 458$	165	8,290	$\pounds 3,798,992.80$	$\pounds 2,804,215$	$\pounds 2,735,024$	
		-2,355	$\pounds 235$	85	7,513	$\pounds 1,763,742$	£862,226	$\pounds 811,352$	$\pounds 458$	165	3,579	$\pounds 1,640,156.89$	$\pounds 1,210,677$	$\pounds 1,160,915$	
		-4,711	$\pounds 187$	67	5,993	$\pounds 1,122,385$	£403,222	$\pounds 358,121$	$\pounds 458$	165	1,224	$\pounds 560,738.94$	£413,908	$\pounds 373,861$	
		-5,061	$\pounds 180$	65	5,767	$\pounds 1,\!039,\!356$	$\pounds 347,305$	$\pounds 302,951$	$\pounds 458$	165	873	$\pounds 400,259.36$	$\pounds 295,\!450$	$\pounds 256,848$	
Bonus No	No advantage	0	£269	81	8,615	$\pounds 2,319,527$	$\pounds 1,285,680$	$\pounds 1,233,284$	£410	123	5,791	$\pounds 2,376,907.73$	£1,681,996	$\pounds 1,629,169$	
	Advantage	$23,\!554$	$\pounds722$	217	23,110	$\pounds 16,\!690,\!046$	$\pounds 13,\!916,\!822$	$\pounds 13,756,647$	$\pounds 410$	123	29,345	$\pounds 12,044,759.96$	$\pounds 8,523,359$	$\pounds 8,398,023$	
		4,711	$\pounds 360$	108	11,514	$\pounds 4,143,132$	$\pounds 2,761,410$	$\pounds 2,695,336$	$\pounds 410$	123	10,502	$\pounds 4,310,478.17$	$\pounds 3,\!050,\!268$	$\pounds 2,982,940$	
		2,355	$\pounds 315$	94	10,065	$\pounds 3,165,673$	$\pounds 1,957,889$	$\pounds 1,899,147$	$\pounds 410$	123	8,146	$\pounds 3,343,692.95$	$\pounds 2,366,132$	$\pounds 2,306,054$	
		-2,355	$\pounds 224$	67	7,166	$\pounds 1,\!604,\!693$	£744,784	$\pounds 697,749$	$\pounds 410$	123	3,436	$\pounds 1,410,122.50$	$\pounds 997,860$	$\pounds 952,284$	
		-4,711	$\pounds 179$	54	5,716	$\pounds 1,021,171$	£335,200	£292,542	$\pounds 410$	123	1,080	$\pounds 443,337.28$	$\pounds 313,723$	$\pounds 275,398$	
		-5,061	$\pounds 172$	52	5,501	$\pounds 945,\!630$	$\pounds 285,519$	$\pounds 243,\!427$	$\pounds 410$	123	730	$\pounds 299,\!603.07$	$\pounds 212,\!011$	$\pounds 174,764$	

Table B.7.: Theoretical model: solution for d = 40

Popus sahomo	Information at	ruoturo						Driging	cohomo					
Donus scheme	mormation st	fucture				Docontroli	rod	Theng	scheme			Controliz	d	
				-		Decentran	aeu ar	~~ ~~		-		Centranze	su ore	~~ ~~
		μ	р	Е	q	R	GP	GP-SC	р	Е	q	R	GP	GP-SC
Bonus Yes	No advantage	0	£337	152	10,795	$\pounds 3,\!641,\!748$	$\pounds 2,346,326$	$\pounds 2,278,550$	£797	359	6,425	$\pounds 5,\!123,\!071$	$\pounds 4,352,092$	$\pounds 4,270,984$
	Advantage	$23,\!554$	$\pounds905$	407	28,957	$\pounds 26,\!204,\!027$	$\pounds 22,729,144$	$\pounds 22,458,307$	£797	359	29,979	$\pounds 23,904,788$	$\pounds20,307,320$	$\pounds 20,057,177$
		4,711	$\pounds 451$	203	14,428	$\pounds 6,504,879$	$\pounds 4,773,565$	$\pounds 4,680,021$	£797	359	11,136	$\pounds 8,879,414$	$\pounds7,543,137$	$\pounds7,428,223$
		2,355	$\pounds 394$	177	$12,\!611$	$\pounds 4,970,231$	$\pounds 3,456,863$	$\pounds 3,377,131$	£797	359	8,780	$\pounds7,001,242$	$\pounds 5,947,615$	$\pounds 5,849,604$
		-2,355	$\pounds 281$	126	8,979	$\pounds 2,519,430$	$\pounds 1,441,955$	$\pounds 1,384,280$	£797	359	4,069	$\pounds 3,244,899$	$\pounds 2,756,569$	$\pounds 2,692,365$
		-4,711	$\pounds 224$	101	7,163	$\pounds 1,603,278$	£743,749	$\pounds 694,319$	£797	359	1,714	$\pounds 1,366,727$	$\pounds 1,161,046$	$\pounds 1,113,746$
		-5,061	$\pounds 215$	97	$6,\!893$	$\pounds 1,\!484,\!676$	$\pounds 657,549$	$\pounds 609,187$	$\pounds797$	359	1,364	$\pounds 1,\!087,\!495$	$\pounds 923,836$	£879,049
Bonus No	No advantage	0	£309	116	9,901	£3,063,154	£1,875,088	£1,817,114	£589	221	6,199	£3,649,690	£2,905,810	£2,843,438
	Advantage	$23,\!554$	£830	311	26,558	$\pounds 22,040,785$	$\pounds 18,853,876$	$\pounds 18,\!653,\!570$	$\pounds 589$	221	29,753	$\pounds 17,517,279$	$\pounds 13,\!946,\!910$	$\pounds 13,780,530$
		4,711	$\pounds 413$	155	13,232	$\pounds 5,471,397$	$\pounds 3,883,563$	$\pounds 3,807,527$	$\pounds 589$	221	10,910	$\pounds 6,423,207$	$\pounds 5,\!114,\!030$	$\pounds 5,030,856$
		2,355	$\pounds 361$	136	11,566	$\pounds 4,180,571$	$\pounds 2,792,620$	$\pounds 2,726,266$	$\pounds 589$	221	8,554	£5,036,448	$\pounds4,009,920$	$\pounds 3,937,147$
		-2,355	$\pounds 257$	97	8,235	$\pounds 2,119,148$	$\pounds 1,130,966$	$\pounds 1,080,073$	$\pounds 589$	221	$3,\!844$	$\pounds 2,262,931$	$\pounds 1,801,701$	$\pounds 1,749,729$
		-4,711	$\pounds 205$	77	6,569	$\pounds 1,348,553$	$\pounds 560,255$	$\pounds 515,141$	$\pounds 589$	221	$1,\!488$	$\pounds 876,172$	$\pounds 697,591$	$\pounds 656,019$
		-5,061	$\pounds 198$	74	6,322	$\pounds 1,\!248,\!794$	$\pounds 490,213$	$\pounds 445,847$	$\pounds 589$	221	$1,\!138$	$\pounds 669,999$	$\pounds 533,\!440$	$\pounds 493,415$

Table B.8.: Theoretical model: solution for d = 50

B.3. Data appendix

Type of Regression Dependent Variable	(1) OLS Qrt Revenues	(2) OLS Qrt Revenues	(3) OLS Qrt Revenues	(4) OLS Qrt Revenues	(5) OLS Qrt Revenues	(6) OLS Qrt Revenues	(7) OLS Qrt Revenues
Linear Time Trend	Yes	No	No	No	No	No	No
Quadratic Time Trend	No	Yes	No	No	No	No	No
Quarter dummy	No	No	Yes	No	Yes	No	Yes
Year dummy	No	No	No	Yes	Yes	No	Yes
Employee dummy	No	No	No	No	No	Yes	Yes
Observations	254	254	254	254	254	254	254
R-squared	0.0075	0.0105	0.0429	0.0698	0.1125	0.3988	0.5160

Table B.9.: ANCOVA

Notes. Revenues are deflated using country-specific CPI with 2005 as base year.
	Percentiles	Smallest		
1%	-39,697	-39,697		
5%	$57,\!057$	20,225		
10%	$110,\!877$	27,592	Obs	64
25%	$241,\!573$	$57,\!057$	Sum of Wgt.	64
50%	407,286		Mean	500,355
		Largest	Std. Dev.	403,359
75%	690,259	1,041,000		
90%	872,090	1,212,018		
95%	1,041,000	$1,\!368,\!864$	Skewness	2
99%	$2,\!554,\!629$	$2,\!554,\!629$	Kurtosis	12

Table B.10.: Percentiles of annual Gross Profits-SC

Notes. Gross Profits-SC is annual, individual-level data. Gross Profits-SC equals annual gross profits minus bonuses, commissions, and annual salaries, which occur to be negative once. Gross Profits-SC is deflated using country-specific CPI with 2005 as base year.

	(1)	(2)	(3)	(4)	(5)
Sample	Pooled	Pooled	Pooled	Pooled	Pooled
Type of Regression	OLS	OLS	OLS	OLS	OLS
Dependent Variable	Qrt Revenues	Qrt Gross Profits	Qrt Gross Profits-SC	Qrt Quantity	Qrt Unit Price
Bonus	237,046***	101,992***	105,058***	2,010***	-86
	(41, 589)	(20, 641)	(20, 614)	(468)	(60)
PD	-162,556	-11,454	-8,301	-2,019	86
	(155, 823)	(66,794)	(64,905)	(1,442)	(76)
Bonus*PD	120,736*	48,644*	44,627	-222	59
	(67, 693)	(27, 906)	(27, 124)	(685)	(52)
Year dummy	Yes	Yes	Yes	Yes	Yes
Quarter dummy	Yes	Yes	Yes	Yes	Yes
Employee dummy	Yes	Yes	Yes	Yes	Yes
Other controls	No	No	No	No	No
Clustering	Employee	Employee	Employee	Employee	Employee
Observations	254	254	254	254	241

Table B.11.: The effect of incentives and price delegation controlling for employee FEs

Notes. All outcomes are deflated using country-specific CPI with 2005 as base year. All specifica-

tions are clustered by employee to allow for serial correlation in individual error terms.

	(1)	(2)	(3)	(4)	(5)
Sample	Pooled	Pooled	Pooled	Pooled	Pooled
Type of Regression	OLS	OLS	OLS	OLS	OLS
Dependent Variable	Qrt Revenues	Qrt Gross Profits	Qrt Gross Profits-SC	Qrt Quantity	Qrt Unit Price
Bonus	179,982**	35,597	37,946*	1,546**	-106**
	(77,703)	(21, 123)	(20, 646)	(616)	(45)
PD	28,007	1,626	-2,495	568	-62
	(59, 232)	(21, 919)	(21, 612)	(550)	(46)
Bonus*PD	-62,291	20,197	23,662	-1,084	164***
	(90, 110)	(28,051)	(26,714)	(679)	(53)
Year dummy	Yes	Yes	Yes	Yes	Yes
Quarter dummy	Yes	Yes	Yes	Yes	Yes
Employee dummy	No	No	No	No	No
Other controls	Yes	Yes	Yes	Yes	Yes
Clustering	Employee	Employee	Employee	Employee	Employee
Observations	254	254	254	254	241

Table B.12.: The effect of incentives and price delegation controlling for target on revenues

Notes. Other controls include the annual target on revenues. All outcomes are deflated using country-specific CPI with 2005 as base year. All specifications are clustered by employee to allow for serial correlation in individual error terms.

	(1)	(2)	(3)	(4)	(5)
Sample	Pooled	Pooled	Pooled	Pooled	Pooled
Type of Regression	OLS	OLS	OLS	OLS	OLS
Dependent Variable	Qrt Revenues	Qrt Gross Profits	Qrt Gross Profits-SC	Qrt Quantity	Qrt Unit Price
Bonus	381,399***	123,911***	120,234***	2,331***	-9
	(111, 505)	(41, 896)	(40,738)	(670)	(39)
PD	-96,338	-53,042	-53,493	81	-115
	(104, 631)	(46, 328)	(45,512)	(752)	(78)
Bonus*PD	13,024	50,822	51,253	-813	202**
	(129, 318)	(50, 284)	(48, 963)	(870)	(80)
Year dummy	Yes	Yes	Yes	Yes	Yes
Quarter dummy	Yes	Yes	Yes	Yes	Yes
Employee dummy	No	No	No	No	No
Other controls	Yes	Yes	Yes	Yes	Yes
Clustering	Employee	Employee	Employee	Employee	Employee
Observations	254	254	254	254	241

Table B.13.: The effect of incentives and price delegation controlling for salary

Notes. Other controls include the annual salary. All outcomes are deflated using country-specific CPI with 2005 as base year. All specifications are clustered by employee to allow for serial correlation in individual error terms.

Sex	Centralized pricing	Delegated pricing	Total
Male	120	96	216
Female	28	12	40
Total	148	108	256

Table B.14.: Gender distribution by pricing group

Variable	Centralized pricing Mean	Delegated pricing Mean
Target on revenues	2,020,750	$1,\!549,\!152$
Annual salary	$33,\!504$	$39,\!990$
Expected annual commissions	11,091	$11,\!318$
Commissions-total compensation ratio	23.03%	21.17%
Min commission rate	0.59%	0.70%
Max commission rate	1.95%	2.38%
Difference between min and max commission rates	1.36%	1.68%
Expected annual commission rate	0.73%	0.86%
Observations	148	106

Table B.15.: Differences in ex-ante characteristics of employment contracts

Notes. Target on revenues, annual salary and expected annual commissions are in 2005 country-currency. Expected commissions are computed plugging expected revenues into compensation scheme. Expected revenues are provided by the firm.

Variable	Centralized pricing Mean	Delegated pricing Mean
Bonus dummy Observations	$59\% \\ 148$	48% 106

Table B.16.: Distribution of bonus scheme across pricing groups

Table B.17.: Size of expected amount of bonus by pricing groups

Variable	Centralized pricing Mean	Delegated pricing Mean
Expected annual bonus Observations	$3,526 \\ 88\dagger$	$3,761 \\ 51\dagger$

Notes. Expected annual bonuses are in 2005 countrycurrency. Expected bonuses are computed plugging expected revenues into compensation scheme. Expected revenues are provided by the firm. † I only include employees that are offered the bonus scheme.

	(1)	(2)	(3)	(4)	(5)	(6)
Sample	Pooled	Pooled	Centralized	Centralized	Delegated	Delegated
Type of Regression	IV: 1st	IV: 2nd	IV: 1st	IV: 2nd	IV: 1st	IV: 2nd
Dependent Variable	$\log(1+\text{Price})$	$\log(1+\text{Quantity})$	$\log(1+\text{Price})$	$\log(1+\text{Quantity})$	$\log(1+\text{Price})$	$\log(1+\text{Quantity})$
Qrt Unit Cost	0.0063***		0.0060***		0.0064^{***}	
	(0.0006)		(0.0008)		(0.0009)	
$\log(1+\text{Qrt Unit Price})$		-1.1248***		-2.0245***		-0.4797***
		(0.2958)		(0.5005)		(0.1473)
F-Stat	104		44		35	
Year dummy	Yes	Yes	Yes	Yes	Yes	Yes
Quarter dummy	Yes	Yes	Yes	Yes	Yes	Yes
Employee dummy	Yes	Yes	Yes	Yes	Yes	Yes
Clustering	Employee	Employee	Employee	Employee	Employee	Employee
Observations	241	241	142	142	99	99

Table B.18.: Price-elasticity of demand

Notes. All variables are in 2005 country-currency.

Channel	Mean	Std.	Dev.	Obs
Centralized	2.90866	1.06	6057	142
Delegated	2.55374	1.11	.427	99
Notes. El	asticities	are	comp	uted
	-			
structurally	pluggir	ng e	quilib	rium
structurally prices and n	pluggir 1arginal c	ng e costs i	quilib nto st	rium atic-

Table B.19.: Price-elasticity for a static monopolist by pricing group

Table B.18 provides the least square regression of the logarithm of quantity on the logarithm of price for the two pricing groups separately. As price is usually endogenous in a supply-demand system¹, I use the unit cost as an instrument for price. The estimates in Table B.18 clearly show a marked difference in the price elasticity across the two pricing groups. A price elasticity that is between -1 and 0 can be problematic in a static context². Alternatively, we can reverse-engineer the price elasticity from the first-order conditions of profit-maximizing static monopolist. For a static monopolist, standard first order conditions would yield:

$$(p - mc)/p = 1/\epsilon \tag{B.1}$$

Table B.19 shows the price-elasticity of demand obtained by plugging prices and unit costs into the first-order condition of a static monopolist and than averaging across individuals for the two pricing groups. A two-sample t-test rejects the null of equal mean-elasticity across pricing groups against the alternative hypothesis that the elasticity is larger for the group of observations where price is centralized at any level of significance³.

¹For example, see Hamilton [1994], chapter 9 for an explanation of the simultaneous equation bias in estimating demand elasticity.

²A price elasticity that is between -1 and 0 can be rationalized by dynamic monopoly pricing in the presence of learning by doing. In this chapter, we do not model dynamic pricing.

³This is just an attempt to test whether price elasticity differs across the two pricing groups.

Sample Type of Regression Dependent Variable	(1) Pooled BiW RREG Qrt Revenues	(2) Pooled BiW RREG Qrt Gross Profits	(3) Pooled BiW RREG Qrt Gross Profits-SC	(4) Pooled BiW RREG Qrt Quantity	(5) Pooled BiW RREG Qrt Unit Price
Bonus	254,125***	86,807***	89,042***	892***	-2
	(64, 567)	(24, 223)	(23, 476)	(289)	(20)
PD	-20,198	-3,044	-6,998	-536**	16
	(49,293)	(18, 493)	(17, 922)	(221)	(15)
Bonus*PD	-58,468	$2,\!688$	6,878	-293	31
	(64,757)	(24, 295)	(23,545)	(290)	(20)
Year dummy	Yes	Yes	Yes	Yes	Yes
Quarter dummy	Yes	Yes	Yes	Yes	Yes
Employee dummy	No	No	No	No	No
Other controls	Yes	Yes	Yes	Yes	Yes
Clustering	BiW RREG	BiW RREG	BiW RREG	BiW RREG	BiW RREG
Observations	254	254	254	254	241

Table B.20.: Impact on outcomes using robust regression with bi-weighting

Notes. In the robust regression with bi-weighting, all cases with a non-zero residual get down-weighted at least a little. All outcomes are deflated using country-specific CPI with 2005 as base year. All specifications are clustered by employee to allow for serial correlation in individual error terms. The list of controls includes the price-delegation dummy (PD) and its interaction with the time dummy.

(1) Pooled Hubert M Qrt Revenues	(2) Pooled Hubert M Qrt Gross Profits	(3) Pooled Hubert M Qrt Gross Profits-SC	(4) Pooled Hubert M Qrt Quantity	(5) Pooled Hubert M Qrt Unit Price
245,729***	79,722***	81,605***	$1,195^{***}$	-20
(62, 378)	(19,760)	(19,021)	(299)	(21)
-21,833	-8,853	-12,778	-209	1
(46,059)	(18, 861)	(17,780)	(232)	(17)
-52,001	$7,\!698$	11,877	-580*	64**
(62, 220)	(25,009)	(23, 977)	(303)	(27)
Yes	Yes	Yes	Yes	Yes
Yes	Yes	Yes	Yes	Yes
No	No	No	No	No
Yes	Yes	Yes	Yes	Yes
Hubert M	Hubert M	Hubert M	Hubert M	Hubert M
254	254	254	254	241
	(1) Pooled Hubert M Qrt Revenues 245,729*** (62,378) -21,833 (46,059) -52,001 (62,220) Yes Yes No Yes Hubert M 254	(1) (2) Pooled Pooled Hubert M Hubert M Qrt Revenues Qrt Gross Profits 245,729*** 79,722*** (62,378) (19,760) -21,833 -8,853 (46,059) (18,861) -52,001 7,698 (62,220) (25,009) Yes Yes No No Yes Yes Hubert M Hubert M 254 254	(1) (2) (3) Pooled Pooled Pooled Hubert M Hubert M Hubert M Qrt Revenues Qrt Gross Profits Qrt Gross Profits-SC 245,729*** 79,722*** 81,605*** (62,378) (19,760) (19,021) -21,833 -8,853 -12,778 (46,059) (18,861) (17,780) -52,001 7,698 11,877 (62,220) (25,009) (23,977) Yes Yes Yes Yes Yes Yes No No No Yes Yes Yes Hubert M Hubert M Hubert M 254 254 254	

Table B.21.: Impact on outcomes using Hubert's robust estimator

Notes. Hubert's robust estimator weights observations depending on the size of the regression residuals (i.e., the larger the size, the smaller the weight). All outcomes are deflated using country-specific CPI with 2005 as base year. All specifications are clustered by employee to allow for serial correlation in individual error terms. The list of controls includes the price-delegation dummy (PD) and its interaction with the time dummy.

Table B.20 shows the estimates of our main specification using robust-regression with bi-weighting and Table B.20 shows the estimates of our main specification using Hubert's robust M-estimator. The intuition behind these estimators is to weight *aberrant* observations in the regression less in order to preserve the importance of the *normal* observations.⁴ Both estimators produce results similar to those presented in our main specification in Table 3.10.

⁴See Blankmeyer [2006] for a detailed description of the high-breakdown robust estimators.

Sample Type of Regression Dependent Variable	(1) Pooled OLS Qrt Revenues	(2) Pooled OLS Qrt Gross Profits	(3) Pooled OLS Qrt Gross Profits-SC	(4) Pooled OLS Qrt Quantity	(5) Pooled OLS Qrt Unit Price
Bonus	$1,028,052^{**}$	291,183	117,373***	7,529**	-25
	(476, 602)	(170, 404)	(41, 251)	(2,931)	(30)
PD	-80,240	-87,442	25,276	1,422	8
	(361, 612)	(167, 752)	(47, 517)	(2,636)	(32)
Bonus*PD	-290,433	63,233	9,561	-4,464	65
	(523, 754)	(205, 252)	(63, 242)	(3,280)	(39)
Year dummy	Yes	Yes	Yes	Yes	Yes
Quarter dummy	No	No	No	No	No
Employee dummy	No	No	No	No	No
Other controls	No	No	No	No	No
Clustering	Employee	Employee	Employee	Employee	Employee
Observations	64	64	64	64	64

Table B.22.: Impact on annual outcomes

Notes. All outcomes are deflated using country-specific CPI with 2005 as base year. All specifications are clustered by employee to allow for serial correlation in individual error terms. Other controls include the price-delegation dummy (PD) and its interaction with the year dummy.

	(1)	(2)	(3)	(4)	(5)
Sample	Pooled	Pooled	Pooled	Pooled	Pooled
Type of Regression	OLS	OLS	OLS	OLS	OLS
Dependent Variable	Qrt Revenues	Qrt Gross Profits	Qrt Gross Profits-SC	Qrt Quantity	Qrt Unit Price
Bonus	257,013**	72,796*	74,673*	1,882**	-77*
	(111, 350)	(39, 812)	(38,768)	(685)	(44)
PD	-20,060	-21,860	-25,584	356	-74
	(84, 484)	(39, 192)	(38, 475)	(616)	(63)
Bonus*PD	-72,608	$15,\!808$	$19,\!681$	-1,116	164^{**}
	(122, 366)	(47, 953)	(46, 197)	(766)	(64)
Year dummy	Yes	Yes	Yes	Yes	Yes
Quarter dummy	Yes	Yes	Yes	Yes	Yes
Employee dummy	No	No	No	No	No
Other controls	No	No	No	No	No
Clustering	Employee	Employee	Employee	Employee	Employee
Observations	254	254	254	254	241

Table B.23.: Impact on outcomes including observations with negative revenues

Notes. All outcomes are deflated using country-specific CPI with 2005 as base year. All specifications are clustered by employee to allow for serial correlation in individual error terms. Other controls include the price-delegation dummy (PD) and its interaction with the time dummy.

Appendix C. Appendix to chapter 4

C.1. Geographic markets

me	Province Na	Province Acronym	Province Code	Region Name	Region Code
RIA	ALESSANDI	AL	001	PIEMONTE	01
STI	А	AT	002	PIEMONTE	01
LA	BIEI	BI	003	PIEMONTE	01
ΕO	CUN	CN	004	PIEMONTE	01
RA	NOVA	NO	005	PIEMONTE	01
NO	TORI	TO	006	PIEMONTE	01
NIA	VERBAI	VB	007	PIEMONTE	01
LLI	VERCE	VC	008	PIEMONTE	01
TA	AOS	AO	009	VALLE D'AOSTA	02
MO	BERGA	BG	010	LOMBARDIA	03
CIA	BRES	BS	011	LOMBARDIA	03
MO	CO	CO	012	LOMBARDIA	03
NA	CREMO	CR	013	LOMBARDIA	03
CO	LEC	LC	014	LOMBARDIA	03
DDI	LO	LO	015	LOMBARDIA	03
VA	MANTC	MN	016	LOMBARDIA	03
NO	MILA	MI	017	LOMBARDIA	03
/IA	PA	PV	018	LOMBARDIA	03
NIO	SONDI	SO	019	LOMBARDIA	03
SE	VARI	VA	020	LOMBARDIA	03
VA	GENC	GE	021	LIGURIA	04
RIA	IMPE	IM	022	LIGURIA	04
ZIA	LA SPE	SP	023	LIGURIA	04
NA	SAVO	SV	024	LIGURIA	04
ΕN	BOLZANO-BOZ	BZ	025	TRENTINO-ALTO ADIGE	05
TO	TREN	TN	026	TRENTINO-ALTO ADIGE	05
NO	BELLU	BL	027	VENETO	06
VA	PADC	PD	028	VENETO	06
GO	ROVI	RO	029	VENETO	06
SO	TREV	TV	030	VENETO	06
ZIA	VENE	VE	031	VENETO	06
NA	VERO	VR	032	VENETO	06
ZA	VICEN	VI	033	VENETO	06
ZIA	GORI	GO	034	FRIULI VENEZIA GIULIA	07
TE	TRIES	TS	035	FRIULI VENEZIA GIULIA	07
NE	UDI	UD	036	FRIULI VENEZIA GIULIA	07
NE	PORDENC	PN	037	FRIULI VENEZIA GIULIA	07
NA	BOLOG	BO	038	EMILIA-ROMAGNA	08
RA	FERRA	FE	039	EMILIA-ROMAGNA	08
LI'	FOI	FC	040	EMILIA-ROMAGNA	08
NA	MODE	MO	041	EMILIA-ROMAGNA	08
MA	PAR	PR	042	EMILIA-ROMAGNA	08
ZA	PIACEN	PC	043	EMILIA-ROMAGNA	08
NA	RAVEN	RA	044	EMILIA-ROMAGNA	08
LIA	REGGIO EMI	RE	045	EMILIA-ROMAGNA	08
INI	RIM	RN	046	EMILIA-ROMAGNA	08
ZO	AREZ	AR	047	TOSCANA	09
ΖE	FIREN	FI	048	TOSCANA	09
то	GROSSE	GR	049	TOSCANA	09
NO	LIVOR	LI	050	TOSCANA	09
CA	LUC	LU	051	TOSCANA	09
SA	MAS	MS	052	TOSCANA	09
ISA	P	PI	053	TOSCANA	09
DIA	PISTO	PT	054	TOSCANA	09
то	PRA	PO	055	TOSCANA	09
NA	SIE	SI	056	TOSCANA	09

Table C.1.: ISTAT classification of regions and provinces in Italy

Province Name	Province Acronym	Province Code	Region Name	Region Code
PERUGIA	PG	057	UMBRIA	10
TERNI	TR	058	UMBRIA	10
ANCONA	AN	059	MARCHE	11
ASCOLI PICENO	AP	060	MARCHE	11
MACERATA	MC	061	MARCHE	11
PESARO	PU	062	MARCHE	11
FROSINONE	\mathbf{FR}	063	LAZIO	12
LATINA	LT	064	LAZIO	12
RIETI	RI	065	LAZIO	12
ROMA	RM	066	LAZIO	12
VITERBO	VT	067	LAZIO	12
CHIETI	CH	068	ABRUZZO	13
L'AQUILA	AQ	069	ABRUZZO	13
PESCARA	PE	070	ABRUZZO	13
TERAMO	TE	071	ABRUZZO	13
CAMPOBASSO	CB	072	MOLISE	14
ISERNIA	IS	073	MOLISE	14
AVELLINO	AV	074	CAMPANIA	15
BENEVENTO	BN	075	CAMPANIA	15
CASERTA	CE	076	CAMPANIA	15
NAPOLI	NA	077	CAMPANIA	15
SALERNO	SA	078	CAMPANIA	15
BARI	BA	079	PUGLIA	16
BRINDISI	BR	080	PUGLIA	16
FOGGIA	FG	081	PUGLIA	16
LECCE	LE	082	PUGLIA	16
TARANTO	TA	083	PUGLIA	16
MATERA	MT	084	BASILICATA	17
POTENZA	PZ	085	BASILICATA	17
CATANZARO	CZ	086	CALABRIA	18
COSENZA	CS	087	CALABRIA	18
CROTONE	KR	088	CALABRIA	18
REGGIO CALABRIA	RC	089	CALABRIA	18
VIBO VALENTIA	VV	090	CALABRIA	18
AGRIGENTO	AG	091	SICILIA	19
CALTANISSETTA	CL	092	SICILIA	19
CATANIA	CT	093	SICILIA	19
ENNA	EN	094	SICILIA	19
MESSINA	ME	095	SICILIA	19
PALEBMO	PA	096	SICILIA	19
RAGUSA	BG	097	SICILIA	19
SIRACUSA	SB	098	SICILIA	19
TRAPAN	TP	099	SICILIA	19
CAGLIAR	CA	100	SARDEGNA	20
SASSARI	SS	100	SARDEGNA	20
NUORC	NU	102	SARDEGNA	20
1,0010	NO OD	101	CADDECNA	20

Table C.2.: ISTAT classification of regions and provinces in Italy

C.2. Sample selection

Variable	Obs	Mean	Std. Dev.	Min	Median	Max
Drawn	2,870,937	2,415	27,400	0	453	9,030,000
EBIT	$2,\!870,\!937$	411	27,565	-1,478,188	47	12,000,000
% EBITDA/Assets	$2,\!870,\!937$	6	97.2	-43,800	6.6	39,500
Granted	$2,\!870,\!937$	3,709	51,000	0	688	18,000,000
HHI	$2,\!870,\!937$	5,717	3,065	0	5,084	10,000
% Leverage	$2,\!870,\!937$	65.7	838.9	-220,300	71.5	482,067
% Liquidity	$2,\!870,\!937$	5.9	9.8	0	1.9	100
No banks	$2,\!870,\!937$	3.3	2.7	1	3	59
NPL	$2,\!870,\!937$	0.05	0.21	0	0	1
Revenues	$2,\!870,\!937$	8,111	151, 137	0	1,251	53,000,000
% ROA	$2,\!870,\!937$	1.8	123	-65,200	3.7	39,500
Total assets	$2,\!870,\!937$	$9,\!974$	$332,\!277$	1	$1,\!466$	84,000,000

Table C.3.: Original sample

Table C.4.: Selected sample

Variable	Obs	Mean	Std. Dev.	Min	Median	Max
Drawn	971,272	1,223	8,537	0	349	1,920,000
EBIT	971,272	127	3,143	-1,478,188	39	559,607
% EBITDA/Assets	$971,\!272$	6	24.4	-5,240	6.3	848
Granted	$971,\!272$	$1,\!686$	10,300	0	518	2,110,000
HHI	971,272	6,417	2,950	0	5,669	10,000
% Leverage	$971,\!272$	67.7	395.8	-69,300	74.1	138,900
% Liquidity	$971,\!272$	5.8	9.6	0	1.8	100
No banks	$971,\!272$	2.6	1.8	1	2	30
NPL	971,272	0.04	0.19	0	0	1
Revenues	$971,\!272$	$3,\!870$	41,375	0	949	5,779,045
% ROA	$971,\!272$	2.7	34	-9,203	3.7	10,851
Total assets	$971,\!272$	4,080	55,419	10	1,166	20,900,000

Notes. Drawn is the total amount of credit in thousands of euros used by a firm in a given half year, while Granted is the corresponding amount granted to the firm. EBIT is annual earnings before interests and taxes. EBIT, Revenues and Total assets are annual firm-level variables expressed in thousands of euros. EBITDA is annual earnings before interests, taxes, depreciation and amortization. EBITDA/Assets is EBITDA over total assets. Leverage is the annual fraction of a firm's cash holdings over its total assets. ROA is returns on assets. EBITDA/Assets, Leverage, Liquidity and ROA are annual firm-level indicators expressed as percentages. The number of banks is the number of lenders a firm borrow from in a given half year. HHI is the Herfindahl-Hirschman Index of a firm's loans portfolio across its lenders in a given half year. HHI can take values between 10,000/No banks when the firm borrows an equal amount from each of its lenders, and 10,000 when the firm only borrows from one banks. NPL is a binary variable that equals 1 if the firm is recorded as non-performing in a given half year and equals 0 otherwise.

C.3. Share of risky bonds



Figure C.1.: Share of risky bonds

The histogram is computed using bank - half-year as unit of observation, for the sample of Italian credit institutions for which we have complete information.

C.4. Validity of the instrumental variable

	(1)	(2)	(3)	(4)
Sample	Selected	Selected	Selected	Selected
Type of Regression	OLS	2SLS-IV	OLS	2SLS-IV
Dependent Variable	EBIT	EBIT	EBITDA/Assets	EBITDA/Assets
Expo. Funding Gap June 2010	1.9164**		0.4588	
	(7.6406)		(0.6581)	
Expo. Funding Gap June 2010*Crisis	7.7659		-2.5359**	
	(1.4894)		(1.0417)	
Expo. Funding Gap		2.5346^{**}		0.6067
		(1.0093)		(0.871)
Expo. Funding Gap*Crisis		1.2225		-3.5047**
		(2.1091)		(1.4385)
F-stat 1		8,531		8,531
F-stat 2 (interaction)		6,128		6,128
Firm dummy	No	No	No	No
Market6-Time dummy	Yes	Yes	Yes	Yes
No of Market6-Time dummies	132,708	132,708	132,708	132,708
Other controls	No	No	No	No
Clustering	Market3-Time	Market3-Time	Market3-Time	Market3-Time
Observations	971.272	971.272	971,272	971,272

Table C.5.: Testing validity of the instrumental variable I

Notes. The unit of observation is a firm - half-year pair. Expo. Funding Gap is the time-variant one-period lagged firm's exposure to its lenders funding gap, and Expo. Funding Gap*Crisis its interaction with the binary variable Crisis which equals 1 for after July 2011 and equals 0 before. Expo. Funding Gap is expressed as a fraction of the total assets (not as a percentage). EBIT is annual earnings before interests and taxes. EBIT, Revenues and Total assets are annual firm-level variables expressed in thousands of euros. EBITDA is annual earnings before interests, taxes, depreciation and amortization. EBITDA/Assets is EBITDA over total assets. Market6-Time dummy represent market-time fixed effects where the time is represented by an half year and a market is defined as a geography-product pair where the geography is a province while the product market is defined using the full 6-digit ATECO 2007 code. Market3 is a geography-product pair where the geography is a province while the product market is defined using the full 6-digit ATECO 2007 code. We always cluster standard errors with respect to Market3 to allow for arbitrary correlations between disturbances within a broad definition of market.

Sample Type of Regression	(5) Selected OLS	(6) Selected 2SLS-IV	(7) Selected OLS	(8) Selected 2SLS-IV
Dependent Variable	HHI	HHI	Leverage	Leverage
Expo. Funding Gap June 2010	-2.069		-1.4354**	
	(1.4992)		(5.6588)	
Expo. Funding Gap June 2010*Crisis	-1.848***		1.0607	
	(1.0887)		(1.9139)	
Expo. Funding Gap	. ,	-2.7365	, , , , , , , , , , , , , , , , , , ,	-1.8985**
		(1.9758)		(7.4962)
Expo. Funding Gap*Crisis		-2.5932***		0.4383
		(1.5293)		(2.6591)
F-stat 1		8,531		8,531
F-stat 2 (interaction)		6,128		$6,\!128$
Firm dummy	No	No	No	No
Market6-Time dummy	Yes	Yes	Yes	Yes
No of Market6-Time dummies	132,708	132,708	132,708	132,708
Other controls	No	No	No	No
Clustering	Market3-Time	Market3-Time	Market3-Time	Market3-Time
Observations	971,272	971,272	971,272	971,272

Table C.6.: Testing validity of the instrumental variable II

Notes. The unit of observation is a firm - half-year pair. Expo. Funding Gap is the time-variant one-period lagged firm's exposure to its lenders funding gap, and Expo. Funding Gap*Crisis its interaction with the binary variable Crisis which equals 1 for after July 2011 and equals 0 before. Expo. Funding Gap is expressed as a fraction of the total assets (not as a percentage). Leverage is the annual share of financial debt over the sum of financial debt plus equity. HHI is the Herfindahl-Hirschman Index of a firm's loans portfolio across its lenders in a given half year. HHI can take values between 10,000/No banks when the firm borrows an equal amount from each of its lenders, and 10,000 when the firm only borrows from one banks. Market6-Time dummy represent market-time fixed effects where the time is represented by an half year and a market is defined as a geography-product pair where the geography is a province while the product market is defined using the full 6-digit ATECO 2007 code. Market3 is a geography-product pair where the geography is a province while the product market is defined using the first 3 digits of the ATECO 2007 code. We always cluster standard errors with respect to Market3 to allow for arbitrary correlations between disturbances within a broad definition of market.

	(9)	(10)	(11)	(12)
Sample	Selected	Selected	Selected	Selected
Type of Regression	OLS	2SLS-IV	OLS	2SLS-IV
Dependent Variable	Liquidity	Liquidity	No banks	No banks
Expo. Funding Gap June 2010	1.4013***		0.6844^{***}	
	(0.3961)		(0.0804)	
Expo. Funding Gap June 2010*Crisis	-0.3075		0.7982^{***}	
	(0.3488)		(0.0564)	
Expo. Funding Gap		1.8533^{***}		0.9052^{***}
		(0.5244)		(0.1048)
Expo. Funding Gap*Crisis		-0.3274		1.1633^{***}
		(0.4812)		(0.0804)
F-stat 1		8,531		8,531
F-stat 2 (interaction)		6,128		$6,\!128$
Firm dummy	No	No	No	No
Market6-Time dummy	Yes	Yes	Yes	Yes
No of Market6-Time dummies	132,708	132,708	132,708	132,708
Other controls	No	No	No	No
Clustering	Market3-Time	Market3-Time	Market3-Time	Market3-Time
Observations	$971,\!272$	971,272	971,272	$971,\!272$

Table C.7.: Testing validity of the instrumental variable III

Notes. The unit of observation is a firm - half-year pair. Expo. Funding Gap is the time-variant oneperiod lagged firm's exposure to its lenders funding gap, and Expo. Funding Gap*Crisis its interaction with the binary variable Crisis which equals 1 for after July 2011 and equals 0 before. Expo. Funding Gap is expressed as a fraction of the total assets (not as a percentage). Liquidity is the annual fraction of a firm's cash holdings over its total assets. The number of banks is the number of lenders a firm borrow from in a given half year. Market6-Time dummy represent market-time fixed effects where the time is represented by an half year and a market is defined as a geography-product pair where the geography is a province while the product market is defined using the full 6-digit ATECO 2007 code. Market3 is a geography-product pair where the geography is a province while the product market is defined using the first 3 digits of the ATECO 2007 code. We always cluster standard errors with respect to Market3 to allow for arbitrary correlations between disturbances within a broad definition of market.

Sample Type of Regression Dependent Variable	(13) Selected OLS ROA	(14) Selected 2SLS-IV ROA	(15) Selected OLS Total assets	(16) Selected 2SLS-IV Total assets
Expo. Funding Gap June 2010	0.9815^{*} (0.5587)		6.8176^{***} (1.8316)	
Expo. Funding Gap June 2010*Crisis	-2.9454^{**} (1.3987)		3.4319 (3.7811)	
Expo. Funding Gap		1.2987^{*} (0.7395)		9.0169^{***} (2.4157)
Expo. Funding Gap*Crisis		-4.038^{**} (1.939)		5.2828 (5.3806)
F-stat 1		8,531		8,531
F-stat 2 (interaction)		6,128		$6,\!128$
Firm dummy	No	No	No	No
Market6-Time dummy	Yes	Yes	Yes	Yes
No of Market6-Time dummies	132,708	132,708	132,708	132,708
Other controls	No	No	No	No
Clustering	Market3-Time	Market3-Time	Market3-Time	Market3-Time
Observations	$971,\!272$	$971,\!272$	971,272	$971,\!272$

Table C.8.: Testing validity of the instrumental variable IV

Notes. The unit of observation is a firm - half-year pair. Expo. Funding Gap is the time-variant one-period lagged firm's exposure to its lenders funding gap, and Expo. Funding Gap*Crisis its interaction with the binary variable Crisis which equals 1 for after July 2011 and equals 0 before. Expo. Funding Gap is expressed as a fraction of the total assets (not as a percentage). ROA is returns on assets (not as percentage). Market6-Time dummy represent market-time fixed effects where the time is represented by an half year and a market is defined as a geography-product pair where the geography is a province while the product market is defined using the full 6-digit ATECO 2007 code. Market3 is a geography-product pair where the geography correlations between disturbances within a broad definition of market.

C.5. Robustness checks

		(1)	(2)	(3)
	Sample	Selected	Selected	Selected
	Period	2006 to 2012	2006 to 2012	12/2010 to $12/2012$
	Type of Regression	2SLS-IV	2SLS-IV	2SLS-IV
	Dependent Variable	NPL	NPL	NPL
Treatment	Expo. Funding Gap	0.0023	0.0172	-0.0058
		(0.02)	(0.0271)	(0.0208)
	Expo. Funding Gap*Crisis	0.0507^{***}	0.0585^{**}	0.0738^{***}
		(0.0181)	(0.0243)	(0.0177)
Banks characteristics	Expo. ROA June 2010	-0.0428***	-0.0537***	-0.0417***
		(0.0034)	(0.0045)	(0.0044)
	Expo. ROE June 2010	-0.0007	0.0002	-0.0009
		(0.0005)	(0.0007)	(0.0006)
	Expo. Tier 2 Ratio June 2010	-1.1508^{***}	-1.4732^{***}	-1.2826***
		(0.1352)	(0.1807)	(0.1668)
	Expo. Risky Share June 2010	-0.0004***	-0.0005***	-0.0004***
		(0.00004)	(0.0001)	(0.0001)
	Expo. NPLs/Assets June 2010	-0.0039***	-0.0036***	-0.0013
	1 /	(0.001)	(0.0012)	(0.0011)
	Expo. Tier 1 Ratio June 2010	0.2229*	0.3312*	0.4172**
	I	(0.1334)	(0.1838)	(0.1681)
	Expo. Total Assets June 2010	-0.0000***	-0.0000***	-0.0000***
	Enpor Total Theorem Valle 2010	(0.0000)	(0.0000)	(0.0000)
Firms characteristics	EBIT	-0.000001***	-0.00002*	-0.0000***
		(0.0000003)	(0.00001)	(0.0000)
	EBITDA / Assets	-0.0007***	-0.0008***	-0.0003*
		(0.0003)	(0.0000)	(0.0001)
	ННІ	-0.000006***	-0.00001***	-0.00001
	11111	(0.000001)	(0.00001)	(0.0000)
	Lovorago	0.000001)	0.000001)	0.00007
	Leverage	(0.000000)	(0.000004	(0.0000)
	Liquidity	0.0007***	0.0007***	0.0000)
	Liquidity	-0.0007	-0.0007	-0.0010
	N. harler	0.0001)	(0.00007)	(0.0001)
	NO DAIKS	-0.0082	-0.0147	-0.0059
	DOA	(0.0008)	(0.0022)	(0.0010)
	ROA	-0.0005	-0.0001	-0.0013
	The definition of the second sec	(0.0003)	(0.0003)	(0.0004)
	10tal assets	(0.00000003°)	0.00001	(0.0000)
		(0.0000001)	(0.00001)	(0.0000)
	F-stat 1	4,140	4,553	52,122
	F-stat 2 (interaction)	6,828	6,066	7,106
	Firm dummy	No	No	No
	Market6-Time dummy	Yes	Yes	Yes
	No of Market6-Time dummies	107,266	69,602	49,279
	Other controls	No	No	No
	Clustering	Market3-Time	Market3-Time	Market3-Time
	Observations	751,930	417,094	$365,\!143$

Table C.9.: Robustness checks

Notes. The unit of observation is a firm - half-year pair. NPL is a binary variable that equals 1 if the firm is recorded as nonperforming in a given period and equals 0 otherwise. Expo. Funding Gap is the time-variant one-period lagged firm's exposure to its lenders funding gap, and Expo. Funding Gap*Crisis its interaction with the binary variable Crisis which equals 1 for after July 2011 and equals 0 before. Expo. Funding Gap is expressed as a fraction of the total assets (not as a percentage). The variables Expo. Var June 2010, for Var equal to ROA, Tier 1 and Tier 2 ratios, Risky Share, NPLs/Assets, and Total Assets is the time-invariant firm's exposure to its lenders Var at June 2010. EBIT is annual earnings before interests and taxes. EBIT and Total assets are annual firm-level variables expressed in thousands of euros. EBITDA is annual earnings before interests, taxes, depreciation and amortization. EBITDA/Assets is EBITDA over total assets. Leverage is the annual share of financial debt over the sum of financial debt plus equity. Liquidity is the annual fraction of a firm's cash holdings over its total assets. ROA is returns on assets. EBITDA/Assets, Leverage, Liquidity and ROA are annual firm-level indicators expressed as number (not percentages). The number of banks is the number of lenders a firm borrow from in a given half year. HHI is the Herfindahl-Hirschman Index of a firm's loans portfolio across its lenders in a given half year. HHI can take values between 10,000/No banks when the firm borrows an equal amount from each of its lenders, and 10,000 when the firm only borrows from one bank. Market6-Time dummy represent market-time fixed effects where the time is represented by an half year and a market is defined as a geography-product pair where the geography is a province while the product market is defined using the full 6-digit ATECO 2007 code. Market3 is a geography-product pair where the geography is a province while the product market is defined using the first 3 digits of the ATECO 2007 code. We always cluster standard errors with respect to Market3 to allow for arbitrary correlations between disturbances within a broad definition of market. Column (2) uses the sample of firms with total assets or total revenues below their respective medians (i.e., small firms). Column (3) uses the selected sample of firms between December 2010 and December 2012.

		(1)	(2)	(2)
	Sample	Selected	Selected	Selected
	Period	Before Lehman	Between Lehman and Italian crisis	After Italian crisis
	Type of Regression	2SLS-IV	2SLS-IV	2SLS-IV
	Dependent Variable	NPL	NPL	NPL
Treatment	Expo. Funding Gap	-0.0196	0.0062	0.0572**
		(0.0306)	(0.0199)	(0.0279)
Banks characteristics	ROA June 2010	-0.0326***	-0.049***	-0.0354***
		(0.0033)	(0.0037)	(0.005)
	Expo. ROE June 2010	-0.00003	-0.001*	-0.0003
		(0.0005)	(0.0005)	(0.0007)
	Expo. Tier 2 Ratio June 2010	-0.7697^{***}	-1.226***	-1.240***
		(0.1298)	(0.1549)	(0.1923)
	Expo. Risky Share June 2010	-0.0002***	-0.0005***	-0.0004***
		(0.00005)	(0.00005)	(0.0001)
	Expo. NPLs/Assets June 2010	-0.006***	-0.0058***	0.0013
		(0.001)	(0.0011)	(0.0013)
	Expo. Tier 1 Ratio June 2010	0.0266	0.1517	0.5172^{***}
		(0.1167)	(0.1584)	(0.1891)
	Expo. Total Assets June 2010	-0.0000	-0.0000***	-0.0000**
		(0.0000)	(0.0000)	(0.0000)
Firms characteristics	EBIT	-0.00000001	-0.000001***	-0.000002***
		(0.000002)	(0.0000003)	(0.000001)
	EBITDA/Assets	-0.0003***	-0.0002*	-0.0003*
		(0.0001)	(0.0001)	(0.0002)
	HHI	-0.00001***	-0.00001***	-0.00001***
		(0.000001)	(0.000001)	(0.000001)
	Leverage	0.00003*	0.00001**	0.00001**
		(0.00002)	(0.000002)	(0.000004)
	Liquidity	-0.0003***	-0.0007***	-0.001***
		(0.00005)	(0.0001)	(0.0001)
	No banks	-0.0109***	-0.0087***	-0.0049***
		(0.0009)	(0.0009)	(0.0012)
	ROA	-0.00003	-0.0013***	-0.0017***
		(0.00003)	(0.0002753663)	(0.0005)
	Total assets	0.00000002	0.0000003*	0.00000001**
		(0.0000002)	(0.0000002)	(0.0000003)
	F-stat	790	22,041	4,133
	Firm dummy	No	No	No
	Market6-Time dummy	Yes	Yes	Yes
	No of Market6-Time dummies	23,931	55,113	28,222
	Other controls	No	No	No
	Clustering	Market3-Time	Market3-Time	Market3-Time
	Observations	$145,\!636$	403,538	202,756

Table C.10.: Lehman collapse vs Italian crisis

Notes. The unit of observation is a firm - half-year pair. NPL is a binary variable that equals 1 if the firm is recorded as nonperforming in a given period and equals 0 otherwise. Expo. Funding Gap is the time-variant one-period lagged firm's exposure to its lenders funding gap, and Expo. Funding Gap*Crisis its interaction with the binary variable Crisis which equals 1 for after July 2011 and equals 0 before. Expo. Funding Gap is expressed as a fraction of the total assets (not as a percentage). The variables Expo. Var June 2010, for Var equal to ROA, Tier 1 and Tier 2 ratios, Risky Share, NPLs/Assets, and Total Assets is the time-invariant firm's exposure to its lenders Var at June 2010. EBIT is annual earnings before interests and taxes. EBIT and Total assets are annual firm-level variables expressed in thousands of euros. EBITDA is annual earnings before interests, taxes, depreciation and amortization. EBITDA/Assets is EBITDA over total assets. Leverage is the annual share of financial debt over the sum of financial debt plus equity. Liquidity is the annual fraction of a firm's cash holdings over its total assets. ROA is returns on assets. EBITDA/Assets, Leverage, Liquidity and ROA are annual firm-level indicators expressed as number (not percentages). The number of banks is the number of lenders a firm borrow from in a given half year. HHI is the Herfindahl-Hirschman Index of a firm's loans portfolio across its lenders in a given half year. HHI can take values between 10,000/No banks when the firm borrows an equal amount from each of its lenders, and 10,000 when the firm only borrows from one bank. Market6-Time dummy represent market-time fixed effects where the time is represented by an half year and a market is defined as a geography-product pair where the geography is a province while the product market is defined using the full 6-digit ATECO 2007 code. Market3 is a geography-product pair where the geography is a province while the product market is defined using the first 3 digits of the ATECO 2007 code. We always cluster standard errors with respect to Market3 to allow for arbitrary correlations between disturbances within a broad definition of market. Column (1) only uses observations before the collpse of Lehman Brothers in Septemebr 2008. Column (2) only uses observations between the collpse of Lehman Brothers in Septemebr 2008 and the start of the Italian crisis in July 2011. Column (3) only uses observations after the start of the Italian crisis in July 2011.