Essays on Debt and Speculation

Janko Heineken

TESI DOCTORAL UPF / Year 2021

THESIS SUPERVISOR
Jaume Ventura
Department of Business and Economics, UPF

Universitat Pompeu Fabra Barcelona

Gewidmet meinen Eltern.

Aknowledgements

I want to thank my advisor Jaume Ventura for his support, generosity, and patience. Your feedback was invaluable in shaping this thesis.

I have spoken to many economists about the projects in this thesis (as well as other projects that did not make it) and I need to thank them for their insightful and challenging comments: Priit Jeenas, Vladimir Asriyan, Victoria Vanasco, Giacomo Ponzetto, David Nagy, Andrea Caggese, Filippo Ippolito, Christian Brownlees, Geert Mesters, Barbara Rossi, Dmitry Kuvshinov, Pablo Kurlat, Xavier Vives, Xavier Freixas, Albert Marcet, Vicente Quadrini, Fabiano Schivardi, Rüdiger Bachmann, Urban Jermann, Joan Monras, Lukas Hoesch, Gianmarco Ruzzier, Daniel Sturm, Davide Debortoli, Jordi Gali. There were many other generous discussants, commenters, seminar participants, and readers along the way. Thank you for making these projects better.

I am thankful for my abundantly talented collaborators Andrea Fabiani, Luigi Falasconi, and Ilja Kantorovitch. I was very privileged to work with you and to learn from you.

I want to thank the people who made it all possible behind the scenes, especially during the market: Libertad Gonzales, Marta Araque, and Laura Augusti in particular, and the whole econ department in general. At CREi Anna Cano, Anna Rios, and Carolina Rojas. At BGSE Marta Ledesma and Coral Ruano. I want to thank Studienstiftung des deutschen Volkes for financial support and the opportunity to be a part of their network.

Thanks to all the great people I met along the way, whether in Germany, in Spain, or anywhere else.

Lastly, I want to thank my family: Tabea, Mama, Papa, Niko, Sina, Mirjam, Oma, Heidi, Abuelo, Omi, and everybody else. I can always rely on you.

Abstract

My dissertation consists of three essays with the common theme of analyzing economic models of financial assets markets featuring heterogeneous market participants and testing the predictions of these models empirically. In the first essay, I present a simple model of heterogeneous traders and demonstrate that different sources of heterogeneity (discount factors vs. return expectations) imply opposite signs for the relationship of turnover and expected returns. This motivates me to analyze the correlational empirical relationship between asset turnover and measures of expected returns, both survey- and model-based, for the US stock and housing markets. I obtain mixed evidence on the sign of the turnover-expected return relationship in the stock market, while in the housing market the relationship is moderately positive.

In the second essay (authored jointly with Andrea Fabiani and Luigi Falasconi), I ask whether monetary policy influences the maturity structure of corporate debt. I answer this question by exploiting: i) time-series and firm-level data on debt maturity for the US corporate sector; ii) several proxies of FED monetary interest rate shocks. The results show that a loosening of the policy rate lengthens corporate debt maturity, an effect entirely driven by the adjustments of very large companies. I explain such findings through a model combining moral-hazard frictions and yield-seeking investors, who increase their demand of long-term debt-securities when the policy rate goes down. Only large and unconstrained companies can accommodate the demand shift. Empirical evidence on the response of corporate bond issuance by large companies and holdings by mutual funds validates the proposed mechanism.

In the third essay (authored jointly with Ilja Kantorovitch), I test the theoretical predictions of the differences-of-opinion literature by analyzing the extensive online discussion on Bitcoin by potential Bitcoin investors to build a time-varying sentiment distribution, defining disagreement as dispersion in sentiment. High disagreement is associated with negative subsequent returns, high turnover growth, and high volatility, confirming the theory's predictions. However, I do not find that an increase in disagreement increases the price, which is seemingly at odds with the theoretical prediction of disagreement leading to overpricing. As the theory predicts, the effect of disagreement

weakens significantly after shorting instruments were introduced at the end of 2017. The results are economically significant: at the monthly frequency, a one standard deviation increase in disagreement leads to a 9.2 percentage points lower cumulative return over the following eight months, and the adjusted R^2 of regressing contemporaneous returns on average sentiment and disagreement is 0.33.

Resum

La meva dissertació consta de tres assajos amb el tema comú d'analitzar l'impacte de la introducció d 'agents no racionals i heterogeni en els models de mercats d'actius financers. En el primer capítol, presento un model simple d'agents heterogenis i dedueixo que diferents causes d'heterogeneïtat (differentes factors de descompte versus differentes expectatives de rendibilitat) impliquen un signe diferent per a la relació entre la velocitat transactional d'actius i els rendiments esperats. Analitzo la relació empírica entre les mesures dels rendiments esperats i la velocitat transactional d'actius per als mercats borsari i de l'habitatge dels EUA i obtinc evidència mixta sobre el signe de la relació entre la velocitat transactional d'actius i els rendiments esperats.

En el segon capítol (escribido juntament amb Andrea Fabiani i Luigi Falasconi) pregunto si la política monetària influeix en l'estructura de venciments del deute corporatiu. Responc a aquesta pregunta utilitzant: i) dades de sèries temporals i a nivell d'empresa sobre el venciment del deute de el sector empresarial nord-americà; ii) diverses aproximacions a les pertorbacions dels tipus d'interès monetaris de la FED. Els resultats mostren que una flexibilització de l'tipus d'interès oficial allarga el venciment del deute corporatiu, un efecte totalment impulsat pels ajustos de les empreses molt grans. Explico aquests resultats a través d'un model que combina friccions de risc moral i inversors que busquen rendibilitat, que augmenten la seva demanda de títols de deute a llarg termini quan el tipus d'interès oficial baixa. Només les empreses grans i sense restriccions poden acomodar el canvi de demanda. L'evidència empírica sobre la resposta de l'emissió de bons corporatius per part de les grans empreses i les tinences dels fons d'inversió valida el mecanisme.

En el tercer capítol (escribido juntament amb Ilja Kantorovitch) poso a prova les prediccions teòriques de la literatura sobre les diferències d'opinió en mercats financers analitzant l'extens debat en línia sobre Bitcoin per construir una distribució de sentiment variable en el temps, definint el desacord com la dispersió en aquest sentiment. Alt desacord s'associa amb rendiments negatius, un alt creixement de la facturació i una alta volatilitat, confirmant les prediccions de la teoria. No obstant això, no trobo que un augment de l'desacord augmenti el preu en el present, el que sembla contradir la

predicció teòrica que el desacord condueix a la sobrevaloració. Com prediu la teoria, l'efecte de desacord es debilita significativament després de la introducció d'instruments de venda a l'descobert a la fi de 2017. Els resultats són econòmicament significatius: en la freqüència mensual, un augment d'una desviació estàndard en el desacord condueix a una rendibilitat acumulada 9.2 menor durant els vuit mesos següents, i el R^2 ajustat de la regressió de les rendibilitats contemporànies sobre el sentiment mitjà i el desacord és de 0.33.

Introduction

The question of how exactly expectations, beliefs, sentiment, and decision-making in asset markets are intertwined, has been of interest to academic economists, as well as to the public at large, for at least a century.[1] In this dissertation, I apply methods from financial economics and applied econometrics to better understand asset markets and especially the role of heterogeneity of their participants. All three chapters adhere to a theme of exploring behavioral assumptions in models of financial markets with heterogeneous participants, and testing the predictions of these models against data, be these data surveys of return expectations, the rich data of the corporate debt markets, or self-collected data consisting of online statements by potential investors.

My overarching goal is to understand how individual choices of heterogeneous market participants, such as firms and investors, drive aggregate variables: prices, volatility, transactions, investment, employment, and liquidity. A topic that I am specifically interested in is the role of market participants' heterogeneous expectations or sentiment in episodes of volatile asset markets: how they take transaction decisions, how they set their prices, and how their beliefs about future prices drive aggregate variables.

One motivation to ask these questions is that asset markets can go through speculative periods, only to then collapse. Asset price booms and busts have real consequences and analyzing these "exuberant" episodes is crucial to precisely understand their drivers, symptoms and consequences. My research methodology is to analyze simple models of financial markets with heterogeneous agents and to take the predictions of these simple models to the data. In all three chapters of this thesis, I carve out empirical phenomena that lend themselves to the analysis through the lens of heterogeneous agent models. By comparing the predictions of these models against the data, I want to contribute to find improved models of financial markets, especially those incorporating heterogeneous expectations and beliefs. Below, I will give a short summary of the three chapters of this

[1] The most prominent seminal example being Keynes (1936) and his notion of 'Animal Spirits', meaning sentiments of fear and exuberance that drive market participants to sell and buy in asset markets. However, there exist ample historical accounts of speculative phenomena, as presented at length in Aliber and Kindelberger (2015).

dissertation, which follow the just-described overarching lines of thought.

In the first chapter of this dissertation, titled "**How do Expected Returns Affect Turnover?**", I analyze a model that features many of the standard elements of models of heterogeneous traders, but is purposefully simple in its design. Risk-neutral agents trade among themselves a single asset, while their valuations for this asset fluctuate. I show that the sign of the empirical relationship of expected returns observed in the market and trading activity (measured as asset turnover) differs under two related sets of assumptions. In the first version of the model, I assume that market participants differ in their discount factors, but are rational. I show that this implies a positive relationship between turnover and expected returns, as a low discount factor leads to high turnover, and implies that the required return to hold the asset must be high. However, the opposite result emerges if the heterogeneity stems from the fact that agents have different expectations about the future asset return. In this case, a high expected return, actually decreases the incentive of agents to engage in transactions.

I analyze two asset markets, namely the US stock market and the US housing market, for which both model-based and survey-based expectation data are available. I find that the evidence on the relationship is in fact mixed: there is no strong relationship between turnover and expectation measures, and the results varies greatly by employed measure of expectations. If anything, I observe a moderately positive relationship in the housing market. On the one hand this might be an unsatisfactory result, on the other hand it can serve as a motivation to analyze expectation data using all its distributional features, in order to improve models of asset pricing. "**How do Expected Returns Affect Turnover?**" thus can be seen as only the initial step into a broader research agenda.

In Chapter 2 "**Monetary Policy and Corporate Debt Maturity**", which I co-authored with Andrea Fabiani and Luigi Falasconi, I find two "surprising" empirical results when analyzing both US monetary policy and the debt maturity structures of listed firms. The first is that if the central bank decreases the policy rate, firms tend to lengthen the maturity structure of their liabilities, the second is that this effect is driven entirely by the very largest firms: the Googles, AT&Ts and Walmarts. This is interesting for two reasons: the previous literature does not provide obvious answers to why this should happen and it gives rise to the suspicion that loose monetary policy, if sustained over long periods of time (such as in the last decade) can tilt the corporate maturity structure (especially of large firms) towards long maturities. Previous literature has suggested that there is an efficient maturity structure, but firms might deviate towards shorter (problematic because of increased roll-over risk) or longer (problematic because of monetary policy being less efficient) structures. The results point towards a possible endogeneity of monetary policy: loose policy can lead to long maturity structures, which then can lead to slow recoveries and thereby trigger even more sustained loose monetary policy. However, in the chapter I do not focus on this endogeneity, which serves as a motivation for the analysis and will be a subject for my future research.

Instead, I first establish that monetary policy has an effect on the corporate maturity structure, that is heterogeneous across firms, and present a suggestive mechanism. To do so, I build a novel model of credit contracts between risk-adverse investors and financially constrained firms. The crucial innovation in this model is that a subset of investors features reach-for-yield motives. When interest rates are low, these investors create demand pressure in the market for high-yield long-term debt, providing an incentive to firms for additional long-term debt issuance. In my framework, firms differ in their endowments and in equilibrium only well-funded (large) firms are unconstrained in their long-term debt issuance. I show that in general equilibrium, if the yield-seeking-motives are strong enough, "large" firms will lengthen their maturity, as they issue new long-term debt, but small firms will not. To lend credibility to the mechanism, I test three distinct predictions, namely that when the interest rate decreases 1) high-yield mutual funds lengthen their asset maturity; 2) large firms issue long-term debt; 3) the cost of finance for large firms decreases more than for small firms.

In Chapter 3 of this dissertation, I explore how to use text data, which consists of online statements by potential investors, to improve our understanding of the relationship of sentiment and asset returns. If researchers analyze the dispersion of investor valuations or investors' beliefs regarding future asset valuations, they would want to know what market participants actually think. However, availability of survey data on expectations and agents' beliefs about assets has been quite sparse, especially for assets other than equity. To increase the understanding of the effect of belief dispersion, I, in

Chapter 3 of this dissertation, titled "**Does Dispersed Sentiment drive Price, Volatility and Turnover for Bitcoin**?", which was authored jointly with Ilja Kantorovitch, look at the utterances of market participants during the Bitcoin boom-bust cycles of the past decade. To collect these data, I scraped millions of comments from the largest online forum on cryptocurrency and scored them on sentiment using a lexical algorithm. I analyze how belief dispersion is related to returns, volatility and turnover of Bitcoin and find that dispersed sentiment does predict negative subsequent returns, a fact that I demonstrate is in line with standard models in the literature on belief dispersion. However, I don't find a negative effect of high dispersion on contemporaneous returns, which one should expect according to the differences-of-opinion literature. I try to rationalize this finding, by exploring which other mechanisms could lead to this result. Interestingly, there is no lasting effect of belief dispersion on turnover, supporting my suspicion that disagreement is not necessarily a driver of high turnover during boom times (as I shown in Chapter 1, there might be reason to believe the exact opposite).

Contents

List of Figures

List of Tables

Chapter 1

HOW DO EXPECTED RETURNS AFFECT TURNOVER?

1.1 Introduction

High trading volume (more sensibly analyzed relative to the overall asset stock, so in terms of turnover) is often cited as an ubiquitous feature of times of exuberant investor sentiment or even asset price "bubbles".[1] However, it is far from obvious why agents would trade more often when they are optimistic, as from a rational perspective trade should be governed by agents' valuations of the asset relative to those of other market participants, which determine the gains from trade. This is a question that lends itself to be explored more deeply in a model of heterogeneous agents. In such a model, if the average market participant or the marginal investor become more optimistic, will there be more trade and which are the crucial assumptions to deliver such a result?

In this paper, I derive the implications of a simple model with risk-neutral heterogeneous agents for the relationship of turnover and return expectations. In fact, this simple model of traders, who are heterogeneous in either their discount factors or their return

[1] For a general motivation of this thinking, see Hong and Stein (2007b).

expectations, predicts that the sign of the turnover-expected return relationship depends on which of the two is the underlying source of heterogeneity. A crucial insight is that if the distribution of the "source" around its mean (which can vary over time) is fixed, a low level of the "source" variable (so either low discount factors, or overall pessimist expectations) implies high turnover, as gains from trade are larger relative to the transaction costs (assumed to be proportional to the price).

However, in equilibrium the average discount factor is the inverse of the expected return, following the logic of market clearing. Thus, in a world where agents differ in their discount factors, high expected returns (which imply low discount factors, as agents demand a high expected return to be compensated for holding the asset despite valuing it so little in the future) lead to high turnover. But in a world in which agents differ in their return expectations, instead of their discount factors, a high average return expectation, optimism, coincides with little trading activity.

Having derived these results, I then take a closer look at the correlational empirical relationship between asset turnover and various measures of expectations, both in the US stock and housing markets. Using both model-based expectations (such as the dividend-price ratio), which impose the assumption of rational expectations, as well as various survey-based expectation measures, I find that in fact there is no significant correlation between turnover and return expectations, neither positive, nor negative, if observed over the long-term. If anything, there is a moderately positive relationship in the housing market, although it is not conclusive.

This paper adds to a growing literature that thinks carefully about the relationship of empirically measured expectations, both model- and survey-based, and asset market phenomena. While most other papers have focused on returns and other phenomena related to asset prices, my paper is unique in its approach as it analyzes the relationship of measures of return expectations and trading activity. There exists an established literature indirectly deriving measures of return expectations from standard models of asset markets, examples of which are Campbell and Cochrane (1999), Lettau and Ludvigson (2001), and Cochrane (2011), the last of which provides a high-level summary the literature. However, turnover does not play a role in these considerations, and neither do heterogeneous expectations. Unsurprisingly, turnover has received some attention in the asset price literature, most prominently by Lo and Wang (2006, 2004) and Lo, Mamaysky, and Wang (2004). These papers focus primarily on the how to employ turnover to derive an additional asset pricing factor and do not dwell on expectations or the sources of traders' heterogeneity.[2]

A more recent literature analyzes the relationship of survey-based expectation measures and returns (as well as the aforementioned model-based expectation measures) and finds that a.) they appear to be measuring something meaningful (not just noise) and, b.) that they are seemingly at odds with model-based measures: e.g. Greenwood and Shleifer (2014) analyze long-running expected return surveys and find that they are negatively associated with model based returns, as well as with future returns. Adam, Mateev, and Nagel (2021) explore whether participants in expectation surveys do not report their beliefs but rather their preferences, and reject this hypothesis. Relative to these papers, my paper focuses on the relationship of heterogeneous expectations and trading activity.

Finally, there exists a literature that associates the occurrence of trading frenzies with bubbly or "exuberant" episodes, the case for which is made in Gallant, Rossi, and Tauchen (1992), Scheinkman and Xiong (2003) and Hong and Stein (2007b). More recently, there have been several theoretical models of bubbly episodes, in which high turnover is predicted, e.g. Barberis et al. (2018) devise a model of extrapolative bubbles, which predicts a) that there is substantial trading volume in bubbly episodes, and b) that there should be a positive relationship between past return and trading volume. Bordalo, Gennaioli, and Kwon (2020) explore a model of excessive optimism about the development of asset fundamentals and link it to asset price bubbles, which come with high turnover. This literature motivates my research question about what a simple model of

[2]A related, but not immediately relevant, strand of the literature is that of Chen, Hong, and Stein (2001) who think about how turnover fares as a predictor of future returns. They show that negative skewness of returns is more pronounced in periods following heavy trading activity. This provides some evidence in favor of the idea of differences-of-opinion: when disagreement is large, there is would be more trading activity and large negative price movements are more likely, as the traders that "sit out" have more negative opinions. However, they and Greenwood, Shleiger, and You (2018) find that high turnover does not forecast expected returns, a result that appears contradictory to the differences-of-opinion literature.

heterogeneous agents can tell us about the relationship of expectations and turnover.

My paper differs from all listed papers in the fact that I derive the relationship of turnover and expectations from a simple model of heterogeneous traders, focus on the role played by different assumptions on the source of heterogeneity and test the clear predictions against both model- and survey-based measures of expectations. The rest of the paper is structured as follows: Section 1.2 presents the prediction of a simple model of heterogeneous traders and time-varying expectations, Section 1.3 presents an empirical test of these predictions. Section 1.4 summarizes the results and concludes.

1.2 A Simple Model of Expected Returns and Turnover

1.2.1 A Model of Heterogeneous and Time-varying Discount Factors

I model a discrete-time economy in which a continuous unit mass of agents trade among themselves a single asset, which is in supply N. Each agent is indexed by $i \in [0, 1]$ and receives a dividend from the asset while being its "owner". Agents are risk–neutral and have rational expectations, discount the future at discount factor β_t^i (which fluctuates over time and differs between agents), and have a unit demand for the asset (so they decide to either hold the asset or not). The asset issues a dividend in each period t, denoted by d_t. In the beginning of each period the dividend is paid out, then agents can decide whether they want to pay price q_t to purchase the asset in a Walrasian market. I will assume that the distribution f_β of the β_t^i is symmetric and fixed around a value β_t, which itself however is allowed to fluctuate over time. If a transaction occurs, both buyer and seller have to pay a transaction costs κ that is proportional to the price q_t. The transaction costs are symmetric between buyer and seller.[3]

We can express this model in terms of the following value functions. The value for individual i of being an owner of an asset in t is

[3]This is merely an expositional assumption and any distribution of the transaction cost among the transaction parties would deliver equivalent insights.

$$V^{o}\left(\beta_{t}^{i},q_{t}\right)= \max\{d_{t}+(1-\kappa)\,q_{t}+\beta_{t}^{i}E_{t}\left[V^{no}\left(\beta_{t+1}^{i},q_{t+1}\right)\right], \quad , \quad d_{t}+\beta_{t}^{i}E_{t}\left[V^{o}\left(\beta_{t+1}^{i},q_{t+1}\right)\right]\}$$

while the value of not being an owner is accordingly

$$V^{no}\left(\beta_{t}^{i},q_{t}\right)=\max\{ -(1+\kappa)\,q_{t}+\beta_{t}^{i}E_{t}\left[V^{o}\left(\beta_{t+1}^{i},q_{t+1}\right)\right], \quad . \quad \beta_{t}^{i}E_{t}\left[V^{no}\left(\beta_{t+1}^{i},q_{t+1}\right)\right]\}$$

The model's timing is as follows: in the beginning of t, owners receive d_t, and then make their trading decisions. In equilibrium, there are thresholds for the individual discount factors separating sellers from non-sellers and buyers from non-buyers. I denote these thresholds as: $\bar{\beta}_t$ for the discount factor above which an agent becomes a buyer, and $\underline{\beta}_t$ for the discount factor below which an agent becomes a seller. The asset price q_t must clear the market in equilibrium:

$$(1-N)\,\bar{\mathrm{p}}_t = N\underline{\mathrm{p}}_t, \tag{1.1}$$

where $\bar{\mathrm{p}}_t$ denotes the share of non-owners that will buy in period t and $\underline{\mathrm{p}}_t$ denotes the share of owners that will sell in period t. For simplicity, I always assume that $N=1/2$, which means that at any time exactly half of the agents will be owners.

Proposition 1.1. *Under the assumption that f_β is a symmetric distribution that is only time-varying in its first moment β_t, the individual β_t^i are drawn from this distribution in an i.i.d fashion, and that N=1/2, we have:*

1. The asset price follows a standard pricing equation

$$q_t = \beta_t E_t\left[d_{t+1}+q_{t+1}\right] \tag{1.2}$$

2. Turnover is determined by

$$\bar{p}_t = 1 - F_\beta\left(\frac{1+\kappa}{E_t\left[R_{t+1}\right]}\right), \tag{1.3}$$

where

$$R_{t+1} = \frac{d_{t+1} + q_{t+1}}{q_t}. \tag{1.4}$$

This proposition implies that: a.) although we have differences in the individual valuations of the asset, they have no effect on the equilibrium price, which is only determined by the average discount factor and the (expected) dividends, and b.) turnover is a positive function of expected returns. The mechanism is simple: in equilibrium the expected return is determined by the average discount factor

$$\frac{1}{\beta_t} = E_t\left[R_{t+1}\right].$$

Trading with another agent presents a more attractive venture if prices are low relative to average expected returns, as transaction costs are proportional to the asset price q_t.

This effect increases with the expected return, as the distribution of individual discount factors F_β is fixed around β_t. Concretely, we have that the threshold discount factor to buy shrinks by a factor larger than one, if the average discount factor falls[4]:

$$\bar{\beta}_t = (1+\kappa)\,\beta_t.$$

1.2.2 Heterogeneous and Time-varying Discount Factors vs Subjective and Time-Varying Expectations

In the version of the model that was presented above, I have employed the assumption of an individual discount factor, which is a.) time-varying (every period agents draw a new β_t^i), and b.) heterogeneous between market participants. This induces trading decisions, driven by what could either be considered fluctuations over time in the discount factor or in their required return for reasons outside the model. In the above model, agents are rational and differ in their discount factors, a set-up which maps empirically to model-based expectation measures, such as the dividend-price ratio. However, I will also empirically test the relationship of asset turnover and survey-based expectations measures. For this it is a sensible assumption that: a.) the market participants have

[4] If the dispersion of β_t^i would increase for lower β_t, this would be a counteracting effect. However, it would be unclear why the dispersion of beliefs would increase, while the expected return increases. However, to analyze the distribution of disaggregated return expectations survey responses is certainly an interesting future research project.

heterogeneous return expectations, and b.) the mean of these return expectations fluctuates over time, which are both empirical phenomena that we observe in surveys on expectations.

In the following, I want to show that the model is essentially identical under this assumption, except for the crucial insight that in this version predicts the exact opposite sign for the turnover-expected return relationship. Consider a second model that only differs in its set-up in the following points:

1. While the discount factor varies over time, it is not heterogeneous across market participants. Every trader discounts at the same β_t.

2. The expectations of the future value of owning the asset relative to the current price $E_t^i\left[\left\{V^o\left(\beta_{t+1}, q_{t+1}\right) - V^{no}\left(\beta_{t+1}, q_{t+1}\right)\right\}/q_t\right]$ are heterogeneous between agents. In fact, it adheres to a distribution $f_{\frac{\mathcal{E}}{q}}$ around a time-varying mean, which is $E_t\left[V^o\left(\beta_{t+1}, q_{t+1}\right) - V^{no}\left(\beta_{t+1}, q_{t+1}\right)/q_t\right]$.

Note that as q_t is known to all market participants in period t

$$E_t^i\left[V^o\left(\beta_{t+1}, q_{t+1}\right) - V^{no}\left(\beta_{t+1}, q_{t+1}\right)/q_t\right] = E_t^i\left[V^o\left(\beta_{t+1}, q_{t+1}\right) - V^{no}\left(\beta_{t+1}, q_{t+1}\right)\right]/q_t.$$

Under these slightly different circumstances, there again will be thresholds:

$\bar{E}_t\left[\left\{V^o\left(\beta_{t+1}, q_{t+1}\right) - V^{no}\left(\beta_{t+1}, q_{t+1}\right)\right\}\right]/q_t$ and $\underline{\mathrm{E}}_t\left[\left\{V^o\left(\beta_{t+1}, q_{t+1}\right) - V^{no}\left(\beta_{t+1}, q_{t+1}\right)\right\}\right]/q_t$

for buying and selling.

I will use $\mathcal{E}_t^i$ and $\mathcal{E}_t$ as a shorthand for $E_t^i\left[V^o\left(\beta_{t+1}, q_{t+1}\right) - V^{no}\left(\beta_{t+1}, q_{t+1}\right)\right]$ and $E_t\left[V^o\left(\beta_{t+1}, q_{t+1}\right) - V^{no}\left(\beta_{t+1}, q_{t+1}\right)\right]$, respectively, as well as $\bar{\mathcal{E}}_t$ and $\underline{\mathcal{E}}_t$ as shorthands for the threshold values for $E_t^i\left[V^o\left(\beta_{t+1}, q_{t+1}\right) - V^{no}\left(\beta_{t+1}, q_{t+1}\right)\right]$ at which market participants buy and sell.

The following proposition summarizes the equilibrium of the model in this case.

Proposition 1.2. *Under the assumption that $f_{\frac{\mathcal{E}}{q}}$ is a symmetric distribution that is only* *time-varying in its first moment $\mathcal{E}_t/q_t$, the individual $\mathcal{E}_t^i/q_t$ are drawn from this distribution in an i.i.d fashion (across time and individual agents), and that N=1/2, we have:*

1. The asset price follows a standard pricing equation

$$q_t = \beta_t E_t \left[d_{t+1} + q_{t+1} \right] \tag{1.5}$$

2. Turnover is determined by

$$\bar{p}_t = 1 - F_{\frac{\mathcal{E}}{q}} \left(\left(1 + \kappa \right) E_t \left[R_{t+1} \right] \right), \tag{1.6}$$

where

$$R_{t+1} = \frac{d_{t+1} + q_{t+1}}{q_t}. \tag{1.7}$$

In fact the threshold for buying is $\bar{\mathcal{E}}_t = (1 + \kappa) \mathcal{E}_t$ and we have again that in equilibrium $1/\beta_t = E_t [R_{t+1}]$. The model's equilibrium behavior is identical to before, with one exception: instead of a positive relationship of turnover and expected returns this model predicts a negative relationship! The only thing that fundamentally changed between the two models was the source of heterogeneity. In the first model the source of heterogeneity is the discount factor, in the second the source of heterogeneity are dispersed expectations about the future value of the asset and thus its return. In the following, I want to explain the mechanism that drives these results and explain what causes the difference.

1.2.3 Interpretation - Understanding the Mechanism

In both versions of the model, that I have presented the mechanism by which heterogeneity generates trade works in a similar manner. If we ordered the market participants according to how much they valued the asset, this order would be reshuffled every period and thus generate a motive for trade. With the simplifying assumption that $N = 1/2$, which means that always half of the market participants will hold the asset, and that the heterogeneous variable changes over time in an i.i.d. fashion, I am analyzing the extreme case: the cue is reshuffled every period, without memory of the past. In a world without transaction costs ($\kappa = 0$), the half of the mass of market participants who value the asset the most would always hold the asset and trade every period would be constant: those who move to the top half buy the asset from those who move down. In the language of the model: without transaction costs the thresholds to both buy and sell are just identical to the mean of the respective distribution: $\beta_t = \bar{\beta}_t = \underline{\beta}_t$ and $\mathcal{E}_t = \bar{\mathcal{E}}_t = \underline{\mathcal{E}}_t$. Transaction costs introduce a wedge between the buying and selling thresholds and create a no-trade region, in which market participants do not buy or sell, because the transaction costs are higher than the gains of trade. In both models the average discount factor and the average expected return vary over time[5], in fact in both models:

$$\beta_t = \frac{1}{E_t\left[R_{t+1}\right]}. \tag{1.8}$$

This equation tells us that the average discount factor must equal the inverse of the expected market return, as agents are risk-neutral. This means that the price will be such that the return equals to the average discount factor. Would it be higher, more agents would want to buy and less sell and the return would adjust downwards, would it be lower, less agents would want to buy and more sell.

The main insight when comparing these models was that in the model, in which the discount factor is heterogeneous, there is a positive relationship between turnover and expectations. In the model, in which there are heterogeneous expectations about future returns, the same relationship is negative. Equation 1.8 gives clues to why this result emerges. In the first model, if we were to measure a higher expected return in the data, it would imply (as agents are perfectly rational) that the average discount factor has decreased. In other words: agents value the future less and are thus demanding higher expected returns to hold the asset. I assume that the distribution around the average discount factor β_t does not change over time, and we know that $\bar{\beta} = (1+\kappa)\,\beta_t$ and therefore the thresholds are moving closer to the average discount factor and the no-trade zone shrinks.

If however the source of heterogeneity are agents' expectations about the asset's future value, the opposite effect occurs. An increase in measured return expectations implies that the average discount factor is lower too, but trade is governed by the dis-

[5]It is well established in the asset pricing literature that one of the two must be time-varying to explain asset price movements.

tribution of individual expectations, not the distribution of the discount factor. In the distribution of return expectations $f_{\frac{\varepsilon}{q}}$, if the mean $\mathcal{E}_t$ increases, the thresholds move further apart and the no-trade zone increases, as $\bar{\mathcal{E}}_t = (1+\kappa)\,\mathcal{E}_t$. Thus, we see less turnover in periods with high return expectations.

Another difference worth highlighting is the interpretation of return expectations between the models. In the model with heterogeneous discount factors, the agents are perfectly rational, but they differ in their preferences and thus demand different minimum expected returns to hold the assets. All agents agree that the expected return in the market is $E_t\left[R_{t+1}\right]$. If we would ask these agents (as return expectation surveys do) what return they expect, there would not be a large dispersion. This model of the world lends itself to comparison with model-based expectation measures. In the model with heterogeneous expectations, however, agents are not rational. They agree-to-disagree about what the expected return next period will be. In such a model the expected return is determined again by the discount factor, and thus by the preferences of agents. However, agents do not understand this and stick with their subjective expectations. This model set-up lends itself to be mapped to survey-based expectations measures, in which respondents clearly agree-to-disagree and display a wide dispersion of opinions.

1.3 Empirical Evidence

In the simple model of Section 1.2, we find that there is a positive relationship between expected returns and turnover when market participants are heterogeneous in their discount factors, and a negative relationship when market participants are heterogeneous in their return expectations. In the following empirical section, I describe the empirical correlational relationship of turnover and different measures of expectations, both model- and survey-based. I focus on simple correlations, as without an assumption on the functional form of the distributions f_β and $f_{\frac{\varepsilon}{q}}$, we do not know the functional form of turnover as a function of return expectations either. It could be non-linear, but we at least know that the model predicts the relationship to be monotonic. Thus, focusing on the simple linear correlation is a reasonable approach to understand the relationship between the variables. I will restrict my analysis to the US stock and housing markets, due to reasons of data availability and quality.

1.3.1 Data

I will use three types of data in the analysis, for both the US housing and the US stock market, respectively: data on aggregate turnover, model-based expected return measures (meaning measures, which according to a model with rational agents approximate expected returns), and survey-based measures, for which individuals (usually investors) are asked about their expectations regarding the developments in the respective markets in the near future (usually over the next year). In the following, I will describe my data sources in detail.

1.3.1.1 Stock Market

Turnover

I use two different series for turnover in the stock market. The first series is value weighted turnover for the US stock market, as calculated from the CRSP data set.[6]. This data set begins as early as 1925, but for reliability reasons I use the time series starting in 1962, as do Lo and Wang (2010). The data set is available on a daily basis, but for the analysis here the highest frequency needed is monthly.
The second data set is the turnover of the 1000 largest shares traded on the New York Stock Exchange, which begins in 1973, and is again available daily, but aggregated to monthly for my purposes.[7] Both measures are highly correlated and in the following I report results from using either one or the other, not both.

Model-based Return Measures

I use three different types of model-based return measures: the realized market returns from both the CRSP (value-weighted) and NYSE data sets (described above for turnover), the value-weighted dividend-price ratio from CRSP, building on the widely accepted idea that fluctuations in the dividend-price ratio are reflecting fluctuations in return expectations, much more so than future dividend growth, see Cochrane (2011), and the consumption-wealth ration, as described in Lettau and Ludvigson (2001).[8] The

[6]See the Lo and Wang (2010) handbook chapter for the detailed calculation.
[7]These data are available from Refinitiv Datastream as “US-DS Market”.
[8]These data are available on Martin Lettau’s website.

rationale behind the consumption-wealth ratio is that if the permanent income hypothesis holds, high prices that are determined by low required returns should imply a lower consumption-wealth ratio. In the model in this paper, in which dividends are consumed on the spot, this idea is directly related to the dividend-price ratio.

While I am interested in the correlation of turnover at time t with the dividend-price ratio and the consumption-wealth ration at time t, the relevant correlation of realized returns and turnover is that of turnover in t and realized returns over the period between t and $t+s$, as under rational expectations realized returns today are an approximation for expected returns in the present. Throughout, I will calculate the correlation for $s = 12$ months ahead.

Survey-based return measures

There are several surveys available that ask participants about their expectations for the future development of the stock market. I use the following surveys: the UBS/Gallup survey, the Graham-Harvey/Duke Fuqua CFO survey[9], the Livingston Survey[10], the survey administered by Robert Shiller[11], and finally the weekly "Investor Sentiment" survey by the American Association of Individual Investors.[12] In Table X the five surveys are described in detail. While some of the measures directly map to expectations of expected returns, others, like the AAII survey, instead ask investors whether they think the market will go up, down, or remain the same. In this case I use the difference between the share of bulls and bears as an approximation for the return. This makes sense for two reasons. Firstly, Greenwood and Shleifer (2014) show that these qualitative measures are highly correlated to survey measures that ask directly for expected returns. Secondly, it maps to the theory: in the second model above there is a monotone relationship between the ratio of bulls and bears and the average expected return (as the distribution of individual expected returns is fixed and symmetric around the average expected return).

[9]Made available on the website of the Duke CFO Global Business Outlook.

[10]Administered and made available by the Federal Reserve Bank of Philadelphia.

[11]Made available at the Yale International Center for Finance as "United States Stock Market Confidence Indices".

[12]Made available on the AAII website.

1.3.1.2 Housing Market

Turnover

For the US housing market, I calculate turnover as the ratio of sales of existing homes relative to the total housing stock. The number of sales for existing homes is collected and made available by the National Association of Realtors, and is reported on a monthly basis since 1970. I take the US housing stock from the American Housing Survey, administered by the Department of Housing and Urban Development and the U.S. Census Bureau.

Model-based Expectation

I calculate the return on housing, using the House Price Index provided by the Federal Housing Finance Agency, which starts in 1970. To calculate the return, rent is treated like a dividend and derived from the Rent of Primary Residence CPI, which is calculated by the Bureau of Labor Statistics. If I just used home price growth instead of treating the average rent payment like a dividend, I would get very similar results, as most of the movement in housing returns is driven by price changes.

Survey-based Expectation Measures

I use the Zillow Home Price Expectation survey, which was initiated in 2010 and is administered on a quarterly basis. Investors and Experts are asked what they expect for the growth of real estate prices year-over-year. As the survey asks for the total change during the calendar year (so not the year-on-year change), I use the expectations that participants report to have for the year-over-year change in the next calendar year.

1.3.2 Empirical Relationship of Asset Turnover and Survey-based Expectations

For both the stock and housing markets, I linearly detrend turnover and compute the correlation coefficient with each of the respective survey-based return expectation series. The results are summarized in Table 1.1. It seems clear that the survey-based expectation measures do not provide much supportive evidence in either direction: for the Shiller, Livingston and AAII measures the correlation coefficient is weakly positive (in the case of Shiller's valuation measure, even moderately positive). However, in the Gallup and Duke CFO surveys the relationship is moderately negative. Furthermore, in the Duke and Livingston surveys the respective relationships are not statistically significant. The co-movement of the series with turnover are displayed in Figures 1.A.1–1.A.10 in the appendix.

Admittedly, the surveys all differ slightly, or even substantially, in their methodology, so different results on the sign could be driven by the differences in sample selection and by the fact that with some surveys, I am restricted to a qualitative measure, such as the share of bullish investors, while with others I use a quantitative return expectation. However, it is striking that especially those surveys that are similar in methodology, like AAII and Gallup, don't seem to be statistically related to turnover in the same way. The most immediate interpretation is that there is simply no robust relationship between survey-expectation and turnover.

In the housing market, we observe a significant positive correlation of 50% between survey-expectations from the Zillow survey and housing turnover. Of course for this example the sample period is relatively short, as the Zillow survey was only initiated in 2010 and during the sample period home sales and prices have done nothing but increase steadily. The corresponding figure is Figure 1.A.11. [13]

1.3.3 Empirical Relationship of Asset Turnover and Model-based Expectations

Table 1.2 summarizes the correlation of stock market turnover with the model-based expectation measures. Again, the evidence for any side is weak: we see no correlation for the value-based realized return over the subsequent 12 months after turnover is measured, only a small positive correlation for the dividend-price ratio, and a negative, but insignificant, correlation for the consumption-wealth ratio. As with the survey-based

[13] An interesting artifact is that despite decent growth of home prices, survey participants are not further increasing their return expectations (as we might expect from extrapolators, which are a common building block in many behavioral models of housing markets). They seem to have settled in at around an expected growth of 3-4% annual nominal increase.

Table 1.1: Correlation of Detrended Stock Market Turnover with Survey-based Expectation Measures

Survey-based Measures	Correlation	Period + Frequency
Shiller One-Year Individual	0.2*** [0.07,0.32]	Biannually from Apr 1999, monthly since Jul 2001.
Shiller One-Year Institutional	0.27*** [0.15,0.38]	Biannually from Oct 1989, monthly since Jul 2001.
Shiller Valuation Individual	0.53*** [0.43, 0.62]	Biannually from Apr 1999, monthly since Jul 2001.
Shiller Valuation Institutional	0.43*** [0.33,0.53]	Biannually from Oct 1989, monthly since Jul 2001.
AAII Member Survey	0.34*** [0.26,0.43]	Weekly, aggregated to monthly Jun 1987- Nov 2020.
Gallup/UBS Survey	-0.59*** [-0.68,-0.48]	Monthly (with gaps) Oct 1996 - Nov 2011
Duke CFO Mean	-0.14 [-0.36, 0.1]	Quarterly from Oct 2000 - Dec 2020
Duke CFO Median	-0.3** [-0.51, -0.06]	Quarterly from Mar 2004 - Dec 2020
Livingston Mean	0.12 [-0.09, 0.31]	Biannually from Jun 1973 - Jun 2020
Livingston Median	0.33*** [0.14, 0.5]	Biannually from Jun 1973 - Jun 2020
	p-value:<1%***, <5%**, <10%*	All Shiller Survey Data up to Nov 2020

measures, there is no smoking gun either way. Figures 1.A.12–1.A.14 display the co-movement of the respective series. As a robustness check, I also tested the correlation while excluding the years 2006-2012 (inclusive) from the sample, which displayed ab-

Table 1.2: Correlation of Detrended Stock Market Turnover with Model-based Expectation Measures

Model-based Measures	Correlation	Period + Frequency
12 Months Realized Returns	-0.06 [-0.13, 0.02]	Monthly from 1962-2020.
Dividend-Price Ratio	0.23*** [0.16, 0.3]	Monthly from 1962-2020.
Consumption-Wealth Ratio	-0.11 [-0.22, 0.02]	Quarterly from Apr 1952.

p-value:<1%***, <5%**, <10%*

normally high turnover, but without a significant change in the results.

In the housing market we again observe a moderately positive correlation between the market return and housing turnover, as can be seen in Figure 1.A.15. Measured at annual frequency, the correlation with the contemporaneous realized return is 43%. If I move to the quarterly frequency, measure turnover per quarter and compute the correlation with the realized return in the subsequent four quarters, we find a correlation coefficient of 44%. In summary, I conclude that the evidence for any kind of significant relationship of turnover and measures of return expectations in the US stock market is weak, and that, if anything, there is evidence for a moderately positive relationship in the US housing market.

1.4 Conclusion

I presented a simple model of heterogeneous trading and showed that different sources of heterogeneity imply different signs for the relationship between trading activity (measured as asset turnover) and return expectations. The model assumes a fixed distribution of the heterogeneous variable around its mean and proportional transaction costs. These ingredients imply that an increase in the mean of the heterogeneous variable decreases the no-trade region and reduces turnover. If the source of heterogeneity is the discount factor of traders, the model predicts a positive relationship between turnover and expected returns, which is driven by the fact that a low average discount factor (which comes with high turnover) implies that the individually required returns to hold the asset, and therefore the market expected return, are high.

However, if the source of heterogeneity in the model is that traders have different expectations of future returns, the relationship between turnover and expected returns is predicted to be negative! The reason is that if average expected returns increase in this version of the model, the no-trade zone widens and turnover decreases.

Empirically, I set out to analyze the relationship between turnover and expected returns in the data, using both model-based and survey-based measures. The evidence on the sign, however, is mixed, with the correlation alternating between negative and positive, depending on the exact measure and the sample period. If anything, there is a moderately positive relationship between turnover and return expectations in the housing market. This result might seem disappointing, but it is in fact interesting, as it provides evidence against the view that trading frenzies are a symptom of asset price bubbles or "exuberant" expectations.

There are avenues for future research: a general model of asset prices and turnover is needed to think about how the additional information added by turnover could help researchers to disentangle the stochastic discount factor, irrational exuberance, and rational bubbles as drivers of volatility in asset markets. Intuitively, the model tells us that if model-based measures of returns are high, but turnover is low, the reason behind the high expected returns cannot be the stochastic discount factor, but a richer model becomes necessary if we want to think about these phenomena more thoroughly when agents are risk-averse.

Lastly, future research can explore the properties of the distribution of survey expectations. In the model, I assumed that the distribution of expectations only changes over time in its first moment. However, this needs to be tested and if there are significant changes in the distribution over time, this could again be exploited to improve our understanding of price volatility in asset markets.

Appendix

1.A Figures

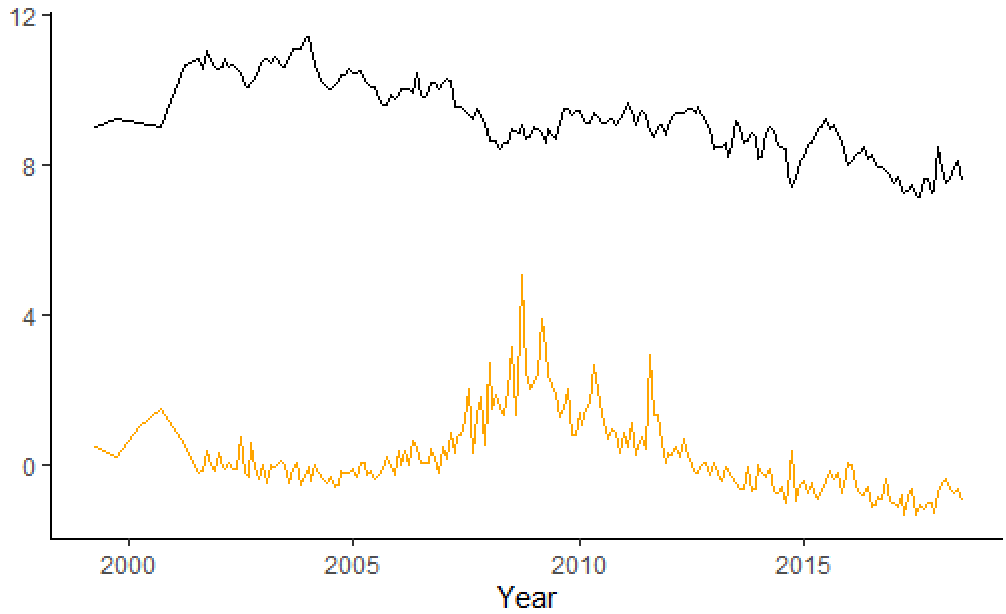

Figure 1.A.1: Shiller Individual Investors One-year Ahead Expectations.

Displayed are the time series depicting the linearly detrended turnover in the US and the time series depicting %bulls-%bears in the Shiller Individual Investor Survey, in which investors are predicting the market return one year ahead. Both series are standardized by dividing through their respective standard deviations. The correlation coefficient is 20%.

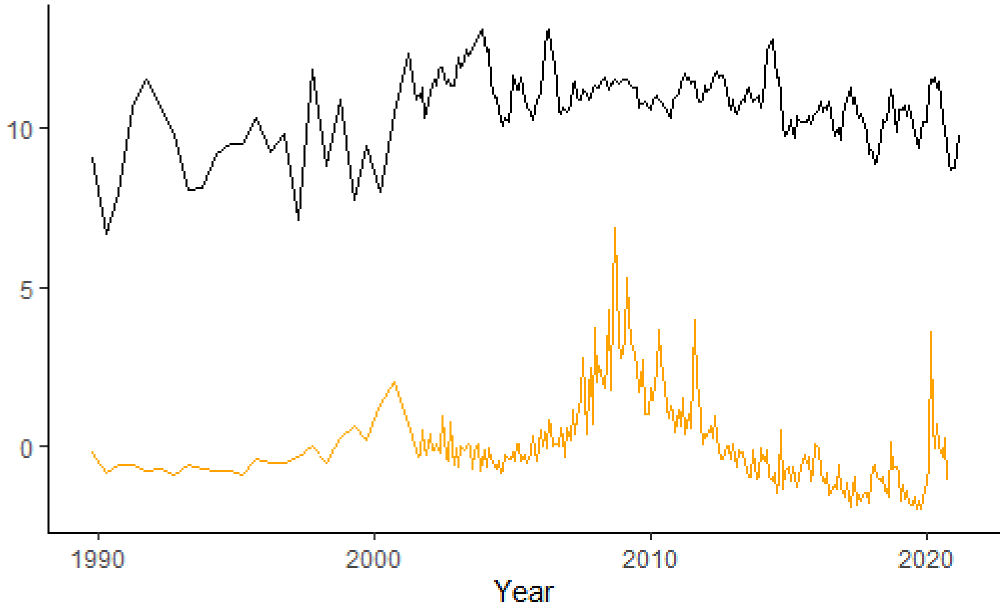

Figure 1.A.2: Shiller Institutional Investors One-year Ahead Expectations.

Displayed are the time series depicting the linearly detrended turnover in the US stock market (orange) and the time series depicting %bulls-%bears in the Shiller Institutional Investor Survey (black), in which investors are predicting the market return one year ahead. Both series are standardized by dividing through their respective standard deviations. The correlation coefficient is 27%.

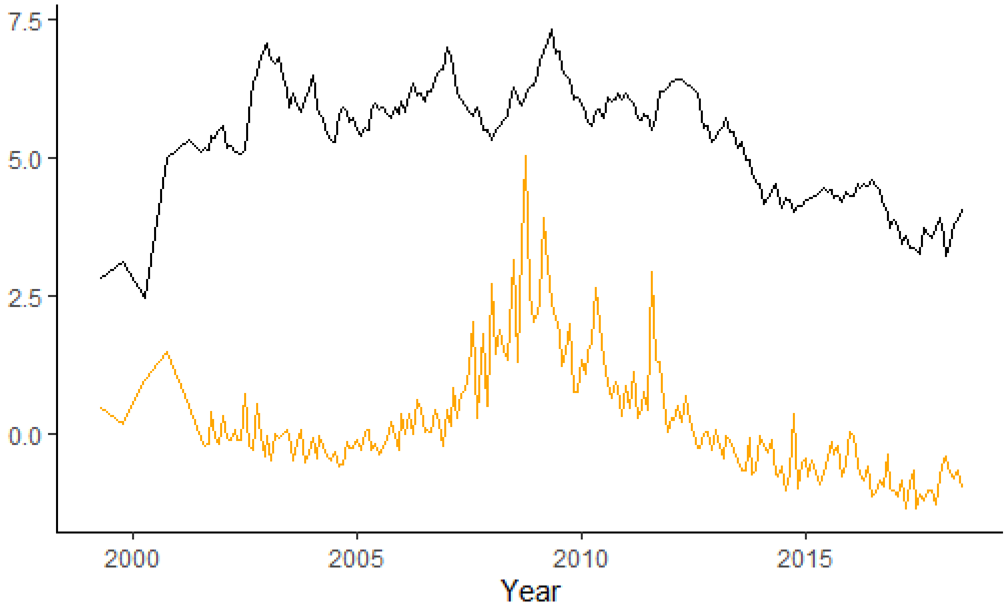

Figure 1.A.3: Shiller Individual Investors Market Valuation Survey.

Displayed are the time series depicting the linearly detrended turnover in the US stock market (orange) and the time series depicting %bulls-%bears in the Shiller Individual Investor Survey (black), in which investors are asked whether they perceive the market to be under- or overvalued. Both series are standardized by dividing through their respective standard deviations. The correlation coefficient is 53%.

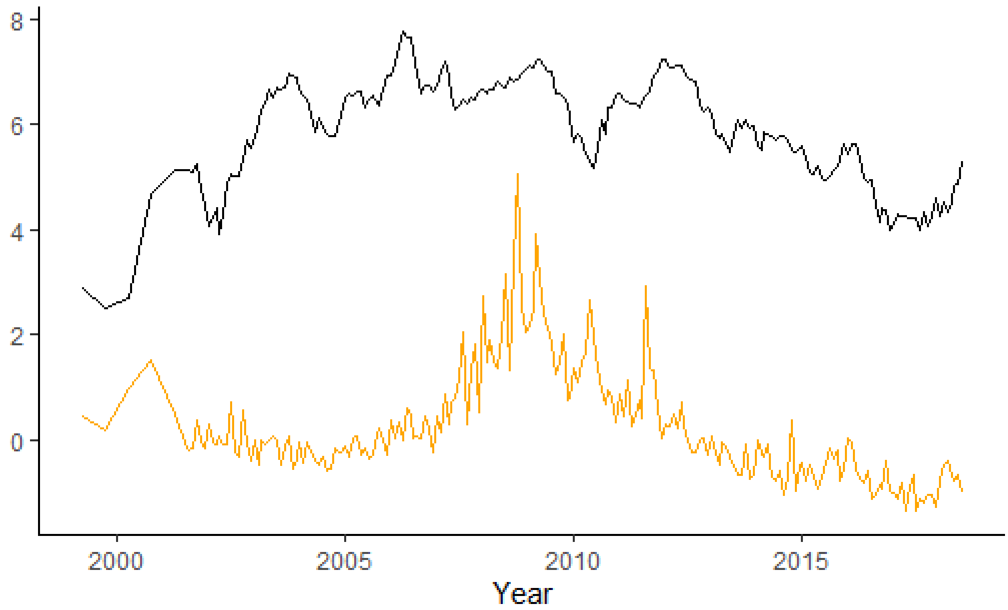

Figure 1.A.4: Shiller Individual Investors Market Valuation Survey.

Displayed are the time series depicting the linearly detrended turnover in the US stock market (orange) and the time series depicting %bulls-%bears in the Shiller Institutional Investor Survey (black), in which investors are asked whether they perceive the market to be under- or overvalued. Both series are standardized by dividing through their respective standard deviations. The correlation coefficient is 43%.

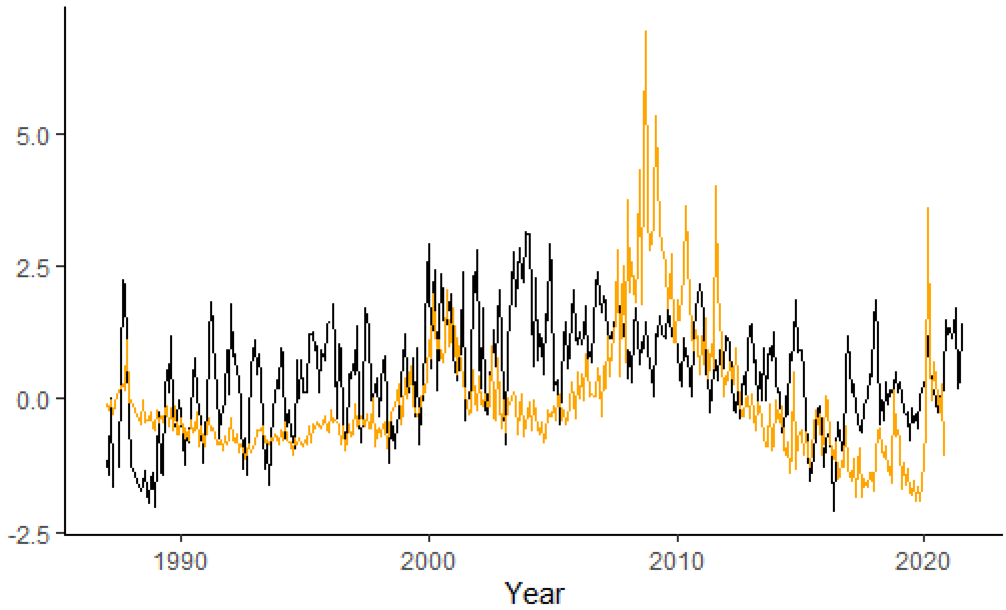

Figure 1.A.5: American Association of Individual Investors Member Survey.

Displayed are the time series depicting the linearly detrended turnover in the US stock market (orange) and the time series depicting %bulls-%bears in the American Association of Individual Investors Member Survey (black), in which AAII members are asked whether they are "bullish" or "bearish". Both series are standardized by dividing through their respective standard deviations. The correlation coefficient is 34%.

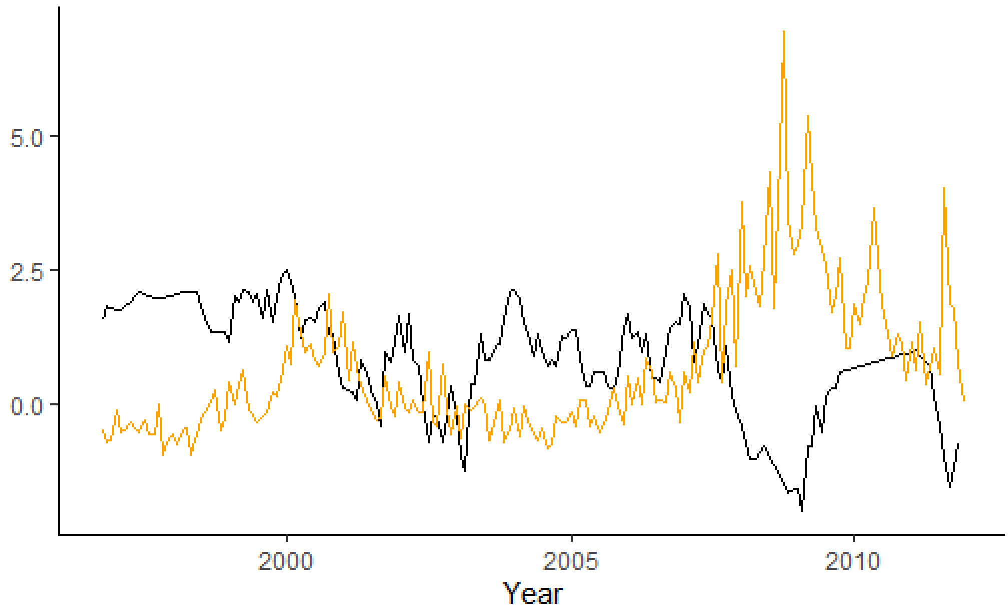

Figure 1.A.6: Gallup/UBS Investor Survey.

Displayed are the time series depicting the linearly detrended turnover in the US stock market (orange) and the time series depicting %bulls-%bears in the Gallup/UBS Investor Survey (black), in which investors are asked whether they are "optimistic" or "pessimistic" with regard to market returns. Both series are standardized by dividing through their respective standard deviations. The correlation coefficient is -59%.

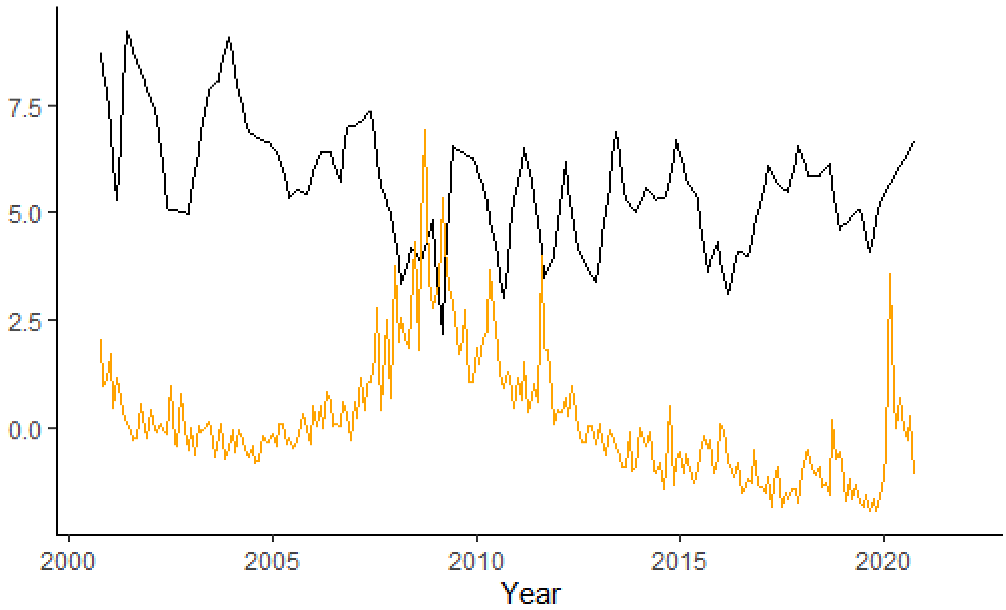

Figure 1.A.7: Duke CFO Survey: Mean One-Year Expectation for the S&P500 Return.

Displayed are the time series depicting the linearly detrended turnover in the US stock market (orange) and the time series depicting the average return expectation in the Duke CFO Survey (black), in which CFOs are asked about what they expect the return of the S&P 500 to be for the next year. Both series are standardized by dividing through their respective standard deviations. The correlation coefficient is -14%.

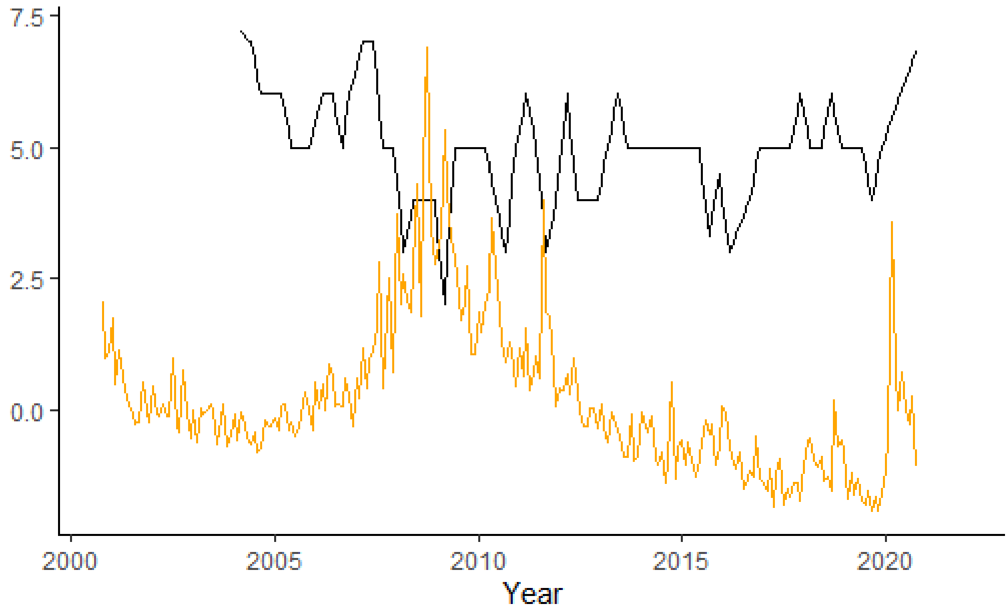

Figure 1.A.8: Duke CFO Survey: Median One-Year Expectation for the S&P500 Return.

Displayed are the time series depicting the linearly detrended turnover in the US stock market (orange) and the time series depicting the median return expectation in the Duke CFO Survey (black), in which CFOs are asked about what they expect the return of the S&P 500 to be for the next year. Both series are standardized by dividing through their respective standard deviations. The correlation coefficient is -30%.

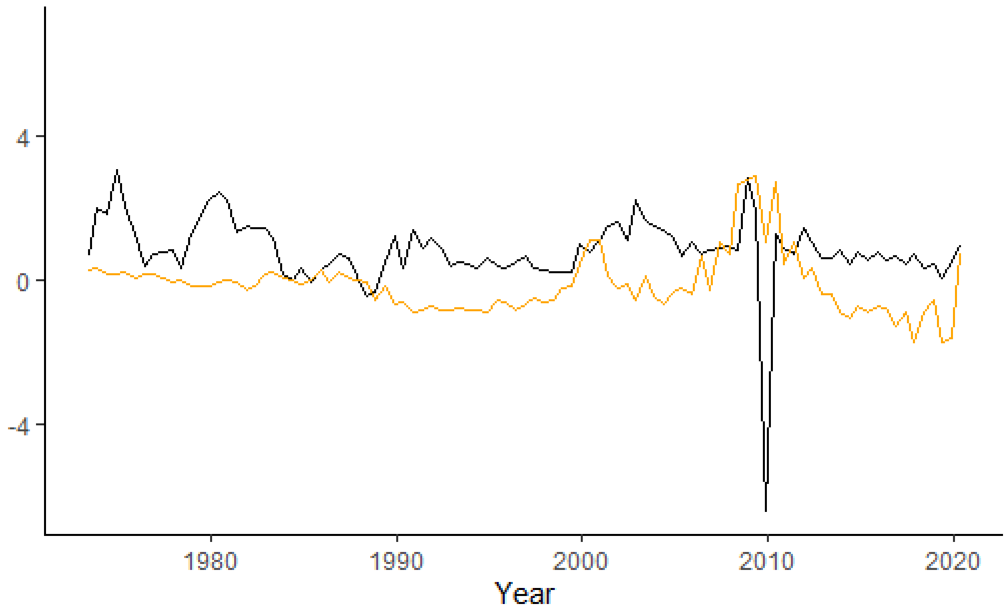

Figure 1.A.9: Livingston Survey, Mean Half-year Expectation for the S&P500 Return.

Displayed are the time series depicting the linearly detrended turnover in the US stock market (orange) and the time series depicting the mean return expectation in the Livingston Survey (black), in which economists are asked about what they expect the return of the S&P 500 to be for the next half-year. Both series are standardized by dividing through their respective standard deviations. The correlation coefficient is 12%.

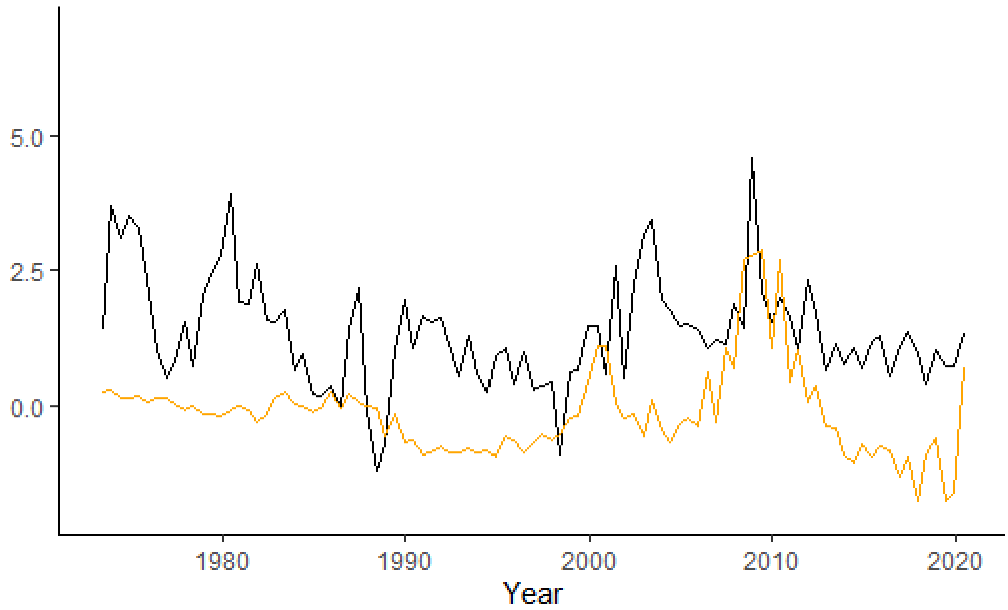

Figure 1.A.10: Livingston Survey, Median Half-year Expectation for the S&P500 Return.

Displayed are the time series depicting the linearly detrended turnover in the US stock market (orange) and the time series depicting the median return expectation in the Livingston Survey (black), in which economists are asked about what they expect the return of the S&P 500 to be for the next half-year. Both series are standardized by dividing through their respective standard deviations. The correlation coefficient is 33%.

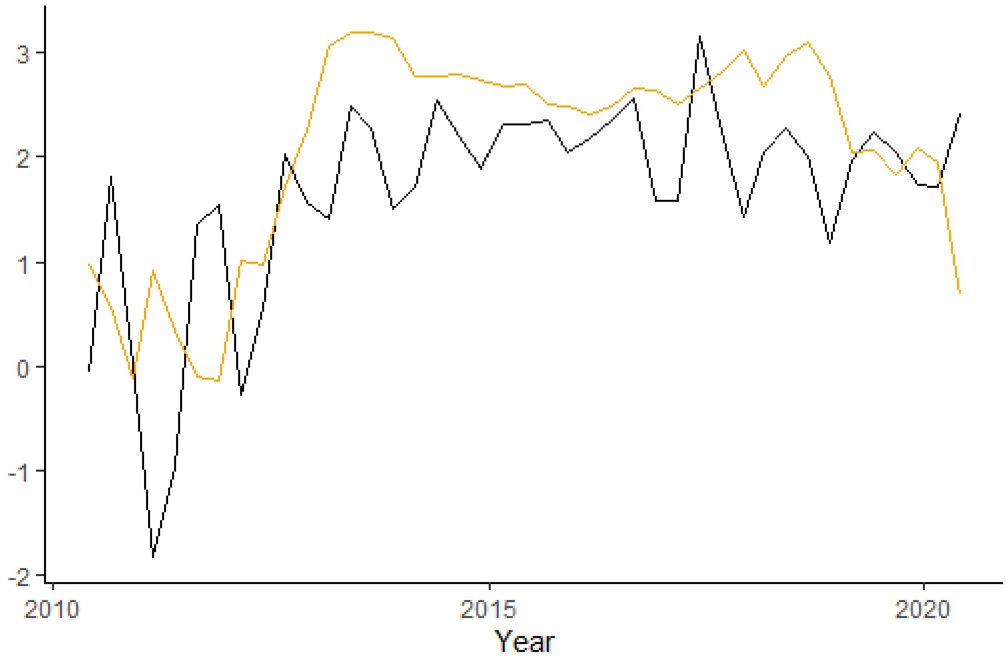

Figure 1.A.11: Zillow Year-over-Year Expectations (Black) and Home Turnover (Orange).

Displayed are the time series depicting the linearly detrended turnover in the US housing market (orange) and the time series depicting the average return expectation in the Zillow Home Price Expectation Survey (black), in which home owners are asked about what they expect the average price growth in the US housing market to be for the coming calendar year. Both series are standardized by dividing through their respective standard deviations. The correlation coefficient is 50%.

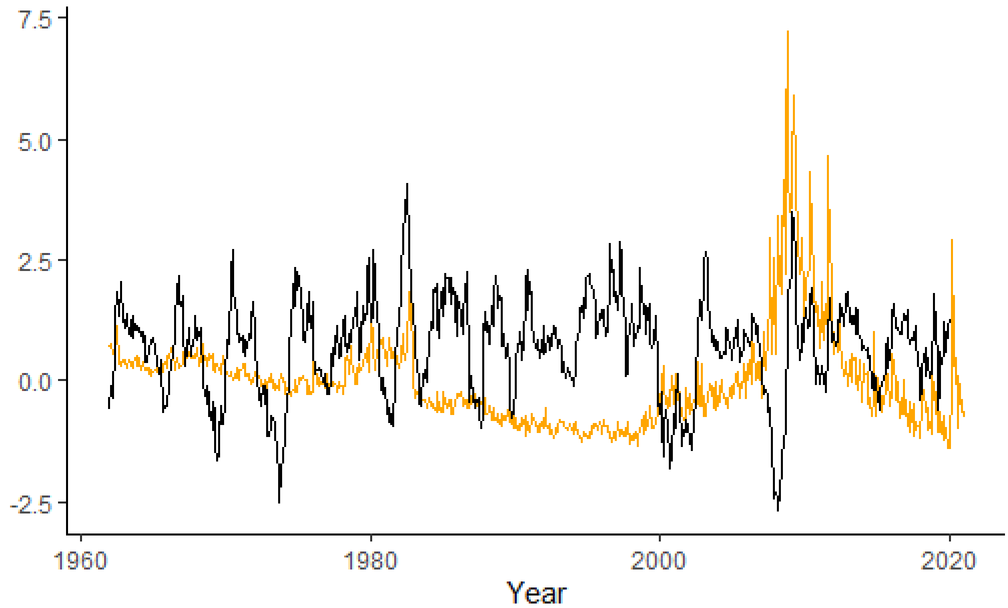

Figure 1.A.12: Subsequent Value-Weighted 12-Month Market Return on US Stocks.

Displayed are the time series depicting the linearly detrended turnover in the US stock market (orange) and the time series depicting the return in the 12 months after the turnover was measured (black). Value-weighted data from CRSP. Both series are standardized by dividing through their respective standard deviations. The correlation coefficient is 6%.

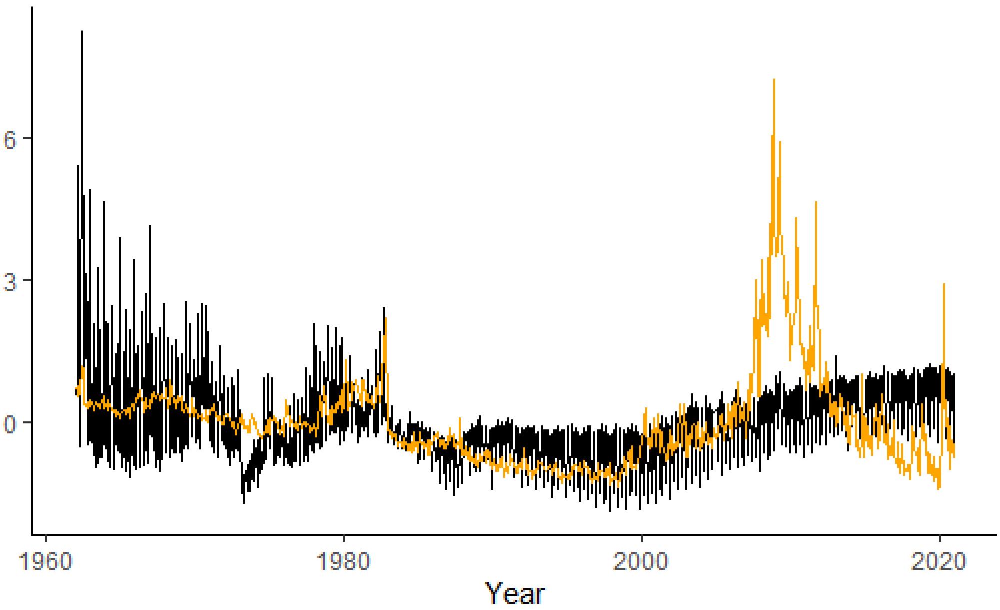

Figure 1.A.13: Value-Weighted Dividend-Price Ratio of US Stocks.

Displayed are the time series depicting the linearly detrended turnover in the US stock market (orange) and the time series depicting the dividend-price ratio (black). Value-weighted data from CRSP. Both series are standardized by dividing through their respective standard deviations. The correlation coefficient is 23%.

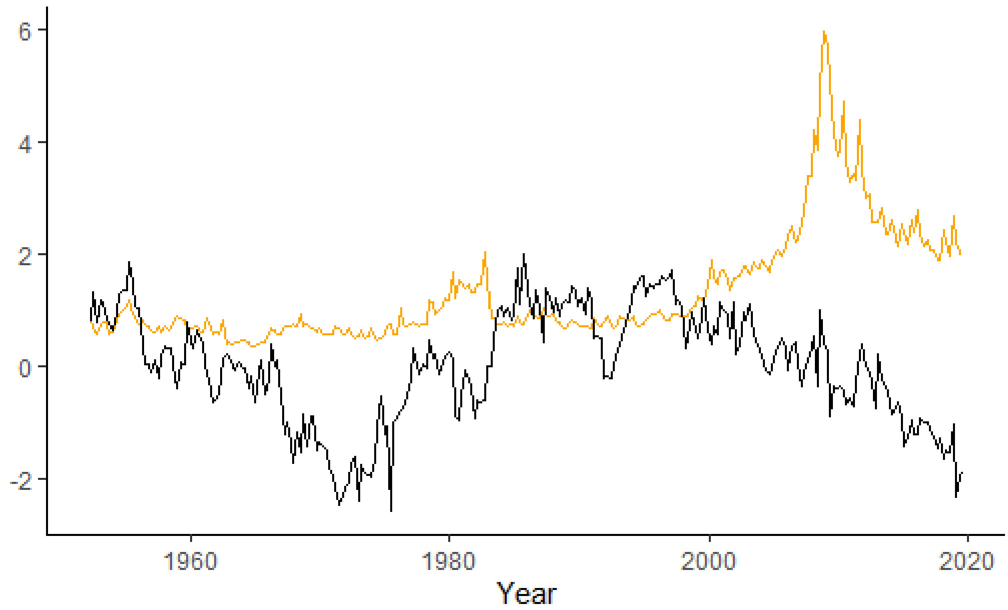

Figure 1.A.14: Consumption-Wealth Ratio.

Displayed are the time series depicting the linearly detrended turnover in the US stock market (orange) and the time series depicting the consumption-wealth ratio (black). Both series are standardized by dividing through their respective standard deviations. The correlation coefficient is -10%.

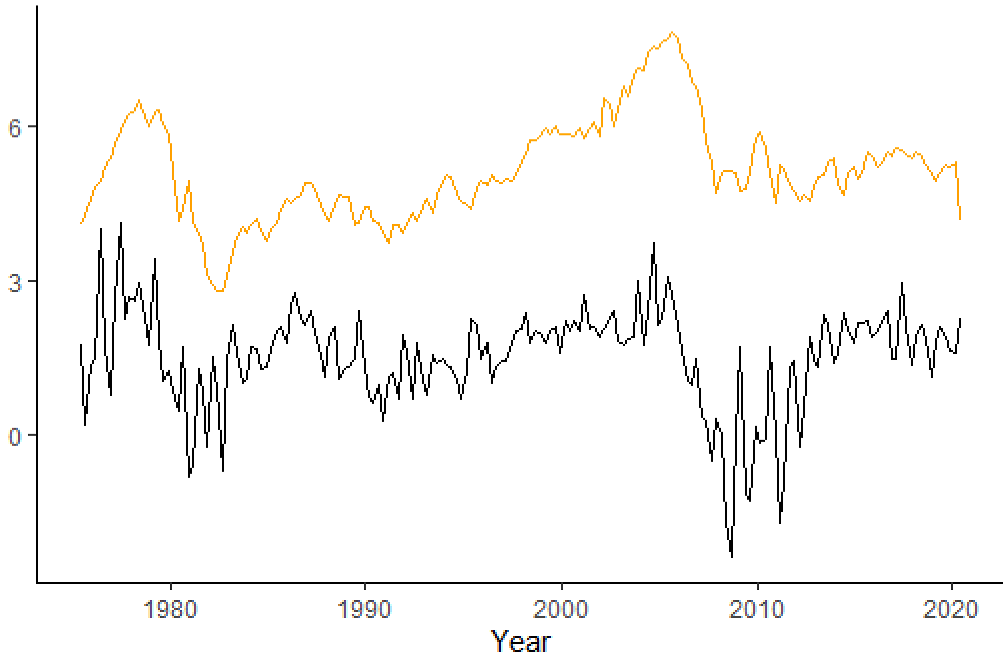

Figure 1.A.15: Market Return on Housing.

Displayed are the time series depicting the linearly detrended turnover in the US housing market (orange) and the time series depicting the market return on housing as computed from the FHFA index and Rent CPI (black). Both series are standardized by dividing through their respective standard deviations. The correlation coefficient is 43% at annual frequency. If we are calculating the returns over the subsequent four quarters at quarterly frequency, we find a correlation coefficient of 44%.

1.B Tables

Table 1.3: Expectation Surveys: Overview

Survey-based Measures	Description
Shiller Individual	Biannual from Apr 1999, monthly Jul 2001-Nov 2020. A random sample of high-income US Americans reports their expectations for the US stock market, including on by how much they think the Dow Jones will have increased in a year and whether they think the market is under- or overvalued. The measures used in the empirical analysis is the share of the sample who expect the Dow Jones Index to increase (One-year) and the share of the sample who think the market is undervalued (Valuation). Sample size is around 100 on average.
Shiller Institutional	Biannual from Oct 1989, monthly Jul 2001-Nov 2020. A random sample of US-based institutional investors is asked the same set of questions as in the “Shiller Individual” sample. Sample size is around 100 on average.
AAII Member Survey	Weekly, aggregated to monthly. Jun 1987-Nov 2020. Conducted via the AAII main publication (today the website). Respondents can choose whether they are "Bullish", "Neutral", or "Bearish". The measure used in the empirical analysis is %bulls-%bears.

Gallup/UBS Survey	Monthly (with notable gaps). Oct 1996 - Nov 2011. Telephone Interviews with 500-1000 participants who have invested >$10000 in the market. They are asked about their level of optimism about the stock market. The measure used in the empirical analysis is %optimists-%pessimists.
Duke CFO Survey	Quarterly from Oct 2000 - Dec 2020. A stable panel of CFOs from a diverse set of companies is asked how they think the return of the S&P 500 will be over the next year.
Livingston Survey	Biannual from 1952 - Jun 2020. Participants are economists in academia, industry, governments, and institutions. They are asked to forecast the level of the S&P 500 "in six months" and "in twelve months". As I can't identify when the response was given, and thus an expected return cannot be inferred, I use the growth of the estimate for the level "in six months" and "in twelve months" as a return expectation for months six to twelve.

Table 1.4: Correlation of Non-Detrended Stock Market Turnover with Survey-based Expectation Measures

Survey-based Measures	Correlation	Period + Frequency
Shiller One-Year Individual	-0.18*** [-0.3, -0.05]	Biannually from Apr 1999, monthly since Jul 2001.
Shiller One-Year Institutional	0.26*** [0.14, 0.37]	Biannually from Oct 1989, monthly since Jul 2001.
Shiller Valuation Individual	0.26*** [0.14, 0.37]	Biannually from Apr 1999, monthly since Jul 2001.
Shiller Valuation Institutional	0.36*** [0.25,0.46]	Biannually from Oct 1989, monthly since Jul 2001.
AAII Member Survey	0.32*** [0.23 0.41]	Weekly, aggregated to monthly Jun 1987- Nov 2020.
Gallup/UBS Survey	-0.64*** [-0.72 -0.55]	Monthly (with gaps) Oct 1996 - Nov 2011
Duke CFO Mean	-0.38*** [-0.56 -0.16]	Quarterly from Oct 2000 - Dec 2020
Duke CFO Median	-0.48** [-0.63 -0.23]	Quarterly from Mar 2004 - Dec 2020
Livingston Mean	-0.12 [-0.31 0.085]	Biannually from Jun 1973 - Jun 2020
Livingston Median	0.015 [-0.19 0.22]	Biannually from Jun 1973 - Jun 2020
	p-value:<1%***, <5%**, <10%*	All Shiller Survey Data up to Nov 2020

1.C Proofs

1.C.0.1 Proposition 1.1

Proof. Under the i.i.d assumption on β_t^i: $\bar{\text{p}}_t = P\left(\beta_t^i > \bar{\beta}_t\right)$ and $\underline{\text{p}}_t = P\left(\beta_t^i < \underline{\beta}_t\right)$. From Equation 1.1 and $N = 1/2$ follows

$$P\left(\beta_t^i > \bar{\beta}_t\right) = P\left(\beta_t^i < \underline{\beta}_t\right).$$

Both the owners and the non-owners will follow trading policies conditional on their respective draws of β_t^i.
Non-owner i buys if

$$\beta_t^i E_t\left[V^o\left(\beta_{t+1}^i, q_{t+1}\right) - V^{no}\left(\beta_{t+1}^i, q_{t+1}\right)\right] > (1+\kappa)\, q_t, \tag{1.9}$$

Similarly, owner i sells if

$$(1-\kappa)\, q_t > \beta_t^i E_t\left[V^o\left(\beta_{t+1}^i, q_{t+1}\right) - V^{no}\left(\beta_{t+1}^i, q_{t+1}\right)\right]. \tag{1.10}$$

Equations 1.9 and 1.10 define the individual thresholds for the discount factor for buying and selling the asset:

$$(1+\kappa)\, q_t = \bar{\beta}_t E_t\left[G\left(\beta_{t+1}^i\right)|\beta_t^i = \bar{\beta}_t\right], \tag{1.11}$$

and

$$(1-\kappa)\, q_t = \underline{\beta}_t E_t\left[G\left(\beta_{t+1}^i\right)|\beta_t^i = \underline{\beta}_t\right], \tag{1.12}$$

where

$$G\left(\beta_{t+1}^i\right) \equiv V^o\left(\beta_{t+1}^i, q_{t+1}\right) - V^{no}\left(\beta_{t+1}^i, q_{t+1}\right).$$

Note that, due to the i.i.d assumption on β_t^i the expected option value of holding the asset in $t+1$ is independent of the individual discount factor in t, and thus:

$$E_t\left[G\left(\beta_{t+1}^i\right)|\beta_t^i = \bar{\beta}_t\right] = E_t\left[G\left(\beta_{t+1}^i\right)|\beta_t^i = \underline{\beta}_t\right] = E_t\left[G\left(\beta_{t+1}^i\right)\right].$$

From the threshold equations 1.11 and 1.12 follow (by respectively adding and subtracting the two equations):

an asset pricing equation, which is

$$q_t = \frac{\bar{\beta}_t + \underline{\beta}_t}{2} E_t \left[G\left(\beta^i_{t+1}\right)\right], \tag{1.13}$$

as well as an equation implicitly determining turnover, which is

$$\bar{\beta}_t - \underline{\beta}_t = \frac{2\kappa q_t}{E_t G\left(\beta^i_{t+1}\right)}. \tag{1.14}$$

The equilibrium price q_t, must be such that the number of owners in the selling region equals the number of owners in the buying region:

$$(1-N)\,\bar{\mathrm{p}}_t = N\underline{\mathrm{p}}_t.$$

Under the stated assumptions that $N = 1/2$ and f_β is symmetric around β_t:

$$\underline{\mathrm{p}}_t = \bar{\mathrm{p}}_t \implies \frac{\bar{\beta}_t + \underline{\beta}_t}{2} = \beta_t \implies q_t = \beta_t E_t \left[G\left(\beta^i_{t+1}\right)\right].$$

This result can be used to understand G:

$$E_t G\left(\beta^i_{t+1}\right) = E_t \left[V^o\left(\beta^i_{t+1}, q_{t+1}\right) - V^{no}\left(\beta^i_{t+1}, q_{t+1}\right)\right]$$

$$= E_t \{ \bar{\mathrm{p}}_{t+1} \left(d_{t+1} + (1-\kappa)\, q_{t+1} + E^*_{t+1}\left[\beta^i_{t+1} | \beta^i_{t+1} < \underline{\beta}_{t+1}\right] E_{t+1}\left[V^{no}\left(\beta^i_{t+2}, q_{t+2}\right)\right]\right)$$

$$+ \left(1 - \bar{\mathrm{p}}_{t+1}\right) \left(d_{t+1} + E^*_{t+1}\left[\beta^i_{t+1} | \beta^i_{t+1} > \underline{\beta}_{t+1}\right] E_{t+1}\left[V^{o}\left(\beta^i_{t+2}, q_{t+2}\right)\right]\right)$$

$$- \bar{\mathrm{p}}_{t+1} \left(-(1+\kappa)\, q_{t+1} + E^*_{t+1}\left[\beta^i_{t+1} | \beta^i_{t+1} > \bar{\beta}_{t+1}\right] E_{t+1}\left[V^{o}\left(\beta^i_{t+2}, q_{t+2}\right)\right]\right)$$

$$- \left(1 - \bar{\mathrm{p}}_{t+1}\right) \left(E^*_{t+1}\left[\beta^i_{t+1} | \beta^i_{t+1} < \bar{\beta}_{t+1}\right] E_{t+1}\left[V^{no}\left(\beta^i_{t+2}, q_{t+2}\right)\right]\right) \},$$

which we can summarize as: [14]

$$E_t G\left(\beta^i_{t+1}\right) = E_t \left[d_{t+1} + 2\bar{\mathrm{p}}_{t+1} q_{t+1} + \left(1 - 2\bar{\mathrm{p}}_{t+1}\right) \beta_{t+1} E_{t+1} G\left(\beta^i_{t+2}\right)\right].$$

[14] Here, the operator E^*_{t+1} denotes the expected value for the individual discount factor, when the mean β_{t+1} and therefore the distribution f_β is known. I use $E^*_{t+1}\left[\beta^i_{t+1} | \beta^i_{t+1} > \bar{\beta}_{t+1}\right] - \beta_{t+1} = \beta_{t+1} - E^*_{t+1}\left[\beta^i_{t+1} | \beta^i_{t+1} < \underline{\beta}_{t+1}\right]$ and $E^*_{t+1}\left[\beta^i_{t+1} | \beta^i_{t+1} > \underline{\beta}_{t+1}\right] - \beta_{t+1} = \beta_{t+1} - E^*_{t+1}\left[\beta^i_{t+1} | \beta^i_{t+1} < \bar{\beta}_{t+1}\right]$, which again follows from the assumption of symmetry on f_β.

By rolling forward we find that $E_t G$ is just the expected value of next period dividend and price:

$$q_{t+1} = \beta_{t+1} E_{t+1} G\left(\beta^i_{t+2}\right),$$

$$\implies \quad E_t G\left(\beta^i_{t+1}\right) = E_t\left[d_{t+1} + q_{t+1}\right] . \tag{1.15}$$

Thus, we find a standard asset pricing equation: $q_t = \beta_t E_t\left[d_{t+1} + q_{t+1}\right]$, which is Equation 1.2.

From Equation 1.14, we know that:

$$\bar{\beta}_t - \underline{\beta}_t = \frac{2\kappa q_t}{E_t\left[d_{t+1} + q_{t+1}\right]}.$$

As $P\left(\beta^i_t > \bar{\beta}_t\right) = P\left(\beta^i_t < \underline{\beta}_t\right)$, and the distribution f_β is symmetric around β_t, it is implied that $\bar{\beta}_t - \beta_t = \beta_t - \underline{\beta}_t$ or $\bar{\beta}_t - \underline{\beta}_t = 2\left(\bar{\beta}_t - \beta_t\right)$.

Thus, $\bar{\beta}_t - \beta_t = \kappa q_t / E_t\left[d_{t+1} + q_{t+1}\right]$ or in terms of returns $\bar{\beta}_t - \beta_t = \kappa / E_t\left[R_{t+1}\right]$. Turnover in this model is identical to the probability for an asset owner to sell, which is the mass of owners with individual discount factors β^i_t larger than $\beta_t + \kappa / E_t\left[R_{t+1}\right]$. Thus turnover is shown, as posited in Equation 1.3, to be:

$$\bar{p}_{t+1} = 1 - F_\beta\left(\frac{1+\kappa}{E_t\left[R_{t+1}\right]}\right).$$

□

1.C.0.2 Proposition 1.2

Proof. In this slightly different set-up, the value functions need to be adapted. The value for individual i of being an owner of an asset in t is

$$\begin{aligned} V^o_t\left(\beta_t, q_t\right) = \quad & \max\{d_t + (1-\kappa)\, q_t + \beta_t E^i_t\left[V^{no}\left(\beta_{t+1}, q_{t+1}\right)\right] \\ & , \\ & d_t + \beta_t E^i_t\left[V^o\left(\beta_{t+1}, q_{t+1}\right)\right]\} \end{aligned}$$

while the value of not being an owner is accordingly

$$V_t^{no}(\beta_t, q_t) = \max\{ \quad -(1+\kappa)\, q_t + \beta_t E_t^i \left[V^o(\beta_{t+1}, q_{t+1})\right],$$

$$\beta_t^i E_t^i \left[V^{no}(\beta_{t+1}, q_{t+1})\right]\}$$

The thresholds for buying and selling, $\bar{E}_t$ and $\underline{\mathrm{E}}_t$ are determined by

$$(1-\kappa)\, q_t = \beta_t \underline{\mathcal{E}}_t$$

and

$$(1+\kappa)\, q_t = \beta_t \bar{\mathcal{E}}_t.$$

These equations imply, following the same steps as in the proof of Proposition 1.1:

$$q_t = \frac{\beta_t}{2}\left(\bar{\mathcal{E}}_t + \underline{\mathcal{E}}_t\right) \tag{1.16}$$

and

$$\beta_t\left(\bar{\mathcal{E}}_t - \underline{\mathcal{E}}_t\right) = 2\kappa q_t. \tag{1.17}$$

For the market to clear, it must be true that

$$(1-N)\, P\left(\mathcal{E}_t^i > \bar{\mathcal{E}}_t\right) = N P\left(\mathcal{E}_t^i < \underline{\mathcal{E}}_t\right),$$

due to the i.i.d. assumptions on $\mathcal{E}_t^i / q_t$. With $N = 1/2$, it must be that $P\left(\mathcal{E}_t^i > \bar{\mathcal{E}}_t\right) \equiv \bar{p}_t = \underline{\mathrm{p}}_t \equiv P\left(\mathcal{E}_t^i < \underline{\mathcal{E}}_t\right)$. As the distribution $f_{\frac{\mathcal{E}}{q}}$ is symmetric, we know that

$$\frac{\bar{\mathcal{E}}_t + \underline{\mathcal{E}}_t}{2} = \mathcal{E}_t.$$

Furthermore,

$$E_t\left[V^o\left(\beta_{t+1}^i, q_{t+1}\right) - V^{no}\left(\beta_{t+1}^i, q_{t+1}\right)\right]$$

$$= E_t\{\bar{\mathrm{p}}_{t+1}\left(d_{t+1} + (1-\kappa)\, q_{t+1} + \beta_{t+1} E_{t+1}^*\left[E_{t+1}^i\left[V^{no}(\beta_{t+2}, q_{t+2})\right] | \mathcal{E}_{t+1}^i < \underline{\mathcal{E}}_{t+1}\right]\right)$$

$$+ \left(1 - \bar{\mathrm{p}}_{t+1}\right)\left(d_{t+1} + \beta_{t+1} E_{t+1}^*\left[E_{t+1}^i\left[V^o(\beta_{t+2}, q_{t+2})\right] | \mathcal{E}_{t+1}^i > \underline{\mathcal{E}}_{t+1}\right]\right)$$

$$-\bar{\mathrm{p}}_{t+1}\left(-(1+\kappa)\, q_{t+1} + \beta_{t+1} E_{t+1}^*\left[E_{t+1}^i\left[V^o(\beta_{t+2}, q_{t+2})\right] | \mathcal{E}_{t+1}^i > \bar{\mathcal{E}}_{t+1}\right]\right)$$

$$-\left(1 - \bar{\mathrm{p}}_{t+1}\right)\left(\beta_{t+1} E_{t+1}^*\left[E_{t+1}^i\left[V^{no}(\beta_{t+2}, q_{t+2})\right] | \mathcal{E}_{t+1}^i < \bar{\mathcal{E}}_{t+1}\right]\right)\}.$$

We can summarize this expression as:[15]

$$E_t\left[V^o\left(\beta_{t+1}, q_{t+1}\right) - V^{no}\left(\beta_{t+1}, q_{t+1}\right)\right] = E_t\left[d_{t+1} + 2\bar{\mathrm{p}}_{t+1} q_{t+1} + \left(1 - 2\bar{\mathrm{p}}_{t+1}\right)\beta_{t+1}\mathcal{E}_{t+1}\right].$$

By rolling forward we find that $E_t G$ is just the expected value of next period dividend and price:

$$q_{t+1} = \beta_{t+1}\mathcal{E}_{t+1},$$

$$\Longrightarrow \quad E_t\left[V^o\left(\beta^i_{t+1}, q_{t+1}\right) - V^{no}\left(\beta^i_{t+1}, q_{t+1}\right)\right] = E_t\left[d_{t+1} + q_{t+1}\right] \; = \mathcal{E}_t. \tag{1.18}$$

Thus, we again find a standard asset pricing equation

$$q_t = \beta_t E_t\left[d_{t+1} + q_{t+1}\right].$$

Plugging the asset pricing equation into Equation 1.17, we find

$$\beta_t\left(\bar{\mathcal{E}}_t - \underline{\mathcal{E}}_t\right) = 2\kappa\beta_t\mathcal{E}_t.$$

Due to symmetry:

$$\bar{\mathcal{E}}_t - \mathcal{E}_t = \mathcal{E}_t - \underline{\mathcal{E}}_t,$$

which implies

$$\bar{\mathcal{E}}_t = (1+\kappa)\,\mathcal{E}_t.$$

Turnover is

$$\begin{aligned}\bar{p}_t = P\left(\mathcal{E}^i_t > \bar{\mathcal{E}}_t\right) P\left(\frac{\mathcal{E}^i_t}{q_t} > \frac{\bar{\mathcal{E}}_t}{q_t}\right) &= 1 - F_{\frac{E}{q}}\left((1+\kappa)\frac{E_t\left[d_{t+1} + q_{t+1}\right]}{q_t}\right) \\ &= 1 - F_{\frac{\mathcal{E}}{q}}\left((1+\kappa) E_t\left[R_{t+1}\right]\right).\end{aligned}$$

$\square$

[15]Here, the operator E^*_{t+1} analogously to the previous proof, denotes the expected value for the individual future expectation, when the mean $\mathcal{E}_{t+1}$, and therefore the distribution $f_{\frac{\mathcal{E}}{q}}$, is known. I use $E^*_{t+1}\left[\mathcal{E}^i_{t+1} | \mathcal{E}^i_{t+1} > \bar{\mathcal{E}}_{t+1}\right] - \mathcal{E}_{t+1} = \mathcal{E}_{t+1} - E^*_{t+1}\left[\mathcal{E}^i_{t+1} | \mathcal{E}^i_{t+1} < \underline{\mathcal{E}}_{t+1}\right]$ and $E^*_{t+1}\left[\mathcal{E}^i_{t+1} | \mathcal{E}^i_{t+1} > \underline{\mathcal{E}}_{t+1}\right] - \mathcal{E}_{t+1} = \mathcal{E}_{t+1} - E^*_{t+1}\left[\mathcal{E}^i_{t+1} | \mathcal{E}^i_{t+1} < \bar{\mathcal{E}}_{t+1}\right]$, which again follows from the assumption of symmetry on $f_{\frac{\mathcal{E}}{q}}$.

Chapter 2

MONETARY POLICY AND CORPORATE DEBT MATURITY

2.1 Introduction

High corporate indebtedness is a major source of vulnerability for many Advanced Economies (IMF 2019a,b; Kaplan 2019).[1] In this respect, the maturity structure is a key feature of corporate debt, influencing firms' reaction to both real and financial shocks (Almeida et al. 2009; Duchin, Ozbas, and Sensoy 2010; He and Xiong 2012b; Kalemli-Ozcan, Laeven, and Moreno 2018; Jungherr and Schott 2020b; Chen, Xu, and Yang 2020). The ongoing Covid-19 crisis is not an exception: in fact, a higher share of maturing obligations has been found to depress non-financial firms' stock returns Fahlenbrach, Rageth, and Stulz (2020) and to limit their access to capital markets Halling, Yu, and Zechner (2020) during the most acute phases of the pandemic. Understanding whether and how monetary policy affects the maturity structure of corporate debt is therefore of utmost importance, as it allows to gauge a potentially relevant bearing of central banks' policy on firms' risk. However, up to our knowledge, existing studies do not systematically explore this question.

[1] See Giroud and Mueller (2017, 2018) on how high firms' leverage boosts business-cycle fluctuations.

We fill the gap by investigating the influence of the interest rate policy by the FED on the maturity structure of the US corporate debt sector. Our focus is over the period 1990-2017 and our empirical exercise exploits: i) different measures of - endogenous and exogenous - variation of the Effective Fed Funds Rate (EFFR); ii) a large variety of time-series, firm-level and security-level datasets.

We find that, at the aggregate level, a reduction of the EFFR *lengthens* corporate debt maturity (i.e., increases the share of total debt with maturity above 1-year). The effect is both statistically and economically significant. A 25 basis points (b.p.) descent in the EFFR triggers a persistent jump in the share of long-term (LT) debt, amounting to roughly 0.42 percentage points (p.p) one year after the shock. For comparison, the average quarterly growth rate of the share of LT debt equals 0.15 p.p.. Next, we look at quarterly balance sheets of US listed firms and find that very large companies - namely those in the top-quartile of their respective industry-wide asset-size distribution - are responsible for the observed aggregate patterns, whereas smaller firms do not adjust.

We explain such findings through a parsimonious model combining financial frictions due to moral hazard Holmström and Tirole (1998, 2000) and short-termist, yield-oriented investors Hanson and Stein (2015), who care about current portfolio yield on top of expected returns and rebalance their portfolios toward LT debt securities when the policy rate descents. The demand shift decreases bond yields, but only large and unconstrained companies can take advantage by issuing LT bonds.

Our model delivers predictions aligned to our aggregate and cross-sectional evidence, and we empirically test its mechanism. We find that yield-oriented corporate bonds mutual funds increase their holdings of corporate bonds (as compared to other funds) when the policy rate goes down, while also tilting their holdings towards longer-term debt securities. Moreover, large companies' likelihood of issuing LT bonds jumps more strongly and the coupon rate reacts with a stronger decline. This suggests that relative fluctuations in bonds issuance by large companies are demand-driven.

Our main contribution is to provide systematic evidence on the relation between monetary policy and the maturity structure of the debt of US non-financial corporations (NFCs). Other papers condition the impact of monetary policy shocks on the ex-ante heterogeneity in firm debt maturity (Ippolito, Ozdagli, and Perez-Orive 2018; Jungherr and Schott 2020a). Differently, we document that the maturity structure of debt endogenously responds to monetary policy shocks. Furthermore, our paper adds to a novel set of studies documenting the role of yield-oriented investors in the transmission of monetary policy (see, e.g., Hanson and Stein 2015; Di Maggio and Kacperczyk 2017; Daniel, Garlappi, and Xiao forthcoming; Lian, Ma, and Wang 2019). We innovate by linking the relation between monetary policy shocks and bond issuance to yield-oriented investors and showing the implications for firms' debt maturity structure.

The rest of this introduction is divided into two parts. First, we provide a detailed preview of the paper. Second, we discuss more thoroughly the related literature and contrast it with our paper.

2.1.1 Detailed Preview of the Paper

We investigate two main research questions. First, we ask whether monetary policy has any impact on the maturity structure of corporate debt. Second, we verify eventual cross-sectional differences across companies in such relation. Our focus rests on *conventional* interest rate policy by the FED over the period 1990-2017 and we limit our attention to the US corporate sector.

We use three different variables for capturing the FED interest rate policy. The baseline exercises use the raw quarterly variation of the EFFR, an endogenous measure of changes in the monetary policy stance displaying large persistence over tightening and loosening cycles (Adrian, Estrella, and Shin 2010). Still, this measure has a direct "real-world" impact on firms' financing cost, and hence we test its influence on the debt maturity structure while controlling for other correlated macroeconomic variables such as GDP growth and inflation rate. Importantly, we also test the robustness of our results to two alternative exogenous measures of (high-frequency) interest rate shocks, borrowed from Gürkaynak, Sack, and Swanson 2005 and Jarocinski and Karadi 2020. Our remaining data come from various sources. To start with, the time-series analysis of corporate debt maturity is based on quarterly data from FED Flows of Funds. Following, among others, Greenwood, Hanson, and Stein (2010), we build the share of LT debt as the ratio between corporate debt with maturity above 1-year and total corporate debt. At the firm-level, we apply an identical measure, retrieved from Compustat quarterly financial data of US listed companies. To investigate our model-based mechanism, we get data on the universe of LT-bond issuance from Mergent FISD and access information on corporate bond mutual funds' holdings from the CRSP Survivor-Bias-Free US

Mutual Fund dataset.

We study the aggregate evolution of LT debt by looking at the change in the LT debt share from the Flows of Funds, which mean equals 0.15 p.p. on a quarterly basis. We employ local projections (Jordà 2005) in a model augmented with other lagged macroeconomic controls. We find that a reduction of the EFFR *lengthens* corporate debt maturity, i.e., it expands the share of debt with maturity above 1-year. The effect is both statistically and economically significant. A 25 basis points (b.p.) descent in the EFFR triggers a persistent jump in the share of LT debt, amounting to roughly 0.42 p.p. one year after the shock. The effect peaks up three years after the shock. Results are robust to using exogenous monetary policy shocks.[2]

Next, we test whether such effect is heterogeneously distributed across companies. For this purpose, we use quarterly balance sheet data for US listed firms from Compustat. The key layer of heterogeneity is firm-size (approximated through total assets), inversely related to the intensity of financial constraints. In particular, we sort firms' in quartiles of the (lagged) industry-level (3-digit SIC classification) size-distribution. We again apply local projections to analyze the dynamic response of the share of long-term debt, in a setting akin to Ottonello and Winberry (forthcoming) and Jeenas (2018). Our interest falls on the interaction between the change of the EFFR and a dummy identifying large companies, i.e. those in the top-quartile of their industry level size-distribution. This allows us to saturate the model with firm and industry*year:quarter fixed effects, controlling for firm-level time-invariant heterogeneity and industry-wide time-varying shocks, respectively. Also, we horse-race this channel against other relevant firm balance-sheet items (fully interacted with the variation of the policy rate). Our results show that *large* companies adjust *more*, that is, they increase (decrease) their share of LT debt relatively more in response to a descent (jump) in the policy rate. Moreover, running separate regressions for companies in different quartiles of the asset-size distribution, we find that smaller companies' debt maturity is generally not responsive to monetary policy. Using exogenous interest rate shocks produces comparable results.

Interestingly, it is not easy to rationalize our results according to standard models of monetary policy transmission to firms, which, in general, do not consider the debt

[2]The effect has similar magnitude to that described for the EFFR variations, though it is less persistent - a difference we impute to the significant autocorrelation characterizing the raw EFFR series.

maturity structure.[3] On the other hand, the corporate finance literature does not directly focus on the policy rate, but rather on the term-spread, delivering counterfactual predictions relative to our findings.[4] Hence, we propose a theory that can account for our aggregate and cross-sectional empirical facts.

We augment a standard model with short and long-term debt and financing frictions due to moral-hazard as in Holmström and Tirole (1998, 2000) with the presence of short-termist, yield-seeking investors as in Hanson and Stein (2015). Such investors are assumed to take long-short positions and care about current portfolio yield and not just expected returns. This modeling assumption reflects short-termist incentives of important classes of investors which, for instance, report each quarter to the stock market and therefore care about current yields on top of total expected returns. As a result, in reaction to a policy rate cut, they rebalance their portfolios toward longer-term debt in an effort to keep their portfolio yield up, ultimately creating buying pressure on the price for LT debt. The boost in demand for LT debt is accommodated by larger NFCs, for which borrowing constraints are not binding.[5]

The model delivers predictions in line with both time-series and cross-sectional evidence and dovetails nicely with both large corporations' and investors' narrative on the link between interest rate policy and firms' debt maturity choices.[6] Nonetheless, we

[3]A classical literature on the credit channel Bernanke and Gertler (1995) of monetary policy exploits either frictions at the level of the company (e.g. Gertler and Bernanke 1989; Bernanke, Gertler, and Gilchrist 1999; Iacoviello 2005; Christensen and Dib 2008; Christiano, Motto, and Rostagno 2014) or at the level of the firms' lenders, typically banks (e.g. Kashyap and Stein 1994; Adrian and Shin 2010; Gertler and Karadi 2011; Borio and Zhu 2012; Dell'Ariccia, Laeven, and Marquez 2014). Under both paradigms, an interest rate cut relaxes credit standards, typically proxied by loan volume and rate, with greater relative benefits for small and constrained companies. Heuristically, such models would likewise predict a greater expansion of debt maturity for smaller companies, and hence may contradict our cross-sectional evidence. More novel contributions - including e.g. Ottonello and Winberry (forthcoming) and Ozdagli (2018) - highlight how larger and less constrained NFCs may respond more to monetary policy shocks but neglect debt maturity.

[4]The reasoning goes as follows: a policy rate cut widens the term-spread (a stable relation documented by, e.g., Adrian and Shin 2010 and which we also verify in our sample) and hence increases the relative convenience of short-term debt issuance.

[5]Importantly, our model leaves room for a standard "balance-sheet channel" of monetary policy, whereby smaller firms benefit more from a relaxation of the monetary conditions. In fact, an interest rate cut relaxes moral-hazard frictions as the value of collateral goes up. Eventually, we prove the existence of equilibria where yield-seeking motives dominate the balance-sheet channel. In this context, the adjustment of the share of LT debt is carried out by the large, unconstrained companies, in line with our evidence.

[6]For instance, a recent article from the Financial Times - commenting the ultra-low interest rate en-

conclude by bringing the model mechanism to the data. First, we check that large firms increase the frequency of issuance of LT bonds when the policy rate decreases (relatively to small ones). The impact is large: a 25 b.p. decline in EFFR implies an additional 28 b.p. jump for large companies in the probability of issuing new bonds at impact, i.e., 4% of the average likelihood of issuing bonds. Moreover, on the intensive margin, the coupon rate at issuance declines substantially more for large companies. Hence, following an interest rate cut, large firms issue more LT bonds and at lower rates, suggesting that their (relative) reaction is demand driven. Finally, we also test that these buying pressures are associated to portfolio rebalancing by yield-seeking investors. We split corporate bond mutual funds (CBMF) into investment-grade (IG) and high-yield (HY); following Choi and Kronlund (2018), the latter group follows a more yield-oriented investment strategy. In line with the theory, we find a bigger increase in corporate bonds holdings and in portfolio's average maturity for HY-funds after a policy rate decline. The usual 25 b.p. decrease in EFFR prompts a 6 p.p. marginal jump in corporate bonds' holdings for HY-funds and a lengthening in the maturity of held debt-securities by 2 p.p..

2.1.2 Contribution to the Literature

Our paper contributes to several strands of literature. To start with, our study is the first - up to our knowledge - to provide a systematic analysis of the relation between monetary policy and the maturity structure of corporate debt.[7] Few other papers exploit the ex-ante heterogeneity in firms' debt maturity to explain the heterogeneous real effects of monetary policy shocks across firms (Ippolito, Ozdagli, and Perez-Orive 2018; Jungherr and Schott 2020b). However, they do not endogenize the response of debt maturity itself to the monetary policy shocks. We do not only provide empirical evidence that such endogenous response is economically meaningful, but also derive a model highlight-

vironment during the Covid-19 pandemic - reports that "companies across the US are taking advantage of low borrowing costs to extend the maturity of their debt, selling longer and longer dated bonds to investors starved of yield. (...) As yields have tumbled and investor appetite for debt has remained unsated, corporate treasurers are now making more opportunistic moves." The article also features related comments from CFO of large corporations such as AT&T and from portfolio managers at different investment firms. The article is available at this link.

[7] Gomes, Jermann, and Schmid (2016) show that nominal long-term debt can generate persistent responses to unanticipated inflation changes (linked to monetary policy shocks) through a debt-overhang mechanism. However, their model does not feature an endogenous firms' optimal debt maturity structure.

ing a mechanism (operating through short-termist, yield-oriented investors), which can eventually be tested in the data. In a closedly related paper to ours, Foley-Fisher, Ramcharan, and Yu (2016) find that a specific unconventional policy by the FED, namely the rebalancing of its portfolio towards longer-term Treasuries, implied a lengthening of the maturity of bond issuance by companies through a gap-filling mechanism (Greenwood, Hanson, and Stein 2010). We differ in two dimensions: first, we look at conventional and regular interest rate shocks rather than at a one-time shock to the maturity profile of the FED's portfolio; second, we lever a different mechanism related to yield-oriented investors.

By doing so, we connect to a growing literature stressing the importance of reach-for-yield in financial markets (Becker and Ivashina 2015) for the reallocation of investment across securities in reaction to monetary policy shocks.[8] We lever a mechanism, which, following Hanson and Stein (2015), sees yield-seeking investors tilting their portfolios towards LT securities after an interest rate loosening (a mechanism validated empirically with evidence on corporate bonds mutual funds, shown to reach-for-yield by Choi and Kronlund (2018)). Our contribution to this literature is to link yield-oriented investors' reaction to monetary policy to the issuance of LT bonds and, ultimately, to the evolution of the maturity structure of corporate debt.

Few novel papers look at financial channels for monetary policy different from a standard credit channel Bernanke and Gertler (1995), i.e., other mechanisms than bank intermediation. Among others, Foley-Fisher, Ramcharan, and Yu (2016), Grosse-Rueschkamp, Steffen, and Streitz (2019), and Giambona et al. (2020) investigate adjustments in the bond market in response to unconventional monetary policy. Similarly to us, Darmouni, Giesecke, and Rodnyansky (2020) exploit frictions in the Eurozone bond markets and find that access to the bond market is linked to greater firms' sensitivity to interest rate shocks. Ottonello and Winberry (forthcoming) and Jeenas (2018) highlight the importance of default risk and liquid assets for the transmission of interest rate shocks to firms. We innovate by focusing on debt maturity and highlighting a bond-

[8]For instance, Di Maggio and Kacperczyk (2017) demonstrate that the FED zero rate policy prompted exit in the money-market funds industry and greater risk-taking by surviving funds, with implications for firms borrowing from such institutions. Bubeck, Maddaloni, and Peydró (forthcoming) find that negative rates in the Euro Area are associated to enhanced risk-taking in banks' securities portfolio, whereas Lian, Ma, and Wang (2019) and Daniel, Garlappi, and Xiao forthcoming obtain homologous findings for individual investors.

channel of conventional interest rate policy connected to reach-for-yield in financial markets.

We contribute as well to the literature on the determinants of the maturity structure of corporate debt Barclay and Smith Jr (1995), Berger et al. (2005), Faulkender (2005), Greenwood, Hanson, and Stein (2010), and Badoer and James (2016), generally placing little emphasis on the policy rate but rather focusing on the term-spread.[9] On the other hand, those works discuss extensively the role of financing frictions, but no evidence exists on their interaction with the prevailing monetary policy stance.[10]

The rest of the paper is organized as follows. In Section 2.2, we describe the data. In Section 2.3, we present the baseline empirical findings. To explain them, we elaborate a model in Section 2.4, which mechanism is tested in Section 2.5. Section 2.6 briefly concludes.

2.2 Data

Our empirical analysis covers the period from 1990Q1 to 2016Q4. We employ several datasets, that we describe separately in this section according to the unit-level of analysis.

2.2.1 Time-series Data

For the time-series analysis of the LT debt share, we use the Federal Reserve Flow of Funds (FoF), tracking financial flows throughout the U.S. economy. We use quarterly data from the credit market liabilities of the non-farm, non-financial, corporate business sector. Our focus rests on the share of total corporate debt with maturity above 1-year, which we label as the share of LT-debt. Following Greenwood, Hanson, and Stein (2010), we define short-term debt as the sum of commercial paper and loans with no longer maturity than 1 year (proxied by adding up the FoF entries "other loans and advances" and "bank loans not elsewhere classified"). On the other hand, long-term debt is given by the sum of corporate bonds, mortgages and industrial revenue bonds.

[9] A notable exception is Baker, Greenwood, and Wurgler (2003), showing that the real short-term rate, among other variables, covaries with the corporate LT-debt share.

[10] Relatedly, Poeschl (2017), Xu (2018) and Mian and Santos (2018) ask how cyclical factors impact debt refinancing policy and maturity: our paper differs as it looks specifically at monetary policy.

The resulting series, given by the fraction of LT debt over the sum of short and LT debt, is depicted in Figure 2.A.1. The black line, referring to such variable in levels, displays a generally increasing trend over the period of interest, with the share of LT-debt increasing from roughly 55% to 70%. Throughout the paper, we look at the impact of shocks to the policy rate on the dynamics of the LT debt share, thereby computing its growth rate over different horizons. In Figure 2.A.1, the solid grey line shows the evolution of the quarterly growth rate, that is labelled as $\Delta LT - Debt_t$ in Table 2.1 and which mean equals 0.15 p.p.. We also report summary statistics for the cumulative growth rate of the LT debt share over longer horizons, used for pinning down impulse response functions through local projections. In general, the variable $\Delta LT - Debt_{t+j}$ is computed as the difference between the LT debt share as of year-quarter $t+j$ and $t-1$, for $j = 0, 1, 2, ..., 20$. For brevity, we show summary statistics only for up to 1-year growth of the LT-debt share. Evidently, both the mean and the volatility of the growth rate increase along with the length of the horizon over which they are computed.

Our first proxy of changes in the policy rate is the quarterly variation in the effective federal funds rate (EFFR), $\Delta EFFR_t$, gathered from FRED. Clearly, $\Delta EFFR_t$ reflects the evolution of business cycle conditions and the connected endogenous response from the FED. Nonetheless, changes in the EFFR have a "real-world" influence on the firms financing costs. Hence, we first show baseline results based on such raw proxy and next verify the robustness of our findings to employing alternative conventional proxies for exogenous monetary (interest rate) policy shocks. In detail, we borrow data from Gürkaynak, Sack, and Swanson (2005), who builds a popular measure of interest rate surprises based on the $\%$ change in FED Funds Futures rate in 30-minute windows around the policy announcement. Next, we additionally retrieve the Jarocinski and Karadi (2020)'s series of "pure" interest rate surprises, i.e. taking out an informational component - attributed to the provision of private FED information on the state of the economy to private agents through the policy announcement - from the simple variation in the FED Funds Futures rate.

We depict the three series in Figure 2.A.2. The post-2009 period is characterized by lower variation in the interest rate policy, associated to the implementation of a zero-rate by the FED in the aftermath of the Great Financial Crisis. For this reason, whenever possible, we check that our results survive if we exclude the period from 2009 onward. Moreover, while the three series display a large extent of correlation, both the exogenous variables are in general an order of magnitude smaller than $\Delta EFFR_t$. Indeed, while a 1 s.d. change in $\Delta EFFR_t$ equals 45 basis points (b.p.), a 1 s.d. change in the Gürkaynak, Sack, and Swanson (2005) and Jarocinski and Karadi (2020) shocks amounts to 10 and 8 b.p., respectively (see Table 2.1).

Finally, we also collect several other FRED macro-economic indicators that we use as controls. In particular, the average annual GDP growth and inflation rates equal 2.46 p.p. and 2.5 p.p., respectively. Moreover, 10% of the year-quarters in our sample are characterized by a recession, as signalled by the dummy Rec_{t-1}. The mean quarterly growth rate for the term-spread and the corporate spread are smaller than 1 b.p., though both variables display a large extent of variability (their s.d. amount to 49 b.p. and 25 b.p., respectively).[11] From Thomson Reuters Datastream, we also download information on the share of Treasuries with maturity above 20-years and compute its quarterly growth rate ($\Delta LT - Treas_{t-1}$).[12]

2.2.2 Firm-level Data

Our primary source for firm-level data is Compustat, a well known database comprehending balance sheet information on the universe of US listed companies. Using Compustat entails pros and cons. On the positive side, it provides balance sheet information on a quarterly basis, whereas most of other large firm-level datasets contain annual balance sheet only. The relatively higher frequency is desirable in that it aligns better to the frequency of the monetary policy revisions by the FED. On the other hand, the information on debt is rather limited. In fact, we can only distinguish the fraction of total debt with maturity above 1-year - in line with our macroeconomic data from FoF - without further data neither on the maturity profile of existing liabilities nor on the relative weight of bank vs bond financing.

Our sample includes 12,655 companies. Once again, we are mainly interested in

[11] The term-spread is defined as the difference between the yield on the 10-year and 3-month benchmark US sovereign bond. The corporate spread reflects risk prempium in the corporate sector and is computed as the difference between the Moody's BAA and AAA Seasoned Corporate Bond Yield.

[12] According to the gap-filling theory (Greenwood, Hanson, and Stein 2010), the share of LT-debt issued by the corporate sector depends negatively on the share of LT-debt issued by the government. Moreover, Badoer and James (2016) show that corporate debt issuance is especially sensitive to variations in very long-term Treasuries issuance. Hence, controlling for changes in the share of government debt with maturity above 20 years should alleviate concerns that our results are driven by a gap-filling mechanism driven by government debt.

the variation over time of the share of debt with maturity above 1-year. The variable $\Delta LT - Debt_{f,t+j}$ indeed represents the variation in firm f's LT Debt share from year-quarter $t-1$ to $t+j$. In Table 2.1 we report summary statistics for $j = 0, 1, ..., 4$. Across the different horizons, the distribution is centered around 0, as suggested by the median value. Nonetheless, the extent of heterogeneity is remarkable (see the high s.d., increasing along the number of quarters of computation of the cumulative growth rate).

An important variable throughout our analysis is firm's asset size, our preferred proxy for financing constraints, i.e. of access to bond financing. There are large differences in firms' asset size (expressed in logs of 1990q1 millions of US$). From the unconditional summary statistics in Table 2.1, a one interquartile variation reflect an increase in asset size by nearly 358 p.p.. Clearly, this figure mixes up both cross-sectional and time-series variation.

However, our interest in asset-size is aimed at understanding the distribution of the relation between interest rate shocks and maturity structure in the cross-section of firms. To this end, we look at the within (3-digit SIC) industry time-varying distribution of total asset size and define a dummy variable, $Large_{f,t-1}$, with value 1 if a company is in the upper quartile and 0 otherwise. The choice is due to the fact that - as we will show in Section 2.3.2 - it is within this class of firms that debt maturity structure responds to changes in the FED interest rate policy. Moreover, Figure 2.A.3 shows how, to start with, the LT debt share is unevenly distributed between firms and increasing in firm size. As a matter of fact, for firms in the first asset-size quartile, the average LT debt share equals roughly 50%, whereas it amounts to nearly 80% for companies in the upper quartiles. Throughout the rest of the paper, we refer to companies in the top-size quartile of their industry distribution as to "large" companies. We also gather additional information from Compustat on other firm-level controls such as leverage, liquid assets and sales growth.

To test our mechanism, we then retrieve data on the issuance of bonds with maturity above 1 year from Mergent FISD. Information from Compustat and Mergent FISD are matched through the (6-digit) issuer CUSIP,[13] resulting in a sample of 2,858 bond is-

[13] In a couple of dozen of cases, multiple companies - typically 2 or 3 - have the same 6-digit CUSIP in Compustat, referring to different subsidiaries of a same group. In such cases, we retain the largest company in Compustat among the ones with same 6-digit CUSIP in an effort to identify the mother company. Excluding all such companies would not affect the results.

suers in Mergent FISD.[14] We start by analyzing bond issuance on the extensive margin through the dummy $\mathbb{K}(Issue)_{f,t+j}$, with value 1 if a firm f issues bonds in year-quarter $t+j$ and 0 otherwise,[15] $j = 0, 1, ..., 20$. On average, the likelihood of a current year-quarter new issuance is $6{,}77\%$, suggesting that bond issuance is relatively lumpy and infrequent. Such average increases slightly but steadily over future horizons, reflecting the fact that older and/or larger companies tend to issue bonds relatively more frequently. Indeed, Figure 2.A.4 looks at the number of bond issuances per year-quarter and splits them depending on whether they are conducted by a large company or not. The share of new issuances by large companies is disproportionately large. In fact, while such firms account (by construction) for roughly 1/4 of the firms, they represent about 60% of new bond issuances. In other terms, this is prima-facie evidence that large companies are much more active in the corporate bond market and hence more likely to react to potential variation in the associated financing costs. In particular, as we are interested in the response of financing costs to monetary policy, we retain data on the annualized coupon rate at issuance, $CouponRate_{f,t+j}$, which is equal to 6% on average.[16]

2.2.3 Corporate Bond Mutual Funds Data

We retrieve data on corporate bond mutual funds (CBMF)[17] holdings from the CRSP Survivor Bias-Free dataset, including information on both surviving and dead funds. Following Choi and Kronlund (2018), we split funds into High Yield (HY) and Investment Grade (IG) funds based on standard Lipper style codes.[18] We label HY-funds as yield-oriented: as shown by Choi and Kronlund (2018), they intuitively invest relatively

[14]Such number of firms refers to the companies in our regression sample. Since we apply firm fixed effects in our regressions, those are companies that issue bonds at least twice throughout our period of analysis.

[15]Assigning 0 to periods in which Mergent FISD does not report a company's bond issuance requires to know whether a company is active or not. To this end, we label a company as "active" if it reports balance sheet information in Compustat, and as "inactive" if it does not. In practice, disapperance from Compustat means that a firm has delisted, implying that it cannot issue bonds anymore.

[16]The large fall in observations with respect to those for the variable $\mathbb{K}Issue_{f,t+j}$ reflects the fact that the distribution of $Coupon_{f,t+j}$ is conditional on $\mathbb{K}Issue_{f,t+j} = 1$.

[17]Specifically, like Choi and Kronlund (2018), we limit the sample of funds to CRSP style categories I, ICQH, ICQM, ICQY, ICDI, ICDS, or IC.

[18]IG funds are classified as those with a Lipper style code of either A, BBB, IID, SII, SID, or USO and HY funds are those coded HY, GB, FLX, MSI, or SFI.

more in longer and riskier debt-securities.[19] Ideally, one would build a measure of fund-specific reach-for-yield, but this requires security-level data on CBMFs' holdings that we do not have access to. Hence, we use the just described second-best, empirically grounded HY-vs-IG funds proxy.

Overall, we analyze 3,487 funds (2,034 are IG and 1,453 are HY) over the data-period 2010q2-2018q2. Table 2.1 describes the related summary statistics. A first outcome variable of interest is the cumulative growth rate of corporate bond holdings through time, $\Delta CB_{m,t+j}$, which displays a large extent of heterogeneity across funds. Second, we are also interested in the changes in the fund's average (weighted) portfolio maturity over time, $\Delta Matu_{m,t+j}$, equally showing significant differences in the cross-section of funds. We gather additional information on fund characteristics, used as controls in our models, including the fund's turnover and expenses ratio, the net asset value and returns.

2.3 Empirical Analysis

In this section, we present the baseline empirical findings of our paper. First, we present the aggregate-level analysis. Next, we investigate cross-sectional differences across firms.

2.3.1 Time-series Analysis

We apply local projections (Jordà 2005) to study the response of the share of LT-debt to changes in the FED's interest rate policy. In particular, we estimate separately the following regressions through OLS:

$$\Delta_h y_{t+h} = \beta_{1,h} \Delta EFFR_t + \Gamma_h X_{t-1} + u_{t,h} \tag{2.1}$$

for $h = 0, 1, ..., 20$. The dependent variable, $\Delta_h y_{t+h}$, is given by the cumulative variation in the share of LT-debt between year-quarters $t-1$ and $t+h$. Hence, plotting

[19]We refer to the data in Table 1, Panel B from Choi and Kronlund (2018) and to the related discussion in the paper. Interestingly, HY funds reach-for-yield relatively more than IG funds, i.e. they invest in securities with higher yields. Importantly, this difference is explained by both risk and maturity, whereas IG funds tend to reach-for-yield more within a given bucket of risk-maturity bucket.

the coefficients $\beta_{1,h}$ provides the impulse-response function of the share of LT debt to a change in the EFFR as of year-quarter t, $\Delta EFFR_t$. Moreover, X_{t-1} is a vector of lagged macro-controls, including variables which might simultaneously have an influence on $\Delta_h y_{t+h}$ and on the current interest rate policy. In particular, X_{t-1} comprehends: the annual GDP growth rate and inflation rate; the quarterly variation in the 10y-3m term-spread, in the corporate spread and in the share of Treasuries with maturity above 20-year; a recession dummy. Finally, $u_{t,h}$ is a robust error-term.

Figure 2.A.5 reports the impulse-response function obtained from the OLS-estimation of the coefficients $\beta_{1,h}$ in Equation 2.1. In particular, the plot assumes a 25 b.p. quarterly negative variation in the EFFR - i.e., a loosening of the short-term policy rate, a convention we maintain throughout the rest of the paper - and also displays the 10% confidence interval around the point estimates. Clearly, an interest rate cut boosts the share of LT-debt. The effect is very persistent and, while effective at impact, peaks up 3 years after the shock, and does not fade away throughout the considered 5-year time-window. Such implausibly large degree of persistence might reflect the endogeneity of the simple raw variation in the EFFR and its significant autocorrelation along monetary policy cycles. That said, the effect is economically meaningful. For instance, a 25 b.p. interest rate descent in year-quarter t implies a cumulative increase in the share of LT-debt by 0.42 p.p. one year after. For comparison, the average growth rate of the LT-debt share equals 0.15 p.p. on a quarterly basis, and 0.83 p.p. on an annual basis.

Importantly, we validate that such result is robust to employing alternative and exogenous monetary policy shocks. That is, in Equation 2.1, we replace $\Delta EFFR_t$ with $\varepsilon_t^{mp,g}$ or $\varepsilon_t^{mp,jk}$, i.e., the high-frequency surprises on FED Funds Futures rates from Gürkaynak, Sack, and Swanson (2005) and the related pure interest rate shocks from Jarocinski and Karadi (2020), respectively. Figure 2.A.6 shows the resulting impulse-response functions, calibrated for a 1 s.d. expansionary exogenous shock. The analysis validates the positive effect of an interest rate loosening on the LT-debt share. Quantitatively speaking, the influence of both shocks is similar, and aligned to that of a 25 b.p. reduction in the raw EFFR. Nonetheless, the shock displays a much less persistent effect, with the impact on the share of LT-debt vanishing in 5 or 14 year-quarters when using the Jarocinski and Karadi (2020) and Gürkaynak, Sack, and Swanson (2005) shocks, respectively.

Finally, we also verify in the Empirical Appendix Figures A1 and A2 that the findings survive to restricting the sample to the period between 1990q1 and 2008q4, which we label as pre-crisis. This is an important robustness check in that most of the variation in the FED's interest rate policy occurs before 2008.

2.3.2 Firm-level Analysis

We investigate US listed firms' quarterly balance-sheets in order to understand cross-sectional differences in the response of the LT-debt share to variations in the interest rate policy. To this end, we borrow the empirical strategy from Jeenas (2018) and Ottonello and Winberry (forthcoming), using a panel version of the Jordà (2005)'s local projections. In practical terms, we estimate by OLS the following set of equations:

$$\begin{aligned}\Delta_h y_{f,t+h} =& \beta_{1,h}\Delta EFFR_t + \beta_{2,h}Large_{f,t-1} + \beta_{3,h}Large_{f,t-1} * \Delta EFFR_t + \\ &+ \Gamma_h X_{f,t-1} + \mu_f + \mu_{s,t} + u_{f,t+h}\end{aligned} \tag{2.2}$$

for $h = 0, 1, ..., 20$. The dependent variable, $\Delta_h y_{f,t+h}$, is given by the cumulative variation of the share of LT-debt of firm f between year-quarters $t-1$ and $t+h$. Most importantly, the model includes the full interaction of the raw quarterly change in EFFR, $\Delta EFFR_t$, and a dummy for large companies, i.e., with value 1 for companies in the upper quartile of the respective industry asset-size distribution, $Large_{f,t-1}$. The coefficient of main interest is $\beta_{3,h}$, capturing the relative response of large companies (as compared to smaller ones) to a variation in the FED short-term policy rate.

We augment the model with a vector of firm controls, which comprehends the (lagged) share of liquid assets, leverage and sales quarterly growth. Eventually, such variables are also fully interacted with $\Delta EFFR_t$. By doing so, we horse-race our channel (based on firms size as a proxy for bond financing constraints) against other layers of heterogeneity which have been found to influence firms' response to monetary policy shocks.[20] Furthermore, we interact $Large_{f,t-1}$ with the usual set of macrocontrols for avoiding that $\beta_{3,h}$ reflects contemporaneous response of large companies to other shocks, which may correlate with the FED's interest rate policy decisions. The

[20] Jeenas (2018) shows that companies with a relatively lower share of liquid assets respond more to monetary policy shocks. Ottonello and Winberry (forthcoming) find that distance to default matters as well for firms' reaction to monetary policy, with leverage being a good proxy for it. Sales growth is meant to capture a firm's profitability.

model is saturated with firm and indyustry*year-quarter fixed effects, i.e. μ_f and $\mu_{s,t}$, respectively. The former set of dummies controls for all observed and unobserved time-invariant heterogeneity at the level of the firm; the latter absorbs time-varying shocks which are common to firms in a given (3-digit SIC) industry. The application of such fixed effects implies that our coefficient of interest $\beta_{3,h}$ is identified by: i) within-firm variation over time, i.e., changes in response of LT-debt share by an otherwise identical firm when it is large as compared to when it was small; ii) cross-sectional variation across firms in a given industry. Finally, $u_{f,t+h}$ is an error term, which we double-cluster at the firm and industry*year-quarter level.

However, the relative adjustment of large companies estimated through Equation 2.2 does not allow to understand the overall response of both large and smaller companies. In fact, Equation 2.2 is saturated with industry*year-quarter fixed effects, which span out completely time-series variation common across all firms.

Hence, we additionally estimate the following model separately for firms in different size-quartiles:

$$\Delta_h y_{f,t+h} = \beta_{1,h} \Delta EFFR_t + \Psi_h X_{f,t-1} + \mu_f + \nu_{f,t+h} \tag{2.3}$$

That is, we estimate a model which exploits just time-variation and hence describes the absolute change in the share of LT-debt after a change in interest rate by the FED.[21] In fact, we do not use year-quarter fixed effects (nor any subtler version of them), while we keep using firm fixed effects to control for unobserved time-invariant firm heterogeneity.

The trade-off across the two models is clear: Equation 2.2 precisely estimates the cross-sectional differences across firms, as it controls for time-varying common heterogeneity within narrowly defined industries. On the other hand, model 2.3 pins down the absolute variation in LT-debt share associated to variation in the policy rate. Hence, it serves the purpose of better understanding the connection between firm-level and time-series findings.

[21] In Equation 2.3, $X_{f,t-1}$ is the usual vector of macro and firm-level controls.

2.3.2.1 Results

Figure 2.A.7 plots the impulse-response function obtained from the estimation of the parameters $\beta_{3,h}$ - for $h = 0, 1, ..., 20$ - from Equation 2.2. Relatively to smaller firms, large companies expand LT-debt more when the policy rate goes down. That is, large companies react more in line with the aggregate-level evidence shown in Section 2.3.1. For understanding the absolute response-level, however, we additionally estimate Equation 2.3 within different size-quartiles. The resulting impulse response functions are displayed in Figure 2.A.9 and suggest that only large companies do adjust, whereas smaller firms' LT-debt share is generally insensitive to monetary policy. Moreover, we replicate both exercises using the by-now familiar exogenous monetary policy shocks from Gürkaynak, Sack, and Swanson (2005) and Jarocinski and Karadi (2020), obtaining specular findings (see Figures 2.A.8 and 2.A.10).

For gauging the economic significance of the just commented effects, we refer to the baseline figures employing the endogenous quarterly change in the EFFR. Once again, Figures 2.A.7 and 2.A.9 are calibrated to a 25 b.p. (expansionary) negative EFFR variation. The jump in the share of LT-debt by large companies peaks up 6 year-quarters after the policy change, when it is comprised between 0.45 and 0.55 p.p. (depending on whether one takes as a reference the adjustment in Figure 2.A.9 or in Figure 2.A.7, respectively).[22] Interestingly, the described size of the effect is comparable - at relevant horizons - with that observed at the aggregate level. At the firm-level, the 6 year-quarter cumulative growth of the LT-debt share equals -1.53 p.p. on average (not shown for brevity in Table 2.1).

Size turns out being the key firm-level attribute to explain cross-sectional differences across firms. In this respect, we report in Table 2.2 additional coefficients from the estimation of the baseline firm-level model (Equation 2.2).[23] In particular, we show the horse-race with the other balance-sheet characteristics employed as firm-level controls. First, companies tend to increase the share of LT-debt when sales jump; nonetheless, the interaction of such dynamics with monetary policy is insignificant. Moreover, intuitively, the share of LT-debt goes down when firms hold relatively more liquid assets,

[22]From a formal perspective, the absolute variation in LT-debt share is pinned down in Figure 2.A.9. Nonetheless, we also refer to Figure 2.A.7 as it estimates precisely the relative adjustment of large companies and the baseline effect on smaller ones can be placed at 0.

[23]Table 2.2 displays coefficients associated to a 1 p.p. (contractionary) increase in the EFFR.

reflecting maturity matching of assets and liabilities. Also in this case, however, the share of held liquid assets does not influence the relation of debt maturity structure and monetary policy. Similarly, leverage has a small and marginally significant influence on such relation at impact, but the effect fades away already 1 year-quarter after, whereas at the macroeconomic level the policy rate has a more persistent impact on the share of LT-debt. Differently, the effect of our dummy for large companies is strongly significant and persistent over time, and resembles well the patterns observed at the aggregate level.[24]

Finally, for robustness, we repeat the exercise over the pre-crisis period (i.e., from 1990 to 2008, included) so to restrict our analysis to a time-window with substantial interest rate shocks. Results are reported in Appendix Figures A3 and A4 and confirm the baseline findings both qualitatively and quantitatively.

2.4 Model

In this section, we present a model which explains our empirical findings: i) the aggregate-level share of LT-debt goes up following an interest rate loosening; ii) such effect is entirely driven by the adjustment of very large companies. Existing theoretical frameworks are not useful in this respect. In fact, models of monetary transmission to firms do not include an explicit discussion on debt maturity, whereas the corporate finance literature does not typically focus on the policy rate.[25] First, we present the model setup. Next, we characterize the equilibrium conditions and perform comparative statics exercises which pin down the relation between the short-term policy rate and the firms' debt

[24] The coefficient on the large company dummy alone turns out being insignificant. This result - apparently in contradiction to the stark differences in the share of LT-debt observed across size-quartiles in Figure 2.A.3 - is due to the inclusion of firm fixed effects in our regressions. Hence, accordingly with our aim of exploring cross-sectional differences, the relevant variation in size captured by our dummy operates mostly between companies, rather than within.

[25] In corporate finance, the focus rather rests on the term-spread, as it describes the relative cost of long-term debt relative to short-term debt. In particular, a jump in the term-spread would predict an increase in the relative cost of LT debt, and hence a related decrease in the issuance of LT bonds - a prediction contradicting our findings. As a matter of fact, a policy rate loosening predicts an expansion of the term spread (see e.g. Adrian and Shin 2010), a robust relation which we document to hold in our sample. In Appendix Figure A5, a simple scatterplot - with the quarterly change in the EFFR on the x-axis and the variations of the term-spread on the y-axis - suggests a strong negative relation. We test such influence more formally in Appendix Figures A6 and A7, which reports the impulse-response function from a model with the term-spread as dependent variable but otherwise identical to that in Equation 2.1.

maturity structure. Finally, we perform few empirical test to validate the mechanism proposed by the model.

2.4.1 Setup

Our economy lasts three periods ($t = 0, 1, 2$) and is populated by a continuum of firms and investors. Each firm is endowed with capital A - a proxy for firm size - and a project. A is heterogeneous across firms and distributed uniformly across firms on the interval $[0, I]$, where I is the initial investment into each firm's project in period 0. Each project also features a stochastic re-investment ρ in period 1, drawn out of an exponential distribution $f(\rho) = \chi e^{-\chi\rho}$, for $\rho \in [0, \infty)$. If the reinvestment need is not met, the project is liquidated and does not generate any payoff in period 2.

Each project generates a riskless short-term pay-out r in period 1. Differently, in period 2, conditional on the reinvestment need being satisfied, the project yields R in case of success, and zero in case of failure. The likelihood of success depends on firms' behavior. We assume that if a firm exerts effort, the project is successful with probability p_h (without loss of generality, we set $p_h = 1$); on the other hand, if the firm shirks, success materializes with probability $p_l < 1$, but it enjoys private benefits B. Additionally, there is aggregate risk: with probability $1 - \delta$, all firms get a pay-out of zero in period 2. Firms do not have a storage technology and are protected by limited liability.

Investors are competitive and a subset of them features reach-for-yield behavior, meaning that they care about the return on their portfolios relative to the current interest rate, as opposed to the series of current and future interest rates. They therefore take more risk than rational investors if interest rates are low.

The short term interest rates between period 0 and 1, and 1 and 2, respectively, are set exogenously by a monetary authority. These exogenous short-term interest rates are denoted as i_1 for the interest rate from period 0 to period 1, and as i_2 for the interest rate from period 1 to period 2. All agents in this model use these short-term rates as discount rates and we treat i_1 as the policy rate. Investors can lend directly at these exogenous rates, while firms receive funds intermediated by the investors. In the following, we lay out formally how firms and investors are modeled.

2.4.1.0.1 Firms.

We follow Holmström and Tirole (1998, 2000) in modeling firm financing as a moral hazard problem in which each individual firm is an agent.

We assume that firms can credibly commit to a contract stating that the project is carried on to period 2 whenever the stochastic re-investment is sufficiently small:

$$\rho \leq \rho^*$$

and terminated otherwise. The continuation threshold ρ^* is a choice variable of the firm and is common knowledge. In equilibrium firms will differ in their choice of ρ^*, thus we use the notation $\rho^*(A)$, denoting the choice of a firm with capital A.

To finance the gap between the initial investment and the endowment, firms issue short and long-term debt. A riskless short-term bond is sold at price $P_s = 1/(1+i_1)$ at time 0 and is promised to yield 1 in expectation in period 1, while a long-term bond is sold at price P_l in time 0 and yields 1 in period 2 if the project is successful. The amount of issued short-term and long-term debt is denoted by d_s and d_l, respectively.

The timing of the financing and execution of the project is as follows. First, in period 1, short-term creditors are compensated out of earnings r, as the firm must assure to repay d_s. Next, the firm draws the re-financing shock from $f(\rho)$. If the decision is not to refinance - i.e., if $\rho > \rho^*$ - then the firm abandons the project and consumes what is left, whereas long term bond-holders do not receive any compensation. If the project is continued - i.e., if $\rho \leq \rho^*$ - and turns out to be successful in period 2, the entrepreneur enjoys $R_b = R - d_l$, while long-term bond-holders receive their compensation d_l. Eventually, if the project is unsuccessful then the firm is again liquidated at value zero and neither the long-term bond-holders, nor the entrepreneur, receive anything. Hence, it follows that, to induce the entrepreneur to exert effort, the following condition must hold:

$$R - \frac{B}{\Delta p} \geq d_l.$$

where $\Delta p = p_h - p_l$. Intuitively, this incentive compatibility constraint means that the repayment in an optimal contract cannot be too large, otherwise the entrepreneur will shirk.

Moreover, limited liability and riskless short-term debt implies:

$$\rho^* \leq r - d_s$$

i.e., the firm cannot be asked to meet the liquidity shock with other funds than those stemming from the project returns. Finally, the firm must raise (and investors must be willing to provide) enough money to finance the project in the first place:

$$\frac{1}{1+i_1}d_s + P_l d_l \geq I - A.$$

Taking stock, the general problem of the firm reads (from now on we will index the choice variable by the endowment):

$$\max_{\rho^*(A),d_s(A),d_l(A)} \frac{r}{1+i_1} - \frac{\int_0^{\rho^*(A)} \rho f(\rho) d\rho}{1+i_1} + \frac{\delta F(\rho^*(A))}{(1+i_1)(1+i_2)} R + \left(P_l(A) - \frac{\delta F(\rho^*(A))}{(1+i_1)(1+i_2)} \right) d_l(A) \tag{2.4}$$

subject to:

$$r - \rho^*(A) \geq d_s(A) \tag{LL}$$

$$R - \frac{B}{\Delta P} \geq d_l(A) \tag{IC}$$

$$\frac{d_s(A)}{1+i_1} + P_l(A) d_l(A) \geq I - A \tag{IR}$$

The objective function represents expected firm profits, discounted as of $t = 0$. The first term gives the risk-free period-1 revenues and the second one subtracts the expected period-1 payments due to the liquidity shock. The third element provides the expected period-2 revenues, influenced by both idiosyncratic liquidity shock and aggregate risk. The last term of the equation collects the net proceeds from the issuance of LT bonds, i.e., the value of liquidity minus total expected repayments.

2.4.1.0.2 Investors. Firms borrow in bond markets featuring a continuum of investors. Investors are heterogeneous and we denote their type as j. Investors have zero initial wealth and construct long-short positions to maximize:

$$E[w^j] - \frac{\gamma}{2} Var[w^j],$$

where w^j is wealth of an investor of type j as of $t = 2$. They purchase a portfolio of LT debt, issued by the firms, and finance this position by rolling over short-term borrowing.

As a result, w^j equals:

$$w^j = d_l^{*j} - \iota(i_1, i_2, j) \int_0^I P_l(A) d_l^j(A) dA.^{26}$$

Here, d_l^{j*} is the realized payoff from holding a portfolio comprising LT debt of all firms, whereas $\iota(i_1, i_2, j)$ is the individual (compound) factor that each investor uses to judge her financing costs. $\iota(i_1, i_2, j)$ is heterogeneous across investor types. In particular, we assume two investor types: $j \in \{R, Y\}$. A fraction $1 - \alpha$ of the investors are "rational" and their compounded discount rate $\iota(i_1, i_2, R)$ is $(1 + i_1)(1 + i_2)$. On the other hand, a fraction α of the investors is of the "yield-seeking" type, whose modelling we borrow from Hanson and Stein (2015). Specifically, such yield-seeking investors compare the expected returns from their investments only with the current interest rate instead of the stream of expected interest rates. Their $\iota(i_1, i_2, Y)$ is $(1+i_1)^2$.[27] The explicit expectation and variance of investor wealth at $t = 2$ are:

$$E[w^j] = \delta \int_0^I F(\rho^*(A)) d_l^j(A) dA - \iota(i_1, i_2, j) \int_0^I P_l(A) d_l^j(A) dA$$

$$Var[w^j] = Var[d_l^{*j}] = \left(\int_0^I F(\rho^*(A)) d_l^j(A) dA \right)^2 \delta(1 - \delta).$$

In the expression for expected wealth, the expected revenues reflect the fact that, for a firm with endowment A, the likelihood of repayment equals $\delta F(\rho^*(A))$. Next, in the variance term, $\int_A F(\rho^*(A)) d_l^j(A)$ is treated like a constant due to full diversification of firms' idiosyncratic risk; in other terms, investors' risk only depends on aggregate shocks.

Investors maximize their wealth by optimally choosing a LT debt portfolio including all firms' debt and, due to the mean-variance utility assumption, they have limited risk-bearing capacity. We assume that there two types of investors. The general problem for

[26] We assume that the mass of firms is I, such that the density of each firm type A is 1 and can be omitted from notations.

[27] The fact that these investors are discounting using an incorrect rate generates the yield-seeking behavior. This modeling choice can be justified by agency or accounting considerations that lead investors to worry about short-term measures of reported performance.

both types is:

$$\max_{d_l^j(A)} \delta \int_0^I F(\rho^*(A)) d_l^j(A) dA - \iota(i_1, i_2, j) \int_0^I P_l(A) d_l^j(A) dA - \frac{\gamma}{2} \left(\int_0^I F(\rho^*(A)) d_l^j(A) dA \right)^2 \delta(1-\delta). \tag{2.5}$$

Finally, we assume that there is an inelastic demand g, originating from preferred-habitat investors into LT debt, that is defined in terms of expected bond payments in $t = 2$.[28]

2.4.2 Discussion of Setup

It is important to discuss the modeling choices and assumptions we made in our model. The model by Holmström and Tirole (1998, 2000) provides a tractable framework to study how the maturity structure interacts with financing constraints and investors demand. In their model, the defining difference between short-term and long-term financing is credit risk that affects only debt of longer maturity. As Holmström and Tirole (1998, 2000), we abstract from other sources of risk, such as duration or rollover risk. Our goal is to endogenize the maturity structure of liabilities of the corporate sector to compare theoretical predictions with the empirical regularities discussed above. The key model ingredients for this are the intermediate income r, which provides firms a cash-flow to service short-term debt, and the incentive compatibility constraint, which restricts firms long-term debt choice. The additional trade-off of more continuation against more short-term debt, adds a margin of adjustment for the maturity structure of constrained firms. Since our investors are risk-averse and thus more complex than those in Holmström and Tirole (1998, 2000), we make some simplifying assumptions, namely that there is no storage technology, and that short-term debt is riskless. The latter can be rationalized when the intermediate cash-flow is large enough, so the firm is sufficiently "cash-rich", then, once its short-term debt has reached $r - \rho^*$ no firm would not want to increase short-term debt above this riskless amount, as the price would deteriorate too fast for an increase in short-term debt to be revenue-increasing. As in Holmström and Tirole (1998, 2000), we abstract from firms rolling short-term debt over to $t = 2$. This allows us to find a well-defined maturity structure and is without loss of generality: as all contingencies and their probabilities are known by firms and investors in $t = 0$, and

[28]Greenwood, Hanson, and Stein (2010) make a similar assumption and describe such investors as pension funds, life insurance companies, endowments, or any institution with an inelastic demand for long-term assets.

the contract specifies how to deal with them, there will be no incentive to refinance once the liquidity shock has realized.

There has long been an established theoretical literature on firms' optimal debt maturity choice such as Flannery (1986, 1994), Diamond (1991), Diamond and He (2014), and He and Milbradt (2016). The downside of the aforementioned papers is that they do not focus on firms' financial constraints. Our empirical evidence suggests that financing constraints are a crucial element of the explanation and thus we use Holmström and Tirole (1998, 2000) as starting point of our analysis. Other papers consider the effect that a given debt maturity has on financial outcomes, such as roll-over risk and credit risk He and Xiong (2012a,b). Relative to these papers, we are exploring the relationship in the opposite direction, in that we are trying to understand how changes in financing conditions affect maturity choices.

Moreover, as we show below, a model without yield-seeking investors does not match these empirical facts. This is our motivation to extend such that investor demand is sensitive to an interest rate change. We achieve this by introducing risk-aversion and reach-for-yield behavior for some investors.[29]

2.4.3 Equilibrium

First we formally characterize the competitive equilibrium of our model.

Definition 2.1. A competitive equilibrium is a set of quantities $\{d_s(A), d_l(A), d_l^R(A), d_l^Y(A)\}_{A\in[0,I]}$, cut-off rules $\{\rho^*(A)\}_{A\in[0,I]}$ and prices $\{P_l(A)\}_{A\in[0,I]}$ such that:

1. $\{d_s(A), d_l(A), \rho^*(A)\}_{A\in[0,I]}$ solve firms' optimization problem 2.4, given $\{P_l(A)\}_{A\in[0,I]}$.

2. $\{d_l^R(A), d_l^Y(A)\}_{A\in[0,I]}$ solve rational and yield-seeking investors' respective maximization problems 2.5.

3. The LT bond market clears:

$$d_l(A) = \alpha d_l^R(A) + (1-\alpha)d_l^Y(A) + \frac{g}{\int_0^I F(\rho^*(A))dA}.^{30} \tag{2.6}$$

[29]There are various approaches to model reach-for-yield behavior, such as Drechsler, Savov, and Schnabl (2018), Acharya and Naqvi (2019), Lu et al. (2019), and Campbell and Sigalov (2020). We follow the approach of Hanson and Stein (2015) who model reach-for-yield as a subset of agents using the current interest rate to discount future income, instead of the path of expected future interest rates.

[30]The inclusion of a large enough g ensures that, under any circumstances, all firms borrow a positive

We start with the first-order conditions for "rational" and "yield-seeking" investors, respectively:

$$P_l(A) = \frac{\delta F(\rho^*(A)) - \gamma F(\rho^*(A))\delta(1-\delta)\int_0^I F(\rho^*(A))d_l^R(A)dA}{(1+i_1)(1+i_2)} \tag{2.7}$$

$$P_l(A) = \frac{\delta F(\rho^*(A)) - \gamma F(\rho^*(A))\delta(1-\delta)\int_0^I F(\rho^*(A))d_l^Y(A)dA}{(1+i_1)^2} \tag{2.8}$$

where $d_l^Y(A)$ and $d_l^R(A)$ are the demand of firm A's bonds by yield-seeking and rational investors, respectively. The two investor types compete in the same market to buy LT debt and face the same price. However, their demand differs due to a different attitude towards interest rates. The first term in the right-hand side of both expressions gives the expected payoff from holding a firm A's LT bonds. Importantly, for such an expected payoff, yield-seeking investors are willing to pay a premium on LT debt if $i_2 > i_1$, as they overreact to changes in i_1, relative to rational investors. The second term suggests that investors, being risk averse, are compensated (through a lower price) for holding risky LT debt. Furthermore, a smaller continuation cutoff $\rho^*(A)$ implies a lower price, as the probability of repayment in $t = 2$ descents.

Rearranging Equations 2.7-2.8 and plugging them into the market clearing condition 2.6 yields the inverse demand for firm A's LT-debt:

$$P_l(A) = \delta F(\rho^*(A)) \frac{1 - \gamma(1-\delta)\left(\int_0^I F(\rho^*(A))d_l(A)dA - g\right)}{(1+i_1)[\alpha(1+i_1) + (1-\alpha)(1+i_2)]}. \tag{2.9}$$

The discount factor in Equation 2.9 is a weighted average of the discount factors of the two types of investors. Moreover, the willingness to pay of the marginal investor decreases in the aggregate volume of LT debt held. Before proceeding further, we state a lemma that will be useful for the derivations of our key results.

Lemma 2.1. *The price for LT debt of a firm with endowment A is unaffected by changes in the supply of LT debt $d_l(A)$ and increases in $\rho^*(A)$.*

Proof. See the Theory Appendix. □

---

amount of LT-debt.

Next, the first order conditions of the firms' maximization program are:

$$\lambda_3(A) = \lambda_1(A) \tag{2.10}$$

$$P_l(A) = \frac{1+\lambda_2(A)}{1+\lambda_3(A)} \frac{\delta F(\rho^*(A))}{(1+i_1)(1+i_2)} \tag{2.11}$$

$$f(\rho^*(A)) \left[\frac{\delta R}{1+i_2} - \rho^*(A) \right] + \left((1+i_1)(1+\lambda_3(A)) \frac{\partial P_l(A)}{\partial \rho^*(A)} - \frac{\delta f(\rho^*(A))}{1+i_2} \right) d_l(A) - \lambda_1(A) = 0 \tag{2.12}$$

where $\lambda_1, \lambda_2, \lambda_3$ are the multipliers linked to the three constraints LL, IC and IR, respectively. Condition 2.10 signals that LL binds if and only if IR does. In condition 2.11, the firm valuation of one unit of LT debt equals the NPV of the risky project times a factor positively (negatively) related to the tightness of the IC (IR) constraint. Finally, in condition 2.12, the optimal liquidation cutoff decreases when the LL constraint binds relatively more.

2.4.4 Unconstrained and Constrained Firms

In this section we analyze the conditions for the existence of unconstrained and constrained firms in equilibrium and characterize their optimal plans. Focusing on such equilibrium, rather than others where all firms are either constrained or unconstrained (in the sense specified below), allows us to relate our theory to the cross-sectional empirical findings in Section 2.3.2. We call a firm "unconstrained" whenever all three constrains are slack and "constrained" if the opposite is true.

Proposition 2.1. *The three constraints of the firm problem only bind concurrently. Unconstrained and constrained firms coexist in equilibrium if:*

$$\bar{A} \in (0, I)$$

where:

$$\bar{A} = I - \frac{r - \frac{\delta R}{1+i_2}}{1+i_1} - \frac{F\left(\frac{\delta R}{1+i_2}\right)}{(1+i_1)(1+i_2)} \delta \left(R - \frac{B}{\Delta P} \right)$$

Proof. See the Theory Appendix. □

$\bar{A}$ denotes the lowest endowment at which the optimal continuation threshold is just feasible if the firm takes on the largest possible level of (short and long-term) debt. In practice, $A = \bar{A}$ is a threshold value for firm size, above (below) which firms can (cannot) implement the optimal continuation value.

We now proceed with the description of the optimal plans for unconstrained and constrained firms. For unconstrained firms $\rho^*(A)$ is at the optimal value:

$$\rho^*(A) = \frac{\delta R}{1 + i_2} \text{ if } A > \bar{A} \tag{2.13}$$

which follows from condition 2.10 and represents the risky project's revenues discounted of aggregate risk. For these firms, the limited liability constraint does not bind, so they are indifferent with respect to taking any amount of short-term debt. Moreover, they collectively invest into LT debt until its price equals their own valuation, as reported in Equation 2.11.

We find it useful to make an assumption on how LT debt is distributed between unconstrained firms, so to match the stylized empirical fact in Figure 2.A.3, i.e., unconstrained firms have a larger LT debt share than constrained firms. Concretely, we assume that unconstrained firms choose a combination of LT debt and short-term debt allowing them to match the highest LT debt ratio among constrained firms. [31]

Differently, a constrained company piles up as much short- and long-term debt as they need, namely for $A < \bar{A}$:

$$d_l(A) = R - \frac{B}{\Delta P}$$
$$d_s(A) = r - \rho^*(A).$$

Finally, the next lemma shows that the continuation cutoff of constrained firms is set below the optimal level.

[31] Details in the Theory Appendix in the proof of Proposition 2.2. This assumption is made possible by a appropriately large preferred-habitat investor demand g. In fact, it implies a relatively large excess-demand for LT-debt, which is filled by unconstrained companies.

Lemma 2.2. *Assuming* $\frac{\chi\delta}{1+i_2}\left(R - \frac{B}{\Delta p}\right) < 1$*, a constrained firm chooses a continuation value* $\rho^*(A) \in \left[0, \frac{\delta R}{1+i_2}\right)$*. It also follows that:*

1. $\frac{\partial \rho^*(A)}{\partial A} > 0$ *if* $A \leq \bar{A}$ *and* $\frac{\partial \rho^*(A)}{\partial A} = 0$ *if* $A > \bar{A}$
2. $\frac{\partial \rho^*(A)}{\partial i_1} < 0$ *if* $A \leq \bar{A}$ *and* $\frac{\partial \rho^*(A)}{\partial i_1} = 0$ *if* $A > \bar{A}$

Proof. See the Theory Appendix. □

Intuitively, firms below $\bar{A}$ need a large amount of debt. In particular, high reliance on short-term financing tightens the limited liability constraint, so that ρ^* is compromised. Increased firms own capital (i.e. size A) reduces the need for short-term finance, therefore increasing the chosen continuation cutoff for constrained companies. Similarly, a descent in the short-term rate i_1 boosts the price of short- and LT debt, thereby loosening the IR constraint, which co-moves with the LL constraint by condition 2.10. A relaxation of the LL constraint increases the feasible continuation value. Clearly, both relations do not apply among unconstrained companies as any firm with $A > \bar{A}$ is already at the optimum.

2.4.5 Effects of Policy Rate Change on the Maturity Structure

In this section, we study how changes to the short-term policy rate i_1 - controlled by the monetary authority - affect firms' debt maturity structure. We first derive predictions in a setting without yield-oriented investors and next in the baseline model presented above to highlight the fact that the inclusion of reach-for-yield motives allows us to replicate the empirical facts documented in Section 2.3.

2.4.5.1 Effect of Policy Rate Change Without Yield-seeking Investors

The absence of reach-for-yield motives represents a limiting case of the above model in which $\alpha = 0$. In such setting, rational investors have no incentive to hold LT debt, because it is risky and firms are just willing to price debt in a risk-neutral fashion.[32] In

[32] This can be seen in condition 2.11. For unconstrained companies, $\lambda_2 = \lambda_3 = 0$, so that the resulting price clearly just discounts aggregate and idiosyncratic risk, without offering any risk-premium.

practical terms, the inverse demand function for LT bonds equals:

$$P_l(A) = \delta F(\rho^*(A)) \frac{1 - \gamma(1-\delta)\left(\int_0^I F(\rho^*(A)) d_l(A) dA - g\right)}{(1+i_1)(1+i_2)}.$$

Market clearing implies:

$$\frac{1 - \gamma(1-\delta)\left(\int_0^I F(\rho^*(A)) d_l(A) dA - g\right)}{(1+i_1)(1+i_2)} = \frac{1}{(1+i_1)(1+i_2)},$$

or

$$\int_0^I P_l(A) d_l(A) dA = \int_0^{\bar{A}} F(\rho^*(A)) d_l(A) dA + \int_{\bar{A}}^I F(\rho^*(A)) d_l(A) dA = g. \quad (2.14)$$

That is, the LT bonds issued by constrained and unconstrained firms have to net out. In this case, in reaction to a descent in the interest rate i_1, by Lemma 2.2, constrained firms increase their continuation value ρ^*. The right hand side of Equation 2.14 is constant, and thus unconstrained firms must decrease their LT debt, therefore shortening debt maturity. This result clashes with the cross-sectional empirical evidence.[33]

2.4.5.2 Effect of Policy Rate Change with Yield-seeking Investors

The next proposition resumes the effects of a policy-rate change in our full model. First, for the sake of exposition we introduce a notion of strength for the yield seeking motive relative to risk aversion, namely the ratio $\varkappa = \alpha/\gamma \in [0, \infty)$.

Given that, we characterize the effect of a change in the interest rate as our main theoretical result in Proposition 2.2.

Proposition 2.2. *If the monetary authority decreases i_1:*

[33] A model with moral hazard frictions and risk-neutral investors, mimicking the original Holmström and Tirole (1998, 2000)'s framework, yields equally counterfactual predictions. Namely, in reaction to an expansionary interest rate shock by the FED, small firms lengthen debt maturity, whereas large ones do not adjust at all. This confirms the intuition that in models with simple credit frictions, constrained companies are in general more reactive to monetary policy, and the introduction of debt maturity just adds one credit margin along which this fact materializes (see discussion in section 2.1.1). A full derivation and description of this model is in the Theory Appendix.

1. *Unconstrained firms do not change* $\rho^*(A)$*. Moreover,* $\exists\,\phi > 0$ *such that if* $\varkappa > \phi$*, unconstrained firms increase their LT debt issuance, i.e.:*

$$\frac{\partial d_l(A)}{\partial i_1} < 0 \textit{ for } A \in (\bar{A}, I].$$

2. *If* $\varkappa \to \infty$*, then* $\dfrac{\partial d_l(A)}{\partial i_1} \to -\infty$ *for* $A \in (\bar{A}, I]$*; and* $\dfrac{\partial \frac{\partial d_s(A)}{\partial i_1}}{\partial \varkappa} = 0\ \forall A.$

3. *Constrained firms increase* $\rho^*(A)$ *and reduce short-term debt. LT debt is unchanged.*

Proof. See the Theory Appendix. □

The first condition on $\varkappa$ requires that yield-seeking motives are strong relative to risk-aversion. From the perspective of the model dynamics, this implies that demand functions for LT bonds react substantially to variations in the policy rate and that it is also relatively elastic. Under this condition, the upward demand shift due to a monetary interest rate loosening creates a mismatch with the existing supply of LT bonds. Constrained companies already issue LT debt at their limit, hence only large unconstrained firms can accommodate the demand shift.

The second result in Proposition 2.2 clarifies that, under $\varkappa$ sufficiently large, the effect on LT debt of unconstrained firms will dominate any adaptation of constrained firms in magnitude. In turn, this means that the model accommodates variations of the LT debt share of large companies which are arbitrarily larger than those of small companies, thereby matching the cross-sectional empirical evidence. Moreover, as both small and large companies adjust in the direction of lengthening the debt maturity, the model also aligns with aggregate-level empirical facts.

2.5 Empirical Evidence on the Model's Mechanism

Our model is consistent with the empirical results on the relation between the FED interest rate policy and corporate debt maturity presented in Section 2.3. Relative to a basic framework with credit frictions only, the key ingredient for aligning the model with the data is the inclusion of yield-oriented investors. In particular, a mechanism arises whereby, in the aftermath of a rate cut, there is a boost in demand for LT bonds by such yield-seeking investors. Next, credit frictions imply that unconstrained (i.e., large) companies can accommodate such upward demand shifts, whereas smaller firms are at their debt-limit already and hence cannot issue LT bonds.

We test such mechanism. In practical terms, this requires showing that, in reaction to a monetary rate loosening, yield-seeking is associated to increased holdings of corporate LT bonds and that large companies increase issuance of such debt-securities relatively more. However, in absence of further evidence, these adjustments may be driven by demand or supply motives. For dissecting their relative contribution, we additionally look at the price of the newly issued debt (i.e., the coupon rate). In this respect, to the extent that large companies also experience a relative stronger reduction in financing rates, it can be argued that their adjustments are relatively more demand-driven.[34]

In the rest of this section, we first present the empirical analysis of corporate bonds mutual funds holdings and then report results on corporate debt issuance.

2.5.1 Monetary Policy and Corporate Bond Holdings of Mutual Funds

We employ once again local projections to analyze the dynamic response of CBMF to monetary interest rate variations:

$$\Delta y_{m,t+h} = \beta_{1,h}\Delta EFFR_t + \beta_{2,h}HY_m + \beta_{3,h}\Delta EFFR_t * HY_m + \Gamma_h X_{m,t-1} + \mu_m + \mu_t + e_{m,t+h} \tag{2.15}$$

The dependent variable, $\Delta y_{m,t+h}$, is given by the growth between year-quarter $t-1$ and $t+h$ of fund m log volume of corporate bond holdings (or log portfolio's average weighted maturity).[35] Our coefficient of interest is $\beta_{3,h}$, loading the interaction between the quarterly EFFR variation, $\Delta EFFR_t$, and a dummy, HY_m, with value 1 if fund m is high-yield, our proxy for yield-seeking mutual funds. Importantly, $\beta_{3,h}$ captures

[34]For bonds, an inverse relation holds between price and interest rate, that is, bond prices go up when the interest (coupon) rate goes down. Hence, a joint increase in issuance and reduction in interest rate describes a demand-driven adjustment, as it also corresponds to a boost in bond prices.

[35]Ideally, we would like to look at the average weighted maturity of the fund's holdings of corporate bonds, rather than of the overall portfolio. Unfortunately, though, only the latter is available in the CRSP Survivor Bias Free dataset.

the relative response of high-yield mutual funds' portfolios (as compared to investment grade ones) to variation in the interest rate policy of the FED, therefore sizing impact of reach-for-yield motives on such relation. $X_{m,t-1}$ is a vector of time-varying fund-level controls, including the interaction of: i) HY_m with several macro-level controls; ii) other lagged fund characteristics (turnover ratio, expense ratio, log asset size and returns) with $\Delta EFFR_t$. Moreover, we augment the model with fund and year-quarter fixed effects (μ_m and μ_{yq}), respectively controlling for time-invariant heterogeneity at the level of the fund and for common shocks across all funds in a given year-quarter. $e_{m,t+h}$ is an error term, double-clustered at the fund and year-quarter level.

To start with, Figure 2.A.11 describes the relative response of HY-funds' corporate bonds holdings to a 25 b.p. cut to the EFFR. The effect is markedly positive at impact and peaks up 6 quarters after the shock. In particular, at impact, HY-funds increase their corporate bonds holdings by 2.87 p.p. (and by 6 p.p. one year after the shock). We find a similar relative jump at impact when using exogenous monetary policy shocks, though the effect appears much less persistent and reverts back to zero 2 quarters after the shock (Figure 2.A.12).

Next, Figure 2.A.13 confirms that, following an interest rate descent, and on top of buying corporate bonds, HY-funds additionally tilt their portfolio towards debt-securities with longer maturity. In detail, one year after the 25 b.p. reduction in the policy rate, portfolio maturity goes up by 2 p.p.. A similar effect emerges in Figure 2.A.14, where we exploit exogenous monetary policy shocks rather than the raw EFFR quarterly variations.

2.5.2 Monetary Policy and Corporate Bond Issuance

2.5.2.1 Extensive Margin: Frequency of Issuance

First, we check whether the likelihood of issuing new LT bonds is differently affected by interest rate changes across small and large companies. To this end, we resort to Model 2.2, i.e., a panel model saturated with firm and industry-time fixed effects and in which the interaction between the large-firm dummy and the interest rate change is horse-raced against different balance-sheet channels. The dependent variable is $\mathbb{1}(Issue)_{f,t+h}$, a dummy variable with value 1 if firm f issues LT bonds in year-quarter $t+h$ and with value 0 if it does not. Hence, at horizon h, the coefficient $\beta_{3,h}$ measures the relative difference in large companies' probability of issuing new debt as of year-quarter $t+h$, induced by a FED revision of the policy rate at t.

Results are displayed in Figure 2.A.15. The comparative response of large companies (vis-a-vis smaller ones) to a 25 b.p. cut in the EFFR is markedly more positive at impact, when it amounts to 28 b.p.. Such increase corresponds to an additional 4% jump relative to the average likelihood of issuing LT bonds as of time t. The effect extends over time, peaking 2 year-quarters after the rate change and vanishing in roughly one year and a half. To understand the absolute impact of monetary policy on the likelihood of issuing LT bonds, we report in Table 2.3 regressions with the dependent measured as of year-quarter t and $t+1$ and which do not include time-varying fixed effects (columns 1 and 4, respectively). Indeed, this exercise suggests that smaller companies also increase the likelihood of issuing bonds when the policy rate falls, however with a magnitude which is twice as small as that observed for larger firms.

We perform different robustness checks. In Figure 2.A.16 we replicate the analysis using the familiar exogenous monetary policy shocks by Gürkaynak, Sack, and Swanson (2005) and Jarocinski and Karadi (2020) and obtain similar results. As we exploit interest rate shocks - more prevalent before the last Global Financial Crisis - we additionally test whether our findings survive the exclusion from our sample of the observations from 2009 onward. Appendix Figures A8 and A9 suggest that the response is qualitatively comparable and, if anything, quantitatively larger.

2.5.2.2 Intensive Margin: Financing Costs

We aim to understand the cross-sectional differences in the reaction of financing costs to variations in the FED's interest rate policy. One issue with this analysis is that only a tiny subset of companies ever issue bonds in two consecutive quarters. In operative terms, this means that if we were to apply a first-differenced model, we would be left with very few observations and, ultimately, a meaningless cross-sectional comparisons across a very small set of companies (highly skewed towards large firms).

Hence, we rather resort to the following model in levels:

$$\begin{aligned} y_{f,t+h} =& \beta_{1,h} EFFR_t + \beta_{2,h} Large_{f,t-1} + \beta_{3,h} Large_{f,t-1} * EFFR_t + \\ &+ \Phi_h X_{f,t-1} + \mu_f + \mu_{s,t} + \xi_{f,t+h} \end{aligned} \tag{2.16}$$

As dependent variable, we use $Coupon_{f,t+h}$, i.e., the coupon rate on firm f's newly issued LT bonds at year-quarter $t + h$. This is regressed against the current level of the policy rate, $EFFR_t$, interacted with the dummy for large firms, $Large_{f,t-1}$. The vector of controls $X_{f,t-1}$ aligns to the previously used empirical models.[36] As usual, we augment the model with firm and industry*year-quarter fixed effects (μ_f and $\mu_{s,t+h}$) and double-cluster the error term, $\xi_{f,t+h}$, accordingly.

Figure 2.A.17 reports the estimated coefficients $\beta_{3,h}$, calibrated to a 1 p.p. lower EFFR-level. Clearly, when the EFFR is lower, large firms' coupon rate is also lower (as compared to that of smaller companies). Quantitatively speaking, a 1 p.p. looser EFFR grants big companies an additional reduction in the coupon rate by roughly 10 b.p.., resulting in a further 1.6% cut in financing costs relatively to their average level. Table 2.4 describes the background regressions for $h = 0, 1$, also shown with different set of fixed effects allowing to evaluate the absolute response of the coupon rate. In columns 1 and 3, in fact, we report a regression which, differently from the model in Equation 2.16, excludes time-varying fixed effects. Both in year-quarter t and $t + 1$, the relation between the coupon rate and the EFFR is generally positive across smaller firms, too. Put differently, based on this exercise, for all companies the coupon rate goes down when the EFFR is lower, but the effect is stronger among large firms.

As clear from Table 2.4, there is a significant loss of observations due to the application of (3-digit SIC)-industry*year-quarter fixed effects. Bond issuance at the firm-level is indeed quite lumpy across time. Therefore, narrowing the comparison within granular industries implies the loss of many (within-industry) singletons. One additional problem is that in this framework the within-industry cross-sectional comparison will comprehend few firms. Hence, for robustness purposes, we apply looser industry definitions and check that our findings go through. In Table 2.5, in columns 1-4, we replicate our analysis, but this time comparing all companies with each other. In this case, the $Large_{f,t-1}$ dummy captures those companies in the top-quartile of the entire sample of NFCs in the US stock market. The result that large companies' coupon rates descent along with the EFFR still goes through. A similar pattern emerges when applying increasingly more granular industry-definition (sectoral-level in columns 5-8 and 2-digit

[36]That is, $X_{f,t-1}$ includes the full interaction of: $Large_{f,t-1}$ and macro controls; firm-level controls and $EFFR_t$. Relatively to previous models, we report few macro-controls (term-spread, corporate-spread and share of Treasuries with maturity above 20-year) in levels rather than in first-differences.

SIC industry in columns 9-12).

2.6 Conclusions

Firms' resilience to shocks crucially depends on the maturity structure of corporate debt. Hence, understanding whether and how monetary policy affects firms' debt maturity is key for gauging the implications for firms' risk of central banks' policies.

This paper provides novel empirical evidence on the relation between the interest rate policy by the FED and the maturity structure of the US corporate (non-financial) sector. Our robust findings suggest that, following a policy rate cut, firms lengthen debt maturity. The effect is entirely driven by very large firms, whereas smaller companies' debt maturity is generally not responsive to monetary policy. Existing theoretical frameworks do not provide an adequate explanation for these findings. Hence, we build a model combining credit frictions due to firms' moral hazard and yield-oriented investors - who increase the demand for long-term bonds when the interest rate goes down. Only large and unconstrained companies can accommodate such upward shift in demand, so that the model aligns with the empirical evidence. We bring the model mechanism to the data and find supportive evidence. As a matter of fact, following a policy rate decline: i) relatively more yield-oriented mutual funds increase their holdings of corporate bonds and tilt their portfolio towards longer-term debt securities; ii) large firms issue more debt and at lower rates, indicating that their adjustment is demand driven.

Ultimately, our work highlights how monetary policy impacts the maturity profile of very large corporations, whose dynamics have significant consequences for the business cycle (Crouzet and Mehrotra forthcoming). An important open question - left for future research - is whether the documented interaction between monetary policy and corporate debt maturity has implications for business cycle and systemic risk.

Appendix

2.A Figures

Figure 2.A.1: % of LT-Debt - Aggregate Level

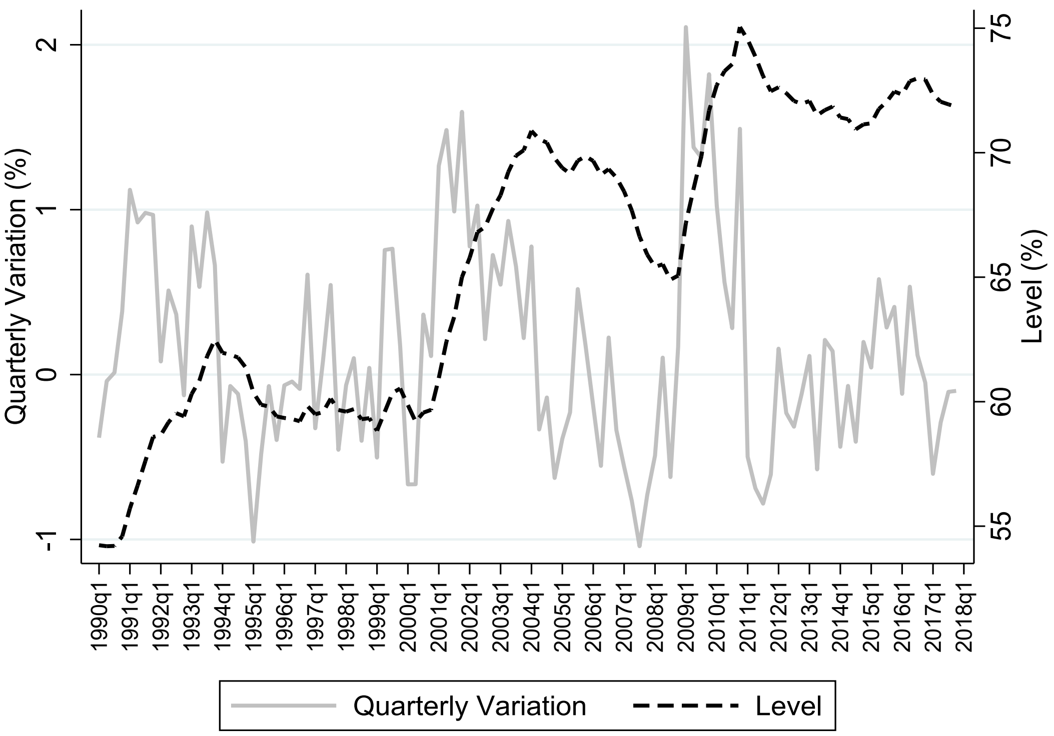

This figure shows the evolution of the aggregate share of LT debt (i.e. with outstanding maturity above 1 year). The black dashed line reports the series in levels, the grey one in first-differences. Following Greenwood, Hanson, and Stein (2010), LT debt is defined as the sum of corporate bonds and mortgages and industrial revenues. The remaining short-term corporate debt is proxied by the sum of short-term loans (and advances) and commercial paper.

Figure 2.A.2: Measures of Changes in the Policy Rate

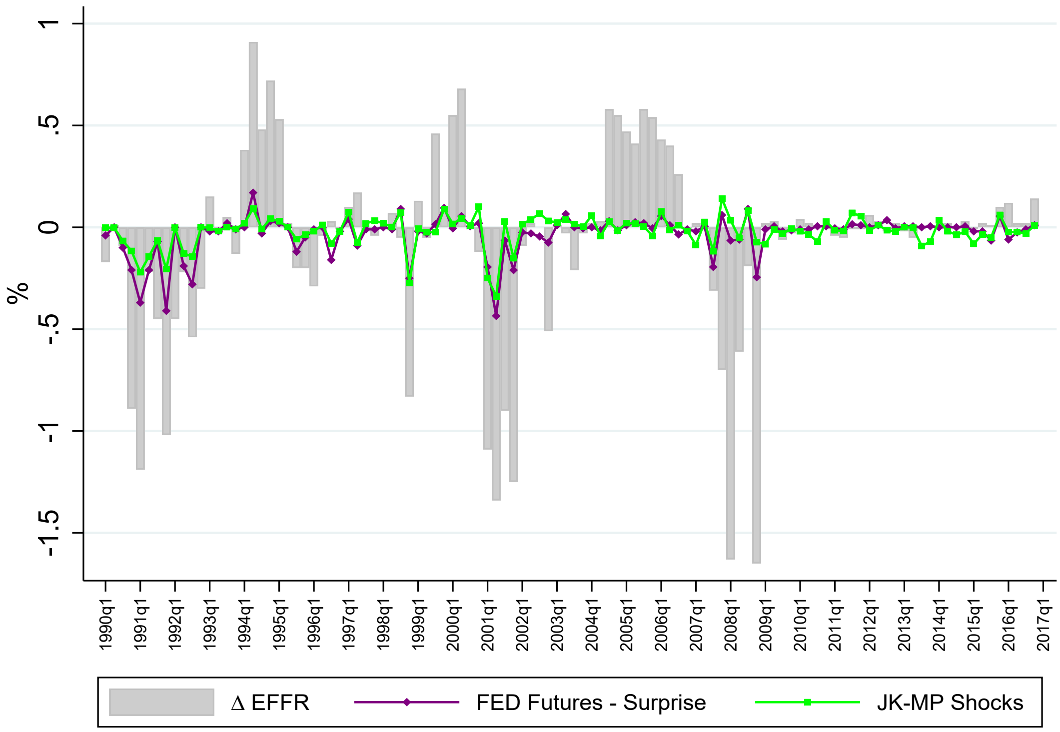

The grey bars show the quarterly variation of the Effective FED Funds Rate. The purple (green) solid line, connected by diamonds (squares), reports the Gürkaynak, Sack, and Swanson (2005) (Jarocinski and Karadi 2020) monetary policy shocks.

Figure 2.A.3: Distribution of % of LT-Debt across Firms - Sorted by Asset-Size Quartiles

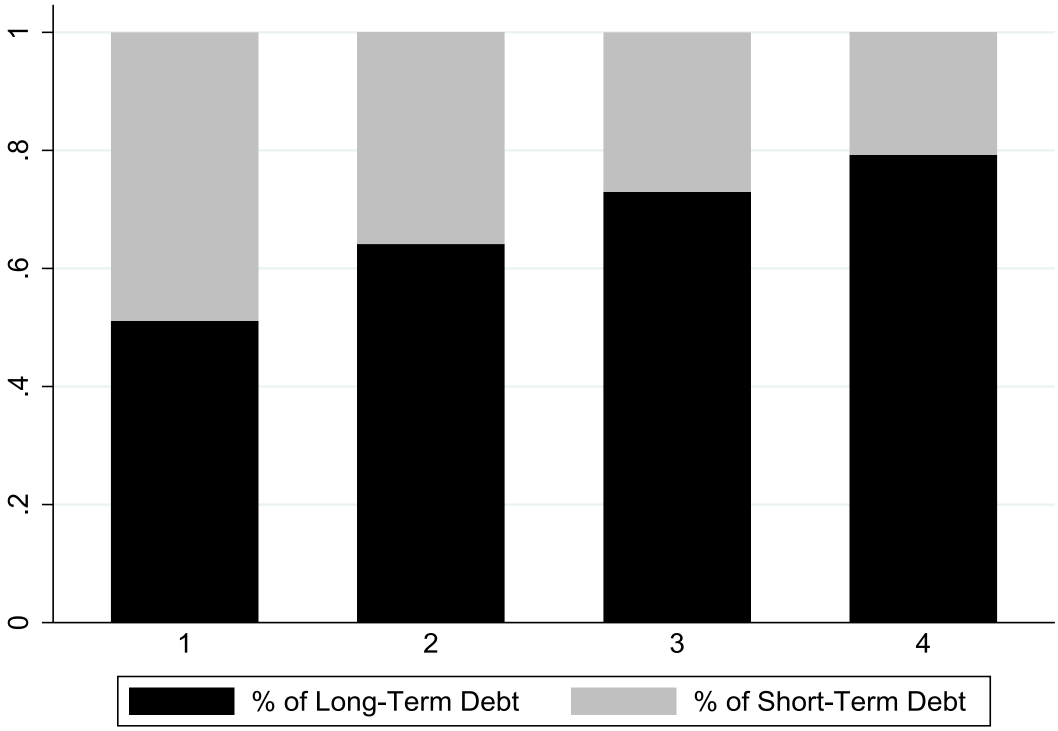

This chart shows the average % of LT-Debt across different groups of companies (black bars); the complement to 100% gives the average % of ST-Debt (grey bars). Firms are sorted according to quartiles of their 3-digit SIC industry asset-size distribution. Quartiles are reported on the x-axis.

Figure 2.A.4: Bond Issuance over Time

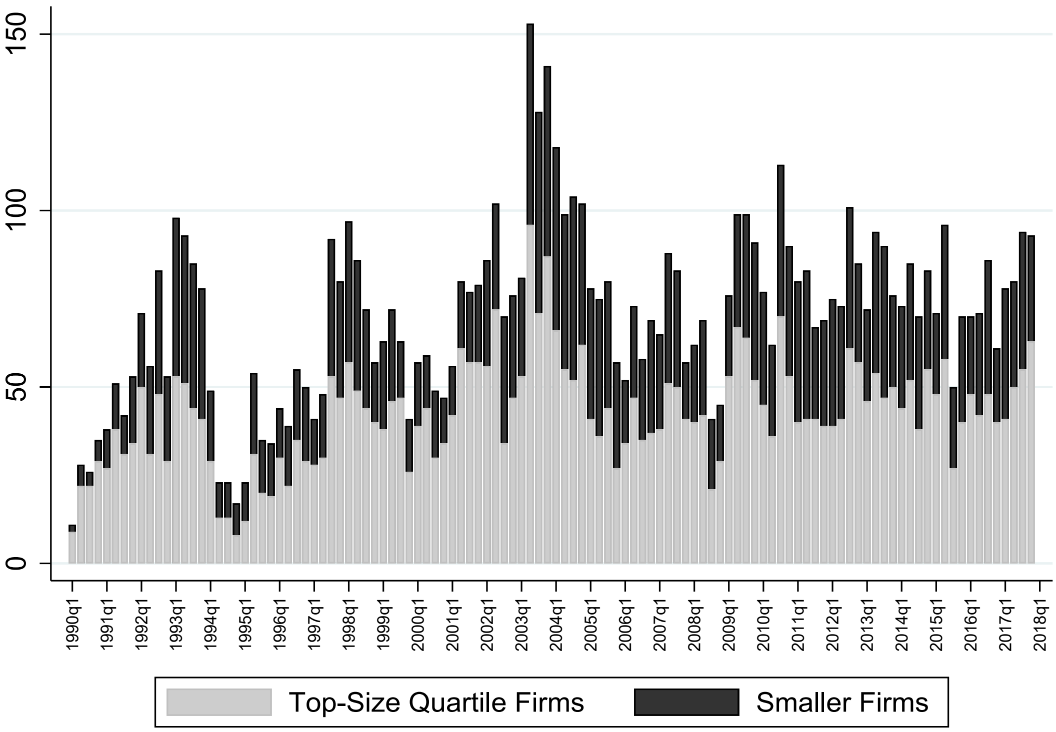

This chart shows the number of bond issuances over time. Companies are sorted based on their (3-digit SIC) asset-size distribution. The grey bars report the number of bond issuances by companies in the upper quartile; the black bars by all other companies. The vertical sum of the grey and black bars provides the total number of bond issuances per year-quarter.

Figure 2.A.5: Monetary Policy and Debt Maturity Structure: Aggregate Response

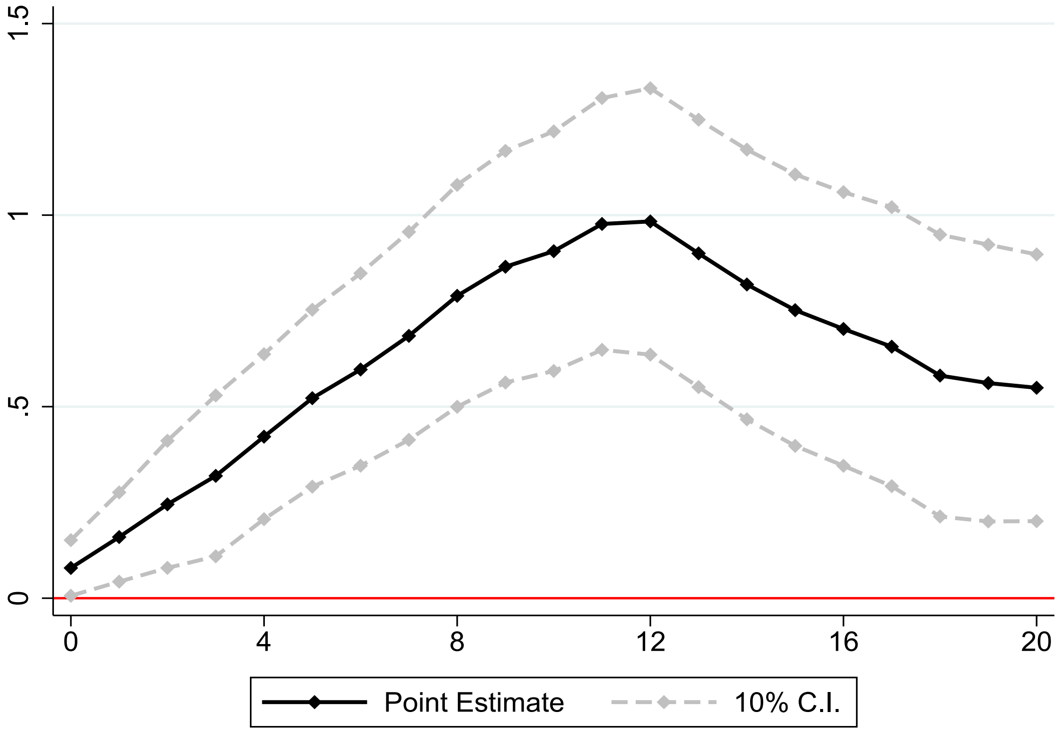

This figure depicts the response of the aggregate-level share of LT debt to a 25 b.p. cut in the EFFR. Formally, it shows the coefficients $\beta_{1,h}$ from the estimation of the following local projection model:

$$\Delta_h y_{t+h} = \beta_{1,h} \Delta EFFR_t + MacroControls_{t-1} + u_{t,h}$$

The dependent variable, $\Delta_h y_{t+h}$, represents the growth of the LT-debt share (expressed in p.p.) from year-quarter $t-1$ to year-quarter $t+h$. $\Delta EFFR_t$ is the quarterly EFFR change. $MacroControls_{t-1}$ is a vector of lagged macroeconomic controls, including annual GDP growth and inflation rate, a dummy for recessions, the quarterly variation in the share of LT treasuries, in the corporate spread and in the 10y-3m term-spread. $u_{t,h}$ is a robust error-term. The x-axis is measured in terms of quarters after the shock. The black solid line reports the point estimates for $\beta_{1,h}$; the dashed grey line the 10% confidence intervals.

Figure 2.A.6: Monetary Policy and Debt Maturity Structure: Aggregate Response using Exogenous Shocks

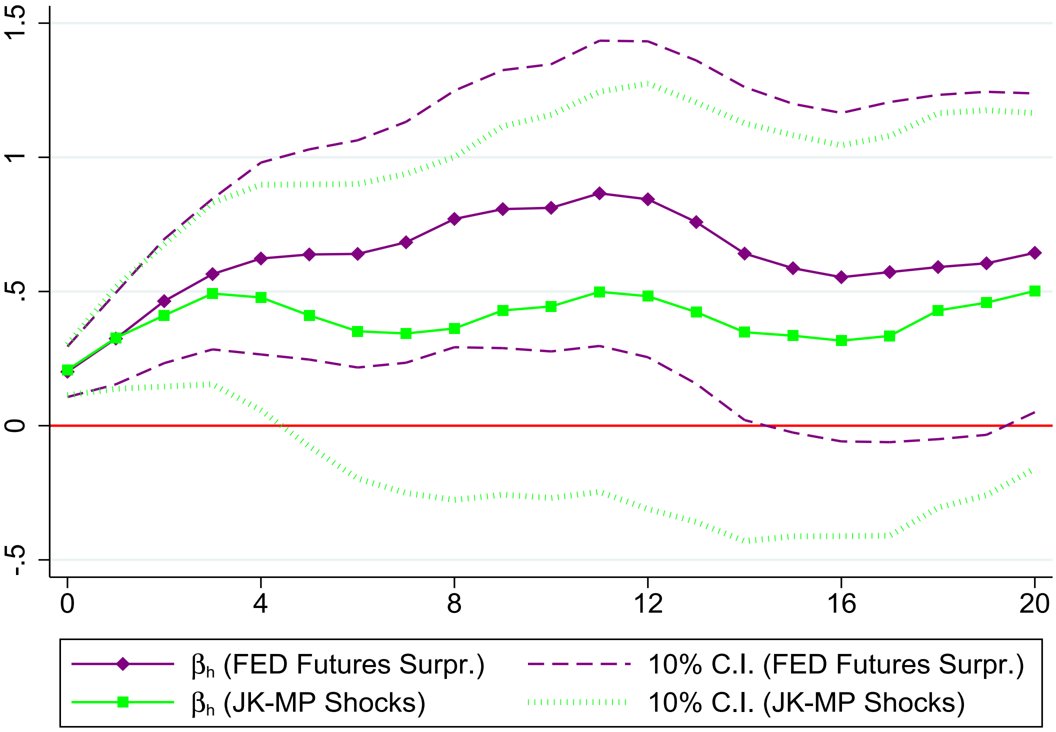

This figure depicts the response of the aggregate-level share of LT debt to a 1 s.d. reduction in the monetary policy shock. Formally, it shows the coefficients $\beta_{1,h}$ from the estimation of the following local projection model:

$$\Delta_h y_{t+h} = \beta_{1,h}\varepsilon_t^{mp} + MacroControls_{t-1} + u_{t,h}$$

The dependent variable, $\Delta_h y_{t+h}$, represents the growth of the LT-debt share (expressed in p.p.) from year-quarter $t-1$ to year-quarter $t+h$. ε_t^{mp} is an exogenous monetary policy shock, gathered from either Gürkaynak, Sack, and Swanson (2005) or from Jarocinski and Karadi (2020). $MacroControls_{t-1}$ is a vector of lagged macroeconomic controls, including annual GDP growth and inflation rate, a dummy for recessions, the quarterly variation in the share of LT treasuries, in the corporate spread and in the 10y-3m term-spread. $u_{t,h}$ is a robust error-term. The x-axis is measured in terms of quarters after the shock. The purple (green) solid line, connected by diamonds (squares), reports the point estimates for $\beta_{1,h}$ using the Gürkaynak, Sack, and Swanson (2005) (Jarocinski and Karadi 2020) shocks and the dashed (dotted) purple (green) line the respective 10% confidence intervals.

Figure 2.A.7: Monetary Policy and Debt Maturity Structure - Relative response of Large firms

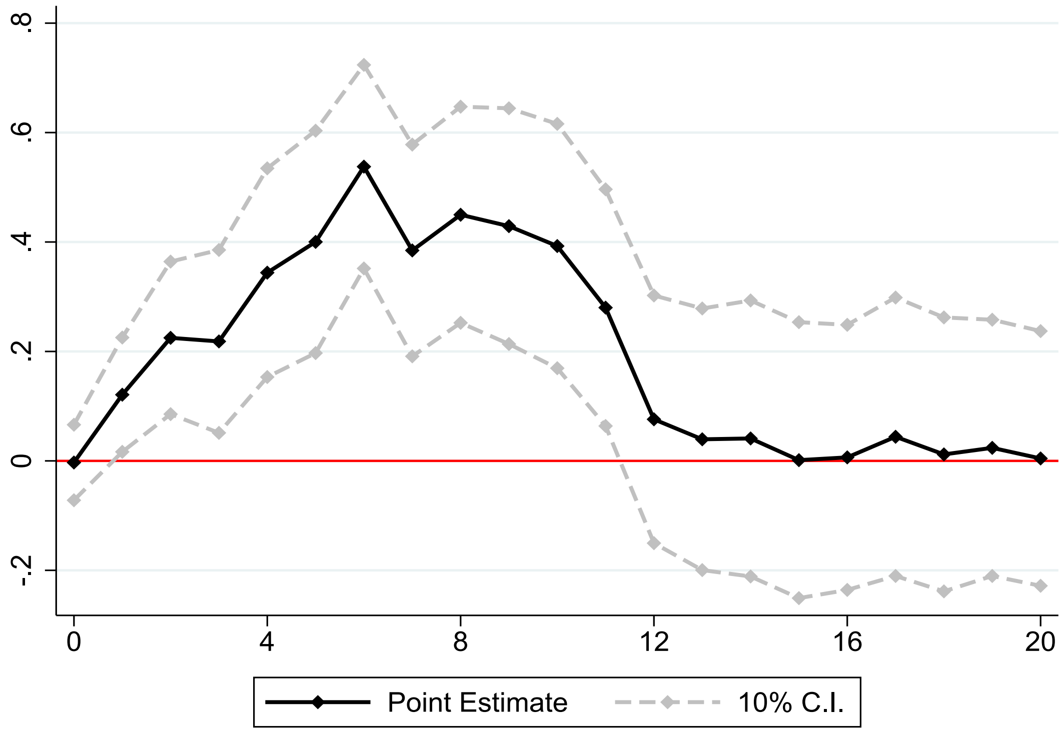

This figure depicts the relative response of companies in the top quartile of the 3-digit SIC industry asset-size distribution to a 25 b.p. cut in the EFFR (as compared to smaller firms). Formally, it shows the coefficients $\beta_{3,h}$ from the estimation of the following local projection model:

$$\Delta_h y_{f,t+h} = \beta_{1,h}\Delta EFFR_t + \beta_{2,h}Large_{f,t-1} + \beta_{3,h}Large_{f,t-1} * \Delta EFFR_t + X_{f,t-1} + \mu_f + \mu_{s,t} + u_{t,h}$$

The dependent variable, $\Delta_h y_{f,t+h}$, represents the growth of the share of LT-debt - expressed in p.p. - from year-quarter $t-1$ to year-quarter $t+h$. $\Delta EFFR_t$ is the quarterly change in the EFFR. $Large_{f,t-1}$ is a dummy variable with value 1 a firm is in the top-quartile of the industry-wide asset-size distribution, and 0 otherwise. $X_{f,t-1}$ is a vector of controls, including the interaction of $Large_{f,t-1}$ with several macro-controls and of $\Delta EFFR_t$ with other firm characteristics, namely lagged sales growth, leverage and liquid assets. μ_f and $\mu_{s,t}$ represent vectors of firm and industry*year-quarter fixed effects, respectively. $u_{t,h}$ is an error-term, double-clustered at the firm and industry*year-quarter level. The x-axis is measured in terms of quarters after the shock. The black solid line reports the point estimates for $\beta_{3,h}$; the dashed grey line the 10% confidence intervals.

Figure 2.A.8: Monetary Policy and Debt Maturity Structure - Relative response of Large firms using Exogenous Shocks

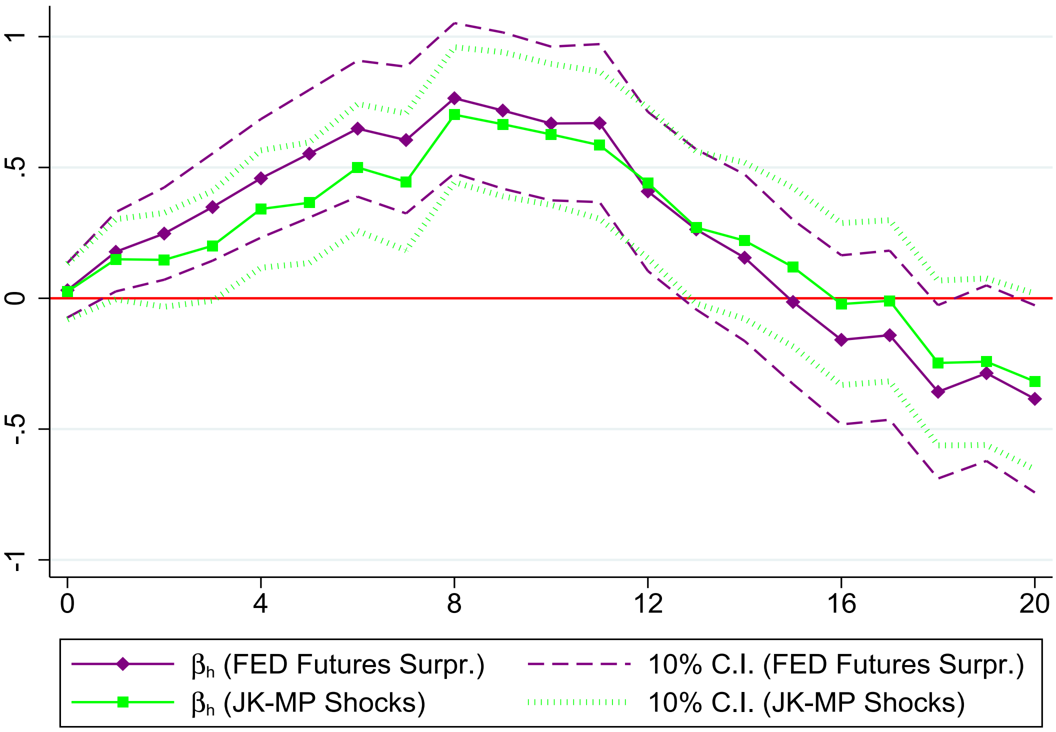

This figure depicts the relative response of companies in the top quartile of the 3-digit SIC industry asset-size distribution to a 1 s.d. b.p.reduction in monetary policy shock (as compared to smaller firms). Formally, it shows the coefficients $\beta_{3,h}$ from the estimation of the following local projection model:

$$\Delta_h y_{f,t+h} = \beta_{1,h}\varepsilon_t^{mp} + \beta_{2,h}Large_{f,t-1} + \beta_{3,h}Large_{f,t-1} * \varepsilon^{mp} + X_{f,t-1} + \mu_f + \mu_{s,t} + u_{t,h}$$

The dependent variable, $\Delta_h y_{f,t+h}$, represents the growth of the share of LT-debt - expressed in p.p. - from year-quarter $t-1$ to year-quarter $t+h$. ε_t^{mp} is an exogenous monetary policy shock, derived either from Gürkaynak, Sack, and Swanson (2005) or from Jarocinski and Karadi (2020). $Large_{f,t-1}$ is a dummy variable with value 1 a firm is in the top-quartile of the industry-wide asset-size distribution, and 0 otherwise. $X_{f,t-1}$ is a vector of controls, including the interaction of $Large_{f,t-1}$ with several macro-controls and of ε_t^{mp} with other firm characteristics, namely lagged sales growth, leverage and liquid assets. μ_f and $\mu_{s,t}$ represent vectors of firm and industry*year-quarter fixed effects, respectively. $u_{t,h}$ is an error-term, double-clustered at the firm and industry*year-quarter level. The x-axis is measured in terms of quarters after the shock. The purple (green) solid line, connected by diamonds (squares), reports the point estimates for $\beta_{1,h}$ using the Gürkaynak, Sack, and Swanson (2005) (Jarocinski and Karadi 2020) shocks and the dashed (dotted) purple (green) line the respective 10% confidence intervals.

Figure 2.A.9: Monetary Policy and Debt Maturity Structure - Absolute firm-level response

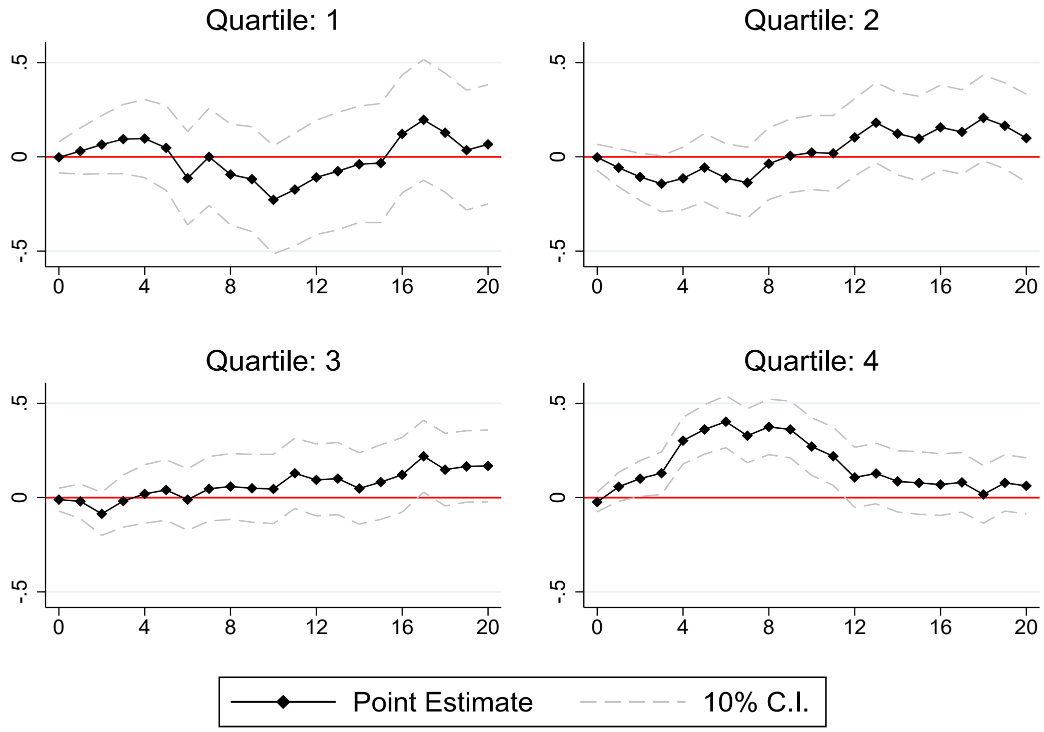

This figure depicts the absolute response of companies in different size quartiles of the 3-digit SIC industry asset-size distribution to a 25 b.p. cut in the EFFR. Formally, it shows the coefficients $\beta_{1,h}$ from the estimation of the following local projection model:

$$\Delta_h y_{f,t+h} = \beta_{1,h} \Delta EFFR_t + MacroControls_{t-1} + X_{f,t-1} + \mu_f + \mu_{s,t} + u_{t,h}$$

The dependent variable, $\Delta_h y_{f,t+h}$, represents the growth of the LT-debt share (expressed in p.p.) from year-quarter $t-1$ to year-quarter $t+h$. $\Delta EFFR_t$ is the quarterly EFFR change. $MacroControls_{t-1}$ is a vector of lagged macroeconomic controls, including annual GDP growth and inflation rate, a dummy for recessions, the quarterly variation in the share of LT treasuries, in the corporate spread and in the 10y-3m term-spread. $X_{f,t-1}$ is a vector of firm-level controls, including lagged sales growth, leverage and liquid assets. μ_f is a vector of firm fixed effects. $u_{t,h}$ is an error-term, double-clustered at the firm and industry*year-quarter level. The x-axis is measured in terms of quarters after the shock. The black solid line reports the point estimates for $\beta_{1,h}$; the dashed grey line the 10% confidence intervals.

Figure 2.A.10: Monetary Policy and Debt Maturity Structure - Absolute firm-level response using Exogenous Shocks

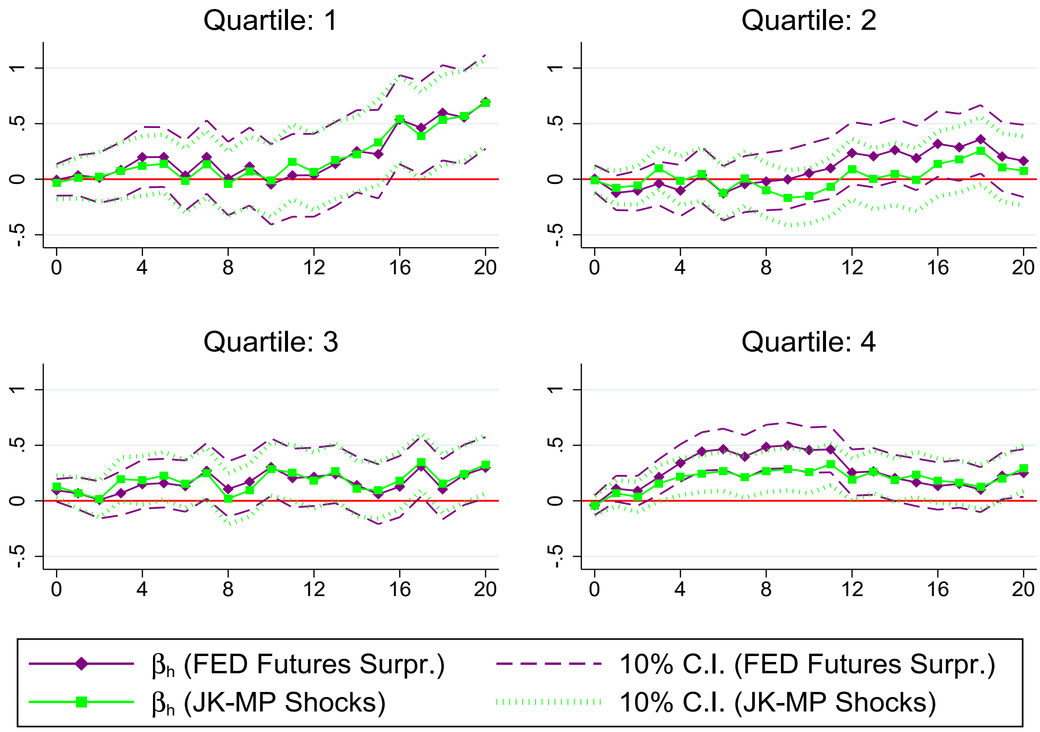

This figure depicts the absolute response of companies in different size quartiles of the 3-digit SIC industry asset-size distribution to a 1 s.d. reduction in monetary policy shock. Formally, it shows the coefficients $\beta_{1,h}$ from the estimation of the following local projection model:

$$\Delta_h y_{f,t+h} = \beta_{1,h}\varepsilon_t^{mp} + MacroControls_{t-1} + X_{f,t-1} + \mu_f + \mu_{s,t} + u_{t,h}$$

The dependent variable, $\Delta_h y_{f,t+h}$, represents the growth of the LT-debt share (expressed in p.p.) from year-quarter $t-1$ to year-quarter $t+h$. ε_t^{mp} is an exogenous monetary policy shock, gathered from either Gürkaynak, Sack, and Swanson (2005) or from Jarocinski and Karadi (2020). $MacroControls_{t-1}$ is a vector of lagged macroeconomic controls, including annual GDP growth and inflation rate, a dummy for recessions, the quarterly variation in the share of LT treasuries, in the corporate spread and in the 10y-3m term-spread. $X_{f,t-1}$ is a vector of firm-level controls, including lagged sales growth, leverage and liquid assets. μ_f is a vector of firm fixed effects. $u_{t,h}$ is an error-term, double-clustered at the firm and industry*year-quarter level. The x-axis is measured in terms of quarters after the shock. The purple (green) solid line, connected by diamonds (squares), reports the point estimates for $\beta_{1,h}$ using the Gürkaynak, Sack, and Swanson (2005) (Jarocinski and Karadi 2020) shocks and the dashed (dotted) purple (green) line the respective 10% confidence intervals.

Figure 2.A.11: Monetary policy and CBMFs' corporate bonds holdings

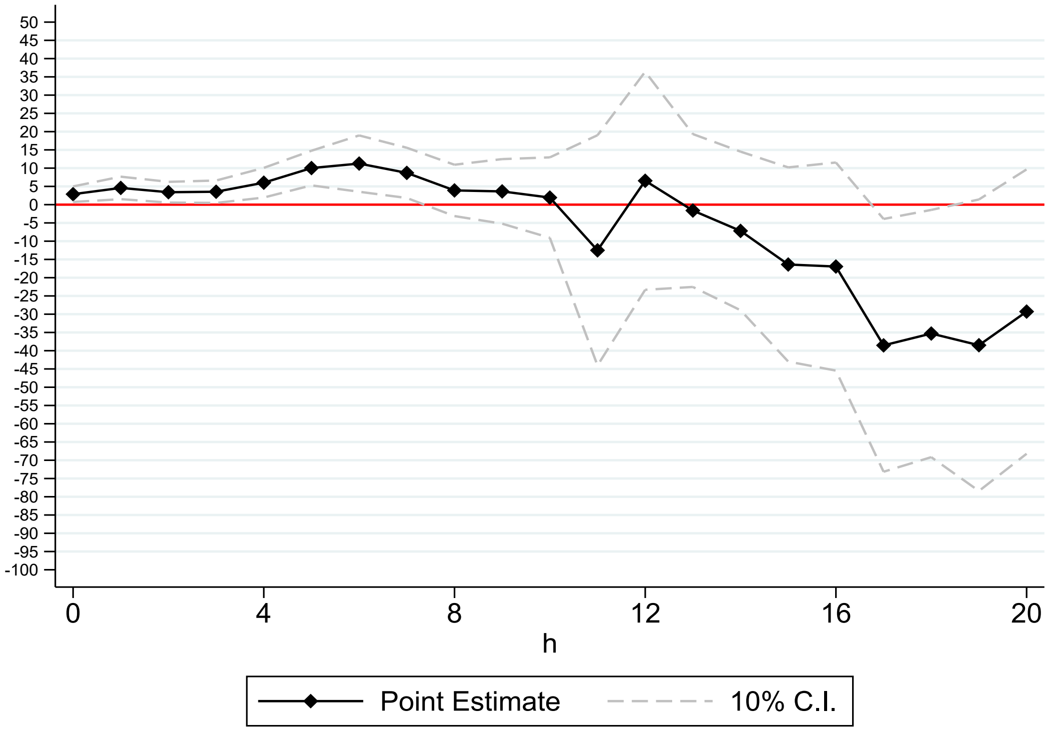

This figure shows the relative response of High-Yield (HY) corporate bonds mutual funds to a 25 b.p. cut in the EFFR (as compared to Investment-Grade funds). Formally, it shows the coefficients $\beta_{3,h}$ from the estimation of the following local projection model:

$$\Delta y_{m,t+h} = \beta_{1,h}\Delta EFFR_t + \beta_{2,h}HY_m + \beta_{3,h}\Delta EFFR_t * HY_m + \Gamma_h X_{m,t-1} + \mu_m + \mu_t + e_{m,t+h}$$

The dependent variable, $\Delta y_{m,t+h}$, is given by the growth between year-quarter $t-1$ and $t+h$ of fund m log volume of corporate bond holdings, expressed in p.p.. $\Delta EFFR_t$ gives the quarterly EFFR variation. HY_m is a dummy with value 1 if fund m is high-yield and with value 0 otherwise. $X_{m,t-1}$ is a vector of time-varying fund-level controls. It includes the full interaction of HY_m with a vector of macro-level controls, namely annual GDP growth and inflation rate, a dummy for recessions, the quarterly variation in the share of LT treasuries, in the corporate spread and in the 10y-3m term-spread. Moreover, we controls for the full interaction of other lagged fund characteristics (turnover ratio, expense ratio, log asset size and returns) with $\Delta EFFR_t$. μ_m and μ_{yq} are fund- and year-quarter fixed effects. $e_{m,t+h}$ is an error term, double-clustered at the fund and year-quarter level. The black solid line reports the point estimates for $\beta_{3,h}$; the dashed grey line the 10% confidence intervals.

Figure 2.A.12: Monetary policy and CBMFs' corporate bond holdings - Using exogenous shocks

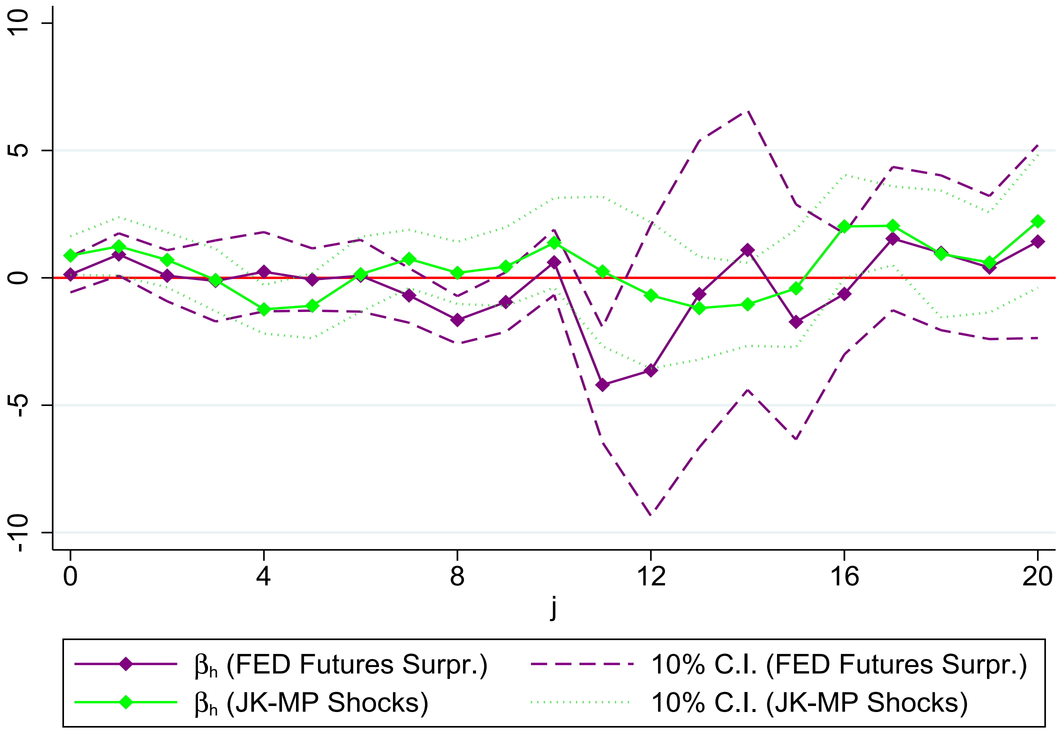

This figure shows the relative response of High-Yield (HY) corporate bonds mutual funds to a 25 b.p. cut in the EFFR (as compared to Investment-Grade funds). Formally, it shows the coefficients $\beta_{3,h}$ from the estimation of the following local projection model:

$$\Delta y_{m,t+h} = \beta_{1,h}\varepsilon_t^{mp} + \beta_{2,h}HY_m + \beta_{3,h}\varepsilon_t^{mp} * HY_m + \Gamma_h X_{m,t-1} + \mu_m + \mu_t + e_{m,t+h}$$

The dependent variable, $\Delta y_{m,t+h}$, is given by the growth between year-quarter $t-1$ and $t+h$ of fund m log corporate bond holdings, expressed in p.p.. ε_t^{mp} is an exogenous monetary policy shock, gathered from either Gürkaynak, Sack, and Swanson (2005) or from Jarocinski and Karadi (2020). HY_m is a dummy with value 1 if fund m is high-yield and with value 0 otherwise. $X_{m,t-1}$ is a vector of time-varying fund-level controls. It includes the full interaction of HY_m with a vector of macro-level controls, namely annual GDP growth and inflation rate, a dummy for recessions, the quarterly variation in the share of LT treasuries, in the corporate spread and in the 10y-3m term-spread. Moreover, we controls for the full interaction of other lagged fund characteristics (turnover ratio, expense ratio, log asset size and returns) with ε_t^{mp}. μ_m and μ_{yq} are fund- and year-quarter fixed effects. $e_{m,t+h}$ is an error term, double-clustered at the fund and year-quarter level. The purple (green) solid line, connected by diamonds (squares), reports the point estimates for $\beta_{1,h}$ using the Gürkaynak, Sack, and Swanson (2005) (Jarocinski and Karadi 2020) shocks and the dashed (dotted) purple (green) line the respective 10% confidence intervals.

Figure 2.A.13: Monetary policy and CBMFs' portfolio maturity

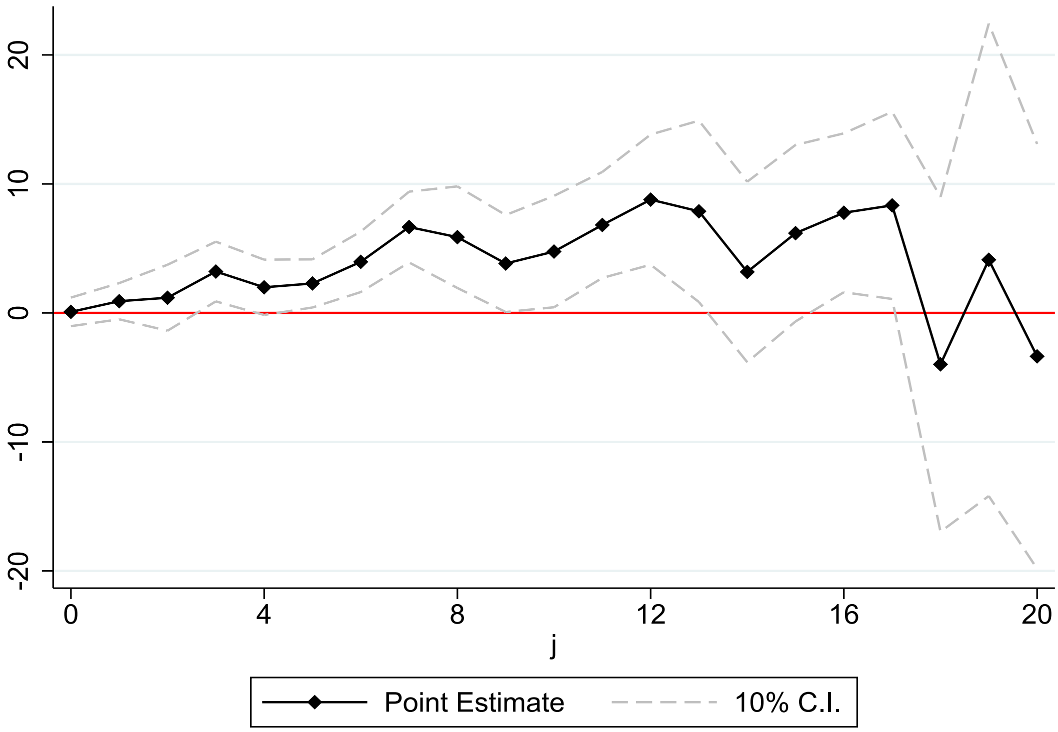

This figure shows the relative response of High-Yield (HY) corporate bonds mutual funds to a 25 b.p. cut in the EFFR (as compared to Investment-Grade funds). Formally, it shows the coefficients $\beta_{3,h}$ from the estimation of the following local projection model:

$$\Delta y_{m,t+h} = \beta_{1,h}\Delta EFFR_t + \beta_{2,h}HY_m + \beta_{3,h}\Delta EFFR_t * HY_m + \Gamma_h X_{m,t-1} + \mu_m + \mu_t + e_{m,t+h}$$

The dependent variable, $\Delta y_{m,t+h}$, is given by the growth between year-quarter $t-1$ and $t+h$ of fund m log (weighted) average maturity, expressed in p.p.. $\Delta EFFR_t$ gives the quarterly EFFR variation. HY_m is a dummy with value 1 if fund m is high-yield and with value 0 otherwise. $X_{m,t-1}$ is a vector of time-varying fund-level controls. It includes the full interaction of HY_m with a vector of macro-level controls, namely annual GDP growth and inflation rate, a dummy for recessions, the quarterly variation in the share of LT treasuries, in the corporate spread and in the 10y-3m term-spread. Moreover, we controls for the full interaction of other lagged fund characteristics (turnover ratio, expense ratio, log asset size and returns) with $\Delta EFFR_t$. μ_m and μ_{yq} are fund- and year-quarter fixed effects. $e_{m,t+h}$ is an error term, double-clustered at the fund and year-quarter level. The black solid line reports the point estimates for $\beta_{3,h}$; the dashed grey line the 10% confidence intervals.

Figure 2.A.14: Monetary policy and CBMFs' portfolio maturity - Using exogenous shocks

This figure shows the relative response of High-Yield (HY) corporate bonds mutual funds to a 25 b.p. cut in the EFFR (as compared to Investment-Grade funds). Formally, it shows the coefficients $\beta_{3,h}$ from the estimation of the following local projection model:

$$\Delta y_{m,t+h} = \beta_{1,h}\varepsilon_t^{mp} + \beta_{2,h}HY_m + \beta_{3,h}\varepsilon_t^{mp} * HY_m + \Gamma_h X_{m,t-1} + \mu_m + \mu_t + e_{m,t+h}$$

The dependent variable, $\Delta y_{m,t+h}$, is given by the growth between year-quarter $t-1$ and $t+h$ of fund m log (weighted) average maturity, expressed in p.p.. ε_t^{mp} is an exogenous monetary policy shock, gathered from either Gürkaynak, Sack, and Swanson (2005) or from Jarocinski and Karadi (2020). HY_m is a dummy with value 1 if fund m is high-yield and with value 0 otherwise. $X_{m,t-1}$ is a vector of time-varying fund-level controls. It includes the full interaction of HY_m with a vector of macro-level controls, namely annual GDP growth and inflation rate, a dummy for recessions, the quarterly variation in the share of LT treasuries, in the corporate spread and in the 10y-3m term-spread. Moreover, we controls for the full interaction of other lagged fund characteristics (turnover ratio, expense ratio, log asset size and returns) with ε_t^{mp}. μ_m and μ_{yq} are fund- and year-quarter fixed effects. $e_{m,t+h}$ is an error term, double-clustered at the fund and year-quarter level. The purple (green) solid line, connected by diamonds (squares), reports the point estimates for $\beta_{1,h}$ using the Gürkaynak, Sack, and Swanson (2005) (Jarocinski and Karadi 2020) shocks and the dashed (dotted) purple (green) line the respective 10% confidence intervals.

Figure 2.A.15: Monetary policy and likelihood of issuing LT bonds - Relative response of large companies

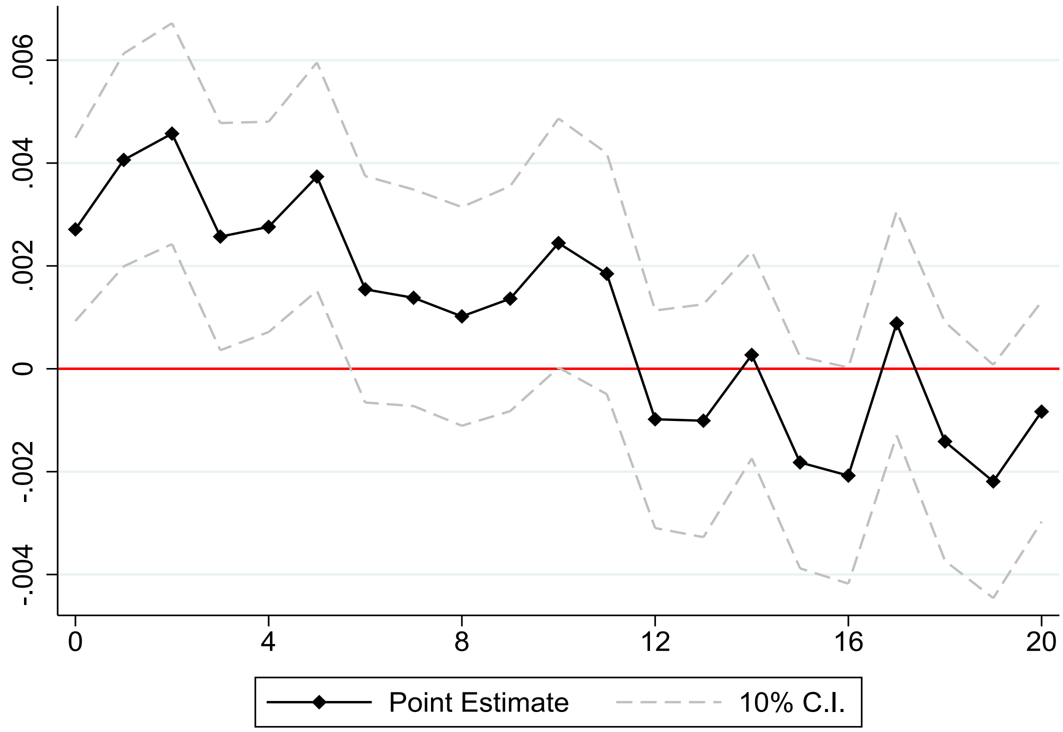

This figure depicts the relative response of companies in the top quartile of the 3-digit SIC industry asset-size distribution to a 25 b.p. cut in the EFFR (as compared to smaller firms). Formally, it shows the coefficients $\beta_{3,h}$ from the estimation of the following local projection model:

$$y_{f,t+h} = \beta_{1,h}\Delta EFFR_t + \beta_{2,h}Large_{f,t-1} + \beta_{3,h}Large_{f,t-1} * \Delta EFFR_t + X_{f,t-1} + \mu_f + \mu_{s,t} + u_{t,h}$$

The dependent variable, $y_{f,t+h}$, is a dummy with value 1 if firms f issues LT bonds in year-quarter $t + h$ and with value 0 if it does not. $\Delta EFFR_t$ is the quarterly change in the EFFR. $Large_{f,t-1}$ is a dummy variable with value 1 a firm is in the top-quartile of the industry-wide asset-size distribution, and 0 otherwise. $X_{f,t-1}$ is a vector of controls, including the interaction of $Large_{f,t-1}$ with several macro-controls and of $\Delta EFFR_t$ with other firm characteristics, namely lagged sales growth, leverage and liquid assets. μ_f and $\mu_{s,t}$ represent vectors of firm and industry*year-quarter fixed effects, respectively. $u_{t,h}$ is an error-term, double-clustered at the firm and industry*year-quarter level. The x-axis is measured in terms of quarters after the shock. The black solid line reports the point estimates for $\beta_{3,h}$; the dashed grey line the 10% confidence intervals.

Figure 2.A.16: Monetary policy and likelihood of issuing LT bonds - Relative response of large companies using exogenous shocks

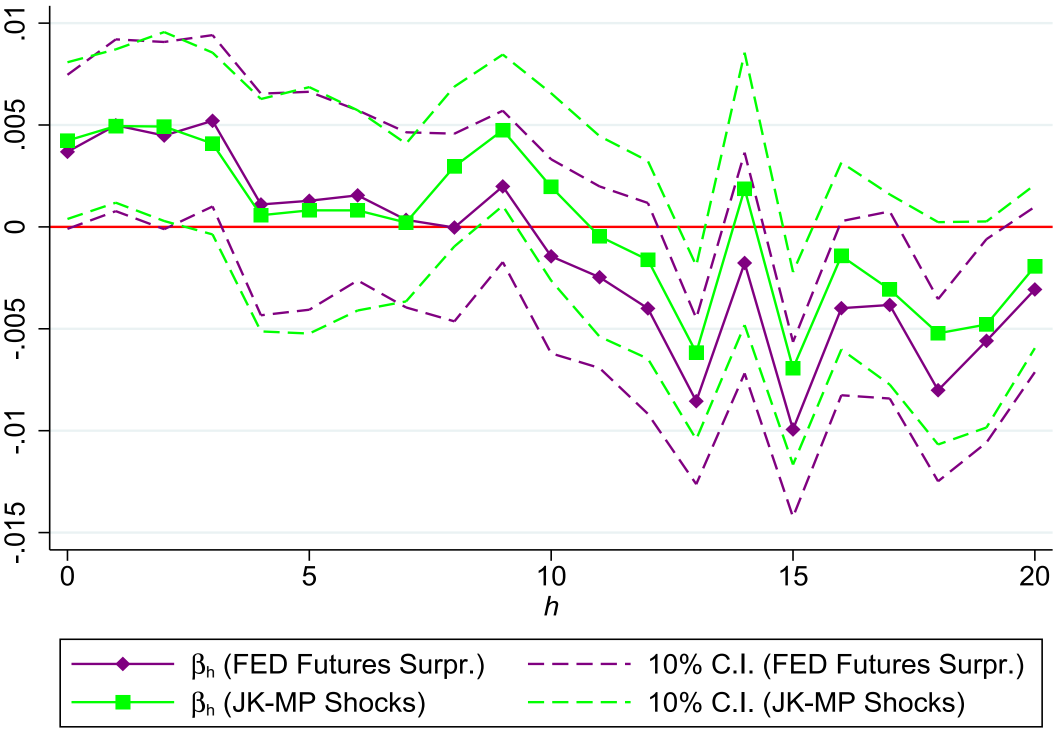

This figure depicts the relative response of companies in the top quartile of the 3-digit SIC industry asset-size distribution to a 25 b.p. cut in the EFFR (as compared to smaller firms). Formally, it shows the coefficients $\beta_{3,h}$ from the estimation of the following local projection model:

$$y_{f,t+h} = \beta_{1,h}\varepsilon_t^{mp} + \beta_{2,h}Large_{f,t-1} + \beta_{3,h}Large_{f,t-1} * \varepsilon_t^{mp} + X_{f,t-1} + \mu_f + \mu_{s,t} + u_{t,h}$$

The dependent variable, $y_{f,t+h}$, is a dummy with value 1 if firms f issues LT bonds in year-quarter $t + h$ and with value 0 if it does not. ε_t^{mp} is an exogenous monetary policy shock, gathered from either Gürkaynak, Sack, and Swanson (2005) or from Jarocinski and Karadi (2020). $MacroControls_{t-1}$ is a vector of lagged macroeconomic controls, including annual GDP growth and inflation rate, a dummy for recessions, the quarterly variation in the share of LT treasuries, in the corporate spread and in the 10y-3m term-spread. $X_{f,t-1}$ is a vector of firm-level controls, including lagged sales growth, leverage and liquid assets. μ_f is a vector of firm fixed effects. $u_{t,h}$ is an error-term, double-clustered at the firm and industry*year-quarter level. The x-axis is measured in terms of quarters after the shock. The purple (green) solid line, connected by diamonds (squares), reports the point estimates for $\beta_{1,h}$ using the Gürkaynak, Sack, and Swanson (2005) (Jarocinski and Karadi 2020) shocks and the dashed (dotted) purple (green) line the respective 10% confidence intervals.

Figure 2.A.17: Monetary policy and financing costs of LT bonds - Relative response of large companies

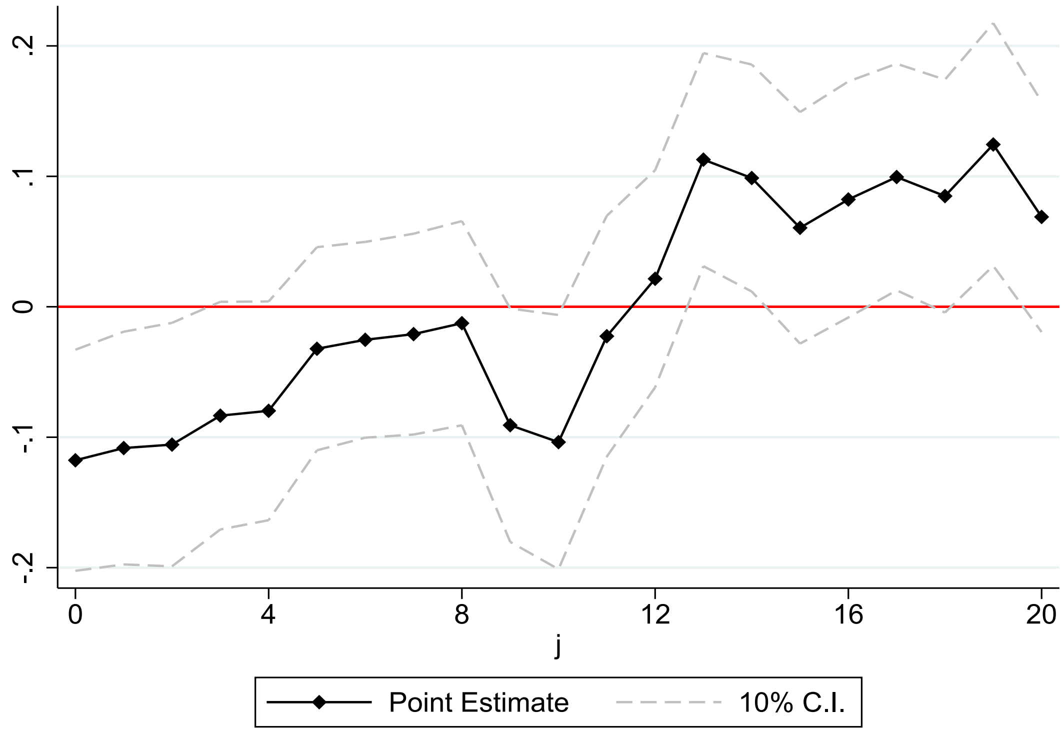

This figure depicts the relative response of companies in the top quartile of the 3-digit SIC industry asset-size distribution to a 1 p.p. lower EFFR (as compared to smaller firms). Formally, it shows the coefficients $\beta_{3,h}$ from the estimation of the following local projection model:

$$y_{f,t+h} = \beta_{1,h}EFFR_t + \beta_{2,h}Large_{f,t-1} + \beta_{3,h}Large_{f,t-1} * EFFR_t + \Phi X_{f,t-1} + \mu_f + \mu_{s,t} + \xi_{t,h}$$

The dependent variable, $y_{f,t+h}$, is the coupon rate on firm f's newly issued LT bonds in year-quarter $t+h$. $EFFR_t$ is the level of EFFR. $Large_{f,t-1}$ is a dummy variable with value 1 a firm is in the top-quartile of the industry-wide asset-size distribution, and 0 otherwise. $X_{f,t-1}$ is a vector of controls, including the interaction of $Large_{f,t-1}$ with several macro-controls and of $\Delta EFFR_t$ with other firm characteristics, namely lagged sales growth, leverage and liquid assets. μ_f and $\mu_{s,t}$ represent vectors of firm and industry*year-quarter fixed effects, respectively. $\xi_{t,h}$ is an error-term, double-clustered at the firm and industry*year-quarter level. The x-axis is measured in terms of quarters after the shock. The black solid line reports the point estimates for $\beta_{3,h}$; the dashed grey line the 10% confidence intervals.

2.B Tables

Table 2.1: Summary Statistics

VARIABLES	(1) Scale	(2) N	(3) mean	(4) p25	(5) p50	(6) p75	sd
Macro-level Variables							
$\Delta LT - Debt_t$	%	112	0.154	-0.360	0.0804	0.553	0.642
$\Delta LT - Debt_{t+1}$	%	111	0.315	-0.467	0.0687	0.992	1.141
$\Delta LT - Debt_{t+2}$	%	110	0.480	-0.590	0.232	1.558	1.602
$\Delta LT - Debt_{t+3}$	%	109	0.651	-0.794	0.123	1.834	2.026
$\Delta LT - Debt_{t+4}$	%	108	0.826	-0.720	0.218	2.554	2.401
$\Delta EFFR_t$	%	112	-0.0638	-0.105	-0.01000	0.0900	0.445
$\varepsilon_t^{mp,g}$	%	112	-0.0340	-0.0325	-0.00750	0.00500	0.0966
$\varepsilon_t^{mp,jk}$	%	112	-0.0192	-0.0396	-0.00759	0.0220	0.0756
ΔGDP_{t-1}	%	112	2.458	1.650	2.600	3.650	1.712
ΔCPI_{t-1}	%	112	2.496	1.714	2.584	3.202	1.284
Rec_{t-1}	0/1 dummy	112	0.0982	0	0	0	0.299
$\Delta LT - Treas_{t-1}$	%	112	-0.0197	-0.231	-0.00926	0.254	0.357
Δi_{t-1}^{10y-3m}	%	112	0.00866	-0.300	-0.0550	0.275	0.485
$\Delta i_{t-1}^{baa-aaa}$	%	112	-0.00205	-0.0900	-0.01000	0.0600	0.252
Firm-level Variables							
$\Delta LT - Debt_{f,t}$	%	327,532	-0.482	-2.051	0	0.800	17.03
$\Delta LT - Debt_{f,t+1}$	%	299,722	-0.915	-4.038	0	1.911	21.77
$\Delta LT - Debt_{f,t+2}$	%	284,670	-1.238	-5.662	0	2.987	24.65
$\Delta LT - Debt_{f,t+3}$	%	270,304	-1.466	-6.823	-0.00232	3.757	26.47
$\Delta LT - Debt_{f,t+4}$	%	259,167	-1.568	-7.964	0	4.795	28.08
$\Delta Sales_{f,t-1}$	%	327,532	0.75	-8.14	1.17	10.4	23
$Liquid\,Assets_{f,t-1}$	%	327,532	12.9	1.39	5.21	16.5	18.1
$Leverage_{f,t-1}$	%	327,532	34.8	12.00	27.4	42.9	45.7
$Size_{f,t-1}$	Log(Mln US$)	327,532	4.965	3.183	4.938	6.764	2.477
$\mathbb{1}(Issue)_{f,t}$	0/1 dummy	118,993	0.0677	0	0	0	0.251
$\mathbb{1}(Issue)_{f,t+1}$	0/1 dummy	110,896	0.0745	0	0	0	0.263
$\mathbb{1}(Issue)_{f,t+2}$	0/1 dummy	106,453	0.0781	0	0	0	0.268
$\mathbb{1}(Issue)_{f,t+3}$	0/1 dummy	102,372	0.0813	0	0	0	0.273
$\mathbb{1}(Issue)_{f,t+4}$	0/1 dummy	99,435	0.0829	0	0	0	0.276
$Coupon_{f,t}$	%	8,046	6.119	4.060	6.125	8	2.846
$Coupon_{f,t+1}$	%	8,262	6.093	4	6.125	8	2.837
$Coupon_{f,t+2}$	%	8,309	6.080	4.016	6.125	7.920	2.824
$Coupon_{f,t+3}$	%	8,316	6.061	4	6.125	7.875	2.826
$Coupon_{f,t+4}$	%	8,240	6.030	4	6.094	7.875	2.801

(continued on next page)

Mutual Fund-level variables							
$\Delta Matu_{m,t}$	%	72,389	-0.412	-3.190	-0.199	2.327	15.55
$\Delta Matu_{m,t+1}$	%	72,072	-0.799	-5.241	-0.731	3.673	19.55
$\Delta Matu_{m,t+2}$	%	71,617	-1.226	-6.860	-1.046	4.437	22.51
$\Delta Matu_{m,t+3}$	%	70,940	-1.614	-8.246	-1.409	5.167	24.40
$\Delta Matu_{m,t+4}$	%	69,924	-1.926	-9.381	-1.816	5.555	26.17
$\Delta CB_{m,t}$	%	71,063	-0.515	-4.449	-0.459	3.214	18.26
$\Delta CB_{m,t+1}$	%	68,201	-1.042	-6.847	-0.974	4.485	23.23
$\Delta CB_{m,t+2}$	%	65,217	-1.644	-8.635	-1.389	5.219	26.46
$\Delta CB_{m,t+3}$	%	62,112	-2.111	-9.789	-1.928	5.607	29.04
$\Delta CB_{m,t+4}$	%	58,889	-2.810	-10.99	-2.422	5.823	30.79
HY_m	0/1 dummy	72,389	0.388	0	0	1	0.487
$TurnoverRatio_{m,t-1}$	%	72,389	0.763	0.440	0.720	1.490	8.429
$ExpenseRatio_{m,t-1}$	%	72,389	0.00979	0.00570	0.00810	0.0119	0.00507
$NAV_{m,t-1}$	Log(Mln US$)	72,389	1.822	1.687	1.806	1.920	0.481
$Returns_{m,t-1}$	%	72,387	0.00960	-0.000325	0.00875	0.0215	0.0234

Macro-level Variables. Period: 1990-2017. $\Delta LT - Debt_{t+j}$ is the change in the aggregate LT-debt share (i.e., fraction of debt with maturity above 1-year) between year-quarter $t-1$ and year-quarter $t+j$, $j = 0, 1, ...4$. $\Delta EFFR_t$ is the quarterly variation in the Effective Funds Rate. $\varepsilon^{mp,g}$ is the 30-minute surprise in FED-Funds futures around policy announcements from Gürkaynak, Sack, and Swanson (2005) (aggregated at the quarterly frequency). $\varepsilon^{mp,jk}$ is the interest rate shock from Jarocinski and Karadi (2020). ΔGDP_{t-1} is the lagged annual GDP growth rate. ΔCPI is the lagged annual inflation rate. Rec_{t-1} is a lagged recession dummy. $\Delta LT - Treas_{t-1}$ is the lagged quarterly change in the share of Treasuries with maturity above 10-year. Δi^{10y-3m}_{t-1} is the lagged quarterly variation of the difference between the 10-year and the 3-month yield on benchmark US Treasuries (term-spread). $\Delta i^{baa-aaa}_{t-1}$ is the lagged quarterly variation of the difference between the BAA and the AAA Moody's Seasoned Corporate Bond Yield (corporate spread).

Firm-level Variables. Period: 1990-2017. Sample: non-financial companies identified as in Ottonello and Winberry (forthcoming). $\Delta LT - Debt_{f,t+j}$ is the change in firm f's LT-debt share (i.e., fraction of debt with maturity above 1-year) between year-quarter $t-1$ and year-quarter $t+j$, $j = 0, 1, ...4$. $\Delta Sales_{f,t-1}$ is the lagged quarterly change in log sales, expressed in p.p.. $LiquidAssets_{f,t-1}$ is the lagged share of liquid assets over total assets. $Leverage_{f,t-1}$ is the lagged ratio between total debt and total assets. $Size_{f,t-1}$ is the lagged log assets size. $\mathbb{1}(Issue_{f,t+j})$ is a dummy variable with value 1 if firm f issues bonds with maturity above 1-year in year-quarter $t+j$ and with value 0 otherwise, $j = 0, 1, ...4$. $Coupon_{f,t+j}$ is the coupon rate on the bonds issued by firm f in year-quarter $t+j$, $j = 0, 1, ...4$.

Mutual Fund-level variables. Period: 2010q2-2018q2. Sample: Corporate Bond Mutual Funds, identified as those with CRSP style categories: I, ICQH, ICQM, ICQY, ICDI, ICDS, or IC. $\Delta Matu_{m,t+4}$ is the change in the log (weighted) average portfolio maturity of fund m between year-quarter $t-1$ and year-quarter $t+j$, $j = 0, 1, ...4$. $\Delta CB_{m,t+4}$ is the change in the log corporate bond holdings between year-quarter $t-1$ and year-quarter $t+j$, $j = 0, 1, ...4$. HY_m is a dummy with value 1 if a fund m is classified as High-Yield, and 0 otherwise. HY funds are those with Lipper style code: HY, GB, FLX, MSI, or SFI. $TurnoverRatio_{m,t-1}$ is the lagged fund m's turnover ratio, corresponding to minimum (of aggregated sales or aggregated purchases of securities), divided by the average 12-month Total Net Assets. $ExpenseRatio_{m,t-1}$ is the lagged fund m's lagged expense ratio, i.e. the ratio of total investment that shareholders pay for the fund's operating expenses. $NAV_{m,t-1}$ is the latest (lagged) fund net asset value, i.e. the value of assets minus liabilities. $Returns_{m,t-1}$ reflects the lagged fund m's quarterly returns, computed as the growth in net asset value from one year-quarter to the next.

Table 2.2: Firm-level Regressions

	$\Delta LT-Debt_{f,t+j}$																				
$j=$	(0)	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)
$Large_{f,t-1}$	0.104	-0.184	-0.333	-0.234	-0.070	-0.263	-0.576	-1.020	-1.229*	-2.048***	-2.676***	-2.945***	-2.159**	-2.954***	-3.599***	-4.471***	-4.923***	-4.808***	-4.424***	-4.177***	-4.189***
	(0.207)	(0.332)	(0.434)	(0.518)	(0.569)	(0.599)	(0.634)	(0.683)	(0.734)	(0.757)	(0.788)	(0.821)	(0.847)	(0.878)	(0.895)	(0.914)	(0.942)	(0.981)	(0.987)	(1.012)	(1.026)
$\Delta EFFR_t * Large_{f,t-1}$	0.006	-0.492**	-0.899***	-0.873**	-1.376***	-1.601***	-2.151***	-1.538***	-1.799***	-1.716***	-1.571***	-1.120**	-0.304	-0.158	-0.164	-0.006	-0.026	-0.177	-0.048	-0.095	-0.018
	(0.163)	(0.242)	(0.306)	(0.368)	(0.399)	(0.421)	(0.435)	(0.454)	(0.470)	(0.484)	(0.489)	(0.498)	(0.508)	(0.524)	(0.536)	(0.547)	(0.547)	(0.549)	(0.550)	(0.538)	(0.544)
$\Delta Sales_{f,t-1}$	0.013***	0.024***	0.005*	0.006**	0.018***	0.021***	0.004	0.005	0.014***	0.018***	0.003	0.003	0.017***	0.021***	0.002	0.001	0.009**	0.012***	-0.003	-0.000	0.016***
	(0.002)	(0.002)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)
$\Delta EFFR_t * \Delta Sales_{f,t-1}$	-0.000	-0.000	-0.004	0.003	0.003	0.006	0.003	-0.001	0.004	0.001	0.005	0.005	0.000	0.014	-0.002	0.004	0.007	0.008	0.009	-0.005	-0.006
	(0.004)	(0.005)	(0.006)	(0.006)	(0.007)	(0.007)	(0.007)	(0.007)	(0.008)	(0.008)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)	(0.009)
$Liquid\ Assets_{f,t-1}$	-0.017***	-0.038***	-0.053***	-0.067***	-0.079***	-0.083***	-0.080***	-0.079***	-0.084***	-0.094***	-0.096***	-0.083***	-0.089***	-0.093***	-0.099***	-0.091***	-0.086***	-0.087***	-0.078***	-0.069***	-0.084***
	(0.004)	(0.006)	(0.008)	(0.010)	(0.012)	(0.013)	(0.014)	(0.015)	(0.016)	(0.017)	(0.018)	(0.019)	(0.020)	(0.021)	(0.022)	(0.023)	(0.024)	(0.024)	(0.024)	(0.025)	(0.026)
$\Delta EFFR_t * Liquid\ Assets_{f,t-1}$	0.002	0.001	0.003	0.019	0.029*	0.020	0.023	0.021	0.025	0.026	0.019	0.008	0.007	0.010	0.011	0.016	0.002	0.005	0.027	0.042*	0.035
	(0.005)	(0.008)	(0.011)	(0.014)	(0.015)	(0.016)	(0.018)	(0.018)	(0.019)	(0.020)	(0.021)	(0.021)	(0.021)	(0.021)	(0.021)	(0.021)	(0.021)	(0.022)	(0.022)	(0.022)	(0.023)
$Leverage_{f,t-1}$	-0.000	-0.001	-0.002	-0.004	-0.007	-0.007	-0.008	-0.010	-0.008	-0.011	-0.014*	-0.017*	-0.019**	-0.026***	-0.031***	-0.032***	-0.031***	-0.035***	-0.040***	-0.045***	-0.047***
	(0.001)	(0.002)	(0.003)	(0.004)	(0.005)	(0.006)	(0.006)	(0.007)	(0.007)	(0.008)	(0.008)	(0.009)	(0.009)	(0.009)	(0.010)	(0.010)	(0.010)	(0.010)	(0.011)	(0.011)	(0.011)
$\Delta EFFR_t * Leverage_{f,t-1}$	-0.003*	-0.002	-0.003	-0.000	-0.002	-0.001	0.002	-0.006	-0.006	-0.013	-0.010	-0.015*	-0.021**	-0.015*	-0.015*	-0.009	-0.012	-0.015*	-0.011	-0.010	-0.016**
	(0.002)	(0.003)	(0.003)	(0.004)	(0.005)	(0.006)	(0.006)	(0.007)	(0.007)	(0.008)	(0.008)	(0.008)	(0.009)	(0.009)	(0.009)	(0.009)	(0.008)	(0.009)	(0.008)	(0.009)	(0.008)
Observations	327,532	299,430	284,233	269,838	258,377	248,691	241,564	232,489	224,179	216,403	210,684	203,171	196,431	190,124	185,563	179,253	173,641	168,203	164,147	158,581	153,430
R-squared	0.093	0.119	0.138	0.154	0.170	0.183	0.192	0.202	0.212	0.223	0.231	0.240	0.248	0.260	0.266	0.276	0.281	0.292	0.300	0.311	0.320

In column j, the dependent variable is $\Delta LT - Debt_{f,t+j}$, i.e., the change in firm f's share of LT-debt between year-quarter $t-1$ and $t+j$. $\Delta EFFR_t$ is the quarterly variation in the Effective FED Funds Rate. $Large_{f,t-1}$ is a dummy with value 1 if firm f is in the top asset-size quartile of the respective (3-digit SIC) industry distribution. $\Delta Sales_{f,t-1}$ is the quarterly variation in firm f's log sales. $Liquid\ Assets_{f,t-1}$ is the share of liquid assets by firm f. $Leverage_{f,t-1}$ is firm f's leverage, defined as total debt over total assets. All firm-level variables are lagged by one year-quarter. Each regression additionally includes the full interaction of $Large_{f,t-1}$ with a vector of lagged macro controls (annual GDP growth and inflation rate; a recession dummy; quarterly variation in term-spread, corporate spread and in share of Treasuries with maturity above 20-year). Furthermore, in each column we apply firm and industry*year-quarter fixed effects. Standard errors are clustered at the Firm and Industry*Year-Quarter level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.3

	(1)	(2)	(3)	(4)	(5)	(6)
VARIABLES		$\mathbb{1}(Issue)_{f,t}$			$\mathbb{1}(Issue)_{f,t+1}$	
$\Delta EFFR_t$	-0.010***			-0.012***		
	(0.003)			(0.004)		
$Large_{f,t-1}$	0.047***	0.046***	0.047***	0.043***	0.044***	0.039***
	(0.005)	(0.006)	(0.006)	(0.005)	(0.006)	(0.007)
$\Delta EFFR_t * Large_{f,t-1}$	-0.013***	-0.008**	-0.011**	-0.018***	-0.013***	-0.016***
	(0.004)	(0.004)	(0.004)	(0.004)	(0.004)	(0.005)
Observations	118,993	118,993	112,669	111,703	111,703	105,198
R-squared	0.089	0.094	0.210	0.091	0.096	0.218
Firm Controsl*$\Delta EFFR_t$	Yes	Yes	Yes	Yes	Yes	Yes
Macro Controls*Large	Yes	Yes	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes
Year:Quarter FE	No	Yes	-	No	Yes	-
Industry*Year-Quarter FE	No	No	Yes	No	No	Yes

In columns (1)-(3), the dependent is variable, $\mathbb{1}(Issue)_{f,t}$ is a dummy variable with value if firm f issues LT bonds in year-quarter t. In columns (3)-(6), the dependent is the same variable, though measuread as of year-quarter $t+1$. $\Delta EFFR_t$ is the quarterly variation in the Effective FED Funds Rate. $Large_{f,t-1}$ is a dummy with value 1 if firm f is in the top asset-size quartile of the respective (3-digit SIC) industry distribution. Firm controls include lagged sales growth, leverage and share of liquid assets. Macro controls are given by annual GDP growth and inflation rate; a recession dummy; quarterly variation in term-spread, corporate spread and in share of Treasuries with maturity above 20-year. Standard errors are clustered at the Firm and Industry*Year-Quarter level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.4: Monetary policy and financing costs through LT bonds

VARIABLES	(1)	(2)	(3)	(4)
	$Coupon_{f,t}$		$Coupon_{f,t+1}$	
$EFFR_t$	0.092**		0.133***	
	(0.043)		(0.045)	
$Large_{f,t-1}$	-1.573***	-0.065	-1.516***	-0.181
	(0.340)	(0.473)	(0.296)	(0.537)
$EFFR_t * Large_{f,t-1}$	0.269***	0.118**	0.231***	0.108**
	(0.041)	(0.052)	(0.039)	(0.054)
Observations	7,310	4,157	7,551	4,215
R-squared	0.706	0.835	0.711	0.836
Firm Controls*EFFR	Yes	Yes	Yes	Yes
Macro Controls*Large	Yes	Yes	Yes	Yes
Firm FE	Yes	Yes	Yes	Yes
Industry*Year:Quarter FE	No	Yes	No	Yes

In columns (1)-(2), the dependent variable, is the coupon rate on firm f's newly issued LT bonds in year-quarter t. In columns (3)-(4), the left-hand side variable is the same, though measured in year-quarter $t+1$. $EFFR_t$ is the level of EFFR. $Large_{f,t-1}$ is a dummy with value 1 if firm f is in the top asset-size quartile of the respective (3-digit SIC) industry distribution. Firm controls include lagged sales growth, leverage and share of liquid assets. Macro controls are given by annual GDP growth and inflation rate; a recession dummy; term-spread, corporate spread and share of Treasuries with maturity above 20-year. Standard errors are double-clustered at the firm and industry*year-quarter level. *** p<0.01, ** p<0.05, * p<0.1.

Table 2.5: Monetary policy and financing costs through LT bonds - Robustness

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
	$Coupon_{f,t}$		$Coupon_{f,t+1}$		$Coupon_{f,t}$		$Coupon_{f,t+1}$		$Coupon_{f,t}$		$Coupon_{f,t+1}$	
$EFFR_t$	-0.187*		-0.148		0.077		0.145***		0.076		0.124**	
	(0.108)		(0.102)		(0.058)		(0.055)		(0.053)		(0.049)	
$Large_{f,t-1}$	-9.106***	-2.410	-9.954***	-1.931	-1.917***	-0.079	-1.945***	0.297	-1.473***	0.277	-1.369***	0.620
	(1.820)	(2.628)	(1.936)	(2.677)	(0.384)	(0.467)	(0.333)	(0.504)	(0.372)	(0.431)	(0.321)	(0.434)
$Large_{f,t-1} * EFFR_t$	0.738***	0.454***	0.710***	0.354**	0.302***	0.196***	0.250***	0.163***	0.285***	0.155***	0.232***	0.109**
	(0.106)	(0.156)	(0.111)	(0.151)	(0.050)	(0.053)	(0.046)	(0.055)	(0.046)	(0.044)	(0.043)	(0.045)
Observations	6,856	6,856	7,057	7,057	6,822	6,704	7,028	6,924	7,278	7,168	7,517	7,416
R-squared	0.722	0.740	0.721	0.739	0.715	0.766	0.715	0.766	0.709	0.759	0.713	0.763
Firm Controls*EFFR	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Macro Controls*Large	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Size-Distribution level	All	All	All	All	Sector	Sector	Sector	Sector	Sic-2	Sic-2	Sic-2	Sic-2
Firm FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year:Quarter FE	No	Yes	No	Yes	No	-	No	-	No	-	No	-
Sector*Year:Quarter FE	No	No	No	No	No	Yes	No	Yes	No	-	No	-
Sic-2*Year:Quarter FE	No	No	No	No	No	No	No	No	No	Yes	No	Yes

In columns (1)-(2), (5)-(6) and (9)-(10), the dependent variable, is the coupon rate on firm f's newly issued LT bonds in year-quarter t. In columns (3)-(4), (7)-(8) and (11)-(12), the left-hand side variable is the same, though measured in year-quarter $t + 1$. $EFFR_t$ is the level of EFFR. $Large_{f,t-1}$ is a dummy with value 1 if firm f is in the top quartile of the asset-size distribution. As reported in the legend row "Size-Distribution level", in columns (1)-(4) the relevant distribution comprehends all companies; in columns (5)-(8), the large dummy is computed within sectors (i.e., 1-digit SIC code); finally, in columns (9)-(12) within 2-digit SIC code industries. Firm controls include lagged sales growth, leverage and share of liquid assets. Macro controls are given by annual GDP growth and inflation rate; a recession dummy; term-spread, corporate spread and share of Treasuries with maturity above 20-year. Standard errors are double-clustered at the: firm and year-quarter level in columns (1)-(4); firm and sector level in columns (5)-(8); firm and (2-digit SIC)-industry*year-quarter level in columns (9)-(12). *** p<0.01, ** p<0.05, * p<0.1.

2.C Empirical Appendix

Figure A1: Monetary Policy and Debt Maturity Structure: Aggregate Response - Pre-Crisis Period

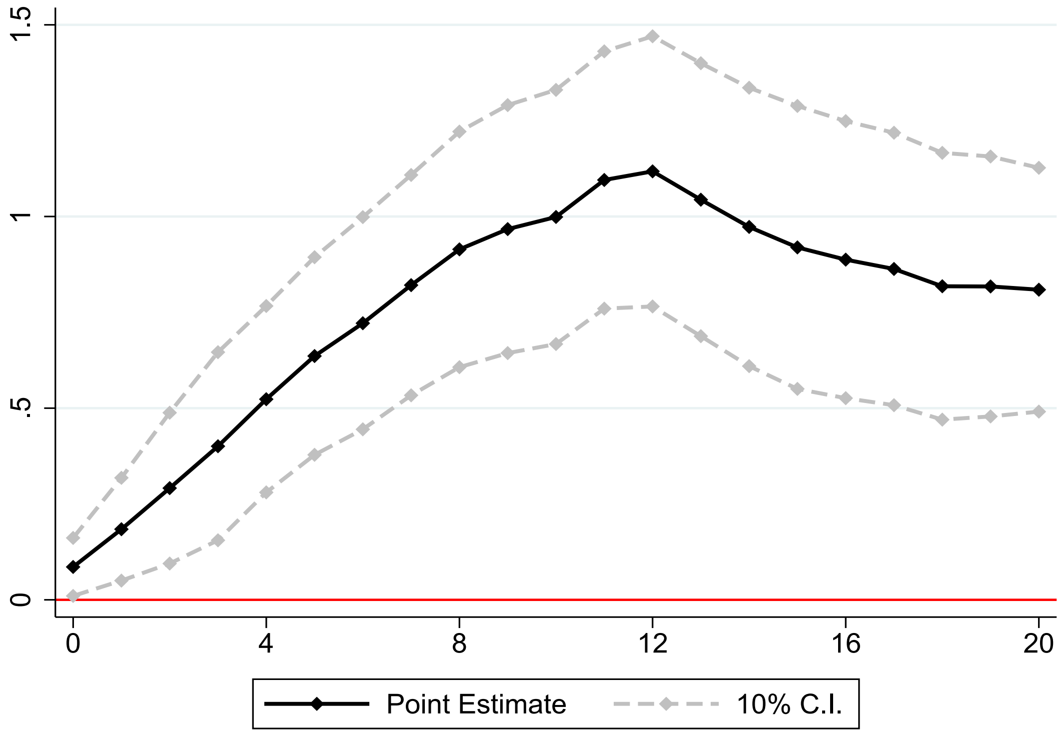

This figure depicts the response of the aggregate-level share of LT debt to a 25 b.p. cut in the EFFR. The sample includes observations from 1990q1 to 2008q4. Formally, it shows the coefficients $\beta_{1,h}$ from the estimation of the following local projection model:

$$\Delta_h y_{t+h} = \beta_{1,h} \Delta EFFR_t + MacroControls_{t-1} + u_{t,h}$$

The dependent variable, $\Delta_h y_{t+h}$, represents the growth of the LT-debt share (expressed in p.p.) from year-quarter $t-1$ to year-quarter $t+h$. $\Delta EFFR_t$ is the quarterly EFFR change. $MacroControls_{t-1}$ is a vector of lagged macroeconomic controls, including annual GDP growth and inflation rate, a dummy for recessions, the quarterly variation in the share of LT treasuries, in the corporate spread and in the 10y-3m term-spread. $u_{t,h}$ is a robust error-term. The x-axis is measured in terms of quarters after the shock. The black solid line reports the point estimates for $\beta_{1,h}$; the dashed grey line the 10% confidence intervals.

Figure A2: Monetary Policy and Debt Maturity Structure: Aggregate Response using Exogenous Shocks - Pre-Crisis Period

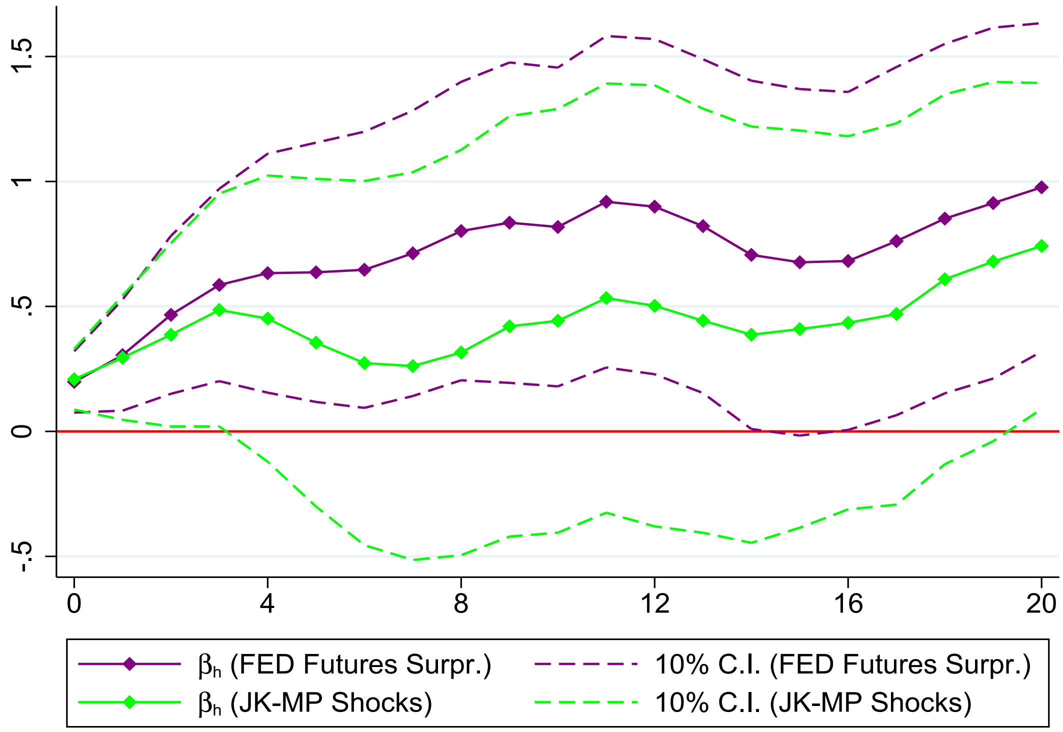

This figure depicts the response of the aggregate-level share of LT debt to a 1 s.d. reduction in the monetary policy shock. The sample includes observations from 1990q1 to 2008q4. Formally, it shows the coefficients $\beta_{1,h}$ from the estimation of the following local projection model:

$$\Delta_h y_{t+h} = \beta_{1,h}\varepsilon_t^{mp} + MacroControls_{t-1} + u_{t,h}$$

The dependent variable, $\Delta_h y_{t+h}$, represents the growth of the LT-debt share (expressed in p.p.) from year-quarter $t-1$ to year-quarter $t+h$. ε_t^{mp} is an exogenous monetary policy shock, gathered from either Gürkaynak, Sack, and Swanson (2005) or from Jarocinski and Karadi (2020). $MacroControls_{t-1}$ is a vector of lagged macroeconomic controls, including annual GDP growth and inflation rate, a dummy for recessions, the quarterly variation in the share of LT treasuries, in the corporate spread and in the 10y-3m term-spread. $u_{t,h}$ is a robust error-term. The x-axis is measured in terms of quarters after the shock. The purple (green) solid line, connected by diamonds (squares), reports the point estimates for $\beta_{1,h}$ using the Gürkaynak, Sack, and Swanson (2005) (Jarocinski and Karadi 2020) shocks and the dashed (dotted) purple (green) line the respective 10% confidence intervals.

Figure A3: Monetary Policy and Debt Maturity Structure - Relative response of Large Companies: Pre-Crisis

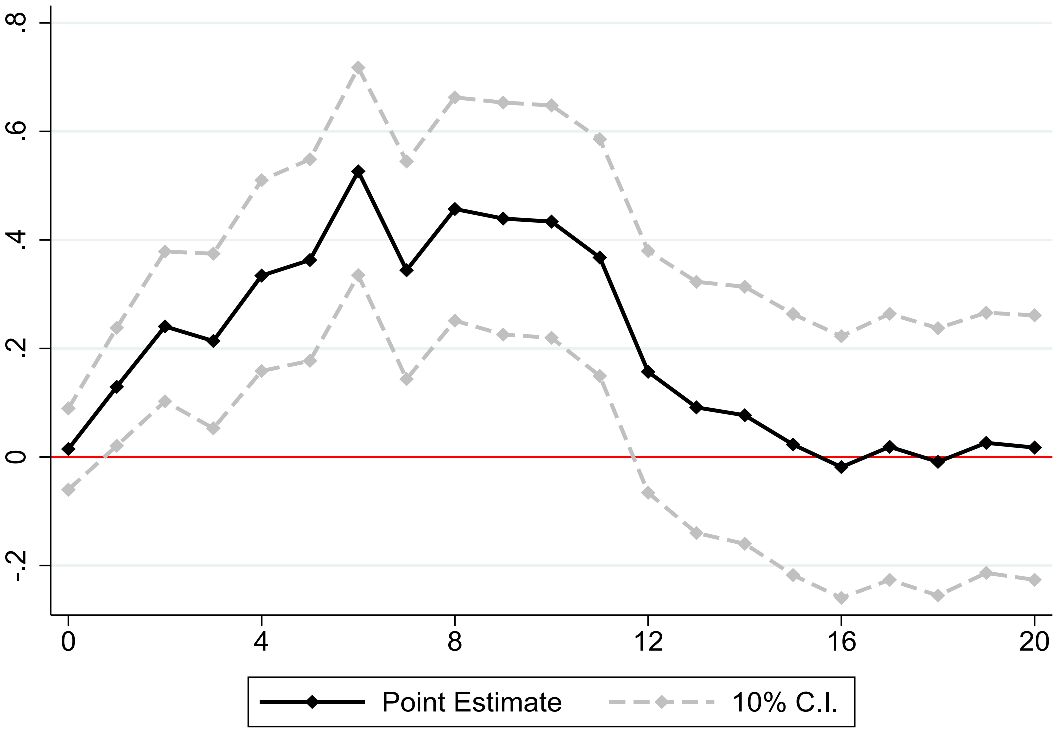

This figure depicts the relative response of companies in the top quartile of the 3-digit SIC industry asset-size distribution to a 25 b.p. cut in the EFFR (as compared to smaller firms). The regression sample goes from 1990q1 to 2008q4. Formally, the picture shows the coefficients $\beta_{3,h}$ from the estimation of the following local projection model:

$$\Delta_h y_{f,t+h} = \beta_{1,h}\Delta EFFR_t + \beta_{2,h}Large_{f,t-1} + \beta_{3,h}Large_{f,t-1} * \Delta EFFR_t + X_{f,t-1} + \mu_f + \mu_{s,t} + u_{t,h}$$

The dependent variable, $\Delta_h y_{f,t+h}$, represents the growth of the share of LT-debt - expressed in p.p. - from year-quarter $t-1$ to year-quarter $t+h$. $\Delta EFFR_t$ is the quarterly change in the EFFR. $Large_{f,t-1}$ is a dummy variable with value 1 a firm is in the top-quartile of the industry-wide asset-size distribution, and 0 otherwise. $X_{f,t-1}$ is a vector of controls, including the interaction of $Large_{f,t-1}$ with several macro-controls and of $\Delta EFFR_t$ with other firm characteristics, namely lagged sales growth, leverage and liquid assets. μ_f and $\mu_{s,t}$ represent vectors of firm and industry*year-quarter fixed effects, respectively. $u_{t,h}$ is an error-term, double-clustered at the firm and industry*year-quarter level. The x-axis is measured in terms of quarters after the shock. The black solid line reports the point estimates for $\beta_{3,h}$; the dashed grey line the 10% confidence intervals.

Figure A4: Monetary Policy and Debt Maturity Structure - Relative response of Large firms using Exogenous Shocks: Pre-Crisis

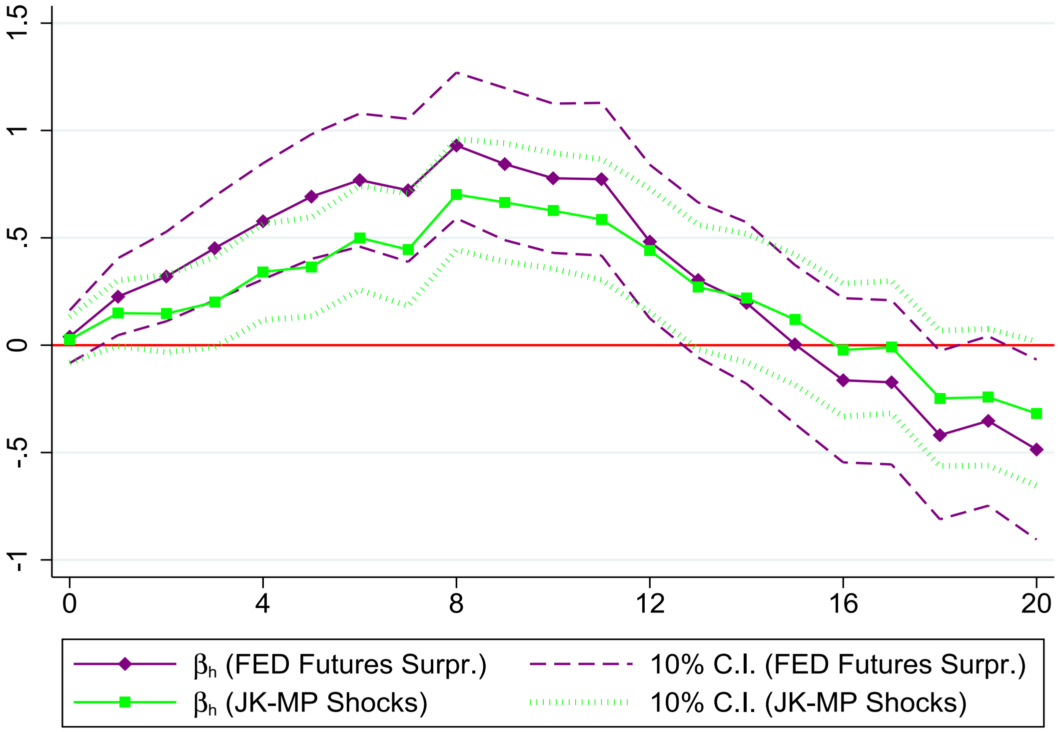

This figure depicts the relative response of companies in the top quartile of the 3-digit SIC industry asset-size distribution to a 1 s.d. b.p.reduction in monetary policy shock (as compared to smaller firms). The regression sample goes from 1990q1 to 2008q4. Formally, the picture shows the coefficients $\beta_{3,h}$ from the estimation of the following local projection model:

$$\Delta_h y_{f,t+h} = \beta_{1,h}\varepsilon_t^{mp} + \beta_{2,h}Large_{f,t-1} + \beta_{3,h}Large_{f,t-1} * \varepsilon^{mp} + X_{f,t-1} + \mu_f + \mu_{s,t} + u_{t,h}$$

The dependent variable, $\Delta_h y_{f,t+h}$, represents the growth of the share of LT-debt - expressed in p.p. - from year-quarter $t-1$ to year-quarter $t+h$. ε_t^{mp} is an exogenous monetary policy shock, derived either from Gürkaynak, Sack, and Swanson (2005) or from Jarocinski and Karadi (2020). $Large_{f,t-1}$ is a dummy variable with value 1 a firm is in the top-quartile of the industry-wide asset-size distribution, and 0 otherwise. $X_{f,t-1}$ is a vector of controls, including the interaction of $Large_{f,t-1}$ with several macro-controls and of ε_t^{mp} with other firm characteristics, namely lagged sales growth, leverage and liquid assets. μ_f and $\mu_{s,t}$ represent vectors of firm and industry*year-quarter fixed effects, respectively. $u_{t,h}$ is an error-term, double-clustered at the firm and industry*year-quarter level. The x-axis is measured in terms of quarters after the shock. The purple (green) solid line, connected by diamonds (squares), reports the point estimates for $\beta_{1,h}$ using the Gürkaynak, Sack, and Swanson (2005) (Jarocinski and Karadi 2020) shocks and the dashed (dotted) purple (green) line the respective 10% confidence intervals.

Figure A5: Contemporaneous Quarterly Variation of Term-Spread and EFFR

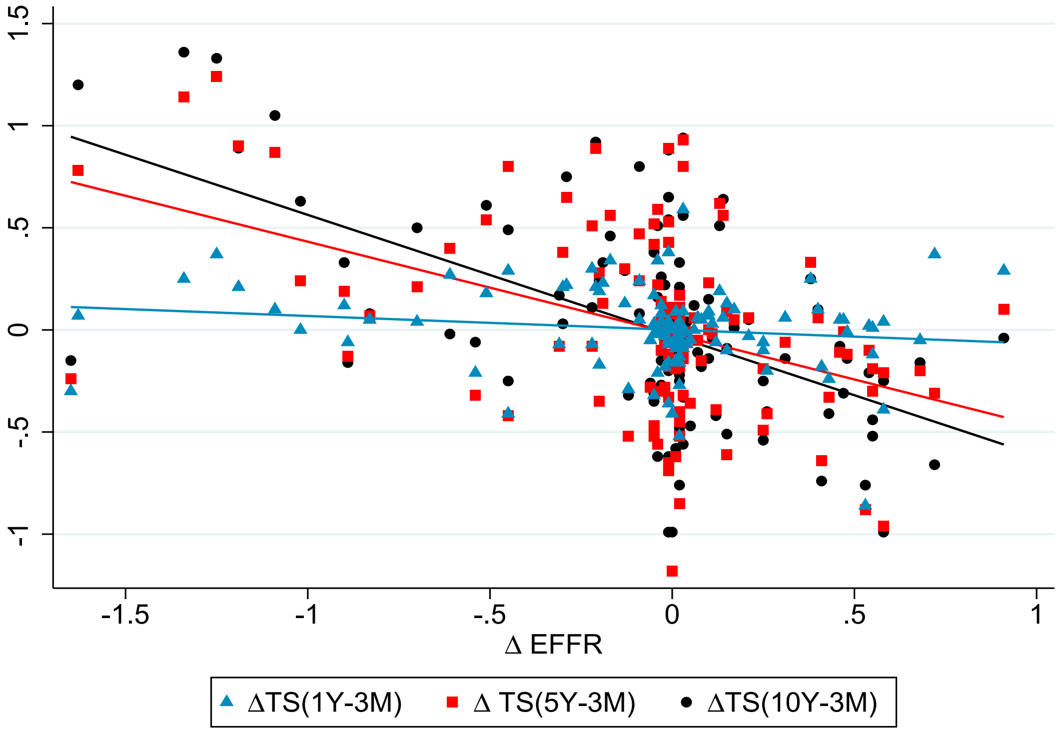

We report the quarterly variation of the Effective FED Funds Rate on the x-axis and the contemporaneous quarterly change in the term-spread on the y-axis. Both measures are expressed in percentage points. We employ several definitions of the term-spread, all based on the benchmark Treasury yields at constant maturity. The 1-year/3-month spread is in light blue. The 5-year/3-month spread is in red, whereas the 10-year/3-month spread is in black. The lines, which are colored accordingly, reflect results from simple bivariate linear-fit regressions.

Figure A6: Monetary Policy and the Term-Spread - Local Projections

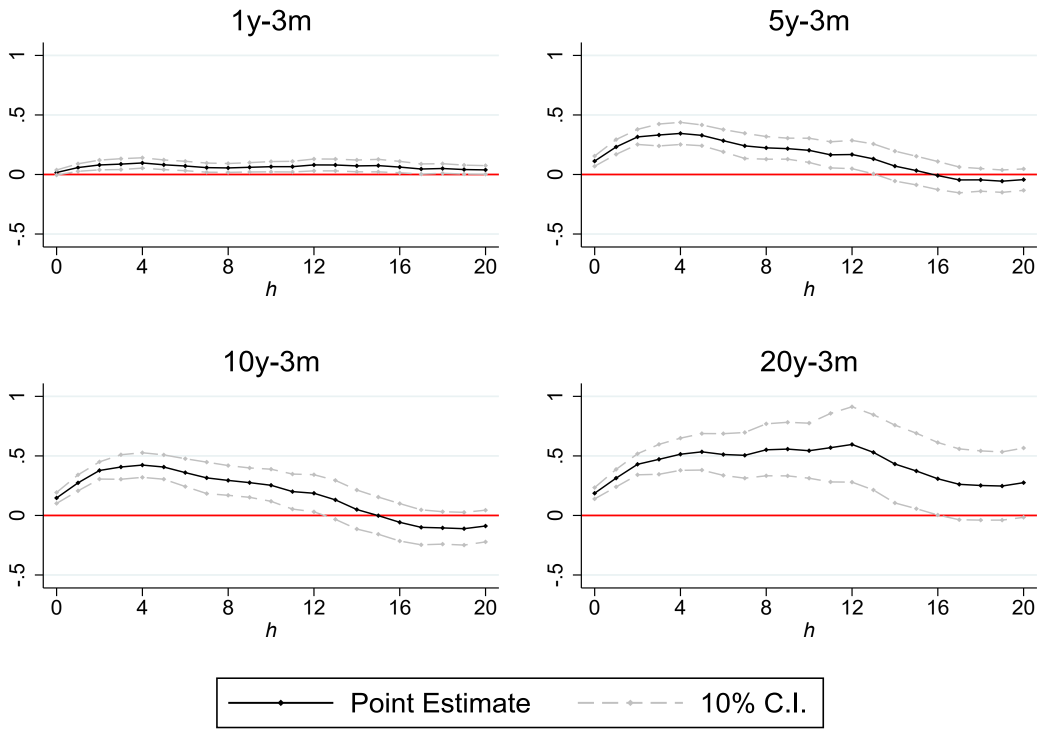

This figure depicts the response of the term-spread to a 25 b.p. cut in the EFFR. We employ several definitions of the term-spread, all based on the benchmark Treasury yields at constant maturity. The 1-year/3-month spread is in the north-west sub-plot. The 5-year/3-month spread is in the north-east sub-plot. The 10-year/3-month spread is in the south-west sub-plot, whereas the 20-year/3-month spread is in the south-east one. Formally, it shows the coefficients $\beta_{1,h}$ from the estimation of the following local projection model:

$$\Delta_h y_{t+h} = \beta_{1,h} \Delta EFFR_t + MacroControls_{t-1} + u_{t,h}$$

The dependent variable, $\Delta_h y_{t+h}$, represents the growth of the term-spread (expressed in p.p.) from year-quarter $t-1$ to year-quarter $t+h$. $\Delta EFFR_t$ is the quarterly EFFR change. $MacroControls_{t-1}$ is a vector of lagged macroeconomic controls, including annual GDP growth and inflation rate, a dummy for recessions, the quarterly variation in the share of LT treasuries and in the corporate spread. $u_{t,h}$ is a robust error-term. The x-axis is measured in terms of quarters after the shock. The black solid line reports the point estimates for $\beta_{1,h}$; the dashed grey line the 10% confidence intervals.

Figure A7: Monetary Policy and the Term-Spread - Local Projections using Exogenous Shocks

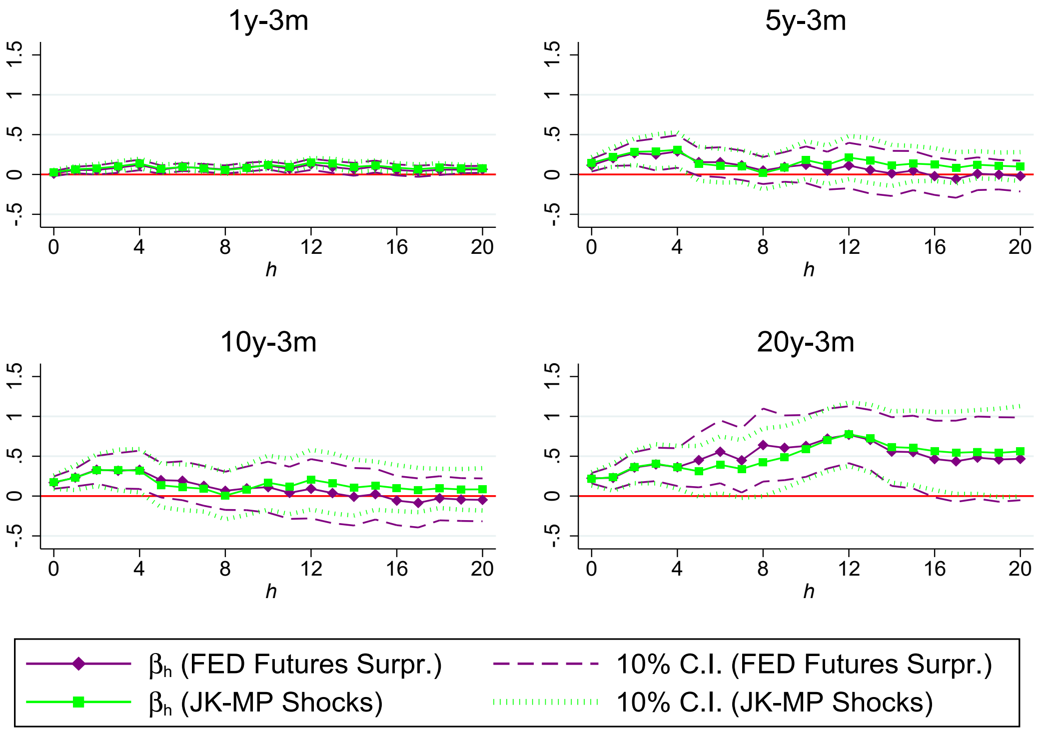

This figure depicts the response of the term-spread to a 1 s.d. reduction in the monetary policy shock. We employ several definitions of the term-spread, all based on the benchmark Treasury yields at constant maturity. The 1-year/3-month spread is in the north-west sub-plot. The 5-year/3-month spread is in the north-east sub-plot. The 10-year/3-month spread is in the south-west sub-plot, whereas the 20-year/3-month spread is in the south-east one. Formally, it shows the coefficients $\beta_{1,h}$ from the estimation of the following local projection model:

$$\Delta_h y_{t+h} = \beta_{1,h}\varepsilon_t^{mp} + MacroControls_{t-1} + u_{t,h}$$

The dependent variable, $\Delta_h y_{t+h}$, represents the growth of the term-spread (expressed in p.p.) from year-quarter $t-1$ to year-quarter $t+h$. ε_t^{mp} is an exogenous monetary policy shock, gathered from either Gürkaynak, Sack, and Swanson (2005) or from Jarocinski and Karadi (2020). $MacroControls_{t-1}$ is a vector of lagged macroeconomic controls, including annual GDP growth and inflation rate, a dummy for recessions, the quarterly variation in the share of LT treasuries and in the corporate spread. $u_{t,h}$ is a robust error-term. The x-axis is measured in terms of quarters after the shock. The purple (green) solid line, connected by diamonds (squares), reports the point estimates for $\beta_{1,h}$ using the Gürkaynak, Sack, and Swanson (2005) (Jarocinski and Karadi 2020) shocks and the dashed (dotted) purple (green) line the respective 10% confidence intervals.

Figure A8: Monetary policy and likelihood of issuing LT bonds - Relative response of large companies: pre-Crisis

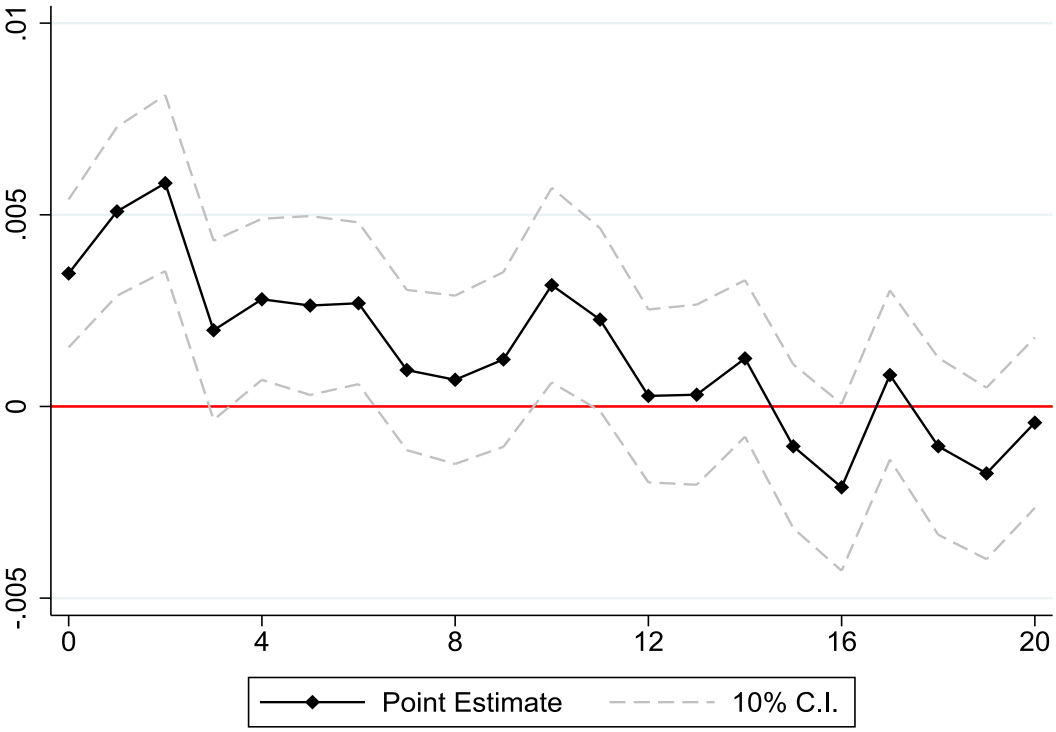

This figure depicts the relative response of companies in the top quartile of the 3-digit SIC industry asset-size distribution to a 25 b.p. cut in the EFFR (as compared to smaller firms). The regression sample goes from 1990q1 to 2008q4. Formally, it shows the coefficients $\beta_{3,h}$ from the estimation of the following local projection model:

$$y_{f,t+h} = \beta_{1,h}\Delta EFFR_t + \beta_{2,h}Large_{f,t-1} + \beta_{3,h}Large_{f,t-1} * \Delta EFFR_t + X_{f,t-1} + \mu_f + \mu_{s,t} + u_{t,h}$$

The dependent variable, $y_{f,t+h}$, is a dummy with value 1 if firms f issues LT bonds in year-quarter $t + h$ and with value 0 if it does not. $\Delta EFFR_t$ is the quarterly change in the EFFR. $Large_{f,t-1}$ is a dummy variable with value 1 a firm is in the top-quartile of the industry-wide asset-size distribution, and 0 otherwise. $X_{f,t-1}$ is a vector of controls, including the interaction of $Large_{f,t-1}$ with several macro-controls and of $\Delta EFFR_t$ with other firm characteristics, namely lagged sales growth, leverage and liquid assets. μ_f and $\mu_{s,t}$ represent vectors of firm and industry*year-quarter fixed effects, respectively. $u_{t,h}$ is an error-term, double-clustered at the firm and industry*year-quarter level. The x-axis is measured in terms of quarters after the shock. The black solid line reports the point estimates for $\beta_{3,h}$; the dashed grey line the 10% confidence intervals.

Figure A9: Monetary policy and likelihood of issuing LT bonds - Relative response of large companies using exogenous shocks: pre-Crisis

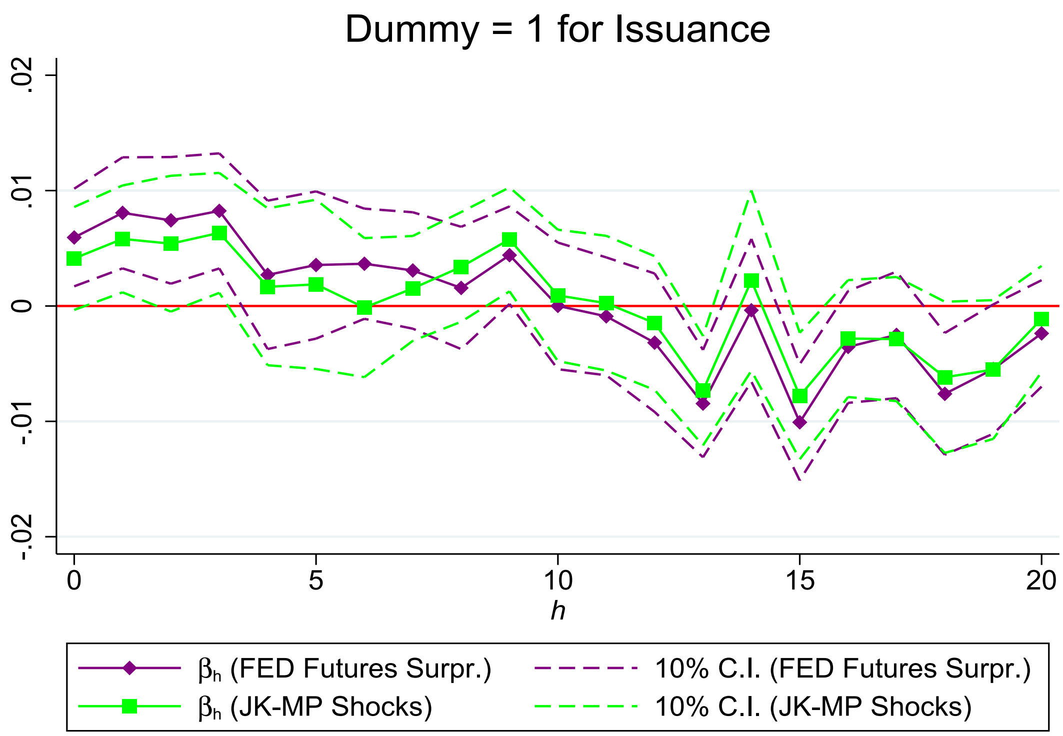

This figure depicts the relative response of companies in the top quartile of the 3-digit SIC industry asset-size distribution to a 25 b.p. cut in the EFFR (as compared to smaller firms). Formally, it shows the coefficients $\beta_{3,h}$ from the estimation of the following local projection model:

$$y_{f,t+h} = \beta_{1,h}\varepsilon_t^{mp} + \beta_{2,h}Large_{f,t-1} + \beta_{3,h}Large_{f,t-1} * \varepsilon_t^{mp} + X_{f,t-1} + \mu_f + \mu_{s,t} + u_{t,h}$$

The dependent variable, $y_{f,t+h}$, is a dummy with value 1 if firms f issues LT bonds in year-quarter $t + h$ and with value 0 if it does not. ε_t^{mp} is an exogenous monetary policy shock, gathered from either Gürkaynak, Sack, and Swanson (2005) or from Jarocinski and Karadi (2020). $MacroControls_{t-1}$ is a vector of lagged macroeconomic controls, including annual GDP growth and inflation rate, a dummy for recessions, the quarterly variation in the share of LT treasuries, in the corporate spread and in the 10y-3m term-spread. $X_{f,t-1}$ is a vector of firm-level controls, including lagged sales growth, leverage and liquid assets. μ_f is a vector of firm fixed effects. $u_{t,h}$ is an error-term, double-clustered at the firm and industry*year-quarter level. The x-axis is measured in terms of quarters after the shock. The purple (green) solid line, connected by diamonds (squares), reports the point estimates for $\beta_{1,h}$ using the Gürkaynak, Sack, and Swanson (2005) (Jarocinski and Karadi 2020) shocks and the dashed (dotted) purple (green) line the respective 10% confidence intervals.

2.D Theoretical Appendix

Proofs

Proof of Lemma 2.1. Firms' endowments are distributed uniformly and continuously on [0,I]. This implies that each firm is atomistic relative to the set of all firms. The expected return of the market portfolio thus does not respond to changes in a single firm's $d_l(A)$ and $\rho^*(A)$:

$$\frac{\partial \int_0^I P_l(A)d_l(A)dA}{\partial d_l(A)} = 0 \qquad \frac{\int_0^I P_l(A)d_l(A)dA}{\partial \rho^*(A)} = 0.$$

Taking derivatives of Equation 2.9 yields

$$\frac{\partial P_l(A)}{\partial d_l(A)} = 0 \qquad \frac{\partial P_l(A)}{\partial \rho^*(A)} = \delta f(\rho^*(A)) \frac{1-\gamma(1-\delta)\left(\int_0^I F(\rho^*(A))d_l(A)dA - g\right)}{(1+i_1)[\alpha(1+i_1)+(1-\alpha)(1+i_2)]}.$$

Note that $\partial P_l(A)/\partial \rho^*(A) > 0$ iff $P_l(A) > 0$, which we know to be true from the first order conditions of the firm, specifically Equation 2.11.

□

Proof of Proposition 2.1. The three constraints of the firms' problem are:

$$r - \rho^*(A) \geq d_s(A) \tag{LL}$$

$$R - \frac{B}{\Delta P} \geq d_l(A) \tag{IC}$$

$$\frac{d_s(A)}{1+i_1} + P_l(A)d_l(A) \geq I - A \tag{IR}$$

First, we will show that these three constraints will only bind simultaneously, such that, notationally, we only need one Lagrange-multiplier $\lambda(A)$. If there exists any one firm that is unconstrained in equilibrium, i.e., if there exists A' such that $0 \leq A' \leq I$ and $\lambda_2(A') = \lambda_3(A') = 0$, then condition 2.11 implies that

$$P_l(A') = \frac{1+\lambda_2(A')}{1+\lambda_3(A')} \frac{\delta F(\rho^*(A'))}{(1+i_1)(1+i_2)} = \frac{\delta F(\rho^*(A'))}{(1+i_1)(1+i_2)}.$$

Condition 2.11 must hold for all possible endowments A and thus

$$\frac{P_l(A)}{\delta F(\rho^*(A))} = \frac{1+\lambda_2(A)}{1+\lambda_3(A)} \frac{1}{(1+i_1)(1+i_2)}$$

and since the LHS is identical for all A, it follows that

$$\frac{1+\lambda_2(A)}{1+\lambda_3(A)} = 1 \qquad \Longleftrightarrow \qquad \lambda_2(A) = \lambda_3(A) \quad \forall A$$

and from condition 2.10 that

$$\lambda(A) \equiv \lambda_1(A) = \lambda_2(A) = \lambda_3(A) \forall A \quad \forall A.$$

Now, we can pinpoint the threshold endowment $\bar{A}$, which is the lowest endowment at which $\lambda(\bar{A}) = 0$. The critical value $\bar{A}$ is the one that makes the constraints bind exactly, for a firm that chooses the optimal cutoff value $\rho^* = \delta R/(1+i_2)$. We find it by plugging the LL and the IC constraints into the IR constraint. From condition 2.11, the price of LT debt takes the form

$$P_l(A) = \frac{\delta F\left(\frac{\delta R}{1+i_2}\right)}{(1+i_1)(1+i_2)}$$

and we get that the threshold endowment must be:

$$\bar{A} = I - \frac{r - \frac{\delta R}{1+i_2}}{1+i_1} - \frac{\delta F\left(\frac{\delta R}{1+i_2}\right)}{(1+i_1)(1+i_2)}\left(R - \frac{B}{\Delta P}\right).$$

Both types, unconstrained and constrained firms, exist concurrently if

$$\bar{A} \in (0, I).$$

□

Proof of Lemma 2.2. Combining conditions 2.11 and 2.12 we see that:

$$\rho^*(A) = \frac{\delta R}{1+i_2} - \lambda(A)\left[\frac{1}{f(\rho^*)} - \frac{\delta d_l(A)}{1+i_2}\right].$$

An increase in $\rho^*(A)$ has two effects for constrained firms. On the one hand, it tightens the LL constraint, such that less short-term debt can be issued. On the other hand, a higher probability of continuation yields a higher price. We will show that our assumption $\chi\delta(1+i_2)^{-1}(R - B/\Delta p) < 1$ assures that the former effect always dominates the latter and that all statements in this lemma follow from this fact. Note that for constrained firms, for which $\lambda(A) > 0$, it is true that:

$$\rho^*(A) < \frac{\delta R}{1+i_2} \qquad \iff \qquad f(\rho^*)\frac{\delta d_l(A)}{1+i_2} < 1.$$

The inequality $f(\rho^*)\frac{\delta d_l(A)}{1+i_2} < 1$ holds due to the assumption that $\frac{\chi\delta}{1+i_2}\left(R - \frac{B}{\Delta p}\right) <$ 1, as

$$f(\rho^*)\frac{\delta d_l(A)}{1+i_2} < f(0)\frac{\delta\left(R - \frac{B}{\Delta p}\right)}{1+i_2} < \frac{\chi\delta}{1+i_2}\left(R - \frac{B}{\Delta p}\right) < 1.$$

Therefore,

$$\rho^*(A) < \frac{\delta R}{1+i_2}.$$

Statements 1. and 2. in the Lemma follow immediately from the IR constraint. As the unconstrained firms, those with $A > \bar{A}$, always choose $\rho^*(A) = \delta R/(1+i_2)$, there will be no effect of a change in either A nor i_1 on their choice of $\rho^*(A)$. For the constrained firms we can derive the change in $\rho^*(A)$ from the constraints.

Recall that for the constrained firm the following must hold:

$$\frac{r-\rho^*(A)}{1+i_1} + \delta F(\rho^*(A))\frac{1-\gamma(1-\delta)\left(\int_0^I F(\rho^*(A))d_l(A)dA - g\right)}{(1+i_1)[\alpha(1+i_1)+(1-\alpha)(1+i_2)]}\left(R - \frac{B}{\Delta P}\right) = I - A.$$

By plugging in the equilibrium price

$$P_l(A) = \frac{\delta F(\rho^*(A))}{(1+i_1)(1+i_2)}$$

this becomes:

$$\frac{r-\rho^*(A)}{1+i_1} + \frac{\delta F(\rho^*(A))}{(1+i_1)(1+i_2)}\left(R - \frac{B}{\Delta P}\right) = I - A.$$

We apply the Implicit Function Theorem to the equality

$$l : \frac{r-\rho^*(A)}{1+i_1} + \frac{\delta F(\rho^*(A))}{(1+i_1)(1+i_2)}\left(R - \frac{B}{\Delta P}\right) - I + A = 0.$$

For $A < \bar{A}$ we get

$$\frac{\partial \rho^*(A)}{\partial i_1} = -\frac{\frac{\partial l}{\partial i_1}}{\frac{\partial l}{\partial \rho^*(A)}} < 0$$

since

$$\frac{\partial l}{\partial i_1} < 0,$$

and

$$\frac{\partial l}{\partial \rho^*(A)} = -\frac{1}{1+i_1} + \frac{\delta f(\rho^*(A))}{(1+i_1)(1+i_2)}\left(R - \frac{B}{\Delta P}\right) < 0,$$

as

$$f(\rho^*)\frac{\delta\left(R - \frac{B}{\Delta P}\right)}{1+i_2} < 1.$$

Moreover,

$$\frac{\partial \rho^*(A)}{\partial A} = -\frac{1}{\frac{\partial l}{\partial \rho^*(A)}} > 0.$$

□

Proof of Proposition 2.2. As the unconstrained firms, those with $A > \bar{A}$, always choose $\rho^*(A) = \delta R/(1+i_2)$, there will be no effect of a change in i_1 on their choice. Furthermore, the total expected revenue from long term debt $\int_A F(\rho^*(A))d_l(A)$ is pinned down by the market clearing equation:

$$\frac{1-\gamma(1-\delta)\left(\int_0^I F(\rho^*(A))d_l(A)dA - g\right)}{[\alpha(1+i_1)+(1-\alpha)(1+i_2)]} = \frac{1}{1+i_2}.$$

We can reformulate this equation as

$$\int_0^I F(\rho^*(A))d_l(A)dA = \frac{\varkappa(i_2-i_1)}{(1+i_2)(1-\delta)} + g.$$

We can then split the left-hand side into the unconstrained and constrained components of LT debt

$$\int_0^{\bar{A}} F(\rho^*(A))d_l(A)dA + \int_{\bar{A}}^I F(\rho^*(A))d_l(A)dA = \frac{\varkappa(i_2-i_1)}{(1+i_2)(1-\delta)} + g$$

and get

$$\int_{\bar{A}}^{I} d_l(A)dA = \frac{1}{F\left(\dfrac{\delta R}{1+i_2}\right)}\left[g + \frac{\varkappa\,(i_2 - i_1)}{(1+i_2)(1-\delta)} - \left(R - \frac{B}{\Delta P}\right)\int_0^{\bar{A}} F(\rho^*(A))dA\right].$$

Thus,

$$\frac{\partial \int_{\bar{A}}^{I} d_l(A)dA}{\partial i_1} = \frac{1}{F\left(\dfrac{\delta R}{1+i_2}\right)}\left[\frac{-\varkappa}{(1+i_2)(1-\delta)} - \left(R - \frac{B}{\Delta P}\right)\frac{\partial \int_0^{\bar{A}} F(\rho^*(A))dA}{\partial i_1}\right]$$

which is negative, as the data suggests, if

$$\frac{-\varkappa}{(1+i_2)(1-\delta)} - \left(R - \frac{B}{\Delta P}\right)\frac{\partial \int_0^{\bar{A}} F(\rho^*(A))dA}{\partial i_1} < 0.$$

Plugging in the explicit expressions into this inequality, we have

$$\frac{\varkappa}{(1+i_2)(1-\delta)} > \left(R - \frac{B}{\Delta P}\right)\int_0^{\bar{A}} f(\rho^*(A))\frac{\dfrac{r-\rho^*(A)}{1+i_1} + \dfrac{\delta F(\rho^*(A))}{(1+i_1)(1+i_2)}\left(R - \dfrac{B}{\Delta P}\right)}{1 - \dfrac{\delta f(\rho^*(A))}{1+i_2}\left(R - \dfrac{B}{\Delta P}\right)}dA.$$

The right-hand side of this inequality is maximized by setting $\rho^*(A) = 0$, so as low as possible, in this case $f(0) = \chi$ and $F(0) = 0$. Thus, in this way, the right-hand side simplifies to

$$\frac{\varkappa}{(1+i_2)(1-\delta)} > \bar{A}\left(R - \frac{B}{\Delta P}\right)\chi\frac{\dfrac{r}{1+i_1}}{1 - \dfrac{\delta\chi}{1+i_2}\left(R - \dfrac{B}{\Delta P}\right)}.$$

Inspecting this result, we can see that if $\varkappa$ is large enough, namely larger than ϕ

$$\phi = \frac{(1+i_2)\chi(1-\delta)\bar{A}\left(R - \dfrac{B}{\Delta P}\right)\dfrac{r}{1+i_1}}{1 - \dfrac{\delta\chi}{(1+i_2)}\left(R - \dfrac{B}{\Delta P}\right)},$$

then we have a sufficient condition for

$$\frac{\partial \int_{\bar{A}}^{I} d_l(A)dA}{\partial i_1} < 0.$$

In order to prove that the aggregate change of LT debt of the unconstrained firms generalizes to individual firm behaviour, namely

$$\frac{\partial \int_{\bar{A}}^{I} d_l(A)dA}{\partial i_1} < 0 \implies \frac{\partial d_l(A)}{\partial i_1} < 0 \text{ for A}\in(\bar{A}\text{,I}],$$

we recall the assumption that unconstrained firms choose at least the minimum LT debt that allows them to equalize the highest LT debt share of the constrained firms, which is the LT debt share of the firm with endowment $\bar{A}$. This minimum component can be expressed as:

$$d_l^{min}(A) = \frac{(1+i_1)(I-A)}{\dfrac{1-\kappa}{\kappa} + \dfrac{\delta F\left(\dfrac{\delta R}{1+i_2}\right)}{1+i_2}}.$$

Here,

$$\kappa = \frac{R - \dfrac{B}{\Delta p}}{R - \dfrac{B}{\Delta p} + r - \dfrac{\delta R}{1+i_2}}$$

is the LT debt share of the firm with endowment $\bar{A}$. Additionally, we impose that for any unconstrained firm $d_l^{min}(A) < d_l(A) < (R - B/\Delta p)$, and that any change in the aggregate LT debt of unconstrained firms is distributed to all unconstrained firms as a change in their LT debt, and that this change, relative to the aggregate change, for any subset of unconstrained firms that has non-zero measure, is larger than zero.

Now, consider an infinitesimal increase in i_1: it will increase the minimum amount of LT debt needed to match the highest LT debt ratio of constrained firms. Concretely:

$$\frac{\partial d_l^{min}(A)}{\partial i_1} = \frac{I-A}{\dfrac{1-\kappa}{\kappa} + \dfrac{\delta F\left(\dfrac{\delta R}{1+i_2}\right)}{1+i_2}}.$$

However, as for all $A \in (\bar{A}, I]$ the choice of LT debt before the increase was strictly higher than the minimum LT debt due to our assumption, the infinitesimal change in the minimum LT debt will not make it surpass the previous amount. Instead, as the aggregate LT debt for the unconstrained firms must decrease, and this decrease is distributed to all unconstrained firms, we find that $\partial d_l(A)/\partial i_1 < 0$ for A∈($\bar{A}$,I]. For an infinitesimal decrease in i_1, the aggregate increases while $d_l^{min}(A)$ decreases, thus in this case $d_l(A)$ must increase for all firms with A∈($\bar{A}$,I]. This concludes the proof of statement 1.

Moving on, from

$$\frac{\partial \int_{\bar{A}}^{I} d_l(A)dA}{\partial i_1} = \frac{1}{F\left(\dfrac{R}{1+i_2}\right)}\left[\frac{-\varkappa}{(1+i_2)(1-\delta)} - \left(R - \frac{B}{\Delta P}\right)\frac{\partial \int_0^{\bar{A}} F(\rho^*(A))dA}{\partial i_1}\right],$$

we can see that if $\varkappa \to \infty$ then

$$\frac{\partial \int_{\bar{A}}^{I} d_l(A)dA}{\partial i_1} \to -\infty.$$

Due to the assumptions that a change in the aggregate must be shared across firms and furthermore each share cannot be trivially small, by the same logic as above

$$\frac{\partial \int_{\bar{A}}^{I} d_l(A)dA}{\partial i_1} \to -\infty \implies \frac{\partial d_l(A)}{\partial i_1} \to -\infty \text{ for A}\in(\bar{A}\text{,I]}.$$

For the constrained firms, $\rho^*(A)$ is pinned-down by the constraints, in which $\varkappa$ plays no role, as it does not affect the price. A change in $\varkappa$ thus has no effect on the choice of $\rho^*(A)$ for firms with $A < \bar{A}$. This concludes the proof of statement 2 of the proposition.

As shown in Lemma 2.2, constrained firms, those with $A \leq \bar{A}$, increase $\rho^*(A)$ in i_1. Their LT debt is determined by the constraint to be $d_l(A) = (R - B/\Delta p)$. This proves statement 3 of the proposition. □

2.E Extensions

2.E.1 Effect of of Policy Rate Change with Risk-neutral Rational Investors

To showcase the effect of our departure from a model with only rational and risk-neutral investors, we analyze a baseline specification, in which we assume that investors are risk neutral and rational: $\gamma = \alpha = 0$. In this case the investor demand function is horizontal, which means that they are willing to hold any amount of LT debt at price

$$P_l(A) = \frac{\delta F(\rho^*(A))}{(1+i_1)(1+i_2)}.$$

This is also the price at which a firm with endowment A would inelastically sell LT debt. Thus, the total amount of LT debt is not pinned down by market clearing. The requirements for constrained and unconstrained firms to exists are the same as derived above, and it is still true in this benchmark case that for constrained firms

$$\frac{\partial \rho^*(A)}{\partial i_1} < 0.$$

However, because $\int_0^I P_l(A) d_l(A) dA$ is not pinned down by a downward sloping investor demand curve, we cannot say whether or how unconstrained firms adjust their maturity structure. Thus, we have that constrained firms take on less short-term debt when the the monetary authority eases, while unconstrained firms have no incentive to change their maturity structure. This is counterfactual, as we see in the data that unconstrained firms should lengthen their maturity structure and do so more than constrained firms.

Chapter 3

DOES DISPERSED SENTIMENT DRIVE RETURNS, TURNOVER, AND VOLATILITY FOR BITCOIN?

3.1 Introduction

A large literature has studied the effect of investor disagreement on returns for different asset classes and periods and with ambiguous results. Generally, the literature discusses two possible opposing mechanisms: (i) in the presence of short-sale constraints, investor disagreement drives up prices, as optimists hold the assets, and returns will be low, and (ii) investor disagreement represents higher uncertainty and thus warrants a higher return for holding the asset.[1] The first mechanism, known as the differences-of-opinion channel, also predicts high turnover and price volatility when investor disagreement is high.

The differences-of-opinion literature is built on the key theoretical insight that if pessimists cannot participate in the market due to high short-sale costs, the asset price will be higher than the fundamental value, leading to subsequent low returns. Furthermore, as opinions fluctuate and trade becomes more likely, disagreement leads to high

[1]For an extensive discussion, see Diether, Malloy, and Scherbina (2002).

volatility and high turnover. These predictions have been derived from a long theoretical literature, starting from Miller (1977) and Harrison and Kreps (1978), later developed into behavioral agree-to-disagree models such as Hong and Stein (2003), Scheinkman and Xiong (2003), Hong, Scheinkman, and Xiong (2006), and more recently Simsek (2013).[2]

A crucial obstacle to testing the predictions of the differences-of-opinion literature has been that investor sentiment is not directly observable. Different proxies have been used, such as analyst opinions or newspaper articles (Sadka and Scherbina 2007). Nowadays, the availability of extensive online discussions about assets allows us to analyze statements and opinions issued by individuals who are potential investors. A seminal paper following this approach is Antweiler and Frank (2004), who analyzed online posts on Yahoo Finance and Raging Bull to predict market volatility and asset returns. Other papers analyzing asset characteristics using different dictionary-based algorithms are Tetlock (2007), Loughran and McDonald (2011), and Jegadeesh and Wu (2013).

In this paper, we exploit the magnitude of online discussion about a highly speculative asset on which opinions are widely divided (Bitcoin) to test a theory of investor disagreement and short-sale constraints. We scrape millions of online comments across a decade of discussion from a Bitcoin-focused online forum and extract sentiment using the lexicon- and rule-based sentiment algorithm called VADER (Hutto and Gilbert 2014), which is specifically trained for online data sets. Our contribution is to explore the joint time-series behavior of this sentiment measure, as well as its dispersion, on the one hand, and Bitcoin's return, turnover, and price volatility, on the other. Our approach allows us to test the predictions of the differences-of-opinion literature in a rich setting of textual data at daily, weekly, and monthly frequency. We argue that Bitcoin is the ideal asset to test these predictions, as it is complicated to judge its value (Bitcoin will never pay dividends). Therefore, opinions on Bitcoin's value differ widely. Moreover, institutionally and due to substantial price volatility, it is difficult to short Bitcoin.

We find that there is a significantly negative predictive relationship between disagreement and the return on holding Bitcoin. Disagreement forecasts negative returns

[2]For a full overview of the differences-of-opinion literature, see Hong and Stein (2007a) and Xiong (2013). Simsek (2021) provides an overview of the macroeconomic implications of investor disagreement.

into the future at the daily, weekly, and monthly frequency. This empirical finding is consistent with the theoretical predictions of the differences-of-opinion literature. The effect is especially strong and predicts low returns several months into the future if sentiment and returns are measured at a monthly frequency, which we interpret as overpricing resolving slowly over time. The association between contemporaneous returns, average sentiment, and disagreement is economically significant as well: the adjusted R^2 is 0.33 at monthly frequency. A one standard deviation increase in disagreement leads to a negative return of about -9.2% over the following eight weeks. This is around 12% of the standard deviation of the eight-week returns for Bitcoin.

Although disagreement predicts low returns, which can be interpreted as a sign of overpricing, disagreement is not positively related to contemporaneous or past returns. This finding seems at odds with the usual understanding of the differences-of-opinion channel, predicting that an increase in disagreement first leads to positive returns and overpricing. However, the literature usually assumes that an asset's fundamental value is independent of investor beliefs or disagreement, but this might not be the case for a purely belief-driven asset such as Bitcoin. In this case, the emergence of disagreement could erode the coordination of beliefs among Bitcoin investors, which is key to the asset's value proposition. A slow adjustment of beliefs on the side of optimists can then lead to a situation in which disagreement predicts low returns in the medium term without initially increasing the price.

We study the consequences of the easing of short-sale constraints for Bitcoin starting in December 2017.[3] We find that, as the literature would predict, the effect of disagreement on returns diminishes significantly towards the end of our sample. However, shorting Bitcoin remains expensive and risky, as margin requirements are high compared to other assets, and Bitcoin's price is extremely volatile.

Extending our analysis to volatility and turnover, we find that disagreement has a strong and significant effect on price volatility and turnover growth. Higher disagreement leads to persistently more trading at the same time as a short-lived increase in volatility. These findings generally are consistent with the predictions of the differences-of-opinion literature. Our findings are also economically significant in this case. In the regression, at a monthly frequency, the adjusted R^2 is 0.06 for turnover growth and 0.34

[3]The CME Group started offering futures contracts for Bitcoin only in December 2017 (CME Group 2017) and options on futures in January 2020 (CME Group 2019).

for volatility.

We contribute by extending the literature about disagreement to a speculative asset with a market capitalization that has increased to over a trillion US dollars since 2010.[4] This makes cryptocurrency assets worth serious scientific attention, despite their quirkiness and novelty. The determinants of their pricing and asset characteristics are interesting in their own right, even from a public policy perspective: a collapse of cryptocurrency prices (e.g., optimists could become disillusioned and leave the market) would destroy immense wealth.

The remainder of the paper is organized as follows. Section 3.2 explains the mechanism behind the results in the differences-of-opinion literature in a stylized model and contrasts it with other possible explanations, namely the idea that disagreement is just a symptom of underlying uncertainty and the that disagreement today is simply driven by low past returns. Section 3.3 details how we collected the data and conducted our sentiment analysis. Then, Section 3.4 tests the derived relationships empirically and finds substantial support for the predictions of the theoretical literature. In Section 3.5 we interpret our results. Section 3.6 concludes.

3.2 Model

3.2.1 Disagreement in a Differences-of-Opinion Model

We present a simple discrete-time model of heterogeneous beliefs and limits to arbitrage[5] to motivate our empirical analysis. There are overlapping generations of risk-neutral traders indexed by i who each live for two periods and maximize end-of-life consumption.[6] The utility function of trader i is

$$U_{it} = \mathbb{E}_{it}\{C_{it+1}\}. \tag{3.1}$$

When young, traders either buy a long-lived asset from the old or invest in a risk-less bond with return $R > 1$. Traders are split into two groups - optimists and pessimists -

[4]As of April 2021. The whole market capitalization of all cryptocurrencies has pushed past two trillion USD.

[5]For a comprehensive review of the theoretical differences-of-opinion literature, see Simsek (2021).

[6]Risk-neutrality is chosen to present the differences-of-opinion channel in the cleanest way. Risk-aversion is considered in Section 3.2.2.

who have diverging beliefs about the asset's value, in this case, Bitcoin. Traders have deep pockets, such that optimists have sufficient wealth to buy up the asset supply, and pessimists must stay out of the market due to short-sale constraints. Therefore, the equilibrium price will be determined by the optimists' beliefs only.

The features of the model's long-lived asset capture the essence of Bitcoin in reduced form. Bitcoin investors believe that a coordinated and permanent shift in beliefs will make Bitcoin valuable as a store of value and currency with some positive probability.[7] We denote this absorbing event as A_t if it takes place in period t. Because Bitcoin's protocol limits its supply, the value of one Bitcoin in this event can be derived according to a quantity theory of money type equation, which we denote as $\bar{P}$.[8] The probability of the collective shift is fixed over time to $\mathcal{P}(A_t) = \phi \in (0, 1)$. Once the event has taken place, all traders agree on Bitcoin's price being fixed to $\bar{P}$ for all future periods. As our focus is on explaining short- to medium-term fluctuations in the price of Bitcoin and not the long-term trend, we assume that ϕ is fixed over time.

However, traders believe that the probability ϕ is time-variant. In particular, they have heterogeneous beliefs about $\mathcal{P}(A_{t+1})$ but agree that the probability is fixed to ϕ from $t + 2$ onward. Within each group, beliefs are homogeneous.[9] The beliefs of group i are distributed as $\phi^i_s \overset{iid}{\sim} G(\phi^i_s)$, where $G(\cdot)$ is the continuous cumulative distribution function of ϕ^i_s over the interval $[0, 1]$.[10] We refer to the group that attributes a higher probability to A_{t+1} as *optimists* and the other group as *pessimists* ($\phi^o_{t+1} > \phi^p_{t+1}$).

Due to perfect competition between optimists, the price in period t is

$$P_t = \frac{1}{R} \left(\phi^o_{t+1} \mathbb{E}(P_{t+1} | A_{t+1}) + (1 - \phi^o_{t+1}) \mathbb{E}(P_{t+1} | \neg A_{t+1}) \right), \tag{3.2}$$

where $\mathbb{E}(P_{t+1}|A_{t+1}) = \bar{P}$. Optimists believe that with probability ϕ^o_{t+1} a belief shift will take place and they will sell the asset at a price $\bar{P}$ when old. Otherwise, they will sell the asset at a price P_{t+1} that depends on the beliefs of tomorrow's optimists. This

[7]For example, as of 9 June 2021, Bitcoin became legal tender in El Salvador.

[8]We do not model the determinants of $\bar{P}$ explicitly, but instead focus on the relationship between disagreement and the price today, while taking $\bar{P}$ as given.

[9]The assumption of homogeneous in-group beliefs is not crucial and can be replaced with heterogeneous beliefs inside the group. Traders agree to disagree and do not learn from the price. This assumption can be relaxed by assuming that traders are overconfident.

[10]The iid assumption is made for simplicity to highlight that prices and beliefs are eventually mean-reverting, which would also hold if beliefs were somewhat persistent over time.

leads us to the main prediction of the model.

Proposition 3.1. *Returns* $\frac{P_{t+1}-P_t}{P_t}$ *are **decreasing** in disagreement* $(\phi^o_{t+1}-\phi^p_{t+1})$ *holding the average belief* $(\phi^o_{t+1}+\phi^p_{t+1})/2$ *constant.*

The intuition for this result is that when disagreement is high, the overoptimism of optimists is more severe, which depresses returns going into the future. Naturally, such overoptimism increases the price initially, which leads to the following Corollary.

Corollary 3.1. *Past returns* $\frac{P_t-P_{t-1}}{P_{t-1}}$ *are **increasing** in disagreement* $(\phi^o_{t+1}-\phi^p_{t+1})$ *holding the average belief* $(\phi^o_{t+1}+\phi^p_{t+1})/2$ *constant.*

The model can also be extended to yield predictions for turnover when introducing convex costs for short-sales instead of assuming that such costs are infinite. For simplicity, assume that the costs are sufficiently high such that the marginal buyer remains an optimist.

Proposition 3.2. *With convex short-sale costs, turnover is **increasing** in disagreement* $(\phi^o_{t+1}-\phi^p_{t+1})$ *holding the average beliefs* $(\phi^o_{t+1}+\phi^p_{t+1})/2$ *constant.*

Convex short-sale costs are chosen for simplicity to relate disagreement and turnover. Generally, increased disagreement positively influences the perceived gains from trade, leading to additional traders entering the market which can increase turnover even when the amount of short-sale is exogenously fixed per trader.

Volatility has a more ambiguous relationship with disagreement than returns or turnover. For example, if traders were long-lived and beliefs fully persistent, prices would be constant for any level of disagreement. As beliefs are short-lived in our stylized model, the relationship between volatility and disagreement is, in principle, nonlinear. If today's optimists are pessimistic relative to the average optimists over time ($\phi^o_{t+1} < \mathbb{E}(\phi^o|\phi^o > \phi^p)$), then higher disagreement can move today's prices closer to the historical average while keeping the average belief constant, decreasing expected volatility.

Still, expected volatility can be positively related to today's disagreement when the distribution of beliefs $G(\phi^i_s)$ is subject to variance shocks. In this case, more disagreement today can indicate more volatile beliefs in the future, which increases expected price volatility. Therefore, we expect volatility to be positively related to disagreement.

Similarly, our predictions for returns and turnover can be extended dynamically if we think about an increase in disagreement stemming from such a mean-reverting variance shock to $G(\phi^i_s)$. Then, disagreement is persistent and overpricing due to disagreement does not resolve immediately, leading to protracted negative returns. Also turnover and volatility remain alleviated for several periods.

Before we turn to our empirical approach, note that the model is deliberately kept simple to provide intuition for the main mechanisms. Moreover, the model is stationary conditional on the base probability of adoption ϕ, whereas asset prices are usually non-stationary as returns follow random walks. Therefore, our simple model is best used to explain fluctuations that happen around the trend in Bitcoin's price. As the price and turnover clearly show non-stationary behavior in Figure 3.3.1, we will use returns, the growth rate of turnover and dispersion in hourly returns as left-hand-side variables in our empirical analysis.[11]

Before moving to our empirical analysis, we present two alternative models through which the relationship between disagreement and returns can be interpreted.

3.2.2 Disagreement as Uncertainty

Large disagreement among traders can be a sign of fundamental uncertainty. When uncertainty increases, risk-averse traders require higher future returns to absorb the risk, which leads to a fall in the price today.

To capture this intuition, consider a model with overlapping generations of representative traders[12] with CARA-utility

$$U_t = 1 - \exp\left(-\gamma W_{t|t+1}\right). \tag{3.3}$$

where $\gamma \geq 0$ is the coefficient of absolute risk aversion and $W_{t|t+1}$ is end-of-period

[11] Although beliefs almost solely drive price movements and trading in Bitcoin, our approach of extracting sentiment from text is not well-suited to explain the long-term price movements in Bitcoin. Sentiment analysis is more appropriate to measure how prevalent relatively positive or negative sentiment is in a given moment in time, which can be an important determinant for short- to medium-run price movements.

[12] The same result would also hold when considering a model with a mass of heterogeneously informed traders as in Grossman and Stiglitz (1980). In that case, an increase in fundamental risk leads to higher disagreement and posterior uncertainty. A decrease in the precision of private signals can have similar effects.

wealth of the representative trader born in period t, who is free to borrow or lend at interest rate $R > 1$. Without loss of generality, initial wealth is normalized to zero. Every period, the old trader sells a single risky asset to the young trader. Otherwise, the asset characteristics are unchanged.

As before, traders believe that the probability of the adoption event A_{t+1} is time-variant. But this time, their beliefs are uncertain. At the beginning of the period, the representative trader draws a belief over ϕ_{t+1} with finite mean and positive variance. Beliefs are iid across generations.

The price P_t is derived from the representative trader being indifferent between holding the asset or lending out P_t at interest rate R,

$$1 - \mathbb{E}_t \left\{ \exp\left(-\gamma P_{t+1}\right) \right\} = 1 - \exp\left(-\gamma R P_t\right), \tag{3.4}$$

leading to

$$P_t = \frac{-\log \mathbb{E}_t \left\{ \exp\left(-\gamma P_{t+1}\right) \right\}}{\gamma R}. \tag{3.5}$$

The price P_t depends on the representative trader's expectations about the probability of adoption ϕ_{t+1}. Due to risk-aversion, a mean-preserving increase in uncertainty about ϕ_{t+1} must lead to a lower price P_t, which is captured in the following proposition.

Proposition 3.3. *Returns* $\frac{P_{t+1}-P_t}{P_t}$ *are* ***increasing*** *in the representative trader's variance of beliefs on the probability of the adoption event* ϕ_{t+1}.

As before, a fall in today's price due to higher uncertainty must also mean that past returns were negative, which leads to the following corollary.

Corollary 3.2. *Past returns* $\frac{P_t-P_{t-1}}{P_{t-1}}$ *are* ***decreasing*** *in the representative trader's variance of beliefs on the probability of the adoption event* ϕ_{t+1}.

It follows that viewing disagreement as a proxy of uncertainty leads to exactly opposite predictions on the relationship between disagreement and returns compared to the differences-of-opinion model. When traders are risk-averse, they require a higher return when absorbing greater risks, leading to falling prices. In contrast, in the presence of overconfidence and short-sale constraints, higher disagreement means that price-setting optimists are increasingly over-optimistic, inflating the price today and leading to low future returns.

3.2.3 Sentiment and Disagreement as a Side-Show

As we set out to derive a proxy for sentiment and disagreement from posts in an online forum, reverse causality is a plausible concern. Such posts may merely react to price movements but not reveal any information that could be useful to predict future returns. In the following, we suppose that sentiment is a function of lagged returns.

$$\text{Sent}_{it} = \alpha_i + \sum_{s=0}^{S} \gamma_{is} \left(\frac{P_{t-s} - P_{t-1-s}}{P_{t-1-s}} \right) + \varepsilon_{it} \tag{3.6}$$

where $S < \infty$ and $\varepsilon_{it} \overset{iid}{\sim} \mathcal{N}(0, \sigma_\varepsilon^2)$. Finally, we use in our analysis average sentiment and the dispersion in sentiments as disagreement, formally

$$\text{Sent}_t = \sum_i \text{Sent}_{it}. \tag{3.7}$$

$$\text{Dis}_t = \sqrt{Var(Sent_{it})} \tag{3.8}$$

Naturally, sentiment should react positively to current, and past returns ($\forall s : \gamma_{is} > 0$) as investors profit from positive returns.[13] Less clear is the relationship between disagreement and past returns. One possible explanation draws on **confirmation bias**. Following this idea, investors in Bitcoin may be likely to disregard information that does not match their prior.

Suppose that investors are split into two groups. The first group consists of dogmatically optimistic traders (high α_o), who do not revise their beliefs in the face of new information ($\gamma_{os} \approx 0$).[14] The second group is composed of less optimistic traders ($\alpha_p < \alpha_o$) with more flexible beliefs ($\gamma_{ps} > 0$). As a result, both groups hold more similar beliefs when returns are positive and disagree more intensely when returns are negative. Therefore, we would expect to see a negative relationship between past returns and disagreement.

To derive predictions about the predictability of returns using sentiment and dis-

[13] A positive relationship may also be plausible when interpreting sentiments as expectations about future returns. Greenwood and Shleifer (2014) show that investors increase their expectations of future returns after positive past returns.

[14] It is also possible to assume that there are dogmatic pessimists, but attributing dogmatism to optimists is in line with anecdotal evidence of a fraction of Bitcoin investors who buy Bitcoin and hold it for extended periods irrespective of news. Moreover, dogmatic pessimists should eventually leave the market.

agreement, we consider two returns processes. First, prices may follow a random walk and returns are white noise. In this case, sentiment or disagreement cannot forecast returns as they are not correlated with the innovations to the price. Second, returns may be autocorrelated, which can lead to sentiment and disagreement predicting future returns.[15] Nonetheless, sentiment and disagreement should lose their predictive power when controlling for lagged returns.

3.3 Data

We use publicly available data from the *Kraken.com* exchange for the opening and closing price of Bitcoin and an aggregated measure of turnover across all major exchanges from *coinmarketcap.com*. We compute daily returns by dividing the difference between closing and opening prices by the opening price and turnover as the daily dollar volume divided by Bitcoin's total market capitalization.[16] We compute a volatility measure as the standard deviation of hourly returns in a given day, week, or month.

We relate Bitcoin's market characteristics to sentiment changes among Bitcoin investors. For this purpose, we scrape the Bitcoin-related online-forum *bitcointalk.org* using the python package *Scrapy*. In particular, we scrape all threads and comments from the *Speculation* subforum, which most closely covers discussions on Bitcoin's price movements and expectations about future price developments. We gathered 1,482,589 comments that were posted between 18 October 2010 and 21 April 2021.

We gathered comments from 54,173 unique accounts, of which 7,183 accounts opened discussion topics. Posting activity follows a power law, as the top ten percent of most active accounts (more than 32 posts) produce more than 84% of all content. In contrast, the median number of posts per account is three. At the same time, no single account dominates the discussion, as the most active account wrote 1.3% of all posts (19,758 in total or five posts per day). Overall concentration is low with a Herfindahl in-

[15]Positive autocorrelation may arise when new information is only gradually incorporated (McQueen, Pinegar, and Thorley 1996), whereas negative autocorrelation may stem from overreaction to new information (Lo and MacKinlay 1990).

[16]Since crypto exchanges are open 24/7, the opening and closing prices are the earliest and latest price available in a specific period according to UTC.

dex of 0.0011.[17] According to *bitcointalk.org*,[18] there are in total over three million registered users, and more than a million page views a day. With this reach, *bitcointalk.org* is an important medium in discussions related to cryptocurrencies.[19]

We use all comments with non-zero valence in our analysis, as the Speculation subforum is already focused on Bitcoin's price movements. The forum allows users to quote other comments in their posts. We filter out such repetitions as quotes and keep only the new part of each post.[20] We run our main specification starting 1 January 2014, as the number of posts per day reaches a higher and more stable level from 2014 on. Figure 3.C.2 provides a word cloud with the most commonly used words of a random sample of 10,000 comments.

Figure 3.3.1 summarizes the time series of the main variables at a weekly frequency: the price level, turnover, price volatility measured as the standard deviation of hourly returns, average sentiment, and dispersion in sentiments. Additionally, the right upper panel displays the number of posts on the Speculation subforum of *bitcointalk.org*. As both the price and turnover display a clear trend, we will use their growth rates. All time series have substantial time variation, which we exploit in our empirical analysis. Bitcoin's price follows a distinct boom and bust cycle. Turnover and volatility increase when Bitcoin's price increases or decreases rapidly (e.g., the boom leading up to December 2017 or the short-term bust in 2019.). Posts per week are also cyclical and peak at around 7,000 posts per week at the beginning of 2018. Finally, although average sentiment and the standard deviation of sentiment is relatively noisy week-to-week, both time series show persistence at lower frequencies.

Table 3.1 provides summary statistics (mean, standard deviation, the first and ninth decile cut-offs, as well as the median) for our main variables of interest: sentiment,

---

[17] We show the time series of the number of active users in Figure 3.C.1a, which is positively correlated to the overall number of posts. In Figure 3.C.1b, we show that the Herfindahl Index is stable over time and decreased during the surge in activity in 2018.

[18] See `https://bitcointalk.org/index.php?action=stats` (accessed 25 June 2021, 20:00).

[19] For example, today's second-largest cryptocurrency *Ether* and its Initial Coin Offering were first announced on *bitcointalk.org* in January 2014: `https://bitcointalk.org/index.php?topic=428589.0` (accessed 25 June 2021, 20:00).

[20] We do not attempt to weigh posts according to importance (e.g., through their number of views or quotes), but instead attribute the same weight to every post. Although an abundance of quotes potentially reflects the greater importance of the quoted post, we find that filtering out quotes increases the explanatory power of our sentiment and disagreement measure in all regressions.

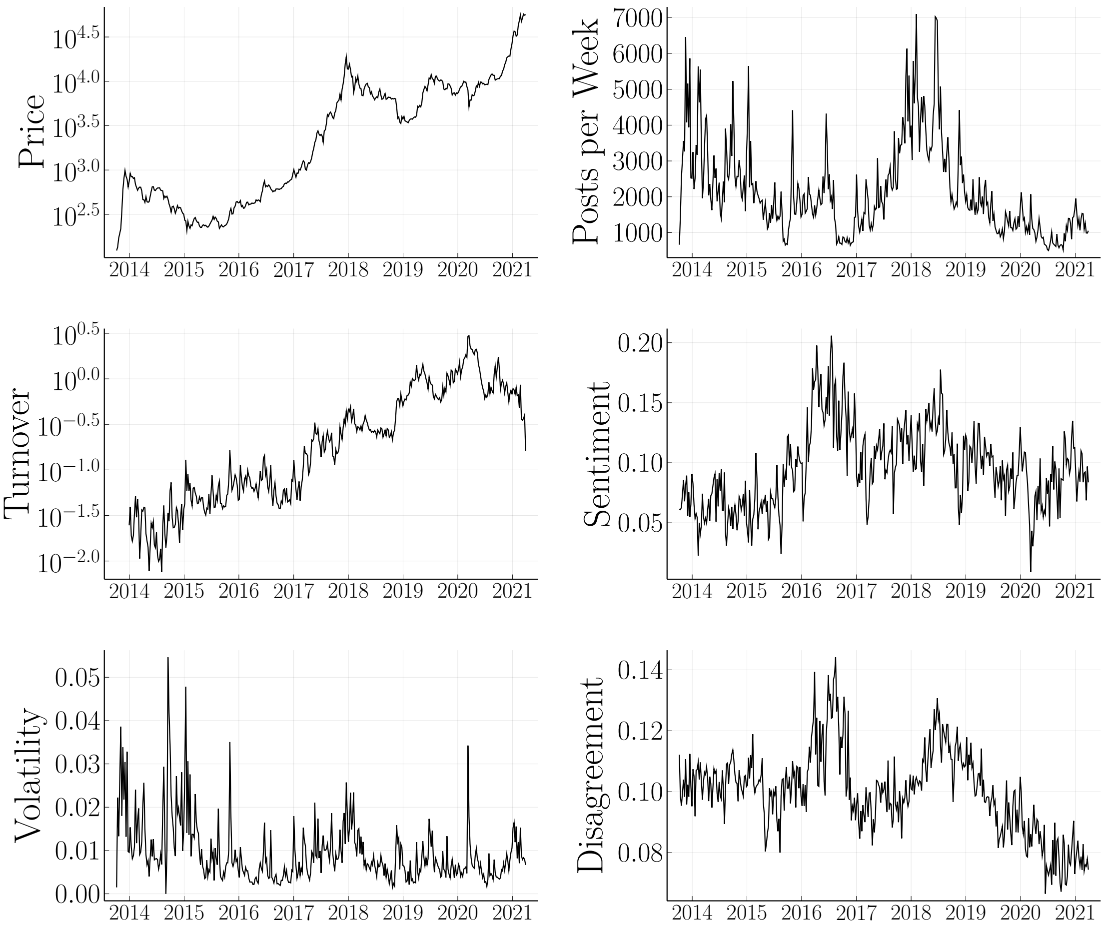

Figure 3.3.1: Overview over the Main Variables at Weekly Frequency.

Notes: On the left: returns (difference between closing and opening price divided by opening price), turnover (volume divided by market capitalization) and volatility (standard deviation of hourly returns). On the right: number of posts, sentiment (the mean of the comment sentiment distribution) and disagreement (the standard deviation of the comment sentiment distribution).

disagreement, returns, turnover growth, and volatility. All statistics are shown at daily, weekly, and monthly frequency. Strikingly, mean returns for Bitcoin are quite high, with 7.6% monthly. Volatility is also high, with an average standard deviation of hourly returns of around 0.9 percentage point.

	Frequency	Mean	SD	Q10	Median	Q90
Sentiment	daily	0.097	0.042	0.044	0.096	0.15
	weekly	0.096	0.034	0.054	0.095	0.14
	monthly	0.096	0.031	0.056	0.094	0.136
Disagreement	daily	0.099	0.019	0.076	0.099	0.122
	weekly	0.1	0.014	0.081	0.1	0.118
	monthly	0.1	0.013	0.08	0.101	0.116
Return	daily	0.3%	4 %	-3.9 %	0.2%	4.4 %
	weekly	1.6%	10.7%	-12.4%	1.1%	15.6%
	monthly	7.6%	23.2%	-18.9%	6.2%	37.4%
Turnover Growth	daily	7 %	49 %	-31.5%	-1.3%	45.2%
	weekly	5.5 %	35.8%	-29.3%	-1.5%	46.5%
	monthly	9.3 %	40.6%	-28.7%	0 %	64.1%
Volatility	daily	0.8 %	0.77%	0.22 %	0.58%	1.57%
	weekly	0.91%	0.68%	0.34 %	0.72%	1.68%
	monthly	0.99%	0.64%	0.42 %	0.78%	1.84%

Table 3.1: Summary Statics of Sentiment, Disagreement, Returns, Turnover Growth, and Volatility of Bitcoin.

Notes: Mean, standard deviation, first decile, median and ninth decile of the main variables. Statistics for returns, volatility, and turnover growth are in percentage points.

3.3.1 Sentiment Analysis using VADER

We use a lexicon and rule-based algorithm called VADER (**V**alence **A**ware **D**ictionary and s**E**ntiment **R**easoner) for the sentiment analysis.[21] The underlying lexicon and algorithm are specialized for the analysis of social media posts (see Hutto and Gilbert 2014,

[21] A detailed description of VADER can be found on Github: `https://github.com/cjhutto/vaderSentiment` (accessed 25 June 2021, 20:00).

for a comparison with other lexica).[22]

A sentiment lexicon is a mapping from "tokens" (words, stems of words, abbreviations, etc.) to a numerical indicator of sentiment. Each token carries a certain valence (negative, neutral, or positive sentiment) irrespective of context. These valence intensities were generated by letting ten independent human raters rate tokens. The final valence is the average of the individual ratings (Wisdom of the Crowd approach). All human raters had been pre-screened, trained, and quality checked. Following this approach, over 9000 tokens were rated on a scale from "[-4] Extremely Negative" to "[4] Extremely Positive" with an option to rate the token as "[0] Neutral." Already existing established lexicons inspired the list of tokens (e.g., LIWC, ANEW, and GI) to which Western-style emoticons (e.g., ":-)"), sentiment-related acronyms and initialism (e.g. "LOL", "ROFL") and commonly used slang (e.g., "nah," "meh") were added. After dropping tokens that ended up with a neutral mean-sentiment rating or a standard deviation of individual ratings higher than 2.5, about 7500 tokens were left and rated on the -4 to $+4$ scale.

Although relying on a lexicon for sentiment analysis, VADER is not a bag-of-words algorithm that neglects the syntax and order of words. Instead, VADER employs five simple rules to improve its sentiment ratings for whole sentences. First, punctuation is included by using the exclamation point (!) as an intensifier. Secondly, capitalization increases the sentiment intensity. Thirdly, modifiers are used to adjust the intensity. With the corresponding valence between -1 (very negative) and 1 (very positive) computed by VADER, "Bitcoin has a bright future" (0.44) is less intense than "Bitcoin has a very bright future" (0.49) and more intense than "Bitcoin has a somewhat bright future" (0.38). Fourthly, the conjunction "but" is used to signal a reversal of semantic orientation. For example, "Bitcoin had a great year, but has a lot of problems" (-0.25) conveys negative sentiment, although the initial statement is positive. Lastly, the three words before a sentiment-laden token are included in the sentiment rating to check for words that flip the semantic orientation. For example, "Bitcoin does not have a great future" (-0.34) conveys negative sentiment, although "great" carries positive sentiment.

To sum up, VADER is an appropriate sentiment analysis tool for the domain of

[22]The authors show that VADER can produce valence ratings with high correlation to human mean-sentiment ratings. In particular, run on a corpus of over 4000 Tweets, sentiment, as calculated by VADER, had the largest correlation (0.88) and R^2 (0.77) to the human mean-sentiment rating.

our investigation. We use the "compound" measure, a weighted average of sentiment normalized to values between -1 (extremely negative) and 1 (extremely positive). As suggested by the package authors, we compute the sentiment index for each comment on the sentence level and use the mean to compute comment-level sentiment. Finally, we aggregate the sentiment data at different frequencies and use the mean to measure the level of daily, weekly, and monthly sentiment. We use the standard deviation of variance as a proxy for disagreement among investors.

3.4 Empirics

Our empirical approach is to extract a sentiment measure from comments on *bitcointalk.org* and use this measure as a proxy for beliefs about the success of Bitcoin (ϕ^i_{t+1} in the model). In particular, we think of our sentiment measure as being relative to some time-variant base level of expectations (e.g., a time-variant ϕ). In that way, high sentiment can be interpreted as expectations of positive returns at any point in time. Henceforth, we refer to the valence measure as computed by VADER from each comment simply as sentiment.

Whereas we capture an average stance of sentiment through the first moment the sentiment distribution, we define disagreement as the dispersion in sentiment. If comments with positive and negative sentiment are posted during the same period, we interpret such dispersion as a sign of high disagreement. We use these measures to analyze the effect of disagreement, conditional on average sentiment, on the return, turnover growth, and volatility of Bitcoin.

All our regressions are summarized by the following equation,

$$X^j_{t+s} = \alpha^j_s + \sum_{l=0}^{L} (\beta^{j,l}_{\mu,s} \text{Sentiment}_{t-l} + \beta^{j,l}_{\sigma,s} \text{Disagreement}_{t-l}) + \varepsilon_{t+s}. \tag{3.9}$$

We use returns, turnover growth, and the dispersion in hourly returns (volatility) as the left-hand-side variable X^j_{t+s} where j stands for each different variable. We run the regression at different leads and, for returns, lags s. Moreover, we also use long-horizon returns with overlapping observations as the left-hand-side variable, in which case X^j_{t+s} stands for the return between the beginning of period $t+1$ and the end of period $t+s$.

We include up to L lags of sentiment and disagreement and estimate Model 3.9 for each variable at daily, weekly, and monthly frequency. Throughout, we apply HAC-robust standard errors following Newey and West (1987). To address concerns due to the persistence of our regressors, we repeat our forecasting regressions with confidence intervals computed according to Campbell and Yogo (2006).[23] Additionally, for the regression with long-horizon returns, we adjust our confidence intervals according to Hjalmarsson (2011), which additionally increases the bandwidth as our forecasting-horizon lengthens.

3.4.1 Return Regressions

We set out to predict returns of Bitcoin through sentiment and disagreement. As a first step, we present evidence that the price of Bitcoin is indeed predictable while not taking a stance on the specific predictor. For this purpose, we use the variance ratio test of Lo and MacKinlay (1988). The idea of the test is that if prices move randomly, the variance of returns should increase linearly in the horizon. If this assumption is violated, returns are not random and can potentially be predicted.

	2 Lags	3 Lags	4 Lags	5 Lags	6 Lags	10 Lags
Daily	-0.58	-0.43	-0.30	-0.22	-0.03	0.45
Weekly	0.83	1.28	1.55	1.41	1.36	1.55
Monthly	1.81*	1.74*	1.96*	1.93*	1.82*	1.54

Table 3.2: Lo and MacKinlay (1988) Variance Ratio Test for Return Predictability.

Notes: We find that Bitcoin's returns are predictable at lower frequencies. Critical values are as for the two-sided t-test. *: p <0:10, **: p <0:05, ***: p <0:01

As shown in Table 3.2, we find evidence that Bitcoin's returns are indeed predictable at monthly frequency, which we confirm in our regression analysis. Before turning to our regression results, we recap the predictions for the relationship between disagreement and returns for each model in Table 3.3.

[23]We provide autocorrelation functions of our variables in Section 3.C.1.

	Past Returns	Future Returns
Differences-of-Opinion	>0	<0
Uncertainty	<0	>0
Side-Show	>0	0

Table 3.3: Predictions for the Relationship between Disagreement and Returns.

The presented models of differences-of-opinion in the presence of short-sale constraints and disagreement as uncertainty yield exactly opposite predictions regarding the relationship of returns and disagreement. If optimists price the asset as in the differences-of-opinion model, an increase in disagreement, while holding average sentiment constant, leads to an increase in overpricing. Such overpricing then is predictive of lower returns in the future due to the mean-reversion of overoptimism. In contrast, an increase in disagreement can be viewed as a sign of uncertainty, which leads to risk-averse traders requiring higher returns, thus lowering the price today. Finally, in the model in which sentiment is simply a reflection of past returns, disagreement should have no predictive power when controlling for past returns.

3.4.1.1 Regressions

As a first step, we estimate the relationship between sentiment and disagreement in period t, and returns in $t-1$, t, and $t+1$ to distinguish between the different models.

Our results are not strictly in line with the proposed models. As seen in Table 3.4, the contemporaneous effect of disagreement on the return is negative, as is the effect one period ahead. For example, a one standard deviation increase in disagreement in month t decreases the return in the subsequent month by 7.7 percentage points. The negative contemporaneous relationship is what we would have expected from the model with risk-averse traders, while the negative predictive effect is in line with the differences-of-opinion model. Therefore, neither model explains the empirical results exactly. To test whether sentiment and disagreement contain information beyond what is reflected by past returns, we report in Table 3.5 the one-period-ahead predictive regression while controlling for lagged returns. If the "disagreement as a side-show" model was true, we would expect that sentiment and disagreement do not predict returns and that past returns significantly forecast future returns. However, this turns out to be incorrect: past

	Returns t-1 daily	Returns t daily	Returns t+1 daily	Returns t-1 weekly	Returns t weekly	Returns t+1 weekly	Returns t-1 monthly	Returns t monthly	Returns t+1 monthly
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Sentiment	0.73***	0.70***	0.09	3.11***	4.59***	1.04*	8.84***	13.38***	4.84**
	(0.08)	(0.09)	(0.08)	(0.48)	(0.55)	(0.61)	(2.18)	(2.02)	(2.22)
Disagreement	−0.45***	−0.42***	−0.08	−2.86***	−3.52***	−1.14**	−9.56***	−12.21***	−7.69***
	(0.07)	(0.07)	(0.07)	(0.49)	(0.47)	(0.56)	(2.16)	(1.92)	(2.41)
Constant	0.93***	0.85***	0.48	13.26***	13.74***	6.88**	52.85***	59.67***	51.59***
	(0.32)	(0.32)	(0.37)	(3.18)	(2.97)	(3.41)	(14.41)	(13.77)	(16.51)
N	2,635	2,634	2,633	378	377	376	88	87	86
R^2	0.03	0.03	0.001	0.09	0.18	0.01	0.17	0.34	0.09
Adjusted R^2	0.03	0.03	−0.0001	0.09	0.18	0.01	0.15	0.33	0.06

Notes: Returns are the growth rate between the opening and closing price in percentage points. Sentiment is the mean of the comment sentiment distribution computed by VADER, and disagreement is the standard deviation of the same distribution. Sentiment and disagreement are normalized by their respective 2014-2021 standard deviation. We generate all variables at daily, weekly, and monthly frequency. Values in parenthesis are HAC-robust standard errors following Newey and West (1987).
*: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$

Table 3.4: Regression of Leads and Lags of Returns on Sentiment and Disagreement.

Notes: Sentiment is positively related with returns, whereas disagreement is negatively related to returns. All relationships are stronger at lower frequencies.

returns have little explanatory power for future returns and the coefficients for sentiment and disagreement remain close to what they where in the regression without lags in Table 3.4. Thus, all in all, none of the three most suggestive models seem to provide a comprehensive explanation for our empirical results.

To this end, we estimate Model 3.9 for many periods ahead and present the results in Figure 3.4.1. We find that disagreement has a strongly persistent negative effect on future returns, which is more pronounced at lower frequencies, cancelling out higher-frequency noise. At monthly frequency, disagreement has a significantly negative effect on returns at 95% confidence up to five months into the future. The plots at the daily and weekly frequency show that this effect is not driven by outliers, but that returns are consistently negative. Thus, disagreement predicts lower returns for up to half a year ahead, which, through the lens of the differences-of-opinion model, suggests that prices take a long time to revert back from overoptimistic levels.[24]

An additional testable prediction whether sentiment and disagreement predict returns further into the future. Indeed, if beliefs and, therefore, disagreement are persistent, overpricing as in the differences-of-opinion model might resolve only slowly. As

[24] This long-lasting effect is also found for other markets. Disagreement in the stock market may forecast lower returns for up to a year (Diether, Malloy, and Scherbina 2002).

	Returns t+1								
	daily	daily	daily	weekly	weekly	weekly	monthly	monthly	monthly
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Sentiment	0.10	0.09	0.08	0.94	0.71	0.68	3.98	4.08	4.10
	(0.08)	(0.09)	(0.09)	(0.59)	(0.59)	(0.62)	(2.61)	(2.58)	(2.71)
Disagreement	−0.08	−0.08	−0.07	−1.07*	−0.87	−0.84	−6.91**	−7.01**	−7.16**
	(0.07)	(0.07)	(0.07)	(0.57)	(0.61)	(0.65)	(2.95)	(2.99)	(3.14)
Return t	−0.02	−0.03	−0.03	0.02	0.02	0.02	0.06	0.06	0.07
	(0.03)	(0.03)	(0.03)	(0.07)	(0.07)	(0.07)	(0.11)	(0.10)	(0.11)
Return t-1		0.003	0.002		0.05	0.05		−0.01	−0.02
		(0.03)	(0.03)		(0.07)	(0.07)		(0.10)	(0.11)
Return t-2			0.01			0.03			−0.03
			(0.02)			(0.06)			(0.04)
Constant	0.48	0.47	0.45	6.62*	5.83	5.55	47.79**	48.32**	49.83**
	(0.36)	(0.37)	(0.37)	(3.49)	(3.75)	(3.92)	(18.80)	(18.79)	(20.16)
N	2,624	2,616	2,608	375	374	373	86	86	86
R^2	0.001	0.001	0.001	0.01	0.01	0.02	0.09	0.09	0.10
Adjusted R^2	−0.0000	−0.0004	−0.001	0.004	0.003	0.002	0.06	0.04	0.04

Notes: Returns are the growth rate between the opening and closing price in percentage points. Sentiment is the mean of the comment sentiment distribution computed by VADER, and disagreement is the standard deviation of the same distribution. Sentiment and disagreement are normalized by their respective 2014-2021 standard deviation. We generate all variables at daily, weekly, and monthly frequency. Values in parenthesis are HAC-robust standard errors following Newey and West (1987).
*: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$

Table 3.5: Regression of Returns on Sentiment and Disagreement Controlling for Lagged Returns.

Notes: The coefficients on all lagged returns are insignificant and including lags reduces the adjusted R^2. The sign on the coefficient of disagreement and sentiment remains stable, but including lags marginally decreases the size and increases the standard errors, leading to a decrease in significance.

a result, disagreement should predict negative returns for multiple periods ahead.

As seen in Figure 3.1, sentiment and disagreement are relatively persistent and may feature stochastic trends. Indeed, Augmented Dickey-Fuller tests on disagreement rejects the null of the series featuring a unit root at the daily frequency, however does not reject the null at the weekly and monthly frequency, highlighting persistent low-frequency movements. As is well known, very persistent regressors can lead to t-statistics that are too large. Therefore, we provide estimates of confidence intervals that take into account the persistence of regressors.

To address these concerns, we employ the methodology in Campbell and Yogo (2006) to compute confidence intervals that are robust to the presence of persistent regressors and show the results in Figure 3.4.2, where the black line shows the central value of the confidence interval. Different to before, we use univariate local projections of returns on disagreement, as Campbell and Yogo (2006) is only applicable to univariate predictive regressions. We find that our results hold at the 90% confidence level.[25]

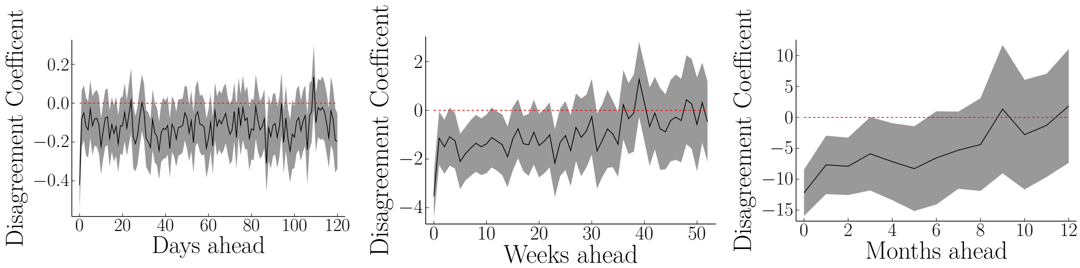

(a) Daily Returns. (b) Weekly Returns. (c) Monthly Returns.

Figure 3.4.1: Local Projections of Returns on Disagreement controlling for Sentiment with HAC-Robust Standard Errors.

Notes: The shown estimates are the coefficients on disagreement when estimating Model 3.9 for leads of returns. Error bands are at 95% confidence and standard errors are HAC-robust according to Newey and West (1987).

[25]We also run our regressions in first differences and show the results in Figures 3.B.1 and 3.B.2. Although significance suffers due to the introduction of additional noise through differencing, the basic results continue to hold. However, we focus on the regression in levels due to its straightforward interpretation.

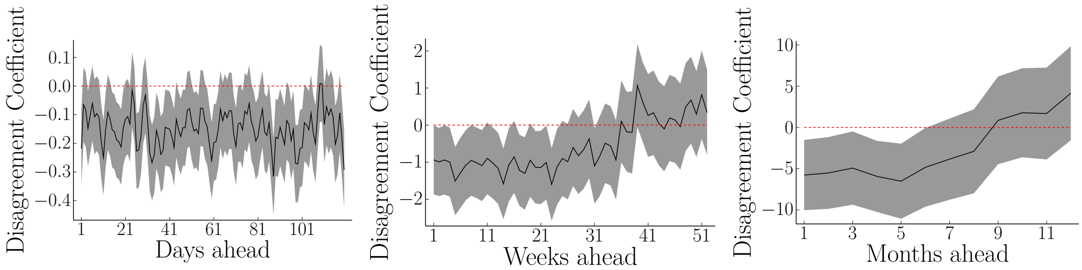

(a) Daily Returns.

(b) Weekly Returns.

(c) Monthly Returns.

Figure 3.4.2: Univariate Robust Local Projection of Returns on Disagreement.

Notes: In a univariate regression, we find that disagreement predicts negative returns for several periods at all frequencies. 90% confidence intervals according to Campbell and Yogo (2006).

Another way to express our findings on the persistence of disagreement shocks is to focus on the cumulative returns over multiple months. Similar to Figure 3.4.2, in Figure 3.4.3 we run univariate predictive regressions with disagreement in period t as the predictor for the cumulative returns between the opening price in $t+1$ and the closing price in $t+s$. We compute first the confidence intervals as in Campbell and Yogo (2006) and additionally widen them by a factor $\sqrt{s}$ as suggested by Hjalmarsson (2011). Without this adjustment, the implied confidence intervals would be too narrow for long-horizon regression with overlapping observations.

We find that the effect of disagreement on cumulative returns is close to being significant at the 90% confidence level for most horizons and significant at the 90% level for some horizons (e.g., five to nine weeks ahead). This loss in significance compared to Figure 3.4.2 is somewhat puzzling, but we attribute it to the conservative computation of the confidence intervals. Note also that these are univariate regressions. Given that disagreement and sentiment are positively correlated, and that sentiment is positively related to returns, we would expect that the effect of disagreement on returns is *biased towards zero* when not controlling for sentiment.

For the estimates that are significant at 90% confidence, we find that a one standard deviation shock on disagreement leads to a eight-week return that is about 9.2 percentage points lower, which corresponds to about 13% of the standard deviation of eight-week returns for Bitcoin. Furthermore, though insignificant, we find that the effect of disagreement on cumulative returns only reverts after more than twelve months.

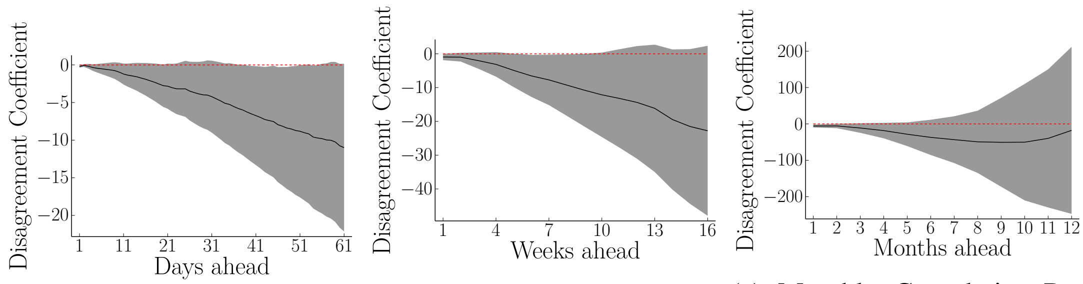

(a) Daily Cumulative Returns. (b) Weekly Cumulative Returns. (c) Monthly Cumulative Returns.

Figure 3.4.3: Long-Horizon Regression of Cumulative Returns on Disagreement.

Notes: We find that a positive one standard deviation shock to disagreement has long-lasting negative effects on returns. Confidence intervals are at 90% according to Hjalmarsson (2011).

Finally, since we hypothesize that the predictive power of disagreement is due to disagreement (and thus overoptimistic beliefs) being persistent, we test whether disagreement remains predictive when controlling for contemporaneous disagreement. In Table 3.B.3, which is shown in the appendix, we run our regression on contemporaneous returns while including contemporaneous sentiment and disagreement, as well as three lags at each frequency. We find that lagged disagreement is insignificant at daily and monthly frequency yet significantly positive at weekly frequency. At the same time, contemporaneous the coefficient of disagreement remains significantly negative at all frequencies.

Again, this finding does not fit well in any of the suggested theories. Whereas the differences-of-opinion model predicts that lagged disagreement inflates yesterday's price with a negative effect on today's return, viewing disagreement as a sign of increasing uncertainty should depress yesterday's price with a positive effect on today's return. We do not find strong evidence for either story.

	Turnover Growth t			Volatiliy t		
	daily	weekly	monthly	daily	weekly	monthly
	(1)	(2)	(3)	(4)	(5)	(6)
Sentiment	-6.59^{***}	-7.20^{***}	-12.23^{***}	-0.22^{***}	-0.36^{***}	-0.42^{***}
	(1.09)	(2.48)	(4.58)	(0.03)	(0.07)	(0.12)
Disagreement	5.11^{***}	4.91^{***}	9.19^{**}	0.10^{***}	0.21^{***}	0.23^{**}
	(0.92)	(1.77)	(4.23)	(0.03)	(0.07)	(0.11)
Constant	-4.68	-9.06	-23.09	0.80^{***}	0.47	0.50
	(3.53)	(8.78)	(27.85)	(0.11)	(0.34)	(0.55)
N	2,644	378	86	2,616	378	87
R^2	0.02	0.04	0.08	0.08	0.25	0.36
Adjusted R^2	0.02	0.03	0.06	0.08	0.24	0.34

Notes: Turnover is total dollar volume across all major exchanges divided by the market capitalization of Bitcoin. Turnover Growth is computed as the growth rate between past period's turnover and current turnover in percentage points. Volatility is the standard deviation of hourly returns over a day, week, or month. Sentiment is the mean of the comment sentiment distribution computed by VADER, and disagreement is the standard deviation of the same distribution. Sentiment and disagreement are normalized by their respective 2014-2021 standard deviation. We generate all variables at daily, weekly, and monthly frequency. Values in parenthesis are HAC-robust standard errors following Newey and West (1987).
*: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$

Table 3.6: Contemporaneous Regressions of Turnover Growth and Volatility on Sentiment and Disagreement.

3.4.2 Turnover and Volatility Regressions

We present our results for the contemporaneous effect of sentiment and disagreement on turnover growth and price volatility in Table 3.6. We find that sentiment is significantly associated with contemporaneous turnover growth and volatility of Bitcoin at all frequencies. Moreover, our results grow in magnitude and explanatory power when looking at lower frequencies, i.e., longer-lasting increases in sentiment or disagreement have greater effects.

Our results are in line with the theoretical predictions of the differences-of-opinion model: disagreement increases trading activity and drives up price volatility. This last finding suggests that increases in disagreement indicate more underlying volatility of beliefs.

On the other hand, we find that sentiment and disagreement do not have much predictive power in explaining turnover growth and volatility one period ahead, as shown in Table 3.7. Disagreement predicts volatility only at the daily and weekly frequency and does not predict turnover growth at all. Note that the lack of mean-reversion in turnover growth means that the effect of disagreement on turnover is relatively persistent. This finding is also confirmed at longer horizons in the local projections in Figure 3.4.4 when focusing on the effect of disagreement. We provide univariate local projects with confidence intervals according to Campbell and Yogo (2006) in Figure 3.C.3. Tables 3.B.1 and 3.B.2 in the appendix show the contemporaneous and one-period-ahead regression in first-differences for volatility and turnover growth.

Although the main focus of our analysis is on the effect of disagreement on returns, turnover growth, volatility, we provide for completeness the corresponding local projections focusing on the effect of sentiment in Figure 3.C.4 with HAC-robust standard errors. Figure 3.C.5 shows the results for the univariate regressions with sentiment as the predictor with confidence intervals computed according to Campbell and Yogo (2006).

	Turnover Growth t+1			Volatiliy t+1		
	daily	weekly	monthly	daily	weekly	monthly
	(1)	(2)	(3)	(4)	(5)	(6)
Sentiment	-0.09	2.02	2.88	-0.20^{***}	-0.25^{***}	-0.25^{*}
	(0.99)	(1.79)	(4.16)	(0.03)	(0.07)	(0.14)
Disagreement	-0.62	-2.41	-3.03	0.08^{***}	0.10^{*}	0.09
	(0.86)	(1.68)	(5.22)	(0.03)	(0.06)	(0.11)
Constant	10.48^{**}	16.78^{*}	23.61	0.81^{***}	0.88^{***}	1.03^{*}
	(4.20)	(10.17)	(35.02)	(0.11)	(0.32)	(0.57)
N	2,643	377	86	2,615	377	86
R^2	0.0002	0.005	0.01	0.06	0.11	0.12
Adjusted R^2	-0.001	-0.001	-0.02	0.06	0.11	0.10

Notes: Turnover is total dollar volume across all major exchanges divided by the market capitalization of Bitcoin. Turnover Growth is computed as the growth rate between past period's turnover and current turnover in percentage points. Volatility is the standard deviation of hourly returns over a day, week, or month. Sentiment is the mean of the comment sentiment distribution computed by VADER, and disagreement is the standard deviation of the same distribution. Sentiment and disagreement are normalized by their respective 2014-2021 standard deviation. We generate all variables at daily, weekly, and monthly frequency. Values in parenthesis are HAC-robust standard errors following Newey and West (1987).
*: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$

Table 3.7: Predictive Regressions of Turnover Growth and Volatility on Sentiment and Disagreement.

Notes: Disagreement predicts lower returns at weekly and monthly frequency. Turnover remains alleviated after an increase in disagreement, whereas the effect of disagreement on volatility disappears after a week.

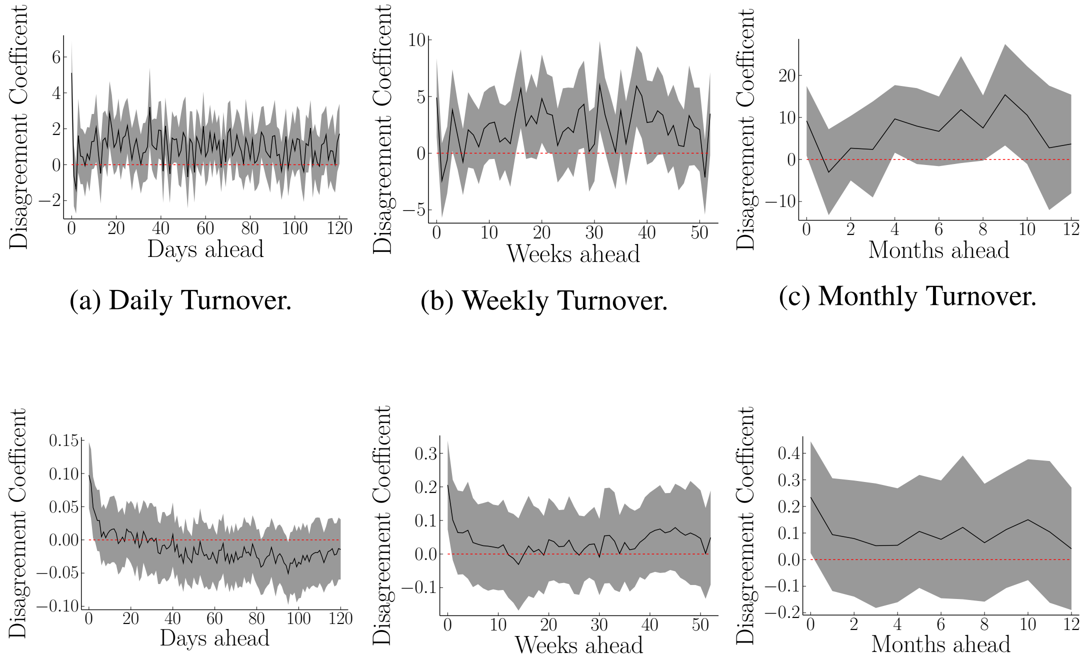

(a) Daily Turnover. (b) Weekly Turnover. (c) Monthly Turnover.

(d) Daily Volatility. (e) Weekly Volatility. (f) Monthly Volatility.

Figure 3.4.4: Local Projections of Turnover Growth and Volatility on Disagreement with HAC-robust Standard Errors.

Notes: The shown estimates are the coefficients on disagreement when estimating Model 3.9 for leads of turnover growth and volatility. Error bands are at 95% confidence and standard errors are HAC-robust according to Newey and West (1987).

3.4.3 Introduction of CME Futures

The presented framework analyzed the effect of disagreement in the presence of short-sale constraints. A major event in this context is the introduction of futures trading contract at the Chicago Mercantile Exchange (CME) on 18 December 2017 (CME Group 2017) and options on futures contract started on 12 January 2020 (CME Group 2019). The introduction of futures contracts and options does not only make markets more complete but should also substantially alleviate short-sale constraints.

The differences-of-opinion model predicts that an easing of short-sale constraints through the introduction of futures and options can eliminate the effect of disagreement, as pessimists can voice their opinion by selling short. To study this prediction, we estimate Model 3.9 contemporaneously and year-by-year. We focus on the contemporaneous specification, as the regression with lagged sentiment and disagreement in Table 3.B.3 suggests that the negative effect of disagreement on future returns stems contemporaneous disagreement. We show the coefficient on disagreement with 95% error bands in Figure 3.4.5. We also report the monthly specification for completeness, although twelve observations per year are arguably too little to draw solid inference.

We find that the coefficient and its error bands on disagreement change over time. In particular, the negative effect of disagreement on returns is particularly large in 2017 and 2018 at the daily and weekly frequency, whereas no effect can be measured in 2015. Potentially, this result can be related to insufficient variance in disagreement and returns in 2015, such that some episodes can be characterized as more or less speculative.

Starting from 2017, the estimate of the coefficient of disagreement for returns tends toward zero. Moreover, the estimate for 2020 is insignificantly different from zero at all frequencies, and the difference between the coefficients in 2016 and 2020 is statistically significant as shown in Table 3.B.4. This finding can be interpreted as short-sale constraints having sufficiently eased since 2018 such that disagreement does not lead to overpricing anymore.[26]

We also study the role of sentiment more generally over time by showing the R^2 of estimating Model 3.9 year-by-year in Figure 3.4.6. Generally, we find that sentiment and disagreement play a larger role at lower frequencies as demonstrated by higher R^2 measures. We also see here that the importance of sentiment changes over time. Although the coefficient on disagreement tends towards zero at the end of the sample, the explanatory power of sentiment and disagreement combined remains high. This is not surprising, as Bitcoin remains a speculative asset also when short-sales are permitted.

3.5 Discussion

We can summarize our main results as follows: disagreement does predict lower subsequent returns for up to five months into the future and the explained variation is not small (but also not suspiciously large) with an adjusted R^2 of around 6% for the one-month-

[26] We conduct a similar analysis for sentiment in Figure 3.C.6. Similarly, we find that the effect of sentiment changes over time, but does not go to zero towards the end of our sample.

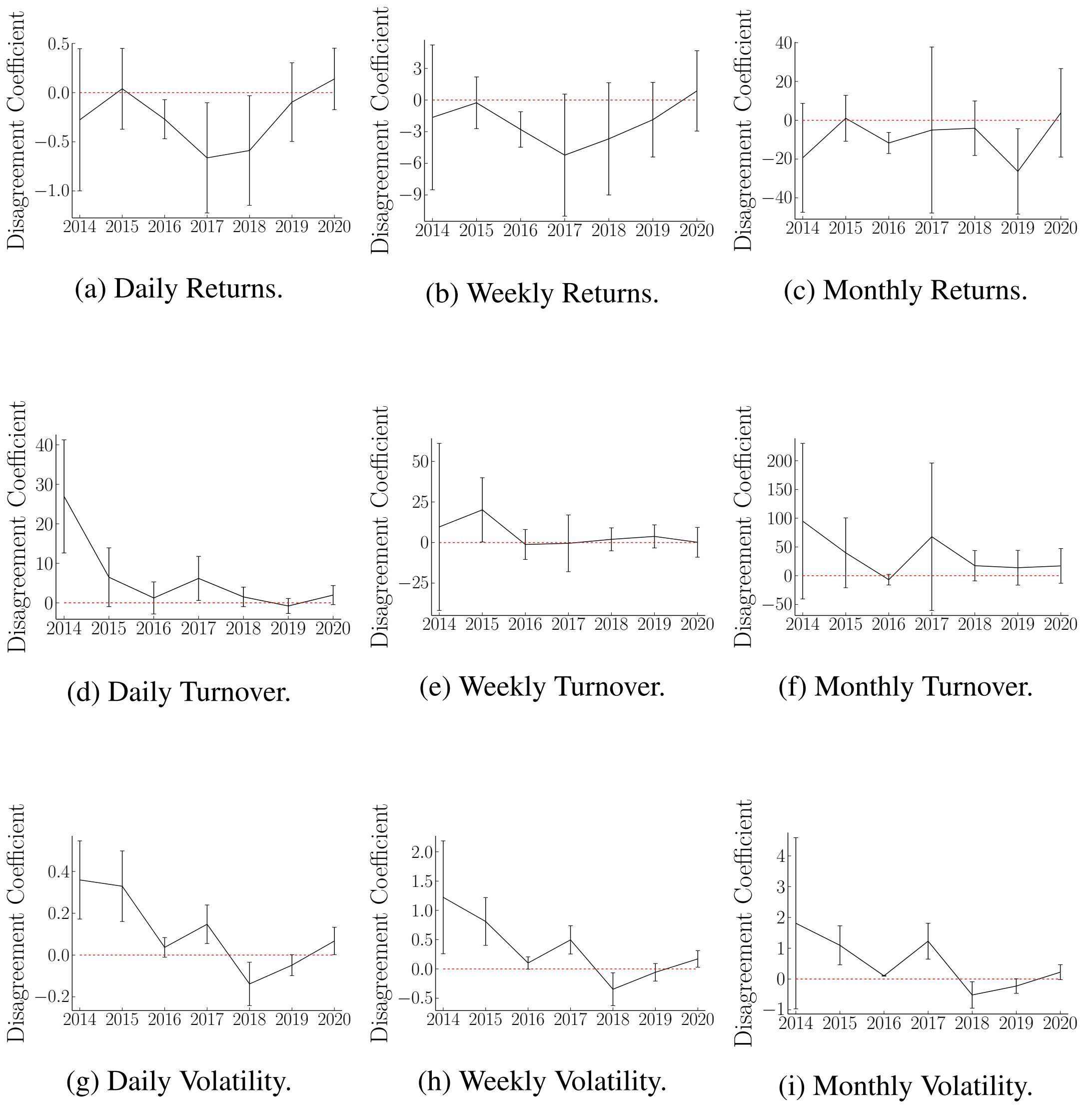

(a) Daily Returns. (b) Weekly Returns. (c) Monthly Returns.

(d) Daily Turnover. (e) Weekly Turnover. (f) Monthly Turnover.

(g) Daily Volatility. (h) Weekly Volatility. (i) Monthly Volatility.

Figure 3.4.5: Year-by-Year Coefficient of Disagreement for the Contemporaneous Regression.

Notes: We find that the contemporaneous effect of disagreement is relatively stable over time. Towards the end of the sample, the negative correlation between disagreement and returns vanishes. Error bands are at 95% confidence and standard errors are HAC-robust according to Newey and West (1987).

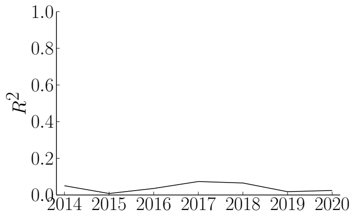

(a) Daily Returns.

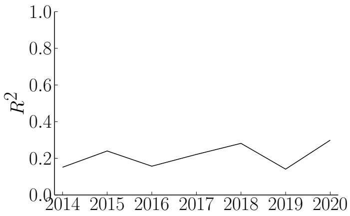

(b) Weekly Returns.

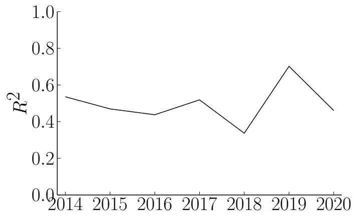

(c) Monthly Returns.

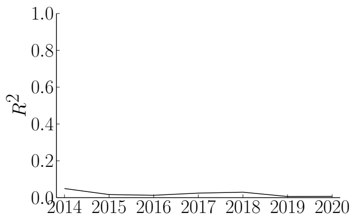

(d) Daily Turnover.

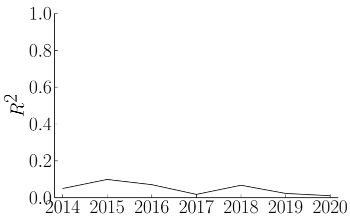

(e) Weekly Turnover.

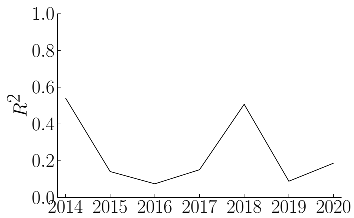

(f) Monthly Turnover.

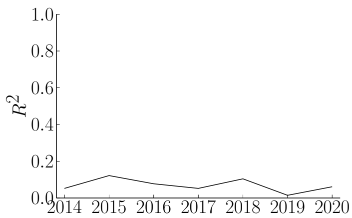

(g) Daily Volatility.

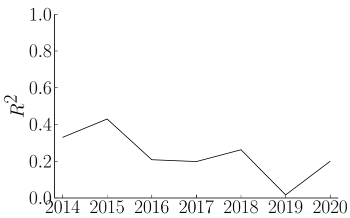

(h) Weekly Volatility.

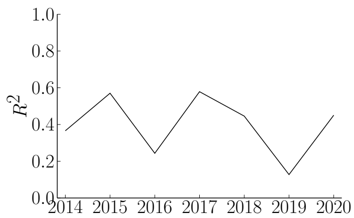

(i) Monthly Volatility.

Figure 3.4.6: Year-by-Year R^2 for the Contemporaneous Regression.

ahead regression,[27] high disagreement does also come with a negative contemporaneous effect on returns. Robustness checks, such as correcting for regressor persistence and controlling for lagged regressors and returns, respectively, do not alter this result. Furthermore, high disagreement comes with contemporaneously higher price volatility and turnover growth, however, for these variables we do not see a strong predictive effect.

3.5.1 Disagreement and Overpricing

Our main result that returns can be predicted by disagreement is in line with the differences-of-opinion argument that we characterize in 3.2.1: we find that high dispersion in our sentiment measure (i.e., disagreement is high) forecasts long-lasting negative returns at the daily, weekly, and monthly frequency. Through the lens of the differences-of-opinion literature, we would interpret this result as buyers' overoptimism decaying slowly, which could only occur if pessimists' ability or willingness to short-sale is limited. Moreover, we also find positive effects of high disagreement on turnover and volatility, which can be interpreted as market participants trading more often when their opinions are more dispersed.

However, our findings also differ from the standard differences-of-opinion story as portrayed in 3.2.1. Following a standard interpretation of the channel, the fundamental value of the asset is orthogonal to investors' beliefs. Therefore, an increase in overoptimism as reflected by high disagreement leads to overpricing. Thus, from the viewpoint of the differences-of-opinion literature, we would have hypothesized that the contemporaneous effect of disagreement on returns is positive. However, we find that it is significantly negative. This result would be expected if disagreement was just a symptom of underlying uncertainty, as we show in 3.2.2, or if disagreement was caused by negative past and contemporaneous returns, as we discuss in 3.2.3. However, with these two explanations we should not see that disagreement predicts negative returns. If the subsequent negative returns are interpreted as a correction of overpricing caused by mean-preserving disagreement, why do we not see a price increase of equal magnitude leading to this overpricing?[28]

[27] The analysis at higher frequencies picks up more noise, which leads to our R^2 being generally largest at monthly frequency.

[28] We only find a positive relationship between past disagreement and contemporaneous returns at weekly returns in Table 3.B.3. Still, the effect is much smaller than the subsequent predicted negative

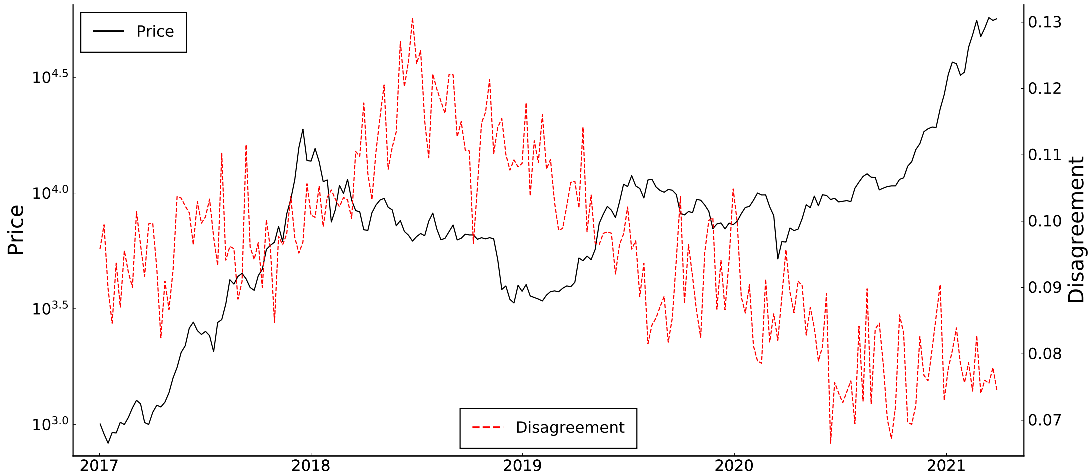

Figure 3.5.1: Disagreement Peaked after Returns Turned already Negative in 2018.

Our empirical results suggest that there has to be some mechanism that goes beyond these straightforward stories that the previous literature has discussed.[29] In the case of an asset entirely supported by beliefs, such as Bitcoin, disagreement could be associated with an erosion of the coordination that makes the asset valuable in the first place. This view reconciles our finding that disagreement is a sign of overpricing with the fact that we do not observe a price increase in the first place. A rise in disagreement lowers the "fair" or "objective" value of the asset while also keeping the price from falling immediately, leading to overpricing.

At the same time, we find that disagreement increases exactly when the price is already falling, for example, as shown in Figure 3.5.1 during the 2018 bust. Therefore, negative returns could themselves increase disagreement. A possible explanation is that traders filter negative news, which coincide with negative returns, heterogeneously. Consider as an extreme example that traders can be split into two groups: dogmatic believers and skeptics. Traders who believe dogmatically in Bitcoin might not correct their beliefs in the face of negative news. In contrast, other, less convinced traders quickly correct their beliefs downward and sell when the price starts falling. The loss of potential buyers and users of Bitcoin leads to a fall in Bitcoin's medium-term value.

This narrative is in sharp contrast to a view of an asset's fundamental value being

returns as in Figure 3.4.1 and 3.4.2.

[29]See Diether, Malloy, and Scherbina (2002) for an overview of the proposed channels.

unaffected by beliefs or disagreement. In general, asset prices may influence a firm's fundamentals in the presence of financial frictions, such that overvaluation due to the optimism of buyers can fix another inefficiency. In our case, the force behind said overvaluation - disagreement - is possibly detrimental to the asset's fundamentals, which leads to further negative returns.

3.5.2 Easing of Short-Sale Constraints

According to the literature on differences-of-opinion, disagreement leads to overpricing in the presence of short-sale constraints. Being able to short allows pessimists to trade on their belief, which reduces asset prices and offsets the influence of optimists. For Bitcoin, short-sales were difficult for two reasons: (i) the lack of financial instruments, especially through established exchanges accessible to institutional investors, and (ii) the extremely high volatility and explosive price behavior. Since the end of 2017, we have seen the gradual introduction of shorting instruments for Bitcoin, which addresses the first point. It is now possible to borrow Bitcoin on large exchanges,[30], the Chicago Mercantile Exchange introduced future contracts for Bitcoin in December 2017 (CME Group 2017) and options on futures contracts in January 2020 (CME Group 2019) which enabled especially institutional investors to bet on a falling price of Bitcoin.

The introduction of CME's futures contracts coincided with a steep fall in the price of Bitcoin, which supports the narrative that over-optimistic buyers inflated Bitcoin's price, and the introduction of futures contracts eased short-sale constraints considerably. This easing should also diminish the effect of disagreement that we find in our analysis. Indeed, if we compare the coefficient of disagreement in the regression on contemporaneous returns in the years 2016 and 2020, we find that the effect is significantly reduced as in Table 3.B.4. However, we find that the effect of disagreement on returns is also small in 2014 or 2015, well before short-sale constraints were eased. These small coefficients can potentially be understood as a sign that our disagreement channel is not strong at all times, as our disagreement measure appears to be noisy without a clear trend in 2014 and 2015.

In Figure 3.5.2 we repeat the local projection from Figure 3.4.2 to see whether the

[30]The annualized interest rates are between around 12% as of April 2021 on *Kraken.com* while requiring 20% collateral in the form of cash or cryptocurrencies.

effect of disagreement on future returns changes after the introduction of futures contracts. In this analysis, we test whether weekly disagreement observed in 2019 predicts returns that stretch until the end of 2020. Indeed, we find that in this time-frame in which futures contract were already well-established, disagreement predicts initially positive returns, which is more in line with the model with risk-averse traders. In reality, aspects of both models are likely to be relevant, and, therefore, the shift of return predictions from negative to positive suggests that relaxed short-sale constraints led to the differences-of-opinion channel being less important.

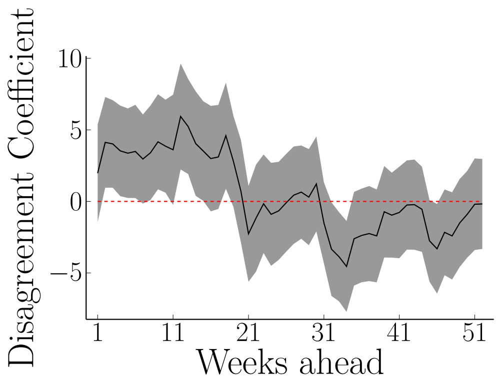

Figure 3.5.2: Univariate Robust Local Projection of Returns on Disagreement from 2019.

Notes: We find that in contrast to our earlier results, disagreement predicts initially positive returns. 90% confidence intervals according to Campbell and Yogo (2006).

Although our analysis suggests that short-sale constraints loosened, shorting Bitcoin remains costly due to limits-to-arbitrage: the maintenance volatility scan exemplifies this for the CME future contract, which as of April 2021 stands around 60%, much higher than on future contracts for other assets.[31] Relatedly, as of April 2021, the CME requires a maintenance amount for a futures contract over five Bitcoin that corresponds to 36% to 40% of the current spot price, which is much larger than the respective 5% for S&P 500 futures at CME. In other words, investors need to lock up larger amounts

[31]The maintenance volatility scan is the highest level of change that is "most likely" to occur with the underlying volatility affecting each future option's price. If the volatility of an asset is high, the margin requirement will be high as well. For comparison, S&P 500 Futures with a duration of seven months have a volatility scan of 25%.

of capital to bet on price movements of Bitcoin than for other assets, which limits the ability of investors to take on larger positions.

3.5.3 Future Research

Our analysis leaves scope for future research. First, we base our measurement of sentiment on the VADER algorithm. This algorithm is trained on online comments but has not been tested on a cryptocurrency domain. Our sentiment measure is positively correlated with returns, suggesting that the algorithm performs reasonably well for our purposes. Moreover, we use this sentiment measure as a proxy for disagreement around a time-variant level of beliefs. Ideally, we would observe beliefs and disagreement directly, which is impossible without detailed surveys. More can be done to tune the sentiment algorithm towards this specific domain.

Secondly, we are observing only a subset of potential Bitcoin investors. The online forum that we analyzed does not contain institutional investors, nor can we be certain that it represents a balanced sample of all bitcoin investors. Investors who visit an online forum and post many comments about Bitcoin might differ from those who trade quietly. Still, relevant for our purposes is only that users of *bitcointalk.org* representatively reflect the sentiments present in the general population of potential bitcoin investors. Future research could seek to analyze the beliefs and trading of institutional investors.

Finally, we did not control for who posted the comments in our regressions: were there some participants who posted many more comments than others and how did the make-up and breadth of the discussion participants change over the years? We only tested that the concentration of comments across individuals at any given time is not too high to introduce an obvious bias to our results (as we report in Figures 3.C.1a and 3.C.1b). Our empirical analysis does not exploit additional information on commenters' identities. However, they represent an interesting topic for future research.

3.6 Conclusion

We performed sentiment analysis on posts from the online forum *bitcointalk.org* to obtain a measure of investors' sentiment regarding the prospects of Bitcoin. We used the dispersion of these sentiment data points as a proxy for disagreement to explore the empirical predictions of the differences-of-opinion literature. We find that disagreement indeed predicts lower returns while being related to higher turnover and volatility, confirming the theoretical predictions of our model. The most striking result is that disagreement predicts negative returns for up to five months into the future, pointing towards a slow correction of large overvaluations. Moreover, sentiment and disagreement play a large role in explaining returns, with an adjusted R^2 of 0.33 in the regression on contemporaneous returns at a monthly frequency.

We study the change in effects of disagreement after CME introduced futures contracts in December 2017 (CME Group 2017). We find that the effect of disagreement significantly diminishes in the years after the introduction, in line with the view that a combination of disagreement and short-sale constraints is necessary to generate overpricing and subsequent predictable lower returns. An important departure from the previous empirical literature is that we find no strong evidence of a positive effect of disagreement on the contemporaneous or past return of Bitcoin. We hypothesize that disagreement may erode the coordination, which is the foundation of Bitcoin's value proposition, leading to a price decrease.

Future research could seek to understand better the effects of changes in beliefs at the investor level. Also, the network structure of online discussions can be exploited to understand better belief formation and the impact of different kinds of discussions on the price of Bitcoin. One could explore, for example, the narratives that generate disagreement between participants through topical analysis. A better understanding of the determinants of investor disagreement can help refine asset pricing theory.

Appendix

3.A Proofs

Proof of Proposition 3.1. Since the price tomorrow is independent of disagreement or the average belief today given that ϕ is constant over time and beliefs are iid across group of traders and time, it is sufficient to show that the optimists' belief ϕ^o_{t+1} and, therefore, P_t are increasing in disagreement $\phi^o_{t+1} - \phi^p_{t+1}$ when holding the average belief $\frac{\phi^o_{t+1}+\phi^p_{t+1}}{2}$ constant. The optimists' belief can be written as

$$\phi^o_{t+1} = \underbrace{\frac{\phi^o_{t+1} + \phi^p_{t+1}}{2}}_{\textit{Average Belief}} + \frac{1}{2}\underbrace{\left(\phi^o_{t+1} - \phi^p_{t+1}\right)}_{\textit{Disagreement}}. \tag{3.10}$$

Indeed, if the right-hand-side increases due to an increase in disagreement and the average belief stays constant, the left-hand-side must increase. The increase in optimists' belief ϕ^o_{t+1} leads to higher price P_t as follows from Equation 3.2 and $\bar{P} > \mathbb{E}(P_{t+1}|\neg A_{t+1})$, which lowers future returns. $\square$

Proof of Corollary 3.1. Follows from the proof of Proposition 3.1, except that disagreement increases today's price and, therefore, increases past returns $\frac{P_t - P_{t-1}}{P_{t-1}}$. $\square$

Proof of Proposition 3.2. Since the price is set by optimists with $\phi^o_{t+1} > \phi^p_{t+1}$ and $\bar{P} > \mathbb{E}(P_{t+1}|\neg A_{t+1})$, it must be that the pessimists' valuation is below the current price P_t, i.e.,

$$V^p_t = \frac{1}{R}\left(\phi^p_{t+1}\bar{P} + (1 - \phi^p_{t+1})\mathbb{E}(P_{t+1}|\neg A_{t+1})\right) < P_t, \tag{3.11}$$

Therefore, the maximization problem of pessimists is

$$\max_{d_t \leq 0} \quad d_t(P_t - V^p_t) - c(d_t) \tag{3.12}$$

with the solution $d_t^* : c'(d_t^*) = P_t - V_t^p$. It follows that pessimists short the asset more if V_t^p is lower due to more pessimistic beliefs ϕ_{t+1}^p, which decreases in disagreement as apparent in:

$$\phi_{t+1}^p = \underbrace{\frac{\phi_{t+1}^o + \phi_{t+1}^p}{2}}_{\textit{Average Belief}} - \frac{1}{2}\underbrace{\left(\phi_{t+1}^o - \phi_{t+1}^p\right)}_{\textit{Disagreement}}. \tag{3.13}$$

Therefore, an increase in disagreement while holding the average belief constant means that pessimists are even more pessimistic, which lowers V_t^p and increases the short positions d_t^* and turnover. □

Proof of Proposition 3.3. Tomorrow's price is independent of the belief of the young representative trader given that ϕ is constant over time and beliefs are iid. It remains to show that the price today is decreasing in the uncertainty of the young representative trader when the average belief is held constant.

The young representative trader views ϕ_{t+1} as a random variable X where $|\mathbb{E}(X)| < \infty$ and $Var(X) \in (0, \infty)$. Denote alternative beliefs that are more uncertain than X but have the same mean as $Z = X + Y$ where $\mathbb{E}(Y) = 0$ and $Var(Y) \in (0, \infty)$ and Y is independent of X. It is sufficient to show that a representative trader with more uncertain beliefs has a lower utility holding the asset than a more certain trader with the same mean belief. Given that the utility function 3.3 is concave in ϕ,

$$\begin{aligned}
\mathbb{E}\left(U_t(Z)\right) &= \mathbb{E}\left(U_t(X+Y)\right) \\
&\overset{\text{L.I.E.}}{=} \mathbb{E}\left(\mathbb{E}\left(U_t(X+Y)|X\right)\right) \\
&\overset{\text{Jensen's}}{<} \mathbb{E}\left(U_t\left(\mathbb{E}\left(X+Y|X\right)\right)\right) \\
&= \mathbb{E}\left(U_t\left(X\right)\right),
\end{aligned} \tag{3.14}$$

where the utility function U_t is written as a function of ϕ. Since a trader with more uncertain beliefs is worse off compared to a trader with more certain beliefs but the same mean, it follows that the price P_t must fall to restore the indifference as in Equation 3.4. As a result, future returns increase in uncertainty. □

Proof of Corollary 3.2. Follows from the proof of Proposition 3.3, except that disagreement decreases today's price and, therefore, decreases past returns $\frac{P_t - P_{t-1}}{P_{t-1}}$. □

3.B Tables

	Return Diff t			Turnover Growth t			Volatiliy Diff t		
	daily	weekly	monthly	daily	weekly	monthly	daily	weekly	monthly
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Sentiment Diff	0.59***	5.97***	15.68***	−6.06***	−10.16***	−15.26***	−0.07***	−0.21***	−0.28***
	(0.12)	(0.88)	(2.56)	(0.98)	(2.27)	(4.11)	(0.01)	(0.04)	(0.05)
Disagreement Diff	−0.20*	−2.32***	−4.31*	4.65***	7.16***	11.83***	0.02*	0.15***	0.14***
	(0.10)	(0.68)	(2.59)	(0.98)	(2.12)	(4.27)	(0.01)	(0.04)	(0.04)
Constant	0.005	0.02	0.28	7.04***	5.52***	10.03***	−0.0004	0.0000	−0.01
	(0.03)	(0.25)	(2.53)	(0.76)	(1.38)	(3.78)	(0.01)	(0.01)	(0.05)
N	2,623	376	87	2,642	378	86	2,595	378	87
R^2	0.01	0.18	0.29	0.02	0.11	0.21	0.01	0.16	0.25
Adjusted R^2	0.01	0.17	0.27	0.02	0.11	0.19	0.01	0.15	0.23

Notes: Returns are the growth rate between the opening and closing price in percentage points. Turnover is total dollar volume across all major exchanges divided by the market capitalization of Bitcoin. Turnover Growth is computed as the growth rate between past period's turnover and current turnover in percentage points. Volatility is the standard deviation of hourly returns over a day, week, or month. Sentiment is the mean of the comment sentiment distribution computed by VADER, and disagreement is the standard deviation of the same distribution. Changes in sentiment and disagreement are normalized by their respective 2014-2021 standard deviation. We generate all variables at daily, weekly, and monthly frequency. We compute the difference in a variable as today's value minus yesterday's value. Values in parenthesis are HAC-robust standard errors following Newey and West (1987).
*: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$

Table 3.B.1: Contemporaneous Regressions in First Differences.

Notes: An alternative to providing robust standard errors is differencing variables until both independent and dependent variables are stationary. We find that the results remain generally unchanged, as changes in disagreement are negatively related to returns and positively related to turnover growth and changes to volatility.

	Return Diff t+1			Turnover Growth t+1			Volatiliy Diff t+1		
	daily	weekly	monthly	daily	weekly	monthly	daily	weekly	monthly
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Sentiment Diff	-0.62^{***}	-4.11^{***}	-11.54^{***}	1.13	2.51^{*}	9.94^{**}	-0.01	0.06^{**}	0.17^{***}
	(0.12)	(0.88)	(2.66)	(0.97)	(1.44)	(4.08)	(0.01)	(0.03)	(0.06)
Disagreement Diff	0.20^{*}	2.47^{***}	2.38	0.97	-1.83	-4.23	0.02	-0.08^{**}	-0.13^{**}
	(0.12)	(0.74)	(2.51)	(0.96)	(1.57)	(3.47)	(0.01)	(0.03)	(0.06)
Constant	-0.003	0.10	0.56	7.01^{***}	5.33^{***}	8.87^{**}	-0.001	-0.003	-0.01
	(0.03)	(0.38)	(2.69)	(0.75)	(1.58)	(4.02)	(0.01)	(0.02)	(0.05)
N	2,622	375	86	2,641	377	86	2,594	377	86
R^2	0.01	0.10	0.16	0.001	0.01	0.07	0.001	0.02	0.12
Adjusted R^2	0.01	0.09	0.14	0.0002	0.002	0.05	-0.0000	0.02	0.10

Notes: Returns are the growth rate between the opening and closing price in percentage points. Turnover is total dollar volume across all major exchanges divided by the market capitalization of Bitcoin. Turnover Growth is computed as the growth rate between past period's turnover and current turnover in percentage points. Volatility is the standard deviation of hourly returns over a day, week, or month. Sentiment is the mean of the comment sentiment distribution computed by VADER, and disagreement is the standard deviation of the same distribution. Changes in sentiment and disagreement are normalized by their respective 2014-2021 standard deviation. We generate all variables at daily, weekly, and monthly frequency. We compute the difference in a variable as today's value minus yesterday's value. Values in parenthesis are HAC-robust standard errors following Newey and West (1987).
*: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$

Table 3.B.2: One-Period-Ahead Regressions in First Differences.

Notes: We find that the effects of sentiment and disagreement revert in comparison to Table 3.B.1. Note that the effect of disagreement of the change in returns one-period-ahead is insignificant and smaller in magnitude than the contemporaneous effect, pointing to a protracted negative effect of disagreement on returns.

	Returns t			Turnover Growth t			Volatiliy t		
	daily	weekly	monthly	daily	weekly	monthly	daily	weekly	monthly
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Sentiment t	1.08***	8.32***	20.29***	−8.74***	−15.87***	−22.62***	−0.16***	−0.38***	−0.49***
	(0.10)	(0.79)	(3.18)	(1.33)	(3.83)	(6.92)	(0.02)	(0.06)	(0.10)
Sentiment t-1	−0.30***	−4.19***	−8.00**	4.49***	16.24***	21.84***	−0.09***	0.02	0.06
	(0.10)	(0.95)	(3.41)	(1.39)	(3.49)	(8.02)	(0.02)	(0.03)	(0.09)
Sentiment t-2	−0.13	1.00	0.15	1.27	−0.79	−10.48	−0.02	0.004	0.08
	(0.10)	(0.97)	(3.68)	(1.36)	(3.36)	(7.18)	(0.02)	(0.05)	(0.07)
Sentiment t-3	−0.24**	−2.32***	−2.10	−1.46	−3.88	5.44	−0.02	0.01	−0.04
	(0.11)	(0.74)	(3.27)	(1.23)	(3.18)	(8.66)	(0.02)	(0.05)	(0.08)
Disagreement t	−0.39***	−4.89***	−11.60***	6.75***	12.34***	23.06**	0.08***	0.29***	0.45***
	(0.08)	(0.87)	(3.93)	(1.19)	(3.84)	(9.91)	(0.02)	(0.07)	(0.12)
Disagreement t-1	0.09	2.30**	−2.26	−2.65**	−11.44***	−20.20**	0.05***	−0.02	−0.11
	(0.09)	(1.02)	(4.67)	(1.20)	(4.31)	(8.29)	(0.02)	(0.05)	(0.11)
Disagreement t-2	0.05	−0.49	0.36	−3.34***	−7.93**	−7.09	0.01	−0.07	−0.07
	(0.09)	(1.00)	(3.94)	(1.22)	(3.14)	(10.61)	(0.01)	(0.06)	(0.10)
Disagreement t-3	−0.09	0.41	3.22	1.90	10.07***	8.43	0.004	−0.02	−0.12
	(0.09)	(0.77)	(3.56)	(1.36)	(3.52)	(8.70)	(0.01)	(0.04)	(0.08)
Constant	1.13***	12.65***	54.00***	3.23	−4.02	−4.46	0.71***	0.65*	1.02*
	(0.42)	(3.42)	(15.92)	(4.07)	(8.52)	(32.69)	(0.16)	(0.39)	(0.61)
N	2,628	377	87	2,638	378	86	2,610	378	87
R^2	0.04	0.27	0.41	0.03	0.14	0.25	0.09	0.26	0.40
Adjusted R^2	0.04	0.26	0.35	0.03	0.12	0.18	0.09	0.24	0.34

Notes: Returns are the growth rate between the opening and closing price in percentage points. Turnover is total dollar volume across all major exchanges divided by the market capitalization of Bitcoin. Turnover Growth is computed as the growth rate between past period's turnover and current turnover in percentage points. Volatility is the standard deviation of hourly returns over a day, week, or month. Sentiment is the mean of the comment sentiment distribution computed by VADER, and disagreement is the standard deviation of the same distribution. Sentiment and disagreement are normalized by their respective 2014-2021 standard deviation. We generate all variables at daily, weekly, and monthly frequency. Values in parenthesis are HAC-robust standard errors following Newey and West (1987).
*: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$

Table 3.B.3: Contemporaneous Regressions Controlling for Lags of Sentiment and Disagreement.

Notes: The results are broadly consistent with mean-reversion. Whereas sentiment in period t is positively related to returns in period t, sentiment in period $t-1$ has a *negative* effect on returns in t, albeit the coefficient is at most half as large as the contemporaneous effect. Disagreement exhibits mean-reversion for returns only at weekly frequency.

	Returns t		
	daily	weekly	monthly
	(1)	(2)	(3)
Sentiment	0.93***	6.78***	20.75***
	(0.11)	(0.90)	(3.29)
Disagreement	−0.42***	−4.90***	−17.46***
	(0.10)	(0.90)	(3.05)
Year 2017	1.60	2.87	−91.41
	(1.45)	(21.14)	(146.18)
Year 2018	0.59	−13.11	−98.69
	(1.80)	(20.73)	(73.90)
Year 2019	−0.91	−19.83	−119.45**
	(1.22)	(12.87)	(59.61)
Year 2020	−1.47*	−31.84**	−132.51**
	(0.83)	(12.98)	(58.94)
Year 2021	2.20	179.96***	108.93
	(3.03)	(44.45)	(79.60)
Disagreement x Year 2017	−0.12	0.35	15.04
	(0.28)	(3.17)	(19.50)
Disagreement x Year 2018	−0.10	1.71	11.17
	(0.30)	(2.59)	(8.53)
Disagreement x Year 2019	0.29	3.20*	16.71**
	(0.23)	(1.81)	(8.17)
Disagreement x Year 2020	0.52***	5.77***	21.37**
	(0.18)	(2.19)	(9.38)
Disagreement x Year 2021	−0.31	−32.35***	−18.16
	(0.77)	(8.07)	(13.48)
Constant	−0.33	13.78**	67.32***
	(0.54)	(6.12)	(20.81)
N	1,917	274	63
R^2	0.05	0.27	0.54
Adjusted R^2	0.04	0.23	0.43

Notes: Returns are the growth rate between the opening and closing price in percentage points. Sentiment is the mean of the comment sentiment distribution computed by VADER, and disagreement is the standard deviation of the same distribution. Sentiment and disagreement are normalized by their respective 2014-2021 standard deviation. We generate all variables at daily, weekly, and monthly frequency. Values in parenthesis are HAC-robust standard errors following Newey and West (1987).
*: $p < 0.10$, **: $p < 0.05$, ***: $p < 0.01$

Table 3.B.4: Contemporaneous Regression with Year-FEs.

Notes: The effect of disagreement on contemporaneous returns is significantly different in 2020 compared to 2016. The sign on the coefficient is positive taking together the base effect and interaction term. However, disagreement has again a more negative effect in 2021.

3.C Figures

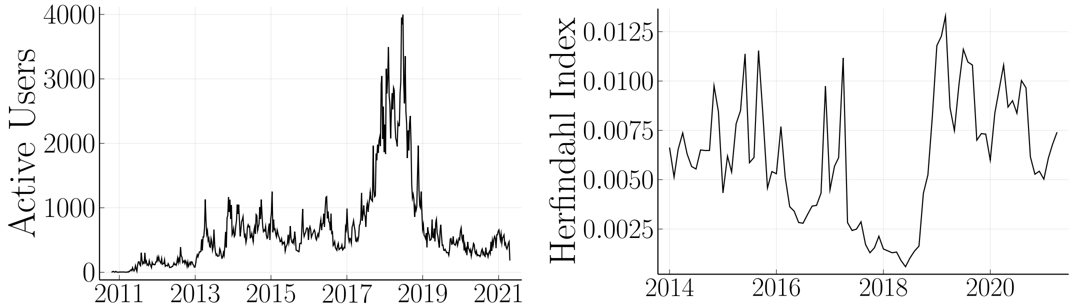

(a) Weekly Active Users.

(b) Monthly Authorship Concentration.

Figure 3.C.1: Activity and Concentration Graphs for *bitcointalk.org*.

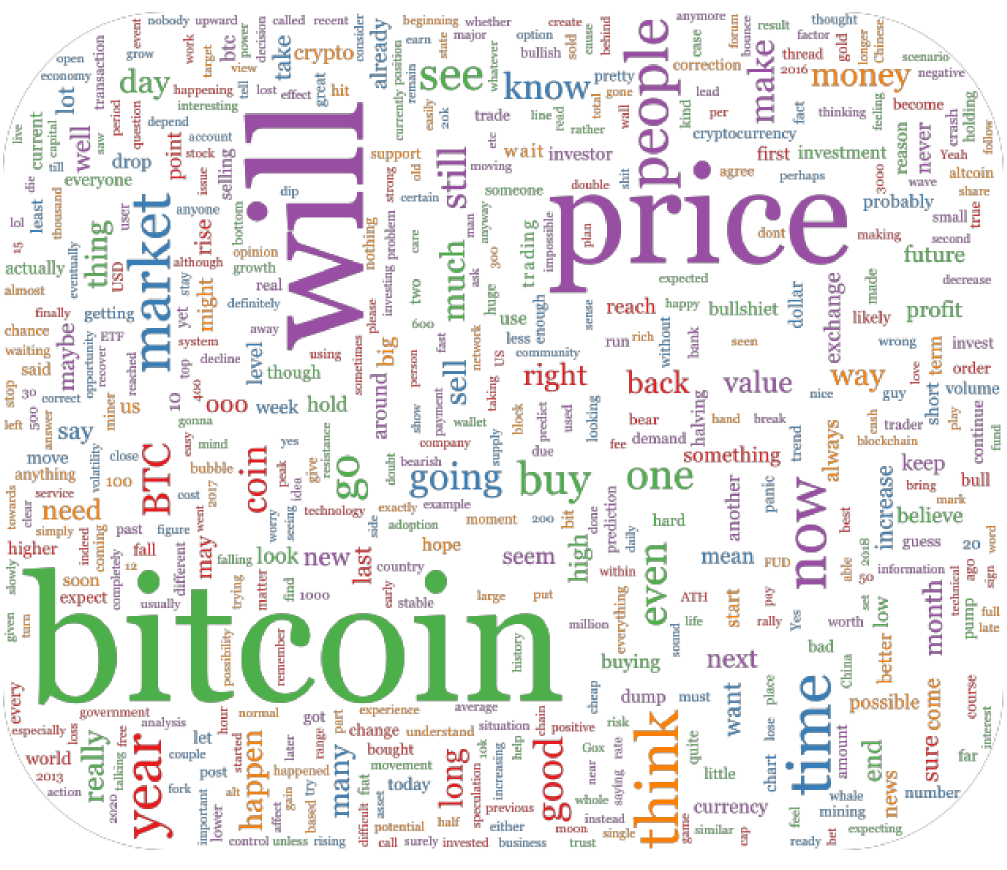

Figure 3.C.2: Most Commonly Used Words out of a Random Sample of 10,000 Comments.

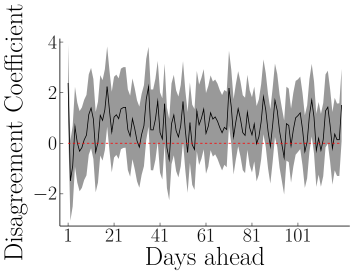

(a) Daily Turnover Growth.

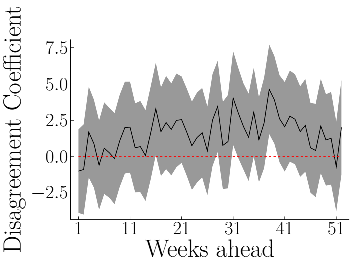

(b) Weekly Turnover Growth.

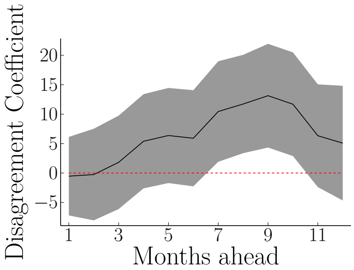

(c) Monthly Turnover Growth.

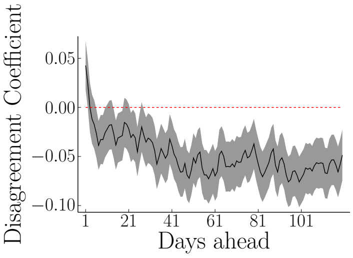

(d) Daily Volatility.

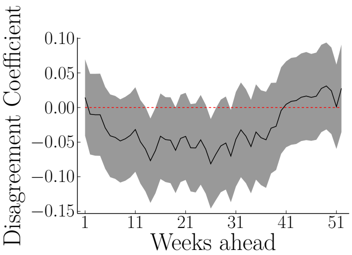

(e) Weekly Volatility.

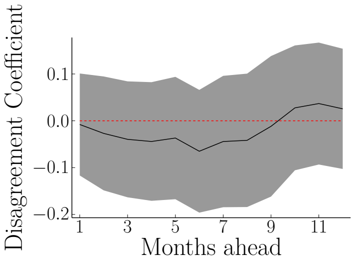

(f) Monthly Volatility.

Figure 3.C.3: Univariate Robust Local Projection of Turnover Growth and Volatility on Disagreement.

Notes: In a univariate regression, we find that disagreement does not predict further turnover growth in the following periods. Moreover, the positive effect of disagreement on volatility is short-lived. The error bands are 90% confidence intervals according to Campbell and Yogo (2006).

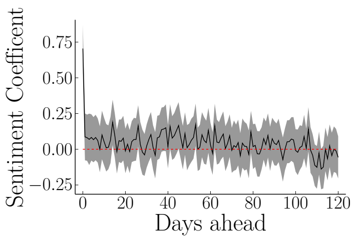

(a) Daily Returns.

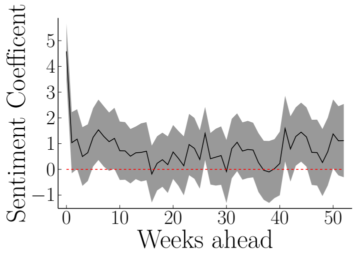

(b) Weekly Returns.

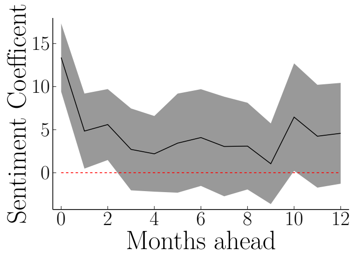

(c) Monthly Returns.

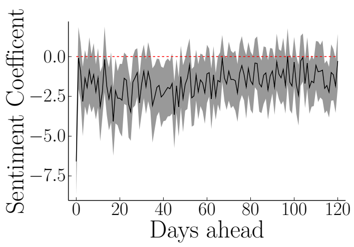

(d) Daily Turnover.

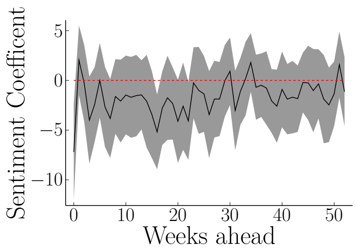

(e) Weekly Turnover.

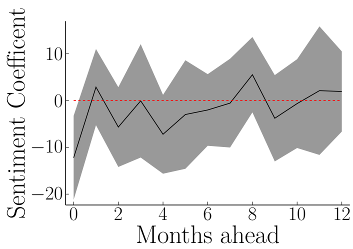

(f) Monthly Turnover.

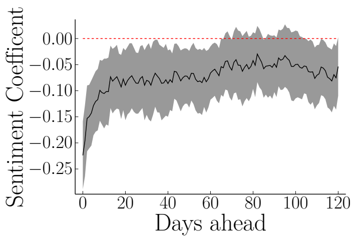

(g) Daily Volatility.

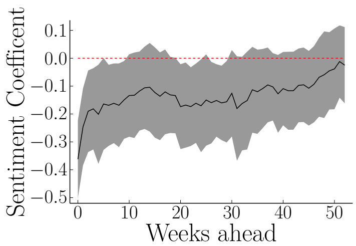

(h) Weekly Volatility.

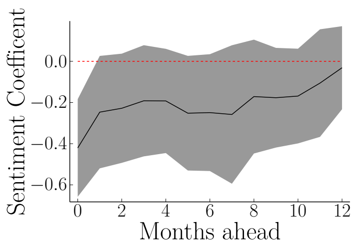

(i) Monthly Volatility.

Figure 3.C.4: Local Projections of Returns on Sentiment Controlling for Disagreement with HAC-Robust Standard Errors.

Notes: The shown estimates are the coefficients on sentiment when estimating Model 3.9 for leads of returns, turnover growth, and volatility. Error bands are at 95% confidence and standard errors are HAC-robust according to Newey and West (1987).

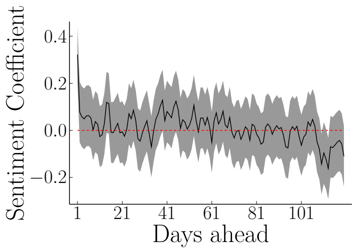

(a) Daily Returns.

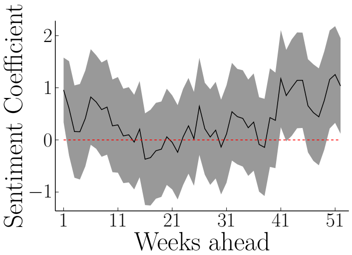

(b) Weekly Returns.

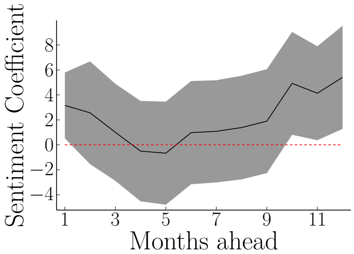

(c) Monthly Returns.

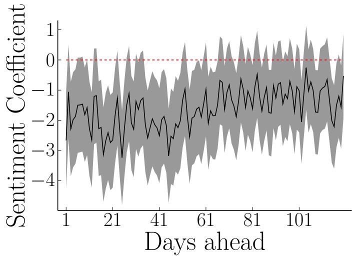

(d) Daily Turnover.

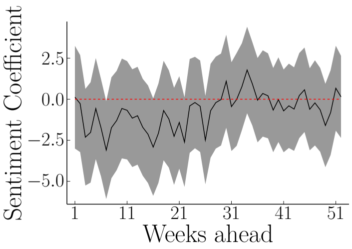

(e) Weekly Turnover.

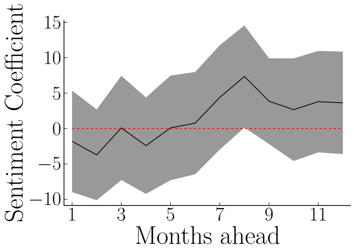

(f) Monthly Turnover.

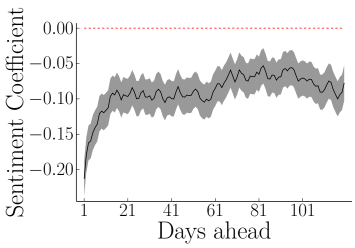

(g) Daily Volatility.

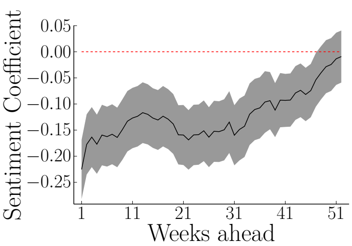

(h) Weekly Volatility.

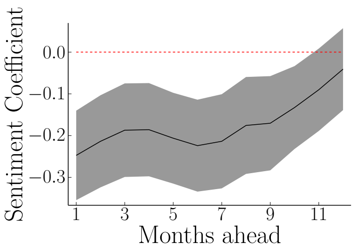

(i) Monthly Volatility.

Figure 3.C.5: Univariate Robust Local Projection of Returns, Turnover Growth, and Volatility on Sentiment.

Notes: In a univariate regression, we find that disagreement predicts negative returns for several periods at all frequencies. The error bands are 90% confidence intervals according to Campbell and Yogo (2006).

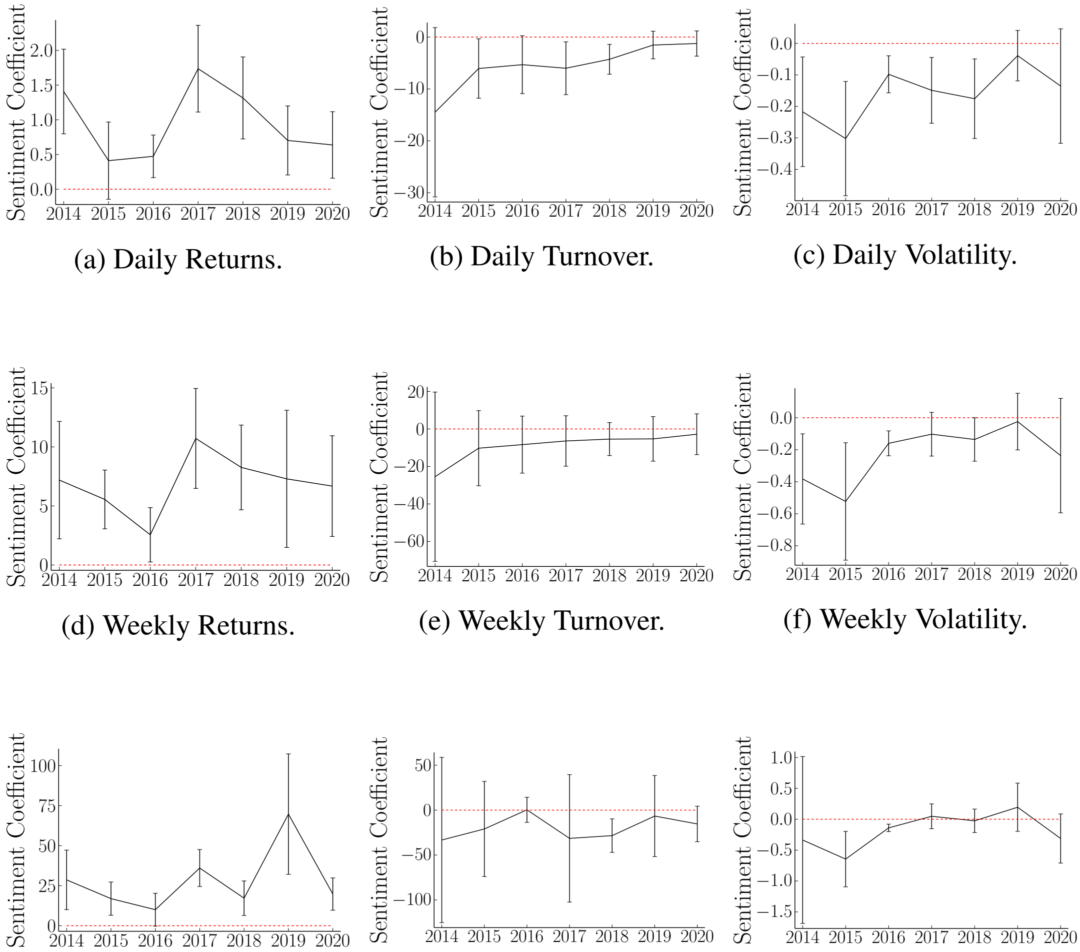

(a) Daily Returns. (b) Daily Turnover. (c) Daily Volatility.

(d) Weekly Returns. (e) Weekly Turnover. (f) Weekly Volatility.

(g) Monthly Returns. (h) Monthly Turnover. (i) Monthly Volatility.

Figure 3.C.6: Year-by-Year Coefficient of Sentiment for the Contemporaneous Regression.

Notes: We find that the contemporaneous effect of sentiment is relatively stable over time. Error bands are at 95% confidence and standard errors are HAC-robust according to Newey and West (1987).

3.C.1 Autocorrelation Plots

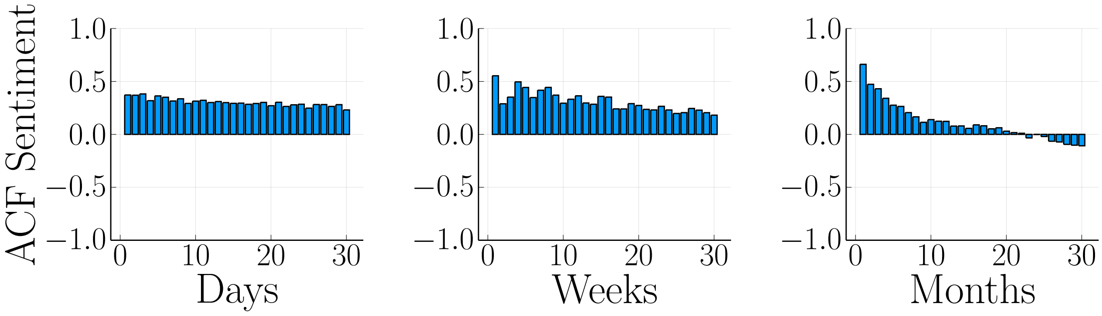

Figure 3.C.7: Autocorrelation Functions Sentiment.

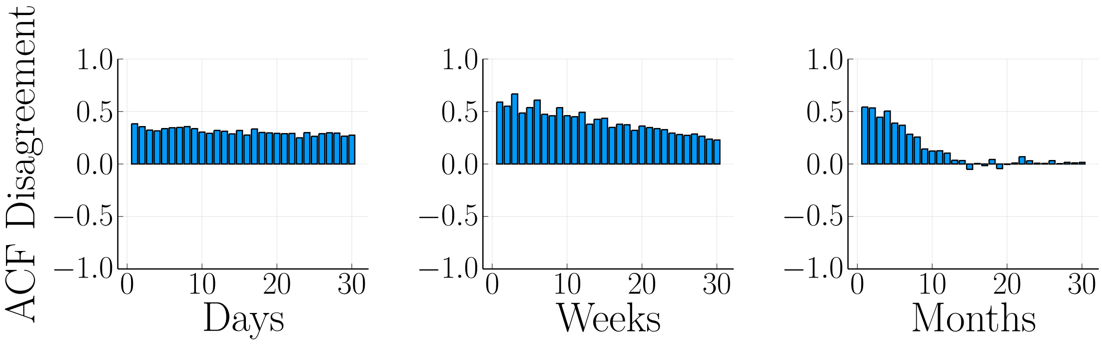

Figure 3.C.8: Autocorrelation Functions Disagreement.

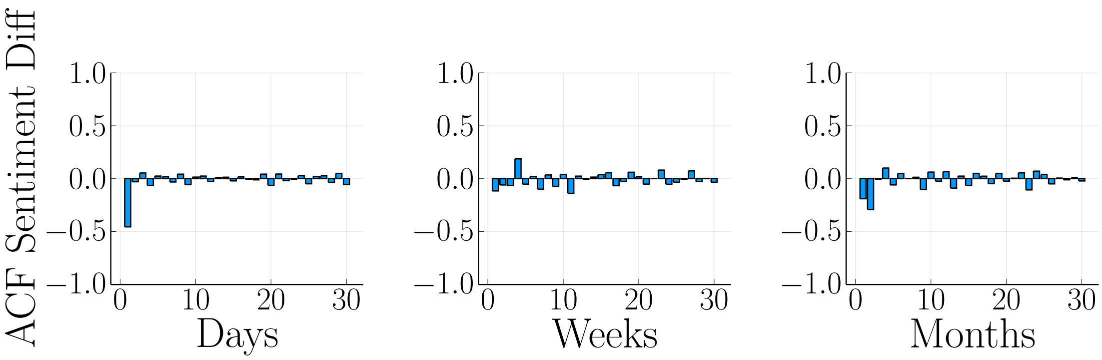

Figure 3.C.9: Autocorrelation Functions Changes in Sentiment.

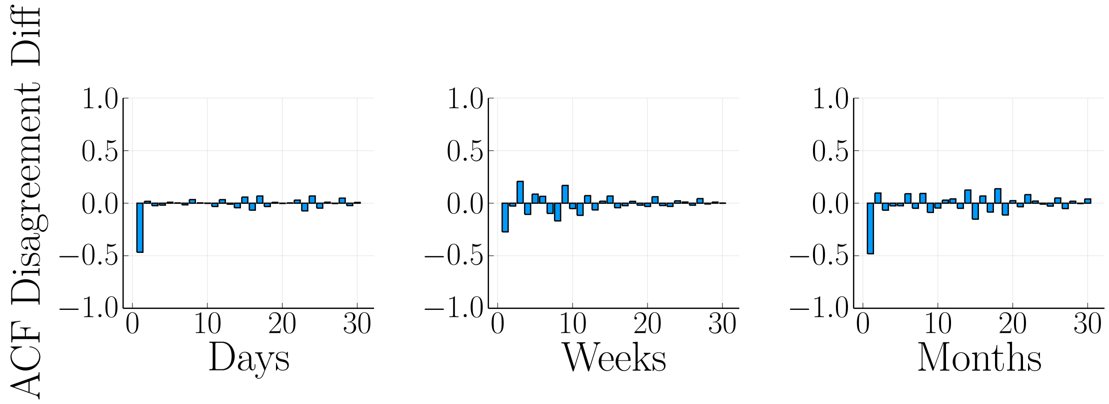

Figure 3.C.10: Autocorrelation Functions Changes in Disagreement.

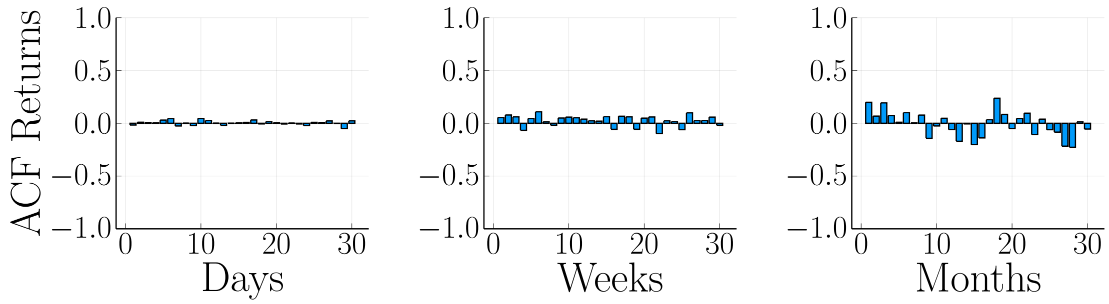

Figure 3.C.11: Autocorrelation Functions Returns.

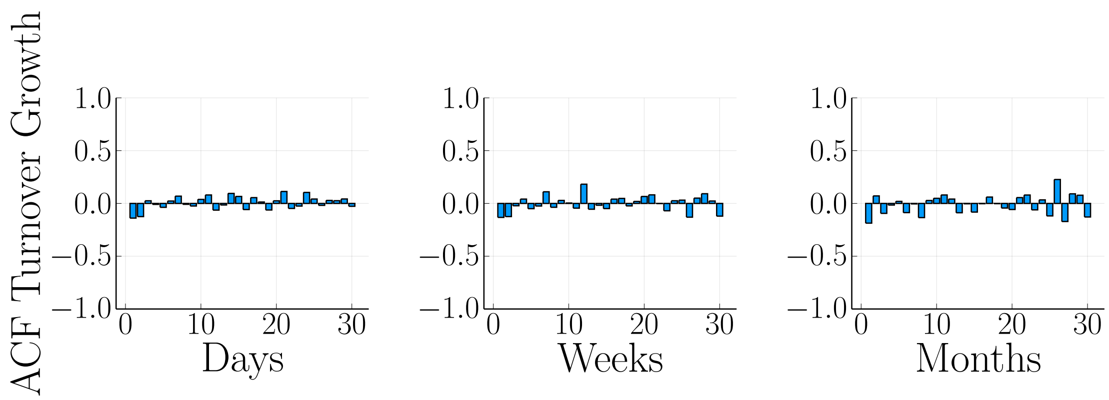

Figure 3.C.12: Autocorrelation Functions Turnover Growth.

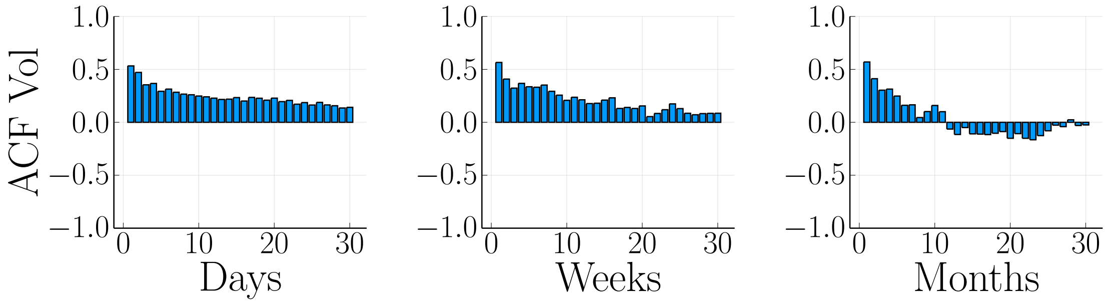

Figure 3.C.13: Autocorrelation Functions Volatility.

Bibliography

Acharya, Viral and Hassan Naqvi (2019). “On reaching for yield and the coexistence of bubbles and negative bubbles”. *Journal of Financial Intermediation* 38, pp. 1–10.

Adam, K., D. Mateev, and S. Nagel (2021). “Do survey expectations of stock returns reflect risk adjustments?” *Journal of Monetary Economics* 117, pp. 723–40.

Adrian, Tobias, Arturo Estrella, and Hyun Song Shin (2010). “Monetary cycles, financial cycles and the business cycle”. *FRB of New York Staff Report* 421.

Adrian, Tobias and Hyun Song Shin (2010). “Financial intermediaries and monetary economics”. *Handbook of monetary economics*. Vol. 3. Elsevier, pp. 601–650.

Aliber, Robert Z. and Charles P. Kindelberger (2015). *Manias, Panics, and Crashes*. London: Palgrave Macmillan UK.

Almeida, Heitor, Murillo Campello, Bruno Laranjeira, and Scott Weisbenner (2009). *Corporate debt maturity and the real effects of the 2007 credit crisis*. Tech. rep. National Bureau of Economic Research.

Antweiler, Werner and Murray Z. Frank (2004). “Is All That Talk Just Noise? The Information Content of Internet Stock Message Boards”. *The Journal of Finance* 59.3, pp. 1259–1294.

Badoer, Dominique C and Christopher M James (2016). “The determinants of long-term corporate debt issuances”. *The Journal of Finance* 71.1, pp. 457–492.

Baker, Malcolm, Robin Greenwood, and Jeffrey Wurgler (2003). “The maturity of debt issues and predictable variation in bond returns”. *Journal of Financial Economics* 70.2, pp. 261–291.

Barberis, N., R. Greenwood, J. Lawrence, and A. Shleifer (2018). “Extrapolation and Bubbles”. *Journal of Financial Economics* 129, pp. 203–27.

Barclay, Michael J and Clifford W Smith Jr (1995). “The maturity structure of corporate debt”. *The Journal of Finance* 50.2, pp. 609–631.

Becker, Bo and Victoria Ivashina (2015). “Reaching for yield in the bond market”. *The Journal of Finance* 70.5, pp. 1863–1902.

Berger, Allen N, Marco A Espinosa-Vega, W Scott Frame, and Nathan H Miller (2005). "Debt maturity, risk, and asymmetric information". *The Journal of Finance* 60.6, pp. 2895–2923.

Bernanke, Ben S and Mark Gertler (1995). "Inside the black box: the credit channel of monetary policy transmission". *Journal of Economic perspectives* 9.4, pp. 27–48.

Bernanke, Ben S, Mark Gertler, and Simon Gilchrist (1999). "The financial accelerator in a quantitative business cycle framework". *Handbook of Macroeconomics* 1, pp. 1341–1393.

Bordalo, P., N. Gennaioli, and S. Y. Kwon (2020). "Diagnostic Bubbles". *Journal of Financial Economics* 141.3, pp. 1060–1077.

Borio, Claudio and Haibin Zhu (2012). "Capital regulation, risk-taking and monetary policy: a missing link in the transmission mechanism?" *Journal of Financial Stability* 8.4, pp. 236–251.

Bubeck, Johannes, Angela Maddaloni, and José-Luis Peydró (forthcoming). "Negative monetary policy rates and systemic banks' risk-taking: Evidence from the Euro area securities register". *Journal of Money, Credit and Banking*.

Campbell, J. Y. and J. H. Cochrane (1999). "By force of habit: , A consumption-based explanation of aggregate stock market behavior." *Journal of Political Economy* 107, pp. 205–51.

Campbell, John Y and Roman Sigalov (2020). *Portfolio choice with sustainable spending: A model of reaching for yield*. Tech. rep. National Bureau of Economic Research.

Campbell, John Y. and Motohiro Yogo (2006). "Efficient tests of stock return predictability". *Journal of Financial Economics* 81.1, pp. 27–60.

CME Group (2017). *CME Group Self-Certifies Bitcoin Futures to Launch Dec. 18 [Press Release]*. URL: https://www.cmegroup.com/media-room/press-releases/2017/12/01/cme_group_self-certifiesbitcoinfuturestolaunchdec18.html.

— (2019). *CME Group Announces Jan. 13, 2020 Launch for Bitcoin Options [Press Release]*. URL: https://www.cmegroup.com/media-room/press-releases/2019/11/12/cme_group_announcesjan132020launchforbitcoinoptions.html.

Cochrane, J. H. (2011). "Presidential address: Discount rates." *Journal of Finance* 66, pp. 1047–108.

Crouzet, Nicolas and Neil Mehrotra (forthcoming). "Small and Large Firms over the Business Cycle". *American Economic Review*.

Chen, Hui, Yu Xu, and Jun Yang (2020). "Systematic risk, debt maturity, and the term structure of credit spreads". *Journal of Financial Economics*.

Chen, J., H. Hong, and J. Stein (2001). "Diagnostic Bubbles". *Journal of Financial Economics* 61, pp. 345–81.

Choi, Jaewon and Mathias Kronlund (2018). "Reaching for yield in corporate bond mutual funds". *The Review of Financial Studies* 31.5, pp. 1930–1965.

Christensen, Ian and Ali Dib (2008). "The financial accelerator in an estimated New Keynesian model". *Review of Economic Dynamics* 11.1, pp. 155–178.

Christiano, Lawrence J, Roberto Motto, and Massimo Rostagno (2014). "Risk shocks". *American Economic Review* 104.1, pp. 27–65.

Daniel, Kent, Lorenzo Garlappi, and Kairong Xiao (forthcoming). "Monetary policy and reaching for income". *The Journal of Finance*.

Darmouni, Olivier, Oliver Giesecke, and Alexander Rodnyansky (2020). "The Bond Lending Channel of Monetary Policy".

Dell'Ariccia, Giovanni, Luc Laeven, and Robert Marquez (2014). "Real interest rates, leverage, and bank risk-taking". *Journal of Economic Theory* 149, pp. 65–99.

Di Maggio, Marco and Marcin Kacperczyk (2017). "The unintended consequences of the zero lower bound policy". *Journal of Financial Economics* 123.1, pp. 59–80.

Diamond, Douglas W. (1991). "Debt Maturity Structure and Liquidity Risk". *The Quarterly Journal of Economics* 106.3, pp. 709–737.

Diamond, Douglas W. and Zhinguo He (2014). "A Theory of Debt Maturity: The Long and Short of Debt Overhang". *The Journal of Finance* 69.2, pp. 719–762.

Diether, Karl B, Christopher J Malloy, and Anna Scherbina (2002). "Differences of Opinion and the Cross Section of Stock Returns". *The Journal of Finance* 57.5, pp. 2113–2141.

Drechsler, Itamar, Alexi Savov, and Philipp Schnabl (2018). "A model of monetary policy and risk premia". *The Journal of Finance* 73.1, pp. 317–373.

Duchin, Ran, Oguzhan Ozbas, and Berk A Sensoy (2010). "Costly external finance, corporate investment, and the subprime mortgage credit crisis". *Journal of Financial Economics* 97.3, pp. 418–435.

Fahlenbrach, Rüdiger, Kevin Rageth, and René M Stulz (2020). *How valuable is financial flexibility when revenue stops? Evidence from the COVID-19 crisis*. Tech. rep. National Bureau of Economic Research.

Faulkender, Michael (2005). "Hedging or market timing? Selecting the interest rate exposure of corporate debt". *The Journal of Finance* 60.2, pp. 931–962.

Flannery, Mark J. (1986). "Asymmetric Information and Risky Debt Maturity Choice". *The Journal of Finance* 41.1, pp. 19–37.

Flannery, Mark J. (1994). “Debt Maturity and the Deadweight Cost of Leverage: Optimally Financing Banking Firms”. *The American Economic Review* 84.1, pp. 320–331.

Foley-Fisher, Nathan, Rodney Ramcharan, and Edison Yu (2016). “The impact of unconventional monetary policy on firm financing constraints: Evidence from the maturity extension program”. *Journal of Financial Economics* 122.2, pp. 409–429.

Gallant, A. R., P. Rossi, and G. Tauchen (1992). “Stock Prices and Volume”. *Review of Financial Studies* 5.2, pp. 199–242.

Gertler, Mark and Ben Bernanke (1989). “Agency Costs, Net Worth and Business Fluctuations”. *American Economic Review* 79, pp. 14–31.

Gertler, Mark and Peter Karadi (2011). “A model of unconventional monetary policy”. *Journal of Monetary Economics* 58.1, pp. 17–34.

Giambona, Erasmo, Rafael Matta, José-Luis Peydró, and Ye Wang (2020). *Quantitative easing, investment, and safe assets: the corporate-bond lending channel*. Tech. rep. Department of Economics and Business, Universitat Pompeu Fabra.

Giroud, Xavier and Holger M Mueller (2017). “Firm leverage, consumer demand, and employment losses during the Great Recession”. *The Quarterly Journal of Economics* 132.1, pp. 271–316.

— (2018). *Firm leverage and regional business cycles*. Tech. rep. National Bureau of Economic Research.

Gomes, Joao, Urban Jermann, and Lukas Schmid (2016). “Sticky leverage”. *American Economic Review* 106.12, pp. 3800–3828.

Greenwood, R., A. Shleiger, and Y. You (2018). “Bubbles for Fama”. *Journal of Financial Economics* 131.1, pp. 20–43.

Greenwood, Robin, Samuel Hanson, and Jeremy C Stein (2010). “A gap-filling theory of corporate debt maturity choice”. *The Journal of Finance* 65.3, pp. 993–1028.

Greenwood, Robin and Andrei Shleifer (2014). “Expectations of Returns and Expected Returns”. *The Review of Financial Studies* 27.3, pp. 714–746.

Grosse-Rueschkamp, Benjamin, Sascha Steffen, and Daniel Streitz (2019). “A capital structure channel of monetary policy”. *Journal of Financial Economics* 133.2, pp. 357–378.

Grossman, Sanford and Joseph Stiglitz (1980). “On the impossibility of informationally efficient markets”. *American Economic Review* 70.3, pp. 393–408.

Gürkaynak, RS, B Sack, and ET Swanson (2005). “Do actions speak louder than words? The response of asset prices to monetary policy actions and statements”. *International Journal of Central Banking* 1.1, pp. 55–93.

Halling, Michael, Jin Yu, and Josef Zechner (2020). "How Did COVID-19 Affect Firms' Access to Public Capital Markets?" *The Review of Corporate Finance Studies*.

Hanson, Samuel G and Jeremy C Stein (2015). "Monetary policy and long-term real rates". *Journal of Financial Economics* 115.3, pp. 429–448.

Harrison, J Michael and David M Kreps (1978). "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations". *Quarterly Journal of Economics* 92.2, pp. 323–336.

He, Zhiguo and Konstantin Milbradt (2016). "Dynamic Debt Maturity". *The Review of Financial Studies* 29.10, pp. 2677–2736.

He, Zhiguo and Wei Xiong (2012a). "Dynamic Debt Runs". *The Review of Financial Studies* 25.6, pp. 1799–1843.

— (2012b). "Rollover risk and credit risk". *The Journal of Finance* 67.2, pp. 391–430.

Hjalmarsson, Erik (2011). "New Methods for Inference in Long-Horizon Regressions". *The Journal of Financial and Quantitative Analysis* 46.3, pp. 815–839.

Holmström, Bengt and Jean Tirole (1998). "Private and public supply of liquidity". *Journal of Political Economy* 106.1, pp. 1–40.

— (2000). "Liquidity and risk management". *Journal of Money, Credit and Banking*, pp. 295–319.

Hong, Harrison, Jose Scheinkman, and Wei Xiong (2006). "Asset Float and Speculative Bubbles". *The Journal of Finance* 61.3, pp. 1073–1117.

Hong, Harrison and Jeremy C Stein (2003). "Differences of Opinion, Short-Sales Constraints, and Market Crashes". *The Review of Financial Studies* 16.2, pp. 487–525.

— (2007a). "Disagreement and the Stock Market". *Journal of Economic Perspectives* 21.2, pp. 109–128.

— (2007b). "Disagreement and the Stock Market." *Journal of Economic Perspectives* 21.2, pp. 109–128.

Hutto, Clayton and Eric Gilbert (2014). "Vader: A Parsimonious Rule-based Model for Sentiment Analysis of Social Media Text". *Proceedings of the International AAAI Conference on Web and Social Media*. Vol. 8. 1.

Iacoviello, Matteo (2005). "House prices, borrowing constraints, and monetary policy in the business cycle". *American Economic Review* 95.3, pp. 739–764.

IMF (2019a). *Global Corporate Vulnerabilities*. Global Financial Stability Report - Chapter 2 (April). International Monetary Fund.

— (2019b). *Vulnerabilities in a Maturing Credit Cycle*. Global Financial Stability Report (April). International Monetary Fund.

Ippolito, Filippo, Ali K Ozdagli, and Ander Perez-Orive (2018). “The transmission of monetary policy through bank lending: The floating rate channel”. *Journal of Monetary Economics* 95, pp. 49–71.

Jarocinski, Marek and Peter Karadi (2020). “Deconstructing Monetary Policy Surprises: The Role of Information Shocks”. *American Economic Journal: Macroeconomics* 12.2, pp. 1–43.

Jeenas, Priit (2018). *Firm balance sheet liquidity, monetary policy shocks, and investment dynamics*. Tech. rep. Working paper.

Jegadeesh, Narasimhan and Di Wu (2013). “Word power: A new approach for content analysis”. *Journal of Financial Economics* 110.3, pp. 712–729.

Jordà, Òscar (2005). “Estimation and inference of impulse responses by local projections”. *American Economic Review* 95.1, pp. 161–182.

Jungherr, Joachim and Immo Schott (2020a). “Corporate Debt Maturity Matters for Monetary Policy”.

— (2020b). “Optimal debt maturity and firm investment”.

Kalemli-Ozcan, Sebnem, Luc Laeven, and David Moreno (2018). *Debt overhang, rollover risk, and corporate investment: Evidence from the european crisis*. Tech. rep. National Bureau of Economic Research.

Kaplan, Robert S (2019). *Corporate Debt as a Potential Amplifier in a Slowdown*. Tech. rep. Federal Reserve Bank of Dallas.

Kashyap, Anil K and Jeremy C Stein (1994). “Monetary policy and bank lending”. *Monetary policy*. The University of Chicago Press, pp. 221–261.

Keynes, John Maynard (1936). *The General Theory of Employment, Interest and Money*. London: Macmillan.

Lettau, M. and S. Ludvigson (2001). “Consumption, aggregate wealth, and expected stock returns.” *Journal of Finance* 56, pp. 815–49.

Lian, Chen, Yueran Ma, and Carmen Wang (2019). “Low interest rates and risk-taking: Evidence from individual investment decisions”. *The Review of Financial Studies* 32.6, pp. 2107–2148.

Lo, A. W., H. Mamaysky, and J. Wang (2004). “Asset Prices and Trading Volume under Fixed Transaction Costs: Implications of an Intertemporal Capital Asset Pricing Model.” *Journal of Political Economy* 112.5, pp. 1054–89.

Lo, A. W. and J. Wang (2004). “Stock Market Trading Volume”. *Handbook of Financial Econometrics: Applications*.

— (2006). “Trading Volume: Implications of an Intertemporal Capital Asset Pricing Model.” *Journal of Finance* 61, pp. 2805–40.

Lo, Andrew W. and A. Craig MacKinlay (1988). "Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test". *The Review of Financial Studies* 1.1, pp. 41–66.

— (1990). "When Are Contrarian Profits Due to Stock Market Overreaction?" *The Review of Financial Studies* 3.2, pp. 175–205.

Loughran, Tim and Bill McDonald (2011). "When is a Liability not a Liability? Textual Analysis, Dictionaries, and 10-Ks". *The Journal of Finance* 66.1, pp. 35–65.

Lu, Lina, Matthew Pritsker, Andrei Zlate, Kenechukwu Anadu, and James Bohn (2019). "Reach for yield by US public pension funds". *FRB Boston Risk and Policy Analysis Unit Paper No. RPA*, pp. 19–2.

McQueen, Grant, Michael Pinegar, and Steven Thorley (1996). "Delayed Reaction to Good News and the Cross-Autocorrelation of Portfolio Returns". *The Journal of Finance* 51.3, pp. 889–919.

Mian, Atif and João AC Santos (2018). "Liquidity risk and maturity management over the credit cycle". *Journal of Financial Economics* 127.2, pp. 264–284.

Miller, Edward M (1977). "Risk, Uncertainty, and Divergence of Opinion". *The Journal of Finance* 32.4, pp. 1151–1168.

Newey, Whitney K and Kenneth D West (1987). "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix". *Econometrica* 55.3, pp. 703–708.

Ottonello, Pablo and Thomas Winberry (forthcoming). "Financial heterogeneity and the investment channel of monetary policy". *Econometrica*.

Ozdagli, Ali K (2018). "Financial frictions and the stock price reaction to monetary policy". *The Review of Financial Studies* 31.10, pp. 3895–3936.

Poeschl, Johannes (2017). *Corporate debt maturity and investment over the business cycle*. Tech. rep. Danmarks Nationalbank Working Papers.

Sadka, Ronnie and Anna Scherbina (2007). "Analyst Disagreement, Mispricing, and Liquidity". *The Journal of Finance* 62.5, pp. 2367–2403.

Scheinkman, Jose A and Wei Xiong (2003). "Overconfidence and Speculative Bubbles". *Journal of Political Economy* 111.6, pp. 1183–1220.

Simsek, Alp (2013). "Belief Disagreements and Collateral Constraints". *Econometrica* 81.1, pp. 1–53.

— (2021). "The Macroeconomics of Financial Speculation". *Annual Review of Economics* 13.

Tetlock, Paul C (2007). "Giving Content to Investor Sentiment: The Role of Media in the Stock Market". *The Journal of Finance* 62.3, pp. 1139–1168.

Xiong, Wei (2013). *Bubbles, Crises, and Heterogeneous Beliefs*. NBER Working Paper No. 18905. National Bureau of Economic Research.

Xu, Qiping (2018). “Kicking maturity down the road: early refinancing and maturity management in the corporate bond market”. *The Review of Financial Studies* 31.8, pp. 3061–3097.