LUISS

Department
of Economics and Finance

Ph.D in Economics

Cycle XXXVI

Essays in structural economics

Lippi Francesco

SUPERVISOR

Reichlin Pietro

CO-SUPERVISOR

Genovese Carmelo

CANDIDATE

Academic Year 2024/2025

Contents

Chapter 1

A structural model for corporate liquidity

Carmelo Genovese[1]

Abstract

I analyse the patterns of the cross-sectional distribution of liquidity across Italian limited liability companies. I document that the level of liquidity is stable over business cycles and that firms react to fluctuations in economic conditions by adjusting their net trade debt position. I show that the general patterns are common across industries. To quantify the impact of cash flow uncertainty and management costs, I estimate an inventory model of liquidity management. The empirical strategy is to match the moments of the cross-sectional distribution of liquidity, which I show to be stable over the business cycle. I use the estimated model to highlight the impact of inflows' uncertainty and opportunity costs on the optimal management policy. I perform comparative statics exercises to understand the impact of different structural changes on liquidity management choices, highlighting direct and indirect effects on the optimal liquidity management policy.

---

[1] Luiss Guido Carli. I would like to thank my supervisor Francesco Lippi, Fabio Braggion, Federico Carlini, Maria Pia Lombardo, Paolo Santucci de Magistris, and Luana Zaccaria.

1.1 Introduction

Liquidity management is one of the core activities at the firm level. Entrepreneurs and managers need to plan the allocation of resources, manage the cash flow, and balance liquidity inflows/outflows. Some firms might be willing to keep liquidity buffers, while others might rely on their supply chain to survive liquidity crunches. Despite the importance of this topic, there is a lack of studies disentangling liquidity dynamics at the firm and the aggregate level, and the literature focuses only on cash holdings, which are only a part of the firms' liquidity. In this paper I focus on the distribution of liquidity across firms, highlighting the importance of trade credit/debt, and propose a structural model to analyze liquidity management decisions that produce results coherent with the empirical findings.

In their seminal paper, Bates et alt.[1], documented a constant increase in US firms' cash holdings from 1980 to 2006. This finding is interesting and puzzling because one expects cash-holding decisions to be affected by business cycle fluctuations and other macroeconomic events. For example, it is not obvious why cash holdings are unaffected by changes in aggregate riskiness or changes in money supply. When the monetary policy is tight, interest rates are high, increasing the opportunity cost of holding liquid assets. This, in theory, gives incentives to reduce cash holdings and invest. Similarly, when the uncertainty level is high, firms might find it optimal to increase their cash holdings to create a buffer to absorb negative shocks. To rationalize the findings in [1], several empirical studies have been proposed, suggesting that precautionary motives, agency costs and financial constraints influence cash holdings decisions.

These studies highlight particular mechanisms influencing liquidity or find some correlation between firms' characteristics and cash holdings choices, but focus only on cash holdings, which is only part of the corporate liquidity and the liquidity management problem. Corporate liquidity depends not only on its cash holdings but also on credit lines and trade credit/debt position. For example, a firm with a credit line granted, which is not used, has a liquidity buffer that is not captured using cash holdings as the only measure of liquidity. Firms with enough liquidity buffers can decide to allow their customers to postpone payments for goods and services, lending their liquidity through the supply line. On the other side, firms with cash deficits might ask for trade debt, and by looking only at cash holdings, the firm seems to have liquidity, while in reality this liquidity is temporally "borrowed" from suppliers, and vice versa for the suppliers. Measuring liquidity at the corporate level is not easy because several factors need to be considered and because liquidity can be stored in multiple forms.

While credit lines have been widely studied as a contingent source of liquidity, only recently the importance of trade credit/debt position has been analysed, mostly because data are rarely available for reduced-form studies. A recent contribution on the importance of credit lines for liquidity management is by Nikolov et al.[2], where they estimate a dynamic model of liquidity management and highlight the importance of collateral assets to rely effectively on credit lines as liquidity buffers. Credit lines have the desirable characteristic of being contingent, meaning that firms use them when they need extra liquidity, but they are costly and not easy to access for most firms.

On the other hand, every firm has commercial relationships within its supply chain, and often B2B transactions are not settled in cash but postponed for some days, raising trade credit/debt. This is a common practice for most firms and is a key component of corporate liquidity management. For example, a firm suffering a liquidity deficit can ask its supplier to delay payments, ask its customers to anticipate payments, or directly settle the sales in cash. Amberg et al. [3] documented this pattern using an exogenous liquidity downfall in Sweden and showed that the trade position is indeed relevant for liquidity management. In another working paper [4], they show evidence that the usage of trade credit is common across firms' characteristics, like size.

On the theoretical side, despite the increasing interest in the topic, there is a scarcity of structural models built to disentangle liquidity decisions or that reproduce a cross-sectional distribution of liquidity coherent with the data. The first key contribution to this literature is Miller, and & Orr [5], which propose an inventory model for cash management decisions. Recently, there has been a growing literature that models firms' investment-liquidity-risk management decisions, like the model proposed by Bolton et al. [6]. However, to the best of my knowledge, those works either do not focus on firms' liquidity, do not consider the trade credit position, or fail to match empirical evidence. In this paper I contribute to this literature, generalizing and estimating the model proposed by Miller & Orr by including stochastic cash flows and asymmetric costs.
In the first part of the paper, I show that the distribution of liquidity across firms is stable over the business cycle if the net trade position is included. The net trade position of each firm includes the standard "Cash and cash equivalents" plus the sum of net trade credit[1]. This measure is closely related to the net working capital. I then perform a descriptive statistic analysis and focus on differences at the industry level. In the second part of the paper, I estimate a fixed-cost model to simulate the decision process of a representative firm managing its liquidity. This model is a generalization of the framework proposed by Miller and Or [5] where the firm minimizes the costs of managing liquidity, which is given by the opportunity cost of holding liquid assets and the fixed cost to be paid to adjust its liquidity position. The model is estimated using Italian S.P.A. balance sheet data. The sample spans from 2013 to 2022, including the COVID-19 period. The solution to the firm's problem is to determine two critical levels of liquidity such that the firm does not change its assets allocation while liquidity is within this critical region. If liquidity ends up outside this region, the firm pays the fixed cost and adjusts its liquidity to the optimal level. The lower level represents the lowest level of liquidity such that the firm prefers to pay the cost and externally raise its liquidity. The upper level represents the maximum level of assets that the firm finds optimal to keep liquid, hence above that level, it prefers paying the fixed cost needed to find the optimal investment opportunity and to invest the excess liquidity, enjoying the investment revenues. Finally, the level of liquidity that the firm desires to hold is the one that minimizes the total expected cost.
The solution is associated with an invariant distribution of liquidity. The empirical strategy of this paper is to match the empirical cross-sectional distributions with the structural distributions to estimate the costs associated with liquidity management. In section 1.3.1, I show that the cross-sectional distribution is stable, therefore it can be compared with the invariant one associated with the structural model. After the estimation, which is based on the generalized method of moments (GMM), I analyse the impact of changes in structural parameters on the optimal policy. In particular, I focus on the relationship between cash flow uncertainty, opportunity costs, and liquidity management. Finally, I state a generalized version of the model.
The paper is organized as follows: In section 1.2 I review the literature, and in section 1.3 I perform the empirical analysis. In section 1.4 I solve an inventory model for cash management and determine the associated distribution of cash holdings. In section 1.4.3 I estimate the model and perform a comparative statics exercise. Before the conclusions, I state a generalized version of the model.

1.2 Literature review

This paper is related to two strands of the literature. The first one is the empirical literature on liquidity management, the second one is the structural corporate finance literature. In the corporate finance literature, one of the most well-known papers is Bates et al.[1], who was the

[1]In Italy, Italian limited liabilities company, S.P.A., are requested to report in their balance sheet short-term trade credit and debt.

first to shed light on the patterns of cash holdings at the firm level. They find that US firms have been increasing their cash holdings since the eighties. They point out that the average firm doubled (or more) its cash/asset ratio. They claim that this was due to the increase in cash flow riskiness, lower inventories and account receivables and the advent of R&D. Acherya et alt. [7] analysed the correlation between cash accumulation and banking lending spreads and concluded that this liquidity accumulation is coherent with precautionary motives for riskier firms. Part of the literature focuses on the different cash management policies between public and private companies. For the US, Gao et al. [8] studied the relationship between agency costs and cash accumulation and found that agency costs are generally associated with higher liquidity. Moreover, higher agency frictions are associated with higher target levels of liquidity and different management choices. Poti et al. [9] found similar results for European firms. They analysed the interconnection between cash management policies and stakeholders' risk attitudes and found evidence that supports the correlation between cash management and tolerance to risk exposure. Bigelli & Sánchez-Vidal [10] analysed the Italian economy and found similar results. In this paper, I propose a different measure of liquidity that takes liquidity provision in the form of trade credit/debt and find that firms keep about 15% of their assets liquid and that there is a substitution pattern between cash and trade credit allowances.
On the theoretical side, liquidity management was first studied in theoretical frameworks by Baumol (1952) and Tobin (1956), who used inventory control theories to model money demand at the household level. Following this literature, Miller & Orr(1966) proposed an inventory model for firm cash holdings with fully stochastic cash flow. Frenkel & Jovanovic [11] developed a generalized version of the model with stochastic flows. Recently, this framework has been used to study cash management and cash decisions at the household level by Alvarez & Lippi [12] and firms' investment decisions by Baley & Blanco [13].
Another closely related strand of the literature focuses on asset allocation, investments and funding. For example, Riddick & Whited [14] linked the level of cash holdings to the investment opportunities of the firm and found that firms accumulate cash holdings to invest when the right investment opportunity arrives. Between income uncertainty and constrained access to external funds, they find that the former plays a major role in determining the firms' saving policy. Bolton, Chen, and Wang [15] extend the analysis by estimating a larger model that includes investment decisions, financing opportunity and cash management a lá Miller&Orr. In their model, the representative firm decides how much to invest, how much to borrow and the optimal payout policy to maximize shareholders' return. Each period the firm receives revenues from production net of the financing decisions. The relationship between corporate liquidity, investment and financing decisions is that the former influences financing and payouts. However, their model implies an invariant distribution of cash holdings that is not coherent with the cross-sectional distribution of cash holdings. I contribute to this literature by solving an inventory model that produces results consistent with empirical evidence and by analysing the impact of changes in structural parameters on liquidity management.

1.3 Empirical evidence

In this section I describe in detail the dataset I used, I motivate my measure of liquidity, and I present the main statistics and empirical findings.
To perform the empirical analysis I use balance sheet data of Italian limited liabilities corporations.[2]

[2] These companies are called "società per azioni" (S.P.A.) or "società in accomandità per azioni" (S.A.P.A.) and correspond to limited corporations whose equity is divided into stocks.

Both public and private companies are included.[3] The sample span is 2013-2022, characterized by expansionary monetary policy, the COVID-19 pandemic, and the post pandemic monetary tightening. In this period there existed about 40.000 S.P.A., however, I excluded firm-year observations if one of the balance sheet entries I need is missing or empty. Additionally, I exclude government firms and corporations whose key activity is a primary sector activity, financial or commodity. I also exclude inactive firms that report zero sales and firms that filed for bankruptcy. After the cleaning, I am left with about 14.000 observations per year. Appendix A includes table 1.9 and figure 1.12, which contain information about the geographical dispersion of the sample. In section 1.3.3 I analyse industry differences.

1.3.1 Measuring liquidity

The first step is to build a measure of corporate liquidity. In household literature, the main measure of liquidity is the amount of cash in hand (or in the bank account). This measure is also used for corporate studies, standardized by total assets to correct for firm size. In particular, this measure is the ratio between "*Cash and cash equivalents*" over "*Total assets*", two balance sheet entries that are publicly available. Calling this measure of liquidity ch_t, the cross-sectional distributions in 2013 and 2022 are shown in the following Figure:

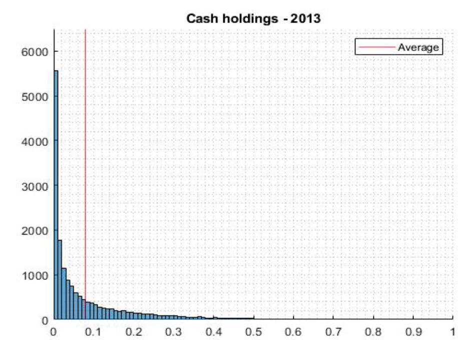

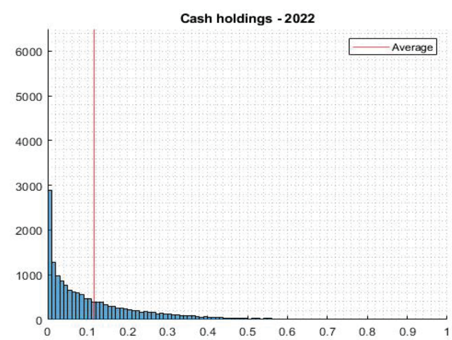

Figure 1.1: The figure shows the cross-sectional distribution of cash holdings in 2013 (left panel) and 2022 (right panel). The X-axis represents the cash holdings ratio, while the Y-axis represents the number of firms in each bin. Each bin represents a percentage point, so that, for example, the first bin contains the number of firms with cash and cash equivalent worth between 0% and 1% of total assets.

From Figure 1.1 it is possible to see how firms used to keep a small fraction of their assets as cash, as the mass is concentrated towards 0. In this sample period, there is a monotonic trend in increasing liquidity, a pattern already documented in the literature. This can be seen both by the shifting of the mass to the right and by the increase of the average cash holding, which happens in every year of the sample. In 2013, half of Italian corporations had less than 2% of their asset held as cash or cash equivalents. In 2022 the number of firms with liquidity below 2% halves, while the average firm now keeps more than 10% assets as cash or cash equivalents. This increase is constant

[3] The data have been downloaded from "AIDA" and "Refinitiv", Bureau van Dijk platforms.

across the sample period. The following table and figure contain details about the distribution of cash holdings across the sample:

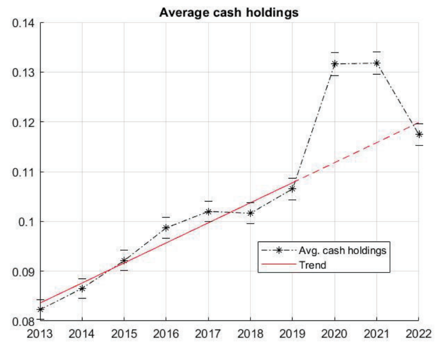

Figure 1.2: Average cash holdings with confidence intervals and trend.

Year	Mean	Stand. Dev.	Median
2013	8.17	12.97	2.76
2014	8.61	13.14	3.13
2015	9.16	13.25	3.72
2016	9.82	13.54	4.44
2017	10.12	13.69	4.76
2018	10.07	13.74	4.68
2019	10.54	14.29	5.01
2020	13.04	14.92	8.06
2021	13.05	14.93	7.95
2022	11.61	14.17	6.68

Table 1.1: Statistics for the ch_t. All variables are in %.

From Table 1.2 it is clear that during the COVID-19 year, there was a sharp increase in cash holdings. One of the reasons is that the Italian government and the European Union authorities adopted a series of policies to facilitate banks' lending and access to credit to small and medium firms. These policies aimed to guarantee that firms had enough resources, especially liquidity, to survive the pandemic.[4] After the pandemic, liquidity returned to the pre-pandemic trend level. The trend was calculated using OLS and observations from 2013 to 2019.

The main question of this paper is to understand how economic and financial friction influences liquidity decisions and why firms have only a small fraction of their asset as liquid, while theoretical works suggest that firms should keep a liquidity buffer to absorb shocks.[5] Before the pandemic, although monetary policy was loose, most firms had only 3-5% of their assets held as cash or cash equivalents. I propose to use a different measure of liquidity because cash holdings are an important component yet an imperfect proxy for corporate liquidity.

Part of the literature defines liquidity as the sum of cash holdings and credit lines, as these financial contracts provide liquidity reserves that can be used when needed. The paper by Nikolov et al. [2] explores the relationship between cash holdings and credit line usage. They find that credit lines are used when cash is needed and external financing is costly, but have the downside that firms need collateral assets to have access to them.

In this paper I take a different perspective and study liquidity coming from the operational activities of the firms, focusing on liquidity inflows/outflows from the working capital. In particular, I use a measure of liquidity that corrects cash holdings for the net exposure on the supply chain, which I measure as the difference between short-term trade credit and debt. The idea that firms use their trade position to manage their liquidity has been proposed and documented by several scholars, including Biais et al.[16], Wilner[17], Cuñat[18], Amberg et al.[3][4]. Nilsen [19] analyses the link between monetary policy, bank lending and trade credit. On the other side, trade credit financing is risky when the underlying trade activity slows down or is harmed, like during the pandemic. Bureau

[4] These policies are DL n. 23 dell'8 aprile 2020 (cd. "Decreto Liquidità"), Legge n. 178/2020 "("Legge di Bilancio 2021"), DL 73 del 25 maggio 2021 (cd. Decreto "Sostegni bis"), and Legge n. 234 del 30 dicembre 2021 (Legge di Bilancio) among others.

[5] For example, in [15], the optimal policy is to keep about 20% of assets liquid.

et al.[6] provides some insights on the impact of COVID-19 on this mechanism.
Following this literature, I propose a measure of liquidity that includes both cash holdings and the net trade position. The net trade position is the difference between commercial credit due by borrowers minus commercial debt due to suppliers, all due within the year.[6] The idea behind this is that firms manage their liquidity by adjusting their trade position. In particular, they can borrow liquidity from their supplier by postponing payments or by asking to receive inputs on credit and, similarly, firms can lend liquidity to their customer by giving them credit. Usually, these credits are short-term and granted to firms of the same production network or supply chains, based on commercial agreements and long-standing relationships. Hence, trade credit/debt is a key and active part of liquidity management. The cost of this liquidity is a premium on the price of the goods/services. To measure liquidity[7] I compute the net trade position and add it to the level of cash holdings:

$$m_t = \frac{Cash\ and\ C.\ Equiv._t + \overbrace{Trade\ credit_t - Trade\ debt_t}^{Net\ trade\ position_t}}{Tot.\ assets_t} \tag{1.1}$$

The following table and figure describe the differences between the standard measure of liquidity (ch_t) and the measure of liquidity including the trade position (m_t):

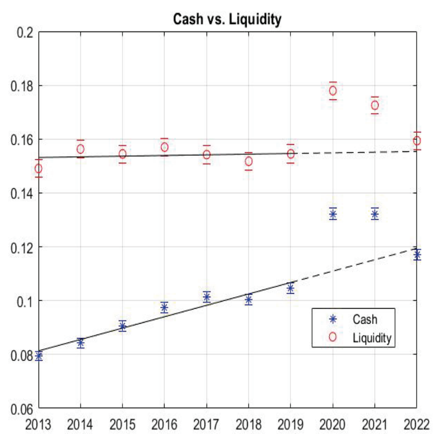

Figure 1.3: Average cash holdings with 95% confidence intervals.

Year	Mean		Stand. dev.		Median	
	ch_t	m_t	ch_t	m_t	ch_t	m_t
2013	7.95	14.91	11.75	21.73	2.93	13.01
2014	8.44	15.63	11.98	21.90	3.35	13.55
2015	9.06	15.45	12.28	21.67	3.99	13.55
2016	9.75	15.70	12.63	21.92	4.78	14.04
2017	10.12	15.42	12.92	22.13	5.13	13.75
2018	10.05	15.18	12.96	22.16	4.96	13.28
2019	10.46	15.45	13.38	22.07	5.60	13.57
2020	13.23	17.80	14.28	21.25	8.63	15.79
2021	13.02	17.26	14.29	20.80	8.46	14.85
2022	11.71	15.94	13.45	20.91	7.10	13.63

Table 1.2: All variables are in %.

Accounting for the short-term net trade position, the liquidity level is higher by about 5%. From Figure 1.3 it is possible to see that firms' liquidity is relatively stable over the business cycle.[8] Analyzing the patterns of ch_t and m_t, it is clear that during the business cycle firms trade off cash

[6] In the Italian balance sheet firms report their commercial credit towards customers due within a year and debt towards supplies due within a year. The name of these entries are "Crediti verso clienti entro l'anno" and "Debiti verso fornitori entro l'anno".

[7] I would like to stress that this measure is an imperfect proxy for firms' liquidity but captures several dynamics missing in ch_t. With this measure, the cross-sectional distribution of cash holdings is comparable with the results in [15].

[8] The Kolmogorov Smirnov test can be found in the appendix.

holdings with the trade position. Interest rates and the cost of credit decreased during the 2013-2019 period and it is possible to conclude that firms, on average, lowered their reliance on credit from suppliers increasing cash, which was easier to get from banks and markets. It is possible to see that in 2020 there was a parallel shift in both measures, but the net trade credit decreased in 2021 while cash holdings did not change. It is also worth noting that m_t allows us to observe firms that have negative liquidity, as the trade debt can be higher than cash and trade credit. With the other measure, ch_t, this is not possible because firms either have 0 or positive cash and cash equivalents. This is reflected in the fact that there is bunching at 0, as it is possible to see from Figure 1.1. In a recent survey by an Italian corporation providing cash flow management services, they found out that "40% of Italian medium-size firms have liquidity deficit at least once a month for more than 50.000€."[9] Firms can have temporarily negative liquidity, for example when they need to make payments to suppliers that are worth more than their cash holdings. In those cases, trade credit/debt becomes an important tool for liquidity management and/or when credit lines are opened. The following figure contains the distribution of m_t:

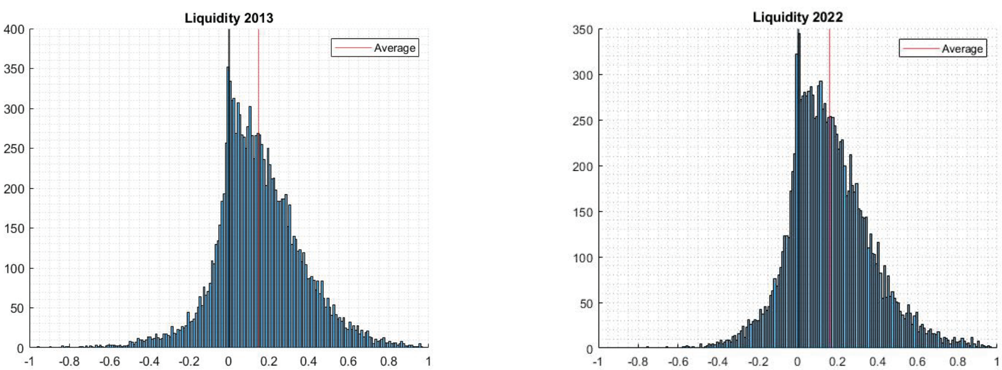

Figure 1.4: The figure contains the cross-sectional distribution of liquidity in 2013 and 2022. The X-axis represents the ratio m_t, while the Y-axis represents the number of firms in each bin. Each bin represents a percentage point.

Differently from cash holdings, which are exponentially distributed with the mass concentrated towards 0, m_t is tent-shaped, with the majority of the mass concentrated between 0 and the average level. Firms keep about 13-15% of their assets as liquid. In the next section, I focus on the relationship between cash and the trade position.

1.3.2 Liquidity decomposition

In this section, I analyze the relationship between cash holdings, liquidity, and the patterns in their dynamics. From table 1.2 it is possible to see that from 2013 to 2019 firms were decreasing their net trade position and increasing their cash holdings. During the COVID-19 pandemic, liquidity rose, and this increase was driven by cash holdings accumulation. At the same time, net trade credit decreased. This was the effect of the government's interventions, that, despite the slowdown of economic activities, sustained corporations' liquidity. The dynamic is clear in the following figure:

[9] Link: "https://agicap.com/it/articolo/studio-costi-nascosti-previsioni-medie-imprese-italiane/".

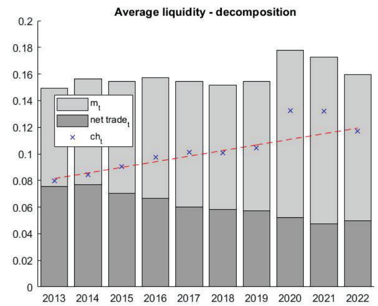

Figure 1.5: This figure contains the average level of liquidity m_t, of cash holdings ch_t and trade position *net* $trade_t$ as a fraction of total asset. The red line is the trend level computed excluding the COVID-19 years

The next question is to understand if the reduction in net trade credit is caused by a reduction in trade credit or an increase in trade debt:

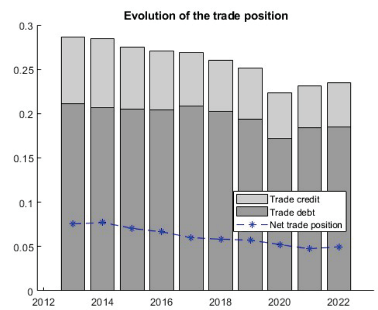

Figure 1.6: This figure contains the average level of trade credit and trade debt per year. The blue line is the net trade position.

From the Figure 1.6 above, it is possible to see that the level of trade debt was stable between 2013 and 2019, with a sudden decrease during the pandemic. On the other side, the level of trade credit is decreasing over the same period, with a sudden decrease during the pandemic. In the appendix, Section 1.6.3, it is possible to see that these changes reflect shifts in the distribution of trade credit and debt across firms, hence it is possible to conclude that the change in the composition of liquidity during the 2013-2019 period was mostly due to a decline of trade credit.[10] This decline, together with the expansionary monetary policy, led to a change in the composition of corporate liquidity in favour of cash holdings. Since the "demand" for trade debt seems to be relatively stable, in the future it would be worth disentangling the decrease in trade credit granted by S.P.A. due to lower demand (substitution effect, as documented in this paper) to the decrease due to other factors, for example integration to the global supply chain (trade credit granted foreign corporations).[11] This

[10] The net trade credit is not in zero net supply because I do not consider other legal types of corporations, like S.R.L., and because Italy is an open economy with corporations having cross-border commercial relationships.

[11] This channel can be relevant because access to alternative liquidity providers can alter bargaining power, and firms' resiliency. See for example Giannetti et al. [20].

substitutability can explain why most firms had low levels of "cash and cash equivalents", especially in the first years of the sample.
To confirm this intuition I perform a panel regression to measure the correlation between the yearly change in cash holdings and the yearly changes in trade credit and debt. The results are summarized in the following table:

	Full sample	No COVID-19	Pre COVID-19
ΔTrade credit	-0.231	-0.221	-0.212
	(0.009)	*(0.010)*	*(0.009)*
ΔTrade debt	-0.000	0.016	0.019
	(0.012)	*(0.008)*	*(0.006)*
Constant	0.001	-0.001	0.002
	(0.000)	*(0.000)*	*(0.000)*
Sector FE	Yes	Yes	Yes
Size FE	Yes	Yes	Yes
Year FE	Yes	Yes	Yes
2020 - 2021	Included	Excluded	Excluded
2022	Included	Included	Excluded
R^2 *(Within)*	6.68%	6.44%	6.42%

Table 1.3: Dependent variable is yearly change in cash holdings Δch. All variables are scaled by total assets and windsorized at the 0.5 and 99.5 percentiles.

There is a negative relation between the change in trade credit granted by the firm and its cash holdings. As firms decrease their trade credit allowances, their cash holdings increase. The inverse is true for trade debt because an increase in trade debt is associated with an increase in cash holdings, but the magnitude is smaller. It is worth pointing out that the constant is negative in the second specification because, as it is possible to see from Figure 1.3, in 2022 the level of cash holdings decreased sharply in response to the end of COVID-19 policies. The negative correlation between changes in cash holdings can be interpreted as evidence of their substitutability in liquidity management.
The same exercise can be performed in levels to understand the correlation between cash holdings and trade credit/debt:

	Full sample	No COVID-19	Full sample	No COVID-19
Constant	0.114	0.107	0.100	0.093
	(0.000)	*(0.003)*	*(0.001)*	*(0.001)*
Trade credit	-0.085	-0.083	-	-
	(0.0124)	*(0.0110)*	-	-
Trade debt	0.020	0.021	-	-
	(0.000)	*(0.012)*	-	-
Net trade position	-		-0.064	-0.0626
	-	-	*(0.011)*	*(0.0010)*
Sector FE	Yes	Yes	Yes	Yes
Size FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
2020 - 2021	Included	Excluded	Included	Excluded
R^2 *(Within)*	0.55%	1.26%	1.24%	0.89%

Table 1.4: Dependent variable is cash holdings *ch*. All variables are rescaled by total assets. Trade credit and trade credit are windsorized at the 0.5 and 99.5 percentiles. Net trade position is the difference between Trade credit and debt.

The results confirm that firms with lower cash holdings are those with higher net trade credit and that changes in cash holdings are negatively correlated with changes in trade credit, confirming that there is a substitution pattern between the two variables.

1.3.3 Industry level analysis

In this section, I highlight differences across industries. Each industry is classified following the NACE2 definition:

Sector	# firms	Sector	# firms
Manufacture	10,094	IT & Communication	1,703
Constructions	2,226	Prof. services	3,626
Trade	5,041	Other. Serv.	621
Transports	1,493	Education	72
Hospitality	569	Health et co.	354
Art et co.	285	N.a.	130

Table 1.5: Number of firms per industry. If a corporation operates in more than one industry, I only consider its main industry. Following the literature, I exclude the following industries: Agricultural, Mining, Quarrying, Natural resources, Finance, and Public Administration.

Most corporations are in the manufacturing industry, in trade and services. Different sectors have different business characteristics that influence operations management and liquidity management. For example, there is heterogeneity in both average cash and cash equivalents and liquidity:

		2013	2014	2015	2016	2017	2018	2019	2020	2021	2022
Manufacture	Liquidity	0.15	0.16	0.16	0.16	0.15	0.15	0.15	0.17	0.16	0.14
	Cash	0.08	0.08	0.09	0.10	0.10	0.10	0.10	0.13	0.13	0.11
Constructions	Liquidity	0.12	0.13	0.13	0.14	0.14	0.14	0.14	0.16	0.15	0.13
	Cash	0.06	0.07	0.07	0.08	0.08	0.08	0.09	0.11	0.11	0.10
Trade	Liquidity	0.14	0.14	0.14	0.14	0.14	0.14	0.14	0.18	0.19	0.17
	Cash	0.08	0.08	0.08	0.09	0.09	0.09	0.10	0.13	0.13	0.12
Transports	Liquidity	0.13	0.14	0.14	0.14	0.14	0.15	0.15	0.16	0.17	0.18
	Cash	0.08	0.09	0.09	0.10	0.10	0.11	0.11	0.12	0.13	0.13
Hospitality	Liquidity	0.04	0.05	0.05	0.06	0.06	0.06	0.06	0.06	0.07	0.08
	Cash	0.06	0.07	0.07	0.08	0.08	0.08	0.08	0.08	0.09	0.10
IT & Comm.	Liquidity	0.22	0.23	0.21	0.22	0.23	0.23	0.21	0.23	0.24	0.23
	Cash	0.11	0.12	0.13	0.13	0.13	0.13	0.14	0.17	0.17	0.16
Prof. services	Liquidity	0.16	0.20	0.17	0.18	0.18	0.18	0.19	0.20	0.21	0.20
	Cash	0.10	0.10	0.11	0.11	0.12	0.12	0.13	0.15	0.15	0.14
Other. Serv.	Liquidity	0.30	0.30	0.31	0.33	0.34	0.33	0.32	0.35	0.34	0.34
	Cash	0.10	0.09	0.10	0.11	0.12	0.11	0.13	0.16	0.16	0.15
Education	Liquidity	0.23	0.21	0.24	0.29	0.29	0.26	0.20	0.16	0.29	0.22
	Cash	0.13	0.11	0.13	0.13	0.14	0.16	0.14	0.16	0.17	0.15
Health et co.	Liquidity	0.16	0.13	0.14	0.15	0.15	0.14	0.14	0.14	0.14	0.12
	Cash	0.08	0.08	0.09	0.09	0.10	0.10	0.11	0.11	0.11	0.09
Art et co.	Liquidity	0.06	0.07	0.06	0.08	0.08	0.09	0.10	0.11	0.13	0.14
	Cash	0.06	0.06	0.07	0.08	0.09	0.08	0.09	0.11	0.12	0.12
N.a.	Liquidity	0.14	0.15	0.14	0.14	0.14	0.15	0.13	0.15	0.15	0.11
	Cash	0.06	0.07	0.08	0.10	0.09	0.11	0.11	0.12	0.13	0.12

Table 1.6: Liquidity and cash holdings as % of total assets.

It is worth noting that in every sector average liquidity is higher than average cash and cash equivalents, suggesting that Italian S.P.A. have, on average, positive net trade credit. This surplus can be both towards other types of corporations or foreign firms, but with the available data, it is not possible to further test this hypothesis and their implications.[12] In the appendix, section 1.6.4, it is possible to find the liquidity distribution by industry. To understand the impact of different business models on trade credit policies I analyse the average changes in trade credit, debt and net trade credit by industry and by year. From section 1.3.2 we already know that the average level of net trade credit decreased in favour of cash holdings. However, it is important to highlight the industry differences. The next figure highlights changes in the average trade credit and debt industry by industry. The core industries, manufacturing, services, and trade decreased their net trade position, while others have less sharp dynamics:[13]

[12] The result holds also considering the median, except for "Other Serv.".

[13] The education industry is more volatile than the other industries because of the lower number of observations for that industry.

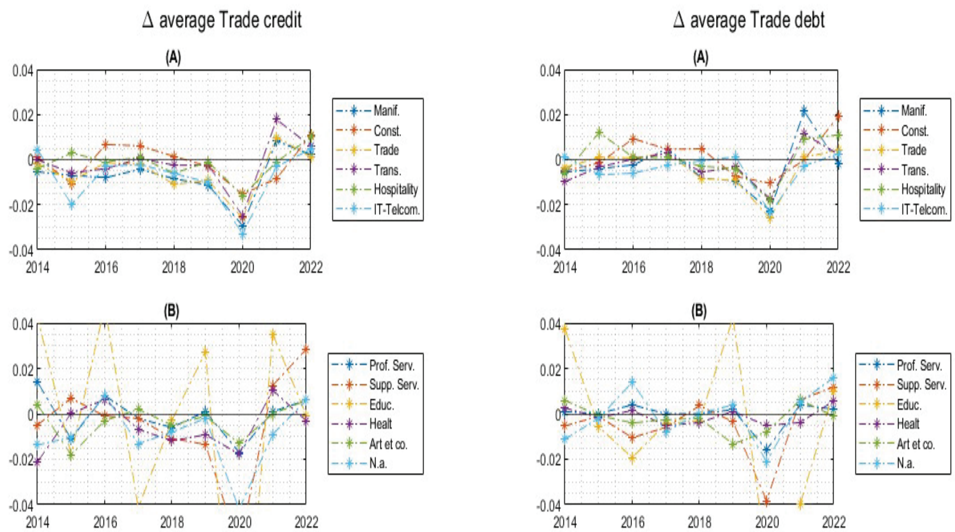

Figure 1.7: Industry by Industry change in the average trade credit (left) and debt level (right). I split industries, by NACE2 order, into panel (A), and panel (B).

Before the pandemic, only the construction industry increased the net trade credit allowance, while the other sectors decreased it. In 2020 both trade credit and debt decreased, but the decrease in trade credit is sharper as it is possible to see from the magnitude. In 2022 the changes are consistent with a transition to the pre-pandemic steady state.

1.4 A model for liquidity management

In this section of the paper, I propose a model for firms' liquidity management. The model belongs to the inventory models class [21] [22], and is a generalization of the one proposed by Miller & Orr [5]. The solution of the model is a steady-state policy for liquidity management that is associated with an invariant distribution. I estimate the model by matching the implied moments with the observed counterparts.

1.4.1 The problem of the firm

A representative firm minimizes the total expected cost of liquidity management. The total cost depends on two costs, the opportunity cost of keeping assets liquid and the fixed costs to be paid to adjust the liquidity level. The opportunity cost $R(m_t)$ is a flow cost and captures the forgone returns of investing liquid assets in longer-term activities. The firm can decide at any moment to raise liquidity, for example by borrowing or opening a credit line, but this is a costly operation, which is modelled as a fixed cost c_1 operation. Similarly, the firm can invest its liquidity in longer-term activities but needs to pay the fixed cost c_2, which captures the cost of finding the right investment opportunity. Sometimes, financing or investment opportunities arrive at the right moment, therefore with some arrival rate λ per unit of time, the adjustments will be free. I assume that the level of liquidity m_t changes with the cash flow[14], which is modelled as a Brownian motion dW_t so that:

$$dm_t = \mu dt + \sigma dW_t \tag{1.2}$$

Given the cost structure and the assumption that liquidity evolves following (1.2), the problem of the firm minimizing the total expected cost of liquidity management is:

$$V(m_{t_0}) = \min_{\{\tau_j, m_t\}} \left\{ \mathbb{E}_{t_0} \left[\sum_{\tau_j} e^{-\rho \tau_j} c(\Delta m_{\tau_j}) \chi(\tau_j) + \int_{t_0}^{+\infty} e^{-\rho t} R(m_t) m_t dt \right] \right\} \tag{1.3}$$

where:

1. τ_j is a stopping time, i.e. the time when the firm decides to raise or invest its liquidity, and Δm_{τ_j} is the change in m at the stopping time;
2. χ_{τ_j} is a dummy variable that is 0 if the adjustment is free (with arrival rate λ);
3. $R(m)$ is the opportunity cost. When $m > 0$, R_2 is the opportunity cost of liquidity, while when $m < 0$, R_1 is the cost of borrowing through trade:

$$R(m) = \begin{cases} R_1 \ \textit{if} \ m < 0 \\ R_2 \ \textit{if} \ m \geq 0 \end{cases}$$

4. $c(\Delta m)$ is the fixed cost to be paid to adjust liquidity:

$$c = \begin{cases} c_1 \ \textit{if} \ \Delta m > 0 \\ c_2 \ \textit{if} \ \Delta m < 0 \end{cases}$$

5. ρ is the discount rate.

From equation (1.3) I derive the associated Hamilton Jacobi Bellman equation:

$$(\rho + \lambda) V(m_t) = \lambda V(m^*) + R(m_t) m_t + \mu V'(m_t) + \frac{\sigma^2}{2} V''(m_t) \tag{1.4}$$

I solve (1.4) by guess and verify. To determine the constants I use standard boundary conditions (optimality of critical points, smooth pasting, and value matching). The system of boundary conditions is not linear, therefore I use a numerical solver. The solution to the problem consists of two critical levels of liquidity $\{m_1, m_2\}$ such that, if liquidity ends up below m_1 or above m_2, the fixed cost $c(\Delta m)$ is paid and the liquidity level is reset to the optimal one m^*. m^* is the level of liquidity that minimizes the problem in (1.3). Details about the solution procedure can be found in the Appendix, section 1.6.5.

What is the intuition behind the solution? For any amount of liquidity, the firm chooses whether to adjust or wait for the next realization of the cash flow, weighing the fixed cost to be paid with the expected opportunity cost. Suppose that the liquidity level is high, then the opportunity cost is high. If the firm thinks that the next period additional resources will flow in, then it might find it optimal to pay the fixed cost, look for the right investment opportunity and decrease the liquidity level to the optimal one. By doing so the firm pays the fixed cost, but for the next periods, the opportunity

[14] In this version of the model I focus on liquidity independently if it arrives (exits) as cash or trade credit (debt).

cost will be lower. m_2 is exactly the level such that, in expectation, investing the excess liquidity $m_2 - m^*$ is equally costly as bearing the expected opportunity cost associated with $m_t = m_2$ rather than $m_t = m^*$.[15] A similar logic applies to m_1. When a good investment or financing opportunity arrives for free, with an arrival rate λ, the firm optimally adjusts without paying any fixed cost. This implies that the size of the adjustment is random and not always equal to either $m^* - m_1$ or $m_2 - m^*$.

1.4.2 Distribution of liquidity

The optimal policy is associated with an invariant distribution of liquidity $\phi(m)$. In particular, given that $dm_t \sim BM(\mu, \sigma)$ and the optimal policy $\{m_1, m^*, m_2\}$, the invariant distribution associated with the solution solves the following Kolmogorov forward equation:

$$\begin{cases} \lambda f(x) = -\mu f'(x) + \frac{\sigma^2}{2} f''(x) & \forall x \in [m_1, m^*) \\ \lambda g(x) = -\mu g'(x) + \frac{\sigma^2}{2} g''(x) & \forall x \in (m^*, m_2] \end{cases} \tag{1.5}$$

The density is piece-wise defined because the Kolmogorov forward equation does not hold at m^*, which is a reinjection point. Solving equation (1.5), the invariant distribution is:

$$\phi(m) = \begin{cases} f(m) = Ee^{\gamma m} + Fe^{\delta m} & \forall m \in [m_1, m^*) \\ g(m) = Ge^{\gamma m} + He^{\delta m} & \forall m \in (m^*, m_2] \\ 0 & \forall m \notin (m_1, m_2) \end{cases} \tag{1.6}$$

Where $\{E, F, G, H, \gamma, \delta\}$ are constants that depend on the parameters of the model. Details about the derivation of the density and the system of equations used to determine the constants can be found in section 1.6.6 of the Appendix. The interpretation of the invariant distribution is that, for any initial t_0, as $t \to +\infty$, the density that the firm will have liquidity level $m_T = k$ is $\phi(k)$. An alternative intuition is that if N firms populate the economy, and they follow the same management policy, then, as time goes on and $N \to +\infty$, $\phi(m)$ is the density of the liquidity distribution across firms. In the next section, I describe the empirical strategy to estimate the model.

1.4.3 Estimation

To estimate the model I follow a generalized method of moments (GMM) procedure. Since the cross-sectional distribution was stationary before COVID-19, I performed the estimation once, using the moments computed on the distribution of all observations between 2013 and 2019. The parameters to be estimated are $\{\rho, R_1, R_2, c_1, c_2, \lambda, \mu, \sigma\}$. I fix ρ to be 1% annual. As a proxy for R_2, the opportunity cost when the firm has liquid resources, I use the average differential between the median return on investment (ROI) and the Italian short-term interest rate[16]. An alternative measure could be the difference between short-term and long-term interest rates, which could be influenced by factors like monetary policy decisions and financial markets' expectations. It is worth noting that in 2013 and 2014 the spreads between long-term and short-term government bond returns were higher than the median ROI. This is the case because Italy was in the aftermath of the EU sovereign debt

[15]This is one of the boundary conditions, the value matching:

$$V(m_2) = V(m^*) + c(m_2)$$

[16]I use the annualized 3 months interest rate reported by the OECD.

crisis and, in 2013, there were some political turmoils, election included. For these reasons I use the median ROI as a proxy for the opportunity cost, but results are similar using the short-term - long-term spread. ROI follows a cyclical behaviour, falling in 2020 and bouncing back in 2021. The estimation of R_1 is more complicated because this parameter captures the cost of having trade debt, or in general the cost of operating with negative net liquidity. From the literature, it is known that there are explicit costs, like the premium on the price of inputs, and implicit costs, for example in terms of bargaining power. To capture these additional costs, I estimate R_1 by assuming that it is proportional to the opportunity cost $R_2 \to R_1 = \iota R_2$. Therefore, the remaining parameters to be estimated are:

$$\theta = \{c_1, c_2, \lambda, \mu, \sigma, \iota\}$$

Since the solution of the model is numerical, I define a grid $\Theta \subset R_6^+$ and perform the estimation over this grid of values. For each $\theta_i \in \Theta$, I compute the distance between the structural moments and their population equivalent. I use the distance to compute the score and search the global minimum θ^* of the criterion over the grid.[17]

To construct the GMM criterion I use the non-centered moments. The structural moment of order n is[18]:

$$\begin{aligned}\mathbb{E}\left[m^n;\theta\right] = &\left[E\frac{e^{\gamma x}}{\gamma}\sum_{j=0}^{n}\left((-1)^j\frac{(m)^{n-j}}{\gamma^j}\frac{j!}{(n-j)!}\right)\right]_{m_1}^{m^*} + \left[F\frac{e^{\delta x}}{\delta}\sum_{j=0}^{n}\left((-1)^j\frac{(m)^{n-j}}{\delta^j}\frac{j!}{(n-j)!}\right)\right]_{m_1}^{m^*} + \\ &+\left[G\frac{e^{\gamma x}}{\gamma}\sum_{j=0}^{n}\left((-1)^j\frac{(m)^{n-j}}{\gamma^j}\frac{j!}{(n-j)!}\right)\right]_{m*}^{m_2} + \left[H\frac{e^{\delta x}}{\delta}\sum_{j=0}^{n}\left((-1)^j\frac{(m)^{n-j}}{\delta^j}\frac{j!}{(n-j)!}\right)\right]_{m^*}^{m_2}\end{aligned} \tag{1.7}$$

The population counterpart is:

$$M^n = \sum_j \left(\frac{m_j}{j}\right)^n \tag{1.8}$$

The GMM criterion is constructed using the percentage deviation of $\mathbb{E}\left[m^n|\theta\right]$ from M^n, giving equal weight to each moment:

$$J(\theta) = \left(\frac{\mathbb{E}\left[m^n;\theta\right] - M^n}{M^n}\right)' I \left(\frac{\mathbb{E}\left[m^n;\theta\right] - M^n}{M^n}\right) \tag{1.9}$$

θ^* is then defined as:

$$\theta^* \equiv \arg\min_{\theta\in\Theta} J(\theta) \tag{1.10}$$

The optimal set of parameters, given that I match the first 7 moments, is:

$\hat{c_1}$	$\hat{c_2}$	$\hat{\lambda}$	$\hat{\mu}$	$\hat{\sigma}$	R_2	$\hat{\iota}$
0.003	0.003	0.015	0.002	0.0274	0.033	1.125

Table 1.7: Optimal set of parameters and the opportunity cost R_2

The minimum of the GMM criterion is $J(\theta^*) = 0.0185$. The associated optimal policy is:

$$\begin{cases} m_1 = -0.667 \\ m^* = 0.044 \\ m_2 = 0.883 \end{cases} \tag{1.11}$$

[17] Over the grid the minimum is unique.

[18] Without loss of generality, I omit θ in the definition.

$[m_1, m_2]$ is the inaction region, where the representative firm adjusts liquidity only if it has a free adjustment opportunity. Upon adjusting, the firm chooses to hold $m^* = 5\%$.[19] The estimated moments are:

n	1	2	3	4	5	6	7
M^n	0.1596	0.0708	0.0279	0.0176	0.0098	0.0074	0.0049
$\mathbb{E}[m^n \vert \theta^*]$	0.1524	0.0683	0.0297	0.0176	0.0103	0.0069	0.0046

Table 1.8: Estimated moments and their empirical counterparts.

In the following figure, I plot the invariant distribution associated with m_1, m^*, m^2. Notice that the procedure does not aim to match the distribution itself, but the first 7 moments only, so that the GMM system is just over-identified.

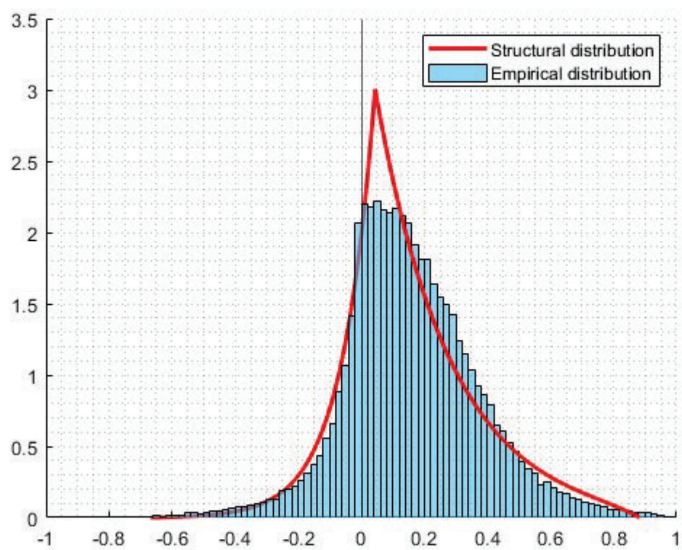

Figure 1.8: Invariant distribution and histogram of the empirical distribution built using all the data from 2013 to 2019. The histogram is normalized to be the empirical probability distribution function.

The invariant distribution has extra mass around the reinjection point, while in the empirical distribution, this mass is slightly shifted to the right.
In the next section, I will analyse the relationship between the fundamental parameters and the optimal policy.

1.4.4 comparative statics

How do firms react to changes in funding costs? and to changes in underlying cash flow uncertainty? To answer this question I depart from θ^* and compare the optimal policy between different steady states, in a

[19]These variables are in percentage of total assets.

comparative statics fashion.
In the literature, it has been shown that the correlation between cash flow uncertainty and cash holdings is increasing, conjecturing that firms prefer to keep more cash buffers. To quantify the impact of cash flow uncertainty on liquidity management, I study the changes in the optimal policy $\{m_1, m^*, m_2\}$ as σ increases from the estimated value $\hat{\sigma}$:[20]

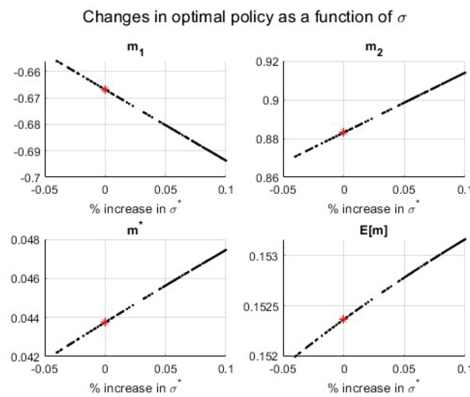

Figure 1.9: Change in $\{m_1, m^*, m_2\}$ and in the average level of liquidity $\mathbb{E}[m]$ as a function of changes in cash flow standard deviation.

Similar to the literature, I find that the average level of liquidity increases with cash flow uncertainty. However, this increase is marginal because the rise in uncertainty has a symmetric effect on both m_1 and m_2.[21]. Hence, an increase in uncertainty leads to a widening of the optimal inaction region, because the firm prefers to bear higher opportunity costs before adjusting as it expects to hit the boundary more often than before. The opposite is true for the free adjustment opportunities:

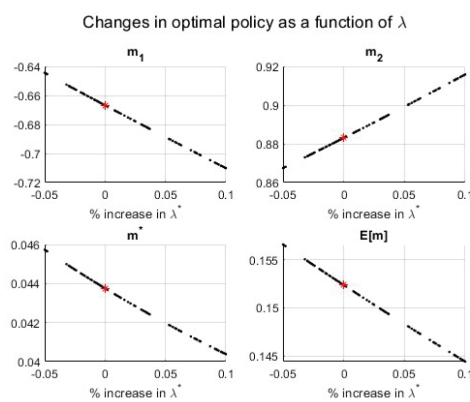

Figure 1.10: Change in $\{m_1, m^*, m_2\}$ and in the average level of liquidity $\mathbb{E}[m]$ as a function of changes in the arrival rate of free adjustment opportunities.

[20] Since the system of boundary condition is not linear, I cannot compute in close form the changes in the optimal policy as a function of σ, e.g. $\frac{\partial m^*}{\partial \sigma}$.
[21] Although it is not clear from the figure, the relations are not linear.

If the firm expects to be able to adjust its liquidity more often, it will find it optimal to wait more before adjusting. This widens the inaction region, as for δ, but decreases the precautionary motives for the firm to keep liquidity buffers: m^* and $E[m]$ decreases. Finally, when the opportunity cost R_2 increases, the firm finds optimal to narrow the inaction region:

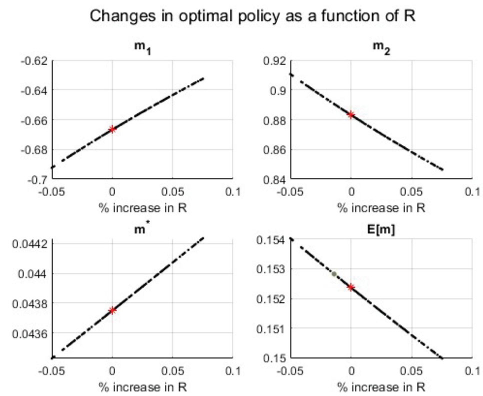

Figure 1.11: Change in $\{m_1, m^*, m_2\}$ and in the average level of liquidity as a function of changes in the opportunity costs.

Given that the inaction region shrinks, the average liquidity decreases too. The optimal level of cash holdings m^* increases because the effect of an increase in R_2 is stronger when liquidity is negative: R_1, the cost of having negative liquidity, is proportionally more by $\hat{\iota}$. If R_1 is fixed to $R_1 = \hat{\iota}\bar{R}_2$ and only R_2 increases, then the optimal level of liquidity decreases too. There is a second-order effect on m_1 too, because the decrease in m^* implies that the firm expects to end up with negative liquidity more often, therefore, it must tolerate bearing a higher liquidity deficit cost.[22]
This analysis shows that changes in financing conditions and cash flow uncertainty have second-order effects that are important to capture for a correct evaluation of the impact of structural changes in the economy. To conclude this section, it is important to highlight that the same observed change in firms' behaviour can be rationalized both with an increase or a decrease in the total cost of liquidity management. For example, we can observe an increase in the average level of liquidity, which is generally associated with an increase in the cost of liquidity management because either σ increased or λ decreased. On the other side, if we observe that firms tolerate more extreme positions in liquidity, this is consistent with both an increase in uncertainty, which leads to an increase in the cost of liquidity management, or with an increase in adjustment opportunities, which is associated with lower costs of management.[23] Future research will be useful to improve the identification of the structural parameters and further disentangle corporate liquidity dynamics.

1.4.5 Generalized version of the model

In this section, I propose a generalized version of the model that accounts for cash holdings and trade credit/debt separately. In the baseline model, the firm chooses its liquidity ratio and decides when to adjust. In this section, I assume that in each period the firm receives a liquidity inflow/outflow. It needs to control its liquidity level by jointly managing cash holdings and the trade position. The inflows/outflows change

[22]It is possible to find the figure in the Appendix, section 1.6.7.

[23]When λ increases the firm expects to reset for free more often, saving on the fixed costs. This reduces, ceteris paribus, the total cost of liquidity management. It is possible to see this from the value function: as λ increases the probability of going to $V(m^*)$ increases.

both variables, and the firm faces different costs. Calling T_t the net trade position, the liquidity level is given by $m_t = ch_t + T_t$. As in the baseline model,

$$dm_t = \mu dt + \sigma dW_t$$

but this time dm_t realizes either as cash flow or as a trade credit/debt:

$$dch_t = \begin{cases} 0 & \textit{With prob. } \eta \\ \mu dt + \sigma dW_t & \textit{With prob. } 1-\eta \end{cases} \tag{1.12}$$

$$dT_t = \begin{cases} 0 & \textit{With prob. } 1-\eta \\ \mu dt + \sigma dW_t & \textit{With prob. } \eta \end{cases} \tag{1.13}$$

So that, in expectation:

$$\mathbb{E}[dch_t] = (1-\eta)\mu dt$$
$$\mathbb{E}[dT_t] = \eta\mu dt$$

The problem of the firm is to decide when to adjust either its trade position, its cash holdings, or both, and the composition of its liquidity upon adjustment:

$$\min_{\{ch_t, T_t, \tau_T, \tau_c, \tau_m\}} \mathbb{E}_t \left[\sum_{\tau_T} e^{\rho\tau_T} c_1 + \sum_{\tau_c} e^{\rho\tau_c} c_2 + \sum_{\tau_m} e^{\rho\tau_m} c_3 + \int_t^{+\infty} e^{\rho t}(R_1 ch_t + R(T_t)T_t)dt \right] \tag{1.14}$$

Where:

1. $\{\tau_T, \tau_c, \tau_m\}$ are the time when the firm decides to adjust the trade position/cash holdings/both respectively;
2. $\{c_1, c_2, c_3\}$ are the fixed costs to be paid to adjust the trade position/cash holdings/both respectively;
3. R_2 is the opportunity cost of cash holdings;
4. $R(T)$ is the opportunity cost of the trade position:

$$R(T) = \begin{cases} R_1 & T < 0 \\ R_2 & T \geq 0 \end{cases}$$

Differently from the baseline model, cash holdings are bounded at 0, which is a reflecting barrier for ch_t. Following the same logic used to derive (1.4) it is possible to compute the HJB associated with (1.14):

$$V(ch_t, T_t) = R_2 dtch_t + R_i(T_t)dtT_t + \frac{1}{1+rdt}\mathbb{E}_t\left[V(ch_{t+d_t}, T_{t+dt})\right] \tag{1.15}$$

From the discrete time approximation:

$$\begin{aligned} \mathbb{E}_t\left[V(ch_{t+d_t}, T_{t+dt})\right] &= V(ch_t, T_t) + \mathbb{E}_t\Big[V'_{ch}(ch_t, T_t)dch_t + V'_T(ch_t, T_t)dT_t + \frac{1}{2}V''_{ch}(ch_t, T_t)dch_t^2 \\ &+ \frac{1}{2}V''_T(ch_t, T_t)d_T^2 + \frac{1}{2}V''_{ch,T}(ch_t, T_t)dch * dT\Big] = \\ &= V(ch_t, T_t) + V'_{ch}(ch_t, T_t)(1-\eta)\mu dt + V'_T(ch_t, T_t)\eta\mu dt + \\ &+ \frac{1}{2}V''_{ch}(ch_t, T_t)(1-\eta)\sigma^2 dt + \frac{1}{2}V''_T(ch_t, T_t)\eta\sigma^2 dt \frac{1}{2}V''_{ch,T}(ch_t, T_t)(1-\eta)\eta\sigma^2 dt \end{aligned} \tag{1.16}$$

Taking limit $dt \to 0^+$ and simplifying terms:

$$\begin{aligned} \rho V(ch_t, T_t) &= R_2 dtch_t + R_i(T_t)dtT_t + V'_{ch}(ch_t, T_t)(1-\eta)\mu + V'_T(ch_t, T_t)\eta\mu + \\ &+ \frac{1}{2}\left[V''_{ch}(ch_t, T_t)(1-\eta)\sigma^2 + V''_{T_t}(ch_t, T_t)\eta\sigma^2\right] \end{aligned} \tag{1.17}$$

where $V''_{ch,T}(ch_t, T_t)$ has been omitted because the cross derivative is 0 by the structure of the problem.

1.5 Conclusions

In this paper I study the liquidity composition of Italian corporations and find that firms keep their liquidity level constant, changing the composition of the assets. I document that firms have on average 15% of their assets liquid, and, during the 2013-2019 period, increased their cash holdings and reduced trade credit allowances, highlighting a substitution pattern. The level of trade debt is stable over the period, while trade credit is declining. When the pandemic shock hit the economy, the net trade position decreased, while cash holdings substantially increased, as the government pursued policies to help firms survive the pandemic.

To disentangle the drivers of liquidity management decisions, and to rationalize the cross-sectional distribution of liquidity, I propose a model of liquidity management based on [5]. I estimate the model matching the structural moments to the data and find that adjustment costs are similar both to rise or to invest liquidity, but the opportunity cost of having negative liquidity is 12.5% higher than the opportunity cost of keeping assets liquid. I confirm the intuition that cash flow uncertainty increases precautionary motives, which causes the firms to increase the average level of cash holdings and to tolerate larger positions before adjusting liquidity. When the opportunity cost increases, the average level of liquidity increases, but the effect is not symmetric because having negative liquidity is associated with a higher cost. Furthermore, changes in structural parameters, have second-order effects on the optimal policy, because, by influencing the optimal level of liquidity, they indirectly influence the tails of the distribution. For example, increasing the opportunity cost only when liquidity is positive has a direct effect on the management policy of firms with an excess of liquidity, but also a second-order effect, because, as the average level of liquidity decreases, firms must tolerate, on average, larger deficits.

1.6 Appendix

1.6.1 Geographical dispersion

The following table 1.9 and figure 1.12 highlight the geographical distribution of the observations.

Figure 1.12: Map of the distribution of firms across Italian regions.

Region	# firms	Region	# firms
Abruzzo	165	Molise	29
Basilicata	51	Piemonte	1257
Calabria	68	Puglia	244
Campania	619	Sardegna	167
Emilia Romagna	1757	Sicilia	275
Friuli Venezia Giulia	361	Toscana	954
Lazio	1049	Trentino Alto Adige	323
Liguria	334	Umbria	180
Lombardia	4993	Valle d'Aosta	26
Marche	362	Veneto	1850

Table 1.9: Table containing the geographical distribution of surviving observations. 5186 did not have a unique geographical identifier.

1.6.2 Kolmogorov Smirnov test

The following table contains the results of the Kolmogorov Smirnov test. For each possible couple of years, I test the hypothesis that the cross-sectional distributions are drawn from two different distributions:

	2022	2021	2020	2019	2018	2017	2016	2015	2014	2013
2022	False	True	True	False	True	False	False	False	False	True
2021	True	False	True	True	True	True	True	True	True	True
2020	True	True	False	True	True	True	True	True	True	True
2019	False	True	True	False	False	False	False	False	False	False
2018	True	True	True	False	False	False	True	False	False	False
2017	False	True	True	False	False	False	False	False	False	True
2016	False	True	True	False	True	False	False	False	False	True
2015	False	True	True	False	False	False	False	False	False	True
2014	False	True	True	False	False	False	False	False	False	True
2013	True	True	True	False	False	True	True	True	True	False

It is possible to see that the test cannot reject the hypothesis that the distributions in 2014-2019 are different.

1.6.3 Trade credit/debt distribution

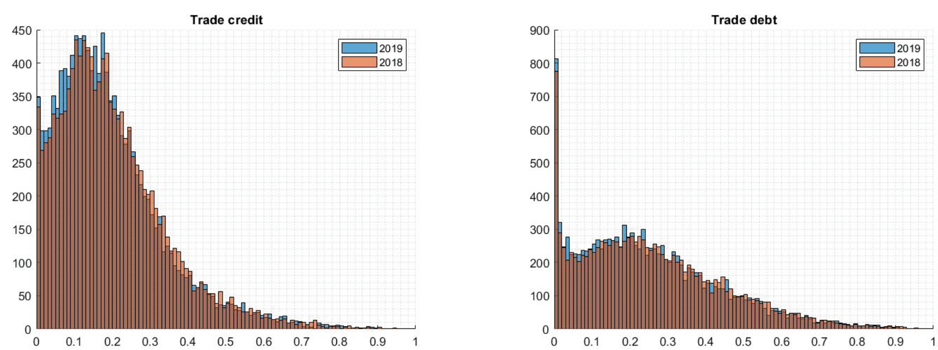

Table 1.10: Trade credit (left) and debt (right) distribution in 2018 and 2019. For trade credit it is possible to see that there is a shift to the left in 2019, meaning that the allowance of trade credit was reduced. Differently, for trade debt, there is no clear shift to the left. Industry averages can be found in Figure 1.6.4.

1.6.4 Liquidity distribution by industries

The next figure contains the distribution of liquidity by sector in a given year. These distributions reflect the differences in the business model. For example, trade companies usually pay their supplier after they have realized their sales, financing themselves using trade credit, as it is possible to see from the 3rd panel of the figure 1.13.

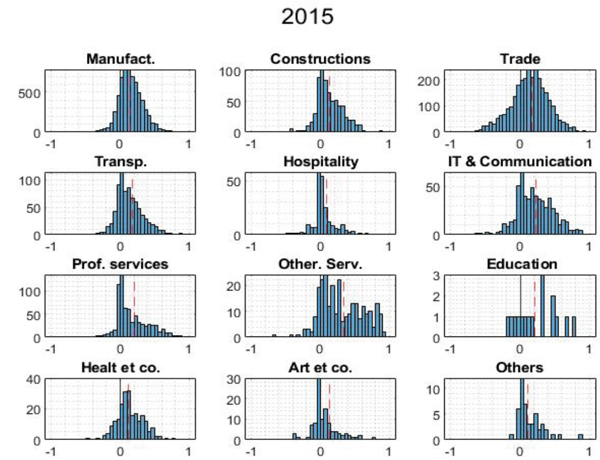

Figure 1.13: An example of the distribution of liquidity across industries in 2015. The histograms are normalized such that on the Y axis there is the number of firms in each bin.

1.6.5 Derivation of the HJB

In this section, I derive the HJB equation from (1.3) The optimal strategy for the firm policy is to adjust only when m_t ends up being too low or too high. Upon adjustment, the firm chooses to set $m_{\tau_j} = m^*$ set such that the expected future costs are minimized. In particular, let us assume that the crucial threshold below which the firm resets its liquidity is m_1, and similarly, m_2 is the crucial threshold value above which the firm prefers to reset. The optimal policy then can be summarized as waiting while $m_t \in [m_1, m_2]$ and paying $c(m)$ when $m_t \notin [m_1, m2]$ to reset to $m^* \in [m_1, m2]$. $\{m*, m1, m2\}$ are unknowns to be determined. To solve for the optima policy let us start with a discrete time formulation of the problem. Inside the inaction region (m_1, m_2), the firm does not control the state unless it can do it for free, with arrival rate λ, and wait for the next period only paying $R(m_t)m_t$. Since the flow cost, $R(m)$, depends on m_t, the value function is piece-wise defined. The following procedure is general for both branches, with the caveat that:

$$V(m_t) = Rm_t + E_t [V(m_{t+1})] \tag{1.18}$$

where m_{t+1} is either $m_t + \Delta m_t$ with probability (1-λ) or m^* with probability λ per unit of time, because the firms adjusts. This implies that:

$$V(m_t) = \Delta R m_t + \frac{1}{1+\Delta r} E_t [\Delta\lambda (V(m^*) - V(m_t)) + (1-\Delta\lambda)V(m_t + \Delta m_t)] \tag{1.19}$$

Taking limit $\Delta \to 0 \Rightarrow \Delta m_t \to dm_t$, and using Itō's lemma to approximate $E_t [V(m_t + dm_t)]$:

$$V(m_t) = dtRm_t + \frac{1}{1+dtr}\left[dt\lambda (V(m^*) - V(m_t)) + (1-dt\lambda)\left[\mu V'(m_t)dt + \frac{1}{2}\sigma^2 dt V''(m_t)\right]\right] \tag{1.20}$$

Rearranging terms and dividing by dt:

$$(r+\lambda)V(m_t) = \lambda V(m^*) + Rm_t + \mu V'(m_t) + \frac{1}{2}\sigma^2 V''(m_t)$$

Which is the Hamilton Jacobi Bellman equation associated with the problem. The solution of the HJB equation will be the value function of the problem. Solving this problem also requires finding the optimal boundaries $\{m_1, m_2\}$ and the optimal resetting point m^*. To solve this differential equation, I use the guess and verify method. In particular, for each branch, I guess that the solution takes the following functional form:

$$V(m) = v + \frac{R}{r+\lambda}m + A_1 e^{a_1 m} + A_2 e^{a_2 m} \tag{1.21}$$

this implies that the first derivative is:

$$V'(m) = \frac{R}{r+\lambda} + A_1 a_1 e^{a_1 m} + A_2 a_2 e^{a_2 m} \tag{1.22}$$

and that the second derivative is:

$$V''(m) = A_1 a_1^2 e^{a_1 m} + A_2 a_2^2 e^{a_2 m} \tag{1.23}$$

substituting (1.21) (1.22) and (1.23) in (1.4) it is possible to find the following equation:

$$(r+\lambda)v+(r+\lambda)(A_1 e^{a_1 m}+A_2 e^{a_2 m}) = \lambda V(m^*)+\frac{R}{r+\lambda}+\mu(A_1 a_1 e^{a_1 m}+A_2 a_2 e^{a_2 m})+\frac{\sigma^2}{2}(A_1 a_1^2 e^{a_1 m}+A_2 a_2^2 e^{a_2 m}) \tag{1.24}$$

Notice that both hand sides have a linear part and an exponential part, thus, matching parts:

$$\begin{cases} (r+\lambda)v = \lambda V(m^*) + \frac{\mu R}{r+\lambda} \\ (r+\lambda)(A_1 e^{a_1 m} + A_2 e^{a_2 m}) = \mu(A_1 a_1 e^{a_1 m} + A_2 a_2 e^{a_2 m}) + \frac{\sigma^2}{2}(A_1 a_1^2 e^{a_1 m} + A_2 a_2^2 e^{a_2 m}) \end{cases} \tag{1.25}$$

Solving the system it is possible to derive the coefficients a_1 and a_2:

$$a_{1,2} = \frac{-\mu \pm \sqrt{\mu^2 + 2\sigma^2(r+\lambda)}}{\sigma^2}$$

which have opposite signs because $\mu^2 \leq \mu^2 + 2\sigma^2(r+\lambda)$ and σ, r and λ are positive. The constant is equal to:

$$v = \frac{\lambda}{r+\lambda}V(m^*) + \frac{\mu R}{(r+\lambda)^2}$$

To solve the HJB there are still 7 unknowns $\{m^*, m_1, m_2\}$ and 2 constants $\{A_1, A_2\}$ for each branches. For simplicity, let us rename A_1, A_2 for the second positive branch as $\{B_1, B_2\}$. To determine the unknowns, let us consider the boundary conditions implied in the problem:

1. Optimality at m^* since it is a minimum point;
2. Value matching at the lower boundary because when $m \leq m_1$ the agents resets to m^*;
3. Value matching at the lower boundary because when $m \geq m_2$ the agents resets to m^*;
4. Smooth pasting at the upper boundary m_1.
5. Smooth pasting at the lower boundary m_2.
6. Continuity at 0.
7. Differentiability at 0.

which are:

1. $$V'(m^*) = 0 \iff \frac{R}{r+\lambda} + A_1 a_1 e^{a_1 m^*} + A_2 a_2 e^{a_2 m^*} = 0$$
2. $$V(m_1) = V(m^*) + c_1 \iff \frac{R_1}{r+\lambda}m_1 + A_1 e^{a_1 m_2} + A_2 e^{a_2 m_2} = \frac{R_1}{r+\lambda}m^* + B_1 e^{a_1 m^*} + B_2 e^{a_2 m^*} + c_1$$

3.

$$V(m_2) = V(m^*) + c_2 \iff \frac{R_2}{r+\lambda}m_2 + B_1 e^{a_1 m_2} + B_2 e^{a_2 m_2} = \frac{R_2}{r+\lambda}m^* + B_1 e^{a_1 m^*} + B_2 e^{a_2 m^*} + c_2$$

4.

$$V'(m_1) = 0 \iff \frac{R_1}{r+\lambda} + A_1 a_1 e^{a_1 m_2} + A_2 a_2 e^{a_2 m_2} = 0$$

5.

$$V'(m_2) = 0 \iff \frac{R_2}{r+\lambda} + B_1 a_1 e^{a_1 m_2} + B_2 a_2 e^{a_2 m_2} = 0$$

6.

$$\lim_{m \uparrow 0} V(m) = \lim_{m \downarrow 0} V(m) \iff v_1 + A_1 + A_2 = v_2 + B_1 + B_2$$

7.

$$\lim_{m \uparrow 0} V'(m) = \lim_{m \downarrow 0} V'(m) \iff \frac{R_1}{r+\lambda} + A_1 a_1 + A_2 a_2 = \frac{R_2}{r+\lambda} + B_1 a_1 + B_2 a_2$$

This is a non-linear system of 7 equations with 7 unknowns that can be solved numerically. These boundary conditions are derived assuming that $m^* > 0$, however, in the code I generalize them such that $m^* < 0$ is possible. The estimated model will always be such that $m_t > 0$.

1.6.6 KFE derivation and solution

The second part of the model consists in estimating an implied cross-sectional distribution. This is possible because of the hypothesis that cash balances evolve following a Brownian motion. This is convenient because, at the steady state, it implies a stationary distribution that can be compared with the empirical cross-sectional distribution. It is possible to derive this distribution by solving the Kolmogorov forward equation (KFE) associated to the firm's problem:

$$\begin{cases} 0 = -\mu f'(x) + \frac{\sigma^2}{2} f''(x) - \lambda f(x) & \forall x \in [m_1, m^*) \\ 0 = -\mu g'(x) + \frac{\sigma^2}{2} g'(x) - \lambda g(x) & \forall x \in (m^*, m_2]) \end{cases} \tag{1.26}$$

where $f(x)$ is the density if $m_t \in [m_1, m^*)$ and $g(x)$ is the density if $m_t \in (m^*, m_2]$. The main intuition behind these equations is that in steady state each instant the probability of m arriving in an infinitesimal interval is equal to the probability of leaving that interval. It is immediate to see that the solutions to the ODEs are exponential functions:

$$f(x) = Ee^{\gamma x} + Fe^{\delta x} \tag{1.27}$$

$$g(x) = Ge^{\gamma x} + He^{\delta x} \tag{1.28}$$

Using the same strategy used to determine $a_{1,2}$, it is possible to compute $\{\gamma, \delta\}$:

$$\{\gamma, \delta\} = \frac{\mu \pm \sqrt{\mu^2 + 2\sigma^2 \lambda}}{\sigma^2}$$

Given $\{\gamma, \delta\}$, the constants to be determined are $\{E, F, G, H\}$. I use the following 4 boundary conditions associated with the problem:

1. No mass at the lower barrier;
2. No mass at the upper barrier;
3. Continuity at m^*;
4. Integrability to 1 since the solution is a density.

The boundary conditions translate into the following system:

$$\begin{cases} f(m_1) = 0 \iff Ee^{\gamma m_1} + Fe^{\delta m_1} = 0 \\ g(m_2) = 0 \iff Ge^{\gamma m_2} + He^{\delta m_2} = 0 \\ f(m^*) = g(m^*) \iff Ee^{\gamma m^*} + Fe^{\delta m^*} = Ge^{\gamma m^*} + He^{\delta m^*} \\ \int_{m_1}^{m^*} f(x)dx + \int_{m^*}^{m_2} g(x)dx = 1 \iff \left[\frac{E}{\gamma}e^{\gamma x} + \frac{F}{\delta}e^{\delta x}\right]_{m_1}^{m^*} + \left[\frac{G}{\gamma}e^{\gamma x} + \frac{H}{\delta}e^{\delta x}\right]_{m^*}^{m_2} = 1 \end{cases} \tag{1.29}$$

1.6.7 Changes in opportunity cost only

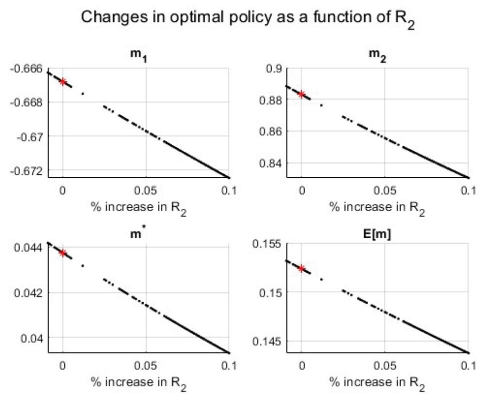

Figure 1.14: Changes in optimal policy as the opportunity cost of holding liquidity changes, keeping the cost of liquidity deficits fixed.

References

[1] T. Bates, K. Kahle, and R. Stulz. "Why Do U.S. Firms Hold so Much More Cash Than They Used to". In: *The Journal of Finance* 64 (Apr. 2009).

[2] B. Nikolov, L. Schmidb, and R. Steri. "Dynamic corporate liquidity". In: *Journal of Financial Economics* 132 (2019), pp. 76–102.

[3] N. Amberg et al. "Curbing Shocks to Corporate Liquidity: The Role of Trade Credit". In: *Journal of Political Economy* 129 (2021), pp. 182–242.

[4] N. Amberg et al. "Why Trade Credit?" Working paper.

[5] M. Miller and D. Orr. "A Model of the Demand for Money by Firms". In: *The Quarterly Journal of Economics* 80.3 (1966), pp. 413–435.

[6] B. Bureau, A. Duquerroy, and F. Vinas. "Activity Shocks and Corporate Liquidity: the Role of Trade Credit". In: *The Review of Corporate Finance Studies* 13 (Aug. 2024), pp. 668–711.

[7] V. Acharya, S. A. Davydenko, and I. A. Strebulaev. "Cash Holdings and Credit Risk". In: *The Review of Financial Studies* 25.12 (Dec. 2012), pp. 3572–3609.

[8] H. Gao, J. Harford, and K. Li. "Determinants of corporate cash policy: Insights from private firms". In: *Journal of Financial Economics* 109.3 (2013), pp. 623–639.

[9] V. Potì, P. Pattitoni, and B. Petracci. "Precautionary motives for private firms' cash holdings". In: *International Review of Economics and Finance* 68 (2020), pp. 150–166.

[10] M. Bigelli and J. Sánchez-Vidal. "Cash holdings in private firms". In: *Journal of Banking and Finance* 36.1 (2012), pp. 26–35.

[11] J. A. Frenkel and B. Jovanovic. "On Transactions and Precautionary Demand for Money". In: *The Quarterly Journal of Economics* 95 (Aug. 1980), pp. 25–43.

[12] F. Alvarez and F. Lippi. "Financial Innovation and the Transactions Demand for Cash". In: *Econometrica* 77.2 (2009), pp. 363–402.

[13] I. Baley and A. Blanco. "Aggregate Dynamics in Lumpy Economies". In: *Econometrica* 89.3 (2021), pp. 1235–1264.

[14] L. A. Riddick and T. M. Whited. "The Corporate Propensity to Save". In: *The Journal of Finance* 64.4 (2009), pp. 1729–1766.

[15] P. Bolton, H. Chen, and N. Wang. "A Unified Theory of Tobin's q, Corporate Investment, Financing, and Risk Management". In: *The Journal of Finance* 13 (Sept. 2011).

[16] B. Biais and C. Gollier. "Trade Credit and Credit Rationing". In: *The Review of Financial Studies* 10(4) (1997), pp. 903–937.

[17] B. S. Wilner. "The Exploitation of Relationships in Financial Distress: The Case of Trade Credit". In: *The Journal of Finance* (2007).

[18] V. Cuñat. "Trade Credit: Suppliers as Debt Collectors and Insurance Providers". In: *The Review of Financial Studies* 20 (Mar. 2007), pp. 491–527.

[19] J. H. Nilsen. "Trade Credit and the Bank Lending Channel". In: *Journal of Money, Credit and Banking* 34 (Feb. 2002), pp. 226–253.

[20] M. Giannetti, N. Serrano-Velarde, and E. Tarantino. "Cheap Trade Credit and Competition in Downstream Markets". In: *Journal of Political Economy* 129.6 (2021), pp. 1744–1796.

[21] W. Baumol. “The Transactions Demand for Cash: An Inventory Theoretic Approach”. In: *The Quarterly Journal of Economics* 66 (4) (1952), pp. 545–56.

[22] J. Tobin. “The Interest-Elasticity of Transactions Demand For Cash”. In: *The Review of Economics and Statistics* 38 (3) (2017), pp. 241–7.

Chapter 2

Algorithmic Pricing and Bank Lending

Carmelo Genovese [1] **Fabiano Schivardi** [2] **Enrico Sette**[3] **Emanuele Tarantino**[4]

2.1 Abstract

We study the impact of artificial intelligence in banking markets. We estimate a structural model of demand and supply for the credit lines market, characterized by information asymmetries and imperfect competition. As in the literature, we found evidence of adverse selection and moral hazard. We then simulate the impact of the adoption of AI pricing algorithms at the bank level on market outcomes in a counterfactual exercise. In traditional concentrated markets, these algorithms learn to sustain collusive prices. However, charging high markups might not be optimal in selection markets. We find that AI generally increases prices also in these markets, but the relative increase strongly depends on the level of asymmetries. Additionally, we experiment with the robustness of the Q-learning algorithm as our application poses challenges to its learning process.

[1] Luiss Guido Carli.

[2] Luiss Guido Carli and EIEF.

[3] Banca d'Italia. Disclaimer: This paper does not necessarily represent the view of the Bank of Italy.

[4] Luiss Guido Carli, EIEF, and European Commission. Disclaimer: This paper does not necessarily represent the view of the European Commission.

We would like to thank Emilio Calvano, Giacomo Calzolari, Cristina Gualdani, Luigi Pascali, and Nicola Pavanini.

2.2 Introduction

AI is considered a general purpose technology, as it is changing our world. Understanding its impact on the economy is a key challenge to address both the issues and the challenges that the adoption of this technology will have on every economic agent. The analysis of the economic impact of AI has already started, for example, in the famous paper by Calvano et al.[1]{CCDP}, which proves that under certain conditions AI can sustain over competitive prices. In this project, we study the impact of AI adoption in financial markets. We focus on the banking market for credit lines, and we simulate AI pricing tools adoption within this market. We think that this is a relevant analysis because banking markets are selection markets, meaning that they are characterized by adverse selection and moral hazard. We study how these failures interact with the adoption of AI, which might generate another failure. The stability of the banking system is crucial for the economy, therefore understanding the impact of this technology on risk-taking and credit rationing can help policymakers to regulate its adoption efficiently.
To simulate the adoption of AI pricing algorithms we first estimate a state-of-the-art model of demand and supply for credit lines, then we perform a counterfactual analysis where we assume that banks adopt AI to price their contracts. In particular, we follow the model proposed by Crawford et al.[2]{CPS.} that explicitly accounts for adverse selection and moral hazard in the equilibrium determination. After the estimation of the model, we run a counterfactual analysis where we introduce AI pricing tools that solve the pricing problem of the adopting bank. AI learns to play the pricing game by strategically interaction with other banks, which might have adopted AI as well. Within the structure of the model we can control what drives the level of asymmetries, and, in general, the fundamentals of the model, and relate the equilibrium with AI adoption to other equilibria. This exercise allows us to disentangle the interaction between AI adoption, moral hazard and adverse selection. We begin with a stylized exercise where the environment is closer to the one studied in [1], and find that when two banks adopt AI, they generally sustain over competitive prices. We then run a comparative analysis with different degrees of adverse selection and find that AI's optimal policy is to decrease the relative cost of credit when asymmetries are higher, not rationing the marginal borrower. However, since our model includes different markets, there is heterogeneity in the results. Differently from the other papers in the literature, in our model banks can have losses when borrowers default. This poses a learning challenge to the AI algorithms because they approximate the optimal policy by gradual learning, and this learning might be harmed if defaults occur frequently. We experiment on the learning process by introducing state-dependent lotteries. We find that the algorithm needs sensibly more time to converge [...]

2.2.1 Literature review

Our work is related to two main strands of the literature, the literature on industrial organization of financial markets and the literature on AI and economics.
In the finance literature, the presence and the impact of asymmetries of information á la Akerlof[3], Rothschild and Stiglitz[4], and Stiglitz and Weiss[4] have been widely studied. These market failures influence access to credit and credit allocation and can harm financial stability. Recently, the focus has shifted toward the interaction between asymmetries of information and other market failures. CPS analysed the interaction between adverse selection and market power and found that banks with higher market power may give up part of their mark-ups to mitigate the consequences of adverse selection. The idea that market structure is a crucial element in the determination of the market equilibrium dates back to Petersen and Rajan [5]. We contribute to this literature by analysing the interaction between market concentration, asymmetric information, and AI adoption. While AI

represents a technological change and not a market failure, it has been pointed out by [1] that its adoption in concentrated markets can lead to tacit collusion, increasing mark-ups. In this paper, we include information asymmetries to understand whether the adoption of AI can worsen market allocation in selection markets. In traditional markets, AI increases markups, while in selection markets this might not be always optimal for the banks, as the analysis in CPS suggests.

2.3 Structural model

To disentangle the interaction between asymmetric information, market power and AI adoption we need a model where all these features are explicitly included. In particular, we need a demand and supply model that allows borrowers to have private information, banks to have market power, and us to introduce AI in the counterfactual analysis. For these reasons we follow the model and the methodology proposed by [2]. Their model combines the methodology to estimate demand in the spirit of Berry et al.[6] with the model proposed by Einav et al.[7]. Following CPS, we focus on first-time borrowers, and we only analyse the market for credit lines, which are not collateralized. Firms would like to borrow a certain quantity from the banks in their market, and the banks propose their price. Competition is assumed to be á la Nash-Bertrand. The borrower faces three choices. First, she chooses from which bank to borrow[1], then, conditional on borrowing from a given bank, how much to borrow and whether to default or not. The decisions will depend on a private component $\boldsymbol{\varepsilon}$ that is not observable by the bank and is the source of asymmetric information. Formally, the problem of the firm i is to choose the lender, how much to borrow, and whether to default maximizing the following utilities:

- Utility of borrowing from bank j (B):

$$U^B_{ijmt} = \bar{\alpha}^B_0 + X'^B_{jmt}\beta^B + \xi^B_{jmt} + \alpha^B P_{ijmt} + Y^B_{ijmt}\eta^B + \varepsilon^B_i + \nu_{ijtm} \tag{2.1}$$

- Utility of using the credit line (U):

$$U^U_{ijmt} = \bar{\alpha}^U_0 + X'^U_{jmt}\beta^U + \alpha^U P_{ijmt} + Y'^U_{ijmt}\eta^U + \varepsilon^U_i \tag{2.2}$$

- Utility from default (D):

$$U^D_{ijmt} = \bar{\alpha}^D_0 + X'^D_{jmt}\beta^D + \alpha^D P_{ijmt} + Y'^D_{ijmt}\eta^D + \varepsilon^D_i \tag{2.3}$$

where j is the bank, m stands for market, and t for year.[2] Each choice depends on observables at the firm level Y_{ijmt} and at the bank level X_{jmt}, the price P_{ijmt} and the unobserved components ε and ξ^B_{ijmt}. The latter represents the bank's characteristics known in the market but not observable in the data, while the former is the firm's private information. ν_{ijmt} is an unobservable shock to demand. Competition on the market is only on prices, while the granted amount is fixed and equal to the one we observe. Conditional on borrowing from a certain bank j, the firm decides the amount of the granted credit loan to use. The source of asymmetry in the loan use choice is ε^U_i. After deciding the lender and how much to borrow, the firm is left with the decision to repay or not the debt. The decision will be to default if the utility from defaulting is positive $U^D_{ijmt} > 0$.
Notice how the three frictions of interest are included in the model. First, banks have market power because they are differentiated and borrowers value their characteristics (X^B_{jmt}, ξ^B_{jmt}) in deciding

[1]The firm can also decide not to borrow as the outside option.
[2]Bank $j = 0$ represent the outside option. Firm i does not borrow if $U_{i0mt} > U_{ijmt} | j \neq 0$.

their lender. Second, asymmetries of information, in particular $\boldsymbol{\varepsilon} = \{\varepsilon_i^B, \varepsilon_i^U, \varepsilon_i^D\}$, lead to adverse selection. The underlying assumption is that each component is 0 mean, as they are unobservable, but can be correlated. Suppose the correlation between the unobserved component in demand (2.1) and in default (2.3) is positive. In that case, there is adverse selection because the firms that are more likely to borrow are also the ones that are more likely to default.[3] In particular, we assume that: $\boldsymbol{\varepsilon} \sim N(\boldsymbol{\mu}, \boldsymbol{\Sigma})$:

$$\begin{pmatrix} \varepsilon_i^B \\ \varepsilon_i^U \\ \varepsilon_i^D \end{pmatrix} \sim N \left(\begin{pmatrix} 0 \\ 0 \\ 0 \end{pmatrix}, \begin{pmatrix} \sigma_B^2 & \rho_{BU}\sigma_B\sigma_U & \rho_{BD}\sigma_B \\ \rho_{BU}\sigma_B\sigma_U & \sigma_U^2 & \rho_{UD}\sigma_U \\ \rho_{BD}\sigma_B & \rho_{UD}\sigma_U & 1 \end{pmatrix} \right)$$

where the adverse selection parameter is ρ_{BD}. $\boldsymbol{\Sigma}$ is jointly estimated with the demand and supply parameters. The last friction is moral hazard, captured by α^D, the sensitivity of the default choice to the price.
On the supply side, banks are differentiated and compete a la Nash-Bertand by choosing the price of the credit line. Their problem is to maximize the expected profits:

$$\Pi_{ijmt} = P_{ijmt}Q_{ijmt}\left(1 - F_{ijmt}(P_{ijmt})\right) - MC_{ijmt}Q_{ijmt} \tag{2.4}$$

where Q_{ijmt} is the expected quantity used, P_{ijmt} is the price charged, MC_{ijmt} is the marginal cost and F_{ijmt} is the probability of default. When choosing the price, each bank knows that the probability of default will be endogenous, but the sources of this endogeneity are two:

1. (Direct) The price charged by the bank influences the default decision (α^D);
2. (Indirect) The price charged by the bank influences the pool of borrowers (ρ_{BD}).

Both channels are taken into account by the banks, which will price the contract accordingly:

$$P_{ijmt} = \frac{MC_{ijmt}}{1 - F_{ijmt} + F'_{ijmt}M_{ijmt}} + \frac{(1 - F_{ijmt})M_{ijmt}}{1 - F_{ijmt} + F'_{ijmt}M_{ijmt}} \tag{2.5}$$

where F'_{ijmt} is the derivative of the default probability with respect to the price, and M_{ijmt} is the markup $-Q_{ijmt}/Q'_{ijmt}$. Equation (2.5) allows us to compute the implied marginal costs since we have P_{ijmt} and we can compute numerically F'_{ijmt} and Q'_{ijmt}.
In the next sections, we discuss in detail the data we use, the variables we include, and the estimation strategy.

2.3.1 Data

To estimate the model we use detailed microdata provided by the Bank of Italy. We use the credit registry data from 2010 to 2018. This dataset contains information about every contractual relationship between firms and banks, conditional on the total exposure being more than 75.000€. For the credit lines, the product we focus on, we observe the amount granted, the amount used by the firm, the stream of payments associated to the line, and delinquecies.[4] We have detailed balance sheet data for firms and banks. For banks, we have financial aggregates at the province level that we use as controls and to construct instruments for the estimation. At the province level, we have the number of deposits, the number of clients, the number of NPLs, information on the costs of deposits

[3] This is the definition of adverse selection following from [8]
[4] Following [4], we impute default to a credit line if there are delinquencies within 3 years.

and banks' demographics. We define a market as a province (m) year (t). A market is active if we observe at least one new borrower using a credit line. The banks active in each market are those who grant new credit lines.
We face a challenge in identifying which firms are non-borrowers in the market for the first credit line. Each year, most of the young firms in a market are non borrowers, but we need to distinguish between non-borrowers who are "active" in the market and non-borrowers who are not interested/capable of borrowing. To screen "active" non-borrowers from others, we impose that they must have more than 1 employee and be younger than 3 years[5]. Furthermore, we exclude from the sample "S.R.L.s.", a particular type of limited liability company that follows special accounting laws. However, after these cleanings, we are left with a non-borrower share ≥95% in many markets. We decided to do random sampling so that the non-borrowers share is at most 66%. We are working on strategies to improve the identification of the "active" non-borrowers. At the moment we are exploiting the fact that some firms are borrowing from banks but are below the 75.000€ threshold.[6] Calling these firms "non-registry" borrowers, our idea is to match the firms that report 0 borrowings with both "non-registry" and standard borrowers and exclude the ones that are relatively more similar to the "non-registry". The rationale is that these firms are not in the market of interest because their characteristics suggest that, if they borrow, they will be granted an amount below the credit registry threshold. Since we are still working to understand the effectiveness of this procedure, in the paper we show results that are obtained with random sampling.
A major caveat is that to estimate the model, we need to observe the prices of every possible bank-firm combination in each market and each year. However, this is not possible because some firm-bank relationships do not exist. Following the literature, we overcome this caveat by constructing prices for bank-firm relationships that we do not observe in the data. First, we use a propensity score matching procedure to match similar borrowers, we assign them with a firm fixed effect[7] and an amount granted, then we use the firm fixed effect to estimate a pricing equation:

$$P_{ijmt} = \gamma_0 + \gamma_1 \mathcal{D}_{ijmt} + \gamma_2 \mathcal{L}_{ijmt} + \lambda_{jmt} + \omega_i + \tau_{ijmt} \tag{2.6}$$

where $\mathcal{D}_{ijmt}$ is the distance between the firm and the bank, $\mathcal{L}_{ijmt}$ are dummies for the size of the credit line, λ_{jmt} is the market year fixed effect, and ω_i is the firm fixed effect. Using this pricing equation, we have all the bank-firm prices needed to estimate the model.
To estimate the model we include several variables at the firm and bank level that influence the borrowing process. At the firm level (Y_{ijmt}), we include total assets, profits, cash flow, trade debt, and the share of intangible assets. We also include the distance between banks and firms.[8] As banks' characteristics (X_{jmt}), we include the number of branches in the market, the share of branches in the market, and the number of years the bank has been active within that market. The idea is that each of these observables can directly influence the borrowing/lending decision. Finally, we include the firm, score[9], and bank/market/year fixed effects.

[5] An age comparable with the median borrower.
[6] We identify these firms because they are not in the credit registry but have positive bank debit in the balance sheet.
[7] We exploit an institutional feature of the Italian economy, where firms usually have multiple lending relationships from their first years.
[8] We use geographical coordinates to measure the distance between the municipality of the closest bank's branch and the municipality of the firm.
[9] Score is a discrete variable 1-10 that indicates the quality of the borrower using balance sheet information. The Bank of Italy provides it.

2.3.2 Estimation

We estimate the model using the maximum simulated likelihood, and we use instruments to identify $\{\alpha^B, \alpha^U, \alpha^D\}$. After including the pricing equation in equation (2.1), we can rewrite it as

$$U_{ijmt} = \underbrace{\delta_{jmt} + \alpha^B \tilde{P}_{jmt}}_{\tilde{\delta}_{jmt}} + \underbrace{Y^B_{ijmt}\tilde{\eta}_{ijmt} + \varepsilon^B_i}_{V^B_{ijmt}} + \zeta_{ijmt} \tag{2.7}$$

where δ_{ijmt} is the standard mean utility in the spirit of [6], and $\alpha^B \tilde{P}_{jmt}$ comes after substituting the pricing equation within the model. $\tilde{\eta}_{ijmt}$ is a combination of η and α^B. Finally, ζ_{ijmt} is the sum of ν_{ijmt} and the pricing error, and we assume that it follows a Type I extreme value distribution. Since α^B enters in $\tilde{\eta}_{ijmt}$ in combination with η, it cannot be estimated directly in the maximum simulated likelihood procedure, therefore it is estimated separately in a two-stage IV regression exploiting the mean utility decomposition $\hat{\tilde{\delta}}_{jmt}$ and using proxies for banks' costs as instruments. The likelihood function depends on the probability of demand, the conditional probability of the amount used and the conditional probability of default. Since $\zeta \sim$ *Type* I, the demand probability takes the standard logit form:

$$Pr^D_{ijmt} = \int \frac{e^{\left(\hat{\delta}_{jmt} + V_{ijmt}\right)}}{1 + \sum_{k \neq \{j,0\}} e^{\left(\hat{\delta}_{kmt} + V_{ikmt}\right)}} f(\varepsilon^B_i)\partial\varepsilon^B_i \tag{2.8}$$

where the utility of the outside option is normalized to 0 and $f(\varepsilon^B_i)$ is the density of ε^B_i. To estimate (2.2) and (2.3) we only focus on the lending relationship that took place, meaning that we do not need to use the pricing equation and that we can directly estimate all the parameters. However, we have to address the omitted variable bias that would reflect on $\{\alpha^U, \alpha^D\}$. For this reason, we use market-year fixed effects and the Hausmann instrument. To estimate the loan use equation, conditional on demand $B = 1$ and utilization L, we use:

$$Pr^U_{ijmt|B=1} = \int \phi_{\varepsilon^U_i|\varepsilon^B_i} \left(\frac{L_{ijmt} - \alpha^U_0 - X^U_{jmt}\beta^U - \alpha^U P_{ijmt} - Y^U_{ijmt}\eta^U - \tilde{\mu}_{\varepsilon^U_i|\varepsilon^B_i}}{\tilde{\sigma}_{\varepsilon^U_i|\varepsilon^B_i}} \right) f(\varepsilon^B_i|B=1)\partial\varepsilon^B_i \tag{2.9}$$

where $\phi_{\varepsilon^U_i|\varepsilon^B_i}$ is the associated probability density function:

$$\varepsilon^U_i|\varepsilon^B_i \sim N\Big(\underbrace{\frac{\sigma_L}{\sigma_D}\rho_{BU}\varepsilon^B_i}_{\tilde{\mu}_{\varepsilon^U_i|\varepsilon^B_i}}, \underbrace{\sigma^2_U\left(1-\rho^2_{BU}\right)}_{\tilde{\sigma}_{\varepsilon^U_i|\varepsilon^B_i}}\Big)$$

A similar logic applies to the probability of default, but this time the density is conditioned on both demand and amount used:

$$Pr^D_{ijmt} = \int \Phi_{\epsilon^D_i|\varepsilon^B_i,\varepsilon^U_i} \left(\frac{\alpha^D_0 + X^D_{jmt}\beta^D + \alpha^D + Y^D_{ijmt}\eta^D - \tilde{\mu}_{\varepsilon^D_i|\varepsilon^B_i,\varepsilon^U_i}}{\tilde{\sigma}_{\varepsilon^D_i|\varepsilon^B_i,\varepsilon^U_i}} \right) f(\varepsilon^B_i|B=1)\partial\varepsilon^D_i \tag{2.10}$$

where $\Phi_{\epsilon^D_i|\varepsilon^B_i,\varepsilon^U_i}$ is the associated cumulative distribution function:

$$\varepsilon^D_i|\varepsilon^B_i,\varepsilon^U_i \sim N\Big(\underbrace{A\varepsilon^B_i + C\varepsilon^U_i}_{\tilde{\mu}_{\varepsilon^D_i|\varepsilon^B_i,\varepsilon^U_i}}, \underbrace{\sigma^2_D - (A\rho_{BD} + C\rho_{UD})}_{\tilde{\sigma}_{\varepsilon^D_i|\varepsilon^B_i,\varepsilon^U_i}}\Big)$$

Calling d_{ijmt} the indicator for the actual credit line choice, and f_{ijmt} the indicator for default, the joint log-likelihood function l is:

$$l = \sum_i d_{ijmt} \left\{ log(Pr^B_{ijmt}) + log(Pr^U_{ijmt}) + f_{ijmt} log(Pr^D_{ijmt}) + (1 - f_{ijmt}) log(1 - Pr^F_{ijmt}) \right\} \quad (2.11)$$

We estimate the log-likelihood using 100 Halton random draws. After the simulation, we use IV to perform a two stage regression and identify α^B. First, using the contraction method proposed by [6], it is possible to find $\hat{\delta}_{jmt}$, which can be decomposed as follows:

$$\hat{\delta}_{jmt} = \bar{\alpha}^B_0 + \alpha^B \tilde{P}_{jmt} + X^B_{jmt} \beta^B + \xi^B_{jmt} \quad (2.12)$$

where ξ^B_{jmt} represents unobserved characteristics of the bank that borrowers value in their decision. We decompose this term using bank, market and year fixed effects:

$$\hat{\delta}_{jmt} = \bar{\alpha}^B_0 + \alpha^B \tilde{P}_{jmt} + X^B_{jmt} \beta^B + \xi^B_j + \xi^B_m + \xi^B_t + \Delta\xi^B_{jmt} \quad (2.13)$$

Finally, we estimate this equation using instrumental variables to address the remaining endogeneity issues. We use the number of deposits, the cost of deposits and the amount of deposits as cost shifters. Further information and details about the structural model and the estimation routine can be found in [2].

2.3.3 Results

Our results are in line with the literature. We report the coefficients of interest in the following table:

Model		Price	Tot. Ass.	Intang./T.A.	Profits	Cashflow	Trade debit	Distance	
Demand		**-11.0530** (3.3101)	0.5797 (0.0312)	0.4600 (0.0362)	0.0043 (0.0060)	0.0110 (0.0144)	-0.3471 (0.0226)	-0.4784 (0.0228)	
Loan use		**0.1266** (0.0042)	0.0032 (0.0009)	0.0043 (0.0016)	-0.0005 (0.0002)	0.0012 (0.0008)	-0.0082 (0.0016)	-0.0339 (0.0047)	
Default		**0.1514** (0.0283)	0.0432 (0.0059)	-0.0042 (0.0095)	-0.0033 (0.0023)	-0.0071 (0.0053)	-0.0340 (0.0096)	-0.2590 (0.0255)	
Σ	Demand	Loan Use	Default	score FE	Sector FE	Amount FE	BMY FE	Bank FE	Market FE
Demand	0.9454 (8.6256)			Yes	Yes	Yes	Yes	No	No
Loan use	0.1138 (0.0414)	0.2984 (0.0026)		Yes	Yes	Yes	No	Yes	Yes
Default	**0.1318** (0.0414)	0.0695 (0.0046)	1 -	Yes	Yes	Yes	No	Yes	Yes

Table 2.1: Main coefficient of interest from equations (2.1), (2.2), and (2.3). The bottom left of the table contains estimated correlations (Σ). Standard errors in brackets. Number of observations: 115226 for demand, 5765 for the used amount and default. Variables are scaled by their 95^{th} percentile.

We find evidence of adverse selection $\rho_{BD} = 13\%$ and moral hazard $\alpha^D = 0.1514$. In particular, firms with a positive propensity to borrow have, on average, a high propensity to default. It is worth noting that $\alpha^U > 0$, while, in theory, the amount used should be decreasing in the price of the line. We are currently working on understanding why this coefficient is positive and significant. In any case, the economic magnitude is small and in the counterfactual analysis, when we vary P_{ijmt}, is

about 2-3% in terms of loan use.
The overall fit of the model is summarized in the following table:

	Demand		Loan use		Default	
	Mean	Std.	Mean	Std.	Mean	Std.
Data	0.0977	0.2969	0.6509	0.3142	0.1275	0.3336
Model	0.0977	0.2204	0.6499	0.1287	0.1332	0.0900

Table 2.2: This table contains the average demand choice ($\mathrm{E}\left[d_{ijmt}\right]$), the average usage rate, and the average default probability for the model and the one observed in the data.

2.4 AI adoption

In this section, we analyse the impact of artificial intelligence on bank lending. We use the estimated model to evaluate the effects of AI pricing tools on credit allocation. This allows us to control the different frictions and their interactions. We assume that banks adopt AI and train their AI simultaneously to find the optimal pricing strategy. This analysis was first proposed by CCDP, which analyse a standard market and do not estimate the underlying equilibrium model. In our study, we use an estimated structural model to quantify the effects of AI adoption since they can have policy relevant implications. Our model is different because the market for credit lines is characterized by asymmetric information that leads to adverse selection and moral hazard.[10] The pricing problem is then substantially different, as a price increase corresponds to a decrease in the quality of the marginal borrower. As the price of the contract increases, only firms with a high propensity to demand will stay in the market, while the others will choose the outside option. However, we have estimated that those firms have, on average, a higher propensity to default. The open question is to understand if AI will ration credit by increasing the prices or if it will decrease the prices and reduce the overall risk associated with the lending activity. To address this question we compare the estimated equilibrium with the equilibrium emerging if the pricing task had been delegated to a pricing algorithm which maximizes profits taking (2.1), (2.2), and (2.3) as given. In the next section, we describe the algorithm of choice, then we present the results.

2.4.1 Qlearning

Following the approach by [9] and [1], we model AI using a reinforcement learning algorithm based on [10], the Q-learning. This algorithm is appealing because it has a clear structure and simple intuition. The purpose of this algorithm is to estimate the Q-matrix, a matrix that maps each possible state of the word and each feasible action into a discounted payoff q:

$$Q(state, action) = \begin{bmatrix} q_{s=1,a=1} & q_{s=1,a=2} & \cdots & q_{s=1,a=A} \\ q_{s=2,a=1} & q_{s=2,a=2} & \cdots & q_{s=2,a=A} \\ \cdots & \cdots & \cdots & \cdots \\ q_{s=S,a=1} & q_{s=S,a=2} & \cdots & q_{s=S,a=A} \end{bmatrix} \tag{2.14}$$

where each entry q represents the discounted expected value of taking action a when the state is s. Estimating this matrix allows the algorithm to find the optimal policy associated with the problem, as the *Argmax* of this matrix is the value function associated with the problem. The estimation is gradual, by trial and error. The algorithm starts playing the game at random, collecting payoffs,

[10]This is a relevant issue for first time borrowers, as the information asymmetries are at the highest level.

and updating the correspondent entry of the Q-matrix. The updated value is a linear combination of the flow payoff (implied by the underlying model) and the old matrix entry.[11] During the learning phase these actions are taken either randomly or following the strategy implied by the current Q-matrix. To ensure that the algorithm updates the payoffs associated with each action in any possible contingency, the algorithms begin by playing random actions and collecting the associated payoff, but, over iterations, the frequency of random actions decreases in favour of the best actions. In particular, whether the algorithm takes a random action (exploration) or what it thinks is the best action (exploitation), depends on a time-dependent probability ε_t:

$$Action = \begin{cases} \text{"}\textbf{\textit{Exploration}}\text{"} \rightarrow random \text{ "}a\text{" } with\ prob. & \varepsilon_t \\ \text{"}\textbf{\textit{Exploitation}}\text{"} \rightarrow best \text{ "}a\text{" } with\ prob. & 1-\varepsilon_t \end{cases} \tag{2.15}$$

where ε_t evolves as $\varepsilon_t = e^{-\beta t}$, given the decay parameter β. Initially, taking random actions means exploring what happens in any contingencies, then, as t increases, the algorithm plays according to the policy it thinks is optimal, but, when it is forced to explore, it enters a different state of the world and updates the payoff accordingly. This deviation can trigger the algorithm to review the current "optimal" policy. Eventually, the algorithm will learn the optimal policy, which is invariant to the exploration. Following the literature, we define convergence as the contingency where the best response in any state of the world does not change for 200.000 iterations.[12]

We cap the training phase up to 60.000.000 iterations of the pricing game. Payoffs are assigned using the estimated model. In particular, at each interaction, we observe the prices $\boldsymbol{p}$ set by the banks in the market and assign the payoff accordingly to the market's equilibrium. In our application, the problem of the firm can be thought of as a repeated game where the firm maximizes its period profits π_t, which depends on the price set by the algorithm and by other agents. Calling $Q(s,p)$ the entry of the Q-matrix corresponding in charging price p when the state is s, the algorithm's learning equation is:

$$q_{t+1}(s,p) = \underbrace{(1-\alpha)q_t(s,p)}_{Past\ knowledge} + \underbrace{\alpha\left[\pi_t + \delta \max_p q_t(s',p)\right]}_{learning\ update} \tag{2.16}$$

where $\alpha \in [0,1]$ is the learning rate.

The theoretical counterpart to $q(s,p)$ is:

$$Q(s,p) = \mathrm{E}\left[\pi|s,p\right] + \delta \mathrm{E}\left[\max_p Q(s',p')|s,p\right] \tag{2.17}$$

Notice that the $argmax_p\{Q(s,p)\}$ is the value function associated with the repeated profit maximization problem. From equation (2.17), it is clear that Q-learning is a way to estimate $Q(s,p)$ and, consequently, the optimal policy.

2.4.2 AI adoption and increase in adverse selection

We study the interaction between AI and adverse selection by analyzing two alternative scenarios. In the first scenario, we let the AI play the game according to the estimated equilibrium, in an alternative scenario we let the AI play the same game but with higher adverse selection (ρ_{BD}). We then compare the results between these two cases and see what is the AI reaction to an increase in

[11] The matrix can be randomly initialized.

[12] Theoretical results on convergence can be found in [11]. In this application, like in [1], there is no theoretical guarantee of convergence, therefore we rely on this definition of convergence.

adverse selection, all else equals. Since we are also interested in understanding the impact of AI in selection markets per sé, we begin our analysis by focusing on duopolies and we assume that both banks simultaneously adopt and train their AI, as in [1]. In our data, a duopoly is a province-year market where only two banks are active. For each borrower in the market and each market, we train two Q-learning algorithms to play the game conditional on the set of characteristics of the borrower. Thus, the state of the game is only the current prices charged by the banks. At each iteration, the algorithms choose the price to charge to the borrower. We assign payoffs following the endogenous expected demand, expected usage and expected default probabilities.[13] We have 32 markets that are duopolies. The number of firms in these markets is 239, and 100 are borrowers. In the baseline scenario, we replicate the results in the literature, but the algorithm does not always converge to a stable best response. In particular, among 100 borrowers, the algorithm exited the learning phase in 57 cases. Among these 57 cases of convergence, only 27 are to a fixed best response, meaning that the algorithms play always the same prices, which are each other best responses. [14] The following figure summarizes the results in the baseline scenario:

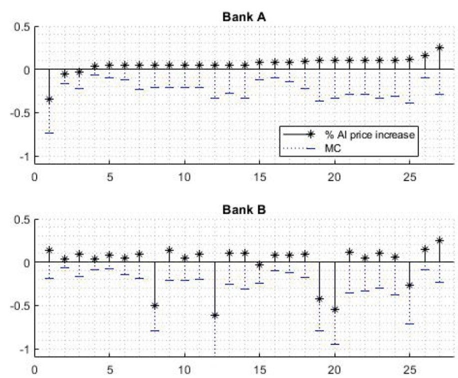

Figure 2.1: Average % price increase from the estimated equilibrium price. Each entry on the X-axis represents a borrower, while the Y-axis is the % deviation from the Nash-Bertrand equilibrium. Ordering is Bank A % price deviation. For example, for the first borrower, bank A decreased the price, while bank B increased it.

In most cases, the algorithm converged to an over-competitive equilibrium, but we occasionally observe a decrease in prices. We also look at the cost of credit resulting from the market equilibrium:

[13]We return to this point in the section "Future developments".

[14]We are investigating why the convergence rate is lower than in CCDP despite the learning being more persistent.

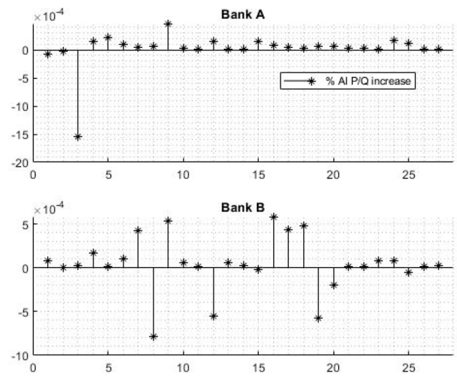

Figure 2.2: Average % cost of credit increase from the estimated equilibrium price. Each entry on the X-axis represents a borrower, while the Y-axis is the % deviation from the Nash-Bertrand equilibrium. Ordering is Bank A % price deviation as in Figure 2.1.

We repeat the same analysis in an alternative scenario where $\rho \uparrow$. The results are summarized in the following figures:

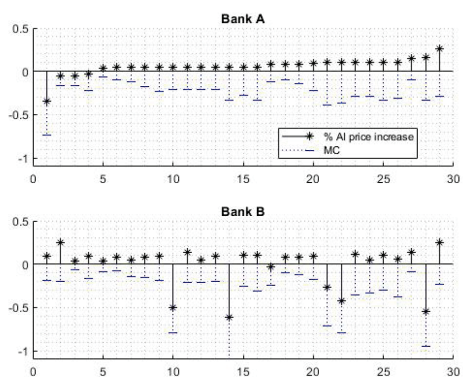

Figure 2.3: Average % price increase from the estimated equilibrium price. Each entry on the X-axis represents a borrower, while the Y-axis is the % deviation from the Nash-Bertrand equilibrium. Ordering is Bank A % price deviation. For example, for the first borrower, bank A decreased the price, while bank B increased it.

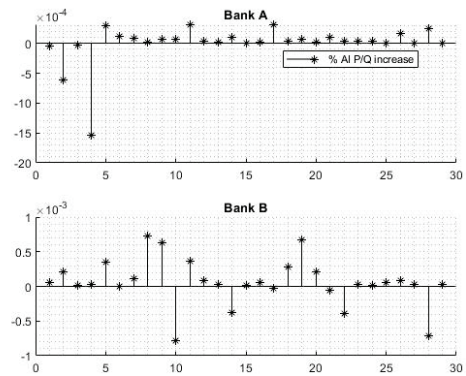

Figure 2.4: Average % cost of credit increase from the estimated equilibrium price. Each entry on the X-axis represents a borrower, while the Y-axis is the % deviation from the Nash-Bertrand equilibrium. Ordering is Bank A % price deviation as in Figure 2.3.

A preliminary analysis suggests that AI tend to decrease the cost of credit when adverse selection increases, however, we see heterogeneous behaviour because the markets are heterogeneous.

2.4.3 Future developments

We are working on extending the analysis in two directions. First, we are currently working on the payoff update procedure. In the baseline case, we assign payoffs using expectations, but, given the nature of the game, another reasonable procedure would be to use lotteries. Since we condition on borrower's characteristics, the idea is to draft two state-dependent lotteries, one that assigns demand and one that determines if the firm defaults. We perform this analysis because we have reasons to believe that the algorithm does not understand the nature of the game, i.e. the selection problem. In particular, we suspect that the AI sees the problem only as a "steeper" demand curve, where an increase in price leads to a relatively higher decrease in expected profits. This is the case because, contrary to standard markets, as $P \uparrow \rightarrow q \downarrow, def \uparrow \rightarrow \mathrm{E}[\pi] \downarrow\downarrow$. We implemented lotteries to investigate if this conjecture is correct. With lotteries, the update of the payoff for the bank winning the demand lottery is either the full profit or the default loss, nudging the algorithms to understand that increasing prices means increasing the number of losses. Ideally, if the number of iterations goes to $+\infty$, the two alternative ways of performing the Q-entry updates should be equivalent, however, since the algorithm learns the optimal strategy in finite time, we might observe differences. We have experimented with the lotteries but the learning process becomes harder. We train the algorithm with the same learning parameters used in section 2.4.2, but, with lotteries, it did not converge. The reason is that with default, the algorithm faces a new learning challenge: π_t is much more volatile and can be either 0 or negative, meaning that every time a state is visited the update $\alpha\pi_t(s,p)$ can sensibly change the Q-entry $Q_{t+1}(s,p)$, changing the best response. This instability can harm the full process, making the algorithm unsuitable for this kind of application. This is a relevant issue that we would like to address, in particular, analyzing convergence with lotteries under different sets of meta parameters α, β, and understanding whether the results with lotteries are equivalent to the one in 2.4.2.

The second direction we are working on is in generalizing the AI adoption results. We want to extend the analysis in the case where the state space of the game includes the characteristics of the borrowers. In the structural model proposed by [2], the demand depends on prices as well as firms' characteristics, but we begin our analysis within borrowers (fixing characteristics) because we wanted to benchmark our results with the literature. Now we would like to depart from the baseline case and we are implementing an algorithm that can learn to play the full game, where the borrower changes over iteration t. The main limitation of this exercise is that standard Q-learning cannot handle the dimension of the problem, as the state space will include 10 variables plus the prices.[15] For this reason, we will implement a different AI: deep Q-learning. With this algorithm, we can train a general AI that can play the pricing game independently of the borrowers' characteristics, and have a general view of the impact of AI on the overall market allocation. In the analysis in section 2.4.2 we can compare differences in pricing strategies on a single borrower, but each borrower is priced by a different AI, while, with a generalized version of the AI we can analyse the impact on credit allocation, rationing, and risk-taking, both in the baseline scenario and in the scenario with higher adverse selection.

2.5 Conclusions

We analyse the impact of artificial intelligence in lending markets. The question is relevant because it is not clear how AI interacts with the typical friction present in financial markets. To answer this question we estimate a state-of-the-art model of demand and supply for credit lines with asymmetric information and later used these estimates as primitives in the AI adoption analysis. The estimates are in line with the results in [2], and the Q-learning implementation led to results similar to the findings [1], however with some key differences. First, the algorithm does not always converge, with the convergence rate being at about 60%. Second, the algorithm mostly plays the Nash-Bertrand prices or increases the markup, but, sometimes, the prevailing equilibrium price can be below Nash. In conclusion, we confirm that AI raises the prices, however, the increase is relatively lower because the algorithm anticipates the selection problem. Finally, we are working on the generalization of these results to study the effect of AI adoption on market outcomes.

[15] Standard Q-learning, also known as tabular Q-learning, suffer the curse of dimensionality: as the state space increases the entry of the Q-matrix increases geometrically, making it unfeasible to explore and learn the optimal policy.

References

[1] E. Calvano et al. "Artificial Intelligence, Algorithmic Pricing, and Collusion". In: *American Economic Review* 110.10 (2020), pp. 3267–97.

[2] G. S. Crawford, N. Pavanini, and F. Schivardi. "Asymmetric Information and Imperfect Competition in Lending Markets". In: *American Economic Review* 108.7 (2018), pp. 1659–1701.

[3] G. A. Akerlof. "The Market for "Lemons": Quality Uncertainty and the Market Mechanism". In: *The Quarterly Journal of Economics* 84.3 (1970), pp. 488–500. ISSN: 00335533, 15314650.

[4] J. E. Stiglitz and A. Weiss. "Credit Rationing in Markets with Imperfect Information". In: *The American Economic Review* 71.3 (1981), pp. 393–410. ISSN: 00028282.

[5] M. A. Petersen and R. G. Rajan. "The Benefits of Lending Relationships: Evidence from Small Business Data". In: *The Journal of Finance* 49.1 (1994), pp. 3–37.

[6] S. Berry, J. Levinsohn, and A. Pakes. "Automobile Prices in Market Equilibrium". In: *Econometrica* 63.4 (1995), pp. 841–890. ISSN: 00129682, 14680262. (Visited on 09/10/2024).

[7] L. Einav, M. Jenkins, and J. Levin. "Contract Pricing in Consumer Credit Markets". In: *Econometrica* 80.4 (2012), pp. 1387–1432. DOI: `https://doi.org/10.3982/ECTA7677`.

[8] Y. Rothschild and J. Stiglitz. "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information". In: *The Quarterly Journal of Economics* 90.4 (1976), pp. 629–649.

[9] L. Waltman and Kaymak. U. "Q-learning agents in a Cournot oligopoly model". In: *Journal of Economic Dynamics and Control* 32.10 (2008), pp. 3275–3293. DOI: `https://doi.org/10.1016/j.jedc.2008.01.003`.

[10] C. J. C Watkins. "Learning From Delayed Rewards". In: (Jan. 1989).

[11] C. J. C. H. Watkins and P. Dayan. "Q-learning". In: *Machine learning* 8 (1992), pp. 279–292. DOI: `https://doi.org/10.1007/BF00992698`.

Chapter 3

Inventory models in monetary economics

Carmelo Genovese[1]

Abstract

I review the use of inventory models in economic theory. These models were introduced in economics during the fifties, and through the years they have been adopted to model the economic behaviour of different kinds of agents. They have risen in popularity in recent years because the availability of detailed microdata allowed researchers to estimate the underlying parameters and understand the mechanisms that drive economic decisions. I review the theoretical contributions, focusing on the main contributions that generalized the early version of the models. For the early works, I focus on the works that contributed to generalising the framework, while, I focus on their recent success for empirical applications.

[1] Luiss Guido Carli.

3.1 Introduction

Inventory models, also known as fixed cost models, are theoretical models developed in operation research during the first half of the twentieth century. These models were created to analyse the problem of firms' inventory keeping, a fundamental source of costs in the fast growing US economy. Examples of early works are Davis (1925, [1]) and Meller (1925, [2]). Inventory modelling was a hot topic during these years, with scientific conferences on logistics and inventory keeping held annually all over the US.[1]

In its simplest formulation, an inventory problem is a problem where an agent needs to manage a stock, x, which is evolving over time. Managing this stock is costly: on one side, keeping the stock costs $A * x$ per period, on the other side, getting (or getting rid) of this stock costs B. For some reason, which will become clear within applications, the agent needs to keep the stock but would like to minimize the flow cost $A * x$, keeping in mind that she will need to pay B to adjust the stock. Hence, the problem of the agent, given the process driving the changes in x, is to balance the amount of stock, decide when to pay the fixed cost to adjust the stock, and, upon adjustments, decide the optimal level of stock x^*. In the case of a firm, x is the amount of goods in the warehouse, A is the cost of keeping the goods in the warehouse, and B is the cost of restocking the stock. The stock of goods decreases when sales are made. When deciding how many units of the good to keep, the firm needs to take into account that increasing the stock will increase the cost of managing the goods, while keeping low inventories implies that the fixed cost will be paid often to procure new goods.

Soon, scholars in economics started to read this literature and understood that a key aspect of inventory management was overlooked: the demand for goods. The models proposed by Arrow et al. (1951, [3]) and by Within (1952, [4]) are the first that jointly analyse the problem of inventory management with the demand for goods driving inventory's depletion. In particular, the contribution of Arrow et al. is to derive the optimal inventory policy in a setting with stochastic demand and both fixed and flow costs associated with inventory management. At the time, the most well-known result was to keep a constant sale/inventory ratio. Their contribution was to show that this policy might not be optimal and that fluctuations should be taken into account when making inventory decisions. Moreover, Within proposed a model where the demand for goods was stochastic, showing that demand uncertainty can lead to periods of depletion, and argued that the optimal management policy for inventory keeping should be based on expected demand. Within also wrote a survey on the evolution of inventory theories up to the fifties (1954 [5]).

These models are important contributions to the general theory of inventory keeping and changed the way researchers without an economic background conjectured the problem. In the following years, researchers have been generalizing the results adding structure to the baseline version of the problem. In particular, engineers and mathematicians have been developing generalized stochastic models for optimal inventory, logistics and transports, which are still relevant topics in operation research.[2]

The work by Within and Arrow et al. gave a second, and perhaps more important "contribution" to the economic literature: the introduction of inventory models in economic research. In the fifties, researchers realized that inventory theory could be extended to non-business applications: every stock/item that is managed as an inventory can be potentially analysed using this theory. The most famous application is for household cash management, which was first conjectured by Allais [7], then formalized independently by Baumol [8] and by Tobin[9]. Baumol himself stated in [8]:

"*T. M. Whitin informs me that the result in question goes back to the middle of the 1920s when it seems to have been arrived at independently by some half dozen writers. [...] Whitin analyzed them*

[1] For example, the RAND corporation used to organize a conference on logistics.

[2] It is possible to look at Perera and Sethi (2023 [6]) for a recent survey.

in his forthcoming [4] which, incidentally, first suggested the subject of this note to me."
In this paper I analyze the evolution of inventory models in economic research, highlighting the main contributions. The survey is organized as follows: in section 3.2 I describe the early application in monetary theory, including the criticisms made to this approach, then, in section 3.3 I summarize the contributions that generalized the baseline framework. Finally, I describe recent empirical applications in section 3.4.

3.2 Inventory models for money demand

The most well-known early application of inventory theories in economics is the famous Baumol-Tobin model for money demand. Both authors realized that the household's cash management problem can be thought of as an inventory problem: keeping the right amount of cash to pay for consumption given that holding and withdrawing cash is costly. The underlying assumption in both models is that households have a cash-on-hand constraint.
When deciding how much cash to hold, households need to balance the fixed cost associated with withdrawing cash[3], b, with the opportunity cost of holding cash, r. This opportunity cost can be thought of as the difference between the return from deposit/invest cash and the (possible) return on cash kept on hand. Notice how this problem is similar to an inventory keeping problem: the cost of raising cash is equivalent to the cost of restocking the inventory of goods, while the opportunity cost is equivalent to the cost of keeping the inventory. The inventory of money decreases as goods are purchased.
In [8] and [9], the household consumes T worth of goods per period, that are paid with cash on hand. This stock of cash decreases at every transaction, eventually running out. At that moment, the household must pay b and make a withdrawal. The problem of the household is then to choose the amount of cash C to withdraw and the frequency of withdrawals, leading to n withdrawals per period.[4] The optimal solution of the model is that when withdrawing, the household chooses:

$$C^* = \sqrt{\frac{2bT}{r}} \tag{3.1}$$

which is known as the "*square root rule*". To highlight the contribution of each paper in this survey, I will show how this optimal rule evolves. Given C^*, the optimal frequency of withdraws f per period is:

$$f^* = \frac{1}{n^*} = \sqrt{\frac{2b}{Tr}} \tag{3.2}$$

While the first to publish these results was Baumol, Tobin realized that this model is a micro foundation of households' money demand.[5] In particular, he realized that the inventory framework could be used to give a micro foundation of the money demand curve as a function not only of income[6] but also of the interest rate and transaction costs. Soon, in macroeconomic models, especially in Keynesian models, the use of interest-sensitive money demand schedules became the standard. For the rest of the paper, I will refer to their model as BT. [7]

[3] In the early works this cost is thought to be the cost of borrowing this cash or to disinvent part of the asset holdings, in modern days this can be thought of as the cost of going to the ATM to withdraw.
[4] For simplicity, it is assumed that T is uniformly consumed over the period.
[5] He stressed this point in the introduction of the paper [8]
[6] Tobin idea is that T, the need for transactions, reflect the level of income.
[7] The success of this theory is due to the articles by Baumol and Tobin, hence the name, but both authors agreed that the real father of this theory is Allais [10].

Following these works, Miller and Orr (1966, [11]), MO henceforth, had the idea to apply this framework to model cash holdings management at the firm level. While this exercise might seem straightforward, the essence of the problem is different. Firms' cash holdings evolve stochastically. They have cash expenditures to pay for costs and inputs and receive cash inflows from sales. Given this difference, they modelled cash holdings as a random Bernoulli variable. This is the opposite approach to BT, where cash was constantly decreasing in a deterministic way. Miller and Orr were well aware that both the assumptions were too extreme and tried to have a version of the model with both random and deterministic components, however, due to mathematical limitations, they focus on the case where the cash flow is a martingale.[8] The second contribution of their paper, which is a direct consequence of the martingale assumption, is that they define the optimal policy in terms of an inaction region. In contrast to BT, in the model of MO, firms' cash holdings can increase, and, consequently, the inventory of cash may become so costly to manage that it is more advantageous to pay the fixed cost b and deposit/invest the extra cash. An approach to solve this problem is to conjecture that the representative firm chooses an inaction region such that, while cash holdings are within this region, it does not pay b. To be specific, it needs to decide a threshold, M_2, such that if the level of cash holdings ends up being above this threshold, then it optimally pays b and invests the cash in excess, enjoying the higher return. On the other side, when cash holdings run out, the firm must withdraw, as in BT. Finally, the last decision of the firm is how much cash to hold when adjusting its cash holdings C^*_{MO}. Notice that uncertainty raises a precautionary motive for the agent: she needs to take into account that with some probability she will be paying b earlier than expected.

Inaction region solutions are also known as steady state solutions because, despite the dynamic nature of the agent's problem, the optimal policy is fixed. [9] In MO the inaction region is $[0, M_2]$, and the optimal amount held upon adjustment is

$$C^*_{MO} = \sqrt[3]{\frac{3b\mu^2 f}{4r}} \tag{3.3}$$

where μ is the size of the period change in cash holdings.[10] The following figure, taken from the paper, gives a graphic intuition behind the model:

[8] Their attempt to have a model with variation around a deterministic trend can be found in the appendix of the paper.

[9] This kind of solution was proposed in mathematics by Karlin [12].

[10] In MO these change is driven by a Bernoulli process that determines whether cash holdings increase or decrease by m.

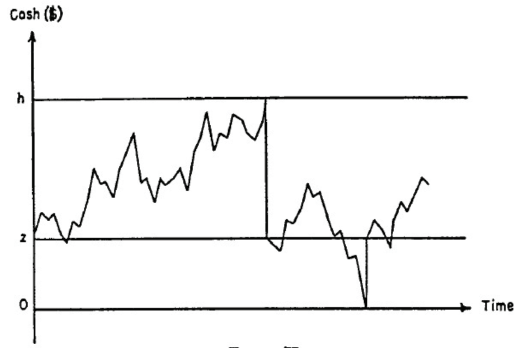

Figure 3.1: Optimal cash holdings policy from the paper by Miller and Orr $\{M_2 = h,\ C^*_{MO} = z\}$. Notice how the cash holding level evolves stochastically but is immediately reset at z when it hits one of the two boundaries $\{0,h\}$.

While the models discussed so far have had a lot of success, several authors distrusted them, highlighting that they were very stylized and that the assumptions behind the model were inconsistent with reality. For example, Sprenkle published an article criticizing this approach to modelling economic behaviours (1969 [13]) (1977 [14]). In this article, he argues against the credibility of modelling cash holdings decisions as inventory decisions, pointing out that several assumptions are too restrictive to have credible results. In particular, supported by the results in some papers trying to test the BT model, like Meltzer (1963 [15]), Sprenkle criticized both the assumption that the transaction flow is deterministic and the cost structure of the model. These critiques are based on two facts. First, these models need detailed and high frequency data to be estimated correctly. Second, even if data were available, the process underlying inventory depletion is too simplistic to capture these dynamics. For example, the elasticity of C^* to r is:

$$\eta_{C^*,r} = -\frac{1}{2} \tag{3.4}$$

Which is independent of other economic quantities. Of course, this result does not hold in reality. While these were valid points, advancements in mathematics as well as the availability of powerful storage and computation machines made it possible to generalize and estimate models much more sophisticated. Nowadays, it is possible to replicate fluctuations and have rich dynamics within these models, but whether a steady state policy based on an inventory model can realistically describe economic behaviours is still an open debate. In the next section, I will discuss the main contributions that addressed these critiques. The focus of these contributions will remain on cash holdings management and the implied money demand. Only later these models will be used for other economic studies.

3.3 From Baumol Tobin to the general stochastic model

While it was clear that the BT model was too stylized to produce results consistent with other economic theories and empirical evidence, its popularity did not fall. Soon, the basic model was enriched, partly addressing the common critiques. These contributions can be divided into two categories. The first one includes contributions aimed at addressing the elasticity problem, while the other includes contributions aimed at generalising the dynamics of the problem. To this extent, the paper by Miller and Orr is an early attempt to generalize the framework by including stochastic stock variation. Despite their similarities, for several years, BT and MO were seen as complements: BT was the model for households' money demand, MO for firms' cash management.

3.3.1 Early contributions: focus on elasticities

One of the first authors that attempted to enrich the baseline versions of the BT model is Johnson (1969 [16]), who proposed a modified version that looked at the effect of interest earnings and fixed costs disbursement on the static budget constraint. In this way, managing costs are reflected in current consumption. The optimal cash holding in his model is:

$$C^*_{Johnson} = \frac{\sqrt{(2+r)bt}}{2r\left(1-\frac{k}{2}\right)} \tag{3.5}$$

where $\frac{k}{2}$ is a proportional cost of withdrawal. The elasticity of $C^*_{Johnson}$ with respect to r, denoted by $\eta_{\{C^*_{Johnson},r\}}$ is not constant as in BT:

$$\eta_{\{C^*_{Johnson},r\}} = -\frac{1}{2+r} \tag{3.6}$$

However, the author himself pointed out that $\eta_{\{C^*,r\}} \approx \frac{1}{2}$ for reasonable levels of r.
A similar motivation is behind the Karni (1973, [17]) version of the BT model. Karni generalized the model by rethinking the transaction level T as dependent on income Y. In particular, inspired by the new value of time economic theories that were taking over in microeconomics, he wanted to integrate an effort dimension in the fixed cost b. His idea was based on the fact that time is valuable, hence the fixed cost to withdraw should include this opportunity cost component:

$$b = \bar{b} + t_e W \tag{3.7}$$

where $\bar{b}$ is the fixed component of the adjustment cost, W is the hourly wage and t_e is the time lost to make a withdraw. Using the same logic, he decomposes the value of transaction T as fraction α of income

$$T = \alpha(\bar{Y} + t_w W) \tag{3.8}$$

where $\bar{Y}$ is fixed income, W is the hourly wage, t_w are hours worked. This attempt to generalize the model does not change the fact that both the demand and the effort enter the cost structure in a deterministic way, therefore, the implied optimal level of cash upon withdrawal still follows the square root role:

$$C^*_{Karni} = \sqrt{\frac{\alpha\left(\bar{Y} + t_w W\right)\left(\bar{b} + t_e W\right)}{2r}} \tag{3.9}$$

However, the elasticity now depends on the relative share of fixed and total income:

$$\eta_{\{C^*_{Karni},r\}} = -\frac{1}{2}(1 - \frac{\bar{Y}}{Y}) \in (-\frac{1}{2}, 0) \tag{3.10}$$

From equation (3.9) it is possible to compute other relevant sensitivities, for example, the elasticity to hours worked t_w or fixed income $\bar{Y}$.
This kind of generalization allowed to extend the money demand theory, and to give microfoundation

to the relationships between cash holdings and other relevant economic quantities, like the mentioned income and wages. This added depth to the framework and its implications for money demand theories.[11]
In this period, the inventory concept was finally applied to a different economic problem. Sheshinski and Weiss (1977, [18]) adopted this framework to model the problem of the firm setting prices.[12] The problem was formulated with the following logic: A firm needs to decide the price P_t of its product. Once the price is set, changing it is costly b, therefore, the firm needs to consider that the price level in the economy will increase[13] and that $P_t < P^*_{t+1}$, where P^*_{t+1} is the price the firm would choose if adjusting P_t was cost-free. This difference generates a profit lost. Hence, the firms need to trade-off the lost in profits with the cost of adjusting the price, which is the typical inventory model trade-off. This is the first version of the famous *"menu cost"* cost model, a macroeconomic model were firms need to optimally choose P given that changing the prices in the menu is costly.[14] Notice that this early version of the menu cost model is in a deterministic setting, like in the BT.

3.3.2 Stochastich and general equilibrium versions of the model

In the eighties, the focus shifted from structure to generalization. So far, one of the crucial limitations of using the BT/MO framework was the inability to have an underlying process which included both a trend, or a deterministic need for transactions (BT), and uncertainty in the sizes of the transactions (MO). Without a general process, the dynamics of the model were far from reality. For example, any household expects to have some expenses over time, but the exact amount fluctuates. The same is true for firms, which have uncertain cash flows. The solution to this problem arrived from progress in mathematics.[15]
The first economists who had the intuition to use stochastic calculus to solve a general inventory model were Frenkel and Jovanovic (1980 [21]).[16] Their model is the core of today's applications. They formulated the inventory problem in continuous time and assumed that the underlying process evolves as a Brownian motion with drift. To solve the model they followed the strategy suggested by MO, assuming that the household follows a steady state policy. Moreover, they assumed that households have a continuous outflow of cash dm, which is stochastic:

$$dm_t = -\mu dt + \sigma dW_t \tag{3.11}$$

where μ is the drift[17] and dW_t is a Brownian motion. This means that the household never knows how much cash will be needed at $t_n > t$, but he knows the distribution:

$$m_{t_n} \sim \mathrm{N}\left(\mu(t_n - t), \sigma^2(t_n - t)\right) \tag{3.12}$$

The essence of the problem is still the same as BT: decide how much cash to hold given the cash in advance constraint. The household pays an opportunity cost r per unit of cash in hand m and each withdrawal costs her b. If cash on hand runs out, she is forced to pay b and withdraw. Of course, upon withdrawing, the household chooses to hold m^*, the optimal level of cash on hand.

[11] This is the freshwater vs. saltwater economics period, where the most important schools in the US were arguing about fundamentals in macroeconomic theories.
[12] They took inspiration from Barro (1972 [19]), which, inspired by inventory models, included a fixed cost to adjust prices in the pricing problem.
[13] They assume a constant and known in advance inflation rate.
[14] Think about a restaurant that needs to print a new menu each time the price of one dish should be changed.
[15] For example, Cox and Miller published in 1966 a book on stochastic calculus [20].
[16] Stochastic calculus was already in use in economic papers, for example in Fisher (1975 [22]).
[17] Usually the drift it is assumed to be negative, so it represents the expected expenses per period. $c = \mu$ can be interpreted as the value of consumption per period.

The optimal policy is $\{m^*, m_2\}$, where m_2 is the level of cash on hand such that the households prefer to pay b and deposit excess liquidity $m_2 - m^*$, like in MO. Several aspects of this contribution deserve a discussion. First, the optimal level of cash holdings m^* is:

$$m^* \approx \sqrt{\frac{2b\sigma^2}{\sqrt{\mu - 2r\sigma^2} - \mu}} \tag{3.13}$$

Which is sensitive to both the cost structure and the parameters of the underlying diffusion process. Second, the model can nest both the BT and the MO models. The model reduces to BT when $\sigma = 0$, while reduces to MO when $\mu = 0$. The implied elasticity of cash holdings to the opportunity cost is:

$$\eta_{m^*,r} = -\sqrt{\frac{r\sigma^2}{4\left(\mu^2 + 2r\sigma^2 - \mu\sqrt{(\mu^2 + 2r\sigma^2)}\right)}} \tag{3.14}$$

which is more convoluted but can still nest the $-1/2$ value by Baumol. From equation (3.13) is it possible to compute several other relevant elasticities. Finally, another benefit of using continuous time modelling is the possibility of computing the invariant distribution of the state variable (cash holdings) associated with the solution. The invariant distribution is the distribution of cash holdings assuming that emerges if the household sticks to the $\{0, m^*, m_2\}$ rule:

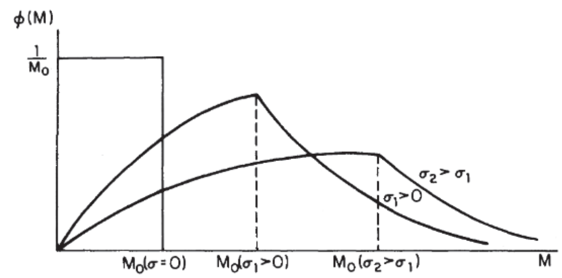

FIGURE II
The Steady-State Probability Density of Money Holdings for Alternative Values of σ

Figure 3.2: Invariant distribution for different values of the parameter σ. Image taken from [21].

This distribution, in theory, can be used to assess whether the implication of the theoretical model is consistent with the data. However, it depends on the Brownian motion parameters $\{\mu, \sigma\}$, which can be estimated only using transaction data at the household level, which were not available yet.[18] The second kind of generalization made in the eighties was to include the inventory style cash management within macroeconomic models. First, Grossman and Weiss (1983 [24]) proposed a dynamic model where utility-maximizing households have a cash-on-hand constraint and can withdraw only at an exogenously given frequency. For simplicity, they do not include the fixed cost component of inventory theory, but they force households to wait between withdrawals, as standard inventory theory implies. This simplified their problem as the households' problem does not include the optimal choice of n. In particular, In their model, there are two types of agents that can withdraw cash from

[18] A good introductory reference to deepen the knowledge in stochastic calculus and its application in economics is Dixit [23].

an interest-earning deposit any other period. These agents have non-financial wealth and can invest their money in a risk-free bond that yields the same return as the deposit. The main result of their analysis is that changes in money supply can lead to real effects because some households cannot adjust their cash holdings right away. Despite this result coming from the exogenous assumption of synchronized adjustments, it was clear that monetary non-neutrality, typical in Keynesian models with sticky prices, could be obtained using fixed cost models. Rotemberg (1983 [25]) is another example. This logic that a fixed cost causes inaction and, consequently, real effects of nominal changes, like monetary policy changes, is behind the "menu cost" literature as well as other literatures.[19]
In the same years, Jovanovic published another important paper. In (1982 [28]), he proposed a model with an endogenous opportunity cost: a general equilibrium version of the inventory model. In particular, in this paper, the opportunity cost r, which is a key component in inventory theory, is determined in equilibrium. Households need cash to pay for consumption, and this cash can be obtained by selling the asset. Of course, while converting assets into cash, the households incur a fixed cost, b. The average cash holdings $\bar{m}$ still follows the square root rule:

$$\bar{m} \approx \sqrt{\frac{bT}{2(\rho + \pi)}} \tag{3.15}$$

where ρ is discount rate and π is inflation.
This general equilibrium version of the inventory model was reformulated by Romer (1986 [29]). He built an overlapping generation (OLG) model where households can convert the risk-free asset into cash by paying a fixed cost. Thus, in planning their consumption stream and their asset holdings, they need to consider the inventory trade-off. A novelty in this model is that the fixed cost is now expressed in utility loss rather than in real terms. The consequence of this assumption is that the frequency of adjustment is independent of wealth. The level of money holdings in this model is:

$$m = \left[\frac{e^{-\pi\tau} + \pi\tau - 1}{(\pi\tau)^2}\right] \left[\frac{(\pi\tau)E}{(1 - e^{-\pi\tau})T}\right] \tau \tag{3.16}$$

Where π is the inflation rate, τ is the time between withdrawals, E is the initial endowment of each agent and T is the lifespan of each individual.[20] Equation (3.16) is the analogue of the square root role, but in a general equilibrium OLG setting. The optimal level of cash holdings is given by the product of three quantities, the ratio of average cash holdings to initial withdraw amount, the average amount used per unit of time, and the interval between adjustments τ. In particular, their counterpart in the BT square root rule are:

$$\left[\frac{e^{-\pi\tau} + \pi\tau - 1}{(\pi\tau)^2}\right] \to \frac{1}{2} \qquad \left[\frac{(\pi\tau)E}{(1 - e^{-\pi\tau})T}\right] \to Y \qquad \tau \to \sqrt{\frac{2b}{rY}}$$

This general equilibrium formulation of C^* implies that the elasticity is now more convoluted, and depending on parameters and assumptions on the functional forms, the elasticity of cash holdings to the nominal interest rate can even be increasing.[21]
Other relevant papers written in this period are Chant (1976 [30]), which highlights the chaotic behaviour of early model away from steady state, Milbourne et al. (1983, [31]), that generalizes and attempts to empirically validate the MO model, and Smith (1986, [32]), then shows how to construct aggregate money demand from individuals inventory problem á la [31]/[21]. Akerlof and

[19] For example see Cooper et al. (1993 [26]) or Bloom (2009 [27]).
[20] To derive analytical results Romer assumed logarithmic utility and no discount of the future.
[21] See section "*Money demand*" of the paper by Romer, pages 667-679, for a detailed discussion.

Milbourne (1980 [33]) analyse money demand when the inventory is subject to random jumps. Romer used his general equilibrium inventory model to study monetary policy transmission, generalizing the results in [24] and [25]. Blanchard returned to the origin of inventory theory and proposed a model to rationalize the empirical evidence coming from the automobile industry (1983, [34]). Finally, Caballero and Engel (1991 [35]) proposed a theory to study aggregate dynamics with agents following steady state policies.

3.4 Contemporaneous applications

In this section of the paper, I will discuss some recent contributions to the BT/MO models and some modern applications of inventory theory to model different economic problems. Before that, I would like to point out that the "*menu cost*" literature, started by [18], can be seen as an application of inventory theories to the price-setting problem. However, this literature has been evolving separately since the eighties, therefore I will not cover it in the remaining of this survey.[22]

The idea that fixed costs lead to inaction and to the violation of the nominal/real dichotomy was discussed in the previous section 3.3. Inspired by this intuition, Alvarez et al. (2009, [37]) proposed a general equilibrium inventory model to explain lumpiness in the reaction of prices and inflation to monetary shocks. This paper marks another important contribution in the general equilibrium inventory model literature because it relaxed the OLG assumption of [29] and extends the findings in [38].

Another relevant contribution is Alvarez and Lippi (2009 [39]), who proposed a version of BT with the possibility of free adjustments. These free adjustments are the source of variation in the model and have a non trivial role in the analysis: they allow the authors to derive equations for elasticities and other relevant statistics, like average cash balances at withdrawals. They model these free adjustment opportunities as the possibility to jump to the optimal inventory level m^* independently of the current level of cash m and the fixed cost b.[23] One immediate consequence of the possibility of free adjustments is that the precautionary motive to hold cash decreases because with some positive probability, cash will be replenished. This also means that the inaction region increases and that the optimal level of cash on hand when withdrawing decreases. In this model, aggregate money demand is:

$$M = \frac{m^*}{1 - e^{\frac{-\theta}{b} m^*}} - \frac{\mu}{\theta} \tag{3.17}$$

Another important point of this paper is that the authors had detailed micro data on households' cash holdings and cash management decisions, like cash holdings at the moment of withdrawing or the size of the withdraws. This was the first empirical application where the critique by Sprenkle was not relevant anymore. Alvarez, Lippi and coauthors are still working on extensions of this framework: [40], [41], and [42].

Other authors were inspired by these contributions. In particular, Bailey and Blanco (2021, [43]) realized that this framework was suited to model firms' investment decisions. They found empirical evidence that firms' investments were lump andy, therefore coherent with a fixed cost model. Their contribution was to formalize the estimation strategy, mapping model implied moments and

[22] One of the workhorses of "*menu cost*" models is Golosov and Lucas (2007 [36]). The reader will find a brief history of that literature and, in note 2, page 174, the main articles that contributed to their formulation of the price-setting problem.

[23] Practically when solving his problem, the households know that with an arrival rate θ he can jump from the current "value" of the problem $V(m_t)$ to the optimal one, $V^(m^*)$, where the "value" of the problem is the value function.

statistics to the empirical counterparts. This allowed them to do a strong impulse response analysis, quantifying the effects of an aggregate productivity shock on the capital to productivity ratio:

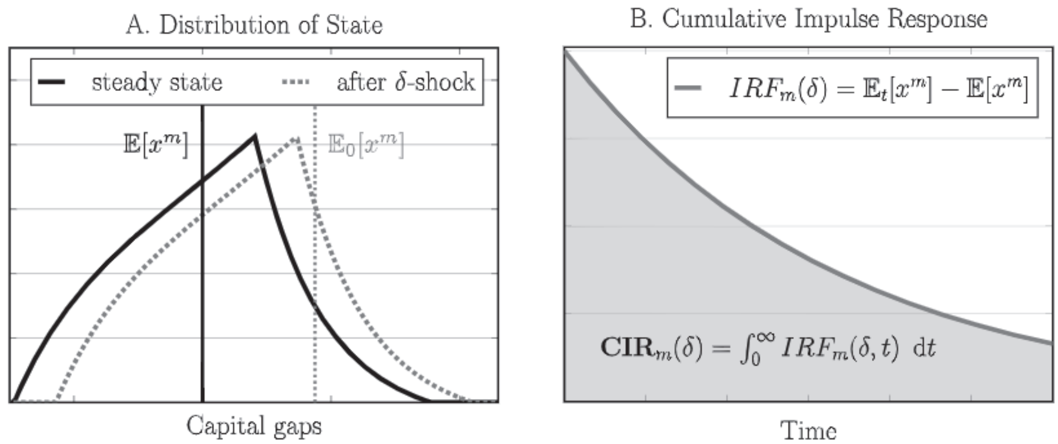

Figure 3.3: This figure, taken from their paper, shows the impact of an aggregate shock δ on the distribution of the state variable x(left panel) and the implied impulse response function of a given moment E$[x^m]$. See the paper for the empirical application on Chilean industrial plants.

The last application of the inventory model I would like to discuss is in corporate finance, in the paper by Bolton et al. (2011, [44]). This is a well established paper in the financial literature that deals with the problem of a firm allocating its assets and deciding the optimal internal/external financing ratio. In particular, the firm needs to decide when to use external financing. The authors include a fixed financing cost that is proportional to capital[24] and assume that liquidity is managed as an inventory, and that can be used instead of external financing. Hence, liquidity has an intrinsic opportunity value. In terms of liquidity management, the implied invariant distribution of liquidity when the firm decides to use external financing is:

[24]Fixed in the sense that does not depend on the financing amount.

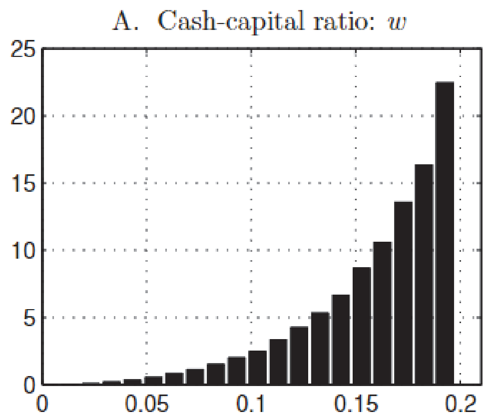

Figure 3.4: Figure taken from the paper. Invariant distribution of cash holdings when the firm is externally financing.

The paper by Bolton et al. touches on different interesting points and has other relevant results which are not shown because they do not belong to the inventory framework.

3.5 Concluding remarks

This is a short and non-exhaustive survey of the use of inventory models in economics. After their introduction by [4] and [3], they become popular tools that to the work by [10] and [9]. that used the framework to give microfoundation to money demand. The basic model was extended and generalized, and nowadays is used for different economic problems. The underlying concept is easy to understand but can generate complex dynamics and interactions in sufficiently complicated models. If an economic problem can be reasonably modelled using inventory theories, the only drawback remaining is the availability of detailed micro data to be able to robust estimations. In their stochastic versions, these models can produce interesting results that can be tested, as it has been shown that it is possible to have a clear mapping between models' and empirical statistics/moments.

References

[1] R. C. Davis. "Methods of Finding Minimum-Cost Quantity in Manufacturing". In: *Manufacturing Industries* (Apr. 1925).

[2] G. F. Mellen. "Practical Lot Quantity Formula". In: *Management and Administration* (Sept. 1925).

[3] K. J. Arrow, T. Harris, and J. Marschak. "Optimal Inventory Policy". In: *Econometrica* 19.3 (1951), pp. 250–272. (Visited on 09/25/2024).

[4] T. M. Whitin. "Inventory Control in Theory and Practice". In: *The Quarterly Journal of Economics* 66.4 (1952), pp. 502–521. (Visited on 09/25/2024).

[5] T. M. Whitin. "Inventory Control Research: A Survey". In: *Management Science* 1 (Oct. 1954), pp. 32–40.

[6] S. C. Perera and S. P. Sethi. "A survey of stochastic inventory models with fixed costs: Optimality of (s, S) and (s, S)-type policies—Discrete-time case". In: *Production and Operations Management* 32.1 (2023), pp. 131–153.

[7] M. Allais. *Economie et interet*. Paris: Imprimerie Nationale, 1947.

[8] W. J. Baumol. "The Transactions Demand for Cash: An Inventory Theoretic Approach". In: *The Quarterly Journal of Economics* 66.4 (1952), pp. 545–556.

[9] J. Tobin. "The Interest-Elasticity of Transactions Demand For Cash". In: *The Review of Economics and Statistics* 38.3 (1956), pp. 241–247.

[10] N. Amberg et al. "Why Trade Credit?" Working paper.

[11] M. H. Miller and D. Orr. "A Model of the Demand for Money by Firms*". In: *The Quarterly Journal of Economics* 80.3 (Aug. 1966), pp. 413–435.

[12] S. Karlin. *Steady State Solutions*. Oxford university press, 1958. Chap. XIV in Arrow, Karlin, and Scarf, Studies in the Mathematical Theory of Inventory and Production.

[13] C. M. Sprenkle. "The Uselessness of Transactions Demand Models". In: *The Journal of Finance* 24.5 (1969), pp. 835–847.

[14] C. M. Sprenkle. "The Uselessness of Transaction Demand Models: Comment". In: *The Journal of Finance* 32.1 (1977), pp. 227–230. ISSN: 00221082, 15406261.

[15] A. H. Meltzer. "The Demand for Money: A Cross-Section Study of Business Firms". In: *The Quarterly Journal of Economics* 77.3 (Aug. 1963), pp. 405–422.

[16] H. G. Johnson. "A Note on the Theory of Transactions Demand for Cash". In: *Journal of Money, Credit and Banking* 2.3 (1970), pp. 383–384. ISSN: 00222879, 15384616.

[17] E. Karni. "The Transactions Demand for Cash: Incorporation of the Value of Time into the Inventory Approach". In: *Journal of Political Economy* 81.5 (1973), pp. 1216–1225.

[18] E. Sheshinski and Y. Weiss. "Inflation and Costs of Price Adjustment". In: *The Review of Economic Studies* 44.2 (June 1977), pp. 287–303. ISSN: 0034-6527.

[19] R. J. Barro. "A Theory of Monopolistic Price Adjustment". In: *The Review of Economic Studies* 39.1 (1972), pp. 17–26. ISSN: 00346527, 1467937X.

[20] D. R. Cox and H.D. Miller. *Theory of Stochastic Processes*. London: Methuen Ltd., 1965.

[21] J A. Frenkel and B. Jovanovic. "On Transactions and Precautionary Demand for Money". In: *The Quarterly Journal of Economics* 95.1 (1980), pp. 25–43. ISSN: 00335533, 15314650.

[22] Fischer S. "The Demand for Index Bonds". In: *Journal of Political Economy* 83.3 (1975), pp. 509–534.

[23] A. Dixit. *The Art of Smooth Pasting*. Harwood Academic Publishers GmbH, 1993.

[24] S. Grossman and L. Weiss. "A Transactions-Based Model of the Monetary Transmission Mechanism". In: *The American Economic Review* 73.5 (1983), pp. 871–880. ISSN: 00028282.

[25] J. J. Rotemberg. "A Monetary Equilibrium Model with Transactions Costs". In: *Journal of Political Economy* 92.1 (1984), pp. 40–58. ISSN: 00223808, 1537534X.

[26] R. Cooper and J. Haltiwanger. "The Aggregate Implications of Machine Replacement: Theory and Evidence". In: *The American Economic Review* 83.3 (1993), pp. 360–382. ISSN: 00028282.

[27] N. Bloom. "The Impact of Uncertainty Shocks". In: *Econometrica* 77.3 (2009), pp. 623–685.

[28] B. Jovanovic. "Inflation and Welfare in the Steady State". In: *Journal of Political Economy* 90.3 (1982), pp. 561–577. ISSN: 00223808, 1537534X.

[29] D. Romer. "A Simple General Equilibrium Version of the Baumol-Tobin Model". In: *The Quarterly Journal of Economics* 101.4 (1986), pp. 663–685. ISSN: 00335533, 15314650.

[30] J. F. Chant. "Dynamic adjustments in simple models of the transactions demand for money". In: *Journal of Monetary Economics* 2.3 (1976), pp. 351–366. ISSN: 0304-3932.

[31] R. D. Milbourne, P. Buckholtz, and M. T. Wasan. "A Theoretical Derivation of the Functional Form of Short Run Money Holdings". In: *The Review of Economic Studies* 50.3 (1983), pp. 531–541. ISSN: 00346527, 1467937X.

[32] G. W. Smith. "A Dynamic Baumol-Tobin Model of Money Demand". In: *The Review of Economic Studies* 53.3 (1986), pp. 465–469. ISSN: 00346527, 1467937X.

[33] G. A. Akerlof and R. D. Milbourne. "The Short Run Demand for Money". In: *The Economic Journal* 90.360 (1980), pp. 885–900. ISSN: 00130133, 14680297.

[34] O. J. Blanchard. "The Production and Inventory Behavior of the American Automobile Industry". In: *Journal of Political Economy* 91.3 (1983), pp. 365–400. ISSN: 00223808, 1537534X.

[35] R. J. Caballero and E. M. R. A. Engel. "Dynamic (S, s) Economies". In: *Econometrica* 59.6 (1991), pp. 1659–1686.

[36] M. Golosov and R. E. Jr. Lucas. "Menu Costs and Phillips Curves". In: *Journal of Political Economy* 115.2 (2007), pp. 171–199.

[37] F. Alvarez, A. Atkeson, and C. Edmond. "Sluggish Responses of Prices and Inflation to Monetary Shocks in an Inventory Model of Money Demand". In: *The Quarterly Journal of Economics* 124.3 (2009), pp. 911–967.

[38] D. Romer. "The monetary transmission mechanism in a general equilibrium version of the baumol-tobin model". In: *Journal of Monetary Economics* 20.1 (1987), pp. 105–122. ISSN: 0304-3932.

[39] F. Alvarez and F. Lippi. "Financial Innovation and the Transactions Demand for Cash". In: *Econometrica* 77.2 (2009), pp. 363–402.

[40] F. Alvarez et al. "Cash: A Blessing or a curse?" In: *Journal of Monetary Economics* 125 (2022), pp. 85–128. ISSN: 0304-3932.

[41] F. Alvarez and F. Lippi. "Cash burns: An inventory model with a cash-credit choice". In: *Journal of Monetary Economics* 90.C (2017), pp. 99–112.

[42] F. Alvarez, F. Lippi, and R. Robatto. “Cost of Inflation in Inventory Theoretical Models”. In: *Review of Economic Dynamics* 32 (Apr. 2019), pp. 206–226.

[43] I. Baley and A. Blanco. “Aggregate Dynamics in Lumpy Economies”. In: *Econometrica* 89.3 (2021), pp. 1235–1264.

[44] P. Bolton, H. Chen, and N. Wang. “A Unified Theory of Tobin’s q, Corporate Investment, Financing, and Risk Management”. In: *The Journal of Finance* 66.5 (2011), pp. 1545–1578.