LONDON SCHOOL OF ECONOMICS AND POLITICAL SCIENCE

Essays in Financial Economics

by

Karamfil Todorov

A thesis submitted to the Department of Finance of the London School of Economics and Political Science for the degree of Doctor of Philosophy

May 13, 2020

Acknowledgements

I would like to thank my advisors, Igor Makarov and Christian Julliard for their generous support, guidance, patience and trust. My numerous interactions with them have significantly improved my academic work. Their personal example helped me develop as a scholar.

In addition, I am deeply grateful to Dimitri Vayanos for his invaluable support and feedback on research, especially during my job market. Special thanks to Jean-Pierre Zigrand and Ron Anderson for their continuous guidance, help and, most importantly, friendship.

Moreover, I have also benefited from countless discussions with faculty members at the Department of Finance at the LSE. I would like to thank Ashwini Agrawal, Georgy Chabakauri, Thummim Cho, Vicente Cuñat, Simeon Djankov, Dirk Jenter, Peter Kondor, Dong Lou, Ian Martin, Martin Oehmke, Daniel Paravisini, Cameron Peng, Christopher Polk, Andrea Tamoni, Michela Verardo, Kathy Yuan and Hongda Zhong.

The five years at the LSE would not have been the same without the great friends I met during my PhD. In particular, I would like to thank the members of my cohort who have made my time at the LSE especially enjoyable: Andreea Englezu, Zongchen Hu, Bruce Iwadate, Francesco Nicolai, Marco Pelosi, Simona Risteska and Yue Yuan. I am also grateful to other fellow PhD students: Lorenzo Bretscher, Fabrizio Core, Brandon Han, Jesus Gorrin, James Guo, Lukas Kremens, Dimitris Papadimitriou, Alberto Pellicioli, Michael Punz, Petar Sabtchevsky, Seyed E. Seyedan and Ran Shi for their friendship and support during all these years.

I would also like to acknowledge the financial support from the ESRC as well as the Finance Department.

Most importantly, I would like to specially thank my parents Neli and Todor, and my sister Julia for their everlasting support. You have been always there for me, giving me all the motivation to grow both personally and professionally. Without you, I would have never become the person I am today.

I could not thank enough my wife Ksenia for being with me through all the ups and downs. Thank you for your love, kindness and patience.

Declaration of Authorship

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

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

I declare that my thesis consists of 75,438 words.

Statement of conjoint work

I confirm that Chapter 3 was jointly co-authored with Christian Julliard, Zijun Liu, Seyed E. Seyedan and Kathy Yuan. I contributed 20% of the work for Chapter 3.

Abstract

In the first paper of my dissertation I study the size and source of exchange-traded funds' (ETFs) price impact in the most ETF-dominated asset classes: volatility (VIX) and commodities. I show that the introduction of ETFs increased futures prices. To identify ETF-induced price distortions, I propose a model-independent approach to replicate the value of a VIX futures contract. This allows me to isolate a non-fundamental component in VIX futures prices, of 18.5% per year, that is strongly related to the rebalancing of ETFs. To understand the source of that component, I decompose trading demand from ETFs into three main parts: leverage rebalancing, calendar rebalancing, and flow rebalancing. Leverage rebalancing has the largest effects. It amplifies price changes and introduces unhedgeable risks for ETF counterparties. Surprisingly, providing liquidity to leveraged ETFs turns out to be a bet on variance, even in a market with a zero net share of ETFs. Trading against leverage rebalancing delivers large abnormal returns and Sharpe ratios above two across markets.

The second paper analyses the impact of the ECB's Corporate Sector Purchase Programme (CSPP) announcement on prices, liquidity and debt issuance in the European corporate bond market. I find that the quantitative easing (QE) programme increased prices and liquidity of bonds eligible to be purchased substantially. Bond yields dropped on average by 30 bps (8%) after the CSPP announcement. Tri-party repo turnover rose by 8.15 million USD (29%), and bilateral turnover went up by 7.05 million USD (72%). Bid-ask spreads also showed significant liquidity improvement in eligible bonds. QE was successful in boosting corporate debt issuance. Firms issued 2.19 billion EUR (25%) more in QE-eligible debt after the CSPP announcement, compared to other types of debt. Surprisingly, corporates used the attracted funds mostly to increase dividends. These effects were more pronounced for longer-maturity, lower-rated bonds, and for more credit-constrained, lower-rated firms.

The third paper (co-authored with Christian Julliard, Zijun Liu, Seyed E. Seyedan and Kathy Yuan) studies the determinants of repo haircuts in the UK market. We find that transaction maturity and collateral quality have first order importance. We also document that counterparties matter in determining haircuts. Hedge funds, as borrowers, receive significantly higher haircuts. Larger borrowers with higher ratings receive lower haircuts, but we find that these effects can be overshadowed by collateral quality. Repeated bilateral relationships also matter and generate lower haircuts. We find evidence supporting an adverse selection explanation of haircuts, but limited evidence in favor of lenders' liquidity position or default probabilities affecting haircuts.

Contents

List of Figures

List of Tables

Chapter 1

Passive Funds Actively Affect Prices: Evidence from the Largest ETF Markets

Karamfil Todorov[1]

1.1 Introduction

Recent years have seen a surge in passive investing.[2] Investors are increasingly putting money into funds that track a given benchmark index instead of actively managing a portfolio. The market for one particular type of these funds, ETFs, has grown considerably. As of 2018, ETFs were managing \$5 trillion globally compared with only \$0.2 trillion in 2004.[3] ETFs are progressively being used by retail and institutional investors to obtain a cost-efficient exposure to portfolios of assets or asset strategies.

[1]I am grateful to Igor Makarov, Christian Julliard, Dimitri Vayanos, Jean-Pierre Zigrand, Ron Anderson, Thummim Cho, Dirk Jenter, Dong Lou, Ian Martin, Martin Oehmke, Daniel Paravisini, Cameron Peng, Andrea Tamoni, Andreas Uthemann and Kathy Yuan for their useful advice and comments. I am also thankful to the scientific committees of the BlackRock Applied Research Award and the 7th SUERF/UniCredit Foundation Research Prize for selecting this paper as the winner in these competitions. I am also grateful to seminar participants at the London School of Economics, the Bank of England, BIS, BlackRock, Boston College, CEMFI, HEC Paris, Indiana University, PIMCO, UniCredit, University of Florida, University of Houston, University of Maryland, University of Notre Dame, and Washington University in St. Louis for their helpful feedback. I thank the Systemic Risk Centre at the LSE for providing the data on variance swaps sourced from Markit under license.

[2]Passive funds do not pick the assets to invest in, but simply follow their benchmarks (as opposed to active funds). Since some benchmarks involve more frequent rebalancing, passive funds are not the same as pure buy-and-hold investors.

[3]Source: Morningstar and own calculations.

On the one hand, commoditization of assets through ETFs makes investing simple and cost-efficient, thereby attracting new capital and possibly increasing liquidity. On the other hand, commoditization could reduce price informativeness and create systemic risks if the presence of large investors with similar objectives leads to synchronized trading, especially during extreme market times. The increasing presence of ETFs in various asset classes has led to a growing number of market participants and academics expressing concerns about the potential distorting impact on underlying assets. The fear is that "too much money is in too few hands".

Assessing the impact of ETFs on prices is difficult because it is hard to distinguish between noise and fundamentals of the underlying asset. The existing literature has almost exclusively focused on equity markets, where fundamental values are complicated to measure. Most papers have tried to quantify non-fundamental price distortions due to ETFs by looking at price reversals or variance ratios. In the research presented here, I use the beneficial setting of the futures market, where non-fundamental price distortions are easier to measure. I construct a unique data set to identify the size and source of the ETF impact on prices in the most ETF-dominated asset classes: volatility (VIX) and commodities.[4]

These ETFs have two beneficial features that make them a useful laboratory to quantify the effects of ETFs on prices. First, ETFs in VIX and commodities hold a much larger share of the market compared to equities. The fraction of ETFs in the market for VIX futures often exceeds 30%, whereas it is less than 2%[5] in the Standard and Poor's (S&P) 500 Index. Several episodes from the VIX market in 2018 (and from the oil market in 2020) suggest that large ETF-induced trading can actively move prices and exacerbate price changes in turbulent times.[6] Second, using the specifics of futures contracts, I directly test whether the ETF-influenced futures price is informative about the fundamental spot value, or is more influenced by less fundamental premiums. The setting of the futures market also allows me to test specific predictions about the price impact of ETFs on the slope of the futures curve.

This paper documents and studies several new ETF-related phenomena. First, I show that ETFs put pressure on prices of underlying assets in VIX and commodity markets.

[4]Some of the exchange-traded products (ETPs) analyzed in this research are structured in the form of an exchange-traded note (ETN) rather than an exchange-traded fund. The institutional differences between the two structures are immaterial since ETPs' exposure is transmitted to the underlying futures market irrespective of the legal structure of the product as I show in section 1.5.3. I use the term ETF (instead of ETP) to refer to a general exchange-traded product throughout the paper as the term is more familiar to the general public.

[5]On average, for the period 2009–2018. The average proportion of the US stock market held by all equity ETFs is close to 6% for the same period.

[6]See, e.g., Pagano et al. (2019), FT (2018), Reuters (2018), FT (2020) and Bloomberg (2020).

Trading demand from ETFs (called ETF demand hereafter) is strongly related to futures prices at a daily frequency. The effects are robust to a large set of controls and to different sub-periods.

Second, I show that ETF price impact is not related to price discovery but manifests itself through an increase in the non-fundamental part of prices. To identify ETF-induced price distortions, I propose a model-independent approach for replicating the fundamental value[7] of a VIX futures contract. I simply use the definition of variance and construct a synthetic futures contract from option prices on the S&P 500 Index and VIX. One advantage of my framework is that I make no parametric or distributional assumptions: the results are also valid in the presence of jumps. This is an important strength of my approach, given that VIX futures often experience large spikes. The synthetic futures contract is not directly influenced by ETF demand since there are no ETFs in the market for options. The price of the replicated contract was close to that of the traded one before the introduction of ETFs but diverged consistently thereafter. I show that the difference between the prices of the two contracts is strongly related to ETF demand and call this difference the **ETF futures gap (EFG)**. The EFG is also related to measures of funding and market liquidity: bid-ask spreads and the TED spread (spread between 3-month LIBOR in USD and the interest rate of Treasury bills). The size of the gap is 18.5% per year, on average.

Third, to study the source of the gap in VIX futures prices, I analyze trading by ETFs and propose a novel decomposition of their demand into three major components: calendar rebalancing, flow rebalancing, and leverage rebalancing. Calendar rebalancing arises because futures are finite-maturity instruments as opposed to stocks, and ETFs have to gradually roll expiring contracts into longer-dated ones to maintain their exposure. ETFs sell portions of the first-month futures and buy portions of the second-month futures on a daily basis, thereby rolling their exposure from the first to the second contract. Flow rebalancing is driven by fund flows: ETFs have to scale up their exposure in case of inflows, and scale it down in case of outflows. Leverage rebalancing arises due to the maintenance of a constant daily leverage by leveraged ETFs and is a new type of mechanic institutional demand. The three types of rebalancing are not specific to futures-based ETFs but can be generalized to ETFs in all asset classes, including equity and fixed income. Calendar rebalancing is analogous to the roll due to benchmark exclusion/inclusion for equity ETFs, or due to maturing bonds for fixed income ETFs. Flow and leverage rebalancing for equity and bond ETFs have similar interpretation to that for VIX and commodity ETFs.

[7]Throughout the paper, fundamental value denotes a value that is a more precise measure of the fundamental spot price at maturity compared to the observed, ETF-influenced futures price.

I show that leverage rebalancing introduces a source of convexity that is not easy to hedge and exposes ETF counterparties (called arbitrageurs hereafter) to variance. This type of rebalancing has the largest impact on the EFG. Calendar rebalancing inherits part of the non-linearity of leverage rebalancing and also mechanically exposes arbitrageurs to the risk of widening price discrepancies since arbitrageurs have to close futures positions before expiration. Calendar rebalancing puts upward pressure on the second-month futures price and downward pressure on the first-month price. Flow rebalancing has a direct effect on prices and also an indirect effect by changing the size of ETFs, and the amount of their calendar rebalancing in future periods. Flow rebalancing moves prices in the direction of flows: inflows increase prices, whereas outflows decrease them. The short-term price impact of the three types of rebalancing translates into longer-term price deviations.

Given that the EFG is most sensitive to leverage rebalancing, I next analyze the risks faced by an arbitrageur who trades against this type of ETF demand. The impact of large leverage-induced trading by institutional investors on systemic risk and prices is an important, yet under-researched question. ETFs provide a useful laboratory to study this question on a daily basis. I show that leverage rebalancing amplifies price changes and introduces unhedgeable risks for ETF counterparties. Leveraged ETFs mechanically have to buy the underlying asset after price increases and sell it after price decreases. This creates a potential feedback channel for prices: ETF demand and price changes reinforce each other, pushing prices away from fundamentals. Trading against leveraged ETFs is, in essence, providing liquidity to investors with short horizons, who follow momentum-like strategy. Due to leverage rebalancing, the potential distorting effect of ETFs on prices can be large even in a market with a zero net share of ETFs. A prominent real-world example of this effect was the VIX market in February 2018. The net market share of ETFs then was close to zero, but the potential price impact due to leverage rebalancing was 60% of the total market size for the first two futures contracts. The amount of leverage rebalancing has been growing in the last decade not only for VIX and commodities but also for some equity and bond indices.

I propose a simple strategy to understand the risks of trading against leverage rebalancing, and document a novel ETF-related anomaly. I form a portfolio that sells short a pair of ETFs with opposite leverages L (e.g., $L = 2$ and $L = -2$), to approximate liquidity provision to leveraged ETFs. A natural guess could be that such a strategy should have a zero return since the profit (loss) from selling the long fund is canceled by the loss (profit) from selling the inverse fund. Surprisingly, I show that the returns on such a strategy are not zero, but are consistently positive across markets. The portfolio delivers annualized returns of 21% in the VIX market, and 43% in the gas market, with

Sharpe ratios of 0.89 and 2.59, respectively. Theoretically, trading against opposite-leveraged ETFs should be negatively exposed to large jumps in the underlying asset. However, empirically, the downside risk is almost never realized. A more puzzling observation is that the strategy is exposed to a "right-way risk" since it benefits from ETF tracking errors in crisis times when liquidity dries up and the market breaks down. To the best of my knowledge, my paper is the first to quantify the risks of trading against leveraged ETFs and the first to show that their tracking errors co-move with extreme market times, thereby benefiting liquidity providers.

The main conclusions of this paper are validated in a large set of robustness tests. I verify the major results using several measures of futures basis and spread to address concerns about time to maturity, the absolute level of prices, or other factors driving the relationship between ETF demand and futures prices. I conduct several tests to verify that the ETF futures gap and its relationship to the trading demand from ETFs are robust features of the data. The empirical evidence shows that the gap is unlikely to be driven by the absence of a continuum of option strikes, truncation errors, differences in margin requirements, or hedging pressure in the options market. A potential concern is that price discovery takes place in the ETF-influenced market, and therefore the gap exists because of fundamental information about the realized spot rather than non-fundamental price pressure due to ETF demand. To address this concern, I test which futures contract is a better predictor of the fundamental spot price at maturity. The empirical evidence shows that the ETF-influenced futures contract is a poor predictor, whereas the synthetic futures constructed from options is a better one.

My main results have several implications. First, they illustrate that price is strongly related to ETF trading demand, when ETFs constitute a large share of the market. In turbulent times, significant leverage-induced rebalancing contributes to extreme market movements and creates a feedback effect on prices. Thus, synchronized trading by ETFs actively moves prices away from fundamentals. This result contributes to the policy debate on the desirability of commoditization. Going forward, the evidence of ETF impact on prices in VIX and commodity markets can be useful for predicting the potential effects of ETFs on stock and bond markets, should these funds develop a larger share of these traditionally studied asset classes.

Second, my results lead to a more nuanced view of the information content of VIX and the VIX futures premium (the average relative difference between the current futures price and the realized spot price at expiration). VIX and its derivatives are often perceived as a barometer of financial stress by large financial institutions and are used as an input in stress-tests and various risk models. The VIX futures premium is frequently interpreted as a measure of investors' risk aversion and future economic uncertainty. However, my results suggest that the prices of VIX futures contracts are significantly disrupted by non-fundamental mechanical ETF demand. I show that the VIX futures premium has increased since the introduction of ETFs, particularly, for the most ETF-influenced futures contracts: one and two months. Prices for these contracts are less informative about the realized spot price and are influenced more by premiums.

Third, my findings show how to decompose trading demand from ETFs and study different aspects of their price impact. I also demonstrate how to quantify the potential distorting impact of leverage rebalancing and how to capture the risk premium of trading against this rebalancing.

The rest of this paper is organized as follows. In section 1.2, I summarize the relevant literature. Section 1.3 describes the data set. Section 1.4 presents the results on ETF impact in the VIX market. Section 1.5 describes the decomposition of ETF demand. Section 1.6 analyzes in detail ETFs' leverage rebalancing. Section 1.7 presents the results for commodity markets, and section 1.8 concludes.

1.2 Literature review

The research presented here contributes to the literature in four main areas: studies on ETFs, VIX and the variance risk premium (VRP), futures markets, and limits to arbitrage. First, it is related to studies on ETFs. A major drawback of the existing ETF literature is that it is almost exclusively focused on equities, where these funds are a relatively small proportion of the market (less than 6%, on average) and where non-fundamental price deviations are hard to measure. Ben-David et al. (2018) study the volatility effects of ETFs in stocks and argue that ETF arbitrage transmits noise trader risk to underlying securities. Malamud (2015) demonstrates that ETFs may create a transmission mechanism for non-fundamental shocks to the underlying securities. Cheng and Madhavan (2009) show that the returns on leveraged ETFs are path-dependent. Tuzun (2012) finds that the rebalancing of leveraged ETFs can increase the volatility of constituent stocks, whereas Ivanov and Lenkey (2018) find evidence that these effects are offset by ETF flows. Bessembinder (2015) studies the impact of ETF flows. Recently, Sushko and Turner (2019) document the increase in the share held by ETFs in several markets and study the impact for liquidity and volatility. The greater presence of ETFs in VIX and commodity markets makes them a natural candidate for studying the impact of ETF demand on prices and risk premiums. And yet, ETFs in VIX and commodity markets have received little to no attention in the literature to date. A related paper by Dong (2016) studies the price impact of VIX ETFs and finds that dealers pass hedging pressure to underlying futures. Similar to the research presented here, Dong (2016) uses a result from Carr and Wu (2006) to estimate the fair value of a VIX futures. However, he does not establish that the fair value relates more closely to fundamentals than the VIX futures price, and does not rule out alternative explanations to ETF demand. The paper also does not study the source of the premium in VIX futures prices. In particular, it ignores the impact of leverage rebalancing. Another gap in the existing ETF literature is that, to the best of my knowledge, none of the studies has decomposed the demand from ETFs to analyze different aspects of their price impact. The research presented here aims to fill this gap by examining different types of trading demand by ETFs in the largest ETF-dominated markets.

Second, my paper contributes to the literature on VIX and the variance risk premium. Cheng (2019) analyzes the VIX premium and finds that, in turbulent times, dealers and asset managers reduce their long volatility positions, whereas hedge funds reduce their short volatility positions. Barras and Malkhozov (2016) find that the VRP inferred from equity prices is different from the one inferred from option prices. They claim that the difference is driven by the financial standing of intermediaries. Dew-Becker et al. (2017) study variance swaps data and find that news about future volatility is unpriced, but investors are willing to pay a large premium to hedge realized volatility. Mixon and Onur (2015) study volatility markets and claim that the long volatility bias of asset managers acts to put upward pressure on VIX futures prices. Alexander and Korovilas (2012) find that ETFs increase the volatility of VIX futures. Eraker and Wu (2017) develop a model with diffusive and jump shocks to explain the large negative returns on VIX ETFs. However, most papers ignore the impact of ETFs on prices of underlying futures or rely on some parametric assumptions to calculate risk premiums. In the research presented here, I show that the mechanics of ETF rebalancing distorts futures prices, and I measure the resulting gap in a model-independent way.

Third, my research adds to the extensive literature on futures markets. In the normal backwardation theory of Keynes (1930) hedgers push down futures price. Fama and French (1987) find evidence of time-varying expected premiums in several commodities. Koijen et al. (2018) explore the predictive power of carry across several asset classes. Mou (2011) studies commodity roll. Tang and Xiong (2012), and Basak and Pavlova (2016) study the financialization of commodities due to institutional flows. Singleton (2013) argues that flows from institutional investors have contributed significantly to the increase in oil prices prior to 2008. In contrast, Gorton et al. (2013) find no evidence that the positions of participants in futures markets predict risk premiums on commodity futures. Most studies on hedging pressure in futures markets use lower-frequency data (quarterly or weekly) on investors' positions to analyze price impact. However, hedging pressure is more likely to be pronounced over short time horizons in the current era of high-frequency trading. Using ETFs to analyze the impact on prices allows the capture of these transitory price effects, since trading demand is observed on a daily basis. The research presented here shows that in markets with a high proportion of ETFs, the demand from these funds is strongly related to the futures premium.

Fourth, my paper also adds to the extensive literature on limits to arbitrage and slow moving capital. Shleifer and Vishny (1997) develop a simple model with noise trader risk and show that arbitrage could persist. Garleanu et al. (2009) show that dealers provide liquidity in option products and charge for the unhedgeable risks they take due to the impossibility of trading continuously, and due to transaction costs. Gromb and Vayanos (2018) develop a theoretical framework in which financially constrained arbitrageurs exploit price-discrepancies across segmented markets. In the research presented here, I document segmentation in the VIX futures market and develop a simple model to illustrate how price discrepancies can exist and persist.

1.3 Data and institutional details

I construct a unique data set on ETFs and their underlying securities in VIX and commodities: US natural gas, silver, gold and oil.[8] The data comes from several sources. Daily prices, flows, holdings, assets under management, volume of trading, and other characteristics of ETFs come from the websites of the sponsors, from the Center for Research in Security Prices (CRSP), and from Bloomberg. Daily data on futures prices, open interest, and volume of trading is from the Chicago Board Options Exchange (CBOE). Daily data on S&P 500 Index options comes from OptionMetrics and ETF borrowing fees are from Interactive Brokers. Weekly data on positions in futures contracts is from the Commodity Futures Trading Commission (CFTC), monthly data on variance swaps quotes comes from Markit Totem. The analyzed period is generally January 2000 to December 2018; however, some of the data is only available for a shorter time period.

[8] In the Appendix, I provide some results for leverage rebalancing of equity, fixed income and foreign exchange ETFs.

1.3.1 The presence of ETFs in different markets

Figure 1.1 shows the proportion of ETFs in the total market capitalization for several markets. The black lines on the graphs show that the proportion of ETFs periodically exceeds 30% of the total market capitalization for VIX and natural gas.[9] The share of ETFs in equities is much smaller and constitutes less than 2%, on average, for most equity indices. Table 1.1 summarizes the proportion of ETFs in the total market capitalization and in daily trading volume across several markets.

[Figure 1.1 and Table 1.1 about here]

The numbers illustrate that, even though ETFs constitute a smaller share of the equity market, the average daily trading volume of ETF shares is large in relative terms: for the S&P 500 Index, it exceeds 18% of the trading volume in constituent shares. In the VIX market, the average daily trading volume exceeds twice that of the underlying futures contracts. In the following analysis, I focus mainly on VIX ETFs and ETFs in the markets for natural gas, silver, gold, and oil.

1.3.2 Institutional details

Unlike most equity ETFs that physically invest in the underlying assets, VIX and commodity ETFs obtain price exposure by entering into positions in futures contracts.[10] Most ETFs follow a benchmark based on the first two futures contracts.[11] They gradually roll their exposure from the first-month contract to the second-month contract (daily for VIX and over a period of 5 days each month for commodity markets). Some ETFs also aim to maintain a constant *daily* leverage ratio, L, which can also be negative (for inverse ETFs). For example, if the benchmark return is 5%, a double-leveraged ($L = 2$) ETF should return 10%, whereas an inverse ETF ($L = -1$) should return -5%. I explain in detail the exact trading motives of ETFs when I analyze the decomposition of ETF demand in section 1.5. I also show that leverage rebalancing can exacerbate market movements and analyze the risks of trading against leveraged ETFs in sections 1.5.2 and 1.6. ETFs are limited to trade in the futures contracts in a mechanical way to minimize their tracking errors. This can lead to a reduction in price informativeness and an increase in the futures premium.

[9] Also for oil in April 2020.

[10] This is how VIX, gas, oil, and many silver and gold ETFs invest. Some silver and gold ETFs hold physical silver or gold: I exclude these from the analysis since they do not rebalance on a daily basis, but physically hold the assets.

[11] There are also VIX ETFs that invest in fourth to seventh-month futures contracts but their share is much lower. I analyze these in section 1.12.10 in the Appendix.

1.3.3 Summary statistics

Figure 1.12 in the Appendix shows the term structure of futures prices for several markets. The term structure of VIX futures is in contango (futures larger than spot) 78% of the time. The picture for the gas market is similar. Generally, after the introduction of ETFs, the futures curve shifts upwards and becomes more concave for VIX and gas.

[Table 1.2 about here]

Table 1.2 presents summary statistics for the VIX market. The distribution of VIX futures is positively-skewed, particularly for short maturities. The return on the largest long VIX ETF (VXX) is -17 basis points (bps) per day: short-selling the ETF generates an annualized return of nearly 54%, consistent with the idea that investors are willing to pay a premium to hedge against volatility shocks. The average slope of the short end of the futures term structure steepened after the introduction of ETFs. The first-month futures basis[12] went up from 0.06 to 0.79, and the spread between the second and the first-month futures contracts increased from 0.46 to 1.10. The increase is also large if I exclude the 2008–2009 financial crisis from the sample. The slopes of other parts of the curve that are not influenced by ETFs, are little changed. The plot in Figure 1.2 shows that the realized VIX futures premium (the return for an investor who sells short a fully collateralized VIX futures contract and holds it until maturity $\frac{F_{t,T}-F_{T,T}}{F_{t,T}}$) has increased since the introduction of ETFs for the most ETF-influenced maturities. The rise is particularly pronounced for one and two month futures contracts. The increase is even more striking if I include the 2008–2009 financial crisis (Figure 1.13 in the Appendix). These results provide initial evidence that the introduction of ETFs is related to the increase in premiums embedded in VIX futures prices.

[Figure 1.2 about here]

1.3.4 Types of traders in the futures market

Empirically, ETFs are usually net buyers of futures contracts, whereas managed money (usually hedge funds) takes the opposite side of the trade. Figure 1.3 shows the weekly positions of different types of investors in VIX and gas markets. The graphs show that

[12] For convenience, I call $F_{t,T} - S_t$ basis throughout the paper. Since the VIX futures term structure is in contango most of the time, it is more convenient to work with $F_{t,T} - S_t$ rather than $S_t - F_{t,T}$. T is maturity, $F_{t,T}$ is time t's price of a futures contract expiring at T, S_t is time t's spot price.

leveraged money (mostly hedge funds) is consistently short VIX futures after 2009, whereas asset managers and dealers are mostly net long.[13] In the gas market, dealers (swap positions) are, on average, net long futures, whereas managed money takes the opposite side and provides liquidity. The bottom left panel of Figure 1.3 shows that dealers' positions for VIX match pretty closely the net positions of ETFs over time: the correlation is around 84%. For gas, ETF positions can be identified with net swap positions as shown in the bottom right panel. Since ETFs constitute a smaller share of the market for silver, gold, and oil, their net positions cannot be identified very precisely with net swap positions. In the next section, I study the impact of ETFs on prices in the largest ETF-dominated market.

[Figure 1.3 about here]

1.4 The impact of ETFs on futures prices in the VIX market

In this section, I analyze the impact of ETFs on futures prices in the VIX market, as this market has the largest proportion of ETFs, and sees ETF trading volume dwarf that of the underlying asset.

1.4.1 Details on VIX

VIX is an important asset for investors because it provides a natural hedge against market downturns. VIX_t^2 is a portfolio of options that measures risk-neutral expectation of realized variance of the S&P 500 Index return (assuming no jumps) over the next month: $VIX_t^2 = \mathbb{E}_t^{\mathbb{Q}}(Rvar_{t,t+30})$.[14] In a world with jumps, VIX_t^2 measures risk-neutral entropy of the S&P 500 Index return as Martin (2015) demonstrates: $VIX_t^2 = \log \mathbb{E}_t^{\mathbb{Q}} R_{t\to t+30} - \mathbb{E}_t^{\mathbb{Q}} \log R_{t\to t+30}$, where $R_{t\to t+30}$ is the gross return on the S&P 500 Index over the next 30 days. Thus, by construction, VIX increases in turbulent times when volatility spikes and aggregate economic uncertainty increases. An important feature of the market for VIX futures, which distinguishes it from traditional futures markets, is that there is no cost-of-carry relationship because the spot asset is, in essence, not physically tradable. Since the portfolio of options underlying the VIX

[13] The names of the different groups of traders are the same as in the classification in the Traders in Financial Futures (TFF) report by the CFTC. Leveraged funds typically consist of hedge funds and other proprietary traders.

[14] As realized variance is a consistent estimator of quadratic variation (see, e.g., Cheng, 2019).

calculation is changing almost continuously, in practical terms it is impossible to trade VIX due to large transaction costs. The simplest way to get exposure is by trading VIX futures.

The market for VIX derivatives has grown considerably over the past decade and has become the largest market for volatility for short maturities (see, e.g., Mixon and Onur, 2015). As of 2018, the total notional value of VIX futures exceeded that of variance swaps for maturities of less than one year. The total dollar volatility exposure (in terms of vega) of VIX futures was close to \$8 billion (bn) per month. Part of the massive increase in volatility investing was due to the rise of ETFs, which provided a simple and cost-efficient way to invest in VIX.

The inception of ETFs was a market innovation that allowed many retail investors who could not easily trade volatility before (due to margin requirements or the need to manage futures expiration dates), to enter the VIX futures market. Data from Thomson Reuters Institutional Holdings shows that the fraction of institutional holdings in VIX ETFs was less than 24%, on average, for the period 2009–2018. Figure 1.14 in the Appendix shows the large increase in open interest of VIX futures after the introduction of ETFs, particularly for the first and second-month contracts. The case of VIX provides a useful laboratory to study the effects of commoditization of trading strategies and the consequences of letting retail investors enter more sophisticated markets. The lessons learned from the VIX market can be useful to predict the effects of larger ETF presence in other markets that were less accessible to retail investors before, e.g., less-liquid corporate bond markets. The excessive inflow of new capital through ETFs could have pushed VIX futures prices away from their fundamental values. This problem is particularly pronounced for VIX given the inability of investors to replicate the futures contract by trading the spot.

1.4.2 Futures price – fundamental or not?

The benefit of studying futures-based ETFs is that I can directly test whether the underlying futures price contains more information about fundamentals, or is biased due to non-fundamental noise. I do this by testing whether the futures price is informative about the fundamental spot price at maturity, or is influenced by premiums.

Historically, the market for VIX futures has been in contango most of the time as shown in Figure 1.12 in the Appendix. First, I analyze whether the contango predicts fundamental spot price changes to maturity, or is more informative about the non-fundamental futures premium. Using the identity $F_{t,T} - S_t = F_{t,T} - F_{T,T} + S_T - S_t$

and without making any assumptions, I check whether time t's basis $F_{t,T} - S_t$ predicts subsequent changes of the spot VIX S_t until maturity T (fundamental information), or the futures $F_{t,T}$ (non-fundamental premium), or both. The left panel of Figure 1.4 illustrates the idea. I run two simple predictive regressions:[15]

$$S_T - S_t = \alpha_1 + \beta_1 \cdot (F_{t,T} - S_t) + \epsilon_{1,t}, \tag{1.1}$$

$$F_{T,T} - F_{t,T} = \alpha_2 + \beta_2 \cdot (F_{t,T} - S_t) + \epsilon_{2,t}. \tag{1.2}$$

By subtracting equation (1.2) from equation (1.1), we see that $\beta_1 - \beta_2$ should be equal to one. Table 1.3 shows the results from the two regressions. The estimates show that a larger basis (higher futures compared with spot) predicts negative change for short-maturity futures, i.e. the futures will drop at expiration. The estimates are the largest in absolute value for one month and are monotonically decreasing for longer maturities. In contrast, the basis for short-maturities does not predict spot changes.

[Figure 1.4 and Table 1.3 about here]

The picture is the opposite for the basis for longer maturities: it positively predicts subsequent spot changes. The front end of the VIX futures' curve could be influenced by ETFs, whereas the long end of the curve would not be, since ETFs trade mostly in the first and the second-month futures. The fact that the basis for one and two months predicts futures but not spot, and in contrast a longer-term basis predicts spot but not futures, shows that short-term basis could be influenced by non-fundamental risk premiums that are unrelated to fundamental realized spot changes. These results are consistent with the findings of Aït-Sahalia et al. (2018) who show that only risk on the short-end of the term structure of variance swaps is priced. Next, I use the high frequency of ETF positions and futures price changes to study whether ETF demand is related to the fact that ETF-influenced futures maturities are less informative about fundamentals.

1.4.3 ETF demand impact on futures prices

To see if the premium in VIX futures prices is related to ETFs, I study the effect of ETF rebalancing demand on different parts of the futures curve. I run the following

[15] Running regressions (1.1) and (1.2) with $S_T - S_t$, $F_{T,T} - F_{t,T}$ and $F_{t,T} - S_t$ scaled by the time to maturity of the futures yields similar results. The results are also unchanged if I control for lags of VIX, time to maturity of the futures, liquidity, open interest, and other factors. To address the concern that the mean-reversion of volatility is driving the results, I run the regressions for the period before ETFs and find statistically significant predictive power of one and two-months basis for spot prices. In the Appendix (Figure 1.15), I present the results of similar regressions for other futures markets.

regression:[16]

$$b_{t,i} = \alpha + \beta_1 D^{\$,all}_{t,i} + \beta_2 b^H_{t,i} + \gamma Ctrl_{t,i} + \epsilon_{t,i}, \tag{1.3}$$

where $b_{t,i}$ is either the absolute basis for maturity one month ($b_{t,1} = F_{t,T_1} - S_t$), or the spread between two subsequent futures (most results are with the second-month spread: $b_{t,2} = F_{t,T_2} - F_{t,T_1}$). In some specifications $b_{t,i}$ is the relative basis or the relative spread ($b_{t,1} = \frac{F_{t,T_1} - S_t}{S_t}, b_{t,2} = \frac{F_{t,T_2} - F_{t,T_1}}{F_{t,T_1}}$). S_t is the spot price, F_{t,T_1} is the price of the first futures contract and F_{t,T_2} is the price of the second one. I use basis as the main dependent variable, instead of the raw futures price, to isolate price movements mostly related to premiums as opposed to the spot price. I analyze separately first-month basis and spread, instead of first and second-month bases (as in Mixon and Onur, 2015) to disentangle the local effects of ETF demand on different parts of the curve. Since the second-month basis is the first-month basis plus spread ($F_{t,T_2} - S_t = F_{t,T_2} - F_{t,T_1} + F_{t,T_1} - S_t = b_{t,2} + b_{t,1}$), using the second-month basis as the dependent variable could capture some of the effects of ETF demand on the first-month basis. Therefore, I focus on spread and control for first-month ETF demand, to isolate the residual price impact between the first- and the second-month contracts. In the Appendix (Table 1.13), I also present the results for the second-month basis.

$D^{\$,all}_{t,i}$ is the net dollar demand from all ETFs for maturity i at time t computed as the sum of changes in dollar holdings from $t-1$ to t for all ETFs. To isolate the effect of a larger share of ETF demand from a pure increase in the size of the overall market, I normalize the demand from ETFs by market capitalization. I use the scaled demand as the main explanatory variable: $\frac{D^{\$,all}_{t,i}}{Mkt\ cap_{t,i}} = \frac{D^{all}_{t,i}}{OI_{t,i}}$, where $D^{all}_{t,i}$ is demand in terms of number of contracts and $OI_{t,i}$ is the open interest for futures with maturity i. Running the main regressions with $D^{all}_{t,i}$ instead of $\frac{D^{all}_{t,i}}{OI_{t,i}}$ also produces statistically and economically significant estimates.

$b^H_{t,i}$ is the basis or spread of a hedge asset. The hedge asset is a synthetic VIX futures contract with the same maturity as the traded one but not influenced by ETF demand – section 1.4.4 explains in detail the exact replication. It absorbs any asset-specific shocks. $Ctrl_t$ are controls for spot price,[17] open interest, days to maturity, return on the benchmark, variance of the benchmark, and liquidity differences between the futures and the hedge asset. I measure liquidity differences by the difference in relative bid-ask spreads between the futures contract and the hedge asset. Other measures of

[16] The results of the augmented Dickey-Fuller (see, e.g., Dickey and Fuller, 1981) and the Kwiatkowski-Phillips-Schmidt-Shin (see, e.g., Kwiatkowski et al., 1992) tests showed that basis, EFG, and demand from ETFs are all stationary.

[17] I perform a robustness test replacing spot VIX with VSTOXX (the VIX analogue for the EURO STOXX 50 index) similar to Mixon and Onur (2015) and find similar results to the ones in the main text.

liquidity produce similar results. When the dependent variable is $b_{t,2}$, I also control for $D^{\$,all}_{t,1}$. If ETF demand has an impact on futures prices, $\beta_1 \neq 0$.

The results are presented in Panel A of Table 1.4. For comparison, I standardize all independent variables. Columns 1 and 6 show that one standard deviation rise in ETF demand as fraction of total market capitalization (2.42% for the first-month futures contract and 5.73% for the second) increases the front-month basis by 0.26 volatility points (33% in relative terms given the mean for the period before ETFs) and the spread by 0.17 volatility points (15% in relative terms). The estimates change slightly once I include several controls as seen from columns 2 and 7. The fact that $\beta_1 > 0$ is evidence of the impact of ETF demand on the price of underlying futures contracts.

[Table 1.4 about here]

One concern about using absolute basis and spread is that, on average, spot VIX and VIX futures were lower in 2014–2018 than they were in 2009–2013. Thus, absolute quantities in different time periods may not be comparable. Columns 3 and 8 of Panel A present the results with absolute basis scaled by S_t and spread scaled by F_{t,T_1} – essentially, measuring the slopes of different parts of the futures curve. Using relative basis and spread would also allow me to compare the results across several markets in section 1.7. The estimates show that one standard deviation rise in ETF demand increases the basis by 1.41% and the spread by 0.48%, respectively. Using absolute ETF demand instead of demand scaled by market capitalization, gives even higher estimates, as shown in columns 4 and 9. Another possible concern is that the relationship between futures prices and ETF demand could be driven by different time to expiration. Scaling the relative basis and spread by days to maturity (columns 9 and 10 of Table 1.14 in the Appendix) shows that this is not the case: the coefficients are still positive and highly statistically significant.

One potential problem with the contemporaneous regressions is that both $b_{t,i}$ and $\frac{D^{\$}_{t,i}}{Mkt\ cap_{t,i}}$ depend on contemporaneous futures returns $r^{F_1}_t$ and $r^{F_2}_t$ (as I show in section 1.5). Columns 5 and 10 present the results of a regression where $D^{\$,all}_{t,i}$ is calculated using lagged returns on the two futures and the ETF benchmark instead of the contemporaneous ones. The coefficients are still positive and statistically significant but lower in magnitude. This is mostly driven by the fact that the benchmark return is not very persistent, and is slightly negatively correlated at the first lag.

The effects of ETF demand are significant only for the respective maturities in which ETFs invest, but there is limited evidence of significant changes in the slopes of other parts of the curve: the estimates for the spreads of third, fourth, fifth and sixth-month futures contracts are mostly insignificant.[18]

1.4.4 The ETF futures gap

Panel A of Table 1.4 illustrates that an increase in ETF demand pushes up short-term basis and spread. Although the OLS regressions control for a large set of observable characteristics, the estimates could be biased due to endogeneity if both ETF demand and futures prices are influenced by a fundamental omitted variable. To address this concern, I disentangle the non-fundamental component of prices and analyze whether the effect of ETF demand manifests itself through an increase in that component.

One of the benefits of analyzing the VIX market is that I can directly measure deviations in the futures price due to ETF demand by constructing a synthetic VIX futures contract from a market with no ETFs, and comparing its price to the price of the ETF-influenced VIX futures contract. The idea is simple. I calculate $\mathbb{E}_t^{\mathrm{Q}}(S_T)$ (Q is the risk-neutral measure) from option prices without making any parametric or distributional assumptions. By comparing $F_{t,T}$ and $\mathbb{E}_t^{\mathrm{Q}}(S_T)$, I can isolate the component of the VIX futures premium that is different between the futures market and the options market. Then, I test directly which of the two futures prices (the ETF-influenced one $F_{t,T}$, or the synthetically constructed one $\mathbb{E}_t^{\mathrm{Q}}(S_T)$) is a less-biased estimate of the fundamental spot price at expiration. To illustrate the approach, note that basis can be decomposed as follows:[19]

$$F_{t,T} - S_t = \underbrace{F_{t,T} - \mathbb{E}_t^{\mathrm{Q}}(S_T)}_{\textbf{ETF futures gap}} + \underbrace{\mathbb{E}_t^{\mathrm{Q}}(S_T) - S_T}_{\textit{Realized VIX premium}} + \underbrace{S_T - S_t.}_{\textit{Spot VIX change}} \tag{1.4}$$

The right panel of Figure 1.4 illustrates the split. The decomposition shows that time t's basis consists of the **ETF futures gap (EFG)**, the realized VIX premium (RVP), and the realized spot VIX change. The element of interest in this decomposition is the EFG. This component is different from zero if there is market segmentation or other frictions. I show that ETF demand manifests itself through an increase in this non-fundamental part of the futures price in sections 1.4.4.3 and 1.4.4.4.

[18]The results for VIX ETFs investing in midterm maturities of the futures (section 1.12.10 of the Appendix) show that their demand is significant for $b_{t,5}$, $b_{t,6}$ but less significant for $b_{t,4}$ and $b_{t,7}$.

[19]Note that spread can be decomposed as follows: $F_{t,T_2} - F_{t,T_1} = F_{t,T_2} - S_t - (F_{t,T_1} - S_t) = (EFG_{t,T_2} - EFG_{t,T_1}) + (RVP_{t,T_2} - RVP_{t,T_1}) + (S_{T_2} - S_{T_1})$.

1.4.4.1 Calculating the EFG

To calculate the EFG for maturity T_1, I measure $\mathbb{E}_t^{\mathrm{Q}}(S_{T_1}) = \mathbb{E}_t^{\mathrm{Q}}(VIX_{T_1\to T_2})$ using the definition of variance:[20]

$$\begin{aligned} \mathrm{Var}_t^{\mathrm{Q}}(VIX_{T_1\to T_2}) &= \mathbb{E}_t^{\mathrm{Q}}\left(VIX^2_{T_1\to T_2}\right) - \left(\mathbb{E}_t^{\mathrm{Q}}(VIX_{T_1\to T_2})\right)^2 \\ \iff \mathbb{E}_t^{\mathrm{Q}}(VIX_{T_1\to T_2}) &= \sqrt{\mathbb{E}_t^{\mathrm{Q}}(VIX^2_{T_1\to T_2}) - \mathrm{Var}_t^{\mathrm{Q}}(VIX_{T_1\to T_2})}. \end{aligned} \tag{1.5}$$

The first term under the square root can be calculated using portfolios of S&P 500 Index options with maturities T_1 and $T_2 = T_1 + 30\ days$ that replicate $VIX^2_{t\to T_1}$ and $VIX^2_{t\to T_2}$, respectively:

$$\begin{aligned} (T_2 - T_1)\mathbb{E}_t^{\mathrm{Q}}(VIX^2_{T_1\to T_2}) &= (T_2 - t)(VIX^2_{t\to T_2}) - (T_1 - t)(VIX^2_{t\to T_1}) \\ \iff \mathbb{E}_t^{\mathrm{Q}}(VIX^2_{T_1\to T_2}) &= \frac{(T_2 - t)VIX^2_{t\to T_2} - (T_1 - t)VIX^2_{t\to T_1}}{T_2 - T_1}. \end{aligned} \tag{1.6}$$

Alternatively, one can use variance swap prices to estimate $\mathbb{E}_t^{\mathrm{Q}}(VIX^2_{T_1\to T_2})$. By using the definition of VIX as a measure of risk-neutral entropy of the S&P 500 Index return, it is straightforward to show that the result is also valid with jumps (the proof is in the Appendix, section 1.12.1). Full details about the empirical calculation of the synthetic VIX futures are in sections 1.12.1 – 1.12.4 of the Appendix.

The second term under the square root in equation (1.5) can be found using a static portfolio of out-of-the-money (OTM) VIX options and applying a result from Breeden and Litzenberger (1978) similar to Martin (2017) (derivation details are in the Appendix, section 1.12.2):

$$\mathrm{Var}_t^{\mathrm{Q}}(VIX_{T_1\to T_2}) = 2R_{f,t\to T_1}\left(\int_{K=0}^{F_{t,T_1}} put_{t,T_1}(K)dK + \int_{K=F_{t,T_1}}^{\infty} call_{t,T_1}(K)dK\right). \tag{1.7}$$

An important point is that the decompositions in (1.5), (1.6), and (1.7) rely on no parametric or distributional assumptions about the S&P 500 Index or VIX.

[Figure 1.5 about here]

Figure 1.5 shows the dynamics of the EFG for maturities at 1–4 months.[21] The graphs illustrate that the EFG for one month and two months was close to zero before 2009, but

[20] Similar to Carr and Wu (2006) and Dong (2016).

[21] Unfortunately, there are often no VIX options with expirations of five, seven, and eight months and it is thus impossible to calculate $\mathrm{Var}_t^{\mathrm{Q}}(VIX_{T_1\to T_2})$ without interpolating the volatility surface. The quality of the options data for six months and nine months is also often poor. For these reasons, the graphs show the EFG for 1–4 months.

fluctuated strongly after that and was mostly positive for two months. The first-month EFG is 13.2% per year, and the second-month EFG is 23.7% per year, on average, for 2007–2018. The weighted average EFG from one and two months futures contracts is 18.5% per year. The plots indicate that the introduction of ETFs, especially leveraged ETFs, is positively correlated with the increase in the short-term gap.

1.4.4.2 Possible explanations for the EFG

Several factors could explain the EFG. First, there is a discretization error when computing $\mathrm{Var}_t^{\mathrm{Q}}(VIX_{T_1 \to T_2})$ in equation (1.7), since a continuum of strikes is not observable in practice, and the integral is approximated with a sum. However, due to the convexity of call and put option prices, this error would bias the risk-neutral variance downwards, pushing the EFG even higher. Therefore, my calculations would underestimate the true gap. I perform several robustness checks to deal with truncation and discretization errors as described in section 1.12.3 of the Appendix.

The second possible explanation is that, in order to replicate $\mathbb{E}_t^{\mathrm{Q}}(VIX^2_{T_1 \to T_2})$ and $\mathrm{Var}_t^{\mathrm{Q}}(VIX_{T_1 \to T_2})$, it is necessary to trade deep OTM options with high bid-ask spreads. Lack of liquidity in those options and higher transaction costs could explain part of the gap. One way to solve this problem is to replace the portfolios of options replicating $VIX^2_{t \to T_1}$ and $VIX^2_{t \to T_2}$ with forward variance swap prices. Using variance swaps instead of option portfolios produces similar estimates of the EFG, even after controlling for jumps. Figure 1.16 in the Appendix displays the results using this approach. Moreover, if liquidity differences were the main reason for the gap, then the gap should have already been positive before the introduction of ETFs. However, the plots show that the gap was close to zero or even slightly negative before ETFs were introduced. Furthermore, if liquidity differences were the main reason for the gap, one would expect longer-dated and less liquid maturities to have more pronounced gaps. The dynamics of EFG for maturities at 3 and 4 months shows that this is not the case. To account for market liquidity, I include relative bid-ask spreads in the EFG regressions.

Another plausible explanation for the gap could be the difference in margin requirements between futures and options markets. The margin-based explanation alone would struggle to explain why the gap takes both positive and negative values, since margins are unlikely to be higher in the futures market than in the options market (typically, futures margins are much smaller). However, funding constraints could explain some of the time variation in the EFG. Shleifer and Vishny (1997) and Gromb and Vayanos (2002) show that when arbitrage capital is scarce, price gaps can exist and persist.

Investors may scale back positions during crisis times and might be reluctant to engage in arbitrage if fearful of undesired liquidation of positions at a loss, in situations where the price discrepancy widened. Garleanu and Pedersen (2011) argue that price discrepancies between two identical assets should depend on the shadow cost of capital. Based on these theories, I include the TED spread (similar to Barras and Malkhozov, 2016) as a control in the EFG regressions to account for funding liquidity constraints.

[Figure 1.6 about here]

The fourth possible reason for the gap is the presence of ETFs in the VIX futures market, but the lack of such funds in the market for options or variance swaps. As I show in section 1.11.3 in the Appendix, the daily rebalancing of ETFs can bias futures prices since ETF counterparties demand a premium for providing liquidity and trading in the opposite direction. Figure 1.6 shows that the EFG and the rebalancing demand from ETFs move together. Due to the mechanics of ETF rebalancing, prices in the futures market can be different from the same prices in the options/variance swaps market before the expiration of the futures. However, at maturity T_1, both $\mathbb{E}^{\mathbb{Q}}_{T_1}(VIX_{T_1 \to T_2})$ and F_{T_1,T_1} are equal to $VIX_{T_1 \to T_2}$.

1.4.4.3 Regressions of the EFG

Panel B of Table 1.4 presents the estimates from a regression of the EFG on ETF demand:

$$EFG_{t,i} = \alpha + \beta_1 D^{\$,all}_{t,i} + \beta_2 EFG_{t-1,i} + \gamma Ctrl_{t,i} + \epsilon_{t,i}. \tag{1.8}$$

I add lagged EFG to account for auto-correlation and TED spread to control for funding liquidity constraints. The results from columns 1 and 5 show that one standard deviation rise in ETF demand as a proportion of market capitalization is related to a contemporaneous increase in the EFG by 1.21% for the first month and by 0.57% for the second month. ETF demand explains 24–26% of the daily variation in the EFG. After controlling for several variables, the regression estimates decrease to 0.97% and 0.45%, respectively (columns 2 and 6), but are still strongly statistically significant. The effects are robust to using demand not scaled by open interest (columns 3 and 7), and are greater in magnitude. One standard deviation rise in the ETF demand calculated with lagged returns predicts 0.77% increase in the first-month gap and 0.38% in the second-month gap as seen from columns 4 and 8. The positive and statistically significant estimates show that the rebalancing of ETFs increases the non-fundamental component of prices. The EFG is related to liquidity changes: the first-month gap to market liquidity, whereas the second-month gap to funding liquidity. A standard deviation rise in relative bid-ask spreads increases the first-month gap by 0.88%, whereas a standard deviation rise in TED spread substantially raises the second-month gap by 1.46%. These results suggest that part of the gap could be due to arbitrageurs' inability to close positions easily in times of crisis when liquidity dries up, or could be due to funding constraints consistent with Garleanu and Pedersen (2011) and Gromb and Vayanos (2002).

Table 1.16 and Table 1.17 in the Appendix present some robustness checks. The results show that ETF demand has a greater effect in periods of high volatility, as seen from columns 1–2 and 4–5 of Table 1.16. The impact of ETF demand is robust to controlling for the returns on the Fama–French five factors and momentum. Columns 7–8 show that the gap was related to liquidity before the introduction of ETFs. Table 1.17 shows that the positive effect of ETF demand is robust to adding lagged demand. There is some evidence of reversals as shown by the negative estimates on the lagged demand from columns 3 and 4.

I also check whether current ETF futures gaps are related to realized futures returns. Table 1.18 in the Appendix shows that a higher EFG decreases realized returns: a 1% increase in the EFG decreases the realized return by 22 bps for the first-month futures contract and by 144 bps for the second.

1.4.4.4 EFG – fundamental or not?

The existence of the ETF futures gap is evidence that the risk-neutral measure imputed from option prices (Q) gives a different forecast of the realized spot price at maturity compared to the risk-neutral measure from the futures market. However, a priori, the mere presence of the gap does not mean that the futures price $F_{t,T}$ is a poor estimate of the fundamental value. It is entirely possible that price discovery takes place in the ETF-influenced market, and therefore, the gap exists because of fundamental information about the realized spot instead of non-fundamental price pressure in the futures market. However, if that was the case, the ETF-influenced futures would be a better predictor of the realized spot price at maturity. Testing this prediction is straightforward.

[Figure 1.7 about here]

The empirical evidence shows that the ETF-influenced futures contract is a poor predictor of realized spot changes, whereas the synthetic futures contract $\mathbb{E}_t^{\mathrm{Q}}(S_T)$ is a better predictor. Figure 1.7 shows the time series of the mean squared error (MSE) from a prediction of realized VIX based on $\mathbb{E}_t^{\mathrm{Q}}(S_T)$ and $F_{t,T}$. The plots show that the MSE of the synthetic futures contract is almost always below that of the traded futures contract. Table 1.19 in the Appendix shows that the basis of the synthetic futures contract is a better predictor of realized spot changes than the basis of the traded futures contract. The coefficient for the first-month futures of 0.25 is highly statistically significant compared to the insignificant estimate of 0.02 from Table 1.3. The estimate for the second futures of 0.55 is also highly statistically significant and larger compared to the insignificant coefficient of 0.27 from Table 1.3. These facts suggest that the EFG arises due to non-fundamental price pressure, consistent with the findings for basis from Mixon and Onur (2015). Compared with their paper, I test the price pressure hypothesis more thoroughly by estimating the non-fundamental part of basis and testing directly whether it reflects information about the fundamental value.

A potential concern is that, even though ETF demand does not directly influence the VIX options market, prices of options could be disrupted if investors hedge the futures exposure with options. To test this, I run regressions of basis $(\mathbb{E}_t^{\mathrm{Q}}(S_{T_1}) - S_t)$ and spread $(\mathbb{E}_t^{\mathrm{Q}}(S_{T_2}) - \mathbb{E}_t^{\mathrm{Q}}(S_{T_1}))$ calculated using the synthetic futures price, on ETF demand. If ETF price pressure affects the synthetic futures price (directly or indirectly), the coefficient on ETF demand should be statistically significant. The results from Table 1.20 in the Appendix show that this is not the case. This fact illustrates that ETF demand manifests itself through an increase in the non-fundamental component of prices (EFG) and has no impact on the more fundamental, synthetic futures contract.

Another concern is that maybe the EFG arises due to price pressure in the VIX options market after 2009, and ETF demand just happens to be correlated with this pressure. To address this concern, I add the hedging pressure from the options market as a control in the EFG regressions (the last two columns of Table 1.17). I measure the pressure in terms of delta-hedging. For each day, I calculate the delta-hedging demand as $-\sum_{j=1}^{M} \Delta_{t,j} OI_{t,j} F_{t,T_i}$, where M is the total number of options on the futures expiring at T_i, $\Delta_{t,j}$ is the Black-Scholes delta of option j, and $OI_{t,j}$ is the total open interest for option j. The minus in front of the expression captures the idea that if the total delta in the options market is positive, the hedging demand would be negative (agents would sell the underlying to hedge the positive-delta option position). I also calculate gamma-hedging demand to account for second-order effects. The estimates from columns 5 and 6 show that the positive and statistically significant effects of ETF demand are robust to including the measures of hedging pressure from the options market.

Anecdotal evidence from my discussions with several hedge fund traders suggests that hedge funds are replicating the VIX futures with a portfolio of options and trading the difference between the two. Figure 1.17 in the Appendix illustrates that leveraged money (mostly hedge funds) is generally short VIX futures after 2011, when EFG starts to be positive. Table 1.21 in the Appendix shows that a 100% increase in the EFG is correlated with a $30 million decrease in hedge funds' positions. In practice, however, hedge funds usually do not construct the exact replicating portfolio. They trade a subset of all options to avoid paying large bid-ask spreads in deep OTM options. This could explain why the ETF demand pressure does not significantly impact the synthetic futures. During turbulent times, the profit made on the hedge portfolio might not be enough to offset the losses from the futures market, and hedge funds could be forced to close down futures positions at a loss. The fact that the gap is strongly related to measures of funding and market liquidity confirms these facts. In section 1.11.3, I show that imperfect replication of the futures contract leads to a higher gap.

The existence of the EFG shows that the VIX futures market is segmented from the S&P 500 Index and VIX options markets. Bardgett et al. (2019) also illustrates that S&P 500 Index options and VIX derivatives have different information about volatility at different time horizons, which could be interpreted as evidence of segmentation between these markets.

1.5 Source of the EFG. Decomposition of ETF demand

To understand the source of the ETF impact on futures prices, I decompose the rebalancing demand from these funds into three major components: calendar rebalancing, leverage rebalancing and flow rebalancing.

1.5.1 Calendar rebalancing

Since futures have an expiration date, to maintain their exposure, VIX and commodity ETFs need to roll out of the maturing contracts before these contracts expire and initiate new positions in longer-maturity contracts. Calendar rebalancing is mechanical and arises exogenously due to futures expiration. Most ETFs in VIX and commodity markets are based on a benchmark that is rolling from the first-month futures contract to the second one over a period of several days. For VIX, the benchmark is a constant-maturity weighted average position: every day, a typical long ETF invests fraction α_t of its wealth in the first-month futures contract, and $1 - \alpha_t$ fraction in the second one, s.t. $\alpha_t T_1 + (1 - \alpha_t) T_2 \approx 21\ days$. T_1 is the time to maturity of the first-month futures contract in business days, T_2 is the time to maturity of the second one, and 21 is the typical number of business days in the rebalancing period (month). For example, suppose that today (t) $T_1 = 21\ days,\ T_2 = 42\ days$ and consider a long ETF: $\alpha_t = 1,\ 1 - \alpha_t = 0$. Tomorrow, both futures contracts are closer to maturity: $T_1 = 20\ days,\ T_2 = 41\ days$, so to keep the duration of the portfolio constant at (roughly) one month, the ETF allocates wealth as follows: $\alpha_{t+1} = \alpha_t - \frac{1}{21} = \frac{20}{21},\ 1 - \alpha_{t+1} = \frac{1}{21}$. After 21 business days, $T_1 = 0\ days,\ T_2 = 21\ days$, the long ETF has completely rolled out of the expiring contract and is 100% invested in the new one month contract, and then the cycle starts again. More information on the benchmark of VIX ETFs is available in S&P500 (2019).

[Figure 1.8 about here]

Calendar rebalancing of ETFs can be seen from the dynamics of open interest in the VIX futures market.[22] Before ETFs were introduced, the change in open interest did not have a clear pattern. The left panel of Figure 1.8 shows typical dynamics before the introduction of ETFs. However, in the post-ETF period, the change in open interest follows a typical pattern, as shown in the right panel. For example, consider the four-month VIX futures in November 2012. Open interest spikes as soon as it becomes a two-months futures contract in January 2013 and ETFs start to buy it. When it becomes a one-month futures contract in February 2013, ETFs start to sell it and open interest declines. The increase in the number of contracts for the two-months futures is roughly equal to the decrease in the number of contracts for the one-month futures. The hump-shaped dynamics can be well identified with net ETF positions.

Since ETFs passively follow the exact rolling rules of the indices they track in order to minimize the tracking error, they have a large hedging demand during the rolling period. If ETF counterparties have a limited capacity to absorb this demand, they would require a premium to meet ETF trades. That premium would be incorporated into futures prices.

1.5.2 Leverage rebalancing

Another important feature of futures-based ETFs in VIX and commodity markets is that many of them are leveraged, or inverse. This means that it is possible to estimate

[22]Calendar rebalancing is not perfectly predictable as the exact rebalancing amount depends on the assets under management (AUM) of the ETF, which in turn depend on the contemporaneous realized return as I show in section 1.5.3.

the impact of leverage-induced trading on prices. Leverage rebalancing is mechanical and arises exogenously due to the maintenance of a constant leverage at a high frequency. Leveraged ETFs use derivatives to provide multiples of the daily return on their benchmarks.[23] Inverse ETFs provide short market exposure for investors who otherwise could be constrained to short-sell. Subsequently, I refer to both leveraged and inverse ETFs as leveraged ETFs since inverse funds can be thought of as a special type of leveraged funds with a negative leverage. It is important to stress that the leverage L is fixed in the prospectus for each fund, and is constant over time. Another crucial point is that leveraged funds seek to deliver their objective on a *daily* basis, and the leverage over longer periods will be different from L due to daily compounding. A leveraged fund could therefore attract mostly short-horizon investors, as noted in a typical fund's prospectus (see, e.g., ProShares, 2019). The proportion of leveraged ETFs in commodity markets, and in VIX, have risen significantly over the last eight years. These ETFs are also quite common in equity and fixed income markets.

With a leverage of L, leveraged ETFs aim to return Lr_{t+1} every day, where r_{t+1} is the daily return on the benchmark from t to $t+1$.[24] AUM A_{t+1} at time $t+1$ should then be $A_{t+1} = A_t(1 + Lr_{t+1})$. An important feature of leveraged ETFs is that, in order to maintain a constant leverage, they always have to rebalance in the same direction as the benchmark. This is true for both long ($L > 0$) and inverse ($L < 0$) ETFs. The derivation is straightforward. At time t, the exposure of a leveraged ETF is LA_t. One period later, the actual exposure is $LA_t(1 + r_{t+1})$, whereas the desired exposure is $LA_{t+1} = LA_t(1 + Lr_{t+1})$. Hence, to maintain a constant leverage, the ETF has to rebalance by

$$\delta_{t+1} = LA_{t+1} - LA_t(1 + r_{t+1}) = L(L-1)A_t r_{t+1}. \tag{1.9}$$

For example, investing \$10 in a double-leveraged ($L = 2$) ETF means that the ETF buys \$20 worth of futures by borrowing another \$10. Suppose that the price goes up by 10%, then the futures position is worth \$22. To maintain the leverage constant at 2, the ETF has to borrow additional \$2 and use it to buy \$2($= L(L-1)A_t r_{t+1} = 2\cdot(2-1)\cdot\$10\cdot 10\%$) of futures contracts. This changes the leverage to $2 = 24/12$.

[23] Most leveraged ETFs use swaps to obtain a levered exposure. The swaps' exposure is transmitted to the futures market by the swap counterparties. Leverage rebalancing throughout the paper focuses on the rebalancing to maintain a constant leverage with respect to the benchmark but ignores the leverage implicit in futures positions for simplicity.

[24] Leveraged funds seek to deliver L multiplied by the daily performance of the benchmark index *before* fees and expenses. With fees and expenses, their effective leverage can be slightly different from L. The analysis is largely unchanged, however, as L can be replaced with $\hat{L} = L(1-\phi)$ where ϕ is the tracking error due to fees and expenses.

As $L(L-1) > 0$ for any leverage $L \notin [0,1]$, equation (1.9) shows that rebalancing demand is of the *same* sign as r_{t+1}. This means that trading demands by long and inverse ETFs do not offset, but instead reinforce each other. Leveraged ETFs could magnify price changes, creating a possible feedback channel for prices. To quantify the potential impact of leverage rebalancing of all N ETFs in a given market, I calculate the leverage rebalancing multiplier $\Gamma_t = \sum_{j=1}^{N} L_j(L_j-1)A_{j,t}$. The red line in Figure 1.1 shows Γ_t as a proportion of the whole market across several assets. This number is around 0.60 for the VIX market in February 2018, which means that if the benchmark spiked by 10%, 6% of the total market capitalization would be the additional buying demand from all ETFs due to leverage rebalancing. The benchmark increased by more than 90% towards market close on 5 February 2018, which means that more than 54% of the market was allocated to buying VIX futures contracts following the spike.

Leverage rebalancing is, essentially, a momentum trade, as it involves buying after price increase and selling after price decrease. Arbitrageurs who trade against ETFs are then contrarian and carry the risk of meeting ETF demand in case of large price changes. As I show in the next sections, arbitrageurs demand a premium for bearing this risk. If there are flows u_{t+1}, the total rebalancing demand by a leveraged ETF from t to $t+1$ becomes $\delta_{t+1} = L(L-1)A_t r_{t+1} + L u_{t+1}$.

1.5.3 Total ETF demand decomposition

To understand the motives of ETFs' trading, I calculate the total daily rebalancing demand by a futures-based ETF with a leverage of L ($L = 1$ corresponds to a non-leveraged ETF) during the rolling period of K days ($K = 21$ for VIX, $K = 5$ for most commodity markets). Full derivation details are in section 1.12.5 of the Appendix. In dollar terms, the total rebalancing demand is:

$$\begin{aligned} D^{\$}_{t+1,1} = & -\underbrace{\frac{L}{K}A_t(1+Lr_{t+1})}_{calendar\ rebalancing} + \underbrace{\alpha_t A_t L(L-1)r_{t+1}}_{leverage\ rebalancing} + \underbrace{(\alpha_t - \frac{1}{K})Lu_{t+1}}_{flow\ rebalancing} + \\ & + \underbrace{\alpha_t(1-\hat{\alpha}_t)LA_t(r^{F_2}_{t+1} - r^{F_1}_{t+1})}_{remainder}, \end{aligned} \tag{1.10}$$

where $r^{F_1}_{t+1}$, $r^{F_2}_{t+1}$ are the net returns on the first-month and the second-month futures contracts, respectively, $\hat{\alpha}_t = \frac{\alpha_t F_{t,T_1}}{\alpha_t F_{t,T_1} + (1-\alpha_t)F_{t,T_2}}$, and $r_{t+1} = \frac{\alpha_t F_{t+1,T_1} + (1-\alpha_t)F_{t+1,T_2}}{\alpha_t F_{t,T_1} + (1-\alpha_t)F_{t,T_2}} - 1 = \hat{\alpha}_t r^{F_1}_{t+1} + (1-\hat{\alpha}_t) r^{F_2}_{t+1}$ is the net return on the benchmark. Equation (1.10) illustrates that the total rebalancing demand can be decomposed into four components: calendar rebalancing due to the roll from the first-month to the second-month futures contract, leverage rebalancing to maintain a constant leverage L, flow rebalancing due to inflows or outflows, and a remainder. Analogously, the total dollar rebalancing demand for the second-month futures contract is:

$$\begin{aligned} D^{\$}_{t+1,2} = & \underbrace{\frac{L}{K}A_t(1+Lr_{t+1})}_{calendar\ rebalancing} + \underbrace{(1-\alpha_t)A_tL(L-1)r_{t+1}}_{leverage\ rebalancing} + \underbrace{(1-\alpha_t+\frac{1}{K})Lu_{t+1}}_{flow\ rebalancing} - \\ & - \underbrace{\hat{\alpha}_t(1-\alpha_t)LA_t(r^{F_2}_{t+1}-r^{F_1}_{t+1})}_{remainder}. \end{aligned} \tag{1.11}$$

The total dollar rebalancing of all N ETFs in a given market is: $D^{\$,\ all}_{t+1,1} = \sum_{j=1}^{N} D^{\$,\ j}_{t+1,1}$, $D^{\$,\ all}_{t+1,2} = \sum_{j=1}^{N} D^{\$,\ j}_{t+1,2}$.

Calendar rebalancing is exactly the opposite for the first-month and the second-month futures contracts. In a market where ETFs are net buyers of futures ($\sum_{j=1}^{N} L_j A_{j,t} > 0$), calendar rebalancing decreases $D^{\$,all}_{t+1,1}$ and increases $D^{\$,all}_{t+1,2}$ (except for extreme realizations of r_{t+1}). Leverage rebalancing is always in the same direction as the realized return on the benchmark. Inflows (rise in flow rebalancing) increase both $D^{\$,all}_{t+1,1}$ and $D^{\$,all}_{t+1,2}$, whereas outflows decrease both of them. The effect of the remainder is due to the fact that the ETF benchmark is a weighted average of the first-month and the second-month futures contracts. Therefore, the return on the second-month futures contract can have an impact on prices for the first (and vice versa) through ETF demand.

On average, in the VIX market, the largest component of rebalancing demand is calendar rebalancing. Flow rebalancing is also large, and sometimes exceeds 75% of the total rebalancing demand from VIX ETFs, as seen from Figure 1.9. Leverage rebalancing has been growing since 2012, and represented more than 40% of total demand at the start of 2018. The remainder has been historically low (less than 5%).

[Figure 1.9 about here]

There are several pieces of evidence that ETFs (irrespective of their legal structure as a fund or note) follow their benchmarks and rebalance in the way described in this section. First, anecdotal evidence from my discussions with several ETF managers and authorized participants suggests that ETFs have no incentive to deviate from the benchmark, as their performance is evaluated based on the tracking error. The compensation for ETF sponsors arises from fees but not from over-performance or under-performance. Second, some of the ETFs make their daily holdings publicly observable. The change in holdings actually seen matches the one predicted by the rebalancing from equations (1.10) and (1.11). Third, the implied weekly net positions of ETFs closely follow the reports from the CFTC, as seen in Figure 1.3, although the match is not perfect because the CFTC data is weekly and the holdings are aggregated for ETFs and other dealers.[25] Eraker and Wu (2017) also conclude that VIX ETFs track their benchmark indices fairly well at a daily frequency.

The decomposition of ETF demand developed in this section is flexible, and can accommodate various types of ETFs. It can accommodate non-leveraged ETF demand by setting $L = 1$ in equations (1.10) and (1.11): for these ETFs, the leverage rebalancing demand vanishes. The decomposition is not a feature of VIX and commodity ETFs, but can be used to analyze the impact of ETF demand in other asset classes. All ETFs have to rebalance due to investor flows and hence, flow rebalancing is present in ETFs across asset classes. The same holds for leverage rebalancing because leveraged ETFs are present in equity, fixed income, and foreign exchange markets, albeit with a smaller proportion. Calendar rebalancing also has a close analogue in equity and fixed income markets. In VIX and commodity markets, this type of demand arises because futures contracts expire and ETFs have to substitute the positions with new contracts. Analogously, equity ETFs have to rebalance in case of inclusions or exclusions of stocks in the benchmark index. Fixed income ETFs also have to rebalance in a similar way when underlying bonds expire, or when there is a change in the benchmark index due to the inclusion or exclusion of bonds. Thus, the effect of different types of ETF rebalancing can also be studied in other markets. However, it is more beneficial to analyze the impact of calendar rebalancing in futures-based ETFs, since we can observe the effects at a daily frequency, than equity or bond ETFs, where this type of rebalancing is much less frequent.

1.5.4 ETFs can affect prices even in a market with a zero net share of ETFs

Equations (1.10) and (1.11) illustrate that the composition of the market (the proportion of ETFs with different leverages) matters in determining the total rebalancing demand and, as a result, the ETF impact on futures prices. For example, consider a market where the size of all long ETFs is exactly equal to the size of all inverse ETFs so that the net share of ETFs is zero ($\sum_{j=1}^{N} L_j A_{j,t} = 0$). However, the net ETF demand in that case will not be zero, even excluding flows. In such a market, there is no

[25] I describe several robustness checks for the computation of leverage rebalancing in section 1.12.6 in the Appendix.

remainder or predictable part of calendar rebalancing ($\frac{L}{K}A_t$). The only sources of rebalancing are leverage rebalancing and the leverage-induced part of calendar rebalancing ($\frac{1}{K}A_tL^2r_{t+1}$), both of which can be quite large despite the equal size of all the long and inverse ETFs. Moreover, since leverage rebalancing is always in the same direction as the benchmark return, even in an equal-sized market the potential amplification of price changes can be substantial.

This observation is in contrast to the ordinary view that ETFs have no price impact if the size of long ETFs is exactly equal to that of inverse ETFs. In fact, providing liquidity in such a market should be compensated by a large risk premium because the potential distorting effects of leverage rebalancing are substantial. For example, a market with \$100 of $L = 1$ ETFs is exactly the same in net demand terms to a market with \$100 in $L = 2$ ETFs (with a total exposure of $2 \cdot \$100$) and \$100 in $L = -1$ ETFs. However, the potential leverage rebalancing in the first market is zero, whereas in the second market, it is *four* times the size of the market ($\$400 = 2 \cdot (2-1) \cdot \$100 + (-1) \cdot (-1-1) \cdot \$100$) multiplied with the realized return on the benchmark. A 10% spike in the benchmark has no feedback effects in the first market, but leads to an additional buying pressure of 40% ($4 \cdot 10\%$) of the whole market size due to mechanical leverage rebalancing in the second market. A prominent real-world example of these effects was the VIX market in the beginning of February 2018. The net share of ETFs then was close to zero, but the potential distorting effect due to leverage rebalancing was 60% of the total market (as shown in Figure 1.1).

1.5.5 Risks posed by ETF demand

Consider arbitrageurs who trade against ETFs in the futures market. If arbitrageurs are competitive and could hedge perfectly (as in a standard Black-Scholes economy), ETF demand pressure would have no effect. However, in practice, arbitrageurs cannot do that as they face incomplete markets because of discrete trading, transaction costs, jumps in the underlying, and other factors (see, e.g., Garleanu et al., 2009). If arbitrageurs cannot perfectly hedge the ETF exposure, they bear non-fundamental risk of ETF demand shocks on three main fronts.

The first and most important one is leverage rebalancing. This is a relatively new and under-researched type of rebalancing by institutional investors with a large market share. An important observation is that arbitrageurs cannot hedge the leverage rebalancing of ETFs by matching long and inverse ETF demands, since the two are of the same sign as r_t (since $L(L-1) > 0$). In section 1.6, I show that leverage rebalancing introduces a source of convexity that is not easy to hedge similar to Garleanu et al. (2009) and hence, exposes arbitrageurs to variance. Intuitively, hedging the exposure would require frequent trading in a rolling position of one-month and two-months options on futures, and rebalancing the position on a daily basis (and even more frequently around market close). Moreover, as I show in section 1.6, in turbulent times the hedge portfolio would still be imprecise since ETF tracking errors are magnified in those times. Thus, trading against leverage rebalancing exposes arbitrageurs to unhedgeable risks due to the impossibility of trading continuously in the benchmark, transaction costs, and tracking errors. Arbitrageurs would require premium for bearing these risks. Leverage rebalancing amplifies price changes by moving prices in the direction of benchmark returns (r_t): positive returns increase F_{t,T_1} and F_{t,T_2}, whereas negative returns decrease both futures prices.

The second area of risk for arbitrageurs is calendar rebalancing. This type of demand depends on realized returns and is not perfectly predictable. Maintaining a constant leverage by leveraged ETFs impacts also calendar rebalancing. The non-linear response of calendar rebalancing arises because leveraged ETFs track benchmark returns multiplied with the respective leverage (the term Lr_{t+1} in the brackets for calendar rebalancing from equations (1.10) and (1.11)). Hence, calendar rebalancing inherits the non-linearity of leverage rebalancing (in L, and, in continuous time, in the realized return r_{t+1}). Thus, part of this rebalancing could also be hard to hedge.

Another feature of calendar rebalancing is that arbitrageurs mechanically bear the risk of widening price discrepancies and cannot wait until expiration. By trading against the calendar demand from ETFs, arbitrageurs would typically sell the two-months futures contract and then buy it back from ETFs once the contract becomes a one-month futures contract. The right graph in Figure 1.8 illustrates that usually the increase in open interest for the second-month contract is similar in size to the decrease in open interest for the first-month contract. This observation shows that the new contract positions initiated by ETFs when the futures has maturity of two months, are closed before expiration, once the futures has maturity of one month. This fact suggests that ETF counterparties also close the futures position before maturity and bear the risk of widening price gaps. Section 1.11.3 shows that, with a short horizon, arbitrageurs would require a premium for bearing this risk. Calendar rebalancing would push up F_{t,T_2} and push down F_{t,T_1} over time. Hence, it would decrease the front-month basis and increase the spread between the first-month and the second-month futures contracts. Some papers have studied the effects of calendar rebalancing on futures prices in commodity markets (see, e.g., Tang and Xiong, 2012; Mou, 2011).

The third area of risk for arbitrageurs is flow rebalancing. The effects of this rebalancing could be pronounced if inflows happen at times when arbitrageurs are more constrained. For example, in the case of VIX ETFs, the underlying asset value increases at times of high marginal utility (when the market crashes). Therefore, inflows to VIX ETFs when VIX spikes would require arbitrageurs to short-sell VIX futures at a time when financial constraints could be binding. Risk-averse investors would require a premium for increasing their short VIX positions at such times. The empirical evidence of fund-by-fund flows suggests that flows are related to benchmark returns, but the coefficient is significant at the 13% level only. However, in extreme times, the relationship is more pronounced. These results are excluded from the paper for brevity.

Flow rebalancing could also have indirect effects through calendar rebalancing. For example, inflows increase ETFs AUM and raise the amount of calendar rebalancing that ETFs perform in future periods. Prices could react in anticipation of these effects. Flow rebalancing would push both F_{t,T_1} and F_{t,T_2} in the direction of flows. Inflows would increase both prices, whereas outflows would decrease them. The impact of flows on prices has been studied before in the mutual fund literature (see, e.g., Lou, 2012; Warther, 1995). As α_t decreases, ceteris paribus, the impact of leverage and flow rebalancing on the second-month futures contract strengthens, whereas the impact on the first weakens. The price impact of the three major types of ETF rebalancing is illustrated in Figure 1.9.

Short-term price impact can translate into longer-term price deviations (futures premium) through at least two channels. First, leverage rebalancing and flows could make arbitrageurs' financial constraints binding. For example, a temporary spike in price due to leverage rebalancing could trigger financial constraints if arbitrageurs have a short position from the previous period. I illustrate this idea in section 1.11.3.1 in the Appendix. In anticipation of this risk, prices can deviate for several periods. Second, both leverage rebalancing and flows ultimately end up as parts of calendar rebalancing since they change the AUM of the ETF and these AUM end up rolling from the first-month to the second-month futures contract. Calendar rebalancing is a lower-frequency component of price impact. Moreover, it mechanically introduces short-termism of arbitrageurs since they cannot wait until expiration as explained above. This short-termism can lead to long-term price deviations as shown in Shleifer and Vishny (1997), and Gromb and Vayanos (2002).

1.5.6 Empirical evidence on the impact of demand components

The estimates from columns 1 and 2 of Table 1.5 show the impact of the three major components of ETF demand on basis and spread. Flows have the highest impact: one standard deviation rise in flows (2.34% for first-month futures contract and 4.12% for the second) increases the basis by 1.19% and the spread by 0.68%. The fact that basis and spread are most strongly related to flows could be because the effect of flows propagates to calendar rebalancing in next periods. Calendar rebalancing has a negative impact on the basis but a positive impact on the spread, which is consistent with the analysis in section 1.5.5. Leverage rebalancing also has a positive impact on prices: one standard deviation rise (2.92% for first-month futures contract and 4.36% for second) increases the basis by 0.41%, and the spread by 0.27%. The signs are as predicted in section 1.5.5.

[Table 1.5 about here]

The estimates from columns 3 and 4 of Table 1.5 show that leverage rebalancing has the largest impact on the EFG in the VIX market: one standard deviation rise is related to an increase of 1.17% in the first-month EFG and 0.40% in the second-month EFG. Calendar rebalancing has a negative impact on the front-month gap and a positive impact on the second. Flow rebalancing has a positive impact on both maturities: one standard deviation rise is related to a 0.13% higher first-month EFG and a 0.34% higher second-month EFG. Given that the EFG is mostly related to leverage rebalancing by ETFs, it is important to understand why market participants require compensation for trading against this type of ETF demand. In the next section, I analyze the risks of leverage rebalancing by ETFs.

1.6 The risk premium of leverage rebalancing

Research shows that leverage-induced trading by retail investors can exacerbate market crashes and push prices away from fundamentals (see, e.g., Bian et al., 2017, 2018). The impact of leverage-induced trading by institutional investors on systemic risk and prices is an important, yet under-researched question. A key constraint of the existing literature is that leverage ratios of institutional investors like mutual funds or hedge funds are rarely publicly observable. Leveraged ETFs provide a useful laboratory to study the effects of leverage-induced trading on a daily basis since these ETFs have a pre-specified constant leverage.

To isolate the effect of leverage rebalancing, consider a market where the size of all long ETFs equals that of all inverse ETFs ($\sum_{j=1}^{N} L_j A_{t,j} = 0$) and there are no flows. In this case, remainder, flow rebalancing, and the predictable part of calendar rebalancing all vanish from equations (1.10) and (1.11), and the only sources of rebalancing are leverage rebalancing and the leverage-induced part of calendar rebalancing. An example of such a setting could be the ETF rebalancing during the trading day. To understand the risks of trading against leverage rebalancing of ETFs in a given market, I construct a portfolio that shorts ETFs with opposite leverages (to approximate trading against ETFs), and I rebalance it at the end of each day to maintain a zero-delta position (to approximate zero net ETF demand $\sum_{j=1}^{N} L_j A_{t,j} = 0$). To study the risks of such a trading strategy, I first derive the dynamics of the AUM for a purely leveraged ETF, that is, an ETF that has no calendar rebalancing, and does not face any flows.

1.6.1 The dynamics of leveraged ETFs' AUM

Consider a simple model with the futures-based benchmark ($\frac{dF_t}{F_t}$) following a geometric Brownian motion (GBM): $\frac{dF_t}{F_t} = \mu dt + \sigma dW_t$, where μ is the instantaneous drift, σ is the instantaneous diffusion, and W_t is a standard Brownian motion. To maintain a constant leverage of L, the ETF is trading continuously:[26] the fund invests a constant fraction of wealth L in the benchmark, and the rest $(1 - L)$ at the constant risk-free rate r_f. I allow for price impact (or tracking errors) and denote the effective leverage by $\hat{L} = L(1 - \phi)$ where $\phi \geq 0$ is the price impact coefficient and $|\hat{L}| \leq |L|$. Section 1.12.7 in the Appendix presents all the derivations. For simplicity, let $r_f = 0$ and let the fee rate of the fund $f = 0$. Then the AUM of a leveraged ETF are:

$$A_T = A_0(\frac{F_T}{F_0})^{\hat{L}} e^{-\frac{\hat{L}(\hat{L}-1)}{2}\sigma^2 T} = A_0(1 + r_T)^{\hat{L}} e^{-\frac{\hat{L}(\hat{L}-1)}{2}\sigma^2 T}, \tag{1.12}$$

As $-\hat{L}(\hat{L} - 1) < 0$ for any leverage $\hat{L} \notin [0, 1]$, both inverse and long leveraged funds would decrease in value if σ^2 is high. The AUM of inverse ETFs will be discounted more than those of long-leveraged ETFs in cases of high σ^2, since $-\hat{L}(-\hat{L} - 1) > \hat{L}(\hat{L} - 1), \forall \hat{L} > 0$. The term $-\frac{\hat{L}(\hat{L}-1)}{2}\sigma^2 T$ is the continuous-time counterpart of the leverage rebalancing in discrete time from equation (1.9).

[26] Empirically, ETFs have to deliver L times the daily return on the benchmark and to minimize tracking errors, leveraged ETFs would ideally like to trade at the closing price of each day. Indeed, most of the trading is concentrated around market closure (see, e.g., Pagano et al., 2019). Usually, the ETF transmits the size of the rebalancing order to the broker dealer 15 minutes before market closure based on the daily return up to that point, and then submits smaller orders at or very close to market closure. The term "continuous", therefore, could be a more accurate description of trading before market close rather than throughout the whole day.

Equation (1.12) shows that leverage rebalancing exposes the ETF to variance and to squared realized returns: $\frac{\partial A_T}{\partial \sigma^2} < 0$ and $\frac{\partial^2 A_T}{\partial r_T^2} > 0$ (for low σ^2).[27] Previous papers on leveraged ETFs have mostly ignored the exposure to squared realized returns, and emphasized that leveraged ETFs are negatively exposed to variance (see, e.g., Cheng and Madhavan, 2009). In practice, however, the exposure is not monotonic in variance. In other words, with continuous rebalancing, leveraged ETFs suffer from high σ^2 but benefit from high r_T^2 during the day: similar to a long-gamma, short-vega options position. The same model can be applied with dt being one day. In that case, leveraged ETFs have a positive exposure to r_T^2 over T days, but a negative exposure to daily σ^2.

1.6.2 Trading against leveraged ETFs

After deriving the dynamics of the AUM for a leveraged ETF, I next analyze the consequences for arbitrageurs who trade against leveraged ETFs.

1.6.2.1 Intuition

Consider an arbitrageur who provides liquidity to a pair of equal-sized ETFs with opposite leverages (e.g., $L = 2$ and $L = -2$) by taking the contrarian position to both ETFs. A natural guess could be that the arbitrageur can perfectly hedge the position by matching the demand from the long fund with that from the inverse fund. Surprisingly, however, such a trading strategy would not deliver the desired hedge, since both the long and the short ETF have a non-linear exposure to the underlying asset: both are implicitly long-gamma, short-vega. As a result, the arbitrageur would acquire a short-gamma, long-vega exposure by trading against a pair of opposite ETFs.

[Figure 1.10 about here]

In other words, shorting a pair of opposite ETFs is not a zero-return strategy, but an implicit bet on variance. If the underlying asset stays still and ends close to the initial value during the trading period, arbitrageurs collect the vega gains. However, if the underlying asset drifts steadily in either direction with little volatility, arbitrageurs could lose money due to the negative exposure to squared realized returns. By trading against the leverage rebalancing of ETFs, arbitrageurs acquire a short position if the price increases, and a long position if the price decreases. Figure 1.18 in the Appendix

[27] For example, with $\hat{L} = 2$ (double-leveraged ETF with no price impact): $\frac{\partial A_T}{\partial \sigma^2} = -TA_0(\frac{F_T}{F_0})^2 e^{-\sigma^2 T} < 0$ and $\frac{\partial^2 A_T}{\partial r_T^2} = 2A_0 e^{-\sigma^2 T} > 0$. The results are similar for other leverages.

illustrates the intuition with a simple example. If the price drifts steadily up, arbitrageurs lose money since they sell the asset and the price keeps increasing. If the price reverts back to the initial value, they make a profit since they sell the asset and the price decreases. Figure 1.10 illustrates the idea on a simple binomial tree. Hedging the variance exposure is presumably not so straightforward, since arbitrageurs would have to trade some portfolio with a long-gamma, short-vega exposure over the course of the trading period.[28] It would not be cost-efficient to hold such a hedging portfolio. Therefore, arbitrageurs would require a premium to bear the variance risks of trading against leveraged ETFs.

1.6.2.2 Theoretical exposure

To proxy for the premium that arbitrageurs earn by trading against ETFs' leverage rebalancing, I construct a portfolio that is short a pair of opposite ETFs. At the beginning of each trading day, I short sell the two ETFs, and then buy them back at the end of the day. The return on the portfolio is similar to the return for an arbitrageur who trades against opposite ETFs during the day. The idea is similar to Nagel (2012) who approximates market makers' liquidity provision with a strategy that sells stocks that outperformed the market and buys stocks that underperformed during the day. Let A be the price (using net asset value (NAV) instead of price gives similar results) of the long-leveraged ETF (leverage L), B be the price of the inverse-leveraged ETF (leverage $-L$). With discrete trading, the return ($r_{SB,T}$) of the strategy that short sells an equal amount of both A and B (sells at prices A_0, B_0 today and buys them back at time T in the future) assuming fully collateralized borrowing is:

$$r_{SB,T} = \frac{A_0 - A_T}{A_0} + \frac{B_0 - B_T}{B_0} = 2 - \left(e^{(\hat{L}\mu - \frac{\hat{L}^2\sigma^2}{2})T + \hat{L}\sigma W_T} + e^{(-\hat{L}\mu - \frac{\hat{L}^2\sigma^2}{2})T - \hat{L}\sigma W_T} \right) \tag{1.13}$$

The expression in the brackets is a sum of two log-normal random variables. The moments of the return distribution are:

$$\begin{aligned}
&\mathbb{E}(r_{SB,T}) = 2 - (e^{\hat{L}\mu T} + e^{-\hat{L}\mu T}) \\
&\mathrm{Var}(r_{SB,T}) \approx \hat{L}^2\sigma^2 T(e^{2\hat{L}\mu T} + e^{-2\hat{L}\mu T} - 2) \\
&Median(r_{SB,T}) = 2 - e^{\frac{-\hat{L}^2\sigma^2 T}{2}}(e^{\hat{L}\mu T} + e^{-\hat{L}\mu T}) \geq \mathbb{E}(r_{SB,T})
\end{aligned} \tag{1.14}$$

[28] For example, in case of intra-day trading, they would have to trade calendar spreads with short-term options expiring every few minutes and longer-term options expiring every day.

The distribution is negatively-skewed. If there was no price impact ($\hat{L} = L$), or the price impact parameter was the same for A and B, $\mathbb{E}(r_{SB,T}) \leq 0$, as $e^x + e^{-x} \geq 2\ \forall x$. Figure 1.20 shows the theoretical distribution for different μ and σ with daily frequency ($T = \frac{1}{252}$). As time T increases, the moments get more extreme: $\lim_{T\to\infty} \mathbb{E}(r_{SB,T}) = -\infty$, $\lim_{T\to\infty} \operatorname{Var}(r_{SB,T}) = \infty$. Hence, if there is no price impact, for longer holding periods the strategy leads to large losses, on average, but with very high variance.

1.6.2.3 Empirical evidence

Contrary to the theoretical prediction in the no-price-impact setup, empirically, the distribution is not only positively-skewed in most markets, but also with a positive mean. The top left panel of Figure 1.11 shows the returns on the strategy for an arbitrageur who sells both an $L = 1$ VIX ETF and an $L = -1$ VIX ETF.[29] The strategy delivers annualized returns of 21% with a Sharpe ratio of 0.89. This is not a feature of the VIX market but a general fact also for other assets. The returns for commodity markets are consistently positive as shown in Figure 1.11: 43% in the natural gas market, and 18% in the silver market, with Sharpe ratios of 2.59 and 1.86, respectively. The annualized alphas with respect to a benchmark of the Fama–French five factors, momentum, variance of the ETF benchmark, and the spot return are 16.6% (6.1 bps per day) for VIX, 42.3% (14 bps per day) for gas, and 18.6% (6.8 bps per day) for silver as shown in Table 1.22 in the Appendix.[30] Constructing a short-both strategy that mimics arbitrageurs' liquidity provision delivers high returns and large Sharpe ratios not only for VIX and commodity ETFs, but also for equity, fixed income, and foreign exchange ETFs as illustrated in Figure 1.21, Figure 1.22 and Figure 1.23 in the Appendix.

[Figure 1.11 about here]

A more puzzling observation is that the returns of the strategy often jump *up* instead of down in case of extreme market movements. For example for VIX, the largest return was realized during the VIX spike on 5 February 2018, when the ETF benchmark increased by more than 80% on a single day, yet the strategy obtained 51% *positive* return. These facts are inconsistent with what one could expect in a simple GBM framework with no price impact.

[29] There is no ETF with $L = -2$, so I take the pair with the largest possible opposite leverage.

[30] I also tried other benchmarks reflecting market-specific factors for each asset and intermediary-based risk factors (see, e.g., He et al., 2017). The alphas were still positive and statistically significant.

1.6.2.4 Explanation

What happens if I relax the GBM assumption? Even if the drift and the volatility are time-varying, the strategy would still have a negative expected return (as $2-e^{\hat{L}\int_0^T \mu_s ds}-e^{-\hat{L}\int_0^T \mu_s ds} < 0$). In case of a general distribution of the benchmark, the short-both strategy would still be negatively exposed to higher-order even cumulants as seen from equation (1.15):

$$-\log(\mathbb{E}(\frac{A_T}{A_0})) - \log(\mathbb{E}(\frac{B_T}{B_0})) = -c(1) - c(-1) = -2\sum_{n\ even}\frac{\kappa_n}{n!}, \tag{1.15}$$

where c and κ_n are the cumulant-generating function (see, e.g., Martin, 2013) and the n-th cumulant of the log AUM at time 0, respectively. Hence, shorting ETFs would still be negatively exposed to large price movements, if ETFs track the benchmark perfectly.

I now relax the assumption of no price impact (or no tracking errors) and allow the funds to have different price impact parameters: $\hat{L}_A = (1-\phi_A)L$, $\hat{L}_B = -(1-\phi_B)L$. As shown in section 1.11.3 in the Appendix, price impact will be the largest when arbitrageurs' risk aversion and variance of the benchmark increase, and leveraged ETFs could have difficulties keeping their leverage constant at L. In that situation, $\mathbb{E}(r_{SB,T})$ can be positive as both long and short ETFs cannot track the benchmark perfectly and they will not have exactly opposite leverages. Figure 1.19 in the Appendix shows that, for example, for VIX, ETFs with opposite leverages do not behave in exactly opposite ways to each other.

Running regressions of daily returns on leveraged and inverse ETFs, against the returns of their respective benchmarks, reveals that most funds underperform the stated return even with daily data (after controlling for fees), as shown in Table 1.6. For example, the regression coefficient for the double-leveraged VIX ETF is 1.93<2. A similar picture is observed for other markets. The coefficients lend support to the hypothesis that $|\hat{L}| < |L|$ and $\phi > 0$. In general, the tracking errors are larger in periods of high benchmark variance as shown in Table 1.6.

[Table 1.6 about here]

Because ETFs fail to track the benchmark perfectly, especially in turbulent times, the jump risk due to high r_T^2 may not be realized. Since the tracking errors get magnified in such times, arbitrageurs benefit, as they are, essentially, short ETFs (they trade in the opposite direction). As empirical evidence shows, the ETF arbitrage mechanism is also vulnerable to breakdowns in times of severe market stress. In other words, the tracking errors of ETFs co-move with bad states of the world. Providing liquidity to ETFs in such market conditions would be compensated by a high premium. The short-both strategy developed in this paper provides a simple and intuitive way to capture these effects by simply selling opposite ETFs. The high Sharpe ratios of the strategy are consistent with the analysis in section 1.5.5 as arbitrageurs face significant unhedgeable risk by trading against leveraged ETFs and are compensated for bearing it.

A potential reason for the spike in strategy returns during turbulent times could be the increased default risk of the ETF sponsor. Many of the leveraged ETFs are structured in the form of an ETN, which is an unsecured debt security issued by the ETN sponsor. Since ETNs are backed by the credit of the issuer, the issuer could have problems maintaining the leverage ratio due to financial distress. To test this prediction, I include CDS spreads of the ETF sponsors in the regression. A limitation of the data is that CDS spreads are available for a small subset of the sponsors. In particular, they are not available for most commodity and equity ETF issuers (ProShares and Direxion) so I run the regression only for the VIX market. The results from column 8 of Table 1.22 in the Appendix show that there is limited evidence that CDS spreads are related to the short-both strategy returns. The coefficients for current and lagged CDS spreads are insignificant at the 10% level. This observation is consistent with the fact that leveraged ETNs in a particular market are a usually a small share of the total debt of the sponsors, and are perhaps unlikely to cause substantial financial distress.

1.6.2.5 Regressions of the short-both strategy

Table 1.7 shows the results of a pooled regression of the short-both strategy returns from 14 markets including some equity and bond indices. The estimates show that the returns are related to the leverage rebalancing multiplier (Γ) in a given market. One standard deviation rise in Γ is correlated with a 7 bps rise (19.3% annualized) in daily returns for the strategy. The intuition for this finding is that, in markets where ETFs are large investors, in situations with large volatility spikes, all leveraged funds have to trade in the same direction, thereby amplifying market movements. The premium for providing liquidity to ETFs is greater in markets where the total demand for liquidity is higher, that is, in markets with a large proportion of leveraged ETFs.

[Table 1.7 about here]

Higher intra-day variance of the benchmark has a positive, but lower-magnitude, effect on returns (5–6 bps). The intuition for this finding is that larger σ^2_{bmk} means that there are more vega gains to be made. The estimate on r^2_{bmk} is negative and significant: one standard deviation rise in r^2_{bmk} decreases returns on the strategy by 2 bps. This is due to negative exposure to squared realized returns. The coefficient on the interaction term between Γ and σ^2_{bmk} is positive, which suggests that the returns on the short-both strategy are larger when variance is high and the potential leverage rebalancing is substantial. The strategy is also related to liquidity: one standard deviation rise in relative bid-ask spreads is related to a 2 bps increase in daily returns on the strategy. Returns are not related to the difference between NAV and ETF price (Premium to NAV), flows, the Fama–French five factors, or momentum.

The analysis in this section shows that the EFG is most sensitive to leverage rebalancing due to the unhedgeable risks faced by arbitrageurs who trade against this type of ETF demand. However, the puzzling observation is that the returns of such a trading strategy are large and are not arbitraged away despite the fact that the jump risk is often in the favorable direction. Instead of losing money in case of jumps, the strategy often delivers its largest gains in those turbulent times due to ETF tracking errors. A prominent example of such situation is the spike in the short-both strategy profits in the case of the VIX market on 5–9 February 2018. To the best of my knowledge, my paper is the first to quantify the risks of trading against leveraged ETFs and the first to show that their tracking errors co-move with extreme market times, thereby benefiting liquidity providers. Explaining the risk-return characteristics of the short-both strategy across markets is a beneficial direction of further research on ETFs.

1.7 The impact of ETFs on futures prices in commodity markets

I next analyze ETFs in several commodity markets with a high proportion of ETFs. I study whether the presence of ETFs in the markets for natural gas, silver, oil, and gold has an impact on futures prices. Similar to the approach for the VIX market in section 1.4, I start the empirical analysis for commodity ETFs by studying the effects of ETF demand on relative basis and spread. Compared with the VIX market, in commodity markets it is harder to construct a synthetic futures contract with exactly the same price at expiration as the traded one. I control for asset-specific fundamental shocks by including in the regression the closest futures contract with no ETFs traded for each commodity. For US natural gas futures traded on the New York Mercantile Exchange (NYMEX), I use Intercontinental Exchange (ICE) gas futures (historical correlation 81% before the ETF introduction). For silver, there was only one liquid futures contract specification before 2011,[31] so I use the closest precious metal: gold (historical correlation 93% before the ETF introduction). For US crude oil, I use Brent oil (historical correlation 99.7% before the ETF introduction). For gold, I use silver futures. I align contracts so that they are expressed in the same units and account for differences in expiration dates.

I use relative basis as the dependent variable, instead of the difference between the ETF-influenced futures price and that of the control contract, due to the concern that some systematic factors have changed the pricing across US and European markets in the post-ETF period. In particular, gas prices in the US have fallen substantially after the increase in shale gas drilling from 2010, and the difference with prices in Europe has widened. US crude oil prices have also diverged from Brent during the period 2010–2013 due to local supply factors.[32] Thus, using absolute levels of futures prices of the control asset to isolate the impact of ETFs could capture other changes in cross-market factors. However, using relative basis and assuming that storage costs have changed in a similar way for both the ETF-influenced and the control contracts (which is a within-market factor) is a more realistic assumption, as anecdotal evidence suggests. Moreover, to account for the above-mentioned changes in cross-market factors, I perform the following analysis. First, I control for the difference in spot prices between US and Europe in the respective basis and spread regressions. This difference captures the systematic change across the two markets, which is not influenced by ETFs as ETFs trade in the futures contract. Second, I construct a synthetic futures contract and estimate the impact of ETF demand in a narrow window that excludes the systematic changes across the two markets as explained in section 1.7.4.

1.7.1 Basis and spread regressions in commodity markets

The results of regression (1.3) for each commodity market are presented in Panel A of Table 1.8. The estimates show that ETF rebalancing is related to commodity futures prices. For natural gas, one standard deviation rise in ETF demand as a proportion of total market capitalization (3.08% for the first-month futures contract and 2.23% for the second) is related to an increase in the spread by 0.34%, but has no impact on

[31] In 2011, silver futures contracts were introduced at the Shanghai exchange. Using these futures as a control asset after 2011 produces similar results to the ones seen with gold. For robustness, I also rerun the regressions for silver and gold using basis and spread of futures contracts in platinum and other metals, where there are little to no ETFs. I did this test to address the concern that both silver and gold undergo the introduction of ETFs, and therefore, might not be a suitable control. The main results were similar.

[32] After 2010, increased volumes of crude oil from North Dakota and Canada flowed into Cushing (where WTI US crude oil is delivered). This led to a build-up in inventories and decreased the price of US crude oil, widening the spread with Brent.

the first-month basis. For silver and oil, one standard deviation rise in ETF demand (1.97% for silver and 0.08% for oil for the two-months contract) is related to an increase in the spread by 0.01% and 0.48%, respectively, but also has no impact on the basis. The lower magnitude for the silver market could be related to the fact that the average relative basis and spread for silver are very small compared with other commodity markets. For gold, ETF demand has no impact on basis or spread.[33] The fact that ETF rebalancing is also related to price changes in commodity markets illustrates that the impact of ETFs on prices is a general observation as opposed to being a feature of the VIX market alone. The observation that ETF demand has no significant impact on basis could be because the spot price is sometimes interpolated from the futures price or because in some commodities, it is easier to perform carry arbitrage compared to VIX.

[Table 1.8 about here]

1.7.2 Decomposition of ETF demand in commodity markets

Flow and leverage rebalancing have the highest impact on basis for commodities (Panel B of Table 1.8). For gas, the coefficient for flows is slightly larger: one standard deviation rise in flow rebalancing is related to an increase in the spread by 0.36%, compared with leverage rebalancing where one standard deviation rise is related to an increase of 0.29%. For silver and oil, the coefficient for leverage rebalancing is larger than the one for flows. Calendar rebalancing is mostly insignificant, except for silver: one standard deviation rise is related to a decrease in the front month basis of 0.01%. The relatively lower effect of calendar rebalancing could be due to the fact that commodity ETFs follow a benchmark with a rolling period of five days compared with the daily rolling in the VIX market.[34] Moreover, the proportion of ETFs in commodity markets is lower than in the VIX market. Therefore, calendar rebalancing could be more predictable and could have less-pronounced effects on prices.

[33] For robustness, I interpolated the spot price similar to Koijen et al. (2018) to address the concern that the data on spot prices from Bloomberg could be of a poor quality. The coefficient for basis was marginally significant in the gas market, for other markets the coefficients were still insignificant.

[34] The rolling period is usually from the 6th to the 10th business day of each month. This is the case for ETFs that follow benchmarks based on S&P Goldman Sachs Commodity Indices. Some ETFs follow Dow Jones Commodity Indices and rebalance from the 5th to the 9th day of each month. For commodity ETFs, the proportion invested in the first-month contract is:

$$\alpha_t = \begin{cases} 1, & \text{if } t < \text{6th business day} \\ \alpha_{t-1} - \frac{1}{5}, & \text{if 6th business day } \leq t \leq \text{10th business day} \\ 0, & \text{if } t > \text{10th business day} \end{cases} \tag{1.16}$$

1.7.3 Difference-in-differences (DD) regressions

In this section, I exploit the staggered nature of the ETF introduction in VIX and commodity markets, and implement DD regression design. The empirical strategy exploits the first ETF trading date to construct a before-after comparison between assets that are part of an ETF benchmark and similar assets that are not influenced by ETF demand.[35] To study the impact of the introduction of ETFs, I run pooled DD regressions. This methodology allows me to avoid confounding the effects of the ETF introduction with unobserved shocks to each specific commodity market. The period after the ETF introduction is defined as the post period (dummy *Post*). A futures contract is defined as being part of the treatment group if it is part of an ETF benchmark. The corresponding dummy variable is *Treat*:

$$b_{t,j} = \alpha_j + \lambda_t + \gamma_{Treat \times Post} \times Treat_j \times Post_t + \epsilon_{t,j}, \tag{1.17}$$

where $b_{t,j}$ is the relative basis or spread of asset j, α_j is an individual (futures) fixed effect and λ_t is a time fixed effect. I consider 10 assets in total (one futures and one hedge asset for VIX, gas, silver, gold and oil) for each maturity (one or two months). I also run a specification with asset-specific trends ($\alpha_{1,j}t$):

$$b_{t,j} = \alpha_{0,j} + \alpha_{1,j}t + \lambda_t + \gamma_{Treat \times Post} \times Treat_j \times Post_t + \epsilon_{t,j}. \tag{1.18}$$

To estimate the effects of the ETF introduction, I consider a 60-day window around the first trading date.[36] To the best of my knowledge, no other significant events influenced exclusively the specific asset, compared with a similar asset with no ETFs introduced around the event date, so the main driver of changes in asset prices was presumably the introduction of the ETF. The estimates from column 2 of Panels A and B in Table 1.9 show that the introduction of ETFs increases the first-month basis by 3.62% (82.5% in absolute value given the control group before-treatment mean of -4.39%), on average, and the spread by 1.78% (59.1% given the control group before-treatment mean of 3.01%). Figure 1.24 in the Appendix shows that the parallel trends assumption of the DD specification is largely satisfied. Column 3 of Table 1.9 shows that including asset-specific trends barely changes the estimates.

[35] Usually, the announcement date and the first date of trading are less than two days apart for most ETFs studied. Among others, Tang and Xiong (2012) run a similar DD regression for correlation measures in commodity markets.

[36] I use the first trading date of the first ETF introduced in a given market. I do not use introduction dates of inverse ETFs since usually the introduction of an inverse ETF coincides with that of a leveraged long ETF, and the net effect on total ETF demand is not clear.

[Table 1.9 about here]

The control group in the regressions from columns 1–3 consists of all futures contracts that did not experience ETF demand at a given point in time. One worry is that futures from different markets could be systematically different and therefore not an appropriate control group. To alleviate this concern, I also run the regression using a matched sample: I interact the $Post$ and $Treat$ dummies with a dummy for each individual asset class. The estimate on the $Treat \times Post$ dummy will then be a weighted average of the individual estimates for each market. The results are presented in column 4. The coefficients are still positive and statistically significant. Maturities longer than two months are not influenced by the introduction of the ETF, on average: the estimates on the third and fourth-month futures spreads are not significant, as shown in column 6 of Panels A and B. That is not surprising since ETFs trade in the first and the second-month futures contracts only.

The magnitude of the coefficients from the DD regressions is consistent with the general estimates for price impact from Table 1.4 and Table 1.8, and even slightly greater for some markets. For example, for VIX, the proportion of ETF rebalancing goes up, on average, by approximately 4.20% of the market capitalization for the first-month futures contract after the introduction of ETFs. The estimates from Table 1.4 would then imply an increase in basis of 2.45% ($4.20/2.42 \cdot 1.41\%$) and in EFG of 1.68% ($4.20/2.42 \cdot 0.97\%$). These are slightly lower than the DD estimates of 3.12% and 1.83%. I also conduct an event-study for realized futures returns in Table 1.23 in the Appendix and find that the introduction of ETFs decreased realized returns by 1–30 bps for various markets.

1.7.4 Synthetic commodity futures contract

To estimate the price deviation caused by ETFs in a similar way to the calculation of the EFG in section 1.4.4, I construct a synthetic futures contract using variables that are successful contemporaneous predictors of the futures price. The two variables that I chose for the synthetic control are spot price and the relative basis of a similar futures contract with no ETFs traded (they explain more than 96% of the daily variation in futures prices across all four commodity markets in the pre-ETF period). To construct the synthetic futures, I use the sample of data that ends six months before the ETF introduction (in-sample period) to determine the optimal weights, and then use these fixed weights in the out-of-sample period to construct the synthetic futures contract. To avoid confounding the effects of the ETF introduction with other factors, I conduct an event study for all commodity markets and VIX in a narrow window of 60 days around the first ETF introduction date. The dependent variable is the relative difference (gap) between the ETF-influenced futures price and the synthetic one. The results from column 5 of Panels A and B in Table 1.9 show that the introduction of an ETF increased the first-month gap by 2.24%, on average, and the second-month gap by 1.16%.

1.8 Conclusion

This paper shows that ETFs put pressure on prices in the most ETF-dominated asset classes: VIX and commodities. The research also proposes a model-independent approach for replicating the fundamental value of a VIX futures. The method makes it possible to isolate a non-fundamental part in VIX futures prices of 18.5 % per year that is strongly related to the rebalancing of ETFs. The paper also provides a decomposition of ETF demand into three main components: calendar rebalancing due to the roll from one futures contract to another, flow rebalancing due to inflow/outflow of money to the fund, and leverage rebalancing due to the maintenance of a constant daily leverage. The framework is flexible to accommodate various types of ETFs, including non-leveraged ETFs, equity and fixed income ETFs.

The research presents a simple strategy to understand the risks of trading against the leverage rebalancing of ETFs. The results show that providing liquidity to leveraged ETFs turns out to be a particular bet on variance, even in a market with a zero net share of ETFs. Trading against leveraged ETFs delivers large abnormal returns and Sharpe ratios for markets with significant ETF presence. The returns spike in turbulent times, when liquidity dries up.

The results from this research show that passive funds actively affect prices of underlying assets in the current era of an increasingly large ETF presence. While ETFs can increase liquidity and trading volume by attracting new capital (e.g., from retail investors), they also withdraw liquidity during extreme market times. These effects could be magnified if ETFs were used by unsophisticated, short-horizon investors. The recent bankruptcy of the largest inverse VIX ETF in 2018 and the extreme events and ETF bankruptcies in the oil market in 2020 are prominent examples of such effects. ETFs are transforming the financial industry and acting more as a “wrapper of views” rather than a “wrapper of assets” by allowing investors to get exposure to various trading strategies across traditional, and alternative asset classes. The problem with fully assessing the consequences of ETFs on underlying assets is that these investment vehicles are relatively new and under-researched. Time and future research will show whether ETFs are the greatest game-changer in the asset-management industry, or simply another "financial weapon of mass destruction".

1.9 Figures

FIGURE 1.1: ETF fractions of total market capitalization (number of futures contracts multiplied by the futures price) and the potential impact of leverage rebalancing for VIX, gas, silver, gold and oil. Monthly averages for the two front contracts. The solid black line shows net ETF fraction (long ETFs minus inverse ETFs) in the total market capitalization of the first and second futures contracts: $\sum_{j=1}^{N} L_j A_{j,t}/Mkt\ cap_t$, where L_j is the leverage of ETF j (L_j<0 for inverse ETFs) and $A_{j,t}$ are its assets under management (AUM) at time t. The dashed red line is $\Gamma_t/Mkt\ cap_t$: a measure of the total rebalancing demand by leveraged ETFs (explained in section 1.5.2) scaled by market capitalization.

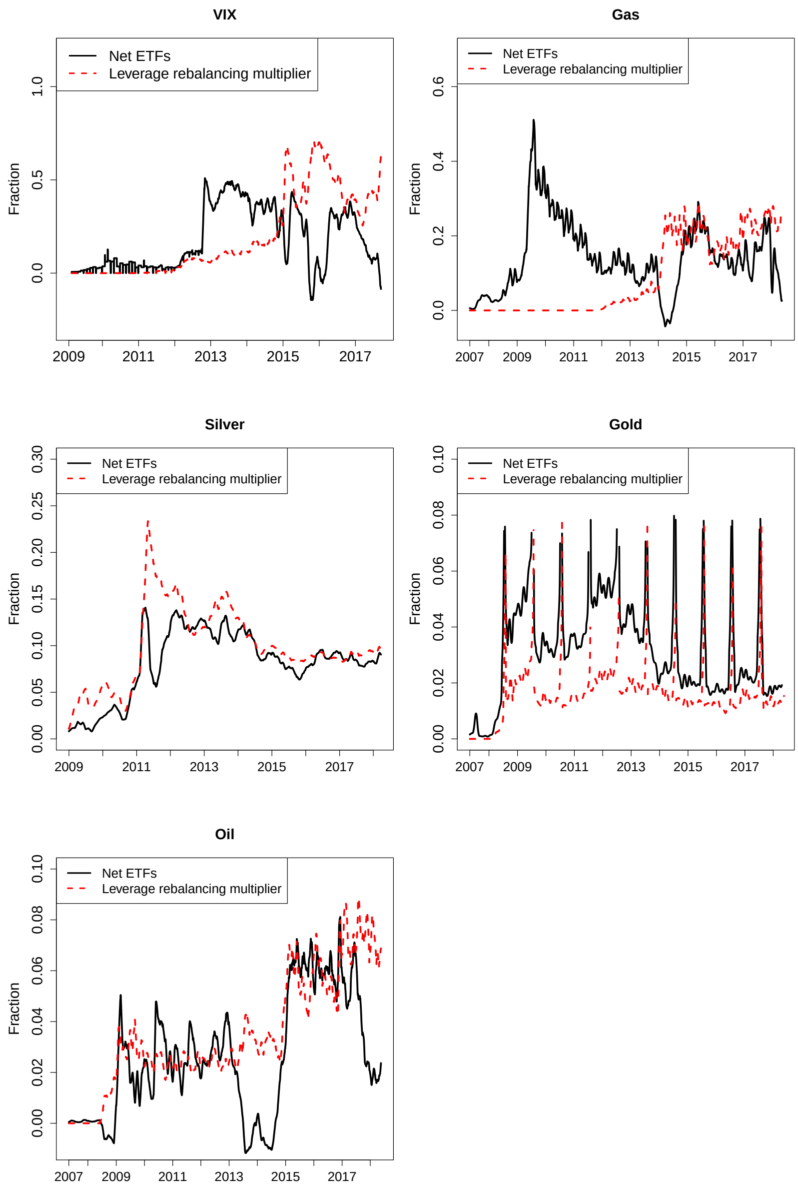

FIGURE 1.2: Realized VIX futures premium before and after ETFs. The chart shows the average size of the VIX futures premium for different maturities before ETFs (excluding the crisis period, June 2004 – September 2008) and after ETFs (January 2009 – February 2018). The premium is calculated as the annualized net return of a short-seller of a VIX futures $\frac{F_{t,T}-F_{T,T}}{F_{t,T}}$, where T (Term) is maturity (in months).

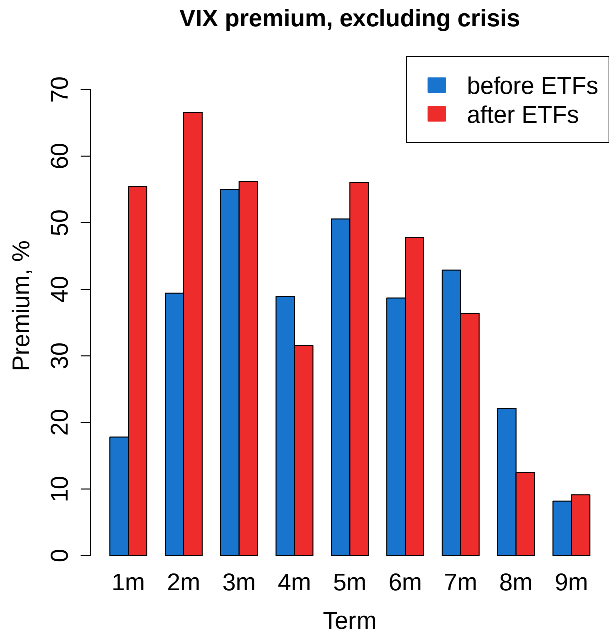

FIGURE 1.3: Positions of traders. The top two panels show net futures positions of different types of traders in the VIX and gas market, respectively. The data is from the Traders in Financial Futures (TFF) reports by the CFTC. The bottom two panels show weekly net ETF positions and net Dealer/Swap positions in the two markets. "Asset Mgr" are asset managers (mostly pension funds, endowments, insurance companies and mutual funds), "Lev Money" are mostly hedge funds and other proprietary traders.

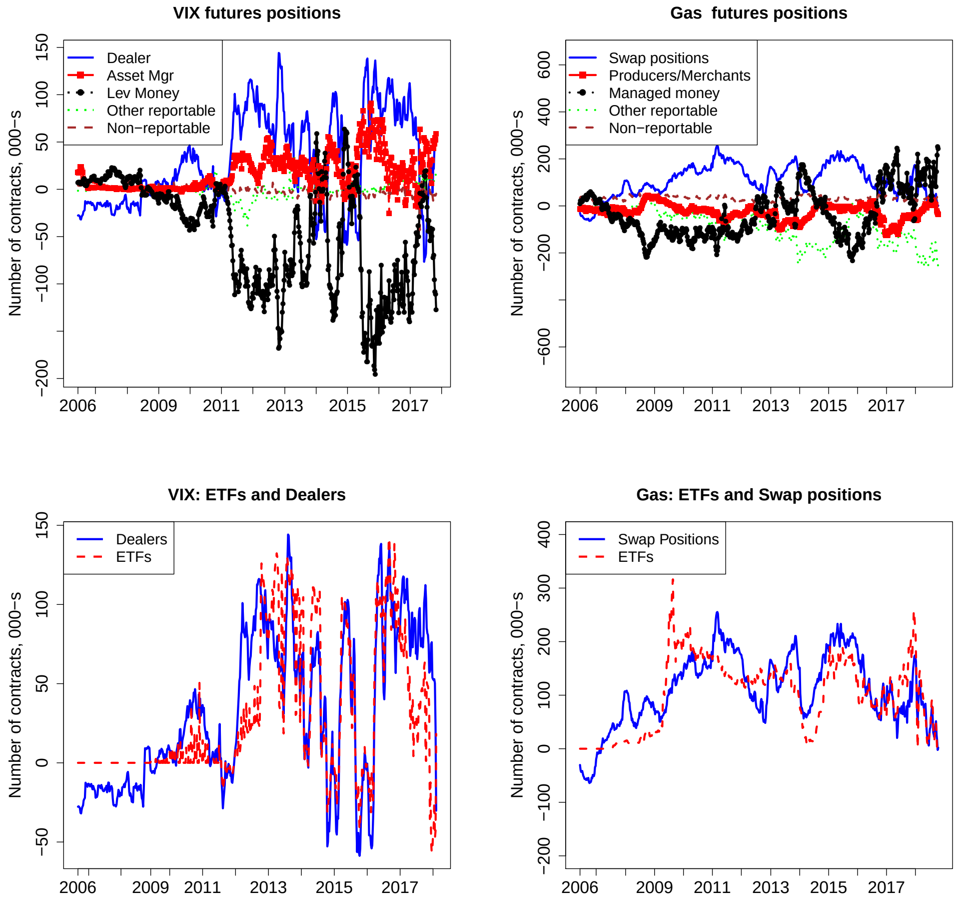

FIGURE 1.4: Basis decomposition. Illustrative example. The left panel shows the decomposition of basis into futures change ($F_{t,T} - F_{T,T}$) and spot change ($S_T - S_t$), the right panel – the decomposition into $EFG_t = F_{t,T} - E_t^{\mathrm{Q}}(S_T)$, Realized VIX premium $= E_t^{\mathrm{Q}}(S_T) - S_T$, and spot change. Maturity is in months.

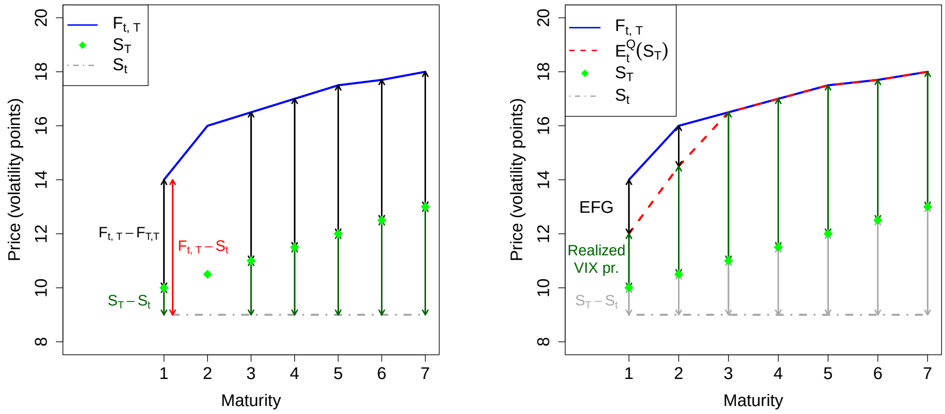

FIGURE 1.5: ETF futures gap. Weekly averages. The figure shows the dynamics of the ETF futures gap (EFG) for maturities at 1–4 months. The first dashed vertical line depicts the date when the first VIX ETF was introduced, the second shows the inception date of the first leveraged VIX ETF.

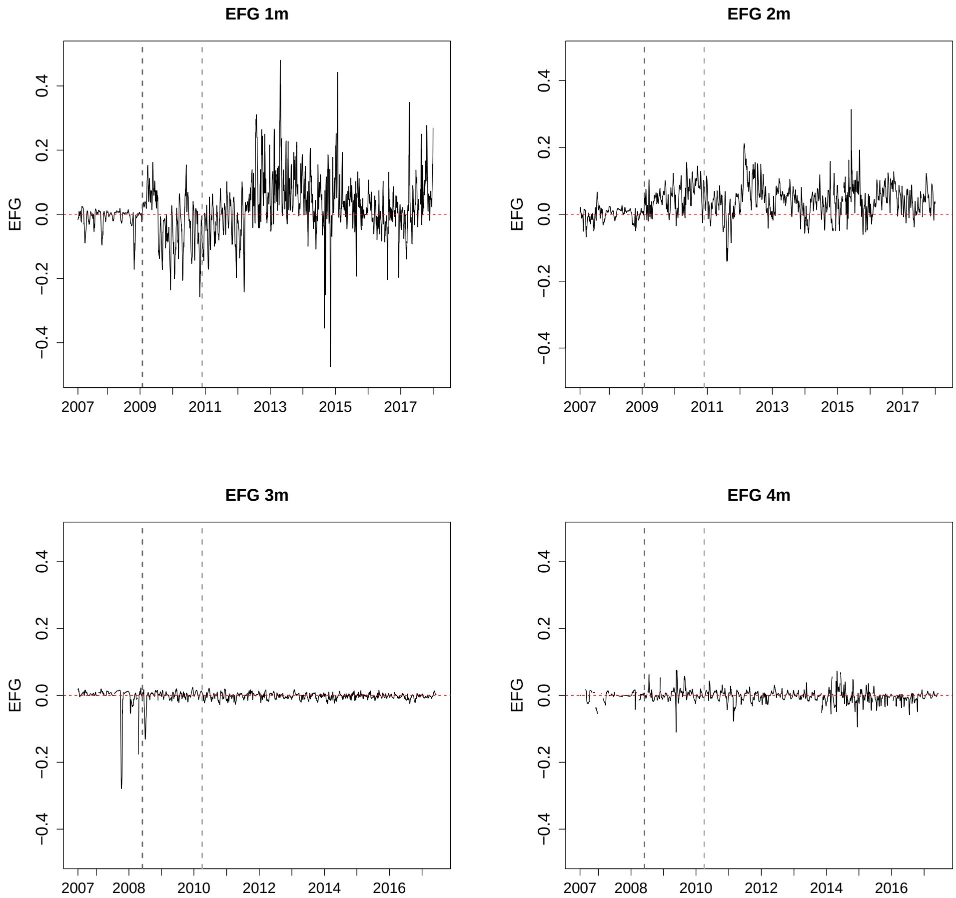

FIGURE 1.6: EFG and rebalancing demand from ETFs. The left panel illustrates the dynamics for the first-month EFG, the right panel for the second-month EFG. Demand is in million USD. The graphs show a representative sample (December 2014 – February 2015 for the first-month gap, and August 2015 – October 2015 for the second-month gap) to illustrate the typical pattern.

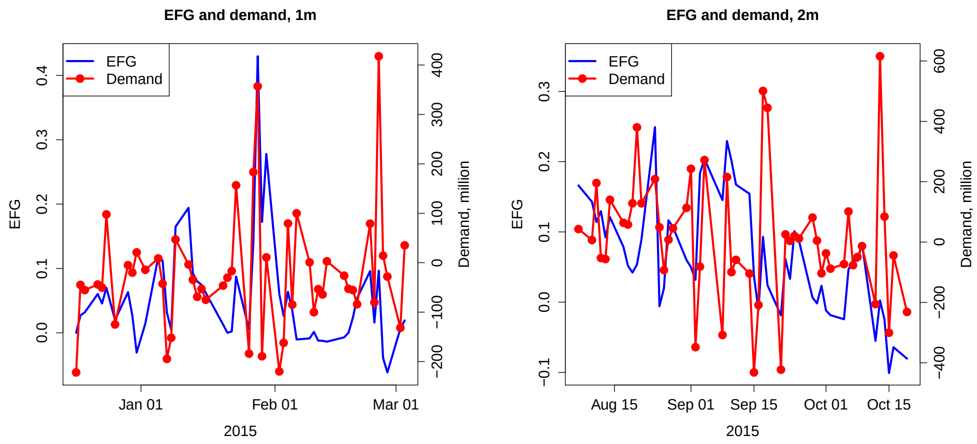

FIGURE 1.7: MSE. The left panel shows the dynamics of the mean squared error from a prediction of realized spot VIX prices (S_T) based on the synthetic futures $\mathbb{E}_t^{\mathrm{Q}}(S_T)$ (red dashed line) and the ETF-influenced futures $F_{t,T}$ (solid black line) for the first-month futures. For each t, I calculate the squared errors as $(S_T - \mathbb{E}_t^{\mathrm{Q}}(S_T))^2$ and $(S_T - F_{t,T})^2$. The right panel illustrates the average MSE by days to maturity for the whole sample. The graphs show that $\mathbb{E}_t^{\mathrm{Q}}(S_T)$ is a better predictor of the fundamental value S_T compared to $F_{t,T}$.

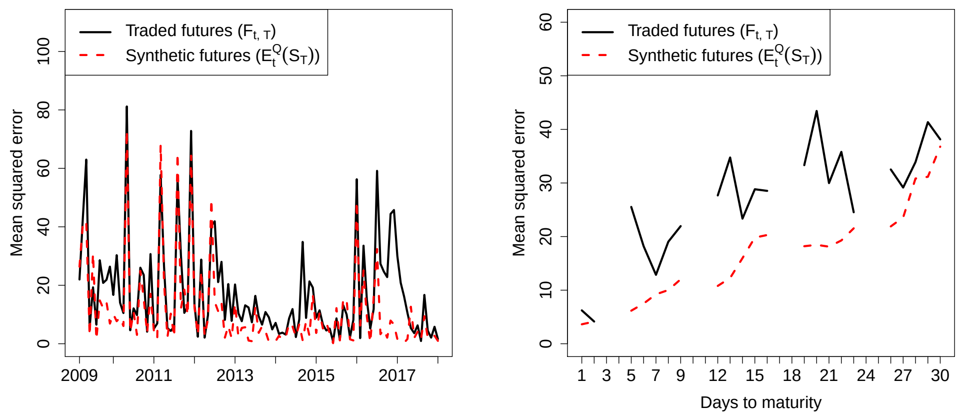

FIGURE 1.8: Open interest dynamics before and after ETF introduction. The left panel shows typical dynamics of open interest for futures maturities at one, two, three and four months for the period before ETFs were introduced. The right panel illustrates typical dynamics after the introduction of ETFs. The emphasized straight lines show the usual cycle of open interest. The graphs show a representative sample (July 2007 – December 2007 and October 2012 – May 2013) to illustrate the typical pattern.

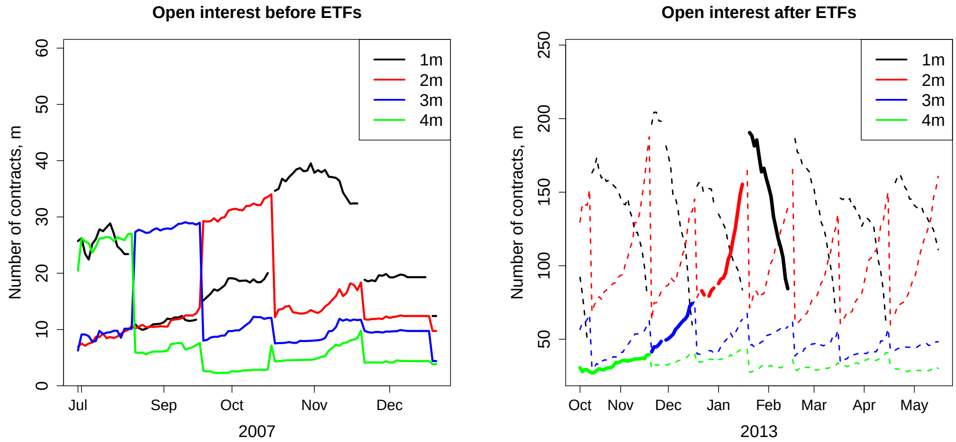

FIGURE 1.9: Decomposition of ETF demand. The top left panel illustrates the dynamics of VIX ETFs' demand decomposition. Demand is in absolute values. The top right panel and the two bottom panels show the effects of change in calendar, leverage, and flow rebalancing on the futures curve: assuming ETFs are net long VIX futures ($\sum_{j=1}^{N} L_j A_{j,t} > 0$). The blue line illustrates the curve before the impact of the rebalancing demand, the red and green ones after it. Maturity is in months, futures price in volatility points.

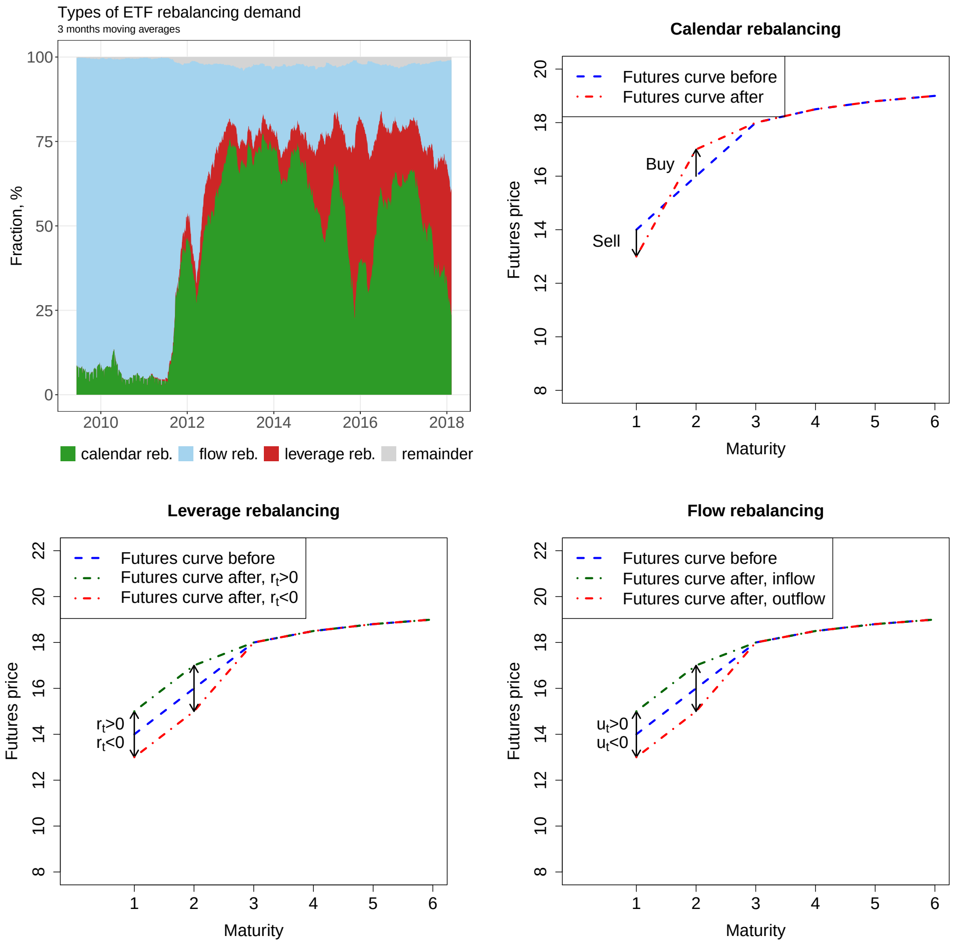

FIGURE 1.10: Trading against opposite ETFs. The figure shows the profit dynamics of liquidity provision to opposite ETFs using a binomial tree example. The graph illustrates the dynamics of the ETF benchmark and the corresponding profits for an arbitrageur who sells short a pair of opposite ETFs ($L = 2$ and $L = -2$). For each period, the parameters of the tree are $u = 1.05$ and $ud = 1$. Red areas indicate nodes where the arbitrageur loses money, and green ones show where the arbitrageur makes profit. More color-intense nodes indicate larger losses or profits.

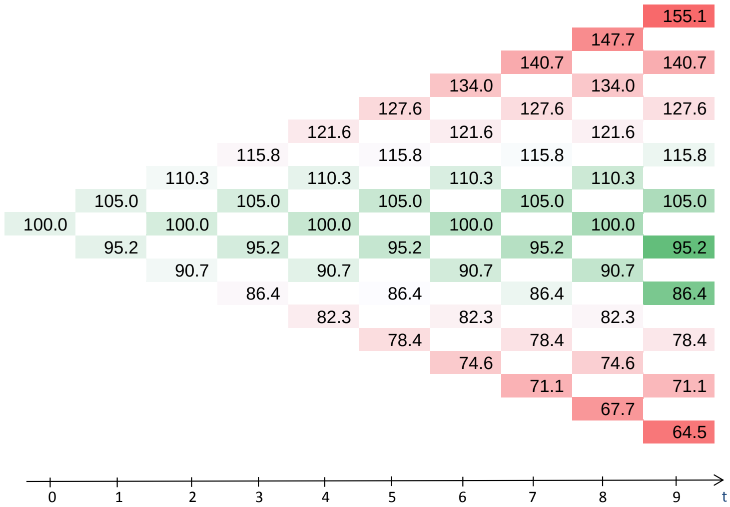

FIGURE 1.11: Intra-day returns on the short-both strategy for several markets. $E(R)$ is the average annualized return, SR is the annualized Sharpe ratio, and $SR,\ fee$ is the annualized Sharpe ratio after subtracting borrowing fees. The dotted blue line shows the cumulative return on the long ETF (leverage L), the dashed red line the return on the inverse ETF (leverage $-L$). The solid black line shows the returns on the strategy that sells short the long and the inverse ETF. VIX: $L = 1$; silver: $L = 2$; gas, gold: $L = 3$.

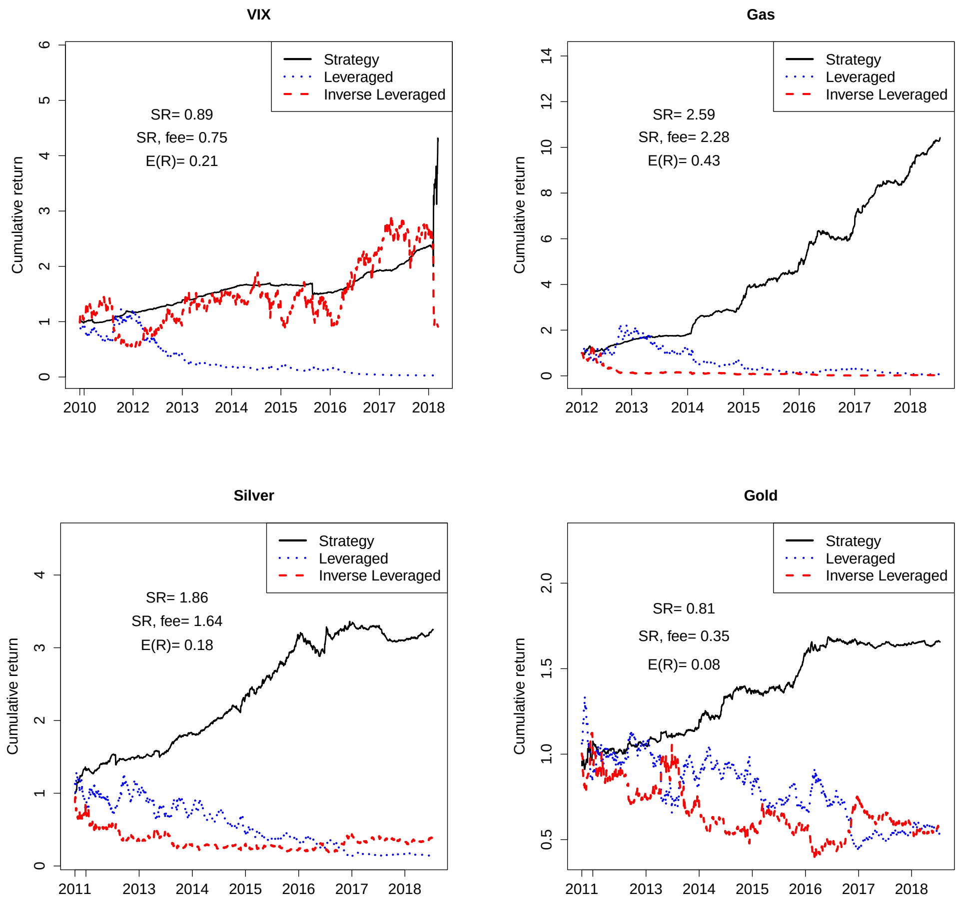

1.10 Tables

TABLE 1.1: Share of ETFs in market capitalization and in trading volume. The table shows the average fraction (over time) of ETFs in total market capitalization and in volume of trading for several markets. The fraction in total market capitalization is calculated as dollar size of all ETFs divided by the dollar capitalization of the benchmark index. ETF fraction in trading volume is calculated as ETF trading volume divided by the trading volume of the benchmark index. The data is at a daily frequency, from the first ETF trading date in a given asset to February 2018.

Market	Long ETFs fraction (%)	Inverse ETFs fraction (%)	Net ETFs fraction (%)	ETFs fraction in trading volume (%)
VIX	40.89	16.42	24.47	206.68
Natural Gas	17.52	3.81	13.71	18.31
Silver	10.12	2.01	8.11	35.11
Gold	4.74	0.40	4.34	9.69
Oil	4.12	1.57	2.55	1.34
Nasdaq	1.96	0.01	1.95	35.98
S&P 500	1.02	0.03	0.99	18.82
Financials	0.14	0.07	0.07	24.83
Real estate	0.14	0.09	0.05	2.80
Basic materials	0.05	0.01	0.04	0.93
Russell 2000	0.07	0.06	0.01	11.85
Utilities	0.00	0.00	0.00	0.07
Value	0.00	0.00	0.00	0.01
Consumer service	0.00	0.00	0.00	0.01
Treasuries 7-10 years	0.00	0.00	0.00	0.00

TABLE 1.2: Summary statistics for the VIX market. S is spot VIX, F_{T_1}, F_{T_2}, ... F_{T_9} are the first, second ... ninth generic futures. All prices are in volatility points. Basis is $F_{T_1} - S$, spread is $F_{T_2} - F_{T_1}$. r^{VXX} is the daily return on the first long VIX ETF (ticker VXX), EFG_{T_1} and EFG_{T_2} are the ETF futures gaps for the first and second-month contracts, respectively. The numbers for r^{VXX}, EFG_{T_1} and EFG_{T_2} are in %, except skewness, kurtosis and number of observations. The pre-last row presents the estimate (β) from a regression on spot VIX. The data is at a daily frequency, from June 2004 (January 2009 for r^{VXX}) to February 2018. The "bef. ETFs" period is prior to 29 January 2009. The "bef. ETFs, excl. crisis" period is prior to 15 September 2008.

	S	F_{T_1}	F_{T_2}	F_{T_3}	F_{T_4}	F_{T_5}	F_{T_6}	F_{T_7}	F_{T_8}	F_{T_9}	Basis	Spread	r^{VXX}	EFG_{T_1}	EFG_{T_2}
Mean	18.57	19.11	20.00	21.01	21.46	21.72	22.07	22.73	22.52	21.75	0.54	0.83	-0.17	1.04	3.61
Mean, bef. ETFs	18.87	18.93	19.38	20.80	21.16	20.91	20.99	22.75	23.00	23.41	0.06	0.46		-0.84	0.09
Mean, bef. ETFs, excl. crisis	16.70	17.15	17.96	19.24	19.68	19.62	19.76	21.05	21.22	21.70	0.45	0.81		-0.51	0.38
Mean, after ETFs	18.42	19.21	20.31	21.10	21.57	22.03	22.45	22.73	22.43	21.27	0.79	1.10	-0.17	1.54	4.44
Std. dev.	9.13	8.34	7.53	7.10	6.67	6.34	6.10	5.98	5.69	5.37	1.79	1.76	3.20	10.28	5.27
Min	9.14	9.60	11.32	12.22	12.97	13.47	13.97	14.43	14.32	10.25	-23.31	-21.10	-14.25	-49.38	-21.10
Max	80.86	67.95	59.77	54.67	50.58	47.07	45.26	45.00	44.45	44.46	4.98	5.45	33.44	43.47	36.00
10%	11.43	12.20	13.13	14.08	14.53	15.12	15.47	16.12	16.70	16.74	-0.64	-0.55	-3.47	-9.51	-2.43
50%	15.60	16.22	17.50	18.88	19.62	19.88	20.25	21.02	20.50	20.00	0.69	1.01	0.00	0.29	3.28
90%	28.27	28.10	29.39	30.48	30.40	30.70	30.80	31.00	30.80	29.62	2.01	2.35	28.33	11.71	9.98
Skewness	2.57	2.26	1.87	1.55	1.32	1.18	1.06	0.96	1.13	1.54	-4.72	-3.96	1.45	0.44	0.28
Kurtosis	11.84	9.50	7.35	5.95	4.88	4.25	3.79	3.50	3.93	5.32	43.71	32.9	13.53	14.09	5.10
β with respect to VIX	1.00	0.90	0.77	0.69	0.63	0.59	0.55	0.51	0.48	0.42					
Observations	3,442	3,442	3,397	3,095	3,053	3,091	3,040	2,771	2,552	1,858	3,442	3,397	2,260	2,550	2,641

TABLE 1.3: Predictive power of basis. The table presents the results from a predictive regression of spot or futures price changes on basis with daily frequency. S_t is spot price, $F_{t,T}$ is futures price for maturity T. Here and in all subsequent tables standard errors are computed using the Newey-West (see, e.g., Newey and West, 1987) estimator with three lags. The major results were unchanged with more lags. *,**, and *** indicate statistical significance at the 10%, 5%, and 1% levels here and in all following tables. Daily frequency, February 2009 – December 2017.

Panel A: Spot VIX on basis: $S_T - S_t = \alpha_1 + \beta_1 \cdot (F_{t,T} - S_t) + \epsilon_{1,t}$

	T=1m	T=2m	T=3m	T=4m	T=5m	T=6m	T=7m	T=8m
β_1	0.02	0.27	0.64***	0.72***	0.84***	0.94***	0.98***	1.01***
	(0.24)	(0.17)	(0.08)	(0.06)	(0.07)	(0.06)	(0.05)	(0.05)
R^2	0.00	0.03	0.09	0.10	0.14	0.18	0.23	0.27
Observations	2,156	2,139	2,115	2,013	2,075	2,051	2,031	1,849

Panel B: VIX futures on basis: $F_{T,T} - F_{t,T} = \alpha_2 + \beta_2 \cdot (F_{t,T} - S_t) + \epsilon_{2,t}$

	T=1m	T=2m	T=3m	T=4m	T=5m	T=6m	T=7m	T=8m
β_2	-0.98***	-0.73***	-0.36***	-0.28***	-0.16***	-0.06**	-0.02	0.01
	(0.23)	(0.15)	(0.07)	(0.06)	(0.07)	(0.06)	(0.06)	(0.05)
R^2	0.14	0.10	0.04	0.04	0.02	0.01	0.00	0.00
Observations	2,156	2,139	2,115	2,013	2,075	2,051	2,031	1,849

TABLE 1.4: Impact of ETF demand in the VIX market. The table presents regression results for the basis, spread, and ETF futures gap (EFG). Panel A shows results for the first-month basis ($b_{t,1}$), and the spread between the first-month and the second-month futures contracts ($b_{t,2}$). Columns 1–2 and 6–7 present the regressions for absolute basis and spread, columns 3–5 and 8–10 for relative basis and spread. $D_{t,i}^{\$,all}$ is the ETF demand for the $i-$th futures contract. Columns 4 and 9 use raw demand, whereas all other columns use demand scaled by total market capitalization. Columns 5 and 10 use ETF demand computed using lagged returns on the benchmark and the two futures contracts (by replacing r_{t+1}, $r_{t+1}^{F_1}$, $r_{t+1}^{F_2}$ with r_t, $r_t^{F_1}$, $r_t^{F_2}$, respectively, in equations (1.10) and (1.11) for a given time $t+1$). Panel B presents the results for one and two-months gaps ($\mathrm{EFG}_{t,1}$ and $\mathrm{EFG}_{t,2}$). Columns 3 and 7 use raw demand, whereas all other columns use demand scaled by total market capitalization. $b_{t,i}^H$ is the relative basis of a hedge asset (synthetic futures contract constructed from options), $\mathrm{r}_{bmk,t}$ is the return on the ETF benchmark, $\sigma_{bmk,t}^2$ is intra-day variance of the ETF benchmark (calculated using 5-minute intervals), $OI_{t,i}$ is open interest for futures i, S_t is spot price. Liquidity ($\mathrm{Liq}_{t,i}$) is the difference in relative bid-ask spreads ($\frac{Ask-Bid}{Mid}$) between the hedge asset and VIX futures contract. TED_t is the spread between 3-month LIBOR in USD and the interest rate of Treasury bills. α_t is the fraction of ETF wealth invested in the front-month futures contract. All independent variables are standardized. Daily frequency, February 2009 – December 2017.

Panel A: Impact on basis and spread

Dependent variables	$b_{t,1}$, abs		$b_{t,1}$, rel			$b_{t,2}$, abs		$b_{t,2}$, rel		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$D_{t,1}^{\$,all}$	**0.26****	**0.32*****	**1.41*****	**1.45*****		-0.04	-0.09	-0.21	-0.29	
	(0.12)	(0.03)	(0.13)	(0.16)		(0.04)	(0.08)	(0.18)	(0.26)	
$D_{t,2}^{\$,all}$						**0.17*****	**0.18*****	**0.48*****	**0.97*****	
						(0.04)	(0.03)	(0.12)	(0.15)	
$D_{t,1}^{\$,all}(r_{t-1})$					**1.25*****					-0.12
					(0.31)					(0.02)
$D_{t,2}^{\$,all}(r_{t-1})$										**0.37****
										(0.18)
$b_{t,i}^H$		1.29***	5.21***	5.22***	5.56***		1.50***	5.64***	5.82***	5.80***
		(0.24)	(1.18)	(1.18)	(1.20)		(0.05)	(1.13)	(1.12)	(1.13)
$\mathrm{r}_{bmk,t}$		−0.42	−1.65	−1.85	−1.84		−0.21*	−0.85	−1.20	−1.13
		(0.36)	(1.60)	(1.66)	(1.66)		(0.12)	(0.62)	(1.03)	(1.00)
$\sigma_{bmk,t}^2$		−0.13***	−0.64***	−0.78***	−0.85***		−0.07**	−0.48**	−0.58***	−0.57**
		(0.04)	(0.19)	(0.21)	(0.22)		(0.03)	(0.19)	(0.21)	(0.23)
$\mathrm{OI}_{t,i}$		0.18***	2.15***	2.16***	1.85***		−0.06	0.61***	0.45***	0.63***
		(0.04)	(0.21)	(0.21)	(0.21)		(0.04)	(0.18)	(0.17)	(0.18)
S_t		−0.85***	−2.75***	−2.69***	−2.99***		−1.03***	−4.03***	−3.97***	−4.20***
		(0.15)	(0.25)	(0.25)	(0.26)		(0.11)	(0.23)	(0.23)	(0.24)
$\mathrm{Liq}_{t,i}$		0.13***	0.85***	0.89***	0.99***		0.28***	1.19***	1.17***	1.29***
		(0.04)	(0.16)	(0.16)	(0.16)		(0.05)	(0.16)	(0.15)	(0.16)
α_t		0.11***	1.00***	1.12***	1.07***		−0.03	−0.17	−0.18	−0.18
		(0.04)	(0.16)	(0.16)	(0.16)		(0.04)	(0.15)	(0.15)	(0.15)
Observations	1,945	1,945	1,945	1,945	1,945	1,922	1,922	1,922	1,922	1,922
R^2	0.07	0.37	0.42	0.42	0.38	0.09	0.48	0.45	0.46	0.44

Panel B: Impact on EFG

Dependent variables	$\mathrm{EFG}_{t,1}$				$\mathrm{EFG}_{t,2}$			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$D_{t,i}^{\$,all}$	**1.21*****	**0.97*****	**2.14*****		**0.57*****	**0.45*****	**0.58*****	
	(0.25)	(0.40)	(0.63)		(0.15)	(0.12)	(0.15)	
$D_{t,i}^{\$,all}(r_{t-1})$				**0.77*****				**0.38****
				(0.26)				(0.12)
$\mathrm{EFG}_{t-1,i}$		6.03***	6.02***	6.02***		3.92***	3.93***	4.09***
		(0.72)	(0.73)	(0.72)		(0.19)	(0.19)	(0.22)
$\mathrm{r}_{bmk,t}$		0.60***	0.47***	0.44***		0.20***	0.15***	0.22***
		(0.12)	(0.10)	(0.10)		(0.05)	(0.05)	(0.05)
$\sigma_{bmk,t}^2$		1.31	1.52	0.85		-0.40	-0.40	-0.66
		(0.89)	(0.94)	(0.83)		(0.57)	(0.57)	(0.58)
$\mathrm{OI}_{t,i}$		1.62	1.92	1.18		-1.14	-1.22	-0.42
		(1.53)	(1.54)	(1.53)		(1.16)	(1.16)	(0.85)
S_t		0.10	0.29	-0.02		-0.36***	-0.36***	-0.21***
		(0.64)	(0.61)	(0.66)		(0.08)	(0.07)	(0.06)
$\mathrm{Liq}_{t,i}$		0.88**	0.92**	0.75**		0.02	0.00	0.16
		(0.41)	(0.41)	(0.37)		(0.11)	(0.11)	(0.11)
TED_t		0.51	0.36	1.28		1.46***	1.42***	1.24***
		(0.97)	(0.96)	(1.06)		(0.41)	(0.42)	(0.40)
α_t		0.62***	0.63***	0.84***		-0.33***	-0.35***	0.18
		(0.24)	(0.25)	(0.25)		(0.12)	(0.12)	(0.11)
Observations	1,898	1,898	1,898	1,895	1,824	1,824	1,824	1,816
R^2	0.24	0.44	0.45	0.44	0.26	0.58	0.58	0.62

TABLE 1.5: Impact of ETF demand components on futures. Calendar rebalancing, leverage rebalancing, flow rebalancing, and remainder are calculated based on equations (1.10) and (1.11). All independent variables are standardized. Controls include basis or spread of a hedge asset (synthetic futures constructed from options), time to maturity, variance of benchmark, return on benchmark, spot price, open interest and liquidity measured by bid-ask spreads. $b_{t,1}$ is relative basis, $b_{t,2}$ is relative spread. $EFG_{t,1}$ is the ETF futures gap for one month, $EFG_{t,2}$ for two months. Dependent variables are in %. All demand components are scaled by market capitalization. The estimates on the four components do not exactly add up to the estimate for total demand from Table 1.4 because the variables are standardized. Daily frequency, February 2009 – December 2017.

Dependent variables	$b_{t,1}$, rel	$b_{t,2}$, rel	$EFG_{t,1}$	$EFG_{t,2}$	$EFG_{t,1}$	$EFG_{t,2}$
	(1)	(2)	(3)	(4)	(5)	(6)
Calendar reb$_{t,i}$	-0.65***	0.31***	-0.28*	0.17*		
	(0.11)	(0.08)	(0.05)	(0.10)		
Leverage reb$_{t,i}$	0.41**	0.27**	1.17***	0.40***		
	(0.20)	(0.13)	(0.32)	(0.10)		
Flow reb$_{t,i}$	1.19***	0.68***	0.13**	0.34***		
	(0.10)	(0.11)	(0.06)	(0.08)		
Remainder$_{t,i}$	0.14	-0.60	-0.31	-0.38		
	(0.14)	(0.52)	(0.30)	(0.42)		
$\sum_{j=1}^{N} L_j A_{j,t-1}$					-0.39**	0.21**
					(0.19)	(0.10)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,936	1,913	1,890	1,819	1,890	1,819
R^2	0.55	0.36	0.48	0.61	0.29	0.38

TABLE 1.6: Regressions of leveraged ETFs' returns on benchmarks. The table shows the estimates of the regression $r_{ETF,t} = \alpha + \beta r_{bmk,t} + \epsilon_t$ for several markets and leverages. The last two rows of estimates present the results for the subsamples of high (above median) and low (below median) variance of the benchmark, respectively. Daily frequency, from the first leveraged ETF inception date in a given market to December 2018.

Dependent variables	VIX		Gas		Silver		Gold	
	L=2	L=-1	L=3	L=-3	L=3	L=-3	L=3	L=-3
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
α	-0.00**	-0.00*	-0.00	-0.00	-0.00	-0.00	-0.00	-0.00
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
β	1.93***	-0.92***	2.85***	-2.84***	2.72***	-2.70***	2.88***	-2.89***
	(0.01)	(0.01)	(0.02)	(0.02)	(0.03)	(0.03)	(0.03)	(0.03)
Observations	1,790	1,803	1,786	1,786	1,853	1,853	1,853	1,853
R^2	0.97	0.73	0.94	0.92	0.87	0.86	0.85	0.86
β, high var	1.91***	-0.90***	2.82***	-2.80***	2.70***	-2.64***	2.82***	-2.84***
	(0.01)	(0.02)	(0.02)	(0.03)	(0.04)	(0.04)	(0.04)	(0.04)
β, low var	1.94***	-0.93***	2.89***	-2.90***	2.75***	-2.75***	2.91***	-2.92***
	(0.01)	(0.01)	(0.03)	(0.03)	(0.06)	(0.06)	(0.07)	(0.07)

TABLE 1.7: Pooled regressions of the short-both strategy returns (r_{SB}). The sample consists of 14 markets, for the period from the first leveraged ETF introduction date in a given market to December 2018 (daily frequency). The markets are: VIX, gas, silver, oil, gold, S&P 500 Index, Nasdaq Composite, financials, real estate, basic materials, Russell 2000, utilities, value stocks and mid-term Treasuries. Standard errors are double-clustered by market and time. All independent variables are standardized. $\Phi = \sum_{j=1}^{N} L_j A_j / Mkt\ cap$ is the net fraction of ETFs scaled by total market capitalization, $\Gamma = \sum_{j=1}^{N} L_j (L_j - 1) A_j / Mkt\ cap$ is ETFs' leverage rebalancing multiplier scaled by total market capitalization. Premium to NAV is the relative difference between ETF price and NAV, Volume is the fraction of ETF trading volume in total trading volume. Flow is the net flows to all ETFs in a given market, Liquidity is based on relative bid-ask spreads. R_f is the risk-free rate, $R_M - R_f$, HML, SMB, CMA, RMW, Mom are the Fama–French five factors and momentum. Refer to Table 1.4 for definitions of other variables.

Dependent variables	r_{SB} (%)			
	(1)	(2)	(3)	(4)
Φ	0.013***	0.012**	0.008*	0.007
	(0.003)	(0.006)	(0.005)	(0.005)
Γ	0.073***	0.068***	0.070***	0.069***
	(0.024)	(0.023)	(0.022)	(0.022)
σ^2_{bmk}		0.053***	0.055***	0.056***
		(0.013)	(0.009)	(0.009)
$\Gamma \times \sigma^2_{bmk}$		0.015**	0.010*	0.010*
		(0.007)	(0.005)	(0.005)
r^2_{bmk}		-0.024***	-0.021***	-0.021***
		(0.001)	(0.003)	(0.003)
r_{bmk}		0.0001	-0.006	-0.008
		(0.006)	(0.007)	(0.006)
Premium to NAV			-0.140	-0.140
			(0.097)	(0.096)
Volume			-0.042*	-0.043*
			(0.025)	(0.024)
Flow			0.003	0.002
			(0.003)	(0.003)
Liquidity			0.021**	0.021**
			(0.009)	(0.009)
R_f				0.053
				(0.052)
$R_M - R_f$				0.009
				(0.007)
HML				-0.010
				(0.009)
SMB				0.013
				(0.010)
CMA				0.011
				(0.008)
RMW				0.013
				(0.019)
Mom				-0.012
				(0.008)
Observations	28,263	25,724	20,468	20,468
R^2	0.367	0.602	0.659	0.640

TABLE 1.8: Impact of ETF demand in commodity markets. Calendar rebalancing, leverage rebalancing, flow rebalancing, and remainder are calculated based on equations (1.10) and (1.11). All independent variables are standardized. $b_{t,1}$ is relative basis, $b_{t,2}$ is relative spread, $b^H_{t,i}$ is the relative basis or spread of a synthetic futures contract: all in %. Controls include time to maturity, variance of benchmark, return on benchmark, spot price, open interest, and liquidity measured by bid-ask spreads. In the regressions for $b_{t,2}$, I control for the demand for the first-month contract. For gas and oil, I also control for the difference in spot prices of the control asset versus the traded contract. Daily frequency, from the first ETF introduction date in a given market to February 2018.

Panel A: Total effect								
Dependent variables	Gas		Silver		Oil		Gold	
	$b_{t,1}$	$b_{t,2}$	$b_{t,1}$	$b_{t,2}$	$b_{t,1}$	$b_{t,2}$	$b_{t,1}$	$b_{t,2}$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$D^{\$,all}_{t,i}$	-0.03	0.34***	0.00	0.01*	0.50	0.48**	-0.01	0.00
	(0.04)	(0.11)	(0.01)	(0.00)	(0.39)	(0.23)	(0.01)	(0.01)
$b^H_{t,i}$	6.92***	4.39***	0.05***	0.03***	1.53***	1.95***	0.03***	0.03***
	(0.92)	(0.19)	(0.01)	(0.01)	(0.42)	(0.06)	(0.01)	(0.01)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,182	2,231	2,096	2,103	2,097	2,083	2,054	2,167
R^2	0.29	0.27	0.83	0.67	0.31	0.53	0.39	0.47

Panel B: Split on components								
Dependent variables	Gas		Silver		Oil		Gold	
	$b_{t,1}$	$b_{t,2}$	$b_{t,1}$	$b_{t,2}$	$b_{t,1}$	$b_{t,2}$	$b_{t,1}$	$b_{t,2}$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Calendar $\text{reb}_{t,i}$	-0.78	0.07	-0.01*	0.00	-0.48	0.56	-0.03	-0.02
	(0.56)	(0.10)	(0.01)	(0.01)	(0.40)	(0.41)	(0.02)	(0.01)
Leverage $\text{reb}_{t,i}$	-1.20	0.29***	0.00	0.02**	0.42*	0.53*	-0.02	0.00
	(1.12)	(0.10)	(0.00)	(0.01)	(0.24)	(0.29)	(0.02)	(0.00)
Flow $\text{reb}_{t,i}$	0.93	0.36***	0.01	0.01***	0.37	0.30**	-0.01	0.00
	(0.68)	(0.10)	(0.01)	(0.00)	(0.29)	(0.12)	(0.02)	(0.01)
$\text{Remainder}_{t,i}$	1.95	0.53	0.00	0.00	-0.56	0.55	0.03*	0.01
	(1.84)	(0.50)	(0.00)	(0.00)	(0.30)	(0.67)	(0.02)	(0.01)
$b^H_{t,i}$	7.12***	4.78***	0.06***	0.03***	1.75***	1.91***	0.03***	0.03***
	(1.21)	(0.50)	(0.01)	(0.01)	(0.52)	(0.28)	(0.01)	(0.01)
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	2,134	2,218	2,067	2,072	2,066	2,037	2,012	2,115
R^2	0.35	0.32	0.67	0.62	0.60	0.67	0.42	0.50

TABLE 1.9: DD estimates. The table shows the estimates from regressions (1.17) and (1.18) with a 60-day window around the first ETF introduction date for VIX, gas, silver, gold, and oil markets. $b_{t,1}$ is relative basis, $b_{t,2}$ is relative spread between the second-month and the first-month futures contracts. $b_{t,3}$ is relative spread between the third-month and the second-month futures contracts, $b_{t,4}$ is relative spread between the fourth-month and the third-month futures contracts. Column 4 is for the matched sample. Column 5 shows the results of an event study for EFG. Since the dependent variable is difference between two prices, I run an event study as opposed to DD regression because the scale of prices for various markets is different. Panel A shows the results for the first-month basis, the first-month EFG and the third-month spread. Panel B shows the results for the second-month spread, the second-month EFG and the fourth-month spread. Standard errors in columns 1–4 and 6 are double-clustered by asset and time. Standard errors in column 5 are computed using the generalized method of moments (GMM) as synthetic futures for commodities are based on estimated regression coefficients. Refer to Table 1.4 for variable definitions.

Panel A: first-month basis, EFG, and third-month spread

Dependent variables		$b_{t,1}$			$EFG_{t,1}$	$b_{t,3}$
	(1)	(2)	(3)	(4)	(5)	(6)
r_{bmk}	-23.88					
	(22.60)					
σ^2_{bmk}	4.64					
	(4.02)					
OI	33.71*					
	(19.55)					
S	-1.13					
	(1.09)					
Liq	25.77					
	(28.31)					
Treat	2.21					
	(1.95)					
Post	-1.83				**2.24***	
	(1.76)				(1.27)	
Treat × Post	**3.87****	**3.62****	**3.12****	**2.75****		1.32
	(1.86)	(1.78)	(1.55)	(1.32)		(1.19)
Asset FE	No	Yes	Yes	No	No	Yes
Time FE	No	Yes	Yes	No	No	Yes
Asset × Time FE	No	No	Yes	No	No	Yes
Observations	1,189	1,198	1,198	1,198	599	1,198
R^2	0.40	0.44	0.44	0.41	0.15	0.13

Panel B: second-month spread, EFG, and fourth-month spread

Dependent variables		$b_{t,2}$			$EFG_{t,2}$	$b_{t,4}$
	(1)	(2)	(3)	(4)	(5)	(6)
r_{bmk}	9.30					
	(27.65)					
σ^2_{bmk}	4.91					
	(7.10)					
OI	7.12					
	(6.41)					
S	-4.41					
	(4.22)					
Liq	2.04					
	(7.14)					
Treat	-5.12					
	(4.97)					
Post	-1.10				**1.16*****	
	(2.04)				(0.35)	
Treat × Post	**1.74****	**1.78****	**1.83****	**1.84****		0.07
	(0.85)	(0.87)	(0.89)	(0.89)		(0.11)
Asset FE	No	Yes	Yes	No	No	Yes
Time FE	No	Yes	Yes	No	No	Yes
Asset × Time FE	No	No	Yes	No	No	Yes
Observations	1,185	1,197	1,197	1,197	592	1,197
R^2	0.24	0.35	0.36	0.34	0.13	0.11

1.11 Appendix A – Model of ETFs and arbitrageurs

In this section, I present a theoretical framework to explain the main empirical findings. I show that ETF demand has an impact on prices in a finite-period economy with a constant absolute risk aversion (CARA) framework for arbitrageurs.

1.11.1 Assets

There is a riskless asset with an exogenous return $r_f > 0$ and two risky assets in zero net supply: futures with a price F_{t,T_i} and a hedge asset with a price H_{t,T_i}. Both risky assets can have a maturity of one or two months: $T_i = \{T_1, T_2\}$ denotes the maturity of the asset. The hedge asset is similar to the futures contract but traded on a market without ETF demand. For example, for VIX, it is the replicating portfolio of options. Assets' payoffs (dividends) each period before maturity ($t < T_i$) are jointly normal:

$$d_{F_i,t} \sim \mathcal{N}(d_i, \sigma^2_{F_i}),$$

$$d_{H_i,t} \sim \mathcal{N}(d_i, \sigma^2_{H_i}),$$

$$corr(d_{F_i,t}, d_{H_i,t}) = \rho_i.$$

At maturity, both assets pay the same amount $d_{F_i,T_i} = d_{H_i,T_i} = S_{T_i}$, where S_{T_i} is the exogenous spot price of the asset underlying the futures contract (I do not model its dynamics for simplicity). Payoffs can be interpreted as the carry of the futures contract for period t. I take the hedge asset's price to be exogenous: $H_{t,T_i} = \sum_{j=t+1}^{T_i} \frac{\mathbb{E}(d_{H_i,j})}{(1+r_f)^j}$.

1.11.2 ETFs

ETFs passively follow their benchmark. There are N ETFs, each with leverage L_j. For simplicity and tractability, I do not model the dynamics of ETF demand $D^{all}_{t,i}$ as a function of realized dividends of the first-month and the second-month futures contract in equilibrium, and treat it as a demand shock similar to Garleanu et al. (2009). The main purpose of the model is to illustrate that demand shocks from ETFs can bias futures prices away from fundamentals but not to solve for the non-linear and less tractable equilibrium dynamics of ETF rebalancing and dividend effects between the two futures contracts.

1.11.3 Arbitrageurs

Arbitrageurs have CARA preferences over wealth in the next period. Consider first a two-period setup: arbitrageurs choose optimal positions at time t, and consume at time $t+1$ when the assets pay off. For simplicity, I derive all results for the general futures contract with maturity T_i.

The maximization problem of an arbitrageur is:

$$\max_{x_{F_i,t},x_{H_i,t}} -\mathbb{E}_t(e^{-\gamma W_{t+1}}) \tag{1.19}$$

subject to a budget constraint:

$$W_{t+1} = x_{F_i,t}d_{F_i,t+1} + x_{H_i,t}d_{H_i,t+1} + (1+r_f)\left(W_t - x_{F_i,t}F_{t,T_i} - x_{H_i,t}H_{t,T_i}\right), \tag{1.20}$$

where W_t is the wealth of an arbitrageur, γ is the risk aversion parameter, $x_{F_i,t}$ is the position in the futures, $x_{H_i,t}$ is the position in the hedge asset. Using log-normality and the fact that $H_{t,T_i} = \frac{\mathbb{E}_t(d_{H_i,t+1})}{1+r_f}$, the maximization problem is equivalent to

$$\begin{aligned} &\max_{x_{F_i,t},x_{H_i,t}} \mathbb{E}_t(W_{t+1}) - \frac{\gamma}{2}\mathrm{Var}_t(W_{t+1}) \\ \iff &\max_{x_{F_i,t},x_{H_i,t}} x_{F_i,t}(\mathbb{E}_t(d_{F_i,t+1}) - F_{t,T_i}(1+r_f)) - \\ &- \frac{\gamma}{2}(x^2_{F_i,t}\sigma^2_{F_i} + x^2_{H_i,t}\sigma^2_{H_i} + 2x_{F_i,t}x_{H_i,t}\rho_i\sigma_{F_i}\sigma_{H_i}). \end{aligned} \tag{1.21}$$

Taking first-order condition (FOC) with respect to $x_{H_i,t}$ gives $x^*_{H_i,t} = \frac{-\rho_i\sigma_{F_i}}{\sigma_{H_i}}x_{F_i,t}$. Substituting back in (1.21) and taking FOC with respect to $x_{F_i,t}$ yields

$$x^*_{F_i,t} = \frac{\mathbb{E}_t(d_{F_i,t+1}) - F_{t,T_i}(1+r_f)}{\gamma\sigma^2_{F_i}(1-\rho^2_i)}.$$

Now, using market clearing ($x^*_{F_i,t} = -D^{all}_{t,i}$), I obtain the equilibrium futures price:

$$F_{t,T_i} = \frac{\mathbb{E}_t(d_{F_i,t+1})}{1+r_f} + \frac{\gamma\sigma^2_{F_i}(1-\rho^2_i)}{1+r_f}D^{all}_{t,i}, \tag{1.22}$$

or

$$F_{t,T_i} = \underbrace{\frac{d_i}{1+r_f}}_{\textit{fundamental value} \, = \, H_{t,T_i}} + \underbrace{\frac{\gamma\sigma^2_{F_i}(1-\rho^2_i)}{1+r_f}D^{all}_{t,i}}_{\textit{ETF price impact}}. \tag{1.23}$$

The EFG is explicitly given by $F_{t,T_i} - H_{t,T_i} = \frac{\gamma\sigma^2_{F_i}(1-\rho^2_i)}{1+r_f}D^{all}_{t,i}$ and is proportional to ETF demand as $\frac{\gamma\sigma^2_{F_i}(1-\rho^2_i)}{1+r_f} \geq 0$. Price impact $\phi_i = \frac{\partial F_{t,T_i}}{\partial D^{all}_{t,i}}$ is larger when arbitrageurs are more risk-averse, when the variance of the futures payoff is larger, and when H is a worse hedge ($|\rho_i| < 1$).

Why do arbitrageurs optimize over wealth next period and cannot wait until the maturity of the futures contract T_i? This assumption is similar to Shleifer and Vishny (1997), where arbitrageurs care about next period returns because they could lose market share, or face outflows if the price discrepancy widens. The ETF framework studied here is an even more pronounced example of a setting where arbitrageurs mechanically have to maximize over wealth next period and cannot wait until maturity. Due to calendar rebalancing, ETFs gradually roll out of the futures positions *before* expiration, and always hold zero contracts at maturity. Therefore, if arbitrageurs take the opposite positions, they also have to gradually close those positions before expiration. Thus, they face the risk of widening price gaps. Another reason for the short horizon of arbitrageurs is the specifics of trading futures contracts. Since futures positions are collateralized and marked-to-market on a daily basis, arbitrageurs could get financially constrained at times when they are short the futures and the EFG increases, if they do not hedge the position perfectly.

1.11.3.1 Financial constraints

Gromb and Vayanos (2002) show that when arbitrageurs' financial constraint binds, arbitrageurs cannot fully absorb demand shocks, and price gaps can persist for more than two periods. To illustrate this idea, consider period $t-1$ and suppose that as of time $t-1$, d_i and $D^{all}_{t,i}$ are uncertain:

$$d_i \sim \mathcal{N}(d_i, \sigma^2_{F_i}),$$

$$D^{all}_{t,i} \sim \mathcal{N}(\bar{y}_i, \sigma^2_y),$$

$$corr(d_i, D^{all}_{t,i}) = \rho_y.$$

The fact that $D^{all}_{t,i}$ and the mean of the dividend are correlated captures the idea that ETF demand shocks are correlated with realized returns. The price as of time $t-1$ is then:

$$F_{t-1,T_i} = \frac{\mathbb{E}_{t-1}(F_{t,T_i})}{1+r_f} = \frac{d_i}{(1+r_f)^2} + \frac{\gamma\sigma^2_{F_i}(1-\rho_i^2)}{(1+r_f)^2}\mathbb{E}_{t-1}(D^{all}_{t,i}). \tag{1.24}$$

Thus, the futures price before T_i reflects not only the expected future value of dividends (the fundamental value), but also the expected price impact of ETFs. Equation (1.24) shows that before T_i, the futures price can be different from the fundamental value due to the risk of bearing ETF demand shocks. The expectations of higher ETF demand shock in the future would push *today's* price away from the fundamental value and give rise to the EFG. The variance of the price is also greater due to the non-fundamental component:

$$Var_{t-1}(F_{t,T_i}) = \frac{1}{(1+r_f)^4}(\underbrace{\sigma^2_{F_i}}_{\text{fundamental var.}} \underbrace{+\gamma^2\sigma^4_{F_i}(1-\rho_i^2)^2\sigma_y^2 + 2\gamma(1-\rho_i^2)\sigma^2_{F_i}\rho_y\sigma_{F_i}\sigma_y}_{\text{non-fundamental var.}}). \tag{1.25}$$

Higher variance would decrease the correlation between the futures contract and the hedge asset, and also increase the volatility of the arbitrageur's return. If arbitrageurs have to meet a financial constraint similar to that expressed in Gromb and Vayanos (2002), higher variance could lead to amplification of price changes. Typically, the financial constraint would be increasing in the demand shock and the volatility of the price.[37] Gromb and Vayanos (2002) show that, due to arbitrageurs' financial constraints, the price gap ($EFG_t = F_{t,T_i} - H_{t,T_i}$) can persist for several periods. The equilibrium outcome depends on whether arbitrageurs close the gap before their financial constraint binds. If the arbitrageurs' initial wealth is not large enough, the financial constraint binds in all periods and arbitrageurs do not fully absorb ETF demand shocks. As a result, the EFG persists over time and is closed only at maturity T_i. The fact that the two-months EFG increases when the TED spread goes up (Table 1.4) is evidence that financial constraints are related to the EFG.

With a financial constraint in place, the risk of trading against ETF demand is amplified. Demand pressure then enters into the pricing kernel similar to Garleanu et al. (2009). Large and unexpected demand shock can tighten the financial constraint and force arbitrageurs to liquidate positions at a loss. Take an arbitrary τ, $0 < \tau < T_i$. Suppose $EFG_{\tau-1} > 0$ and arbitrageurs enter period τ with a position $x_{F_i,\tau-1} < 0$ in the futures contract. Then, suppose that there is a large and positive realization of the fundamental dividend $d_{F_i,\tau}$. Large and positive $d_{F_i,\tau}$ would increase the ETF demand shock $D^{all}_{\tau,i}$ mostly due to leverage rebalancing. That would push up $F_{\tau,i}$ further away from $H_{\tau,i}$ and could make the arbitrageurs' financial constraint binding since arbitrageurs lose money (they sold the futures). In that case, arbitrageurs would have to liquidate their positions to meet the constraint. The liquidation would decrease arbitrageurs' wealth and increase the EFG further. Arbitrageurs' selloff could amplify the price change, which would increase the leverage-induced demand by ETFs and further tighten the budget constraint. Thus, the presence of ETFs in a given market creates a potential feedback channel and pushes *today's* price away from the fundamental value,

[37] This is the case in Gromb and Vayanos (2002), Gromb and Vayanos (2018) and several other papers. With normally distributed dividends, risk-free loans are not possible, so one can use a similar constraint that is increasing in volatility, instead of using the maximum possible loss for an arbitrageur.

increasing the futures premium. Compared to traditional amplification mechanisms (see, e.g., Gromb and Vayanos, 2002), the feedback channel in the case with ETFs arises not only because of financial constraints but also due to mechanically induced short-horizon momentum trades by leveraged ETFs.

1.11.3.2 Transaction costs

Empirically, another risk for arbitrageurs in the VIX market is the risk of widening bid-ask spreads. As seen from Figure 1.25, the bid-ask spreads of the synthetic futures contract are larger than those of the traded VIX futures contract. Therefore, hedging the futures position involves significant transaction costs. Moreover, bid-ask spreads spike during periods of high volatility. This would decrease the correlation between the prices of the two contracts and amplify arbitrageurs' losses if they were to close the position.

Assume that, in order to hedge her position in F_{t,T_i} with a position in H_{t,T_i}, the arbitrageur incurs a transaction cost c. The cost is paid both when the arbitrageur buys the asset, and when she sells it (e.g., bid-ask spreads). Without loss of generality and for analytical tractability, I assume that transaction costs are paid per unit of F_{t,T_i}. The maximization problem of an arbitrageur is then slightly changed:

$$\max_{x_{F_i,t},x_{H_i,t}} x_{F_i,t}(\mathbb{E}_t(d_{F_i,t+1}) - F_{t,T_i}(1+r_f)) - \frac{\gamma}{2}(x^2_{F_i,t}\sigma^2_{F_i} + x^2_{H_i,t}\sigma^2_{H_i} + \\ + 2x_{F_i,t}x_{H_i,t}\rho_i\sigma_{F_i}\sigma_{H_i}) - c|x_{F_i,t}|. \tag{1.26}$$

Taking first-order condition (FOC) in the two cases ($x_{F_i,t} > 0$ and $x_{F_i,t} \leq 0$) and using market clearing), I obtain the equilibrium futures price:

$$F_{t,T_i} = \frac{d_i}{1+r_f} + \frac{\gamma\sigma^2_{F_i}(1-\rho^2_i)}{1+r_f}D^{all}_{t,i} - c, \text{if } D^{all}_{t,i} < 0,$$

$$F_{t,T_i} = \frac{d_i}{1+r_f} + \frac{\gamma\sigma^2_{F_i}(1-\rho^2_i)}{1+r_f}D^{all}_{t,i} + c, \text{if } D^{all}_{t,i} \geq 0.$$

The fact that the one-month EFG increases when bid-ask spreads rise (Table 1.4) is evidence that transaction costs affect the EFG.

The results from this section show that arbitrageurs who trade against ETFs would charge a premium for bearing unhedgeable risks. Investors can hedge fundamental shocks to the futures contract by trading a similar futures contract with no ETF demand (e.g., a replicating option portfolio for the case of VIX, or non-US gas futures for gas). If, however, the hedge asset is not perfectly correlated with the futures contract ($|\rho_i| \neq 1$), the price impact will be positive. Hedging the futures exposure with H_{t,T_i} is not a pure textbook arbitrage since it entails the risk of widening price discrepancy between the two securities before maturity. Arbitrageurs absorb ETF demand as a liquidity service, rather than to meet their own portfolio needs, and are compensated in equilibrium for liquidity provision. The compensation comes in the form of a temporary price impact that pushes prices away from fundamental values.

1.12 Appendix B – derivations and additional robustness checks

1.12.1 Forward VIX derivations

$R_{T_1\to T_2}$ is the gross return on the S&P 500 Index, $R_{f,t\to T} = e^{r_f(T-t)}$ is the constant gross risk-free rate. Using $R_{T_1\to T_2} = \frac{R_{t\to T_2}}{R_{t\to T_1}}$ and $\mathbb{E}_t^{\mathrm{Q}} R_{T_1\to T_2} = R_{f,T_1\to T_2}$:

$$\begin{aligned}
&\mathbb{E}_t^{\mathrm{Q}}(VIX^2_{T_1\to T_2}) = \frac{2}{T_2-T_1}\mathbb{E}_t^{\mathrm{Q}}\Big(\log \mathbb{E}_t^{\mathrm{Q}} R_{T_1\to T_2} - \mathbb{E}_t^{\mathrm{Q}} \log R_{T_1\to T_2}\Big)\\
&= \frac{2}{T_2-T_1}\mathbb{E}_t^{\mathrm{Q}}\Big((T_2-t)r_f - (T_1-t)r_f - (\mathbb{E}_t^{\mathrm{Q}} \log R_{t\to T_2} - \mathbb{E}_t^{\mathrm{Q}} \log R_{t\to T_1})\Big)\\
&= \frac{2}{T_2-T_1}\Big(\log \mathbb{E}_t^{\mathrm{Q}} R_{t\to T_2} - \mathbb{E}_t^{\mathrm{Q}} \log R_{t\to T_2} - (\log \mathbb{E}_t^{\mathrm{Q}} R_{t\to T_1} - \mathbb{E}_t^{\mathrm{Q}} \log R_{t\to T_1})\Big)\\
&= \frac{1}{T_2-T_1}\Big((T_2-t)VIX^2_{t\to T_2} - (T_1-t)VIX^2_{t\to T_1}\Big).
\end{aligned} \tag{1.27}$$

1.12.2 Calculating $\mathrm{Var}_t^{\mathrm{Q}}(VIX_{T_1\to T_2})$

Based on a result from Breeden and Litzenberger (1978), the price of any function $g(S_T)$ satisfies:

$$\frac{1}{R_{f,t\to T}}E_t^{\mathrm{Q}}(g(S_T)) = \frac{1}{R_{f,t\to T}}g(E_t^{\mathrm{Q}}(S_T)) + \int_{K=0}^{F_{t,T}} g''(K)put_{t,T}(K)dK + \int_{K=F_{t,T}}^{\infty} g''(K)call_{t,T}(K)dK. \tag{1.28}$$

Take $g(S_T) = S_T^2$, then:

$$\begin{aligned}
&\frac{1}{R_{f,t\to T}}\Big(E_t^{\mathrm{Q}}(S_T^2) - (E_t^{\mathrm{Q}}(S_T))^2\Big) = 2\left(\int_{K=0}^{F_{t,T}} put_{t,T}(K)dK + \int_{K=F_{t,T}}^{\infty} call_{t,T}(K)dK\right).\\
&\mathrm{Var}_t^{\mathrm{Q}}(S_T) = 2R_{f,t\to T}\left(\int_{K=0}^{F_{t,T}} put_{t,T}(K)dK + \int_{K=F_{t,T}}^{\infty} call_{t,T}(K)dK\right).
\end{aligned} \tag{1.29}$$

Another way to get the same equation is by using $R_T = \frac{S_T}{S_t}$ in equation (11) of Martin (2017).

Take $T = T_1,\ T_2 = T_1 + 30\ days,\ S_{T_1} = VIX_{T_1\to T_2}$, then:

$$\mathrm{Var}_t^{\mathrm{Q}}(VIX_{T_1\to T_2}) = 2R_{f,t\to T_1}\left(\int_{K=0}^{F_{t,T_1}} put_{t,T_1}(K)dK + \int_{K=F_{t,T_1}}^{\infty} call_{t,T_1}(K)dK\right), \tag{1.30}$$

where F_{t,T_1} – time t's price of a futures on $VIX_{T_1\to T_2}$ with maturity T_1,
$put_{t,T_1}(K)$ – time t's price of a put option on $VIX_{T_1\to T_2}$ with maturity T_1 and strike K,

$call_{t,T_1}(K)$ – time t's price of a call option on $VIX_{T_1 \to T_2}$ with maturity T_1 and strike K.
The underlying asset of the call and put options is the futures on the VIX since at maturity T_1, $F_{T_1,T_1} = S_{T_1} = VIX_{T_1 \to T_2}$, and there are no dividends.

1.12.3 Details on the synthetic VIX futures calculation

$VIX^2_{t \to T_1}$, and $VIX^2_{t \to T_2}$ are calculated using the exact same procedure as outlined in the CBOE VIX White Paper. The correlation between the calculated VIX and the CBOE-quoted one for a maturity of one month is 99.8%.

Empirically, sometimes no S&P 500 Index options expire at the exact same time as VIX futures. VIX futures typically expire in the morning of the Wednesday before the third Friday of the month (the settlement value is calculated using special opening quote values). There are always options expiring 30 days after, since these are used to calculate the settlement price of VIX. S&P 500 Index options (weeklys) also cease trading on Wednesday, but are p.m.-settled and expire at 4:00 p.m. I deal with this issue in several ways. First, I compute $\mathbb{E}^{Q}_t(VIX^2_{T_1 \to T_2})$ using the evening quotes for options. Second, I interpolate in the volatility space to get prices of options expiring in the morning, and compute $\mathbb{E}^{Q}_t(VIX^2_{T_1 \to T_2})$ using these prices. With both approaches, the estimates for the EFG shared similar patterns as in Figure 1.5. Sometimes, especially before weeklys were introduced, there were no S&P 500 Index options expiring at the same time as VIX futures. For those, I interpolate $VIX^2_{t \to T_1}$ using the nearest (usually within 1–2 days) expiring option contracts.

For the computation of $\text{Var}^{Q}_t(VIX_{T_1 \to T_2})$, I need the futures price F_{t,T_1}. I tested three different ways to estimate it. First, by finding the strike for which put and call prices are closest (analogous to the calculation of VIX). Second, by implementing an iterative procedure to find $F_{t,T_1} = \mathbb{E}^{Q}_t(VIX_{T_1 \to T_2})$ that makes equation (1.5) hold. Third, by using the traded VIX futures price. The main results were similar with each of the three estimates.

A potential concern for the computation of $VIX^2_{T_1 \to T_2}$ is the existence of discretization errors. As a robustness check, similar to Aït-Sahalia et al. (2018), I calculated forward VIX by interpolating the volatility surface and finding option prices for all strikes following the methodology of Carr and Wu (2009). The EFG was less volatile and smaller in magnitude, on average, but the main results of the paper were unchanged. As another robustness test, I calculated the EFG using minimum price instead of mid-point price for each option in equation (1.7) similar to Kadan and Tang (2020). The main results of the paper were unchanged but the EFG was underestimated even more using this approach.

Another concern is that the specific rules used by the CBOE for selecting options to calculate VIX could lead to instabilities in the intra-day value of the index, especially during extreme market movements as noted by Andersen et al. (2011). However, I use the CBOE methodology to compute forward VIX on a daily basis. These instabilities should be less severe than on an intra-day basis (see, e.g., Aït-Sahalia et al., 2018).

1.12.4 Variance and volatility

One concern with the replicating portfolio that gives the option-implied VIX futures price is that we can replicate variance by using it, but the VIX futures settles at volatility. However, by choosing the number of contracts in a smart way, we can easily eliminate the risk and make sure to profit from the EFG. Denote the portfolio under the square root in equation (1.5) $X_t^2 = \mathbb{E}_t^{\mathrm{Q}}(VIX^2_{T_1\to T_2}) - \mathrm{Var}_t^{\mathrm{Q}}(VIX_{T_1\to T_2})$. For example, suppose that $F_{t,T_1} > \mathbb{E}_t^{\mathrm{Q}}(S_{T_1}) = \sqrt{X_t^2}$ so the EFG is positive (this is the case for the most of the sample). We can make sure to profit from the gap by selling $2X_t$ units of F_{t,T_1} and buying one unit of X_t^2 at time t. At maturity T_1, $F_{T_1,T_1} = X_{T_1} = S_{T_1}$. Assume, for simplicity that $r_f = 0$ (with $r_f > 0$, the profit is even larger). The total profit from t to T_1 is then $2X_tF_{t,T_1} - X_t^2 + S_{T_1}^2 - 2X_tS_{T_1} = 2X_tF_{t,T_1} - 2X_t^2 + S_{T_1}^2 - 2X_tS_{T_1} + X_t^2 = 2X_t(F_{t,T_1} - X_t) + (S_{T_1} - X_t)^2 > 0$ as $F_{t,T_1} > X_t$. The logic for the case $F_{t,T_1} < \mathbb{E}_t^{\mathrm{Q}}(S_{T_1})$ is similar.

1.12.5 Derivations of ETF demand decomposition

At time t, the ETF has a dollar position of $L\alpha_tA_t$ in the first-month futures contract and $L(1-\alpha_t)A_t$ in the second. It holds $\frac{L\alpha_tA_t}{F_{t,T_1}}$ units of the first-month contract and $\frac{L(1-\alpha_t)A_t}{F_{t,T_2}}$ units of the second. At time $t+1$, the ETF holds $\frac{L(\alpha_t-\frac{1}{K})A_{t+1}}{F_{t+1,T_1}}$ units of the first-month contract and $\frac{L(1-\alpha_t+\frac{1}{K})A_{t+1}}{F_{t+1,T_2}}$ of the second. The total rebalancing demand (in number of contracts) for the first-month contract from t to $t+1$ is then:

$$
\begin{aligned}
D_{t+1,1} &= \frac{L(\alpha_t - \frac{1}{K})A_{t+1}}{F_{t+1,T_1}} - \frac{L\alpha_t A_t}{F_{t,T_1}} \\
&= \frac{1}{F_{t+1,T_1}}\left(\alpha_t\Big(LA_t(1+Lr_{t+1}) + Lu_{t+1} - LA_t(1+r^{F_1}_{t+1})\Big) - \frac{L}{K}A_{t+1}\right) \\
&= \frac{1}{F_{t+1,T_1}}\left(\alpha_t\Big(LA_t(Lr_{t+1} - r_{t+1} + r_{t+1} - r^{F_1}_{t+1}) + Lu_{t+1}\Big) - \frac{L}{K}A_t(1+Lr_{t+1}) - \frac{L}{K}u_{t+1}\right) \\
&= \frac{1}{F_{t+1,T_1}}\Bigg(- \frac{L}{K}A_t(1+Lr_{t+1}) + \alpha_t A_t L(L-1)r_{t+1} + (\alpha_t - \frac{1}{K})Lu_{t+1} + \\
&\quad + \alpha_t(1-\hat{\alpha}_t)LA_t(r^{F_2}_{t+1} - r^{F_1}_{t+1})\Bigg).
\end{aligned}
\tag{1.31}
$$

I used the fact that $r_{t+1} - r^{F_1}_{t+1} = (1-\hat{\alpha}_t)(r^{F_2}_{t+1} - r^{F_1}_{t+1})$, where $r^{F_1}_{t+1}$ is the net return on the first-month futures contract, $r^{F_2}_{t+1}$ is the net return on the second, $\hat{\alpha}_t = \frac{\alpha_t F_{t,T_1}}{\alpha_t F_{t,T_1} + (1-\alpha_t)F_{t,T_2}}$, and $r_{t+1} = \frac{\alpha_t F_{t+1,T_1} + (1-\alpha_t)F_{t+1,T_2}}{\alpha_t F_{t,T_1} + (1-\alpha_t)F_{t,T_2}} - 1 = \hat{\alpha}_t r^{F_1}_{t+1} + (1-\hat{\alpha}_t)r^{F_2}_{t+1}$ is the net return on the benchmark. In dollar terms, the rebalancing demand is:

$$
\begin{aligned}
D^{\$}_{t+1,1} &= D_{t+1,1}F_{t+1,T_1} \\
&= \underbrace{-\frac{L}{K}A_t(1+Lr_{t+1})}_{\textit{calendar rebalancing}} + \underbrace{\alpha_t A_t L(L-1)r_{t+1}}_{\textit{leverage rebalancing}} + \underbrace{(\alpha_t - \frac{1}{K})Lu_{t+1}}_{\textit{flow rebalancing}} + \\
&\quad + \underbrace{\alpha_t(1-\hat{\alpha}_t)LA_t(r^{F_2}_{t+1} - r^{F_1}_{t+1})}_{\textit{remainder}}.
\end{aligned}
\tag{1.32}
$$

Analogously, the total dollar rebalancing demand for the second futures is:

$$
\begin{aligned}
D^{\$}_{t+1,2} &= D_{t+1,2}F_{t+1,T_2} \\
&= \underbrace{\frac{L}{K}A_t(1+Lr_{t+1})}_{\textit{calendar rebalancing}} + \underbrace{(1-\alpha_t)A_t L(L-1)r_{t+1}}_{\textit{leverage rebalancing}} + \underbrace{(1-\alpha_t + \frac{1}{K})Lu_{t+1}}_{\textit{flow rebalancing}} - \\
&\quad - \underbrace{\hat{\alpha}_t(1-\alpha_t)LA_t(r^{F_2}_{t+1} - r^{F_1}_{t+1})}_{\textit{remainder}}.
\end{aligned}
\tag{1.33}
$$

Equations (1.32) and (1.33) can be rewritten in a way to isolate the terms multiplying L^2. For equation (1.32):

$$
D^{\$}_{t+1,1} = \underbrace{(\alpha_t - \frac{1}{K})}_{\geq 0} A_t L^2 r_{t+1} - \alpha_t A_t L r_{t+1} - \frac{L}{K}A_t + (\alpha_t - \frac{1}{K})Lu_{t+1} +
$$

$$
+\alpha_t(1-\hat{\alpha}_t)LA_t(r^{F_2}_{t+1} - r^{F_1}_{t+1}).
$$

Running the main regressions with the non-linear terms instead of calendar and leverage rebalancing still shows that the predictable part of calendar rebalancing ($\frac{L}{K}A_t$) is statistically significant. Table 1.5 also confirms this fact.

1.12.6 Alternative calculations of leverage rebalancing

A possible concern is that ETFs have higher tracking errors when trying to maintain a constant daily leverage due to non-linear price responses as shown in section 1.6. In particular, the exact size of leverage rebalancing could be slightly different. To alleviate this concern, I tried two alternative ways to estimate leverage rebalancing. First, I calculated it using β from Table 1.6 instead of L – essentially, using the average effective leverage. Second, I estimated leverage rebalancing as the difference between the total rebalancing demand, and the sum of flow rebalancing, remainder, and the part of calendar rebalancing that is linear in L. The latter three components of ETF demand are less likely to be influenced by non-linear price responses. In both alternative approaches, the main results of the paper were unchanged: the EFG was most sensitive to leverage rebalancing, and the proportion of leverage rebalancing increased steadily over time. The relative size of the leverage rebalancing component in the two robustness tests was 3% lower than in the main analysis, on average. Using β from Table 1.6 instead of L for net calendar rebalancing, flow rebalancing, and remainder barely changed the results as the errors canceled out for long and inverse ETFs (in the linear terms). I also performed a robustness test, excluding from the analysis episodes with extremely large tracking errors. The main results of the paper were unchanged. These tests are excluded from the paper for brevity but available on request.

1.12.7 Leverage rebalancing. Derivations of the futures-based benchmark dynamics in a GBM setting

Consider a simple model where the spot follows a geometric Brownian motion (GBM): $\frac{dS_t}{S_t} = \mu_S dt + \sigma dW_t$. W_t is a standard Brownian motion, μ_S and σ are the instantaneous drift and diffusion of the spot.[38] The dynamics of the price of a futures with maturity i is then $\frac{dF_{t,T_i}}{F_{t,T_i}} = (\mu_S - (r_f + y_{t,T_i}))dt + \sigma dW_t$, where y_{t,T_i} is the per dt convenience yield and r_f is the risk-free rate.

[38] They can be non-deterministic but for simplicity I assume they are constant. The results with time-varying volatility are qualitatively similar.

The dynamics of the futures-based benchmark ($\frac{dF_t}{F_t}$) under the physical measure P is then:

$$
\begin{aligned}
&\frac{dF_t}{F_t} = \hat{\alpha}_t \frac{dF_{t,T_1}}{F_{t,T_1}} + (1-\hat{\alpha}_t)\frac{dF_{t,T_2}}{F_{t,T_2}} \\
&= \hat{\alpha}_t \Big((\mu_S - (r_f + y_{t,T_1}))dt + \sigma dW_t \Big) + (1-\hat{\alpha}_t)\Big((\mu_S - (r_f + y_{t,T_2}))dt + \sigma dW_t \Big) \\
&= \underbrace{\Big(\mu_S - (r_f + y_{t,T_2} - \hat{\alpha}_t (y_{t,T_2} - y_{t,T_1})) \Big)}_{\mu} dt + \sigma dW_t \\
&= \mu dt + \sigma dW_t.
\end{aligned}
\tag{1.34}
$$

Then, with a management fee f, the AUM of a leveraged ETF evolve as::

$$
\begin{aligned}
\frac{dA_t}{A_t} &= \hat{L}\frac{dF_t}{F_t} + ((1-\hat{L})r_f - f)dt \\
d\log A_t &= (\hat{L}\mu - \frac{\hat{L}^2\sigma^2}{2} + (1-\hat{L})r_f - f)dt + \hat{L}\sigma dW_t \\
\iff A_T &= A_0 e^{(\hat{L}\mu - \frac{\hat{L}^2\sigma^2}{2} + (1-\hat{L})r_f - f)T + \hat{L}\sigma W_T} \\
\iff A_T &= A_0 (\frac{F_T}{F_0})^{\hat{L}} e^{((-f+(1-\hat{L})r_f)T - \frac{\hat{L}(\hat{L}-1)}{2}\sigma^2 T)}.
\end{aligned}
\tag{1.35}
$$

The last line is obtained from the previous one by adding and subtracting $\frac{\hat{L}\sigma^2}{2}T$ in the power of e.

1.12.8 Derivations of the moments of the short-both strategy

$$
\begin{aligned}
&\mathrm{Var}(r_{SB,T}) = \mathrm{Var}(\frac{A_T}{A_0}) + 2\mathrm{Cov}(\frac{A_T}{A_0}, \frac{B_T}{B_0}) + \mathrm{Var}(\frac{B_T}{B_0}) \\
&= e^{2L\mu T}(e^{L^2\sigma^2 T} - 1) + 2(e^{(L\mu - \frac{L^2\sigma^2}{2})T + L\sigma W_T + (-L\mu - \frac{L^2\sigma^2}{2})T - L\sigma W_T} - e^{L\mu T}e^{-L\mu T}) + \\
&+ e^{-2L\mu T}(e^{L^2\sigma^2 T} - 1) \\
&= (e^{L^2\sigma^2 T} - 1)(e^{2L\mu T} + e^{-2L\mu T}) + 2(e^{-L^2\sigma^2 T} - 1) \approx L^2\sigma^2 T(e^{2L\mu T} + e^{-2L\mu T} - 2).
\end{aligned}
\tag{1.36}
$$

1.12.9 Potential offsetting effect of flows

In cases of extreme market movements, both long and short leveraged ETFs have to trade in the same direction. However, flows could offset these effects as noted in Ivanov and Lenkey (2018). If flows to ETFs were contrarian, this could cancel the effects of leverage rebalancing. For example, suppose the benchmark jumps up on a given day. If investors take out money from long ETFs and invest it in inverse ETFs, that would decrease the flow rebalancing and mitigate the feedback channel from ETF demand to prices. However, Table 1.10 shows that there is limited evidence in favor of this suggestion. Only for gas are flows negatively correlated with changes in leverage rebalancing. For silver and oil, there is even evidence of positive relation between the two components of ETF demand: flows could reinforce leverage rebalancing.

TABLE 1.10: Flows on leverage rebalancing. The table presents regression results of flows on leverage rebalancing. The sample period is from the first leveraged ETF introduction date in a given market to February 2018 (daily frequency).

Dependent variables	VIX		Gas		Silver		Oil	
	Flow, F_{t,T_1}	Flow, F_{t,T_2}	Flow, F_{t,T_1}	Flow, F_{t,T_2}	Flow, F_{t,T_1}	Flow, F_{t,T_2}	Flow, F_{t,T_1}	Flow, F_{t,T_2}
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Lev reb, F_{t,T_i}	-0.01	0.05	-0.30	-0.29*	0.04*	0.00	0.11**	0.00
	(0.04)	(0.05)	(0.25)	(0.18)	(0.02)	(0.01)	(0.05)	(0.00)
Observations	1,687	1,716	1,608	1,687	1,901	2,016	1,595	1,674
R^2	0.01	0.01	0.26	0.12	0.03	0.01	0.03	0.01

1.12.10 Mid-term VIX ETFs

Table 1.11 shows the results of regression (1.3) for mid-term VIX ETFs. These ETFs invest one third of their AUM in the fifth-month futures contract, one third in the sixth-month one, and roll one third from the fourth-month to the seventh-month futures contract on a daily basis. The results from Table 1.11 illustrate that ETF demand has an impact only on the fifth-month spread and the sixth-month spread, mainly due to leverage rebalancing and flows. The effects of calendar rebalancing are less pronounced, probably because mid-term VIX ETFs constitute a lower fraction of open interest compared to short-term VIX ETFs. One standard deviation rise in leverage rebalancing increases the fifth-month spread by 0.25%, and the sixth-month spread by 0.33%.

Table 1.11: Mid-term VIX ETFs. The table presents the results of regression (1.3) for mid-term VIX ETFs. Columns 1–4 correspond to the relative spread of 4–7 months futures. Panel B shows the estimates for ETF demand components. All independent variables are standardized. Refer to Table 1.4 for variable definitions. Daily data, February 2009 – December 2017.

Panel A: Total effect

Dependent variables	$b_{t,4}$, rel (1)	$b_{t,5}$, rel (2)	$b_{t,6}$, rel (3)	$b_{t,7}$, rel (4)
$D^{\$,all}_{t,i}$	-0.17	0.13**	0.12*	0.01
	(0.11)	(0.06)	(0.06)	(0.02)
$b^H_{t,i}$	4.21***	4.35***	4.98***	5.30***
	(1.00)	(0.94)	(0.97)	(1.44)
$\mathrm{r}_{bmk,t}$	-0.06	-0.02	-0.02	-0.06*
	(0.05)	(0.06)	(0.07)	(0.03)
$\sigma^2_{bmk,t}$	-0.18*	-0.15***	-0.12	-0.01
	(0.10)	(0.03)	(0.09)	(0.08)
$\mathrm{OI}_{t,i}$	0.15	0.78***	0.82***	0.45***
	(0.21)	(0.23)	(0.18)	(0.10)
S_t	-2.47***	-1.61***	-0.96***	-0.75***
	(0.27)	(0.30)	(0.28)	(0.18)
$\mathrm{Liq}_{t,i}$	-0.48**	-0.10***	-0.70*	-0.91***
	(0.24)	(0.02)	(0.41)	(0.16)
α_t	-0.18**	-0.03	0.07	0.10
	(0.08)	(0.07)	(0.07)	(0.06)
Observations	1,732	1,720	1,724	1,731
R^2	0.50	0.38	0.38	0.27

Panel B: Split on components

Dependent variables	$b_{t,4}$, rel (1)	$b_{t,5}$, rel (2)	$b_{t,6}$, rel (3)	$b_{t,7}$, rel (4)
Calendar $\mathrm{reb}_{t,i}$	-1.58*			0.26
	(0.84)			(0.81)
$\mathrm{Remainder}_{t,i}$	1.40			-0.24
	(1.38)			(0.27)
Leverage $\mathrm{reb}_{t,i}$	-0.25	0.25**	0.33**	0.17
	(0.31)	(0.12)	(0.13)	(0.28)
Flow $\mathrm{reb}_{t,i}$	-0.30**	0.10*	0.15	-0.81
	(0.12)	(0.06)	(0.151)	(0.63)
Observations	1,697	1,685	1,694	1,692
R^2	0.58	0.41	0.35	0.30

1.13 Appendix C – Additional regressions, figures and tables

FIGURE 1.12: Futures term structure. Summary statistics. S is the spot, F1, F2, ... F9 are the first, second, ... ninth generic futures contracts. The red dot corresponds to the mean, the bold horizontal line corresponds to the median price of a given futures maturity. The bottom and top borders of each bar show the first and third quartile of the price distribution, respectively. The green bars indicate futures maturities in which ETFs invest. The data are at a daily frequency, from the first ETF introduction date in a given market to February 2018.

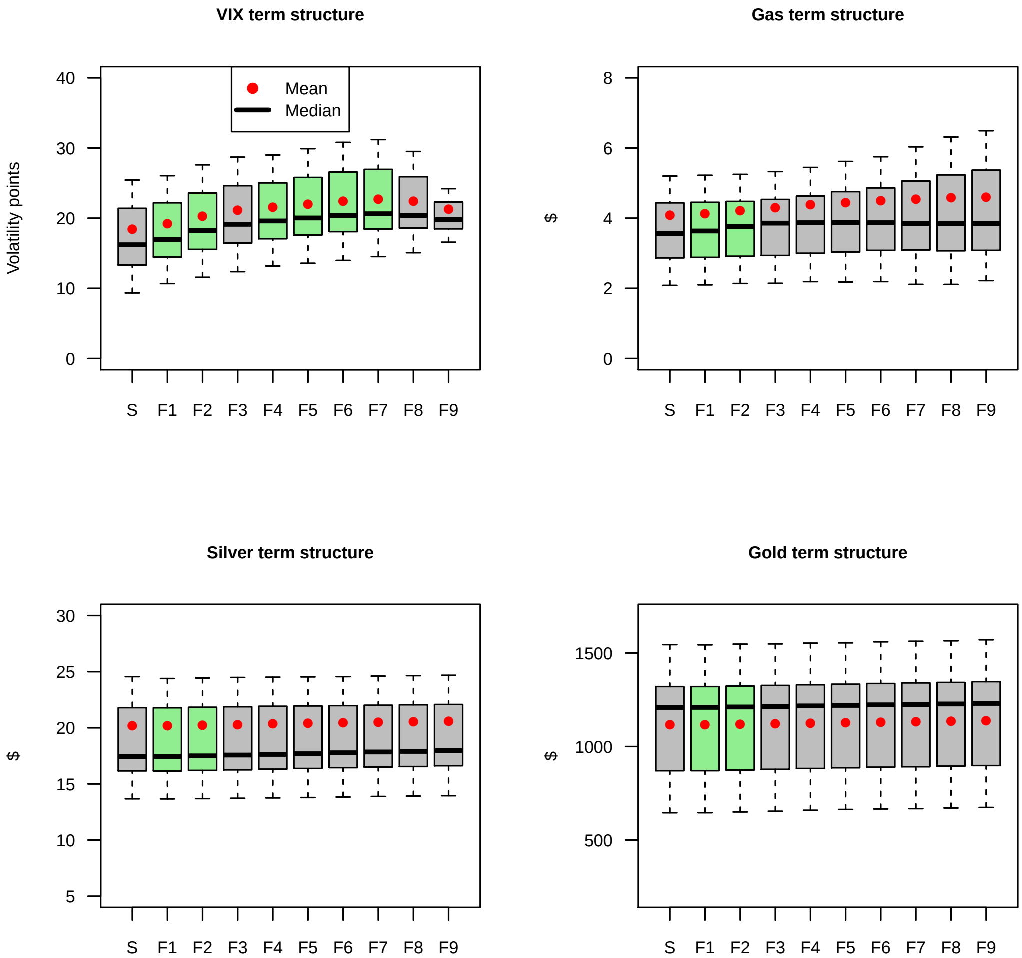

FIGURE 1.13: Realized VIX futures premium before and after ETFs. The chart shows the average size of the VIX futures premium for different maturities before ETFs (including the crisis period, June 2004 – January 2009) and after ETFs (January 2009 – February 2018). The premium is calculated as the annualized net return of a short-seller of a VIX futures $\frac{F_{t,T}-F_{T,T}}{F_{t,T}}$, where T (Term) is maturity.

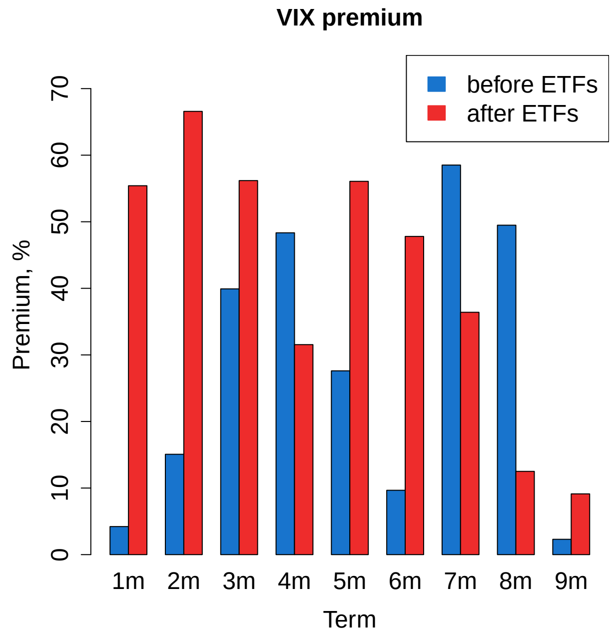

FIGURE 1.14: Open interest dynamics. The left panel shows the dynamics of open interest (monthly averages) in the VIX market for futures with maturities at 1–8 months. The dotted vertical line indicates the date when the first ETF was introduced. The right panel shows the fraction of ETFs in open interest (monthly averages) for futures with maturities at 1–8 months.

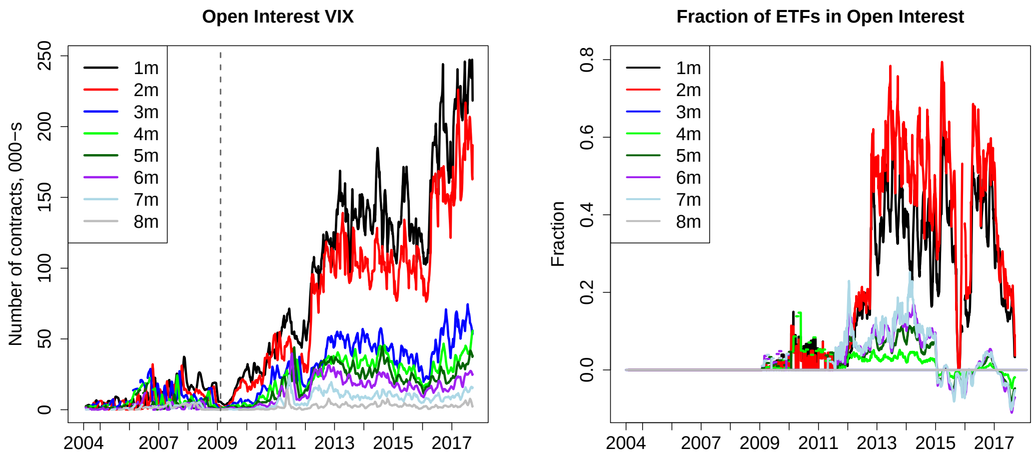

FIGURE 1.15: Predictive power of basis in commodity markets. The figure depicts β_1 and β_2 from regressions (1.1) and (1.2) significant at the 5% level for gas, silver, gold and oil for maturities at 1–9 months.

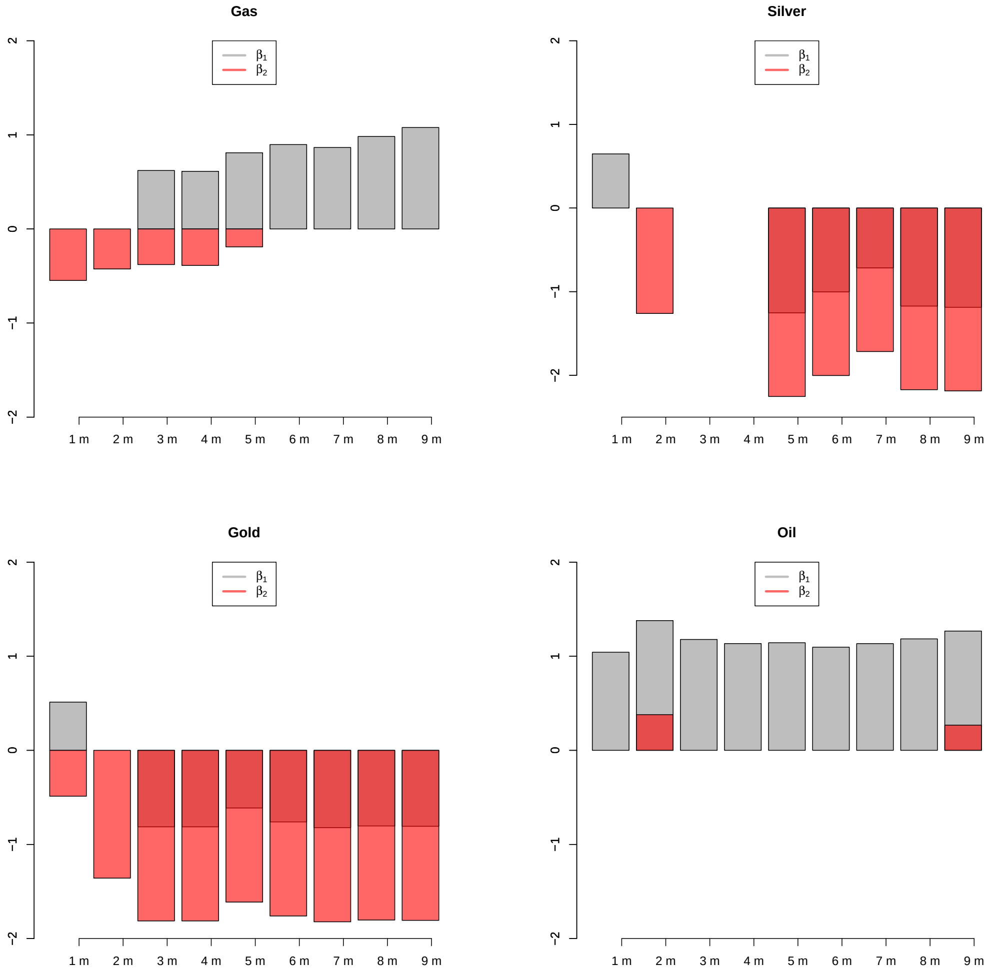

FIGURE 1.16: Monthly ETF futures gap imputed from variance swaps. The left panel shows the dynamics of the EFG for one month, the right panel for two months.

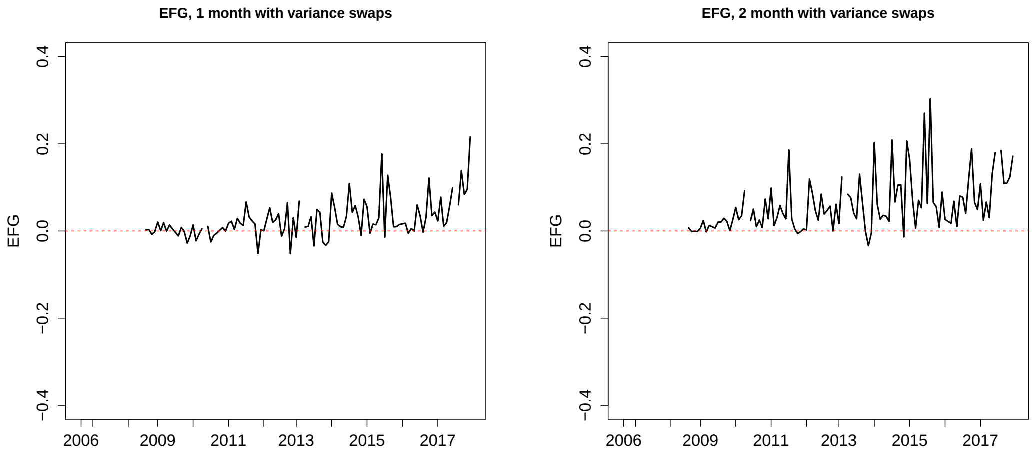

FIGURE 1.17: Hedge fund positions and the EFG. The top panel shows the dynamics of weekly net positions in VIX futures contracts for the main groups of investors. The bottom panel illustrates the weighted-average ETF futures gap between one and two months. The plots suggest hedge funds are mostly short VIX futures contracts after the EFG starts to be positive.

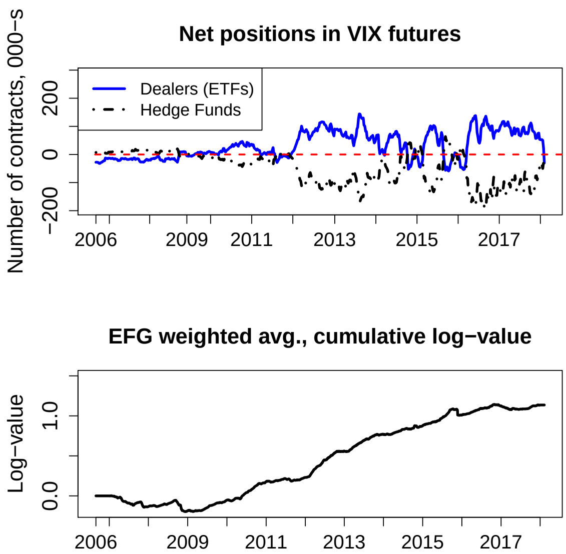

FIGURE 1.18: Example of trading against the leverage rebalancing of ETFs. At time t, the arbitrageur is trading against a double-long and a double-short ETFs, both of which have AUM of 10 (the exposure is 20 for the long and -20 for the short ETF). At time $t+1$, the price of the underlying asset increases from 10 to 11, so both ETFs have to buy, and the arbitrageur acquires a net short position of 8. The left panel shows the marked-to-market loss of an arbitrageur from $t+1$ to $t+2$, if the price drifts up from 10 to 11, to 12 over 3 periods. The right panel shows the marked-to-market profit of an arbitrageur from $t+1$ to $t+2$, if the price reverts back to the initial value (from 10 to 11, to 10) over 3 periods.

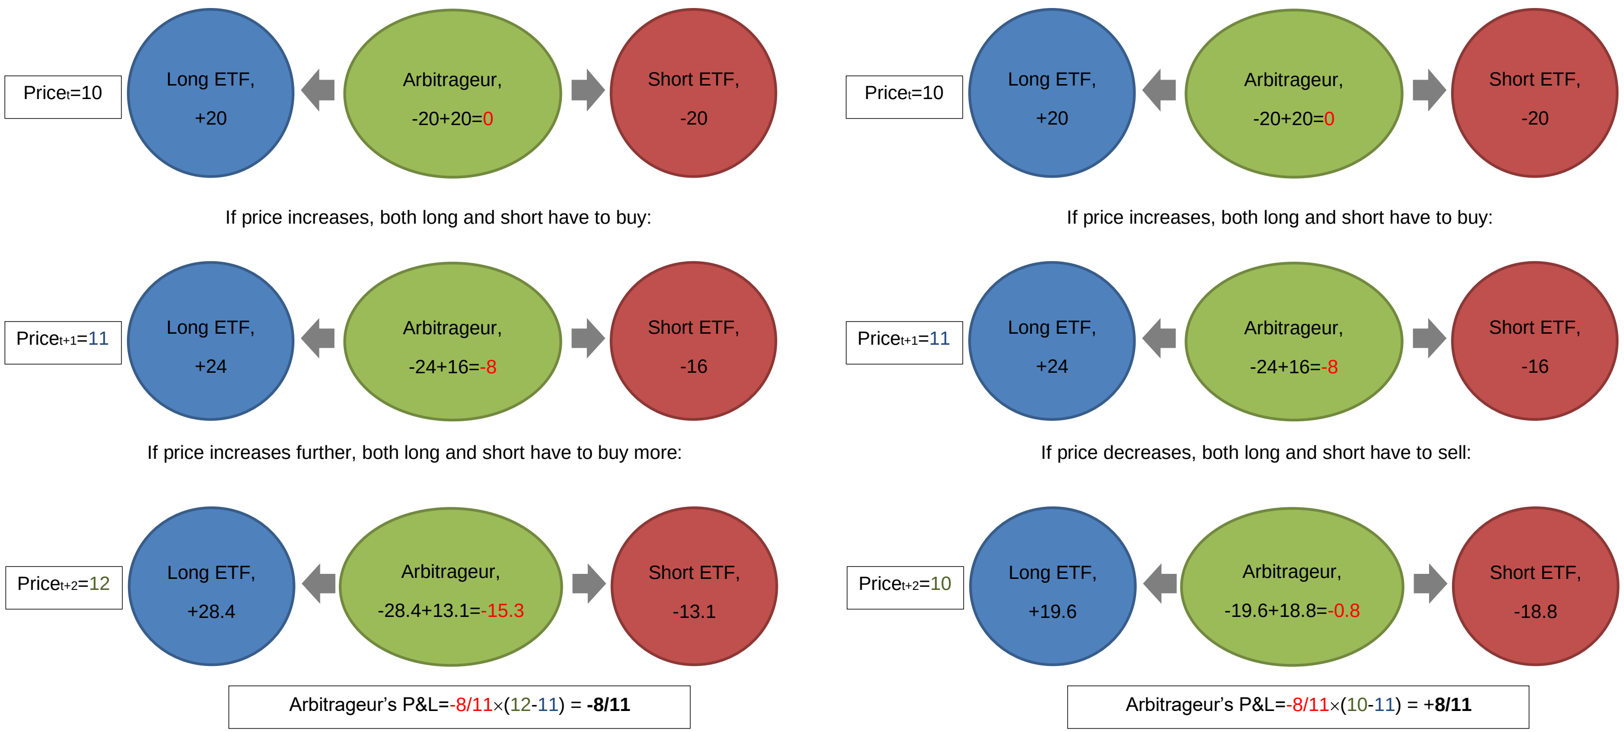

FIGURE 1.19: VIX ETFs: log-values. The figure shows the log-values of a long ($L = 1$) and inverse ($L = -1$) VIX ETFs. The plot illustrates that the two ETFs are not exact mirror images of each other.

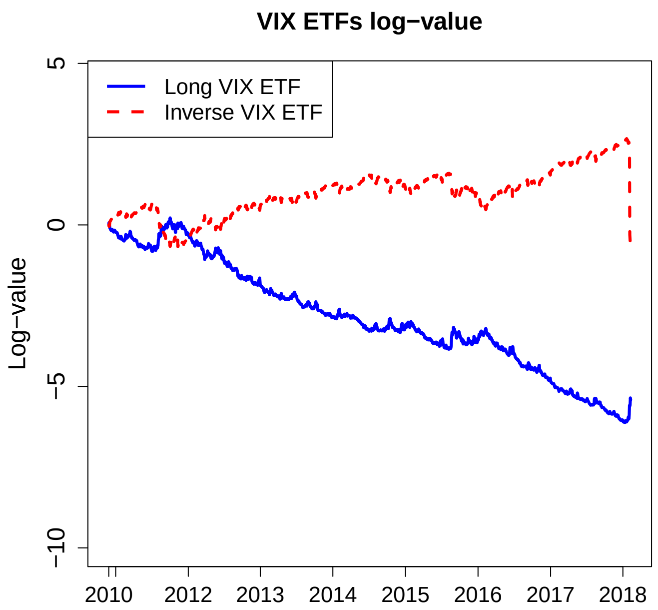

FIGURE 1.20: Theoretical distribution of the short-both strategy return ($r_{SB,T}$) at daily frequency for different values of μ and σ (both annualized). Mean is the average daily return, SR is the annualized Sharpe ratio. Risk-free rate for calculating the Sharpe ratio is set to 1%. Skew is skewness.

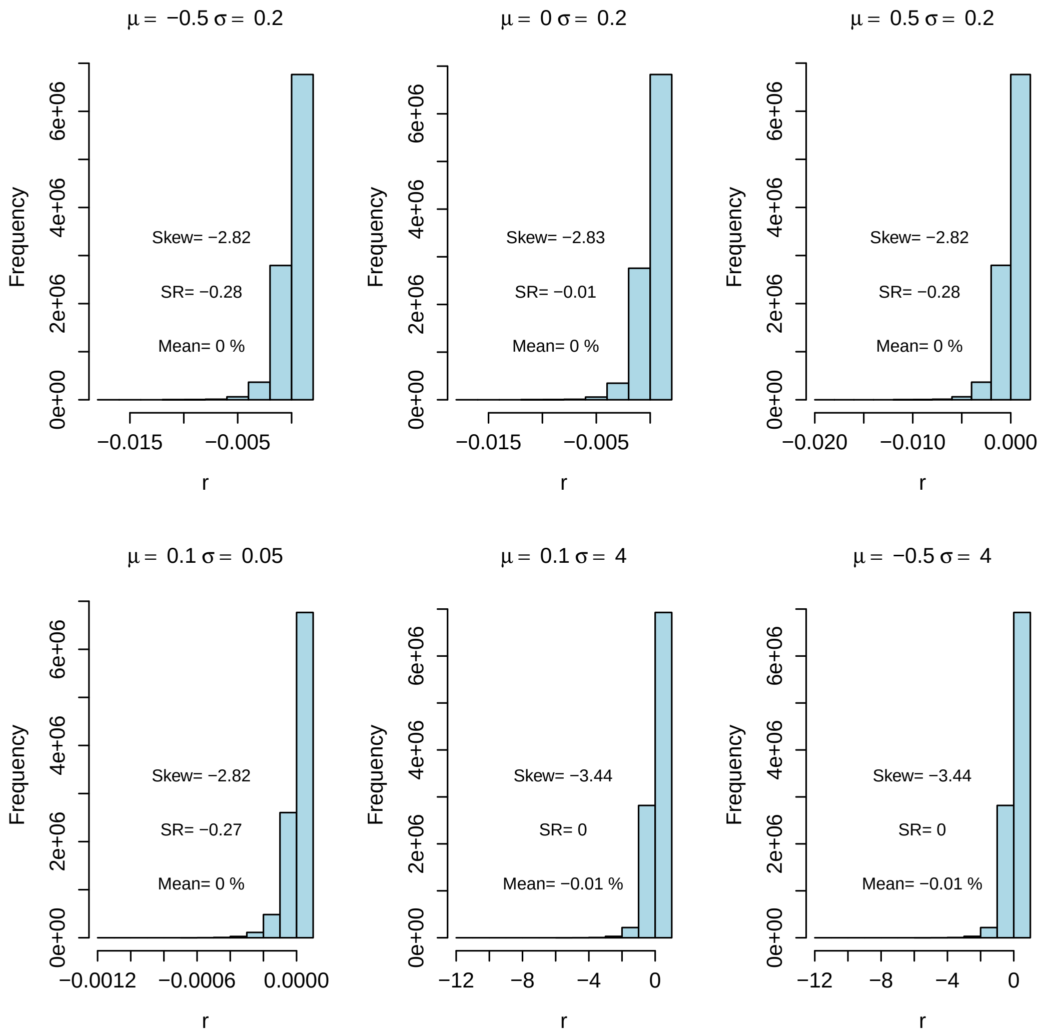

FIGURE 1.21: Empirical distribution of intra-day $r_{SB,T}$ at daily frequency. Mean is the average daily return, SR is the annualized Sharpe ratio. Risk-free rate for calculating the Sharpe ratio is from French's website. Skew is skewness.

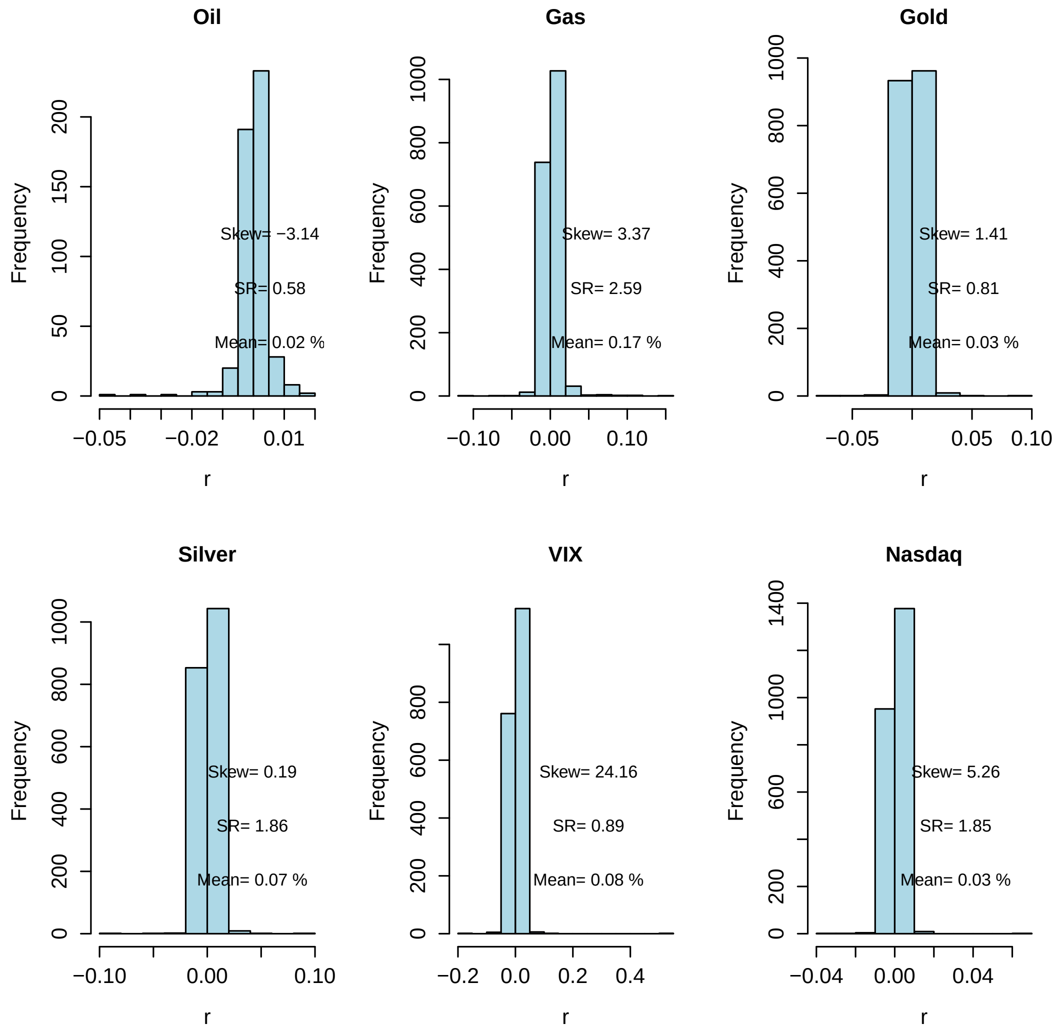

FIGURE 1.22: Empirical distribution of $r_{SB,T}$ at daily frequency using returns from market's open to close. Equity indices.

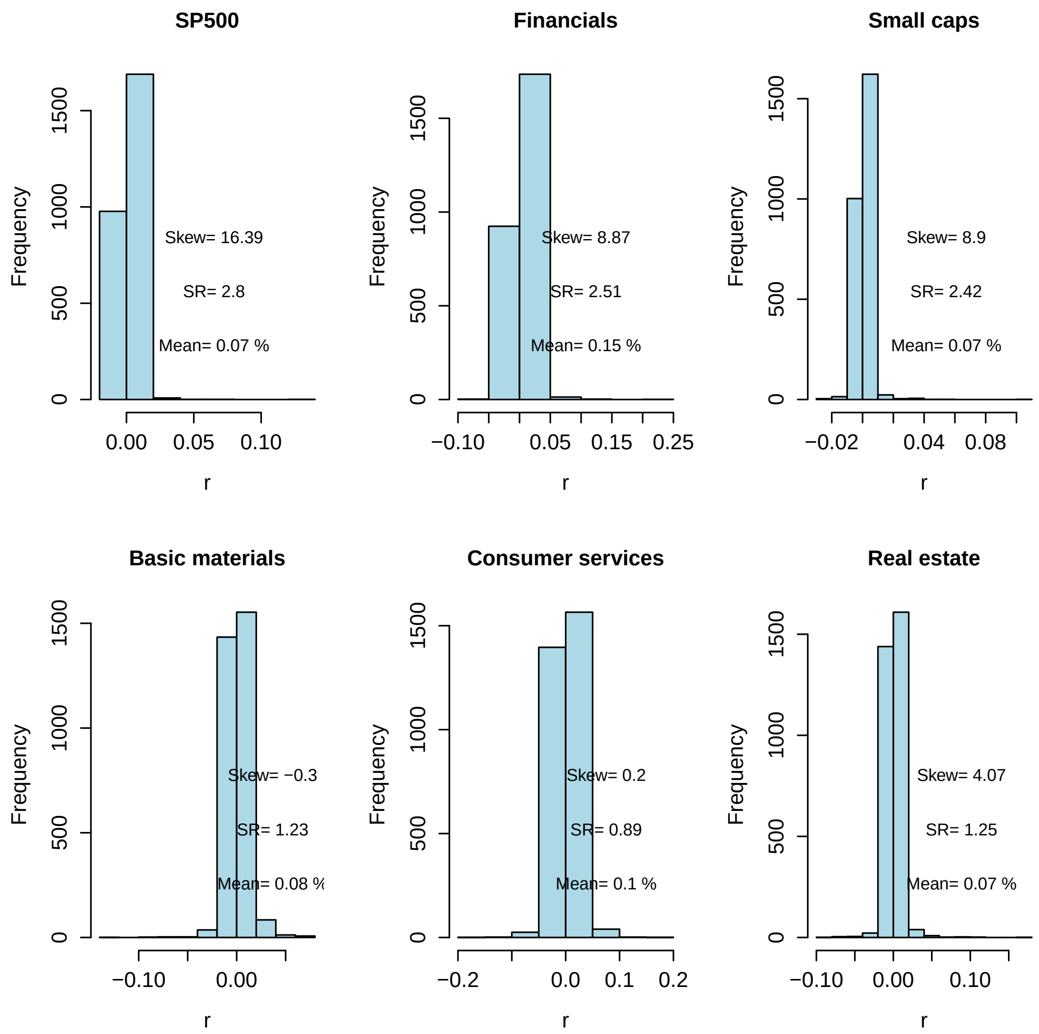

FIGURE 1.23: Empirical distribution of $r_{SB,T}$ at daily frequency using returns from market's open to close. Equity, foreign exchange and bond indices.

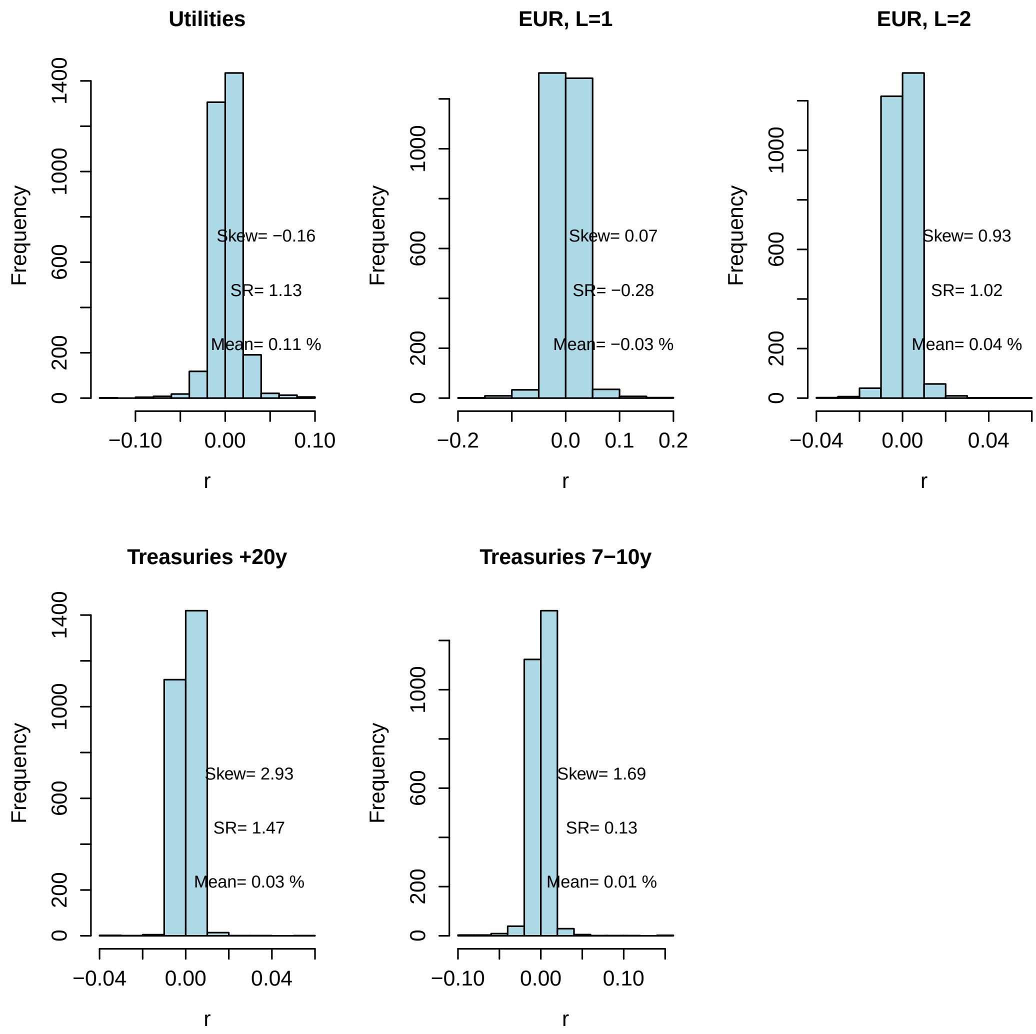

FIGURE 1.24: DD estimates, testing the parallel trends assumption. The figure presents estimates of the DD effect from regression (1.17) with a 60-day window. The x-axis shows days from the event date (introduction of ETF). Red dots illustrate the point estimates, black vertical lines mark 2 standard deviations.

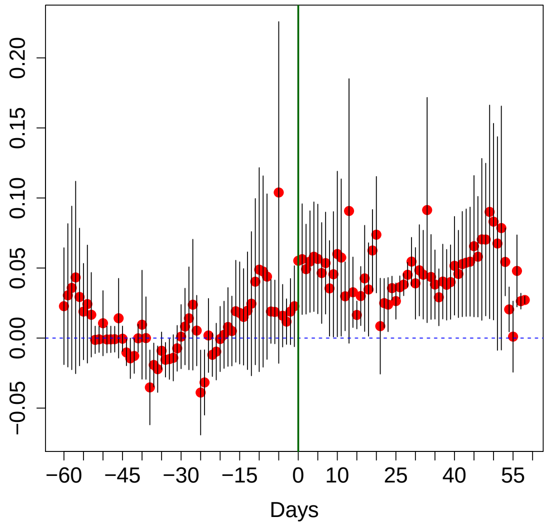

FIGURE 1.25: Bid-ask spreads. The figure shows the dynamics of relative bid-ask spreads of the synthetic futures and traded VIX futures for two months maturity.

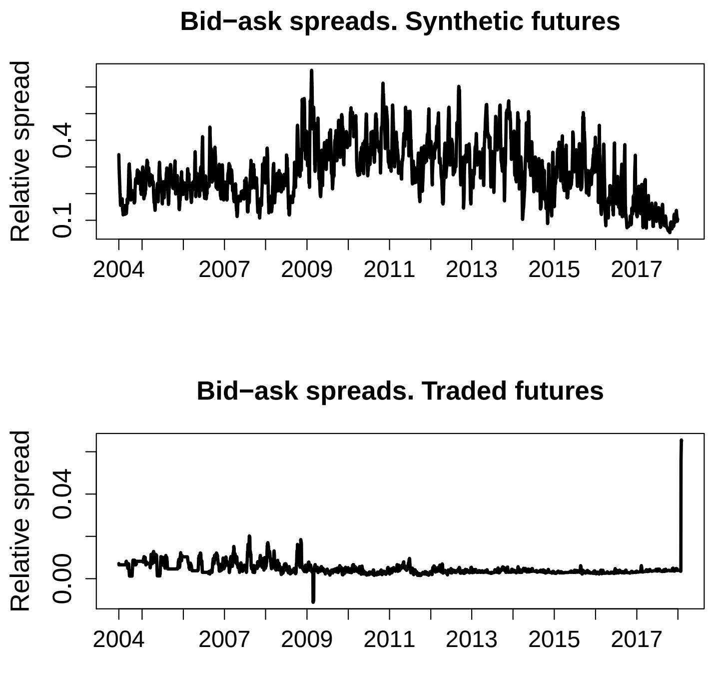

TABLE 1.12: Predictive power of basis – longer sample. Panels A and B present the results from a predictive regression of spot or futures price changes on basis. Daily frequency, June 2004 – February 2018.

Panel A: Spot VIX on basis: $S_T - S_t = \alpha_1 + \beta_1 \cdot (F_{t,T} - S_t) + \epsilon_{1,t}$								
	T=1m	T=2m	T=3m	T=4m	T=5m	T=6m	T=7m	T=8m
β_1	0.06	0.33*	0.62***	0.63***	0.74***	0.85***	0.91***	0.94***
	(0.24)	(0.19)	(0.09)	(0.07)	(0.08)	(0.08)	(0.08)	(0.08)
R^2	0.00	0.03	0.09	0.10	0.14	0.18	0.23	0.27
Observations	3,247	3,235	2,953	2,973	2,991	2,936	2,627	2,389

Panel B: VIX futures on basis: $F_{T,T} - F_{t,T} = \alpha_2 + \beta_2 \cdot (F_{t,T} - S_t) + \epsilon_{2,t}$								
	T=1m	T=2m	T=3m	T=4m	T=5m	T=6m	T=7m	T=8m
β_2	-0.94***	-0.68***	-0.38***	-0.37***	-0.26***	-0.15**	-0.09	-0.06
	(0.22)	(0.16)	(0.07)	(0.06)	(0.07)	(0.06)	(0.06)	(0.06)
R^2	0.14	0.10	0.04	0.04	0.02	0.01	0.00	0.00
Observations	3,247	3,235	2,953	2,973	2,991	2,936	2,627	2,389

TABLE 1.13: Second basis regressed on ETF demand. In columns 1–4 and 6, demand ($D^{\$,all}_{t,2}$ and $D^{\$,all}_{t,2}(r_{t-1})$) is scaled by market capitalization; in column 5, it is not. Columns 1–2 use absolute basis ($F_{t,T_2} - S_t$), 3–6 use relative basis ($\frac{F_{t,T_2}-S_t}{S_t}$). All independent variables are standardized. Daily frequency, February 2009 – December 2017.

Dependent variables	absolute basis		relative basis			
	(1)	(2)	(3)	(4)	(5)	(6)
$D^{\$,all}_{t,2}$	0.39***	0.24***	2.12***	1.77***	1.89***	
	(0.08)	(0.03)	(0.40)	(0.18)	(0.20)	
$D^{\$,all}_{t,2}(r_{t-1})$						0.63***
						(0.17)
$b^H_{t,2}$		2.59***		8.33***	8.31***	8.41***
		(0.14)		(0.41)	(0.41)	(0.41)
$r_{bmk,t}$		−0.07		−0.94***	−1.03***	−0.53***
		(0.05)		(0.19)	(0.20)	(0.17)
$\sigma^2_{bmk,t}$		−0.38		−2.90***	−2.90***	−3.27***
		(0.24)		(0.71)	(0.73)	(0.73)
$OI_{t,2}$		−0.19***		0.28	0.05	0.28
		(0.06)		(0.29)	(0.29)	(0.30)
S_t		−0.49***		−2.50***	−2.58***	−2.60***
		(0.10)		(0.26)	(0.27)	(0.27)
Liq_t		0.16***		0.85***	0.87***	0.94***
		(0.04)		(0.21)	(0.21)	(0.21)
α_t		0.08		1.09***	1.09***	1.05***
		(0.05)		(0.23)	(0.23)	(0.23)
Observations	1,922	1,907	1,922	1,854	1,847	1,853
R^2	0.12	0.81	0.13	0.75	0.75	0.74

TABLE 1.14: Impact of ETFs on basis and EFG. Regressions in first differences (columns 1–8). Relative basis scaled by days to maturity (columns 9–10). Controls include time to maturity, variance of benchmark, return on benchmark, spot price, open interest, and liquidity measured by bid-ask spreads. All independent variables are standardized. Daily frequency, February 2009 – December 2017.

Dependent variables	$\Delta b_{t,1}$, abs	$\Delta b_{t,1}$, rel	$\Delta b_{t,2}$, abs	$\Delta b_{t,2}$, rel	$\Delta EFG_{t,1}$	$\Delta EFG_{t,2}$	$\Delta EFG_{t,1}$	$\Delta EFG_{t,2}$	$b_{t,1}$, rel, days	$b_{t,2}$, rel, days
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$D^{\$,all}_{t,i}$	0.18^{*}	0.48^{***}	0.15	0.34^{*}	0.93^{**}	0.28^{**}			0.10^{***}	0.04^{**}
	(0.10)	(0.07)	(0.15)	(0.19)	(0.38)	(0.11)			(0.02)	(0.02)
Calendar $\text{reb}_{t,i}$							-0.47^{*}	-0.03		
							(0.27)	(0.11)		
Leverage $\text{reb}_{t,i}$							2.32^{**}	0.53^{***}		
							(1.09)	(0.16)		
Flow $\text{reb}_{t,i}$							0.94^{**}	0.25		
							(0.41)	(0.16)		
$\text{Remainder}_{t,i}$							-0.44	-0.21		
							(0.78)	(0.158)		
Controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,944	1,944	1,921	1,921	1,897	1,823	1,889	1,818	1,932	1,908
R^2	0.70	0.52	0.33	0.26	0.35	0.26	0.341	0.26	0.47	0.42

TABLE 1.15: Impact of ETF fractions on prices and EFG in the VIX market. The table presents regression results for the basis, spread, and EFG regressed on the net ETF fraction ($\Phi_{t,i} = \sum_{j=1}^{N} L_j A_{j,t,i}/Mkt\ cap_{t,i}$) in futures with maturity i. All independent variables (except $\Phi_{t,i}$ in columns 3 and 6 of Panel B) are standardized. Panel A shows the result for the first-month basis, and the spread between the first and the second futures. Columns 1 and 3 present the regressions for absolute basis and spread, columns 2 and 4 for relative basis and spread. Panel B presents the results for one and two-months EFG. Columns 3 and 6 use non-standardized demand. Controls include return on the ETF benchmark, variance of the ETF benchmark, open interest, spot price, liquidity and time to maturity. Daily frequency, February 2009 – December 2017.

Panel A: Basis and spread on ETF fraction

Dependent variables	$b_{t,1}$, abs.	$b_{t,1}$, rel.	$b_{t,2}$, abs.	$b_{t,2}$, rel.
	(1)	(2)	(3)	(4)
$\Phi_{t,i}$	0.19***	**0.42****	0.70***	**0.61*****
	(0.05)	(0.17)	(0.13)	(0.12)
$b^H_{t,i}$	1.35***	5.20***	1.50***	5.63***
	(0.24)	(1.17)	(0.05)	(1.12)
Controls	Yes	Yes	Yes	Yes
Observations	1,946	1,946	1,923	1,923
R^2	0.39	0.49	0.47	0.49

Panel B: EFG on ETF fraction

Dependent variables		$EFG_{t,1}$			$EFG_{t,2}$	
	(1)	(2)	(3)	(4)	(5)	(6)
$\Phi_{t,i}$	2.07***	1.82***	**0.97*****	0.53***	0.41***	**1.29*****
	(0.37)	(0.32)	(0.22)	(0.14)	(0.13)	(0.28)
$EFG_{t-1,i}$		6.64***	6.63***		4.04***	4.04***
		(0.57)	(0.58)		(0.17)	(0.17)
Controls	Yes	Yes	Yes	Yes	Yes	Yes
Observations	1,899	1,899	1,899	1,825	1,825	1,825
R^2	0.28	0.47	0.47	0.29	0.61	0.61

Table 1.16: Robustness: periods of high and low variance, Fama–French five factors. Columns 1–2 and 4–5 show the impact of ETF demand in periods of high (above median) and low (below median) variance of the benchmark. Columns 3 and 6 present the results of a regression with the Fama–French five factors and momentum. Columns 7–8 show the results for the period before ETFs. $D^{\$,all}_{t,i}$ is scaled by market capitalization. $R_{M,t} - R_{f,t}$, HML_t, SMB_t, CMA_t, RMW_t, Mom_t are the Fama–French five factors and momentum. Daily frequency, February 2009 – December 2017.

Dependent variables	$\text{EFG}_{t,1}$		$\text{EFG}_{t,1}$	$\text{EFG}_{t,2}$		$\text{EFG}_{t,2}$	$\text{EFG}_{t,1}$	$\text{EFG}_{t,2}$
	high var	low var		high var	low var		before ETFs	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$D^{\$,all}_{t,i}$	1.15***	0.71*	0.96***	0.48***	0.36***	0.43***		
	(0.23)	(0.37)	(0.38)	(0.18)	(0.10)	(0.11)		
$\text{EFG}_{t-1,1}$	5.40***	4.47***	6.07***	3.91***	3.12***	3.93***	6.52***	3.27***
	(0.58)	(1.06)	(0.72)	(0.30)	(0.18)	(0.19)	(0.26)	(0.15)
$\text{r}_{bmk,t}$	0.69***	0.49	0.45***	0.39***	-0.06	0.18***	-0.55**	-0.33***
	(0.14)	(0.30)	(0.11)	(0.08)	(0.14)	(0.06)	(0.23)	(0.12)
$\sigma^2_{bmk,t}$	0.99	0.68*	1.49	-0.70	-0.12	-0.32	-0.27	-0.00
	(1.29)	(0.35)	(0.96)	(0.90)	(0.11)	(0.57)	(0.22)	(0.12)
$\text{OI}_{t,i}$	1.76	2.44	1.69	-1.14	-1.28	-1.11	-0.89***	0.01
	(1.80)	(1.99)	(1.49)	(1.25)	(1.25)	(1.16)	(0.32)	(0.17)
S_t	-0.57	0.63***	0.31	-0.60***	-0.17***	-0.31***	-0.60*	0.045
	(1.12)	(0.15)	(0.58)	(0.13)	(0.04)	(0.08)	(0.33)	(0.18)
$\text{Liq}_{t,i}$	0.97*	1.04**	0.91**	-0.21	0.22	-0.01	0.33**	0.14**
	(0.51)	(0.53)	(0.37)	(0.20)	(0.14)	(0.11)	(0.16)	(0.06)
TED_t	1.60	-1.17	0.46	2.33**	0.92**	1.44***	0.38	-0.15
	(1.86)	(1.05)	(0.94)	(1.08)	(0.44)	(0.40)	(0.33)	(0.19)
α_t	0.96***	0.65*	0.66***	-0.11	-0.04	-0.31**	-0.22	0.05
	(0.35)	(0.37)	(0.24)	(0.19)	(0.18)	(0.12)	(0.28)	(0.15)
$R_{M,t}$-$R_{f,t}$			0.15*			1.61		
			(0.09)			(1.24)		
HML_t			0.03			-0.04		
			(0.34)			(0.12)		
SMB_t			-0.03			0.01		
			(0.24)			(0.10)		
CMA_t			-0.13			0.01		
			(0.26)			(0.08)		
RMW_t			-0.11			0.08		
			(0.25)			(0.09)		
Mom_t			0.09			0.17		
			(0.22)			(0.10)		
Observations	949	949	1,870	912	912	1,801	464	481
R^2	0.50	0.36	0.44	0.60	0.53	0.58	0.65	0.53

TABLE 1.17: Robustness: adding lagged demand and hedging pressure from the VIX options market. Columns 1–2 show the results for relative basis, 3–6 for EFG. $D_{t,i}^{\$,all}$ is scaled by market capitalization. $\text{Delta}_{t,i}$ and $\text{Gamma}_{t,i}$ are measures of hedging pressure in the VIX options market. $\text{Delta}_{t,i}$ is the sum of all Black-Scholes deltas multiplied with the open interest and the futures price for all options on the first-month or second-month futures. $\text{Gamma}_{t,i}$ is the sum of all Black-Scholes gammas multiplied with the open interest and the squared price for all options on the first-month or second-month futures. All independent variables are standardized. $R_{M,t}-R_{f,t}$, HML_t, SMB_t, CMA_t, RMW_t, Mom_t are the Fama–French five factors and momentum. Daily frequency, February 2009 – December 2017.

Dependent variables	$b_{t,1}$, rel	$b_{t,2}$, rel	$EFG_{t,1}$	$EFG_{t,2}$	$EFG_{t,1}$	$EFG_{t,2}$
	(1)	(2)	(3)	(4)	(5)	(6)
$D_{t,1}^{\$,all}$	0.89***	0.39***	1.22***	0.60***	1.37***	0.48***
	(0.11)	(0.11)	(0.36)	(0.10)	(0.41)	(0.14)
$D_{t-1,1}^{\$,all}$	0.70***	0.15**	-0.37***	-0.17**	-0.42***	-0.19*
	(0.11)	(0.07)	(0.08)	(0.08)	(0.09)	(0.11)
$b_{t,i}^{H}$	5.42***	5.19***				
	(1.22)	(1.14)				
$\text{EFG}_{t-1,i}$			6.12***	4.15***	6.27***	4.04***
			(0.60)	(0.16)	(0.82)	(0.17)
$\text{r}_{bmk,t}$	-0.97	-0.51	1.86***	1.26***	2.14***	1.51***
	(0.86)	(0.65)	(0.59)	(0.22)	(0.69)	(0.25)
$\sigma^2_{bmk,t}$	-1.17***	-0.76***	0.55	-0.47	0.54	-0.21
	(0.28)	(0.16)	(0.82)	(0.50)	(0.82)	(0.18)
$\text{OI}_{t,i}$	1.66***	0.15	-1.56	-0.36	-1.14	-0.27
	(0.27)	(0.24)	(1.35)	(0.29)	(1.47)	(0.28)
S_t	-2.79***	-3.89***	-0.09	-0.31**	-0.05	-0.32**
	(0.32)	(0.26)	(0.28)	(0.12)	(0.49)	(0.15)
$\text{Liq}_{t,i}$	1.22***	1.59***	0.73***	0.08	0.70***	0.07
	(0.21)	(0.37)	(0.17)	(0.14)	(0.15)	(0.10)
TED_t			-0.01	1.26***	0.21*	0.93***
			(0.17)	(0.39)	(0.12)	(0.30)
α_t	0.93***	-0.21*	0.46***	-0.34*	0.63**	-0.41***
	(0.20)	(0.12)	(0.17)	(0.18)	(0.31)	(0.13)
$R_{M,t}-R_{f,t}$			0.13	1.14*	0.36	1.53**
			(0.29)	(0.66)	(0.48)	(0.71)
HML_t			-0.21	-0.02	-0.08	-0.15
			(0.22)	(0.10)	(0.39)	(0.11)
SMB_t			-0.36	-0.03	-0.21	-0.16
			(0.24)	(0.10)	(0.28)	(0.11)
RMW_t			-0.10	-0.01	-0.15	0.03
			(0.19)	(0.09)	(0.27)	(0.09)
CMA_t			-0.06	0.04	-0.04	0.08
			(0.18)	(0.07)	(0.34)	(0.10)
Mom_t			-0.14	0.17	-0.08	0.06
			(0.18)	(0.11)	(0.24)	(0.09)
$\text{Delta}_{t,i}$					-0.62*	-0.02
					(0.36)	(0.13)
$\text{Gamma}_{t,i}$					-0.09	-0.12
					(0.40)	(0.14)
Observations	1,942	1,920	1,895	1,822	1,842	1,797
R^2	0.49	0.45	0.51	0.63	0.50	0.63

TABLE 1.18: Predictive regressions of futures returns on EFG in the VIX market. Columns 1 and 2 show the results of monthly predictive regressions of the realized futures returns on the EFG. For one-month futures contract, I use returns from the date when a two-months contract becomes a one-month contract, to expiration. For two-months futures contract, I use returns calculated from 45 days before maturity, to expiration. $r^{F_i}_{t,T_i} = \frac{F_{T_i,T_i} - F_{t,T_i}}{F_{t,T_i}}$. Columns 3 and 4 present daily predictive regressions. Return on futures is already excess return because the collateral earns the risk-free rate of interest: all futures positions are fully collateralized, with the collateral invested in three-month Treasury bills. $R_{M,t} - R_{f,t}$, HML_t, SMB_t, CMA_t, RMW_t, Mom_t are the Fama–French five factors and momentum. The data sample is February 2009 – December 2017.

Dependent variables	$r^{F_1}_{t,T_1}$ (1)	$r^{F_2}_{t,T_2}$ (2)	$r^{F_1}_{t}$ (3)	$r^{F_2}_{t}$ (4)
$\mathrm{EFG}_{t,i}$	-0.22^{**}	-1.44^{***}	-0.31^{***}	-1.28^{***}
	(0.10)	(0.51)	(0.06)	(0.20)
$\mathrm{r}_{bmk,t}$	0.03	0.005	0.01	0.02^{**}
	(0.03)	(0.05)	(0.01)	(0.01)
$\sigma^2_{bmk,t}$	0.02	0.01	0.004	0.01
	(0.02)	(0.02)	(0.02)	(0.03)
$\mathrm{OI}_{t,i}$	-0.03	0.02	-0.001	-0.04^{***}
	(0.03)	(0.04)	(0.01)	(0.01)
S_t	-0.08^{**}	-0.08^{***}	-0.03^{***}	-0.10^{***}
	(0.03)	(0.03)	(0.01)	(0.01)
$\mathrm{Liq}_{t,i}$	0.004	0.003	0.01	0.01
	(0.02)	(0.02)	(0.01)	(0.01)
TED_t	0.07	0.09^{**}	0.07^{***}	0.12^{***}
	(0.05)	(0.04)	(0.02)	(0.02)
α_t	-0.02	0.03	0.01	-0.004
	(0.02)	(0.02)	(0.01)	(0.01)
$R_{M,t} - R_{f,t}$	0.01	0.08	0.003	-0.001
	(0.03)	(0.07)	(0.01)	(0.01)
SMB_t	0.03	-0.07^{**}	-0.003	0.005
	(0.02)	(0.03)	(0.005)	(0.01)
HML_t	-0.04	-0.01	0.005	0.02^{*}
	(0.05)	(0.03)	(0.01)	(0.01)
RMW_t	0.08^{***}	0.02	-0.002	0.01
	(0.03)	(0.04)	(0.01)	(0.01)
CMA_t	0.02	-0.04^{*}	-0.005	-0.000
	(0.02)	(0.02)	(0.01)	(0.01)
Mom_t	-0.04	0.03	-0.002	0.01
	(0.04)	(0.04)	(0.005)	(0.01)
Observations	102	68	1,893	1,822
R^2	0.25	0.27	0.15	0.21

Table 1.19: Regression results of synthetic futures and spot on synthetic basis. Panel A presents the results from a predictive regression of spot price changes on synthetic basis: $S_T - S_t = \alpha_1 + \beta_1 \cdot (\mathbb{E}_t^Q(S_T) - S_t) + \epsilon_{1,t}$. Panel B presents the results from a predictive regression of synthetic futures changes on synthetic basis: $\mathbb{E}_T^Q(S_T) - \mathbb{E}_t^Q(S_T) = \alpha_2 + \beta_2 \cdot (\mathbb{E}_t^Q(S_T) - S_t) + \epsilon_{2,t}$. Daily frequency, February 2009 – December 2017.

Panel A: Spot VIX on basis

	T=1m	T=2m	T=3m	T=4m
β_1	0.25***	0.55***	0.64***	0.73***
	(0.04)	(0.18)	(0.10)	(0.11)
Observations	2,137	2,119	2,009	1,840
R^2	0.06	0.05	0.07	0.09

Panel B: Synthetic VIX futures on basis

	T=1m	T=2m	T=3m	T=4m
β_2	-0.75***	-0.45***	-0.36***	-0.27***
	(0.05)	(0.19)	(0.11)	(0.09)
Observations	2,137	2,119	2,009	1,840
R^2	0.41	0.14	0.50	0.33

Table 1.20: Impact of ETF demand on synthetic basis $(\mathbb{E}_t^Q(S_{T_1}) - S_t)$ and spread $(\mathbb{E}_t^Q(S_{T_2}) - \mathbb{E}_t^Q(S_{T_1}))$. All independent variables are standardized. Daily frequency, February 2009 – December 2017.

Dependent variables	$\mathbb{E}_t^Q(S_{T_1}) - S_t$ (1)	$\mathbb{E}_t^Q(S_{T_2}) - \mathbb{E}_t^Q(S_{T_1})$ (2)	$\mathbb{E}_t^Q(S_{T_1}) - S_t$ (3)	$\mathbb{E}_t^Q(S_{T_2}) - \mathbb{E}_t^Q(S_{T_1})$ (4)
$D_{t,i}^{\$,all}$	-0.05	-0.06	0.03	-0.07
	(0.07)	(0.06)	(0.08)	(0.06)
$b_{t,i}^H$	1.87***	1.47***		
	(0.19)	(0.13)		
$r_{bmk,t}$	-0.44***	0.24**	-0.74***	0.27***
	(0.09)	(0.09)	(0.09)	(0.09)
$OI_{t,i}$	-0.04	-0.40***	-0.20	-0.34***
	(0.12)	(0.10)	(0.16)	(0.10)
$\sigma_{bmk,t}^2$	-1.15***	0.26	-1.92***	0.41
	(0.33)	(0.30)	(0.39)	(0.30)
S_t	-1.05***	-0.28	-2.18***	0.02
	(0.23)	(0.18)	(0.29)	(0.16)
$Liq_{t,i}$	-0.05	0.37	1.08**	0.33
	(0.37)	(0.32)	(0.43)	(0.32)
α_t	0.16	-0.17**	0.11	-0.15*
	(0.10)	(0.08)	(0.11)	(0.09)
Observations	1,872	1,815	1,872	1,815
R^2	0.38	0.30	0.28	0.22

TABLE 1.21: Regressions of positions of leveraged money in VIX futures. The table presents weekly regressions of the positions of leveraged money (mostly hedge funds) on the ETF futures gap. $b_{t,1}$ and $b_{t,2}$ are absolute basis and spread. Column 3 is with raw variables, the rest with standardized ones. Weekly frequency, September 2006 – December 2017 (some data is missing).

Dependent variables	Weekly Hedge Funds' net positions, million USD				
	(1)	(2)	(3)	(4)	(5)
EFG_t	**-7.46****	**-7.54****	**-30.62****		
	(3.51)	(3.47)	(14.20)		
ETF positions_t	-38.46***	-38.45***	-0.78***	-44.69***	-32.88***
	(4.17)	(4.17)	(0.08)	(3.06)	(4.57)
$\sigma^2_{bmk,t}$		-12.57**	-39.63	-7.68	-7.06
		(5.23)	(24.02)	(4.94)	(6.87)
$b_{t,1}$	-2.03	-3.11	-1.54	-0.47	-3.48
	(2.97)	(3.01)	(2.09)	(2.13)	(2.53)
$b_{t,2}$	-12.60**	-13.03**	-6.64***	-12.11***	-12.40**
	(5.07)	(5.25)	(2.29)	(3.79)	(5.35)
S_t	2.58	2.82	0.57	0.21	-1.63
	(5.28)	(5.27)	(0.63)	(4.33)	(5.59)
Observations	416	416	416	452	452
R^2	0.59	0.60	0.60	0.72	0.55

TABLE 1.22: Short-both strategy returns across markets. The table shows α-s of excess returns on the short-both strategy $r_{SB,t}$ for each market with respect to the Fama–French five factors, momentum, spot return, and variance of the benchmark σ^2_{bmk}. The last column shows the estimates from a regression that includes five year CDS quotes of the long and the inverse ETF sponsors for VIX – Barclays and Credit Suisse, respectively. $r_{SB,t}$ are in %. $R_{M,t} - R_{f,t}$, HML_t, SMB_t, CMA_t, RMW_t, Mom_t are the Fama–French five factors and momentum. Interpretation of the coefficients, e.g. HML_t for Financials: 1% increase in HML_t increases $r_{SB,t}$ by 0.1220%. The sample period is from the first inverse ETF introduction date in a given market to December 2018 (daily frequency).

Dependent variables	$r_{SB,t}$							
	VIX (%)	Gas (%)	Silver (%)	Oil(%)	Gold (%)	S&P 500 (%)	Financials (%)	VIX (%)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
α	0.06***	0.14***	0.07***	0.03*	0.03*	0.08***	0.16*	0.06**
	(0.02)	(0.03)	(0.01)	(0.01)	(0.02)	(0.02)	(0.10)	(0.03)
S_t	10.57	1.20	4.25	10.30	1.04	7.76	2.22	10.22
	(17.11)	(2.25)	(4.54)	(11.37)	(1.46)	(7.28)	(3.59)	(16.94)
$R_{M,t}$-$R_{f,t}$	20.90	0.53	−1.43	−1.84	0.75	7.32	8.60	21.35
	(12.74)	(2.76)	(2.52)	(1.83)	(5.84)	(4.86)	(14.08)	(17.97)
SMB_t	−18.36*	−2.19	1.35	−1.82	−0.40	−5.34	−16.15	−18.47*
	(10.10)	(4.82)	(3.23)	(2.29)	(5.34)	(7.26)	(10.63)	(10.13)
HML_t	−1.43	−1.05	−2.29	−2.54	6.67	−11.42	12.20***	−1.41
	(2.97)	(6.41)	(4.27)	(3.07)	(9.42)	(13.61)	(2.97)	(2.84)
RMW_t	11.64	−2.36	9.10	1.45	−0.55	9.35	9.71	11.87
	(11.17)	(6.57)	(5.66)	(5.82)	(6.78)	(12.21)	(10.54)	(11.20)
CMA_t	9.68	2.99	−2.28	0.77	−0.47	23.92	−17.46**	9.64
	(7.79)	(8.19)	(5.45)	(4.26)	(16.50)	(21.65)	(7.53)	(7.77)
Mom_t	0.06	0.03	−0.00	0.01	0.07**	−0.03	−0.07	0.05
	(0.05)	(0.04)	(0.03)	(0.02)	(0.03)	(0.06)	(0.04)	(0.05)
$\sigma^2_{bmk,t}$	0.01*	0.02	0.01**	0.00**	-0.00***	0.01***	0.00	0.01
	(0.01)	(0.01)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.01)
CDS_t, long sponsor								0.01
								(0.01)
CDS_t, inverse sponsor								0.01
								(0.01)
CDS_{t-1}, long sponsor								-0.01
								(0.01)
CDS_{t-1}, inverse sponsor								-0.01
								(0.01)
Observations	1,890	1,708	1,793	419	1782	2,535	2,530	1,890
R^2	0.02	0.00	0.01	0.01	0.03	0.04	0.02	0.02

TABLE 1.23: ETF impact on realized futures returns. The table presents the results of an event-study for realized futures returns to see if the steepening of the curve caused by ETF demand, pushed realized returns down. It shows the estimates of the regression $r_t^{F_i} = \alpha + \beta Post_t + \gamma r_t^{H_i} + \epsilon_{t,j}$ with a window of 60 days around the first ETF introduction date. $r_t^{F_i}$ is the realized return on the futures until maturity T_i ($r_t^{F_i} = \frac{F_{T_i,T_i} - F_{t,T_i}}{F_{t,T_i}}$), α is an intercept, $r_t^{H_i}$ is the return on the hedge asset. All returns are in %.

Dependent variables	VIX		Gas		Silver		Gold		Oil	
	$r_t^{F_1}$	$r_t^{F_2}$	$r_t^{F_1}$	$r_t^{F_2}$	$r_t^{F_1}$	$r_t^{F_2}$	$r_t^{F_1}$	$r_t^{F_2}$	$r_t^{F_1}$	$r_t^{F_2}$
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
$Post_t$	-0.63	-0.30***	-0.08	-0.02**	0.01	-0.03**	-0.01	-0.00	-0.01	-0.01**
	(0.41)	(0.07)	(0.06)	(0.01)	(0.12)	(0.02)	(0.03)	(0.00)	(0.01)	(0.00)
$r_t^{H_i}$	0.65**	0.67**	0.47***	0.42***	0.52***	0.41***	0.47***	0.38***	1.04***	0.98***
	(0.32)	(0.33)	(0.09)	(0.15)	(0.15)	(0.13)	(0.16)	(0.13)	(0.37)	(0.38)
Observations	120	120	119	119	120	120	120	120	120	120
R^2	0.48	0.47	0.24	0.22	0.23	0.22	0.34	0.31	0.93	0.92

Chapter 2

Quantify the Quantitative Easing: Impact on Bonds and Corporate Debt Issuance

Karamfil Todorov[1]

2.1 Introduction

Recent years have seen a surge in the use of unconventional monetary policy tools across the developed world. Following the 2008 financial crisis, major central banks have exhausted the traditional monetary toolkit and have started using new instruments to spur economic activity and tackle low inflation. Being constrained by the (near zero) lower bound of interest rates, national regulators have expanded their traditional measures for stimulating the economy by conducting large-scale asset purchases known as quantitative easing (QE). Billions of dollars, euros, pounds, and yen have been injected into the economy by the Federal Reserve (Fed), the ECB, the Bank of Japan (BoJ), and the Bank of England (BoE) in order to buy bonds and other asset classes as part of QE. These massive purchases represent a substantial fraction of the Gross

[1]I am grateful to the editor (William Schwert) and the referee (Bo Becker) for their excellent comments and suggestions. The referee's recommendations to disentangle the different QE channels and to use notional value of new issued bonds were extremely helpful. I would also like to thank Ron Anderson, Simeon Djankov, Daniel Ferreira, Christian Julliard, Dong Lou, Igor Makarov, Martin Oehmke, Daniel Paravisini, Andrea Tamoni, Dimitri Vayanos, Kathy Yuan and Jean-Pierre Zigrand for their useful advice and comments. I am also thankful to seminar participants at the London School of Economics and Political Science, London Business School, Goethe University Frankfurt, and Paris Dauphine University for their helpful feedback.

Domestic Products (GDP) of these countries,[2] but despite its large scale, little research has been devoted to the consequences of QE for bond prices, liquidity, and for corporates in general. In this paper, I exploit a unique data set on bond transactions and firms' financial indicators, which allows me to study questions as yet unanswered in the QE literature. I quantify the very heterogeneous effects of the European QE programme on bonds and on corporate debt issuance. Surprisingly, I find that the main effects of QE for corporates were transmitted through increase in dividend payments rather than through changes in investment.

2.1.1 Empirical evidence

The analysis developed in this research is based on the Corporate Sector Purchase Programme (CSPP) launched by the ECB in 2016 as part of QE. The focus of the study is on the ECB because of all the central banks implementing QE, the European institution is the largest holder of corporate bond debt.[3] This paper exploits the exogenous nature of the CSPP announcement to construct a before-after comparison between bonds that are eligible for purchase and bonds that are not. I make use of a proprietary data set on bond transactions between February 2015 and June 2016 from Euroclear, combined with pricing information from Bloomberg and with firms' financial data from the Bureau van Dijk, and employ a difference-in-differences (DD) research design to estimate the causal impact of QE.

The first broad topic of this paper is the impact of QE on bond prices and liquidity. The results show that QE significantly increased prices: yields of eligible bonds declined by 30 bps (8%) after the CSPP announcement. The policy also had a large positive impact on liquidity, which I quantify using two main metrics: trading activity (turnover) and cost of trading (bid-ask spreads). After the CSPP announcement, tri-party (repo) turnover increased on average by 8.15 million USD (29% rise given the control group before-QE mean of 28.14 million USD) and bilateral (buyer-seller) turnover[4] went up by 7.05 million USD (72% increase given the control group before-QE mean of 9.69 million USD). The bid-ask spread of yields went down by 46%, indicating massively reduced costs of a round-trip transaction. However, the positive impact on liquidity

[2] Around 20% for the Fed and the BoE, 40% for the ECB, and more than 90% for the BoJ as of May 2017. Sources – the Fed, the ECB, the BoE, the BoJ.

[3] 170.38 billion euros as of September 2018. Source – the ECB.

[4] A bilateral trade takes place when two member firms provide matching instructions to Euroclear to transfer ownership from one counterparty to the other in return for an agreed amount of money. The tri-party repo transactions are repurchase transactions arranged by a tri-party repo agent whereby one party can obtain short-term funds by providing a security it owns for a sale to a counterparty with a matching repurchase agreement set for a later date.

was mostly concentrated in the segment of QE-eligible bonds. The effects on ineligible bonds were less pronounced, and in the tri-party repo market, liquidity of these bonds even decreased after QE. Moreover, the initial spike in trading activity gradually dissipated. In contrast, the impact on yields and the reduction in trading costs were more persistent. There was also some evidence of spillover effects to yields of ineligible bonds.

As a next step, the paper studies which bonds benefited the most from the QE intervention. The drop in yields is decomposed into different risk premia by applying the methodology of Krishnamurthy and Vissing-Jorgensen (2011) (KVJ). I run DD regressions, splitting the sample of bonds into several maturity-rating-liquidity risk buckets to isolate the duration risk channel, the default risk channel, and the liquidity channel of QE. The results show that the greatest effects of the CSPP were observed for bonds with higher duration and default risks. The estimates reveal an inverse relationship between the rating of a bond and the decline in yields. A positive relationship is observed between the maturity of the bond and the reduction in yields: longer-duration bonds experienced the largest drop in yields. On the other hand, the liquidity impact of QE was only marginally different for bonds with different duration and default risks.

The fact that eligible lower-rated, longer-maturity bonds close to the threshold were the main drivers of the decrease in yields shows that the ECB's intervention had a higher positive impact on riskier debt instruments within the treated set of bonds. These findings suggest that the European QE was successful in reducing the duration, the default, and the liquidity premiums of eligible bonds and in re-allocating risks. The logic is as follows. In times of financial turmoil investors typically require a larger premium for holding riskier bonds, causing a decrease in bond prices. As a consequence, banks holding these risky illiquid bonds could become constrained by their risk bearing capacity as in Vayanos and Vila (2009), or face Value at Risk (VaR) constraints. By implementing QE, the central bank steps in, inflates bond prices, and improves liquidity by making it easier for investors to sell these risky illiquid assets as part of the bond buying programme, thereby reducing the risk premium and lowering bond yields. The drop in yields and the increase in liquidity could potentially reduce risks in the system by removing problematic assets from banks' balance sheets and by contributing to deleveraging (see, e.g., Woodford, 2012). The effects of the European QE documented in this research are consistent with such a story.

The second broad part of this paper studies the reaction of corporates to QE. The analysis shows an overall increase in the issuance of QE-eligible bonds compared to ineligible ones both in the absolute number of bonds, and in the par amount outstanding. However, firms issuing eligible bonds might be systematically different from those issuing ineligible bonds. To address this concern, I saturate the firm dimension by restricting the sample to firms that can issue bonds in both segments. If there was a switch to issuing more QE-eligible debt, these firms would be more likely to issue a larger amount of their debt liabilities in euro-denominated bonds that satisfy the CSPP criteria, and less in other types of bonds. This is indeed what I find. Corporates issued 1.39 QE-eligible bonds more per week (a 38% increase given the mean before QE for ineligible bonds of 3.70), after the CSPP announcement, compared to ineligible bonds. This corresponds to 1.94 billion EUR, or a 49% rise in newly issued QE-eligible debt each week. Firms also increased the fraction of bond debt in their total liabilities.

Next, the study focuses on the heterogeneous effects of QE for firms with different risks. A number of papers (see, e.g., Kashyap et al., 1993; Rauh and Sufi, 2010) show that corporates become more credit-constrained and reduce bank borrowing during financial crises. However, it is unclear whether QE would alleviate the problem by boosting new debt issuance and easing borrowing for more financially constrained firms. It is also debatable whether these firms would use the funds attracted through bond issuance to increase investment. To answer these questions, I conduct the following analysis. Given that longer-duration, lower-rated bonds experienced the largest drop in yields, one could hypothesize that firms would be more likely to issue bonds with longer maturities and that riskier firms would make wider use of the programme. To test these hypotheses, I split the eligible segment in different rating and maturity groups and run several DD regressions. The results show that corporates were indeed more likely to issue bonds with longer maturities. Moreover, the estimates illustrate a monotonic inverse relationship between the rating of a firm and the amount of issued QE-eligible debt. This fact shows that lower-rated corporates made wider use of the programme.

Surprisingly, however, the funds attracted through QE, were not used to increase investment, but were mostly utilized through dividend payments. The results of a DD regression on a subset of firms with financial data before and after the CSPP announcement show that firms issuing QE-eligible bonds did not change cash holdings, investments in working capital, property, plant, and equipment (PPE), or research and development (R&D). To the contrary, these firms increased dividends four times compared to the before-QE mean of firms not issuing QE-eligible bonds. These new empirical findings suggest that QE was successful in providing cheap liquidity to more credit-constrained corporates by lowering the cost of market debt issuance. However, the programme had greater impact on firms' dividend payments and little impact on investment.

The main conclusions of this paper are verified in a series of robustness checks. To address the issue of heterogeneity in bond characteristics, I construct a matched sample, restricting the treatment group to bonds for which there is a suitable control bond with similar pre-treatment characteristics (par amount outstanding, rating, time to maturity). I also try different variations of the control group and extensions of the post-announcement period. Furthermore, I collapse the data by time and run pure time-series tests to verify the panel regression estimates. The main findings of the study are robust to these modifications. To check that the sample selection of bonds in treatment and control groups does not mechanically produce the main results, the study performs two placebo tests. In the first, the sample period is shifted backwards by one year to remove any effects of seasonality. In the second, USD-denominated bonds that satisfy all the CSPP eligibility criteria are used as a treatment group. Both placebo tests show no statistically significant effects and rule out the possibility that the main conclusions of the research are due to seasonality, or due to mechanical sample selection.

2.1.2 Related literature

The analysis developed in this paper is related to the growing literature on quantitative easing. Most articles in the QE literature have focused on the effects of the Fed's QE programme. Krishnamurthy and Vissing-Jorgensen (2011) study the impact of the Federal Reserve's QE programmes on interest rates and outline several important channels through which QE affects bond yields. In a recent paper, Song and Zhu (2018) analyze the implementation of QE in the US and study which bonds are more likely to be purchased by the Fed. Christensen and Gillan (2018) find that the Federal Reserve's second QE programme improved market liquidity. However, there is little research on the impact of the ECB's QE on bond yields, liquidity, and on corporate debt issuance, partly because there is no analogue of the Trade Reporting and Compliance Engine (TRACE) database for European bonds. Given the significance of the corporate bond market in Europe, however, the impact of the ECB's QE is an important concern for European investors and market participants. To the best of my knowledge, the research presented here is the first one to analyze the impact of the European QE on bond yields and bond liquidity, and the first one to study firms' reaction to QE in general.

This paper is also related to the literature on liquidity measures in financial markets. Asquith et al. (2013) assess the impact of a mandatory transparency regulation on bond liquidity using volume of trading and turnover ratio. Amihud (2002) develops a new price-volume indicator to measure liquidity in the stock market. Huang and

Wang (2009) show that illiquidity can lead to crashes in asset prices and increases in volatility. Bao et al. (2011) confirm the latter conclusion and show that liquidity of corporate bonds decreased significantly during the financial crisis of 2008. In recent research, the TABB Group (2016) summarizes the most commonly used measures of liquidity and shows that liquidity increased in the US corporate bond market after the financial crisis. In the research presented here, I contribute to the literature by quantifying the impact of QE on liquidity of European corporate bonds using turnover, bid-ask spreads, and the Amihud illiquidity measure.

The paper presented here also builds on the extensive literature on corporate debt issuance. Bolton and Freixas (2000) present a model in which, because of the greater flexibility of bank credit, corporates with higher default probability would prefer bank financing to bond financing during financial distress. Becker and Ivashina (2014) find evidence that firms substitute bank credit with bonds as credit conditions tighten. Rauh and Sufi (2010) find similar effects for US rated firms: corporates substitute bank credit with bond financing as their credit quality increases. The conclusions of De Fiore and Uhlig (2015) are in line with these findings. Their paper shows how banking distress can lead to a surge in corporate bond issuance. The results of the research presented here complement the literature on corporate debt issuance. Consistent with the conclusions of Adrian and Shin (2014), Becker and Ivashina (2014), and De Fiore and Uhlig (2015), I find that there is a switch to issuing more corporate debt in the aftermath of the financial crisis. The new result is that these conclusions continue to hold even during a QE intervention and in a low-interest rate environment. Moreover, I show that riskier firms issued more debt in response to QE implementation. I also find that corporates were more likely to issue longer-maturity bonds since QE reduced yields of these bonds the most.

Financial regulators and international finance associations have expressed differing views on the potential impact of the European corporate bonds purchase programme on debt issuance. For example, in a note, the ICMA (2016) argues that cheaper funding opportunities due to the CSPP would not have a significant influence on firms' debt issuance, since "many corporates are already awash with cash." The note also argues that the ECB's intervention might in fact push investors towards riskier QE-ineligible bonds in the search for yield, and reduce liquidity in eligible bonds. On the other hand, Demertzis and Wolff (2016) and Sparks (2017) claim that there is evidence that corporates took advantage of the CSPP to issue a greater number of securities. There is no clear consensus on the effects of the CSPP for corporate debt issuance, bond prices, bond liquidity, and the use of QE funds. The research presented here aims to fill in these gaps.

The rest of this paper is organized as follows. In section 2.2, I describe the data set, the research framework, and give details of the CSPP. Section 2.3 presents the results for the impact of the ECB's intervention on bond prices and liquidity. Section 2.4 describes the corresponding results for corporate debt issuance and the use of QE funds. Section 2.5 presents two placebo tests and section 2.6 concludes.

2.2 Data and research framework

2.2.1 Data

The data set used in this research consists of information on the total volume of trading (turnover) and the holdings of individual bonds reported by Markit based on reports supplied by Euroclear Plc. (the largest securities depository in the Eurobond market and a major clearing house). I combine these data with information on the issue, the issuing entity, and the price of each bond from Bloomberg, and with data on firms' financial indicators from the Bureau van Dijk.

The Euroclear data set reports total daily transactions on individual securities from February 2015 to June 2016. For each security (identified by International Securities Identification Number (ISIN)) I observe the total value of transactions (expressed in USD equivalents) in the bilateral and tri-party repo markets. The entire data set consists of 109,421 different ISINs and 330 trading days. As many bonds do not trade on a day-to-day basis, I construct total weekly turnover of a given bond, which is consistent with the literature on bond liquidity (see, e.g., Asquith et al., 2013). Since the focus of this paper is the impact of the CSPP on European corporate bonds, I filter out all non-corporate bonds, and all bonds outside the EU. This leaves a total of 6,590 bonds traded during the sample period. The next section provides details of the corporate bonds purchase programme.

2.2.2 The CSPP

The ECB started a series of QE-programmes following ECB President Mario Draghi's famous "whatever it takes" speech in 2012. The first programmes, the Asset-Backed Securities Purchase Programme (ABSPP) and Covered Bond Purchase Programme (CBPP) were aimed at buying long-term EU sovereign bonds. These programmes improved economic conditions in the euro area and contributed to growth (see, e.g., Demertzis and Wolff, 2016), but were not enough to achieve the inflationary target set by the ECB. It became evident that the European regulator should go beyond government bonds and target riskier assets as part of its QE programme.

On 10 March 2016, the ECB announced a further expansion of its asset purchase programme including the new CSPP. This measure was introduced with its main purpose being to "further strengthen the pass-through of the Eurosystem's asset purchases to the financing conditions of the real economy" and was expected to last until at least March 2017.[5] The decision to include investment grade non-bank corporate bonds in the QE programme caught the market by complete surprise. The CSPP was a very novel measure since the ECB entered the private sector financing market, whereas previously it had bought almost exclusively government bonds as part of QE. However, the full details of the CSPP, including the eligibility criteria, were not specified, with the promise of making them public "in due course." In fact, this promise left more questions than answers, since it did not explicitly define the range of ratings, set of maturities, or firm eligibility criteria.

On 21 April 2016, the ECB published a full set of criteria for a bond to be purchased under the new programme. To be eligible, a bond had to satisfy the following conditions:

- be denominated in euros
- have a minimum rating of BBB- or equivalent
- have a remaining maturity of 6 months – 30 years at the time of purchase
- be issued by a corporation established in the euro area; credit institutions not eligible.

The programme was scheduled to start at the beginning of June 2016.

2.2.3 Summary statistics

Table 2.1 and Figure 2.1 present summary statistics of the QE-eligible (546) and ineligible (6,044) bonds.

[Table 2.1 and Figure 2.1 about here]

Figure 2.2 shows the age-rating distribution of all European corporate bonds for comparison. As we can see from Table 2.1, QE-eligible bonds (column QE) are, on average, younger and have larger amounts outstanding, compared to the ineligible ones (column

[5]The programme was extended during the ECB meeting in March 2017.

non-QE). This is true for the mean and the median, and largely for the whole distribution (as seen from the deciles). In addition, QE-eligible bonds are always rated above BBB. This heterogeneity might be driving the differences in bond prices and liquidity. I address this issue in section 2.3.

[Figure 2.2 about here]

Figure 2.3 presents monthly summary statistics of the bilateral trading volume for all European corporate bonds and for the QE-eligible segment separately. As we can see, the distribution of turnover is highly positively skewed, the mean being higher than the median. This is more pronounced for the QE-eligible segment where the difference is of the magnitude 2–4 times. Inspecting the quartiles also confirms the positive skewness, since even the third quartile is often lower than the mean. Table 2.8 in the Appendix illustrates that a similar picture is observed for the tri-party repo market where the positive skewness is even stronger.

[Figure 2.3 about here]

Figure 2.3 illustrates also that the difference between the mean and the median trading volume starts to diverge for the last three bars, particularly for the segment of QE-eligible bonds. This is the first evidence of increased trading in large quantities around March 2016. Presumably, it was the newly launched ECB's QE programme that drove these results.

[Figure 2.4 about here]

The plots of bond yields and trading volume before and after the CSPP in Figure 2.4 support this hypothesis. Before the QE announcement, both bilateral and tri-party repo trading volumes move in parallel trends. However, after the CSPP, we see a large spike in turnover of QE-eligible bonds for both markets. At the same time, we see a drop in yields at the time of the CSPP announcement and in the post-announcement period which suggests QE increased bond prices. To prove that the ECB's intervention caused these changes, however, I test the causality in a more rigorous way using a regression approach.

2.2.4 Research framework

The empirical strategy in this research exploits the CSPP announcement to construct a before-after comparison between bonds that are eligible for purchase and bonds that are not. To study the impact of QE, I take a window around the announcement and employ a difference-in-differences research design. This methodology allows me to avoid confounding the effects of QE with unobserved shocks to the European corporate bond market. Figure 2.4 confirms that the parallel trends assumption for the DD specification is largely satisfied.

The sample period analyzed in this study is January 2016 to June 2016 (23 weeks). To the best of my knowledge, no other significant events influenced exclusively the European corporate bond market in the sample period, so the main driver of changes in liquidity and debt issuance in this market was presumably the ECB's intervention. The period between the announcement of the CSPP on 10 March 2016, and 21 April 2016 (the day when the criteria were revealed) is defined as the interim period, and the period after 21 April 2016 as the post-announcement period. The corresponding regression dummies are *Inter* and *Post*. A bond is defined as being part of the treatment group if it satisfies all the CSPP eligibility criteria. The corresponding dummy variable is *Treat*. The basic regression specification is:

$$y_{it} = \alpha_i + \lambda_t + \gamma_{Treat\times Inter} \times Treat_i \times Inter_t + \gamma_{Treat\times Post} \times Treat_i \times Post_t + \epsilon_{it}, \tag{2.1}$$

where y_{it} is the variable measuring yields (Y_{it}), liquidity (e.g., $Turnover_{it}$, $Turnover\ ratio_{it}$, $Bid\text{-}ask\ spread_{it}$), firms' debt issuance (e.g., $Number\ of\ bonds_{it}$), or firms' financial indicators (e.g., D_{it}). α_i is an individual (bond/firm) fixed effect and λ_t is a time fixed effect. In later specifications, the basic regression is modified by adding controls and rating dummies. The next section studies the impact of the QE programme on bond prices and liquidity.

2.3 Impact of QE on bond prices and liquidity

This part analyzes the first broad topic of the paper, namely, the effects of the ECB's QE on bond prices and liquidity. To estimate the impact on prices, I look at changes in bond yields of eligible versus ineligible bonds after QE. To quantify the effect on liquidity, I measure it within two main dimensions: trading activity and cost of trading. The first one is estimated using weekly turnover (volume of trading) and turnover ratio (volume of trading scaled by the par amount of the bond) as they are commonly used indicators of trading activity in the bond literature (see, e.g., Asquith et al., 2013). The cost of trading aspect of liquidity is quantified using bid-ask spreads as they represent the cost of a round-trip transaction in raw terms.

Table 2.9 in the Appendix also provides regression results using scaled bid-ask spreads: divided by either the bid price (sometimes referred to as the liquidity cost score $LCS = \frac{Ask-Bid}{Bid}$) or by the mid-price (the effective spread $ES = \frac{Ask-Bid}{Mid}$). It also presents the results for price impact indicators (the Amihud illiquidity measure) to quantify liquidity within a third dimension: the ease with which large quantities can be traded on the market. The main conclusions of this research also hold with these measures of liquidity.

First, I analyze the impact of QE on bond yields. Columns 1–3 of Table 2.2 estimate the basic regression (2.1) and a modification without time and bond fixed effects:

$$\begin{gathered} y_{it} = \alpha + \gamma_{Treat} \times Treat_i + \gamma_{Inter} \times Inter_t + \gamma_{Post} \times Post_t + \\ \gamma_{Treat \times Inter} \times Treat_i \times Inter_t + \gamma_{Treat \times Post} \times Treat_i \times Post_t + \epsilon_{it}. \end{gathered} \tag{2.2}$$

The dependent variable is yields (Y_{it}). The DD estimates remain robust to adding different control variables and reveal the causal impact of the ECB's intervention on bond prices. Yields of eligible bonds were reduced in the post-announcement period by 30 bps (an 8% decrease given the control group before-QE mean of 371 bps). The reduction of 22 bps in the interim period is also statistically significant. The negative coefficients on $Inter$ and $Post$ from column 2 suggest that there was an overall decrease in yields for all bonds of 21 bps in the interim period and of 31 bps in the post-announcement period. However, these coefficients are only marginally significant. Using the estimates from the second column, the total effect on yields of QE-eligible bonds was approximately 43 (22+21) bps in the interim period and 60 bps in the post-announcement period. These are significant reductions (11% and 16%, respectively).

The effects on liquidity are also statistically and economically significant. Columns 6–9 of Table 2.2 show that liquidity improved in terms of trading activity. Bilateral turnover went up in the post-announcement period by 7.05 million USD (a 72% rise given the control group before-QE mean of 9.69 million USD). Tri-party repo market turnover went up by 8.15 million USD (a 29% increase given the control group before-QE mean of 28.14 million USD) during the post-announcement period. The larger increase (in relative value) in the bilateral market suggests that investors mostly engaged in buy-sell transactions around the time of the announcement, and a relatively smaller fraction of bonds was used as collateral for borrowing. However, the increase in trading volume might have been driven by the heterogeneity of trading large and small bonds (as measured by the par outstanding). To address this concern, I use turnover scaled by the total par amount as an alternative liquidity measure. Columns 8–9 of Table 2.2 present the result of regression (2.1) using this approach. Now, the increase in the turnover ratio for the repo market also becomes significant for the interim period.

QE also decreased the cost of trading bonds. The estimates from column 10 of Table 2.2 show that bid-ask spreads were reduced by 2 bps in the interim period and by 6 bps in the post-announcement period. These are massive reductions of 15% and 46%, respectively, given the control group before-QE mean of 13 bps. Overall, the results from Table 2.2 support the hypothesis that the ECB's intervention increased bond prices and liquidity significantly.

[Table 2.2 about here]

2.3.1 Robustness checks and alternative liquidity measures

So far, the control group consisted of all bonds that did not meet any of the eligibility criteria set by the ECB. However, since the treatment and control groups are different by design, this heterogeneity might have driven my results. To address this issue, I construct a matched sample, restricting the treated sample to bonds for which there is a suitable control bond with similar pre-treatment characteristics.

The bond characteristics used to construct the matched sample are: par amount outstanding, credit rating just before the announcement, and time to maturity just before the announcement. The whole sample is divided by par amount outstanding into four quartiles. For the rating, it is split into four groups. The first group consists of only AAA and AA+ rated bonds, the second of AA, AA-, A, A-, A+ rated bonds, the third of BBB, BBB-, and BBB+ rated bonds, and the last consists of all bonds with ratings lower than BBB-. For the time to maturity, bonds are divided on above and below the median time to maturity. This division results in 32 potential cells. Obviously, there will be no QE-eligible bonds in rating group four, which leaves 24 potential groups.

The estimates for the matched sample DD yield regression are in column 4 of Table 2.2. To control for bond attributes, I add a dummy variable for each cell to the regression (2.2), and interact the cell dummy with $Post$, $Inter$, and $Treat$ dummies. Their inclusion means that the estimates are a weighted average of the within-cell DD estimates. The coefficients for both the interim and the post-announcement period remain negative and significant and have roughly the same magnitude as before.

Another robustness check that I conduct is to define the control group as all bonds that satisfy the CSPP eligibility criteria (except for those in EUR). This leaves a total of 1,991 bonds. The results are in column 5 of Table 2.2. As we can see, the main conclusions still hold. Thus, the significant effects of the QE programme documented in column 3 of Table 2.2 are robust to alternative specifications of the control group.

In the Appendix (Table 2.10), I also present similar robustness checks for the liquidity measures. The estimates for the turnover measures and the bid-ask spreads also remain statistically significant in these regressions. Table 2.9 in the Appendix shows the results from estimating the impact of QE on alternative liquidity measures. As we can see from the tables, the bid-ask spread of prices, the effective spread, and the liquidity cost score also go down, indicating that the cost of a round-trip transaction has decreased as a result of QE. The Amihud illiquidity measure illustrates that both bilateral market illiquidity and tri-party repo market illiquidity decreased substantially. This fact shows that liquidity also improved in terms of market depth.

In a final robustness check, I collapse the data by time period and run pure time-series tests. The results are presented in the Appendix, Table 2.11. The estimates from columns 1–6 show that the main conclusions from this section still hold and that the estimates are similar in magnitude.

2.3.2 Does liquidity revert back after the initial shock?

The plot of trading volume in Figure 2.4 shows that the initial increase for QE-eligible bonds at the time of the CSPP announcement started to dissipate in the following few weeks. This raises the concern that liquidity changes could reflect temporary responses as new information causes a short-lived spike in trade.[6] To trace whether the trading volume reverts back, I extend the post-announcement analysis by four months.[7] Figure 2.9 in the Appendix presents the results. It shows that turnover gradually declined after the initial spike. The reversion was more pronounced for the bilateral market, where the jump in trading volume following the CSPP announcement was larger in magnitude. By August 2016, the bilateral trading volumes of QE-eligible and QE-ineligible bonds were quite close. The picture for the repo market was slightly different as QE-eligible bonds were traded more frequently than ineligible bonds for a longer time period. Eventually, the trading volumes converged around the end of September 2016.

These results are consistent with the following buy-and-hold story. Following the announcement of the programme, investors started trading QE-eligible bonds intensively, causing a spike in trading volume. However, once eligible bonds were purchased, there was a gradual dissipation of trading which suggests that investors did not sell eligible

[6] I would like to thank the referee for pointing this out.

[7] A possible caveat is that extending the analyzed period further away from the initial shock could cast doubt on the conclusion that the observed differences in turnover are only due to QE.

bonds or did not trade these bonds in a significantly different manner compared to ineligible ones. These effects were more pronounced for the bilateral market as the turnover reversion was much faster there, compared to the repo market. Another explanation for the decline in trading volume could be the fact that, since the ECB was buying eligible bonds, these bonds were removed from the market, and therefore turnover decreased.

[Table 2.3 about here]

The results from Table 2.3 confirm the gradual decline in turnover. The magnitude of the estimates is smaller compared to Table 2.2, which suggests that there was a decrease in trading after the initial shock. The estimates for the bilateral turnover are insignificant, but the coefficients for the repo turnover in the post-announcement period and the coefficients for the turnover ratios are still significant. This fact shows that the overall improvement in liquidity (measured by trading volume) of QE-eligible bonds documented from before, is robust to extending the post-announcement period, albeit with much weaker effects. In contrast, the estimates for prices (columns 5 and 6 of Table 2.3) are very similar to the ones from Table 2.2, which shows that the effects on prices are more robust. Yields of eligible bonds were reduced by 26 bps in the post-announcement period compared to 30 bps from before. The results for bid-ask spreads are even slightly stronger – the coefficient of the $Treat \times Post$ dummy suggests a 7 bps decline in bid-ask spreads over the longer post-announcement period. This finding is again consistent with improved liquidity (measured by the cost of trading) of QE-eligible bonds. I also ran the regression for bonds with different ratings and for alternative liquidity measures over the extended post-announcement period. The results were similar to the ones for the main period analyzed and were excluded from the paper for brevity.

2.3.3 Does QE reduce risks?

This part studies the heterogeneous QE effects for bonds with different risks. I decompose the reduction in yields into different risk premia by applying the methodology of Krishnamurthy and Vissing-Jorgensen (2011).[8] Similar to them, I use DD regressions to isolate the duration risk channel, the default risk channel, and the liquidity channel of QE. To disentangle the different risks, I split the sample of bonds into four maturity groups: 0–2 years, 2–5 years, 5–10 years, and more than ten years,[9] and then, within

[8]I would like to thank the referee for this excellent suggestion.

[9]The results are similar if I split bonds into quartiles or into three groups with a similar number of bonds in each: 0–2 years, 2–5 years, more than five years. The split of more than ten years is to isolate the effect on the longest-maturity corporate bonds.

each maturity group, into three rating buckets: AA, A, and BBB. All bonds within a particular maturity-rating bucket then have similar duration and default risks. To test for heterogeneous effects of QE on bonds with different risks, I run the following regression for each maturity-rating group:

$$y_{it} = \alpha_i + \lambda_t + \gamma_{Treat \times Post2} \times Treat_i \times Post2_t + \epsilon_{it}, \qquad (2.3)$$

where $Post2$ is a dummy equal to one after 10 March 2016 ($Inter$ + $Post$). The dependent variable is yields or bid-ask spreads.

To quantify the default risk channel of QE, I compare the yield spread of eligible versus ineligible bonds after the CSPP, within each maturity bucket. Analogously, to isolate the duration risk channel, I compare the yield spread of eligible versus ineligible bonds after QE, within each rating group. To isolate the liquidity channel, I look at the change in bid-ask spreads for eligible versus ineligible bonds after QE, within each maturity-rating bucket. The results are summarized in Table 2.4. Rows are different default risk groups, columns – different duration risk groups. Panel A shows the effect on yield spreads, Panel B on bid-ask spreads. The numbers in the table are the estimates of regression (2.3) for bonds within a particular maturity-rating group. For example, -0.381 is the coefficient on the $Treat \times Post2$ dummy for the sample of bonds rated AA, with a maturity of more than ten years. It shows a 38.1 bps reduction in yields for QE-eligible bonds in this category, compared to ineligible bonds in the same category, after QE. In relative terms (Table 2.13 in the Appendix), this is a 15.25% reduction. Firm and time fixed effects are used in the regressions to account for individual and time-specific risks. F-statistics and number of observations are in Table 2.14 in the Appendix.

[Table 2.4 about here]

The point estimates in Table 2.4 illustrate an interesting pattern. If we fix a row, we can see that within each rating group, the statistically significant coefficients are monotonically decreasing from bonds with lower maturity to bonds with higher maturity. This result illustrates that QE decreased yields of longer-term bonds more, compared to those of short-term bonds. The fact is consistent with KVJ's prediction and shows that QE had a positive effect on reducing the duration risk of bonds. The magnitude is the largest for lower-rated bonds close to the eligibility threshold (BBB). Moreover, if we fix a column, we can see that the significant estimates within each maturity group are also monotonically decreasing from bonds with higher ratings to bonds with lower ratings. This fact shows that QE had a positive effect on reducing the default risk premium. In other words, the safety premium channel during the ECB's QE worked in the opposite direction to KVJ's prediction for the Federal Reserve's QE: the estimates in Table 2.4 show that the CSPP reduced yields of safer bonds less, compared to yields of riskier bonds. However, this result is not surprising given the differences between the ECB's and the Fed's programmes. The prediction of the KVJ study is for QE involving very safe assets, whereas the CSPP included riskier corporate bonds.

To investigate the conclusion that lower-rated bonds experienced a larger drop in yields, I restrict the sample to only treated (QE-eligible) bonds and run a difference-in-difference-in-differences (DDD) regression within each maturity group:

$$y_{it} = \alpha_i + \delta_{(Treat_i=1)t} + \gamma_{Post2_{BBB}} \times Treat_i \times Post2_t \times BBB_i + \epsilon_{it}. \tag{2.4}$$

Hence, the reference category in each maturity group is treated bonds with ratings higher than BBB. If the conclusion is true, the coefficients on the interaction terms should be negative. The results are presented in the last row of Panel A. As we can see, the estimates are negative and monotonically decreasing across maturity groups which shows that within each duration risk group, eligible lower-rated bonds close to the threshold experienced the largest reduction in yields.

The statistically significant bid-ask spread estimates from Panel B suggest that BBB-rated bonds also experienced the largest increase in liquidity within most maturity groups. However, the difference in liquidity for AA-rated compared to A-rated bonds is not clear. Moreover, there is no obvious pattern of liquidity improvement for bonds with different maturities within each rating group – the estimates are monotonically decreasing for A-rated bonds only. Furthermore, the coefficients for BBB-rated bonds within the sample of QE-eligible bonds (the last row of Panel B) show that there are minor statistically significant differences in liquidity for longer-maturity bonds. Running the DDD regression for all BBB-rated bonds (combining all maturities) shows that there are no significant liquidity changes to these bonds compared to AA-rated or A-rated bonds. The results are excluded for brevity.

In the Appendix (Table 2.15), I repeat the same analysis but split the sample of bonds into three dimensions: duration risk, default risk, and also liquidity risk. The split in liquidity is done to isolate possible pre-existing liquidity differences for bonds with similar duration and default risks. Within each maturity-rating risk bucket bonds are further sorted into illiquid (above median bid-ask spread) and liquid (below median bid-ask spread) based on their bid-ask spreads before the QE announcement. There are fewer bonds within each maturity-rating-liquidity risk bucket and less statistically significant results but the general picture is unchanged. This fact shows that the effects in Table 2.4 are not due to pre-existing differences in liquidity. Table 2.16 in the Appendix presents the p-values of t-tests on coefficient differences for various rating-maturity groups. The tests strongly reject the hypothesis that the coefficient for BBB-rated bonds in the yield regressions is the same as the coefficient for AA-rated bonds for maturities 2–5 years, 5–10 years, and more than ten years. Panel B of the table shows that we cannot reject the hypothesis of equal liquidity changes for bonds with different ratings within the longest-maturity group.

The results from this subsection show that QE had heterogeneous effects on bonds with different risks. The largest impact of QE was observed for prices of bonds with higher duration and default risks. The liquidity effects, on the other hand, were only slightly different for bonds with different duration and default risks.

2.3.4 Discussion of the results

The substantial decrease in duration and default risks and the increase in liquidity of QE-eligible bonds documented in this section means that QE was successful in reducing the bond risk premium. This is an important result for market participants as during market crashes, this premium can become a substantial component of yields, causing a decrease in bond prices. As a consequence, banks holding risky illiquid assets could become constrained by their risk bearing capacity as in Vayanos and Vila (2009), or face VaR constraints. Hence, their ability to invest in riskier instruments would be reduced and bank supply would go down. The idea of QE is then to boost economic activity by expanding the list of eligible collateral and by stimulating investors to switch to riskier investments (portfolio rebalancing channel). As the central bank steps in and implements asset purchases, it eliminates risk by reallocating the risky illiquid assets from private banks' balance sheets to the central bank's balance sheet, thereby relaxing the VaR constraint and decreasing the market price of risk.

The decrease in yields has two main beneficial effects. First, it allows investors to dispose of the risky QE-eligible corporate bonds by selling them to the central bank. Essentially, the increase in asset prices that QE generates is akin to a capital injection for leverage-constrained institutions. Moreover, it would allow new investors to increase their holdings of QE-eligible securities because of the reduced duration and default premiums and the fact that these securities can be sold to the central bank in the future. The results of this section show that both the price impact measures and the cost of trading of QE-eligible bonds went down indicating that investors could dispose of these bonds easier and incur lower costs in doing so. The surge in trading volume shows increased activity in QE-eligible bonds which might be related to both disposition effects (see, e.g., Koijen et al., 2017) and to new inflows into QE-eligible bonds. However, in contrast to QE-eligible bonds, the impact of QE on premiums of ineligible instruments was much weaker. Since the portfolio rebalancing channel towards riskier instruments was one of the main goals of QE, we can conclude that the ECB's intervention achieved this goal in the subset of eligible instruments, but that the effect on ineligible bonds was less pronounced.

Second, the QE intervention would directly benefit credit-constrained corporates who could finance more cheaply on the bond market by issuing eligible bonds. The fact that the QE effects on yields were very heterogeneous across different maturity-rating groups but liquidity effects were roughly similar suggests that price changes had a first order effect on corporates' decisions to issue new bonds, whereas liquidity changes had a second order impact. The result that most of the decrease in yields within the QE-eligible segment was driven by lower-rated bonds suggests that the CSPP had a larger positive impact on riskier debt instruments and for lower-rated corporates. As we will see in the next section, these corporates made wider use of the programme.

2.4 Impact of QE on corporate debt issuance

This section analyzes the reaction of corporates to QE. When the central bank intervenes by buying large amounts of corporate bonds, this presumably makes it easier for companies to raise capital. Essentially, the ultimate purpose of QE in this respect is to boost aggregate investment by lowering the cost of market debt issuance. The policy would allow firms that were previously unable to receive financing due to either bad quality, or asymmetric information problems, to become eligible for direct financing through the ECB. In other words, bonds issued by these firms could now be sold to the ECB, and the boost in demand by the central bank should have incentivized credit-constrained firms to start issuing more euro-denominated bonds that fulfilled the CSPP criteria after the QE announcement.

2.4.1 Issuance of new bonds

To detect an issuance of a new bond, I compute the age of each security on each date in the sample period. In each week, I then filter for bonds which are less than one week old; and this will reveal the newly issued bonds. Figure 2.5 shows the evolution of the number and the par of newly issued bonds in the sample period. As we can see, there was an increase both in the absolute number and in the notional amount of newly issued QE-eligible bonds after the CSPP announcement.

[Figure 2.5 about here]

To formally test the hypothesis that firms made use of the new cheaper funding opportunity resulting from the CSPP, I run regression (2.3) with the number and the par of newly issued bonds as the dependent variables. The results are presented in Table 2.5. As we can see from the first column of Panel A, firms issued 1.18 QE-eligible bonds more per week compared to ineligible bonds, after the CSPP. This is a 12% increase, given the mean before QE for ineligible bonds of 9.70 new bonds per week. The coefficient from column 1 of Panel B suggests that firms started issuing more QE-eligible debt after the CSPP announcement not only in number of bonds but also in notional amount. The estimate shows a 2.19 billion EUR rise of newly issued eligible debt compared to ineligible one in the post-announcement period. This is a significant increase of nearly 25%. These results suggest that corporates started to issue more QE-eligible bonds after the CSPP announcement compared to other types of debt.

[Figure 2.6, Figure 2.7 and Figure 2.8 about here]

However, firms issuing QE-eligible bonds might be systematically different from those issuing QE-ineligible bonds. To rigorously test the hypothesis that firms made use of the cheaper funding opportunity, and switched to issuing more euro-denominated QE-eligible bonds, I saturate the firm dimension in the following way. I search for firms that issue bonds in different currencies (at least one of them being the euro). There are 254 such firms (out of 1,550). If there was a switch to issuing more QE-eligible debt, these firms would be more likely to issue a larger amount of their debt in euro-denominated bonds that satisfy the CSPP criteria, and less in other types of bonds. To see the cumulative effect on young bonds, I plot the total number of young bonds before the CSPP announcement, in the interim period, and following it. Figure 2.6 illustrates a clear increase in the number of newly issued QE-eligible bonds, which was mostly caused by issuance from BBB-rated firms. Figure 2.7 shows that a similar picture is observed for the par amount. Figure 2.8 presents the results with the euro sample split into QE-eligible and QE-ineligible bonds to show the differential effects of QE. Figure 2.10 in the Appendix does the same for quarterly frequency to see the cumulative effect.

The figures show that firms started issuing more debt in euro-denominated bonds both in absolute numbers, and in the total amount of debt. There is an increase in the par amount of QE-eligible euro bonds in the post-announcement period, but at the same time a decline in issuance of euro QE-ineligible bonds. The total amount of euro-denominated new debt increased slightly, and the difference with other currencies went up from the start of the interim period. These facts serve as additional evidence in favour of the hypothesis that firms issued more euro-denominated QE-eligible debt as a result of the CSPP.

The DD estimates confirm these observations. The second column of Panel A in Table 2.5 shows that firms issuing in several currencies, switched to more QE-eligible debt. Corporates issued 1.39 QE-eligible bonds more per week after the CSPP, compared to the number of ineligible bonds. This is a 38% increase given the mean before QE for ineligible bonds of 3.70. The coefficient from column 2 of Panel B illustrates that the increase in notional value was also significant – firms issued 1.94 billion EUR more of QE-eligible debt compared to ineligible one, which is a 49% increase. Column 3 of both panels shows that corporates issued more euro-denominated bonds in general.

Next, I investigate the hypothesis that more credit-constrained firms would issue more bonds as a result of the programme. The rating of the firm is used as a proxy for credit constraints. Although not a perfect measure, rating could serve to check which firms made wider use of the cheap financing opportunity. The logic is as follows: companies in rating categories below A that can issue eligible bonds (mostly BBB, BB) face more severe asymmetric information problems since there is a larger firm heterogeneity in these rating categories. Moreover, these companies obtain credit under less favourable conditions compared to firms rated AAA, AA, or even A. The fraction of firms within each rating category and the fraction of QE-eligible bonds issued by a firm in a certain category is in Table 2.17 in the Appendix. As we can see, roughly 97.71% of the QE-eligible bonds were issued by firms rated AA, A, or BBB. The regression estimates from columns 4–6 in Panels A and B of Table 2.5 support the hypothesis that firms with lower ratings (more credit-constrained) made wider use of the cheap funding opportunity: the coefficients rise from A to BBB firms.[10] BBB-rated firms issued on average 0.59 eligible bonds more each week compared to ineligible ones, after QE. In terms of par amount, they issued 0.77 billion EUR more in eligible debt. These are massive increases: 65% in number of bonds and 76% in notional terms compared to the mean of ineligible BBB-rated bonds before QE. Next, I run a DDD regression for the sample of treated

[10]The estimates for BB-rated firms were not significant and are excluded for brevity. Running the regressions for the whole sample of firms (Table 2.18 in the Appendix) produces similar results.

bonds:

$$y_{it} = \alpha_i + \lambda_t + \sum_{k=1}^{3} \gamma_{Treat \times Post2,k} \times Treat_i \times Post2_t \times Rating_k + \epsilon_{it}. \quad (2.5)$$

The results are even stronger. The coefficients from columns 11–13 show that the increase for AA-rated firms becomes significant and that the estimates rise from AA to A to BBB firms. Since most of the firms that can issue QE-eligible bonds are rated AA, A, or BBB, and only 2.29% are rated BB, this fact shows that there is a monotonic inverse relationship between the rating of a firm and the amount of issued QE-eligible debt. BBB-rated firms are the ones that issued the most after the CSPP.

These results are not surprising given the evidence from Table 2.4. Since lower-rated eligible bonds experienced the largest drop in yields, one could expect firms with lower ratings to make wider use of the programme by issuing more bonds. This is indeed what we observe in Table 2.5. Another conclusion from Table 2.4 was that yields of longer-maturity bonds decreased the most. One could hypothesize then that firms would be more likely to issue bonds with longer maturities. To test for this, I run regression (2.3) for the number and the par of bonds within each maturity bucket. The coefficients from columns 7–10 in Panels A and B of Table 2.5 suggest that firms were more likely to issue longer-maturity bonds. The estimates for bonds with short and medium maturities (0–2 years, 2–5 years) are insignificant but the ones for longer maturities are positive and statistically significant. The coefficient for bonds with maturities of 5–10 years is the largest and shows that firms issued 0.58 (36%) eligible bonds more per week compared to ineligible ones. In terms of notional, they issued 1 billion EUR (63%) of new QE-eligible debt more each week compared to ineligible one. These facts show that the duration risk channel of QE incentivized firms to issue more bonds with longer maturities since these bonds experienced the largest drop in yields.[11]

[Table 2.5 about here]

A possible caveat is that, although the regression results show that firms made use of the CSPP and issued more QE-eligible debt, it might also be that firms took even more bank credit, or used other sources of debt financing. That would cast doubt on the story developed so far. However, I rule out this possibility by looking at the ratio

[11] I also performed t-tests on coefficient differences for various rating and maturity groups similar to the ones in Table 2.16 in the Appendix. The t-tests strongly rejected the hypothesis of equal increase in issuance between A and AA bonds, and between BBB and AA at the 1% significance level for the sample of QE-eligible bonds. The tests also rejected the hypothesis of equal increase in issuance between bonds with short maturities (0–2 years) and long maturities (more than ten years) at the 5% level. These tests were excluded from the paper for brevity.

of bond debt outstanding to total debt for each firm. The average ratio at the end of the first quarter of 2016 was 73.41%, whereas at the end of the second quarter of 2016 it was 81.07%. This observation suggests that firms increased not only their total amount of bond debt, but also the fraction of bond debt in total liabilities. This is consistent with the findings of Gerba and Macchiarelli (2016): banks surveyed in their paper claimed that "the extra liquidity they[banks] receive has basically no impact on their decisions to grant (or not) loans." The next step in the analysis is to see what corporates do with the financing raised through the issuance of new QE-eligible bonds. The answer to this question is not as straightforward because there are many possible caveats and it is hard to claim that the observed effects are only due to the ECB's intervention. I provide some analysis that attempts to answer this question, but I am aware of the different effects at play and of the possible limitations.

2.4.2 What do corporates do with the funds?

The results in this section are based on data from the Bureau van Dijk on firms' dividends (D), total assets (A), fixed assets (FA), tangible fixed assets (TFA), long-term debt (LTD), R&D expenses (RD), property, plant, and equipment (PPE), working capital (WC), and cash holdings.[12] Out of the 1,550 distinct firms in my sample, only 417 have data on these characteristics and even less have data for the time period around the CSPP-announcement. A further limitation of the firms' financial data is that it is only observed at quarterly frequency compared to the weekly frequency of the Euroclear data. Another effect is that the firms in the sample are different along many dimensions (size, age, parent/daughter companies, etc.) and it is hard to argue that the only difference in observed responses is due to the issue of QE-eligible bonds. With these caveats in mind, I proceed to the empirical analysis.

I use the sample period from the first quarter of 2016 to the third quarter of 2016 and define the post period as the second and third quarters of 2016. Firms that have issued at least one QE-eligible bond are in the treatment group. The results are robust to taking only the third quarter of 2016 and to taking only firms issuing after the announcement. The basic regression (2.1) is run with firms' financial data as the dependent variable. As Table 2.6 shows, most $Treat \times Post$ coefficients are insignificant, except the one for dividends. It might be true that firms that issue QE-eligible bonds are just larger, on average, which could partly explain the difference in the change of dividends in the post-announcement period. To address this issue, I run the regression

[12] Unfortunately, data on capital expenditures and share repurchases were available for only a tiny set of firms. There were no statistically significant effects for these indicators and the results are excluded for brevity.

on a sample of firms matched by total assets by splitting the firms in four quartiles based on their total assets in the first quarter of 2016. The regression coefficient (column 2) is now significant at the 1% level and shows a four times increase in dividends for firms that issued QE-eligible bonds, compared to the before-QE mean of firms not issuing QE-eligible bonds.

[Table 2.6 about here]

These findings suggest that corporates used the funds attracted through the issuance of QE-eligible bonds for one-off dividend payments to current shareholders. They did not, however, change cash holdings or invest the funds in working capital, PPE, or R&D. One of the goals of the QE programme was to incentivize corporates to use the attracted funds to increase investment in the real economy. To the contrary, the figures in Table 2.6 show that the funds were not used for this purpose in the sample of firms studied. However, these results should be interpreted with caution, given the limitations of the Bureau van Dijk data and the caveats outlined above.

Overall, the analysis in this section suggests that corporates made use of QE and issued more euro-denominated, QE-eligible debt. These results complement the literature on corporate debt issuance. Consistent with the conclusions of Adrian and Shin (2014), Becker and Ivashina (2014), and De Fiore and Uhlig (2015), I find that there is a switch to issuing more corporate debt in the aftermath of the financial crisis. The new result is that these conclusions continue to hold even during a QE intervention and in a low-interest rate environment. Moreover, I show that more credit-constrained firms made wider use of the bond financing opportunity during a QE implementation. These new empirical findings suggest that QE was successful in providing cheap liquidity to more credit-constrained corporates, but at the same time had greater impact on firms' dividend payments and little impact on investment.

2.5 Placebo tests

As a final check, this paper performs two placebo tests to verify that the sample selection of bonds in treatment and control does not mechanically produce the above results. In the first test, the sample period is shifted backwards by one year to remove any effects of seasonality. A bond belongs to the treatment group if it satisfies all the QE-eligibility criteria. Similarly, the post-announcement period is defined as after 21 April 2015, as though there was a CSPP announcement on that day. Regression (2.1) is estimated using yields, bid-ask spreads, bilateral turnover ratio, tri-party turnover ratio, number and par of newly issued bonds as dependent variables. The results of the estimation are in Table 2.7, Panel A. As we can see, none of the DD estimators are statistically significant.

The second placebo test uses USD-denominated bonds that satisfy all the CSPP eligibility criteria as a treatment group. These bonds share similar characteristics with the ones classified as QE-eligible earlier (maturity of less than 30 years, rating above BBB, issued by a non-credit institution, etc.) and are also similar in total numbers (479 compared to the 546 euro-denominated, QE-eligible bonds). However, the USD bonds should not be affected by the ECB's announcement and hence should not exhibit any significant pattern around the CSPP, unless the effects I found for the QE-eligible bonds were due to some other characteristic common for all bonds satisfying the eligibility criteria. The results of regression (2.1) using yields, bid-ask spreads, bilateral turnover ratio, tri-party turnover ratio, number and par of newly issued bonds as dependent variables are in Table 2.7, Panel B. As we can see, the DD estimates are again not statistically significant.

[Table 2.7 about here]

The results of the second placebo test are robust to alternative specifications of the treatment group. I ran tests with different currencies, with a broader set of USD bonds, relaxing some of the eligibility criteria, and even with all USD-denominated bonds. None of the DD estimates turned out to be statistically significant. These tests are excluded from this paper for brevity. Overall, the placebo tests rule out the possibility that the positive effects on bond yields, liquidity, and firms' debt issuance, found in the main text, were due to seasonality, or mechanical sample selection.

2.6 Conclusion

This paper analyzed the impact of an unconventional QE programme on bond prices, bond liquidity, and corporate debt issuance. The results show that the ECB's CSPP decreased yields and increased liquidity in the European corporate bond market, especially in the QE-eligible segment. Yields of eligible bonds decreased by 30 bps (8%) in the post-announcement period and bid-ask spreads went down by 6 bps (45%). Tri-party repo turnover increased on average by 8.15 million USD (29%) in the post-announcement period, and bilateral turnover increased by 7.05 million USD (72%).

All these effects are statistically and economically significant. They were particularly pronounced for QE-eligible, lower-rated, longer-maturity bonds. These bonds experienced the largest drop in yields, which means that the ECB's intervention had a higher positive impact on riskier debt instruments. However, there was limited evidence for a significant shift towards ineligible, riskier bonds during the analyzed period. Since the portfolio rebalancing channel towards riskier instruments was one of the main goals of QE, I can conclude that the ECB's intervention achieved this goal in the subset of eligible instruments, but that the effect on ineligible bonds was less pronounced. The increases in prices and liquidity show that QE was successful in reducing the bond risk premium.

As to firms' reaction to QE, the study finds that after the significant reduction in the cost of bond issuance, corporates did make use of the cheaper financing opportunity. Firms issuing in several currencies switched to more QE-eligible debt and euro-denominated debt in general, after the CSPP announcement, compared to debt in other currencies. These effects were more pronounced for more credit-constrained firms and for longer-maturity bonds. Surprisingly, corporates used the funds attracted through the issuance of QE-eligible bonds mostly for increasing dividend payments, but did not change cash holdings or invest the funds in working capital, PPE, or R&D. These results are new and show that the QE programme achieved its goal of providing cheaper credit to more credit-constrained corporates, but did not incentivize firms to increase investment. Possible extensions to this paper might include analyzing the impact of QE on firm's default probability for firms that can issue QE-eligible debt. It would also be interesting to see whether the ECB's intervention actually forced investors "searching for yield" (e.g., pension funds, insurance companies) to buy more ineligible bonds in the long run.

There are many important under-researched questions about the economic effects of QE. Our knowledge of the impact of such unconventional monetary policy tools is constrained because they are relatively new and unexplored. Some critics even believe that the policy acts in the opposite direction: it provides monetary stimulus in the short run, but creates instability and exacerbates market distortions in the long run. The analysis presented here provides evidence on the relative efficiency of QE in the short term. However, as time goes by and new information becomes available, we would be in a better position to assess the impact and to fully understand the consequences of the era of unconventional monetary policy in which we are currently living.

2.7 Figures

FIGURE 2.1: Fraction of bonds by rating category. The figure shows the rating distribution of QE-eligible and QE-ineligible bonds just before the CSPP announcement on 10 March 2016. The number of bonds is 6,590 (546 eligible and 6,044 ineligible). Rating group junk includes bonds rated below B.

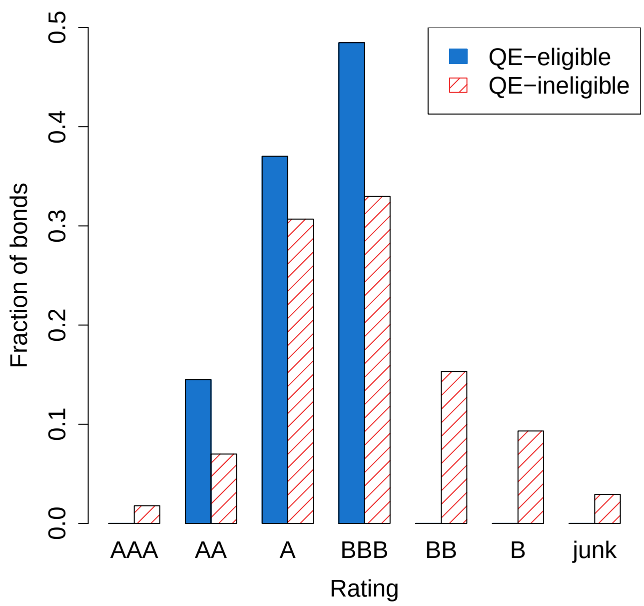

FIGURE 2.2: Age-rating distribution of all European corporate bonds just before the CSPP announcement on 10 March 2016. The number of bonds is 6,590 with a total face value of 4.42 trillion EUR. Age is in years. Rating group junk includes bonds rated below B. The area of each rectangle represents the fraction of bonds in the respective age-rating group. The area of the entire square is 100%.

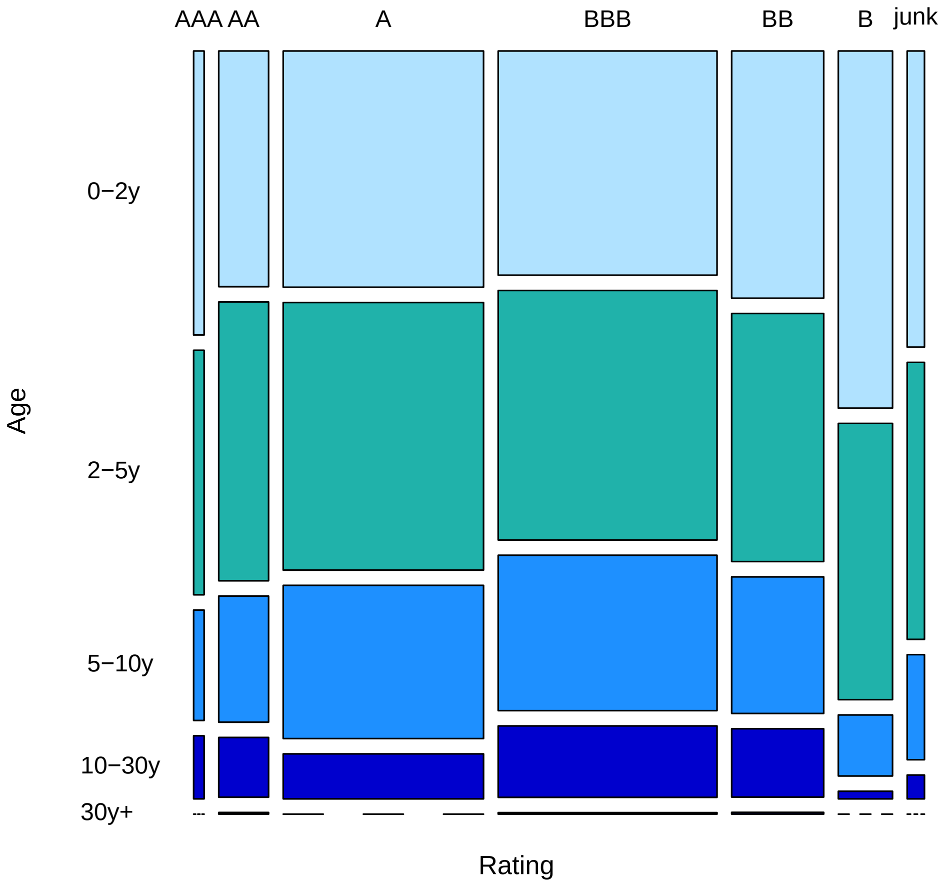

FIGURE 2.3: Monthly bilateral turnover. Quantiles' dynamics: July 2015–June 2016. The top panel shows the turnover of all European corporate bonds (6,590 bonds), the bottom panel illustrates the turnover of QE-eligible bonds only (546 bonds). The dot corresponds to the mean, and the bold horizontal line corresponds to the median bilateral turnover of a given month. The bottom and top borders of each bar show the first and the third quartile of the monthly bilateral turnover, respectively. The last three bars indicate the CSPP period.

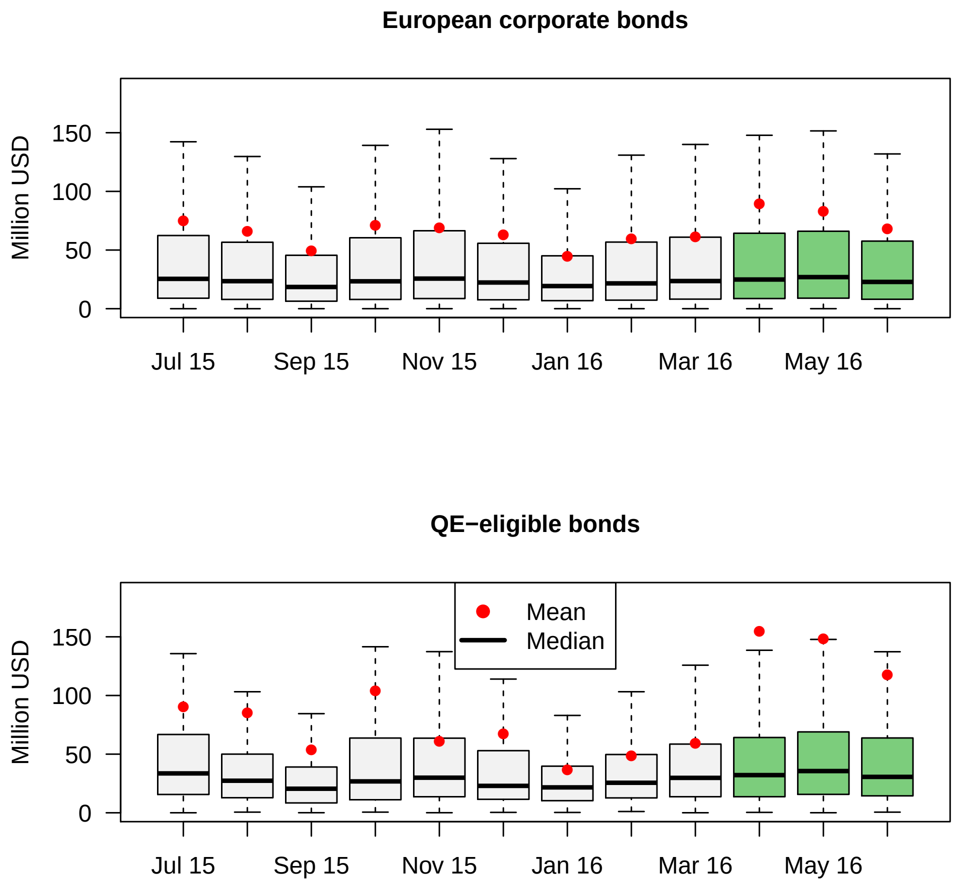

FIGURE 2.4: QE-eligible and QE-ineligible bonds: January 2016–June 2016. The top panel shows the daily time series of mean bond yields. The middle and the bottom panels illustrate the weekly dynamics of mean turnover. The first vertical dashed line indicates the day after 10 March 2016 (CSPP announced), the second dashed line shows 21 April 2016 (eligibility criteria made public). Units for turnover: USD.

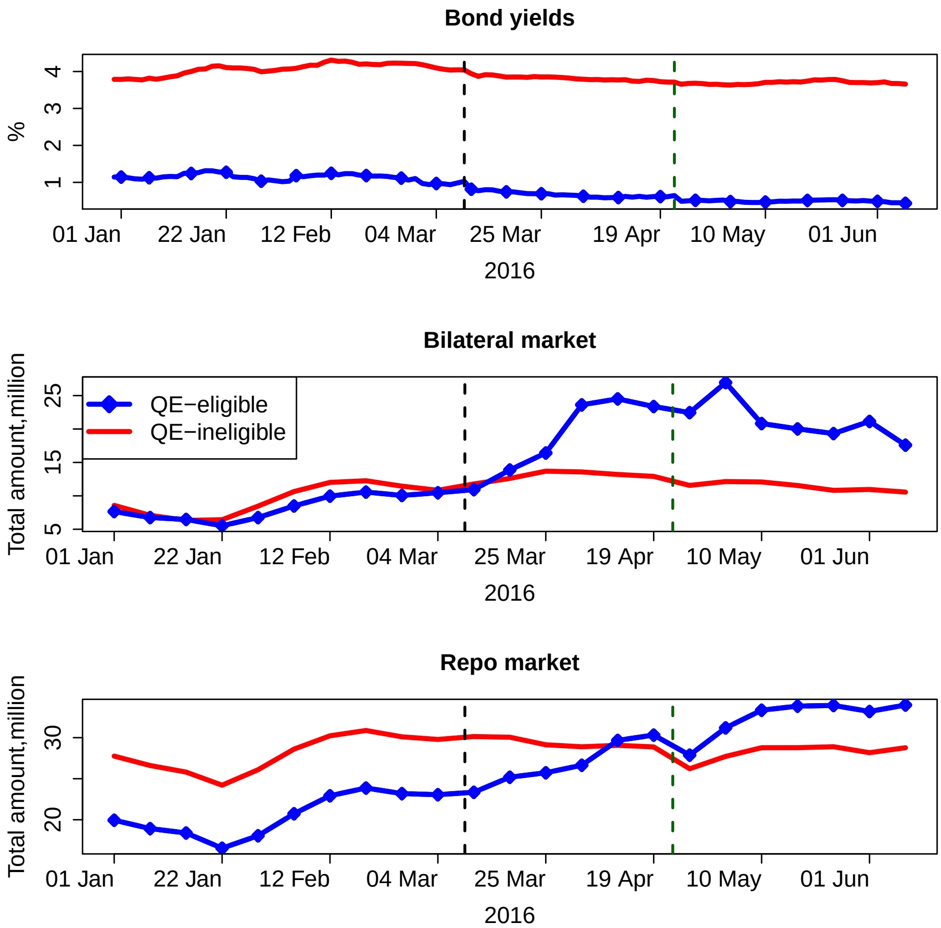

FIGURE 2.5: Weekly dynamics of the number and the par of newly issued bonds: QE-eligible compared to QE-ineligible. Sample period: January 2016–June 2016. The first vertical dashed line indicates the day after 10 March 2016 (CSPP announced), the second dashed line shows 21 April 2016 (eligibility criteria made public). Par is the notional amount of bonds outstanding, measured in billion EUR.

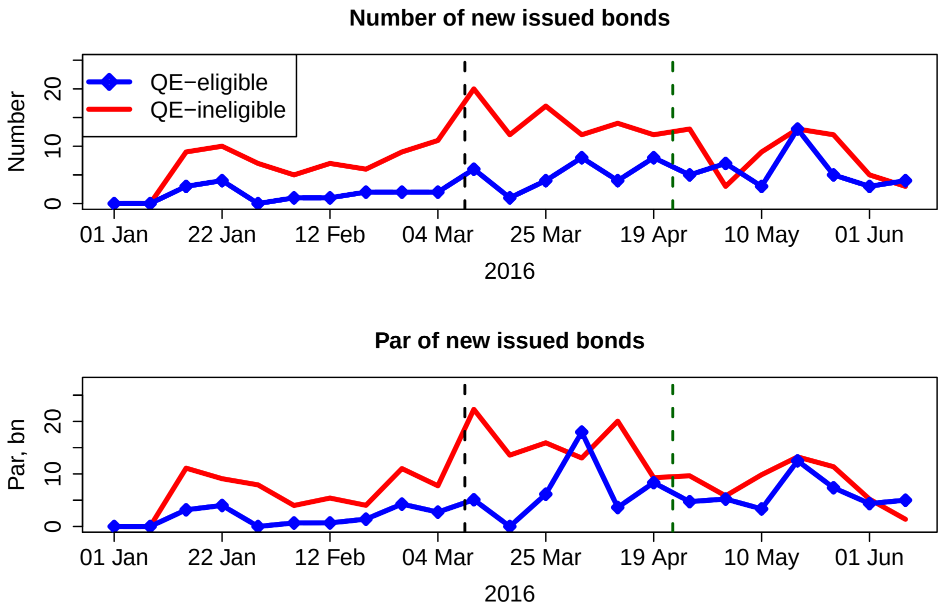

FIGURE 2.6: Number of new bonds issued by firms issuing in several currencies: January 2016–June 2016. The top panel shows the total number of newly issued eligible and ineligible bonds. The mid and bottom panels present the results split by rating categories. The first dashed line indicates the day after 10 March 2016 (CSPP announced), the second dashed line shows 21 April 2016 (eligibility criteria made public).

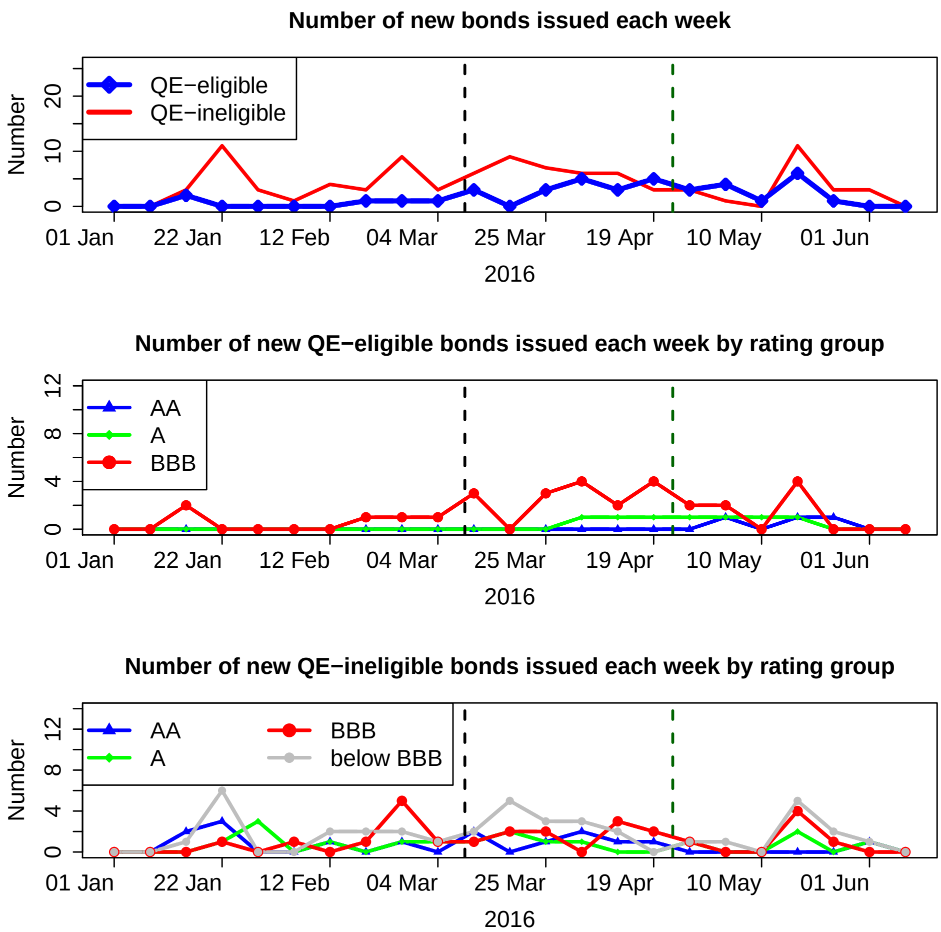

FIGURE 2.7: Par of new bonds issued by firms issuing in several currencies: January 2016–June 2016. The top panel shows the total par of newly issued eligible and ineligible bonds. The mid and bottom panels present the results split by rating categories. The first dashed line indicates the day after 10 March 2016 (CSPP announced), the second dashed line shows 21 April 2016 (eligibility criteria made public). Par is the notional amount of bonds outstanding, measured in billion EUR.

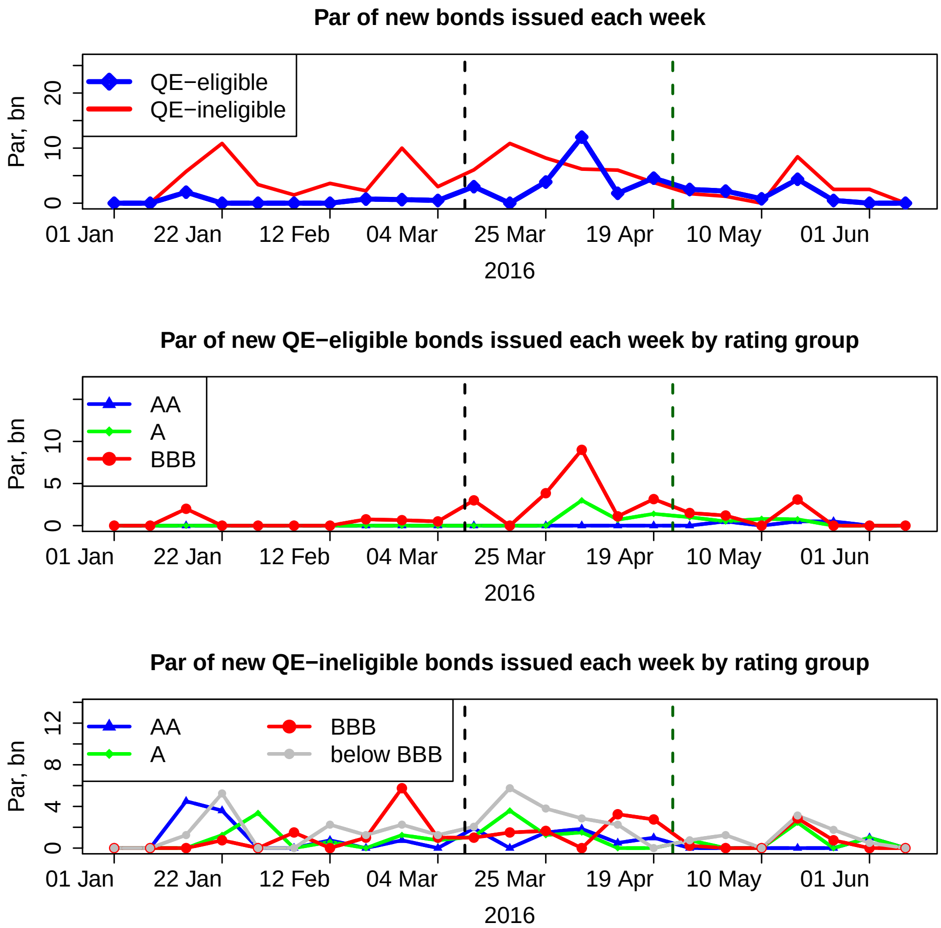

FIGURE 2.8: New bonds issued by firms issuing in several currencies: splitting the sample of euro-denominated bonds. Sample period: January 2016–June 2016. The first dashed line indicates the day after 10 March 2016 (CSPP announced), the second dashed line shows 21 April 2016 (eligibility criteria made public). Par is the notional amount of bonds outstanding, measured in billion EUR.

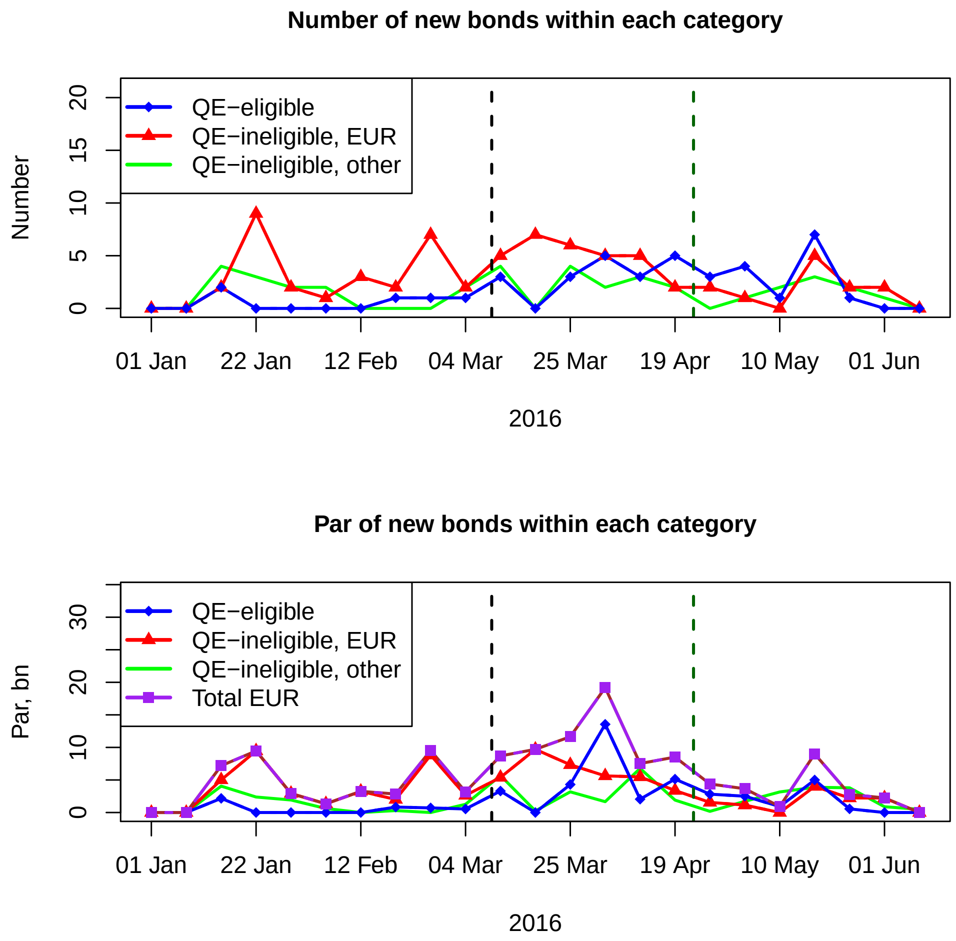

2.8 Tables

TABLE 2.1: Summary statistics. The table shows summary statistics for eligible (QE) and ineligible (non-QE) bonds just before the CSPP announcement on 10 March 2016. The total number of eligible bonds is 546, the total number of ineligible ones is 6,044.

	Age, years		Maturity, years		Coupon, %		Par, million EUR	
	QE	non-QE	QE	non-QE	QE	non-QE	QE	non-QE
Mean	3.19	4.31	5.75	6.97	2.54	4.17	834.44	657.25
Standard deviation	2.67	4.17	3.00	20.64	1.48	2.37	433.20	544.13
Min	0.00	0.00	0.88	0.05	0.00	0.00	98.56	0.03
Deciles								
20%	1.12	1.41	3.04	1.87	1.25	1.90	528.76	262.55
50%	2.60	2.99	5.42	4.21	2.12	4.12	728.98	549.11
70%	3.64	4.90	7.10	6.52	3.12	5.38	899.99	799.66
100%	13.79	31.38	17.72	999.62	8.12	14.00	3,068.05	10,255.61

TABLE 2.2: Impact of QE on bond prices and liquidity. Column 1 estimates regression (2.2) for bond yields (Y). Column 2 estimates the same regression with several controls (bond coupon, rating dummies, age, par amount). Junk is a dummy for all ratings below B. Column 3 estimates regression (2.1). Column 4 estimates regression (2.2) for a sample of bonds matched on par outstanding, rating, and time to maturity. Column 5 estimates regression (2.1) with a control group consisting of all bonds that satisfy the CSPP eligibility criteria. Columns 6–7 estimate regression (2.1) for the bilateral turnover (BT) and the tri-party repo turnover (RT), respectively. Column 8 estimates the same specification for the bilateral market turnover ratio – BTR (here and in all subsequent regression tables measured in %) and Column 9 for the tri-party repo market turnover ratio (RTR). Column 10 estimates the same regression specification for bid-ask spreads of yields (BA). The coefficients $Treat \times Inter$ and $Treat \times Post$ are the DD estimates of the effect of the CSPP announcement. Here and in all subsequent tables yields and bid-ask spreads are observed daily, whereas turnover and turnover ratios are weekly. Standard errors (shown in parentheses) are heteroskedasticity robust and double-clustered by firm and time (1,550 · 23 clusters). *,**, and *** indicate statistical significance at the 10%, 5%, and 1% levels. Units for Y, BTR, RTR, and BA: %. Units for BT, RT: million USD.

Dependent variables	Y	Y	Y	Y	Y	BT	RT	BTR	RTR	BA
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Coupon		0.90***								
		(0.14)								
AA		-2.14***								
		(0.39)								
A		-2.02***								
		(0.45)								
BBB		-2.07***								
		(0.53)								
BB		(-0.89)								
		(0.71)								
B		2.08***								
		(0.95)								
Junk		19.63***								
		(2.85)								
Age		-0.16***								
		(0.04)								
Par		-0.01								
		(0.01)								
Treat	-2.86***	-0.61***		-5.31***	0.20**					
	(0.12)	(0.10)		(0.78)	(0.10)					
Inter	-0.24*	-0.21*		-0.04	-0.33***					
	(0.14)	(0.16)		(0.34)	(0.02)					
Post	-0.37*	-0.31*		-1.18	-0.41***					
	(0.21)	(0.17)		(1.82)	(0.03)					
$Treat \times Inter$	**-0.22*****	**-0.22*****	**-0.22*****	**-0.21*****	**-0.13*****	**2.40**	**1.52**	**0.46****	**0.39****	**-0.02****
	(0.04)	(0.04)	(0.04)	(0.03)	(0.02)	(1.56)	(1.14)	(0.21)	(0.18)	(0.01)
$Treat \times Post$	**-0.29*****	**-0.29*****	**-0.30*****	**-0.28*****	**-0.23*****	**7.05*****	**8.15*****	**1.17*****	**1.51*****	**-0.06*****
	(0.07)	(0.07)	(0.07)	(0.04)	(0.04)	(1.75)	(2.25)	(0.28)	(0.27)	(0.01)
Bond fixed effects and time dummies	No	No	Yes	No	Yes	Yes	Yes	Yes	Yes	Yes
Observations (Bond-Time)	611,356	611,356	611,356	611,356	101,672	115,174	107,875	115,048	107,758	611,356
F-statistics	1111.33	15,459.10	16,523.86	360.87	368.15	125.86	76.43	81.04	52.47	38.67

TABLE 2.3: Liquidity dissipation. The table presents the estimates of regression (2.1) with the post-announcement period extended by four months – until 08 October 2016. Columns 1–4 present the results for the turnover measures, column 5 for yields, and column 6 for bid-ask spreads. Standard errors (shown in parentheses) are heteroskedasticity robust and double-clustered by firm and time. *,**, and *** indicate statistical significance at the 10%, 5%, and 1% levels. See Table 2.2 for variable and unit definitions.

Dependent variables	BT	RT	BTR	RTR	Y	BA
	(1)	(2)	(3)	(4)	(5)	(6)
$Treat \times Inter$	-1.65	-1.63	0.46**	0.38**	-0.22***	-0.02**
	(1.84)	(1.04)	(0.19)	(0.16)	(0.04)	(0.01)
$Treat \times Post$	3.29	4.11**	0.50***	1.30***	-0.26***	-0.07***
	(2.34)	(1.88)	(0.15)	(0.20)	(0.09)	(0.02)
Bond fixed effects and time dummies	Yes	Yes	Yes	Yes	Yes	Yes
Observations (Bond-Time)	195,226	183,295	195,088	183,137	1,091,051	1,091,051
F-statistics	20.36	81.41	88.70	127.21	1892.16	331.62

Table 2.4: Isolating QE channels. The table presents the coefficients for the $Treat \times Post2$ interaction dummy in regression (2.3) estimated for a particular maturity-rating bucket of bonds. Panel A presents the results for yields, Panel B for bid-ask spreads. The last row in each panel estimates DDD regression (2.4) only for the sample of treated bonds. Each regression has firm and time fixed effects. Standard errors (shown in parentheses) are heteroskedasticity robust and double-clustered by firm and time. *,**, and *** indicate statistical significance at the 10%, 5%, and 1% levels. Units for all estimates: %.

Panel A: Yields				
	0–2 years	2–5 years	5–10 years	>10 years
AA	-0.052	0.007	-0.126***	-0.381***
	(0.09)	(0.008)	(0.043)	(0.034)
A	0.022	-0.144***	-0.292***	-0.476***
	(0.038)	(0.030)	(0.035)	(0.075)
BBB	-0.108	-0.148*	-0.532***	-0.824***
	(0.082)	(0.076)	(0.092)	(0.140)
BBB, QE	-0.148*	-0.203***	-0.354***	-0.368**
	(0.075)	(0.064)	(0.088)	(0.146)

Panel B: Bid-ask spreads				
	0–2 years	2–5 years	5–10 years	>10 years
AA	-0.288	-0.023***	-0.011***	-0.024***
	(0.203)	(0.007)	(0.003)	(0.004)
A	0.085	-0.015**	-0.019***	-0.021***
	(0.076)	(0.007)	(0.003)	(0.004)
BBB	-0.241	-0.025***	-0.025***	-0.024***
	(0.160)	(0.005)	(0.005)	(0.007)
BBB, QE	-0.264	0.003	-0.008***	-0.006
	(0.167)	(0.011)	(0.003)	(0.005)

Table 2.5: Impact of QE on the issue of new bonds. Panel A estimates regression (2.3) for the number of bonds issued each week, Panel B – the same regression for the total par of bonds issued each week. For both panels, column 1 presents the results for the sample of all firms, column 2 for the sample of firms issuing in more than two currencies (254 firms). Column 3 shows the estimates for bonds issued in euros. The reference category for this column is bonds issued by the 254 firms in currencies other than the euro. Columns 4–6 present the results for the sample of 254 firms split into different rating groups, columns 7–9 for the same sample split according to maturity of new bonds issued. Columns 11–13 present the results for regression (2.5) where the QE-eligible sample of firms is split into different rating groups. The reference category for this regression is firms rated BB. Standard errors (shown in parentheses) are heteroskedasticity robust and double-clustered by firm and time. *,**, and *** indicate statistical significance at the 10%, 5%, and 1% levels. Units for the estimates in Panel A: number of bonds. Units for the estimates in Panel B: billion EUR.

Panel A: Number of bonds													
Dependent variables	All	All	EUR	AA	A	BBB	0–2y	2–5y	5–10y	>10y	AA, QE	A, QE	BBB, QE
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
$Treat \times Post2$	1.18**	1.39*	2.87*	0.31	0.55*	0.59**	-0.03	0.48	0.58**	0.56**	0.23**	0.54***	0.92***
	(0.60)	(0.79)	(1.67)	(0.42)	(0.30)	(0.28)	(0.22)	(0.57)	(0.28)	(0.26)	(0.12)	(0.14)	(0.27)
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Week dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No
Treat × Week dummies	No	No	No	No	No	No	No	No	No	No	Yes	Yes	Yes
Observations (Bond-Week)	8,970	3,082	2,185	714	1,029	1,287	650	776	920	736	897	897	897
F-statistics	11.48	5.34	2.58	0.51	4.59	4.71	0.02	0.59	4.56	3.51	10.65	10.65	10.65

Panel B: Par of bonds													
Dependent variables	All	All	EUR	AA	A	BBB	0–2y	2–5y	5–10y	>10y	AA, QE	A, QE	BBB, QE
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
$Treat \times Post2$	2.19**	1.94*	2.73*	0.45	0.48*	0.77**	0.03	0.62	1.00*	0.53**	0.12**	0.63***	0.84***
	(1.10)	(1.11)	(1.59)	(0.50)	(0.27)	(0.37)	(0.26)	(0.63)	(0.58)	(0.25)	(0.06)	(0.23)	(0.35)
Firm fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Week dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No	No	No
Treat × Week dummies	No	No	No	No	No	No	No	No	No	No	Yes	Yes	Yes
Observations (Bond-Week)	8,970	3,082	2,185	714	1,029	1,287	650	776	920	736	897	897	897
F-statistics	11.22	5.59	1.80	0.70	4.72	5.01	0.01	0.80	3.75	5.10	7.31	7.31	7.31

TABLE 2.6: Firms' financial indicators. The table presents the results of regression (2.1) using firms' financial indicators as dependent variables. Column 1 uses dividends, column 2 – dividends for firms matched on total assets. Columns 3–8 use tangible fixed assets, long-term debt, R&D expenses, property, plant, and equipment, working capital, and cash holdings, respectively. The analyzed period is first-third quarter of 2016. Standard errors (shown in parentheses) are heteroskedasticity robust and double-clustered by firm and time. *,**, and *** indicate statistical significance at the 10%, 5%, and 1% levels. Units for the estimates: 100,000.

Dependent variables	D	D-matched	TFA	LTD	RD	PPE	WC	Cash
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
$Treat \times Post$	5.75**	20.64***	-15.17	-2.81	-1.92	-15.32	-2.57	-0.72
	(2.54)	(7.21)	(42.07)	(14.35)	(1.85)	(30.94)	(9.50)	(6.78)
Firm and quarter fixed effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations (Firm-Quarter)	264	264	364	360	370	364	279	369
F-statistics	12.46	35.31	9.48	5.12	8.26	9.48	10.21	5.46

Table 2.7: Placebo tests. In Panel A, I shift the post-announcement period backwards by one year. In Panel B, I use USD instead of EUR bonds. Columns 1 and 2 estimate regression (2.1) for yields and bid-ask spreads, respectively. Columns 3 and 4 estimate the same regression for the bilateral and tri-party repo turnover ratio, respectively. Columns 5 and 6 present the estimates of the same regression for the number (NB) and the par of newly issued bonds. Standard errors (shown in parentheses) are heteroskedasticity robust and double-clustered by firm and time. *,**, and *** indicate statistical significance at the 10%, 5%, and 1% levels. Units for all estimates in columns 1–4: %. Units for the estimates in column 6: million EUR.

Panel A: Placebo test I						
Dependent variables	Y	BA	BTR	RTR	NB	Par
	(1)	(2)	(3)	(4)	(5)	(6)
$Treat \times Post$	-1.33	-2.62	-0.28	0.96	0.00	-0.29
	(1.37)	(2.89)	(0.19)	(1.17)	(1.29)	(0.71)
Bond fixed effects and week dummies	Yes	Yes	Yes	Yes	No	No
Firm fixed effects and week dummies	No	No	No	No	Yes	Yes
Observations (Bond-Week)	509,161	509,161	86,204	80,339	4,784	4,784
F-statistics	5.33	5.42	68.32	37.59	2.01	2.35
Panel B: Placebo test II						
Dependent variables	Y	BA	BTR	RTR	NB	Par
	(1)	(2)	(3)	(4)	(5)	(6)
$Treat \times Inter$	-0.02	-0.03	-0.03	-0.01		
	(0.04)	(0.02)	(0.10)	(0.22)		
$Treat \times Post$	-0.09	-0.06	0.10	0.05		
	(0.07)	(0.04)	(0.16)	(0.43)		
$Treat \times Post2$					0.76	0.39
					(1.41)	(1.52)
Bond fixed effects and week dummies	Yes	Yes	Yes	Yes	No	No
Firm fixed effects and week dummies	No	No	No	No	Yes	Yes
Observations (Bond-Week)	611,356	611,356	115,068	107,781	3,536	3,536
F-statistics	89.79	113.03	107.20	63.16	1.27	3.07

2.9 Appendix

FIGURE 2.9: Extended post-announcement period. QE-eligible and QE-ineligible bonds. The top panel shows the daily time series of mean bond yields. The middle and the bottom panels illustrate the weekly dynamics of mean turnover. The first vertical dashed line indicates the day after 10 March 2016 (CSPP announced), the second dashed line shows 21 April 2016 (eligibility criteria made public). Units for turnover: USD.

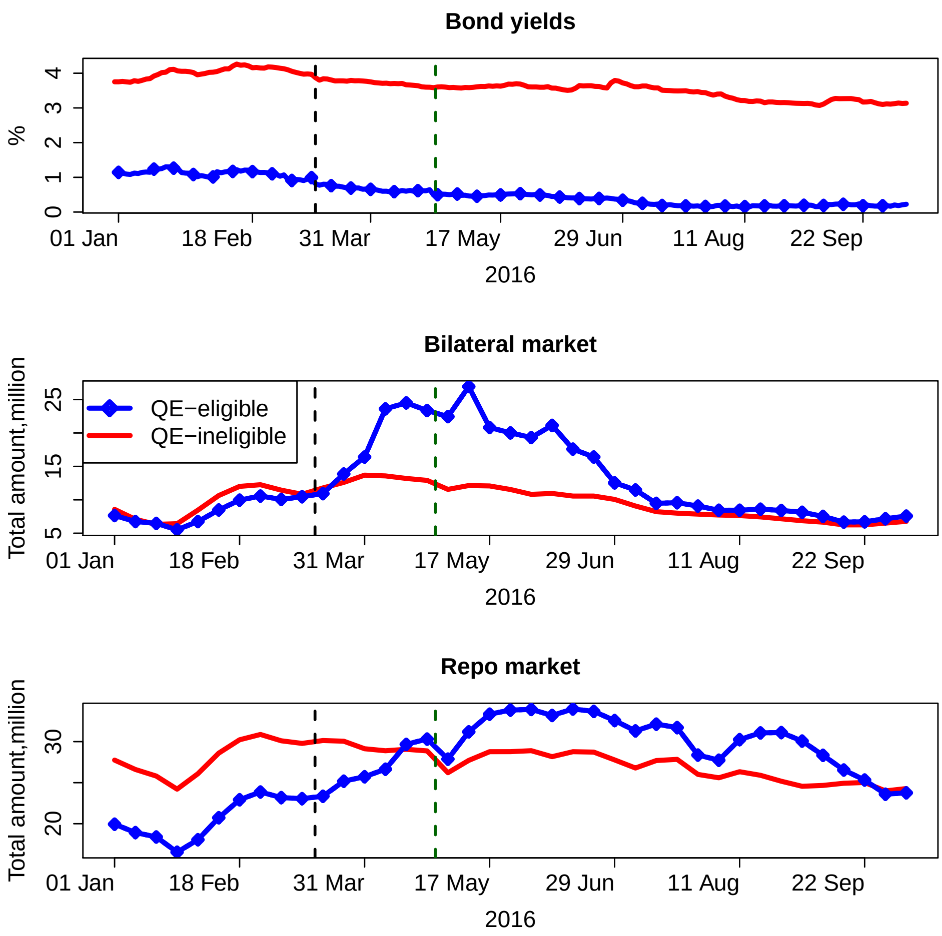

FIGURE 2.10: New bonds issued by firms issuing in several currencies: splitting the euro sample. Quarterly aggregation. The first dashed line indicates the day after 10 March 2016 (CSPP announced), the second dashed line shows 21 April 2016 (eligibility criteria made public). Par is the notional amount of bonds outstanding, measured in billion EUR.

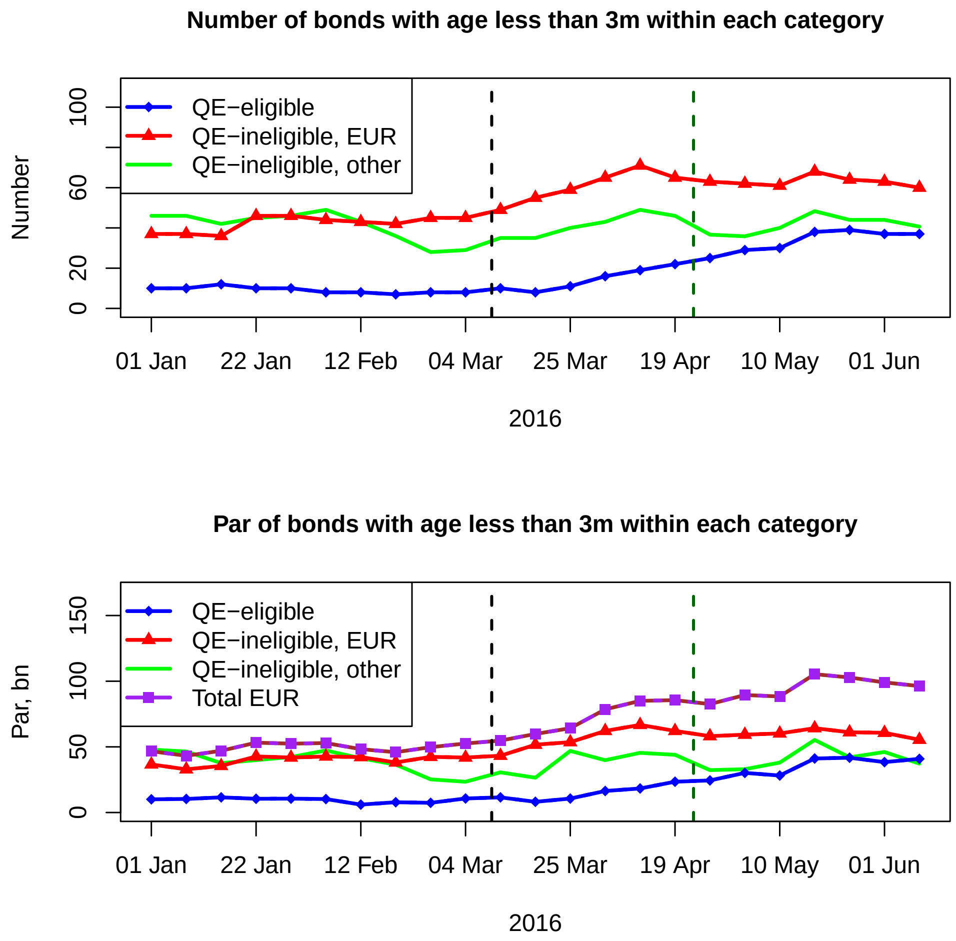

Table 2.8: Trading volume. The table presents summary statistics of the monthly turnover for all bonds from June, 2015 to June, 2016. Panel A shows the numbers for the bilateral market, Panel B – for the tri-party repo market. Panel C compares QE-eligible and QE-ineligible bonds before and after the CSPP announcement. Units: million USD.

Panel A: Bilateral market

	06.2015	08.2015	10.2015	12.2015	02.2016	03.2016	04.2016	05.2016	06.2016
Mean	65	66	71	63	60	61	89	83	68
Standard deviation	197	244	247	205	252	167	382	293	226
Observations	4,327	4,434	4,529	4,744	4,754	4,820	4,991	5,062	5,031
Deciles									
50%	25	23	23	22	22	24	25	27	23
70%	50	47	50	46	47	49	54	54	47
100%	5,130	6,901	5,097	5,021	11,442	5,068	8,494	7,175	4,996

Panel B: Tri-party repo market

	06.2015	08.2015	10.2015	12.2015	02.2016	03.2016	04.2016	05.2016	06.2016
Mean	199	208	195	166	166	168	183	180	175
Standard deviation	375	398	353	310	291	308	346	343	326
Observations	4,093	4,205	4,354	4,474	4,588	4,596	4,692	4,693	4,696
Deciles									
50%	78	87	84	66	63	68	70	66	69
70%	178	189	179	143	147	151	161	153	152
100%	5,944	10,675	7,532	4,140	5,287	8,194	7,052	6,231	4,687

Panel C: Bond turnover before and after QE

	Bilateral market				Tri-party (repo) market			
	Mean	20%	50%	80%	Mean	20%	50%	80%
Turnover (million USD), QE-eligible								
Before QE	8.51	1.64	4.07	9.74	20.74	4.15	10.91	28.04
After QE	19.84	1.98	5.28	12.58	30.13	5.03	14.23	39.59
Turnover (million USD), QE-ineligible								
Before QE	9.69	0.78	3.35	11.22	28.14	2.18	10.40	40.27
After QE	12.03	0.90	3.91	12.80	28.57	2.02	10.30	39.58
Turnover ratio (%), QE-eligible								
Before QE	0.91	0.23	0.51	1.11	2.52	0.57	1.31	3.44
After QE	2.13	0.30	0.62	1.40	3.43	0.60	1.76	4.71
Turnover ratio (%), QE-ineligible								
Before QE	1.24	0.21	0.53	1.43	4.61	0.32	1.53	5.81
After QE	1.44	0.21	0.60	1.61	4.47	0.33	1.53	5.68

TABLE 2.9: Alternative measures of liquidity. Impact on Price. The table presents the results of regression (2.1) using price measures and the Amihud illiquidity measure. Columns 1-5 show the estimates for the daily bid-ask spread, effective spread, liquidity cost score, mid and ask prices as dependent variables. Column 6 presents the estimates for the Amihud illiquidity indicator in the bilateral market, column 7 – in the tri-party repo market. The Amihud illiquidity measure for a period p is $A_{ip} = \frac{1}{D_{ip}} \sum_{t=1}^{D_{ip}} \frac{|r_{it}|}{Dvol_{it}}$, where $Dvol_{it}$ is the daily trading volume in asset i, r_{it} is the daily return of the asset, and D_{ip} is the number of days in the period over which the averaging is done. Standard errors (shown in parentheses) are heteroskedasticity robust and double-clustered by firm and time. *,**, and *** indicate statistical significance at the 10%, 5%, and 1% levels. Units for the estimates in columns 1-3: %. Units for the estimates in columns 6-7: 10^{-9}.

Dependent variables	BA, prices	ES	LCS	Mid	Ask	A, bilateral	A, tri-party
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$Treat \times Inter$	-0.35***	-0.49***	-0.55***	4.22***	4.04***	-1.62	-24.97
	(0.04)	(0.04)	(0.10)	(0.48)	(0.46)	(1.47)	(23.56)
$Treat \times Post$	-0.32***	-0.49***	-0.59***	3.29***	3.13***	-2.87**	-22.92*
	(0.03)	(0.04)	(0.12)	(0.52)	(0.54)	(1.46)	(12.76)
Bond Fixed Effects and Time Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations (Bond-Week)	539,951	539,951	539,951	539,951	539,951	238,442	238,442
F-statistics	60.07	13.84	4.22	25.72	24.29	4.63	4.13

Table 2.10: Impact of QE on turnover. Column 1 estimates regression (2.2) for the bilateral turnover (BT). Column 2 estimates the same regression with several controls (bond coupon, rating dummies, age, par amount). Column 3 estimates regression (2.2) for a sample of bonds matched on par outstanding, rating and time to maturity. Column 4 estimates regression (2.1) with a control group consisting of all bonds that satisfy the CSPP eligibility criteria. Columns 5-6 present the results for the same regressions as in columns 3-4, respectively, but for the bilateral turnover ratio (BTR). Columns 7-12 estimate the same regression specifications as in columns 1-6 for the tri-party repo market. Columns 13-14 estimate the same regression specifications as in columns 3-4, respectively, but for bid-ask spreads. Standard errors (shown in parentheses) are heteroskedasticity robust and double-clustered by firm and time ($1{,}550 \cdot 23$ clusters). *,**, and *** indicate statistical significance at the 10%, 5%, and 1% levels. Units for BTR, RTR and BA: %. Units for BT, RT: million USD.

Dependent variables	BT	BT	BT,CG	BT,ACG	BTR, CG	BTR, ACG	RT	RT	RT,CG	RT,ACG	RTR, CG	RTR, ACG	BA, CG	BA, ACG
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)	(14)
Coupon		-1.07***						0.11						
		(0.23)						(0.35)						
AA		0.01						-1.64						
		(1.05)						(1.49)						
A		-0.86						-0.51						
		(0.64)						(0.91)						
BBB		0.78						1.42						
		(0.73)						(0.95)						
BB		-0.24						0.63						
		(0.62)						(1.06)						
B		1.16						0.34						
		(1.54)						(1.16)						
junk		-1.91**						-0.27						
		(0.74)						(1.56)						
Age		-1.11***						-0.63***						
		(0.13)						(0.17)						
Par		0.01						0.01						
		(0.01)						(0.01)						
Treat	-1.47*	-3.84***					-9.95***	-10.21***						
	(0.85)	(0.90)					(1.24)	(1.23)						
Inter	3.41***	3.36***					1.26	1.32*						
	(0.90)	(0.91)					(0.80)	(0.80)						
Post	1.61*	1.55***					0.10	0.17						
	(0.84)	(0.83)					(0.87)	(0.88)						
Treat × *Inter*	**3.31***	**3.30***	**2.20**	**2.04**	**0.39**	**0.24**	**2.75****	**2.75***	**1.21**	**1.23**	**0.37**	**0.06**	**-0.02**	**-0.02***
	(1.74)	(1.75)	(2.87)	(1.84)	(0.29)	(0.19)	(1.45)	(1.43)	(1.45)	(1.58)	(0.26)	(0.17)	(0.02)	(0.01)
Treat × *Post*	**7.66*****	**7.40*****	**6.70*****	**6.42*****	**1.03*****	**0.85*****	**8.45*****	**8.37*****	**7.45*****	**7.75*****	**1.37*****	**1.16*****	**-0.05*****	**-0.04*****
	(2.28)	(2.25)	(2.08)	(3.17)	(0.38)	(0.48)	(2.05)	(2.04)	(1.92)	(2.61)	(0.32)	(0.31)	(0.01)	(0.01)
Bond Fixed Effects and Time Dummies	No	No	No	Yes	No	Yes	No	No	No	Yes	No	Yes	No	Yes
Observations (Bond-Time)	115,174	115,174	115,174	38,216	115,068	38,193	107,875	107,875	107,781	36,103	107,875	36,092	611,186	101,672
F-statistics	65.09	92.23	87.57	47.98	88.97	63.38	41.63	49.82	33.59	41.29	64.52	66.60	88.19	399.63

Table 2.11: Data collapsed by time period. The table presents the results of the main regressions (except those from Table 2.12) if the data are collapsed by time (week or day). All estimates are from pure time-series regressions. Column 1 presents the results for yields, column 2 – for bid-ask spreads, columns 3-6 – for the turnover measures, columns 7-13 – for the number of newly issued bonds by a firm's rating category and by bond maturity. Columns 14-21 show the estimates for firms' financial indicators. Standard errors (shown in parentheses) are heteroskedasticity and autocorrelation robust. *,**, and *** indicate statistical significance at the 10%, 5%, and 1% levels.

Dependent variables	Y	BA	BT	BTR	RT	RTR	AA	A	BBB	0-2y
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Treat × *Inter*	-0.20***	-0.01	2.62	0.46***	1.72	0.39***				
	(0.02)	(0.03)	(2.12)	(0.11)	(1.33)	(0.13)				
Treat × *Post*	-0.30***	-0.06*	7.70***	1.19***	8.40***	1.45***				
	(0.02)	(0.03)	(0.98)	(0.10)	(0.55)	(0.11)				
Treat × *Post*2							0.32	0.55**	0.57**	0.01
							(0.50)	(0.28)	(0.28)	(0.20)
Observations (Day/Week/Quarter)	114	114	23	23	23	23	23	23	23	23
F-statistics	101.1	2.31	22.74	64.78	91.27	76.06	2.04	5.41	5.80	0.06

	2-5y	5-10y	>10y	D	D-matched	TFA	LTD	RD	PPE	WC	Cash
	(11)	(12)	(13)	(14)	(15)	(16)	(17)	(18)	(19)	(20)	(21)
Treat × *Post*				5.97**	22.47**	-14.85	-6.13	-2.43	-14.28	-4.91	0.60
				(1.05)	(8.40)	(26.43)	(5.78)	(2.12)	(9.96)	(3.75)	(6.41)
Treat × *Post*2	0.44	0.57**	0.55**								
	(0.47)	(0.28)	(0.27)								
Observations (Day/Week/Quarter)	23	23	23	3	3	3	3	3	3	3	3
F-statistics	0.75	3.95	3.20	542.20	685.21	14.12	13.76	5.75	15.15	0.53	0.00

TABLE 2.12: Isolating QE channels: data collapsed by time period. The table presents the coefficients for the $Treat \times Post2$ interaction dummy in a time series regression estimated for a particular maturity-rating bucket of bonds. Panel A presents the results for yields, Panel B – for bid-ask spreads. The last row in each panel estimates the time-series regression only for the sample of treated bonds. Standard errors (shown in parentheses) are heteroskedasticity and autocorrelation robust. *,**, and *** indicate statistical significance at the 10%, 5%, and 1% levels. Units for all estimates: %.

Panel A: Yields				
	0–2 years	2–5 years	5–10 years	>10 years
AA	-0.328	-0.184	-0.142***	-0.305***
	(0.281)	(0.144)	(0.050)	(0.022)
A	-0.162	-0.191***	-0.226***	-0.298***
	(0.150)	(0.054)	(0.029)	(0.061)
BBB	-0.139	-0.284***	-0.473***	-0.822***
	(0.082)	(0.116)	(0.123)	(0.133)
BBB, QE	-0.097***	-0.178***	-0.316***	-0.356***
	(0.026)	(0.061)	(0.122)	(0.120)

Panel B: Bid-ask spreads				
	0–2 years	2–5 years	5–10 years	>10 years
AA	-0.295	0.098	-0.011***	-0.024***
	(0.230)	(0.122)	(0.005)	(0.002)
A	-0.015	-0.023***	-0.023***	-0.020***
	(0.015)	(0.002)	(0.001)	(0.002)
BBB	-0.031	-0.023***	-0.026***	-0.023***
	(0.022)	(0.002)	(0.002)	(0.002)
BBB, QE	-0.173	0.032	-0.008***	0.017
	(0.150)	(0.032)	(0.002)	(0.014)

TABLE 2.13: Isolating QE channels: relative changes. The table presents the relative change in yields or bid-ask spreads for a particular maturity-rating bucket from Table 2.4. The numbers are the coefficients from Table 2.4 divided by the control group before-QE mean of bonds within a particular bucket. Panel A presents the quantities for yields, Panel B – for bid-ask spreads. *,**, and *** indicate statistical significance at the 10%, 5%, and 1% levels. Units for all estimates: %.

Panel A: Yields				
	0–2 years	2–5 years	5–10 years	>10 years
AA	-9.40	0.71	-10.85***	-15.25***
A	2.05	-8.62***	-13.97***	-13.83***
BBB	-8.42	-8.36*	-17.89***	-19.79***
BBB, QE	-19.31*	-27.65***	-30.51***	-30.68**
Panel B: Bid-ask spreads				
	0–2 years	2–5 years	5–10 years	>10 years
AA	-226.77	-22.43***	-12.94***	-27.27***
A	22.73	-13.16**	-16.96***	-20.59***
BBB	-87.96	-14.97***	-21.36***	-20.17***
BBB, QE	-98.31	26.68	-9.09***	-8.05

TABLE 2.14: Isolating QE channels: number of observations and F-statistics. The table presents the number of observations and F-statistics (in brackets) for each regression for a particular maturity-rating bucket from Table 2.4. Panel A presents the quantities for yields, Panel B – for bid-ask spreads.

Panel A: Yields				
	0–2 years	2–5 years	5–10 years	>10 years
AA	9,260	1,5015	10,349	6,203
	(18.97)	(243.75)	(272.46)	(154.73)
A	33,699	56,294	46,352	25,442
	(88.12)	(1,706.81)	(1,795.82)	(965.57)
BBB	34,240	53,925	54,059	31,934
	(393.01)	(506.87)	(2,553.16)	(695.40)
BBB, QE	3,771	17,868	21,211	3,617
	(156.326)	(2,269.38)	(2,334.11)	(1,079.98)

Panel B: Bid-ask spreads				
	0–2 years	2–5 years	5–10 years	>10 years
AA	9,260	15,015	10,349	6,203
	(6.43)	(22.37)	(94.18)	(59.62)
A	33,699	56,294	46,352	25,442
	(1.54)	(39.02)	(215.84)	(96.17)
BBB	34,240	53,925	54,059	31,934
	(166.08)	(276.16)	(64.52)	(6.56)
BBB, QE	3,771	17,868	21,211	3,617
	(15.72)	(245.26)	(945.99)	(521.24)

TABLE 2.15: Isolating QE channels: splitting bonds on liquid and illiquid before QE. The table presents the coefficients for the $Treat \times Post2$ interaction dummy in regression (2.3) estimated for a particular maturity-rating bucket for the sample of liquid bonds (below median bid-ask spread before QE) and illiquid bonds (above median bid-ask spread before QE). Panels A and B present the results for liquid bonds, Panels C and D – for illiquid bonds. Panels A and C present the results for yields, Panels B and D – for bid-ask spreads. Empty cells indicate that there are no bonds within the particular maturity-rating bucket. Each regression has firm and time fixed effects. Standard errors (shown in parentheses) are heteroskedasticity robust and double-clustered by firm and time. *,**, and *** indicate statistical significance at the 10%, 5%, and 1% levels. Units for all estimates: %.

Panel A: Yields. Lliquid

	0–2 years	2–5 years	5–10 years	>10 years
AA		0.110	-0.209***	-0.376***
		(0.041)	(0.063)	(0.034)
A	-0.054		-0.516***	-0.479***
	(0.081)		(0.115)	(0.063)
BBB	-0.262	0.414	-0.785***	-0.792***
	(0.274)	(0.855)	(0.179)	(0.130)

Panel B: Bid-ask spreads. Lliquid

	0–2 years	2–5 years	5–10 years	>10 years
AA		-0.009**	0.024***	-0.024***
		(0.004)	(0.005)	(0.004)
A	0.324**		-0.018***	-0.016***
	(0.143)		(0.005)	(0.002)
BBB	-0.149	0.035	-0.025***	-0.022***
	(0.253)	(0.047)	(0.006)	(0.004)

Panel C: Yields. Illiquid

	0–2 years	2–5 years	5–10 years	>10 years
AA	0.012	0.059	-0.101***	
	(0.059)	(0.036)	(0.030)	
A	0.036	-0.123***	-0.203***	-0.609***
	(0.039)	(0.027)	(0.032)	(0.079)
BBB	-0.048	-0.126*	-0.258***	
	(0.094)	(0.070)	(0.045)	

Panel D: Bid-ask spreads. Illiquid

	0–2 years	2–5 years	5–10 years	>10 years
AA	-0.133	-0.017*	-0.013***	
	(0.144)	(0.010)	(0.002)	
A	-0.033	0.023	-0.019***	-0.011**
	(0.042)	(0.022)	(0.002)	(0.005)
BBB	-0.254	-0.022***	-0.045***	
	(0.192)	(0.006)	(0.016)	

TABLE 2.16: Isolating QE channels: t-tests. The table presents the p-values of t-tests on the hypothesis that the regression coefficients are equal for a set of bonds within a particular rating-maturity group. Each number in the table reports the p-value from such a test. Panel A presents the quantities for yields, Panel B – for bid-ask spreads. For example, the first row of Panel A presents the p-value from a test of whether the coefficient for BBB-rated bonds is the same as the coefficient on AA-rated bonds within each maturity group.

Panel A: Yields				
	0–2 years	2–5 years	5–10 years	>10 years
BBB vs AA	0.37	0.01	0.00	0.00
BBB vs A	0.03	0.94	0.00	0.00
A vs AA	0.59	0.00	0.00	0.16
BBB, QE, 0–2 years		0.05	0.00	0.00
BBB, QE, >10 years		0.00	0.51	
Panel B: Bid-ask spreads				
	0–2 years	2–5 years	5–10 years	>10 years
BBB vs AA	0.02	0.48	0.00	0.95
BBB vs A	0.00	0.00	0.01	0.85
A vs AA	0.00	0.27	0.00	0.55
BBB, QE, 0–2 years		0.00	0.00	0.00
BBB, QE, >10 years		0.00	0.03	

TABLE 2.17: Percentage of QE-eligilbe bonds issued by firms with different ratings. The table shows the fraction of firms in each rating category and the corresponding proportion of QE-eligible bonds issued by these firms. Units of the numbers: %.

Firm's rating group	Percentage of firms	Percentage of QE-bonds issued
AA	10.19	9.92
A	29.30	40.20
BBB	35.08	47.58
BB	22.54	2.29
junk	2.89	0.00

TABLE 2.18: Impact of QE on the issue of new bonds: whole sample of firms. Panel A estimates regression (2.3) for the number of bonds issued each week, Panel B – the same regression for the total par of bonds issued each week. For both panels, columns 1-3 present the results for the sample of all firms split into different rating groups, columns 4-7 – for the same sample split according to maturity of new bonds issued. Standard errors (shown in parentheses) are heteroskedasticity robust and double-clustered by firm and time. *,**, and *** indicate statistical significance at the 10%, 5%, and 1% levels. Units for the estimates in Panel A: number of bonds. Units for the estimates in Panel B: billion EUR.

Panel A: Number of bonds							
Dependent variables	AA	A	BBB	0-2y	2-5y	5-10y	>10y
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$Treat \times Post2$	0.10	0.59*	0.61*	0.37	0.37	0.28*	0.48**
	(0.39)	(0.34)	(0.37)	(0.31)	(0.78)	(0.16)	(0.24)
Firm Fixed Effects and Week Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations (Bond-Week)	1,840	2,651	3,055	1,860	2,302	2,698	2,110
F-statistics	0.07	10.94	10.49	1.51	0.18	10.13	11.81

Panel B: Par of bonds							
Dependent variables	AA	A	BBB	0-2y	2-5y	5-10y	>10y
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
$Treat \times Post2$	0.62	0.57*	0.91**	0.45	-0.77	3.49***	3.15***
	(0.78)	(0.33)	(0.45)	(0.56)	(1.13)	(1.27)	(1.10)
Firm Fixed Effects and Week Dummies	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations (Bond-Week)	1,840	2,651	3,055	1,860	2,302	2,698	2,110
F-statistics	0.71	11.76	11.16	0.71	0.40	15.95	16.06

Chapter 3

What Drives Repo Haircuts? Evidence from the UK Market

Christian Julliard, Zijun Liu, Seyed E. Seyedan, Karamfil Todorov, Kathy Yuan[1]

3.1 Introduction

The repurchase agreement (repo) market is a major tool for short-term funding of financial institutions. Although there are no definitive data about the size of this market, the International Capital Market Association suggests that the value of the global commercial market can be up to 15 trillion EUR.[2] During the recent financial crisis repo markets experienced various disruptions and potentially contributed to the severity of the crisis. For example, Copeland et al. (2010) show that, during the days prior to bankruptcy, the amount of collateral Lehman Brothers financed in tri-party repo fell drastically. Gorton and Metrick (2012) argue that the repo market experienced a run during the crisis, manifested in a rise of haircuts, which exacerbated the crisis.

Given the importance of the repo market and its contribution to the systemic risk of the financial system – especially in the wake of the recent 2008 crisis – there is ample interest from academics, policy makers and members of the public in better understanding and monitoring this market. However, due to the over-the-counter nature of repo transactions, repo contract terms are rarely disclosed. Adrian et al. (2013) provide an

[1]We are grateful to seminar participants at the London School of Economics, the European Meeting of the Econometric Society, and the RiskLab/BoF/ESRB Conference on Systemic Risk Analytics for their helpful feedback. The views and opinions expressed in this article are those of the authors and do not necessarily reflect the official policy or position of the Hong Kong Monetary Authority.

[2]ICMA (2013).

overview of the sources that provide information for the US repo market and conclude that, although some sources provide data on interest rates and notional values used in repo trades, very little is known about haircuts, collaterals and counterparties.

The systemic importance of the repo market and the shortage of micro-data prompted the UK regulator to require banks to disclose transaction-level data on their repo books. We were given the opportunity to work with this unique regulatory data set to analyze the structure of the UK repo market. We have access to all trade level repo data such as notional value, maturity, counterparty, collateral, and haircut, except for repo rates. To our knowledge, this is the only database that covers transaction-level haircut information for a rich set of different collaterals and counterparties. Given the importance of haircuts and the fact that they control the amount of inside liquidity generated by the shadow banking system, we aim to answer the question of what factors drive their magnitude using transaction-level data. Furthermore, we examine the structure and attributes of the repo market network and assess their influence on haircuts.

A priori it might appear puzzling why repo loans feature both interest rate and haircut. Recent theoretical work such as Ozdenoren et al. (2018) shows that while both repo rates and haircuts are affected by the demand and supply for funding liquidity, riskiness of repo loans drives the former and the severity of adverse selection that lenders face influences the latter. Our survey of several trading houses in London has revealed that while repo rates are determined at the trading desk, haircuts are set by the credit department of the corresponding firm. The observation of the separate roles played by rates and haircuts motivates us to formulate testable hypotheses to study haircut determinations in details empirically.

In particular, we build testable hypotheses to study haircut determinations empirically based on the existing theoretical work on collateralized borrowing and repo runs. The theoretical work on collateralized borrowing can be categorized into two streams. One is based on the difference of opinion approach in a general equilibrium setting (see, e.g., Geanakoplos, 1997; Simsek, 2013). The other is based on contractual and/or information frictions (see, e.g., Dang et al., 2011, 2013; Gottardi et al., 2017; Ozdenoren et al., 2018). The repo run literature focuses on coordinations either extending Diamond and Dybvig (1983) to the repo setting (see, e.g., Martin et al., 2014), short-term borrowing (see, e.g., Acharya et al., 2011), or endogenous information acquisition (see, e.g., Gorton and Ordonez, 2014), or to adverse selection and inter-temporal coordination (see, e.g., Ozdenoren et al., 2018).

In our empirical investigation, we find that transaction maturity has a first order importance in setting haircuts. Haircuts are also increasing in the VaR of collaterals and collateral concentration. This set of findings indicates that collateral quality and liquidity are important determinants of haircuts. We also find that counterparties matter in haircut determinations: one or two banks in our sample receive a significant share of repo trades with zero haircuts, hedge funds are charged at higher haircut, larger borrowers with higher ratings receive lower haircuts. However, we do find that collateral quality can overshadow counterparty characteristics. Furthermore, there is evidence that borrowers with lower ratings use higher quality collateral to receive a lower haircut. Hence, the influence of counterparty attributes is concealed.

We also find that bilateral relationships matter in haircuts. Banks charge higher haircuts when they transact with non-bank institutions. This is supportive of the difference of opinion explanation of haircuts since it is likely that banks and non-bank financial institutions have different valuation models about collateral. However, it may also support the adverse selection explanation of haircuts since there are information frictions between different types of business. Furthermore, we find significant pairwise borrower-lender relationships: some borrowers receive consistently lower haircuts when interacting with certain counterparties, and a few bilateral pairs conduct a large portion of zero haircut trades in our sample. These findings are difficult for the difference of opinion theory to explain since these bilateral pairs are often from different lines of business. They are, however, supportive of the adverse selection theory since relationship banking lowers information frictions.

We find little evidence that lenders' liquidity position or default probabilities affect haircuts, suggesting that the traditional bank run mechanism cannot explain repo runs. This lends support to the inter-temporal feedback/coordination explanations of repo runs.

Finally, we examine the structure and attributes of the repo market network and assess if the network structure has an influence over haircuts. We observe that the banks with higher centrality measures ask for lower haircuts on reverse repos and pay lower haircuts on repos. We interpret this set of findings as supportive of the demand-and-supply theory for funding liquidity since the unique market position of central network players affects the terms of bilateral repo contracts.

The rest of this paper is organized as follows. Section 3.2 provides a brief description of repurchase agreements and summarizes the relevant literature. Section 3.3 outlines the main hypotheses that we test in the data. Section 3.4 describes the data. Section 3.5 analyzes the determinants of haircuts and presents the testable hypotheses. Section 3.6 concludes.

3.2 Background information on repurchase agreements and related literature

3.2.1 Background information on repurchase agreements

A repurchase agreement is the simultaneous sale of, and forward agreement to repurchase, securities at a specific price, at a future date (Duffie, 1996). In effect, a repo is a collateralized loan, where the underlying security serves the collateral role. The party that borrows cash and delivers collateral is said to be doing a repo, and the party that lends cash and receives collateral is doing a reverse repo. The difference between the original loan value and the repayment specifies the repo rate. The haircut or margin, on the other hand, is determined by the difference between the loan and the collateral value. Usually, the borrower has to post collateral in excess of the notional amount, and the haircut is defined as $h = 1 - F/C$ with collateral value C and notional amount F (Krishnamurthy et al., 2014). For example, if a borrower receives \$98 against \$100 value of collateral, the haircut is 2%.

In Europe, the legal title to the collateral is transferred to the cash lender by an outright sale. In the US this is not the case, but the repo collateral is not subject to an automatic stay and can be sold by the lender should the borrower default (ICMA, 2013).

Repurchase agreements are broadly classified into two categories. Tri-party repo is a transaction for which post-trade services such as collateral management (e.g. selection, valuation, and verifying eligibility criteria), payment, margining, etc. are outsourced to a third-party agent which is a custodian bank.[3] A tri-party agent settles the repos on its book, but in a bilateral repo, settlement usually occurs on a delivery versus payment basis, and the cash lender must have back-office capabilities to receive and manage the collateral (Adrian et al., 2013).

A growing number of repos are cleared via central (clearing) counterparties (CCPs). CCPs place themselves between the two sides of a trade, leading to a less complex web of exposures (Rehlon and Nixon, 2013). They provide benefits such as multilateral netting and facilities to manage member defaults in an orderly manner, but can also pose systemic risks to the financial system. CCPs always receive a haircut, whether in a reverse repo or repo. Banks doing a reverse repo with a CCP will need to give a haircut, which amounts to a negative value for the haircut.

[3] In Europe, the main tri-party agents are Clearstream, Euroclear, Bank of New York Mellon, and SegaInterSettle.

3.2.2 Related literature

The financial crisis rekindled interest in the theoretical and empirical study of the short-term funding market. The theoretical work on collateralized borrowing can be categorized into two streams. One is based on the difference of opinion approach in a general equilibrium setting such as in Geanakoplos (1997), Geanakoplos and Zame (2002), Geanakoplos (2003), Fostel and Geanakoplos (2012) and Simsek (2013). The other is based on contractual and/or information frictions such as in Dang et al. (2011), Dang et al. (2013), Gottardi et al. (2017) and Ozdenoren et al. (2018). We will discuss the theoretical literature in details when forming testable hypotheses in the next section of the paper. There is also a body of literature that models crisis and runs in the repo market. One approach is based on the classical setting in Diamond and Dybvig (1983) extending to the repo setup as in Martin et al. (2014). In this setup, the liquidity needs of the lender, the capital position of the borrower, and the market microstructure of the repo market play important roles in determining the magnitude of the run. Acharya et al. (2011) model freezes in the market for short-term financing in the form of a sudden collapse in debt capacity of collateral in an information-theoretic framework. Gorton and Ordonez (2014) focus on the information in-sensitivity of debt contracts and show how a sudden switch of information environment might trigger a deep discount and collateral crisis. Ozdenoren et al. (2018) emphasize the inter-temporal feedback of (expected) future asset price and the decisions of today's borrowers and lenders. Dynamic mis-coordination might lead to a run in the repo market.

The empirical studies of repurchase agreements have been mostly focused on the US repo market. Several papers have studied developments in this market during the financial crisis. Broadly speaking, two distinct phenomena can be identified in the US bilateral and tri-party repo markets. In the bilateral market, as argued by Gorton and Metrick (2012), a run occurred in the form of rapid increases in haircut levels. This is further supported by multiple hedge funds failing due to margin calls (Adrian et al., 2013). Adrian and Shin (2010) empirically show that repo transactions have contributed the most to the procyclical adjustments of the leverage of banks. From this perspective, the rapid increase of haircuts in bilateral repos during the crisis can also be viewed as (forced) deleveraging of broker-dealers (Adrian et al., 2013). In contrast, in the tri-party market haircuts moved very little and the amount of funding remained fairly stable but, instead, lenders refused to extend financing altogether to the most troubled institutions – namely, Bear Stearns and Lehman Brothers (Copeland et al., 2010). Krishnamurthy et al. (2014) argue that there was a run in the tri-party market but only for non-agency mortgage-backed securities (MBS)/asset-backed securities (ABS), which constituted a relatively small and insignificant part of the short-term debt market. In the tri-party market, tension seemed to affect specific institutions rather than the broad collateral classes, except maybe for the private-label securitized assets (Adrian et al., 2013). Martin et al. (2014) relate the differences between the behavior of these two markets with respect to their microstructure: in the tri-party market, haircuts are fixed in custodial agreements that are revised infrequently, but this is not the case in the bilateral market.

There is a limited number of empirical studies on repos. Most US studies on repos are on tri-party repos starting with Copeland et al. (2014), Krishnamurthy et al. (2014) and Hu et al. (2019). They generally find that the market is quite segmented and market power, collateral concentration and fund families might play important roles. To our knowledge, empirical studies on bilateral repos are rare. Therefore, the work by Gorton and Metrick (2012) using a proprietary database is important for the understanding of repo transaction, where various types of collaterals and counterparties are present. The repo studies in the European area are mostly conducted on general collateral repos or through CCPs, where regulations play a very important role (Mancini et al., 2016). To the best of our knowledge, the repo haircut database used in this paper is the only database that covers a significant part of a bilateral repo market.

3.3 Testable hypotheses on haircuts

Collateralized borrowing is an ancient financial institution. It serves an important economic function and has been used for a long time, and under very different institutions. For example, pawnshop loan records from China circa 662-689 A.D. show that silk garments were used as collateral (see, e.g., Goetzmann and Rouwenhorst, 2005). The popularity of collateral-backed lending is often attributed to its abilities to mitigate information frictions. In practice, producing information about borrowers or their actions can be very costly (due to credit registries, monitors, courts, etc.). Collateral allows the flow of credit while economizing on costly information acquisition with the haircut. However, according to the above pawn shop logic, the haircuts on collateral should be determined by the quality of collaterals only, not by the identity of the borrowers. Intuitively, the volatility or the illiquidity of the collateral asset matters in determining the amount of loan extended because in the event of default, the lender may not be able to recover the full market price (valued at the initial lending date) of the collateral. This leads to our first testable hypothesis.

Hypothesis 1 (collateral quality): The repo haircut is larger when the collateral is of lower quality and/or illiquid.

Collateral quality can be measured using VaR, maturity, rating, or asset types. Transaction maturity should matter since as the maturity of repo debt is longer, the loss from worsening collateral quality is greater. We use data from Bloomberg to calculate VaR based on the time series of prices before the date when the asset was used as a collateral in the repo/reverse repo contract. VaR (for 5-10 days) is used because most financial intermediaries need a certain holding period when finding a trading counterparty.

However, the pawnshop logic stops short in explaining the impact of counter-party quality and relationship banking on the magnitude of the haircuts in repo contracts. The empirical evidence has shown that the former matters. For example, Dang et al. (2011) show that repo by hedge fund borrowers have higher haircut than bank borrowers, on average. There are mainly two strands of the recent theory developments that study collateralized borrowing and hence, have implications for haircuts on repo contracts: those based on belief disagreement in a general equilibrium framework, and those based on contractual and/or information frictions. Geanakoplos (2003) is the first to propose a general equilibrium framework with difference in opinions to study leverage constraints and hence, haircuts on repos. The mechanism works as follows: optimists borrow from pessimists to speculate on the collateral. Since pessimists do not value the collateral as much as optimists do, they are reluctant to lend, which constrains optimists' ability to borrow and results in a haircut, which means that the face value of the loan is lower than the market value of the asset. Simsek (2013) emphasizes that only the belief disagreement about the probability of the downside states has a significant effect on haircut and asset prices. Since it is difficult to measure difference of opinion, we conjecture that when borrowers are from a different line of business from lenders, the potential belief disagreement is larger. This leads to our second testable hypothesis.

Hypothesis 2 (counterparty types): The repo haircut is larger when the counterparties in the contracts are from different lines of business.

The second strand of the literature uses the principal-agent models of borrowing constraints. As demonstrated in Simsek (2013), there is an equivalence of the principal-agent framework and the general equilibrium framework proposed by Geanakoplos (2010) as long as the optimistic borrowers have all the bargaining power. The principal-agent framework can be extended to include frictions other than belief disagreements such as costly state verification, moral hazard, or adverse selection (see, e.g., Dang et al., 2011; Ozdenoren et al., 2018). In these cases, the credit quality of the counterparty matters rather than the difference in types. This leads to our third testable hypothesis.

Hypothesis 3 (counterparty's quality): The repo haircut is larger when the default probability (credit quality) of borrower is higher (lower), or when the borrower is better privately informed about the quality of the collateral.

There is a strand of literature that models coordinations and runs, which have implications for repo haircuts. Gorton and Ordonez (2014) find that endogenous information acquisition can cause a sudden increase in haircut and a collateral crisis, hence, lenders' characteristics might matter. Similarly, in a dynamic sequential trade model, Dang et al. (2011) find that the haircut size is increasing in the liquidity needs of the lender, and in the default probability of the lender in a subsequent repo transaction. Similarly, in a series of dynamic Diamond and Dybvig (1983) models with an asset collateral market, Martin et al. (2014) find that collateral and liquidity constraints matter and hence, the liquidity of lenders matters in the haircut determination. This leads to our fourth testable hypothesis.

Hypothesis 4 (lender's quality and liquidity): The repo haircut is larger when the default probability and/or liquidity need of the lender is higher.

In contrast, Ozdenoren et al. (2018), in a dynamic adverse selection model, do not find that lenders' credit quality or liquidity constraints matter in haircuts. They find instead that the severity of adverse selection matters. This indicates that the bilateral relationship between borrower and lender should matter in haircuts since it lowers the information friction. This leads to our fifth testable hypothesis.

Hypothesis 5 (bilateral relationship): Haircuts are lower for bilateral parties with banking relationship.

Ozdenoren et al. (2018) also show that there are other ways to lower adverse selection. For example, a portfolio of collateral assets will have a larger borrowing capacity if it includes some safe asset. The idea is that the safe collateral convinces the lender to fund the borrower to invest in the risky collateral assets since the lender can recover the loan backed by the safe collateral. This initial investment, in turn, increases the prices of risky assets, and allows borrowers to borrow more against their risky collaterals, creating an unravelling effect and generating more liquidity. This leads to our last testable hypothesis.

Hypothesis 6 (portfolio repos): Risky assets in a portfolio repo with safe assets have lower haircut than purely risky asset repos.

We turn next to the description of the data, empirical strategy, and present hypotheses test results.

3.4 Overview of the data

The regulatory data set is a snapshot of the repo books of six banks that are major players in the UK repo market. The total size of their repo books – the sum of repos and reverse repos – is around 511 billion GBP (including CCP transactions) as at the end of 2012.[4] According to Financial Stability Board (2013), the UK-resident deposit-taker banks hold around 2.1 trillion GBP in gross repo activity on their balance sheets, hence our data set accounts for around 24% of the total repo activity in this market. The majority of this activity is with non-UK resident banks, including the activity between UK and foreign branches of the same consolidated group, and is highly concentrated (Financial Stability Board, 2013).

Each of the six banks reports its outstanding repo transactions as at the end of 2012, including the gross notional, maturity, currency, counter-party, haircuts and collaterals. We have supplemented this data set with additional data on securities, counter-parties, and the reporting banks from Datastream and Bloomberg. In what follows we report information and results for reverse repos (REVR) and repos (REPO) separately. This classification is from the point of view of the reporting banks, hence *in a reverse repo the reporting bank is lending* to a counter-party, and in a repo the reporting bank is borrowing money from a counter-party.

3.4.1 General sample

Table 3.1 and Table 3.2 present an overview of our data set in terms of key variables. They show the breakdown of the data along four categories: maturity, currency, counterparty type, and collateral type (Panels A, B, C, and D, respectively). The breakdown

[4]The actual reporting periods differ slightly across the banks but all are toward the end of 2012.

is only for the deals that have no missing information on haircut. For each category, we report the sum of the notional amounts of deals for each subcategory in Table 3.1, and the weighted average of haircuts for each subcategory in Table 3.2. Table 3.1 also shows the percentage of each category in terms of the notional values. The average haircuts in Table 3.2 are weighted by the gross notional of transactions. Both values and haircuts are reported for reverse repos and repos separately. Since repo indicates bank borrowing, we denote the repo values with negative numbers.

[Table 3.1 about here]

By comparing the values of reverse repos and repos, we find that the reporting banks are net borrowers in the repo market (see the row labeled "Total" in Table 3.1). Panel A of Table 3.1 shows that most of the borrowing and lending transactions for these reporting banks have maturities of less than three months. While borrowing exceeds lending for overnight contracts, lending is larger for transactions with maturities of less than three months. This observation suggests that the reporting banks conduct maturity transformation, to some extent. However, for maturities longer than one year they are still net borrowers. Panel B of the same table shows that the reporting banks in our sample borrow more in GBP and EUR followed by USD. They lend mostly in GBP followed by EUR and USD. In net terms, they borrow mostly in GBP and lend in currencies such as EUR, USD, GBP, followed by JPY.

Panel C of Table 3.1 shows that the reporting banks, in aggregate, borrow from counter-parties such as central banks and governments, other banks, money-market funds and broker-dealers, and lend to counter-parties such as other asset managers, insurance companies and pension funds, and through CCPs. This is in line with our general understanding of the money flow pattern in the wholesale funding market.[5] Finally, Panel D in Table 3.1 shows the breakdown based on collateral types. It shows that when the six banks borrow, only a small percent of their repo collaterals is US government bonds. Hence, it appears that the reporting banks use relatively worse collaterals when borrowing than lending in the repo markets. They intermediate in (and borrow against) relatively worse collaterals such as securitisation products and corporate debt. UK government bonds are the most common collateral used both in repo and in reverse repo contracts.

[Figure 3.1 and Figure 3.2 about here]

[5] The first row in Panel C describes the values when counter-party is a reporting bank. The reporting banks report on a UK consolidated basis, but counter-parties are reported on a global basis. Therefore, there may be discrepancies between the reverse repos and repos with the reporting banks.

Inspecting the maturity-currency relationship (Figure 3.1 and Figure 3.2), we see that the majority of contracts (frequency, not notional values) are in EUR and USD followed by GBP and JPY. Most of the contracts have maturity of less than 3 months across all currency groups and only a very small fraction of the contracts have maturity of more than half a year within each currency category. GBP has a relatively higher fraction of reverse repo contracts within 3 to 6 months, compared to other currencies. Repo and reverse repo transactions in JPY and other currencies happen almost exclusively with maturity of less than 1 month.

Panel A of Table 3.2 shows that, except for very long maturities, the reporting banks are able to borrow at slightly lower haircuts than they lend. This observation means that they can use the collateral they receive in a reverse repo to obtain more funding. A similar pattern exists for different currencies as shown in Panel B.

[Table 3.2 about here]

Panel C makes it clear that the above-mentioned haircut advantage for reporting banks arises from trades with hedge funds, other asset managers and, to a lesser extent, with other banks and broker-dealers. In the transactions with these counter-parties, the banks can receive funding at significantly lower margins. This advantage disappears when they trade with central banks and government agencies, insurance and pension funds and other reporting banks.

Finally, Panel D in Table 3.2 shows the breakdown based on collateral types. It displays how margins depend on the quality of collaterals. For example, both repos and reverse repos for German government bonds have a low average haircut, while haircuts for corporate debt and securitisation are higher. The numbers also show that the six reporting banks are able to borrow at a lower haircut compared to the one they charge for the same type of collateral. This is true for all collateral types, except securitized debt. Note that the UK government collateral commands a relatively high haircut, but this is largely due to the longer maturity of the collateralized assets.

3.4.2 Zero-haircut sample

There are a lot of zero haircuts in the data as illustrated by the histogram of haircuts in Figure 3.3: over 35% of the whole sample. Some of these zero haircuts are due to the way haircuts are reported in CCP trades as explained in section 3.5, but even excluding CCP trades, zero-haircut trades are still quite common. This finding is not surprising and has been confirmed by other data collections undertaken at the global level. A summary of the zero-haircut trades trades is presented in Table 3.3. The table shows that the vast majority of contracts are with other banks and are denominated in EUR. Most of the zero-haircut contracts are overnight (84% for the repo sample, 72% for the reverse repo sample), as shown in Figure 3.4.

[Figure 3.3, Figure 3.4 and Table 3.3 about here]

The network graphs in Figure 3.5 and Figure 3.6 illustrate the topology of the zero-haircut trades. The size of each node reflects the number of counterparties with which it has at least one zero-haircut deal. Edge widths show the total number of zero-haircut trades between two given nodes. The figures show that the zero haircut observations from the repo and the reverse repo samples are generated mostly by one or two entities. In the repo market, one of the banks (bank A in Figure 3.6) receives more than 98% of all the zero-haircut trades. This borrower has 89 counterparties who are willing to lend at zero haircut, but it does most zero haircut borrowing with one particular counterparty (24% of all trades) – C697 in Figure 3.6. In the reverse repo market, another bank (bank B in Figure 3.5) is involved in 95% of all the zero-haircut trades. The top 10 counterparties account for 68% of all zero-haircut repo trades and 71% of all zero-haircut reverse repo trades, which shows that a small number of counterparties contribute to the majority of zero-haircut observations. These facts suggest that there are important borrower-lender relationships among the determinants of the zero-haircut trades, supporting our fifth testable hypothesis highlighted above. We investigate the role of bilateral relations further in later sections.

[Figure 3.5 and Figure 3.6 about here]

3.5 The Determinants of haircuts

We now analyze what explanatory variables govern haircuts and in what ways these variables affect them. For this purpose, we ran multiple regressions on reverse repo and repo data separately, with different specifications as described below.

For the most part of the regression analysis, we focus on the sample excluding the trades with CCPs. In practice, CCPs often calculate haircuts (or initial margin requirements) on a portfolio basis. That is, the over-collateralization of repo positions is calculated at the portfolio or netting set level, without applying haircuts on individual transactions.

In our data set, firms still report a transaction-level haircut, but this is often zero given that the ‘true’ haircut is applied at the portfolio level. In such cases, it is not meaningful to look at haircuts on individual transactions that are centrally cleared. In addition, there is only one CCP in our sample, which uses a fixed schedule of haircuts. Therefore, we focus on the sample that excludes CCP transactions.

In order to make sure that the multitude of zero haircuts does not distort our results, in addition to the ordinary least square regressions, we perform two sets of regressions. We use the Tobit model with truncation at zero, and use the logit transform to generate more variation in haircuts and to run logistic regression.

We split the data and consider separately repo and reverse repo transactions since they are different samples: one has reporting banks as borrowers, and the other has the reporting banks as lenders. Moreover, we observe heterogeneity in the counterparties in the two types of transactions, which allows us to conduct a more detailed analysis of the haircut determinants.

Table 3.4 presents all the explanatory variables used in different regressions. We have dummy variables for currencies, collateral types, counterparty types, bank-counterparty pairs and a dummy for collateral bundled in a portfolio with a very safe asset. Other than dummy variables, we use trade-specific variables, collateral rating and maturity, and counterparty characteristics. We also have two measures for counterparty and collateral concentration. Counterparty concentration measures the share of transactions with a specific counterparty in total, evaluated using the notional amount of transactions. It represent how systemically important that counterparty is to the bank. Similarly, collateral concentration is measured by the share of transactions against a specific collateral in total, evaluated using the notional amount of transactions. We also include an interaction term between collateral rating and counterparty rating. The logic behind this term is to find whether counterparty and collateral quality can compensate for each other as a conditional effect.

[Table 3.4 about here]

Table 3.5 shows summary statistics for haircuts and non-dummy explanatory variables for the sample used in the baseline regressions. Except collateral and counterparty ratings which are categorical, other variables in this table are continuous. The summary statistics are represented separately for reverse repos and repos in Panels A and B, respectively, given that haircut practices can potentially differ significantly between the two instruments. Variables have been winsorized at 0.5% level.

Even though haircuts can have a value as high as 46%, the weighted average of haircuts is about 6% for reverse repos and about 2% for repos. Notional values are log-transformed. Maturity values, both for transactions and collateral, are in years. The weighted average of maturity for the transactions is about 22-29 days, while the mean is around 26-29 days. Average collateral maturity used is between 7.5 and 12 years. Collateral and counterparty ratings are modified into numeric scale from 1 to 20, with 20 being the highest rating. The average collateral quality in this scale is about 14, while the average counterparty rating is between 14 and 15.

[Table 3.5 about here]

The summary statistics for counterparty return on assets (RoA), leverage, credit default swap (CDS) spread, and cash ratio are also presented in Table 3.5, and the respective definitions are in Table 3.4. The logic for including RoA is to see how profitability of the counterparty can affect haircuts, and the cash ratio is intended to proxy for liquidity needs. Overall, the summary statistics for reverse repos and repos are not significantly different.

In Table 3.6–Table 3.11 we present the main results of this study. These tables show regression results in order to understand what factors might determine haircuts. The dependent variable is haircut in all tables and explanatory variables are listed in the second column. We have classified explanatory variables into several categories. These categories are shown in the first column. The columns that are labeled with numbers display regression coefficients for different sets of explanatory variables. All continuous explanatory variables are standardized in order to simplify the comparison of coefficients for different variables. Standard errors, which are not reported, are clustered at the reporting bank level. One, two and three stars denote 10%, 5% and 1% significance levels, respectively. The tables present the results for Tobit, OLS, and Logistic regressions for reverse repos and repos.

[Table 3.6, Table 3.7 and Table 3.8 about here]

The results in Table 3.6–Table 3.8 are for reverse repo transactions. In these transactions the reporting bank lends cash and receives collateral, and the counterparty borrows money and delivers collateral to the bank. Hence, counterparty characteristics correspond to borrower characteristics in these transactions. Table 3.6 presents the outcome of the Tobit regression, and Table 3.7 and Table 3.8 show the OLS and Logistic regressions, respectively. The main results that we emphasize below are robust with respect to the choice of models.

We present analogous results for repos in Table 3.9–Table 3.11. In these transactions the reporting bank borrows cash and delivers collateral, and the counterparty lends money and receives collateral. Hence, counterparty characteristics correspond to lender characteristics in these transactions.

[Table 3.9, Table 3.10 and Table 3.11 about here]

Column (1) in all these tables reports the result when the smallest set of explanatory variables is used. In this column, we include currency dummies, notional and maturity of transaction, collateral characteristics (rating and maturity), collateral type dummies, and dummies for counterparty type, but we leave out counterparty characteristics. In column (2) we add counterparty characteristics and concentration measures for counterparties and collateral. Columns (3) and (4) are similar to column (1), but they also include network centrality measures described in section 3.5.2. Analogously, columns (5) and (6) are similar to column (2) but include network centrality measures.

In columns (1) and (2) we do not include the reporting bank characteristics, instead we look for haircut determinants by assessing the effects of explanatory variables within transactions conducted by each reporting bank. This is achieved by including reporting bank fixed effects in the regressions. To account for special relationships in the repo and reverse repo samples, we add a set of dummies for each bank-counterparty pair. We describe the results for these dummies in the next section.

The next section elaborates on the main results presented in Table 3.6–Table 3.11 in light of the six hypotheses formulated in section 3.3.

3.5.1 Tests of hypotheses

Test 1 (collateral quality): The repo haircut is larger when the collateral is of lower quality and/or illiquid.

As aforementioned, collateral quality can be measured using VaR, maturity, rating, and/or asset types. Transaction maturity is also a proxy because the longer the maturity, the riskier the underlying collateral becomes. Furthermore, when the collateral concentration ratio increases, the collateral portfolio pool becomes riskier. To test hypothesis 1, we include VaR of each asset, collateral rating, maturity, asset types in terms of corporate debt, securitisation products, transaction maturity, collateral concentration, notional value in all baseline regressions. We compute the VaR using two approaches. First, the measure is obtained using the historic approach (using the quintiles of the historical return distribution). We calculate simple returns and take the 5-days, 5% VaR as our main measure.[6] Second, we also computed VaR using the parametric approach (using the deciles of the normal distribution). The results are largely similar to the results obtained using the historic approach. In the main text, we provide the results obtained with the historic VaR.

The results from Table 3.6–Table 3.11 show that VaR has a positive impact on the haircut both in the repo market, and in the reverse repo market. Table 3.7 and Table 3.10 show that one standard deviation increase in the 5-day, 5% VaR is correlated with 9 bps increase in the repo haircut and 5 bps increase in the reverse repo haircut. The estimates from the Logit and Tobit regressions confirm the positive and statistically significant results. The effect is robust to adding different controls – the estimates in columns 1-6 barely change.

Similar results are obtained for transaction maturity and securitisation products. Transaction maturity has a significant positive and robust effect on haircuts across all specifications: one standard deviation rise in maturity increases haircuts by 83-103 bps for reverse repos and by 24-47 bps for repos. Securitized collateral increases haircut by 20-64 bps when the reporting banks are lending, and by 9-14 bps when the same banks are borrowing. The notional value of transactions also increases haircuts: one standard deviation increase in notional is correlated with 4-9 bps rise in haircuts for reverse repo transactions, and with 4-6 bps rise in repos.

For the repos, higher collateral concentration – another measure for the riskiness of the collateral portfolio – increases the haircut. Therefore, our reporting banks are charged significantly higher haircut when borrowing relatively large sums against the same collateral. On the other hand, collateral concentration measures do not exhibit any notable effect on haircuts in reverse repo transactions. This might reflect the fact that our reporting banks are relatively larger than their counterparties and are able to absorb a large amount of the same collateral when trading with these smaller counterparties.

Other results on collateral quality and liquidity depend on whether the tests are undertaken with the reverse repo or repo sample, that is, whether banks are lending via reverse repo or borrowing via repo. When banks are lending, they lower the haircut if the collateral rating is higher. When they are borrowing, their lenders require higher haircuts when collateral is of longer maturity and corporate debt. This might reflect

[6] Using 1% or 10 days produces similar results.

the fact that banks in our sample use predominantly corporate debt as collateral assets to borrow.

In general, there is strong evidence that collateral quality and liquidity variables are important determinants of repo haircuts.

Test 2 (counterparty types): The repo haircut is larger when the counterparties in the contract are from different lines of business.

To test hypothesis 2, we define a dummy variable for all non-bank counterparties in our sample (broker-dealers, hedge funds, asset managers, insurance companies, pension funds, central banks, governments and all others). Since all these counterparties are from different lines of business compared to the six reporting banks, the point estimate on the dummy shows the average effect on haircuts when the counterparties are from different business types. In order to see how haircuts applied between a bank and a non-bank entity differ from the haircuts between two banks, we ran analogous regressions to those in Table 3.6–Table 3.11, except that there is only one dummy variable for counterparty type which takes value of 1 if the counterparty is not a bank, and 0 otherwise. The results from Table 3.17 and Table 3.18 in the Appendix show that haircut increases both in the repo market, and in the reverse repo market. For contracts where banks deal with non-bank counterparties, the haircut increases by 9-13 bps in the reverse repo market and by 6-8 bps in the repo market. The estimates from the Logit and Tobit regressions (excluded for brevity) confirm the positive and statistically significant effects.

These results suggest that when banks trade with institutions similar to themselves, on average, they charge lower haircuts, controlling for all observables (counterparty or collateral rating, maturity, etc.). Similar institutions use comparable models, and therefore it is more likely that two banks have less disagreement than two completely different entities, for example, a bank and a hedge fund, hence the higher haircuts for non-bank counterparties in our sample. This might also be due to the fact that there is lower information friction and hence, less adverse selection between counterparties of similar types. This evidence supports both the difference of opinion framework started with Geanakoplos (1997) and the adverse selection framework as in Ozdenoren et al. (2018).

However, these findings should be interpreted with caution. Even though the haircut increases when banks interact with non-bank entities on average, there are heterogenous effects for different non-bank institutions. Table 3.6–Table 3.11 show that banks charge higher haircut when lending to hedge funds and asset managers, but lower haircut when lending to broker-dealers, insurance companies and pension funds. Banks receive higher haircut when borrowing from insurance companies and pension funds, but lower haircut when borrowing from central banks and governments, and broker-dealers. These observations illustrate that despite the positive effect on average, sometimes the haircut can be lower when banks interact with non-bank entities which goes against the hypothesis.

Test 3 (counterparty's quality): The repo haircut is larger when the default probability (credit quality) of borrower is higher (lower), or when the borrower is better privately informed about the quality of the collateral.

To test hypothesis 3, we use the rating and the leverage ratio of the borrower in the reverse repo sample. The results from Table 3.6–Table 3.8 show that higher-rated (lower default probability) borrowers are charged a lower haircut: one unit increase in rating decreases the haircut by 8-21 bps. However, the coefficient is less statistically significant in the Tobit and the Logit regressions and sometimes switches sign, especially in the specifications including network centrality measures. A possible reason for this is the collinearity between the counterparty rating and the centrality measures: the correlation between the two variables is close to 40%.

Using the counterparty's leverage ratio produces more robust results. The coefficients are positive and significant, which shows that riskier counterparties are charged a higher haircut. The OLS estimates show that one standard deviation rise in leverage increases the haircut by 53-79 bps, which is a massive increase. The coefficients from the Tobit and Logit specifications confirm the positive effects.

Removing the bank-counterparty interaction dummies from the regressions shows that the coefficient on rating is more statistically significant and negative across all specifications. Higher-rated counterparties receive a lower haircut in these regressions which shows that some of the rating effects are absorbed by the bank-counterparty interaction dummies. These results are excluded from the paper for brevity. Overall, there is evidence that riskier borrowers are charged a higher haircut.

Table 3.6–Table 3.11 show that among the counterparty types, hedge funds receive massively higher haircuts in all specifications, relative to the baseline haircut received by banks: they are charged 99-157 bps higher haircut, on average. When banks borrow from hedge funds, there is no significant change in the charged haircut as seen from the

coefficients for the repo sample. Broker-dealers both receive and charge a lower haircut in most specifications. Similar effects are observed for central banks and government agencies. Other asset managers are charged higher haircuts, but give lower ones in a contract with the reporting banks. Insurance companies and pension funds charge massively higher haircuts as a lender (90-103 bps more) but receive lower haircuts as a borrower (23-33 bps less).

The results in columns (2), (5), (6) of Table 3.6–Table 3.8 show that larger counterparties are charged lower haircut: one standard deviation increase in size massively reduces the haircut by 93-139 bps. The results for the repo sample are less significant and indicate that larger lenders charge a higher haircut. Higher counterparty CDS increases the haircut both for repos and for reverse repos, but the effect is less significant. Counterparties with missing data on size, rating, CDS, etc. charge a higher haircut as lenders but receive a lower haircut as borrowers. The majority of these counterparties are small banks and some hedge funds. For reverse repos, there are relatively more other asset managers and less broker-dealers with missing data on size, rating, CDS, etc. compared to the general sample.

An important question about haircuts is how collateral risk and counterparty risk interact. There is a significant and positive coefficient on the interaction term between counterparty and collateral rating for the reverse repos. Excluding this interaction term from the regression weakens the magnitude and significance of the effect of counterparty characteristics. This observation means that collateral quality can overshadow counterparty characteristics. It seems that borrowers with lower ratings try to use higher quality collateral to receive a lower haircut, and as a consequence the influences of counterparty attributes are concealed. After accounting for this interaction we can observe that larger counterparties and borrowers with higher ratings are charged lower haircuts.

Test 4 (lender's quality and liquidity): The repo haircut is larger when the default probability and/or liquidity need of lenders is higher.

We use lender's rating to account for default probability in the repo sample. To proxy for liquidity needs, we use lender's cash ratio. The evidence from Table 3.9–Table 3.11 is mixed. The estimates for rating are only marginally significant and positive, which goes against the hypothesis. The estimates for cash ratio are insignificant but negative, which supports the hypothesis. Overall, there is mixed evidence in favor of this hypothesis.

Test 5 (bilateral relationship): The repo haircut is lower for bilateral parties with banking relationship.

Table 3.12 shows the percentage of significant bank-counterparty interaction dummies in column (2) of Table 3.7 and Table 3.10. Figure 3.7 and Figure 3.8 present a network graph of all the bank-counterparty interaction dummies, significant at the 1% level. Red color means the interaction coefficient is negative (lower haircut if the two nodes form a contract). Blue color means the coefficient is positive (higher haircut if the two nodes form a contract). The thickness of the edge between two nodes shows the magnitude of the coefficient on the interaction dummy. The size of each node reflects the number of significant interactions involving the node. The figures are consistent with the hypothesis that relationships matter in haircut determination. The effect is particularly pronounced for the repo market, where one of the banks (E in the figure) receives significantly lower haircuts from most of its counterparties. In the reverse repo market two other banks (B and F) charge lower haircuts in deals with a subset of counterparties. On the other hand, another bank (D on the graph) consistently requires a higher haircut.

[Figure 3.7 and Figure 3.8 about here]

For robustness, we also split the general sample into only hedge funds and only banks. As we see from columns (3) and (4) of Table 3.12, the relationship dummies are significant for the subsample of hedge funds. However, even more of them are significant if we split the sample to only banks (columns (5) and (6)): more than 70% are significant at the 5% level for reverse repo deals. These facts suggest that the special relationships are mostly driven by bank-to-bank effects.

[Table 3.12 about here]

Test 6 (portfolio repos): Risky assets in a portfolio repo with safe assets have lower haircut than purely risky asset repos.

To implement this test, we define a dummy equal to one if an asset is a part of portfolio which contains at least one highest-rated asset (AAA). The coefficient on the dummy for collateral bundled in a safe-asset portfolio from Table 3.6–Table 3.8 shows that lower-rated assets in a portfolio with a safe asset have a lower haircut compared to the same assets in a standalone arrangement. The estimates from Table 3.7 show that combining lower-rated asset in a portfolio with a high-rated asset reduces the haircut by 5-16 bps. A more detailed analysis of the safe-asset portfolios shows that lower-rated counterparties are more likely to bundle assets in such portfolios. Hedge funds are the counterparties with the largest fraction of portfolios bundled with a safe asset.

3.5.2 Network effects

The financial crisis has shown the importance of the interconnectedness of the banking system and the need to analyze risk not by looking at individual institutions in isolation, but by assessing network structure and interplay between institutions. As a result, various studies have used network analysis tools to study the interbank and inter-dealer markets (see, e.g., Denbee et al., 2014; Li and Schürhoff, 2018).

In this part we aim to examine the network structure of the UK repo market using our data set. We use network centrality measures borrowed from the literature on network analysis and employed by Li and Schürhoff (2018). Table 3.13 provides summary statistics of these measures (for definitions see Li and Schürhoff, 2018).

[Table 3.13 about here]

Figure 3.9 displays the repo market network plot. The network plot shows the reporting banks in yellow color and size of the nodes is proportional to total degree measure. In order to see if the network structure affects haircuts in the repo market, we use the principal component of the unweighted and weighted centrality measures in the explanatory regressions. The results are presented in columns (3)–(4) and columns (5)–(6) of Table 3.6–Table 3.11 for reverse repos and repos. We see that the banks with higher centrality measures ask for less haircuts on reverse repos and also pay lower haircuts on repos. The results using weighted or unweighted measures are virtually the same.

[Figure 3.9 about here]

In unreported regressions we use the entire sample including the CCP deals. None of the results mentioned above changes significantly, with two notable exceptions. First, with CCP transactions, the two network measures are not significant in any case, hence we do not observe any meaningful network effect when CCP transactions are included.

Second, including CCP transactions attenuates the impact of counterparty concentration on increasing the haircuts. Overall, given the issues described at the beginning of the section, it seems that including CCP transactions introduces some noise in the way that the architecture of the market affects haircuts, and it is to be expected that the results related to the network measures and counterparty concentration become less significant.

3.6 Conclusion

In this study we analyze the structure of the UK repo market using a novel data set collated by the UK regulator. We examine the maturity structure, collateral types and different counterparty types that engage in this market, and test six theoretical hypotheses of haircut determination. We aim to answer the question of what variables determine haircuts using transaction-level data. We find that collateral quality measured by transaction maturity and VaR have a first order importance in setting haircuts. Banks charge higher haircuts when they transact with non-bank institutions. In particular, hedge funds as borrowers receive a significantly higher haircut even after controlling for measures of counterparty risk. Larger borrowers with higher ratings receive lower haircuts, but this effect can be overshadowed by collateral quality, because weaker borrowers try to use higher quality collateral to receive a lower haircut. Finally, we examine the structure and attributes of the repo market network to assess if the network structure has an influence over haircuts. We find evidence of important borrower-lender relationships. Banks with higher centrality measures ask for more haircuts on reverse repos and pay lower haircuts on repos.

3.7 Figures

FIGURE 3.1: Currency versus maturity of the contracts for the sample of reverse repos. The area of each rectangle represents the fraction of contracts (frequency, not notional values) within a particular maturity-currency group. The area of the entire square is 100%.

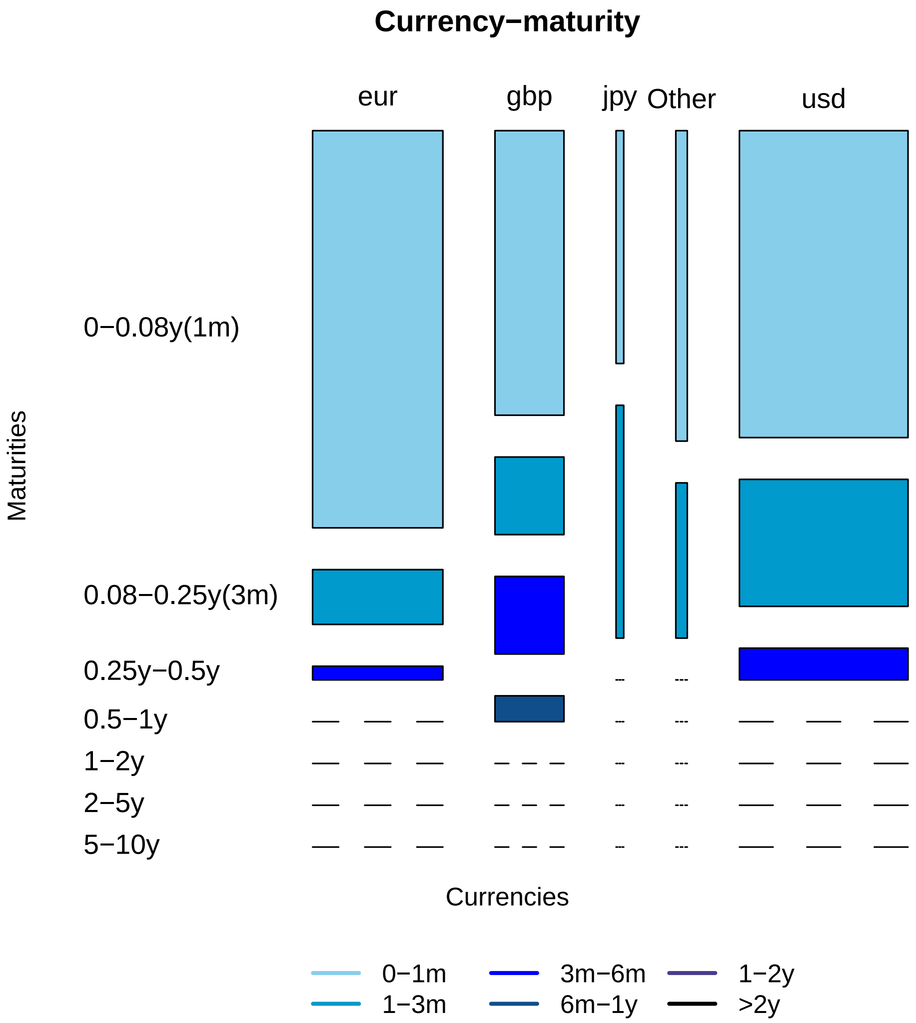

FIGURE 3.2: Currency versus maturity of the contracts for the sample of repos. The area of each rectangle represents the fraction of contracts (frequency, not notional values) within a particular maturity-currency group. The area of the entire square is 100%.

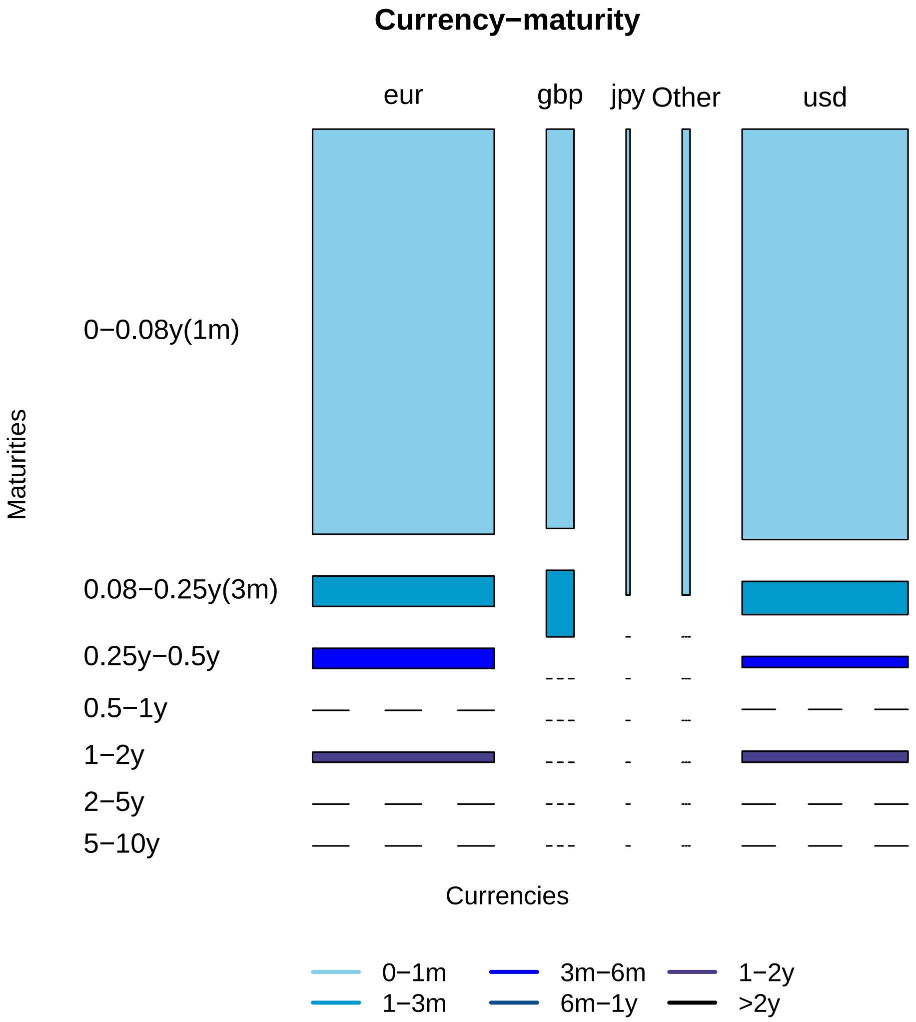

FIGURE 3.3: Histogram of haircuts. The figure shows the density of haircuts.

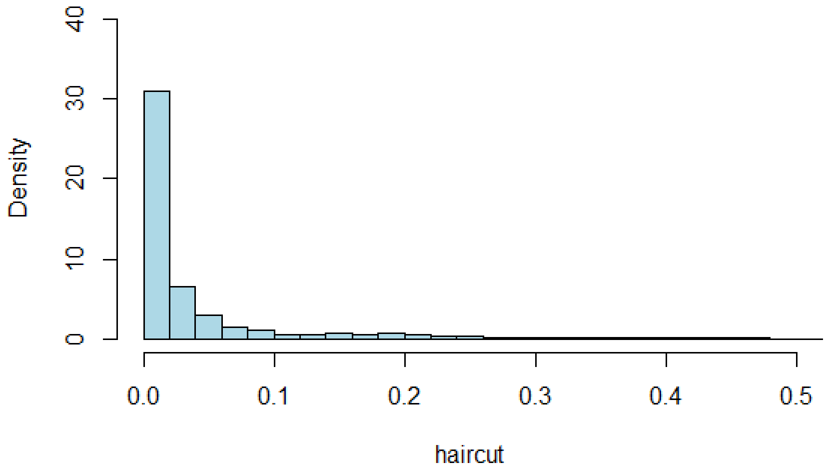

FIGURE 3.4: Zero-haircut sample. Contract maturities. The figure shows the number of zero-haircut contracts for each maturity. The top panel shows the distribution of reverse repos, the bottom – of repos.

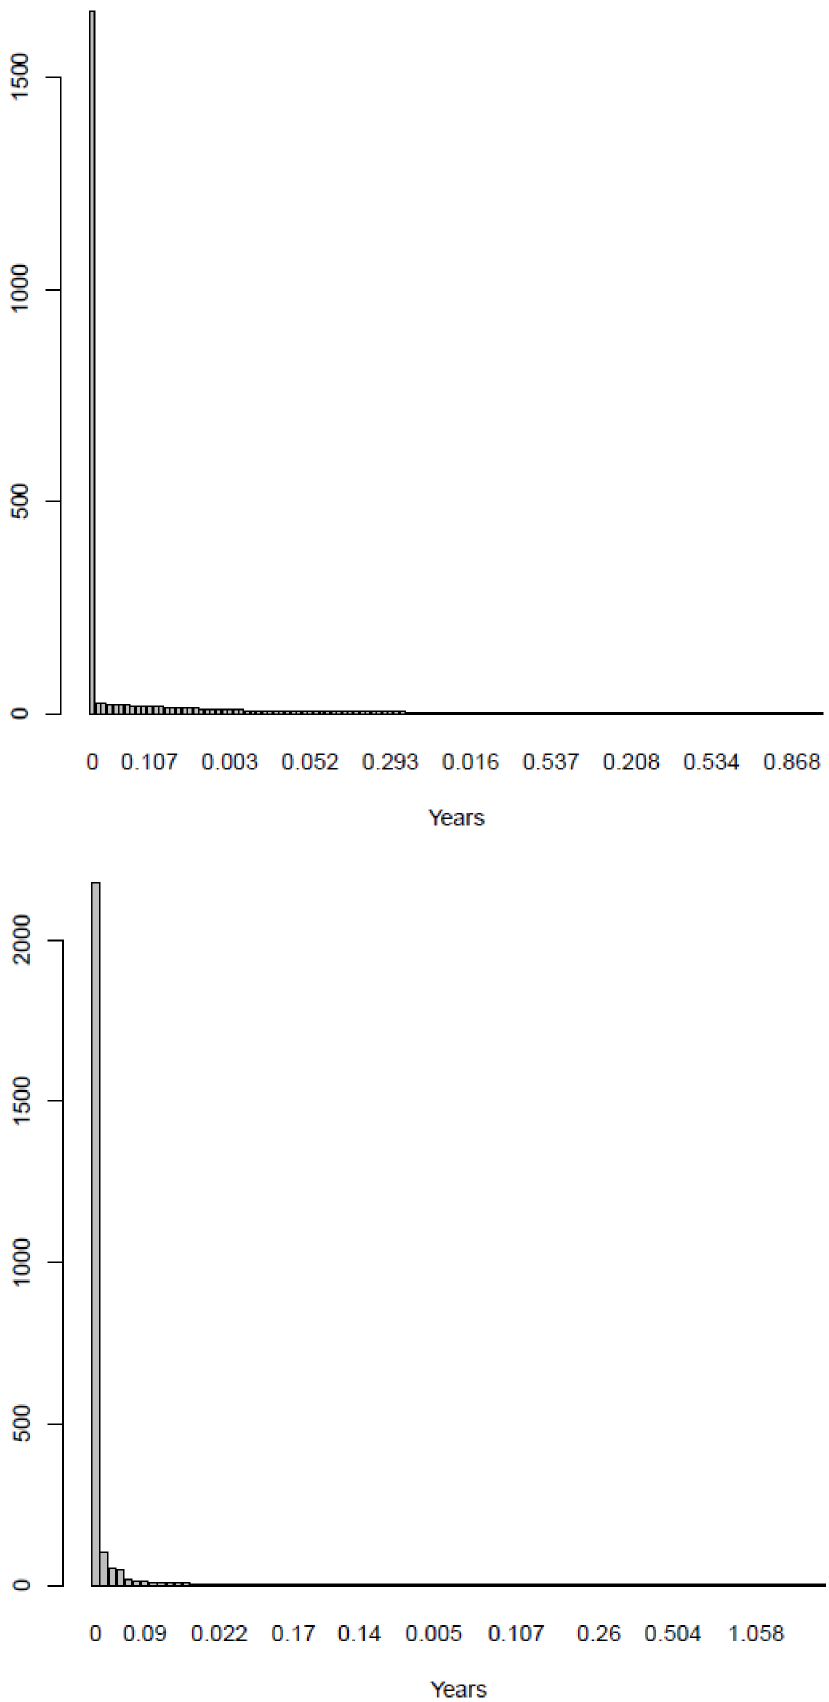

Figure 3.5: Zero-haircut network for reverse repos. The size of each node reflects the number of counterparties with which it has at least one zero-haircut deal. Edge widths show the total number of zero-haircut trades between two given nodes. For Figures 3.5, 3.6, 3.7 and 3.8 nodes labeled A, B, D, E, F, G denote the six reporting banks, and nodes labeled with C and numeric denote their counterparties.

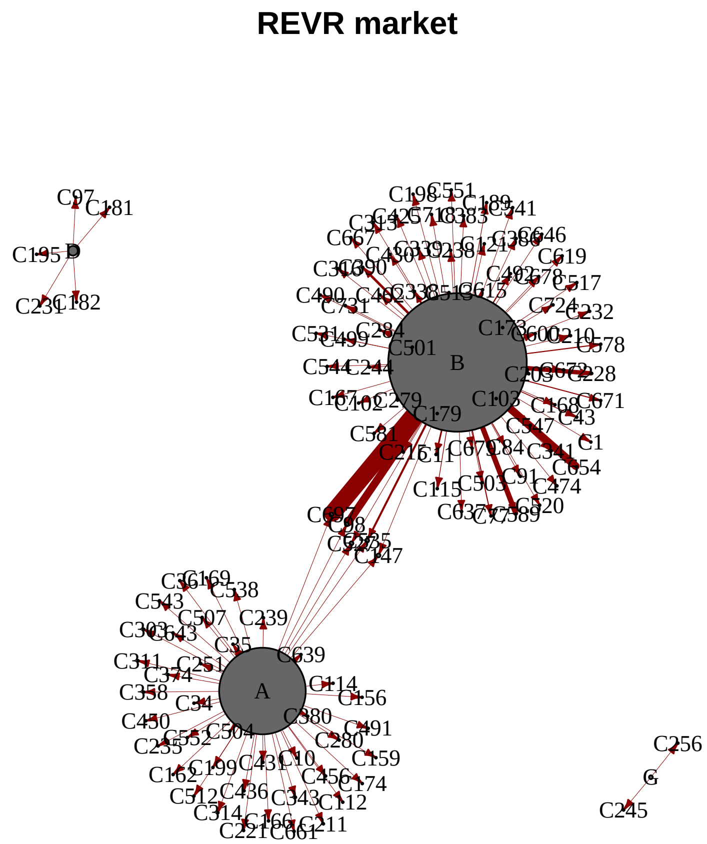

FIGURE 3.6: Zero-haircut network for repos. The size of each node reflects the number of counterparties with which it has at least one zero-haircut deal. Edge widths show the total number of zero-haircut trades between two given nodes.

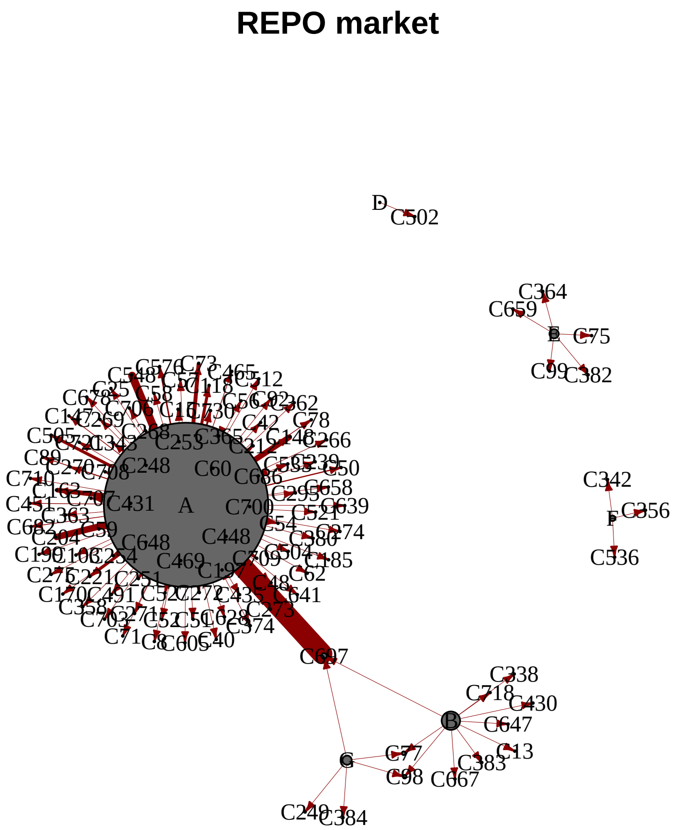

FIGURE 3.7: Significant relationships in the reverse repo market. The figure shows the significant bank-counterparty interaction dummies at the 1% significance level from the OLS regression specification. Red color means the interaction coefficient is negative (lower haircut if the two nodes form a contract). Blue color means the coefficient is positive (higher haircut if the two nodes form a contract). Edge width shows the absolute magnitude of the coefficient on the interaction dummy. The size of each node reflects the number of significant interactions involving the node.

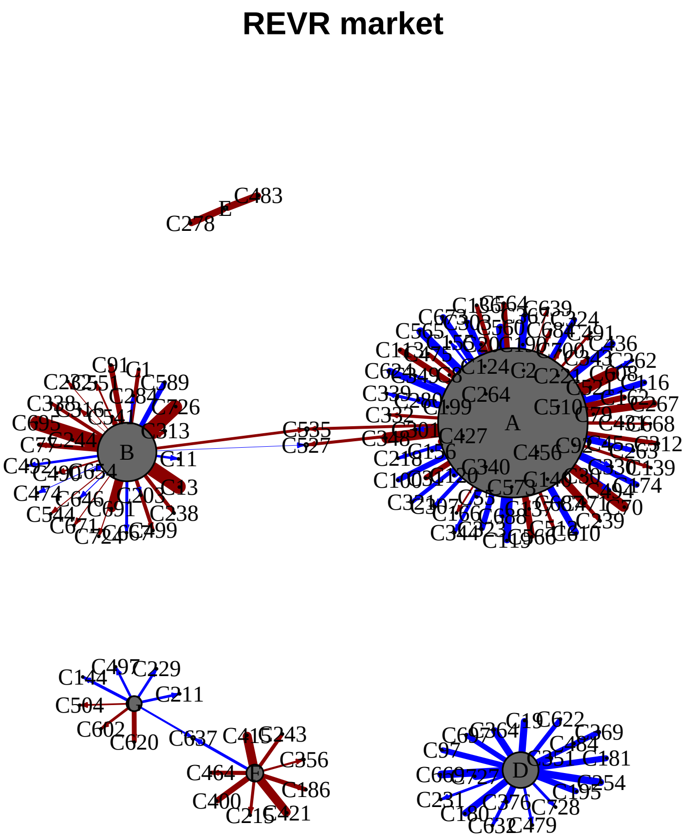

Figure 3.8: Significant relationships in the repo market. The figure shows the significant bank-counterparty interaction dummies at the 1% significance level from the OLS regression specification. Red color means the interaction coefficient is negative (lower haircut if the two nodes form a contract). Blue color means the coefficient is positive (higher haircut if the two nodes form a contract). Edge width shows the absolute magnitude of the coefficient on the interaction dummy. The size of each node reflects the number of significant interactions involving the node.

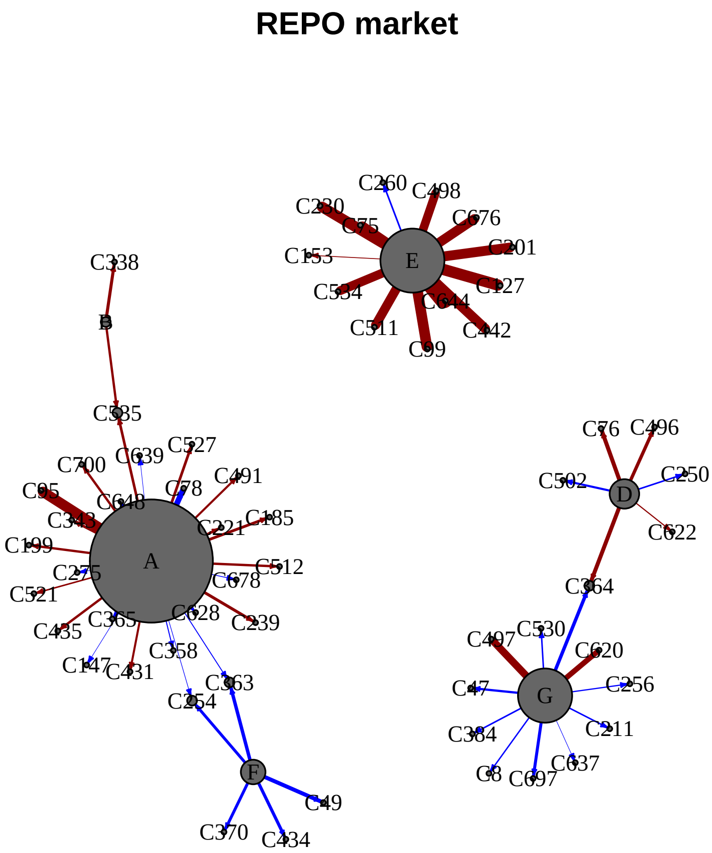

FIGURE 3.9: Network flows plot. The figure shows the flow of money for the sample of six reporting banks.

3.8 Tables

TABLE 3.1: The breakdown of value of contracts by maturity, currency, counterparty type, and collateral type. The table presents the breakdown of the deals (in bn GBP) by maturity, currency, counterparty type, and collateral type (Panels A, B, C, and D, respectively). For each category, it shows the value of the trades in bn GBP and the percentage of total trades for the reverse repos and repos, respectively. The total values in Panels A, B, C and D are based on the data from the six reporting banks that report haircut and collateral information. Discrepancies in row Total between the Panels are due to missing information.

	REVR		REPO		
	Value	Percent	Value	Percent	Net
A. Maturity					
Overnight	29.7	12.2%	-38.1	14.3%	-8.5
1 day-3m	140.7	57.6%	-130.7	48.9%	10.0
3m-1y	65.8	26.9%	-78.1	29.2%	-12.3
1y-5y	8.0	3.3%	-18.5	6.9%	-10.5
5y+	0.0	0.0%	-1.7	0.6%	-1.6
Total	244.2	100.0%	-267.0	100.0%	-22.8
B. Currency					
GBP	110.2	45.1%	-149.8	56.1%	-39.6
EUR	90.6	37.1%	-86.7	32.5%	4.0
USD	30.5	12.5%	-26.8	10.0%	3.7
JPY	6.0	2.5%	-1.6	0.6%	4.4
Other	6.9	2.8%	-2.1	0.8%	4.8
Total	244.2	100.0%	-267.0	100.0%	-22.8
C. Counterparty type					
Another reporting bank [a]	8.2	3.4%	-10.2	3.8%	-2.0
Other banks	29.3	12.0%	-43.6	16.3%	-14.3
Broker-dealer [b]	15.0	6.1%	-15.8	5.9%	-0.8
Hedge fund	15.1	6.2%	-15.5	5.8%	-0.4
MMFs	0.0	0.0%	-1.9	0.7%	-1.9
Other asset managers [c]	11.5	4.7%	-8.3	3.1%	3.2
CCP	145.5	59.6%	-131.3	49.3%	14.2
Insurance and pension	9.5	3.9%	-8.5	3.2%	1.0
Central bank and government	5.5	2.3%	-28.6	10.7%	-23.0
Other [d]	4.4	1.8%	-2.8	1.0%	1.6
Total	244.1	100.0%	-266.6	100.0%	-22.5
D. Collateral type					
US govt	10.9	6.0%	-5.4	2.9%	5.5
UK govt	83.1	45.8%	-111.7	59.1%	-28.6
Germany govt	25.5	14.0%	-19.1	10.1%	6.4
France govt	16.9	9.3%	-7.2	3.8%	9.7
GIIPS [e]	4.1	2.2%	-4.4	2.3%	-0.3
Other sovereign	31.6	17.4%	-16.0	8.4%	15.7
Corporate debt	7.5	4.1%	-11.7	6.2%	-4.2
Securitisation	2.0	1.1%	-13.5	7.1%	-11.5
Other	0.0	0.0%	0.0	0.0%	0.0
Total	181.6	100.0%	-188.9	100.0%	-7.3

[a] The reporting banks report on a UK-consolidated basis, but counterparties are reported on a global basis. Therefore, there may be discrepancies between the reverse repos and repos with the reporting banks.

[b] Broker-dealers are mostly securities firms that are subsidiaries of large banks. [c] Non-leveraged non-MMF mutual funds – asset managers that are not hedge fund or MMF. [d] Includes corporations, schools, hospitals and other non-profit organizations. [e] Greece, Italy, Ireland, Portugal, and Spain government bonds.

TABLE 3.2: The breakdown of average haircuts by maturity, currency, counterparty type, and collateral type. The table presents the breakdown of the deals by maturity, currency, counterparty type, and collateral type (Panels A, B, C, and D, respectively). For each category, it shows the average haircut for the reverse repos and repos, respectively. The averages are weighted by the gross notional of the transactions.

	REVR	REPO
A. Maturity		
Overnight	1.9%	0.7%
1 day-3m	3.2%	1.4%
3m-1y	0.6%	0.5%
1-5y	0.0%	0.7%
5y+	0.0%	0.0%
B. Currency		
GBP	1.4%	0.8%
EUR	1.5%	1.4%
USD	2.6%	0.9%
JPY	0.1%	0.0%
Other	0.2%	0.1%
C. Counterparty type		
Another reporting bank [a]	0.1%	0.2%
Other banks	1.9%	1.4%
Broker-dealer [b]	0.9%	0.6%
Hedge fund	1.4%	0.1%
Other asset managers [c]	1.0%	0.1%
Insurance and pension	0.3%	0.5%
Central bank and government	0.0%	0.3%
Other [d]	0.3%	0.0%
D. Collateral type		
US govt	0.4%	0.0%
UK govt	1.0%	0.4%
Germany govt	0.1%	0.1%
France govt	0.1%	0.1%
GIIPS [e]	0.2%	0.1%
Other sovereign	1.1%	0.2%
Corporate debt	1.1%	0.6%
Securitisation	0.5%	0.8%
Other	0.0%	–
Overall average	1.2%	0.7%

[a] The reporting banks report on a UK-consolidated basis, but counterparties are reported on a global basis. Therefore, there may be discrepancies between the reverse repos and repos with the reporting banks.

[b] Broker-dealers are mostly securities firms that are subsidiaries of large banks. [c] Non-leveraged non-MMF mutual funds – asset managers that are not hedge fund or MMF. [d] Includes corporations, schools, hospitals and other non-profit organizations. [e] Greece, Italy, Ireland, Portugal, and Spain government bonds.

Table 3.3: Summary of the zero-haircut sample excluding deals with CCPs. The table presents breakdown of deals by currency, counterparty and collateral type for the sample of deals with zero haircut, excluding the deals with CCPs. Percentages represent frequency of deals.

Category	Subcategory	REVR	REPO
Currency	GBP	33.6%	6.3%
	USD	22.1%	40.0%
	EUR	40.5%	51.0%
	JPY	1.6%	0.9%
	Other	2.2%	1.9%
Counterparty type	Another reporting bank	4.3%	2.2%
	Other banks	53.4%	68.7%
	Broker-dealer	6.1%	9.5%
	Hedge fund	0.9%	0.0%
	Other asset managers	6.4%	16.0%
	Insurance and pension	11.6%	1.4%
	Central bank and govt	2.2%	1.6%
	Other	15.2%	0.5%
Collateral type	Sovereign	36.7%	44.2%
	Corporate debt	63.0%	43.9%
	Securitization	0.3%	11.9%
	Other	0.0%	0.0%

Table 3.4: Description of the explanatory variables.

Variable	Description
gbp	Dummy variable = 1 if transaction is in GBP.
eur	Dummy variable = 1 if transaction is in EUR.
jpy	Dummy variable = 1 if transaction is in JPY.
othercurrency	Dummy variable = 1 if transaction is not GBP, EUR or JPY.
notional	Log notional of the transaction in millions GBP.
maturity	Maturity of the transaction in years.
collrating	Rating of the collateral: 20 is the highest and 1 is the lowest.
collmaturity	Maturity of the collateral in years.
corpdebt	Dummy variable = 1 if collateral is corporate bond.
securitisation	Dummy variable = 1 if collateral is securitisation.
VaR	Historical 5-day, 5% Value-at-Risk of the asset.
asset in safe portf	Dummy variable = 1 if the asset is in a portfolio with at least one asset rated AAA.
brokerdealers	Dummy variable = 1 if counterparty is a broker-dealer.
hedgefund	Dummy variable = 1 if counterparty is hedge fund.
othermanager	Dummy variable = 1 if counterparty is other asset manager.
ccp	Dummy variable = 1 if counterparty is CCP.
insur&pension	Dummy variable = 1 if counterparty is insurance company or pension fund.
cb&govt	Dummy variable = 1 if counterparty is central bank or government.
other	Dummy variable = 1 if counterparty is other type.
nonbank	Dummy variable = 1 if counterparty is not a bank or broker-dealer.
cptysize	Log size of the counterparty in millions GBP.
cptyroa	RoA of the counterparty.
cptyrating	Rating of the counterparty: 20 is the highest and 1 is the lowest.
cptyleverage	Leverage ratio of the counterparty (risk-weighted assets over equity).
cptycds	CDS spread of the counterparty.
cptycashratio	Cash ratio of the counterparty (cash over short-term debt).
nocptydata	Dummy variable = 1 there is no counterparty data.
cptycon	Concentration of the counterparty measured by the share of transactions with that counterparty in total: higher number indicates more concentration.
collcon	Concentration of the collateral measured by the share of transactions against that collateral in total: higher number indicates more concentration.
cpty&collrating	Interaction term between counterparty rating and collateral rating.
pcu	Principal component of the network centrality measures for unweighted network.
pcw	Principal component of the network centrality measures for weighted network.

TABLE 3.5: Summary statistics for the sample excluding deals with CCPs. The table shows the summary statistics of variables used in the regressions excluding the deals with CCPs, for repo and reverse repo transactions. Variables have been winsorized at 0.5% level. Rating scale is 1–20, with 20 being the highest rating.

Variable	Obs	Mean	Std dev	Min	Max	Average [a]
			A. REVR			
Haircut	8,754	6.25%	10.13%	0.00%	46.15%	6.15%
Notional	10,435	6.25	0.86	3.45	8.32	6.25
Maturity	10,435	0.07	0.14	0.00	3.00	0.06
Collateral maturity	7,085	11.88	10.42	0.22	43.18	12.01
Collateral rating	5,729	14.54	4.83	3.00	20.00	14.60
Ctpy size	6,512	5.17	0.70	3.57	6.25	5.16
Ctpy RoA	6,506	0.29	0.41	-1.26	1.98	0.29
Ctpy leverage	6,469	5.56	1.33	2.97	11.00	5.56
Ctpy CDS	5,593	0.01	0.01	0.01	0.04	0.01
Ctpy cash ratio	6,484	-0.01	5.48	-81.44	4.37	-0.03
Ctpy rating	6,495	14.59	1.28	8.00	20.00	14.60
			B. REPO			
Haircut	7,386	2.37%	5.82%	0.00%	46.15%	2.36%
Notional	11,896	6.18	0.79	3.45	8.32	6.21
Maturity	11,905	0.08	0.35	0.00	3.00	0.08
Collateral maturity	8,993	7.50	7.81	0.22	43.18	7.50
Collateral rating	8,629	14.34	4.99	3.00	20.00	14.33
Ctpy size	8,380	5.37	0.62	3.57	6.25	5.37
Ctpy RoA	8,367	0.36	0.39	-1.26	1.98	0.36
Ctpy leverage	7,300	5.87	1.42	2.97	11.00	5.86
Ctpy CDS	5,908	0.02	0.01	0.01	0.04	0.02
Ctpy cash ratio	8,160	0.01	6.63	-81.44	4.37	0.01
Ctpy rating	8,445	15.19	1.94	8.00	20.00	15.19

[a] Average is weighted by the gross notional of transactions.

TABLE 3.6: Reverse repo Tobit regressions excluding CCPs. The table shows Tobit regression results for reverse repos excluding deals with CCPs, where the Tobit model with truncation at zero is used. The dependent variable is haircut and explanatory variables are listed in the second column. The first column shows the category of explanatory variable. The columns that are labeled with numbers display regression coefficients for different explanatory variables. All quantitative variables (notional, maturity, collrating, collmaturity, VaR, cptysize, cptyroa, cptyleverage, cptycds, cptycashratio, cptycon, collcon, cpty & collrating, pcu and pcw) are standardized. Standard errors (not reported) are clustered at reporting bank level. One, two and three stars denote 10%, 5% and 1% significance levels, respectively.

Category	Variable	(1)	(2)	(3)	(4)	(5)	(6)
Deal var	notional	0.009**	0.006*	0.015**	0.018**	0.016**	0.019**
	maturity	0.157***	0.146***	0.150***	0.098***	0.147***	0.097***
Collateral var	collrating	-0.012***	-0.015***	-0.011***	-0.011***	-0.013***	-0.013***
	collmaturity	-0.0004	0.003	-0.002	-0.0003	-0.001	0.001
	corpdebt	-0.010*	-0.014**	-0.024**	-0.027**	-0.027**	-0.030**
	securitisation	0.037***	0.013	0.099***	0.088***	0.086***	0.074***
	VaR	0.004***	0.004***	0.006***	0.004***	0.005***	0.005***
	asset in safe portf	-0.009*	-0.010**	-0.037***	-0.035***	-0.037***	-0.035***
Cpty type	brokerdealers	-0.008	-0.005	-0.048***	-0.049***	-0.026**	-0.042***
	hedgefund	0.126***	0.065***	0.177***	0.176***	0.175***	0.174***
	othermanager	0.030**	-0.011	0.038**	0.043**	0.046**	0.053***
	insur&pension	0.011	-0.022*	-0.072***	-0.063***	-0.063***	-0.051***
	cb&govt	-0.019	-0.001	-0.092***	-0.079***	-0.131***	-0.110***
	other	0.033***	-0.006	-0.062***	-0.043***	-0.053***	-0.032***
Cpty var	cptysize		-0.166**			-0.208**	-0.216**
	cptyroa		-0.006			-0.041***	-0.036***
	cptyrating		-0.025**			0.015**	0.010**
	cptyleverage		0.127***			0.078***	0.075***
	cptycds		0.001			0.018**	0.018**
	cptycashratio		0.005			-0.016**	-0.009
	nocptydata		-0.230**			0.081	-0.040
Misc	cptycon		0.014***			0.009	-0.010
	collcon		0.005			0.009	0.008
	cpty&collrating		0.001***			0.0004***	0.0004***
Network var	pcu			-0.017***		-0.020***	
	pcw				-0.060***		-0.064***
	Bank FE	Yes	Yes	No	No	No	No
	Bank-Cty FE	Yes	Yes	Yes	Yes	Yes	Yes
	Currency FE	Yes	Yes	Yes	Yes	Yes	Yes
	Observations	3,925	3,907	3,925	3,925	3,907	3,907
	Pseudo R^2	2.893	2.952	2.891	2.891	2.948	2.949

Table 3.7: Reverse repo OLS regressions excluding CCPs. The table shows OLS regression results for reverse repos, excluding deals with CCPs. The dependent variable is haircut and explanatory variables are listed in the second column. The first column shows the category of explanatory variable. The columns that are labeled with numbers display regression coefficients for different explanatory variables. All quantitative variables (notional, maturity, collrating, collmaturity, VaR, cptysize, cptyroa, cptyleverage, cptycds, cptycashratio, cptycon, collcon, cpty & collrating, pcu and pcw) are standardized. Standard errors (not reported) are clustered at reporting bank level. One, two and three stars denote 10%, 5% and 1% significance levels, respectively.

Category	Variable	(1)	(2)	(3)	(4)	(5)	(6)
Deal var	notional	0.003	0.004**	0.006**	0.007**	0.009**	0.009***
	maturity	0.095***	0.103***	0.090***	0.083***	0.097***	0.091***
Collateral	collrating	-0.008***	-0.012***	-0.008***	-0.007***	-0.011***	-0.011***
var	collmaturity	-0.001	0.002	-0.0004	0.001	0.002	0.004**
	corpdebt	-0.008*	-0.009*	-0.013*	-0.011*	-0.015*	-0.012*
	securitisation	0.036***	0.020**	0.064***	0.057***	0.052***	0.046***
	VaR	0.005**	0.005***	0.005**	0.005*	0.005**	0.005**
	asset in safe portf	-0.005*	-0.006**	-0.015***	-0.015***	-0.016***	-0.016***
Cpty	brokerdealers	0.003	0.007	-0.020***	-0.024***	-0.014**	-0.027***
type	hedgefund	0.139***	0.099***	0.157***	0.134***	0.140***	0.111***
	othermanager	0.022**	0.009	0.028**	0.023**	0.031**	0.022**
	insur&pension	0.006	-0.003	-0.026***	-0.032***	-0.023***	-0.033***
	cb&govt	0.008	0.019**	-0.024***	-0.023***	-0.017***	-0.012*
	other	0.017***	0.005	-0.009*	-0.003	-0.009	-0.006
Cpty	cptysize		-0.093**			-0.139**	-0.134**
var	cptyroa		-0.003			-0.017***	-0.010***
	cptyrating		-0.021***			-0.008***	-0.011***
	cptyleverage		0.079***			0.065***	0.053***
	cptycds		-0.003			0.006**	0.006**
	cptycashratio		0.006**			0.001	0.007***
	nocptydata		-0.164***			-0.129***	-0.195***
Misc	cptycon		0.005			-0.001	-0.003
	collcon		0.002			0.004	0.005
	cpty&collrating		0.001***			0.001***	0.001***
Network	pcu			-0.021***		-0.023***	
var	pcw				-0.028***		-0.028***
	Bank FE	Yes	Yes	No	No	No	No
	Bank-Cty FE	Yes	Yes	Yes	Yes	Yes	Yes
	Currency FE	Yes	Yes	Yes	Yes	Yes	Yes
	Observations	3,925	3,907	3,925	3,925	3,907	3,907
	R^2	0.615	0.650	0.637	0.633	0.664	0.658

TABLE 3.8: Reverse repo Logistic regressions excluding CCPs. The table shows Logistic regression results for reverse repos, excluding deals with CCPs. The dependent variable is logit-transformed haircut and explanatory variables are listed in the second column. The first column shows the category of explanatory variable. The columns that are labeled with numbers display regression coefficients for different explanatory variables. All quantitative variables (notional, maturity, collrating, collmaturity, VaR, cptysize, cptyroa, cptyleverage, cptycds, cptycashratio, cptycon, collcon, cpty & collrating, pcu and pcw) are standardized. Standard errors (not reported) are clustered at reporting bank level. One, two and three stars denote 10%, 5% and 1% significance levels, respectively.

Category	Variable	(1)	(2)	(3)	(4)	(5)	(6)
Deal var	notional	0.084*	0.049	0.006**	0.165**	0.146**	0.155***
	maturity	1.480***	1.201***	0.090***	1.181***	1.002***	1.060***
Collateral	collrating	-0.134**	-0.144**	-0.008**	-0.119**	-0.133**	-0.123***
var	collmaturity	0.059*	0.065*	-0.0004	0.104***	0.064*	0.110***
	corpdebt	-0.009	-0.032	-0.013**	-0.081	-0.131*	-0.085
	securitisation	0.336**	0.132	0.064***	0.805***	0.681***	0.705***
	VaR	0.030***	0.033**	0.005**	0.027**	0.027**	0.027**
	asset in safe portf	-0.134**	-0.130**	-0.015***	-0.344***	-0.318***	-0.338***
Cpty	brokerdealers	-0.123	-0.275	-0.020***	-0.758***	-0.561**	-0.751***
type	hedgefund	1.485***	0.779***	0.157***	1.392***	1.921***	1.270***
	othermanager	0.459***	-0.154	0.028***	0.413***	0.585***	0.367***
	insur&pension	0.106	-0.467*	-0.026***	-1.235***	-1.033***	-1.304***
	cb&govt	-1.021***	-1.361***	-0.024***	-1.944***	-2.253***	-2.305***
	other	0.654***	0.024	-0.009*	-0.089	-0.063	-0.131
Cpty	cptysize		-2.252**			-2.826**	-2.556**
var	cptyroa		-0.111			-0.383***	-0.215***
	cptyrating		-0.318***			0.104**	0.035
	cptyleverage		1.619**			1.364**	0.991**
	cptycds		0.082			0.214**	0.206**
	cptycashratio		0.159**			-0.018	0.041
	nocptydata		-4.268***			-0.614	-1.697
Misc	cptycon		0.205***			-0.008	0.027
	collcon		0.133**			0.134*	0.136
	cpty&collrating		0.004***			0.003**	0.003***
Network	pcu			-0.021**		-0.446**	
var	pcw				-0.374**		-0.368**
	Bank FE	Yes	Yes	No	No	No	No
	Bank-Cty FE	Yes	Yes	Yes	Yes	Yes	Yes
	Currency FE	Yes	Yes	Yes	Yes	Yes	Yes
	Observations	3,925	3,907	3,925	3,925	3,907	3,907
	R^2	0.582	0.617	0.595	0.590	0.643	0.638

TABLE 3.9: Repo Tobit regressions excluding CCPs. The table shows Tobit regression results for reverse repos excluding deals with CCPs, where the Tobit model with truncation at zero is used. The dependent variable is haircut and explanatory variables are listed in the second column. The first column shows the category of explanatory variable. The columns that are labeled with numbers display regression coefficients for different explanatory variables. All quantitative variables (notional, maturity, collrating, collmaturity, VaR, cptysize, cptyroa, cptyleverage, cptycds, cptycashratio, cptycon, collcon, cpty & collrating, pcu and pcw) are standardized. Standard errors (not reported) are clustered at reporting bank level. One, two and three stars denote 10%, 5% and 1% significance levels, respectively.

Category	Variable	(1)	(2)	(3)	(4)	(5)	(6)
Deal var	notional	0.004^{**}	0.002^{*}	0.005^{**}	0.007^{**}	0.003^{**}	0.006^{**}
	maturity	0.047^{**}	0.027^{**}	0.042^{**}	0.052^{**}	0.020^{**}	0.034^{**}
Collateral var	collrating	-0.001^{***}	0.001	-0.002^{***}	-0.002^{***}	0.001	0.001
	collmaturity	0.002	0.002	0.003^{**}	0.003^{**}	0.003	0.003
	corpdebt	0.004	0.009^{***}	0.005	0.006^{*}	0.010^{**}	0.010^{***}
	securitisation	0.008	0.011^{**}	0.015^{***}	0.018^{***}	0.019^{**}	0.022^{**}
	VaR	0.011^{***}	0.011^{***}	0.008^{**}	0.009^{***}	0.008^{**}	0.009^{**}
	asset in safe portf	0.011	0.010	0.010	0.010	0.011	0.011
Cpty type	brokerdealers	-0.044^{***}	-0.038^{***}	-0.038^{***}	-0.044^{***}	-0.031^{**}	-0.039^{**}
	hedgefund	-0.020	-0.015	0.014	0.007	0.015	0.011
	othermanager	0.003	-0.008	-0.035^{***}	-0.028^{***}	-0.045^{**}	-0.036^{**}
	insur&pension	0.122^{***}	0.121^{***}	0.114^{***}	0.098^{***}	0.118^{**}	0.107^{**}
	cb&govt	0.007	-0.005^{*}	-0.012^{**}	-0.014^{**}	-0.021^{*}	-0.022^{*}
	other	-0.002	-0.012	-0.056	-0.040	-0.062	-0.044
Cpty var	cptysize		0.008			0.018	0.004
	cptyroa		0.008^{***}			0.008^{**}	0.009^{**}
	cptyrating		0.003			0.003	0.003
	cptyleverage		-0.049^{**}			-0.012	-0.003
	cptycds		-0.003			0.002	0.006
	cptycashratio		0.005			-0.005	-0.003
	nocptydata		-0.091			0.025	0.005
Misc	cptycon		0.017^{***}			-0.002	-0.002
	collcon		0.009^{***}			0.011^{**}	0.010^{**}
	cpty&collrating		-0.0003^{***}			-0.0003^{**}	-0.0003^{**}
Network var	pcu			-0.021^{***}		-0.022^{**}	
	pcw				-0.023^{***}		-0.023^{**}
	Bank FE	Yes	Yes	No	No	No	No
	Bank-Cty FE	Yes	Yes	Yes	Yes	Yes	Yes
	Currency FE	Yes	Yes	Yes	Yes	Yes	Yes
	Observations	3,028	2,915	3,028	3,028	2,915	2,915
	Pseudo R^2	-0.969	-0.933	-0.971	-0.970	-0.932	-0.932

TABLE 3.10: Repo OLS regressions excluding CCPs. The table shows OLS regression results for repos, excluding deals with CCPs. The dependent variable is haircut and explanatory variables are listed in the second column. The first column shows the category of explanatory variable. The columns that are labeled with numbers display regression coefficients for different explanatory variables. All quantitative variables (notional, maturity, collrating, collmaturity, VaR, cptysize, cptyroa, cptyleverage, cptycds, cptycashratio, cptycon, collcon, cpty & collrating, pcu and pcw) are standardized. Standard errors (not reported) are clustered at reporting bank level. One, two and three stars denote 10%, 5% and 1% significance levels, respectively.

Category	Variable	(1)	(2)	(3)	(4)	(5)	(6)
Deal var	notional	0.005***	0.004***	0.005***	0.006***	0.004***	0.005***
	maturity	0.047***	0.029***	0.043***	0.049***	0.024***	0.033***
Collateral	collrating	-0.001	0.001	-0.0001	-0.0002	0.001*	0.001*
var	collmaturity	0.002	0.002	0.003*	0.003**	0.003**	0.003**
	corpdebt	0.004	0.008***	0.006**	0.007**	0.009***	0.009***
	securitisation	0.002	0.004	0.009**	0.012***	0.012***	0.014***
	VaR	0.009**	0.009***	0.007*	0.008**	0.007**	0.007**
	asset in safe portf	0.003	0.003	0.003	0.003	0.003	0.003
Cpty	brokerdealers	-0.012***	-0.005	-0.014***	-0.018***	-0.006	-0.011***
type	hedgefund	-0.005	-0.001	0.0004	-0.003	-0.0004	-0.002
	othermanager	-0.009	-0.015*	-0.045***	-0.039***	-0.049***	-0.042***
	insur&pension	0.096***	0.099***	0.099***	0.090***	0.103***	0.097***
	cb&govt	-0.009	-0.016*	-0.023***	-0.023***	-0.028***	-0.028***
	other	0.003	-0.005	-0.046	-0.034	-0.050	-0.037
Cpty	cptysize		0.023**			0.024**	0.017
var	cptyroa		0.002			0.001	0.001
	cptyrating		0.006***			0.006***	0.006***
	cptyleverage		-0.025***			-0.004	0.003
	cptycds		0.0001			0.005	0.007**
	cptycashratio		0.001			-0.006*	-0.005
	nocptydata		0.041			0.123***	0.109***
Misc	cptycon		0.014**			-0.001	-0.001
	collcon		0.006**			0.008***	0.008***
	cpty&collrating		-0.0002***			-0.0002***	-0.0002***
Network	pcu			-0.013***		-0.014***	
var	pcw				-0.017***		-0.016***
	Bank FE	Yes	Yes	No	No	No	No
	Bank-Cty FE	Yes	Yes	Yes	Yes	Yes	Yes
	Currency FE	Yes	Yes	Yes	Yes	Yes	Yes
	Observations	3,028	2,915	3,028	3,028	2,915	2,915
	R^2	0.572	0.589	0.572	0.572	0.589	0.589

Table 3.11: Repo Logistic regressions excluding CCPs. The table shows Logistic regression results for repos, excluding deals with CCPs. The dependent variable is logit-transformed haircut and explanatory variables are listed in the second column. The first column shows the category of explanatory variable. The columns that are labeled with numbers display regression coefficients for different explanatory variables. All quantitative variables (notional, maturity, collrating, collmaturity, VaR, cptysize, cptyroa, cptyleverage, cptycds, cptycashratio, cptycon, collcon, cpty & collrating, pcu and pcw) are standardized. Standard errors (not reported) are clustered at reporting bank level. One, two and three stars denote 10%, 5% and 1% significance levels, respectively.

Category	Variable	(1)	(2)	(3)	(4)	(5)	(6)
Deal var	notional	0.320***	0.242***	0.005***	0.397***	0.266***	0.321***
	maturity	0.505***	0.262*	0.043***	0.419***	0.143	0.242*
Collateral var	collrating	-0.043***	-0.020	-0.0001	-0.027*	-0.012	-0.003
	collmaturity	0.138**	0.143**	0.003*	0.167***	0.163***	0.164***
	corpdebt	0.482***	0.663***	0.006**	0.553***	0.687***	0.676***
	securitisation	0.380**	0.449**	0.009**	0.748***	0.759***	0.857***
	VaR	0.331***	0.335***	0.007*	0.303***	0.273**	0.287**
	asset in safe portf	0.101	0.139	0.101	0.101	0.139	0.139
Cpty type	brokerdealers	-1.026***	-0.901***	-0.014***	-1.146***	-0.666***	-0.815***
	hedgefund	-0.108	-0.116	0.0004	-0.027	0.011	-0.083
	othermanager	-0.030	-0.174	-0.045***	-1.182***	-1.313***	-1.333***
	insur&pension	1.440***	1.219***	0.099***	1.301***	1.329***	1.167***
	cb&govt	-0.145	-0.404	-0.023***	-0.707***	-0.794**	-0.839***
	other	0.133	-0.062	-0.046	-1.839	-2.124	-1.988
Cpty var	cptysize		0.301			0.715	0.427
	cptyroa		0.090			0.039	0.040
	cptyrating		0.033			0.160*	0.165*
	cptyleverage		-0.500			0.009	0.192
	cptycds		0.107			0.236*	0.268**
	cptycashratio		0.112			-0.141	-0.115
	nocptydata		-0.609			4.105***	3.748***
Misc	cptycon		0.273***			0.004	0.040
	collcon		0.104***			0.122***	0.125***
	cpty&collrating		-0.006***			-0.005***	-0.005***
Network var	pcu			-0.013***		-0.298***	
	pcw				-0.276***		-0.240***
	Bank FE	Yes	Yes	No	No	No	No
	Bank-Cty FE	Yes	Yes	Yes	Yes	Yes	Yes
	Currency FE	Yes	Yes	Yes	Yes	Yes	Yes
	Observations	3,028	2,915	3,028	3,028	2,915	2,915
	R^2	0.641	0.658	0.641	0.641	0.658	0.658

TABLE 3.12: Percentage of significant interactions in the REVR and REPO OLS regressions. The table presents the percentage of significant bank-counterparty interaction dummies in the OLS regressions, column 2 from Table 3.7 and Table 3.10. The first two columns show the results for the general sample, columns 3 and 4 for the subsample of hedge funds only, and columns 5 and 6 for the subsample of banks only.

Significance level	REVR	REPO	REVR, HF	REPO, HF	REVR, Banks	REPO, Banks
	(1)	(2)	(3)	(4)	(5)	(6)
10%	68.1%	57.0%	62.5%	42.1%	77.4%	60.6%
5%	60.6%	50.6%	50.0%	31.6%	71.9%	47.9%
1%	49.7%	34.2%	43.8%	31.6%	65.4%	38.0%

TABLE 3.13: Centrality measures summary

Network type	Measure	Mean
Unweighted	in degree	66.14
	out degree	67.07
	eigenvector centrality	-0.22
	betweenness	15701.22
	closeness out	0.19
	closeness in	0.05
	kcore in	3.67
	kcore out	4.17
	clustering coefficient	0.04
Weighted	in degree (trade number)	151.03
	out degree (trade number)	1931.42
	in degree (value)	4.09 bn
	out degree (value)	38.64 bn
	eigenvector centrality (trade number)	-0.27
	eigenvector centrality (value)	-0.24

3.9 Appendix

TABLE 3.14: The breakdown of value of contracts by maturity, currency, counterparty type, and collateral type. Sample of six banks excluding CCPs. The table presents the breakdown of the deals (in bn GBP) by maturity, currency, counterparty type, and collateral type (Panels A, B, C, and D, respectively) for the sample of six banks excluding CCPs. For each category, it shows the value of the trades in billions GBP and the percentage of total trades for the reverse repos and repos, respectively. The total values in Panels A, B, C and D are based on the data from the six reporting banks that report haircut and collateral information. Discrepancies in row Total between the Panels are due to missing information.

	REVR		REPO		
	Value	Percent	Value	Percent	Net
A. Maturity					
Overnight	23.4	23.7%	-33.0	24.4%	-9.6
1 day-3m	51.6	52.4%	-58.6	43.3%	-7.0
3m-1y	21.8	22.1%	-27.5	20.3%	-5.7
1y-5y	1.8	1.8%	-14.5	10.7%	-12.7
5y+	0.0	0.0%	-1.7	1.2%	-1.6
Total	98.6	100.0%	-135.3	100.0%	-36.7
B. Currency					
GBP	26.9	27.3%	-41.0	30.3%	-14.2
EUR	31.4	31.9%	-65.4	48.3%	-33.9
USD	27.4	27.8%	-25.2	18.6%	2.2
JPY	6.0	6.1%	-1.6	1.2%	4.4
Other	6.9	7.0%	-2.1	1.6%	4.8
Total	98.6	100.0%	-135.3	100.0%	-36.7
C. Counterparty type					
Another reporting bank [a]	8.2	8.3%	-10.2	7.6%	-2.0
Other banks	29.3	29.7%	-43.6	32.2%	-14.3
Broker-dealer [b]	15.0	15.2%	-15.8	11.7%	-0.8
Hedge fund	15.1	15.3%	-15.5	11.5%	-0.4
Other asset managers [c]	11.5	11.7%	-8.3	6.2%	3.2
Insurance and pension	9.5	9.7%	-8.5	6.3%	1.0
Central bank and government	5.5	5.6%	-28.6	21.1%	-23.0
Other [d]	4.4	4.5%	-2.8	2.1%	1.6
Other	0.0	0.0%	-1.9	1.4%	-1.9
Total	98.6	100.0%	-135.3	100.0%	-36.7
D. Collateral type					
US govt	10.2	15.3%	-5.4	6.7%	4.8
UK govt	14.5	21.7%	-17.6	21.9%	-3.1
Germany govt	5.4	8.0%	-12.9	16.0%	-7.5
France govt	4.9	7.3%	-4.7	5.9%	0.1
GIIPS	3.9	5.8%	-3.9	4.8%	0.0
Other sovereign	18.9	28.4%	-10.8	13.4%	8.2
Corporate debt	7.0	10.5%	-11.7	14.5%	-4.7
Securitization	1.9	2.9%	-13.5	16.8%	-11.6
Other	0.0	0.1%	0.0	0.0%	0.0
Total	66.7	100.0%	-80.4	100.0%	-13.8

[a] The reporting banks report on a UK-consolidated basis, but counterparties are reported on a global basis. Therefore, there may be discrepancies between the reverse repos and repos with the reporting banks.

[b] Broker-dealers are mostly securities firms that are subsidiaries of large banks. [c] Non-leveraged non-MMF mutual funds – asset managers that are not hedge fund or MMF. [d] Includes corporations, schools, hospitals and other non-profit organizations. [e] Greece, Italy, Ireland, Portugal, and Spain government bonds.

TABLE 3.15: The breakdown of reverse repos. This table exhibits a finer breakdown of the reverse repo contracts. The numbers are in percentage points and indicate the percentage of notional value in each category. The data is double sorted by counterparty type (columns) and maturity, currency and collateral type in Panels A, B, and C, respectively. Columns 1–8 refer to the following counterparty types:
1. Another reporting bank
2. Other banks
3. Broker-dealer
4. Hedge fund
5. Other asset managers
6. Insurance and pension
7. Central bank & govt, and 8. Other

	Counterparty type								
	1	2	3	4	5	6	7	8	Total
A. Maturity									
Overnight	1.4	18.8	8.0	4.0	2.0	2.1	0.0	2.2	38.4
1 day-3m	0.81	17.5	9.3	10.1	5.6	5.5	2.6	2.2	53.9
3m-1y	0.3	1.7	0.3	0.3	2.5	1.6	0.5	0.5	7.6
1-5y	0.0	0.1	0.0	0.0	0.0	0.0	0.0	0.0	0.1
5y+	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Total	2.5	38.2	17.7	14.4	10.1	9.2	3.1	4.9	100.0
B. Currency									
GBP	1.1	2.8	1.5	2.6	6.3	5.8	0.1	2.6	22.8
EUR	0.6	16.1	2.9	6.3	1.4	3.0	1.3	1.2	32.6
USD	0.7	15.6	11.1	4.0	2.2	0.2	0.7	0.9	35.6
JPY	0.0	1.5	0.9	1.3	0.3	0.0	0.0	0.2	4.0
Other	0.1	2.3	1.3	0.2	0.0	0.1	1.0	0.1	5.0
Total	2.5	38.2	17.7	14.4	10.1	9.2	3.1	4.9	100.0
C. Collateral type									
US govt	0.2	3.1	6.2	0.9	1.4	0.0	0.8	0.0	13.0
UK govt	0.1	0.6	0.9	0.3	7.4	4.9	0.2	2.4	16.8
Germany govt	0.3	1.2	0.4	0.6	0.6	0.6	1.1	0.1	4.9
France govt	0.0	1.7	0.2	0.4	0.3	1.1	0.1	0.2	4.0
GIIPS	0.0	0.2	0.0	3.6	0.1	0.2	0.4	0.0	4.6
Other sovereign	0.6	14.2	3.9	1.5	1.1	0.6	1.7	0.9	24.4
Corporate debt	1.0	10.9	3.3	4.8	1.8	1.9	0.1	2.6	26.4
Securitization	0.1	1.7	1.4	1.4	0.2	0.4	0.1	0.5	5.5
Other	0.0	0.0	0.2	0.0	0.0	0.0	0.0	0.0	0.2
Total	2.3	33.7	16.5	13.6	12.9	9.6	4.5	6.8	100.0

TABLE 3.16: The breakdown of repos. This table exhibits a finer breakdown of the repo contracts. The numbers are in percentage points and indicate the percentage of notional value in each category. The data is double sorted by counterparty type (columns) and maturity, currency and collateral type in Panels A, B, and C, respectively. Columns 1–8 refer to the following counterparty types:
1. Another reporting bank
2. Other banks
3. Broker-dealer
4. Hedge fund
5. Other asset managers
6. Insurance and pension
7. Central bank & govt, and 8. Other

	Counterparty type								
	1	2	3	4	5	6	7	8	Total
A. Maturity									
Overnight	3.5	25.6	10.7	4.8	5.8	1.0	1.7	0.4	53.2
1 day-3m	0.8	10.3	5.8	7.3	2.7	3.9	4.4	0.8	36.3
3m-1y	0.2	2.4	0.8	0.5	0.2	0.5	2.1	0.0	6.7
1-5y	0.3	1.7	1.5	0.0	0.0	0.3	0.0	0.0	3.8
5y+	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Total	4.8	40.0	18.8	12.6	8.7	5.7	8.2	1.2	100.0
B. Currency									
GBP	0.6	1.9	2.2	2.3	2.3	2.8	2.2	0.4	15.1
EUR	1.4	20.9	7.3	6.8	4.5	0.9	4.9	0.5	46.9
USD	2.0	15.5	8.3	3.0	1.8	2.0	0.9	0.3	33.6
JPY	0.8	0.2	0.0	0.2	0.0	0.0	0.0	0.0	1.4
Other	0.0	1.5	1.0	0.2	0.1	0.0	0.1	0.0	2.9
Total	4.8	40.0	18.8	12.6	8.7	5.7	8.2	1.2	100.0
C. Collateral type									
US govt	0.5	1.9	0.6	0.1	0.2	0.0	0.4	0.0	3.7
UK govt	0.3	0.7	0.2	0.7	2.0	1.0	1.9	0.4	7.9
Germany govt	0.4	4.1	0.6	1.9	0.5	0.0	2.2	0.1	10.0
France govt	0.1	2.0	0.2	0.9	0.6	0.0	0.7	0.0	4.4
GIIPS	0.0	1.0	0.5	2.4	0.3	0.0	0.8	0.0	5.0
Other sovereign	2.2	8.3	4.1	2.5	0.8	0.3	2.1	0.3	20.5
Corporate debt	1.3	15.6	7.5	2.9	5.2	3.8	1.0	0.1	37.1
Securitization	0.6	6.5	2.9	0.2	1.1	0.2	0.1	0.0	11.4
Other	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Total	5.3	40.0	16.6	11.7	10.8	5.5	9.2	0.9	100.0

TABLE 3.17: Reverse repo OLS regressions excluding CCPs with nonbank dummy. The table shows OLS regression results for reverse repos, excluding deals with CCPs. The dependent variable is haircut and explanatory variables are listed in the second column. The first column shows the category of explanatory variable. The columns that are labeled with numbers display regression coefficients for different explanatory variables. All quantitative variables (notional, maturity, collrating, collmaturity, VaR, cptysize, cptyroa, cptyleverage, cptycds, cptycashratio, cptycon, collcon, cpty & collrating, pcu and pcw) are standardized. Standard errors (not reported) are clustered at reporting bank level. One, two and three stars denote 10%, 5% and 1% significance levels, respectively.

Category	Variable	(1)	(2)	(3)	(4)	(5)	(6)
Deal var	notional	0.003	0.004**	0.006**	0.007**	0.009**	0.009***
	maturity	0.095***	0.103***	0.090***	0.083***	0.097***	0.091***
Collateral	collrating	-0.008***	-0.012***	-0.008***	-0.007***	-0.011***	-0.011***
var	collmaturity	-0.001	0.002	-0.0004	0.001	0.002	0.004**
	corpdebt	-0.008*	-0.009*	-0.013*	-0.011*	-0.015*	-0.012*
	securitisation	0.036***	0.020**	0.064***	0.057***	0.052***	0.046***
	VaR	0.005**	0.005***	0.005**	0.005*	0.005**	0.005**
	asset in safe portf	-0.005*	-0.006**	-0.015***	-0.015***	-0.016***	-0.016***
Cpty	nonbank	0.090*	0.131**	0.090*	0.090*	0.131**	0.131**
Cpty	cptysize		-0.093**			-0.139**	-0.134**
var	cptyroa		-0.003			-0.017***	-0.010***
	cptyrating		-0.021***			-0.008***	-0.011***
	cptyleverage		0.079***			0.065***	0.053***
	cptycds		-0.003			0.006**	0.006**
	cptycashratio		0.006**			0.001	0.007***
	nocptydata		-0.164***			-0.129***	-0.195***
Misc	cptycon		0.005			-0.001	-0.003
	collcon		0.002			0.004	0.005
	cpty&collrating		0.001***			0.001***	0.001***
Network	pcu			-0.021***		-0.023***	
var	pcw				-0.028***		-0.028***
	Bank FE	Yes	Yes	No	No	No	No
	Bank-Cty FE	Yes	Yes	Yes	Yes	Yes	Yes
	Currency FE	Yes	Yes	Yes	Yes	Yes	Yes
	Observations	3,925	3,907	3,925	3,925	3,907	3,907
	R^2	0.615	0.650	0.637	0.633	0.664	0.658

TABLE 3.18: Repo OLS regressions excluding CCPs with nonbank dummy. The table shows OLS regression results for repos, excluding deals with CCPs with. The dependent variable is haircut and explanatory variables are listed in the second column. The first column shows the category of explanatory variable. The columns that are labeled with numbers display regression coefficients for different explanatory variables. All quantitative variables (notional, maturity, collrating, collmaturity, VaR, cptysize, cptyroa, cptyleverage, cptycds, cptycashratio, cptycon, collcon, cpty & collrating, pcu and pcw) are standardized. Standard errors (not reported) are clustered at reporting bank level. One, two and three stars denote 10%, 5% and 1% significance levels, respectively.

Category	Variable	(1)	(2)	(3)	(4)	(5)	(6)
Deal var	notional	0.005***	0.004***	0.005***	0.006***	0.004***	0.005***
	maturity	0.047***	0.029***	0.043***	0.049***	0.024***	0.033***
Collateral	collrating	-0.001	0.001	-0.0001	-0.0002	0.001*	0.001*
var	collmaturity	0.002	0.002	0.003*	0.003**	0.003**	0.003**
	corpdebt	0.004	0.008***	0.006**	0.007**	0.009***	0.009***
	securitisation	0.002	0.004	0.009**	0.012***	0.012***	0.014***
	VaR	0.009**	0.009***	0.007*	0.008**	0.007**	0.007**
	asset in safe portf	0.003	0.003	0.003	0.003	0.003	0.003
Cpty	nonbank	0.080***	0.067***	0.080***	0.080***	0.067***	0.067***
Cpty	cptysize		0.023**			0.024**	0.017
var	cptyroa		0.002			0.001	0.001
	cptyrating		0.006***			0.006***	0.006***
	cptyleverage		-0.025***			-0.004	0.003
	cptycds		0.0001			0.005	0.007**
	cptycashratio		0.001			-0.006*	-0.005
	nocptydata		0.041			0.123***	0.109***
Misc	cptycon		0.014**			-0.001	-0.001
	collcon		0.006**			0.008***	0.008***
	cpty&collrating		-0.0002***			-0.0002***	-0.0002***
Network	pcu			-0.013***		-0.014***	
var	pcw				-0.017***		-0.016***
	Bank FE	Yes	Yes	No	No	No	No
	Bank-Cty FE	Yes	Yes	Yes	Yes	Yes	Yes
	Currency FE	Yes	Yes	Yes	Yes	Yes	Yes
	Observations	3,028	2,915	3,028	3,028	2,915	2,915
	R^2	0.572	0.589	0.572	0.572	0.589	0.589

Bibliography

Acharya, V. V., Gale, D., Yorulmazer, T., 2011. Rollover risk and market freezes. The Journal of Finance 66 (4), 1177–1209.

Adrian, T., Begalle, B., Copeland, A., Martin, A., 2013. Repo and securities lending. In: Risk Topography: Systemic Risk and Macro Modeling. University of Chicago Press.

Adrian, T., Shin, H. S., 2010. Liquidity and leverage. Journal of Financial Intermediation 19 (3), 418–437.

Adrian, T., Shin, H. S., 2014. Procyclical leverage and Value-at-Risk. Review of Financial Studies 27 (2), 373–403.

Aït-Sahalia, Y., Karaman, M., Mancini, L., 2018. The term structure of variance swaps and risk premia. Working paper.

Alexander, C., Korovilas, D., 2012. Understanding ETNs on VIX futures. Working paper.

Amihud, Y., 2002. Illiquidity and stock returns: Cross-section and time series effects. Journal of Financial Markets 5, 31–56.

Andersen, T., Bondarenko, O., Gonzalez-Perez, M., 2011. A corridor fix for VIX: Developing a coherent model-free option-implied volatility measure. Working paper.

Asquith, P., Covert, T., Pathak, P., 2013. The effects of mandatory transparency in financial market design: Evidence from the corporate bond market. Working paper No. 19417. National Bureau of Economic Research.

Bao, J., Pan, J., Wang, J., 2011. The illiquidity of corporate bonds. Journal of Finance 66, 911–946.

Bardgett, C., Gourier, E., Leippold, M., 2019. Inferring volatility dynamics and risk premia from the S&P 500 and VIX markets. Journal of Financial Economics 131 (3), 593–618.

Barras, L., Malkhozov, A., 2016. Does variance risk have two prices? Evidence from the equity and option markets. Journal of Financial Economics 121, 79–92.

Basak, S., Pavlova, A., 2016. A model of financialization of commodities. Journal of Finance 71 (4), 1511–1556.

Becker, B., Ivashina, V., 2014. Cyclicality of credit supply: Firm level evidence. Journal of Monetary Economics 62, 76–93.

Ben-David, I., Franzoni, F., Moussawi, R., 2018. Do ETFs increase volatility? Journal of Finance 73 (6), 2471–2535.

Bessembinder, H., 2015. Predictable ETF order flow and market quality. Journal of Trading 10 (4), 17–23.

Bian, J., Da, Z., Lou, D., Zhou, H., 2017. Leverage networks and market contagion. Working paper.

Bian, J., He, Z., Shue, K., Zhou, H., 2018. Leverage-induced fire sales and stock market crashes. Working paper.

Bloomberg, 2020. https://www.bloomberg.com/news/articles/2020-04-22/how-etfs-new-whales-of-the-oil-market-are-roiling-crude-prices. Last accessed: April 2020.

Bolton, P., Freixas, X., 2000. Equity, bonds, and bank debt: Capital structure and financial market equilibrium under asymmetric information. Journal of Political Economy 108 (2), 324–351.

Breeden, D., Litzenberger, R., 1978. Prices of state contingent claims implicit in options prices. Journal of Business 51, 621–651.

Carr, P., Wu, L., 2006. A tale of two indices. Journal of Derivatives 13 (3), 13–29.

Carr, P., Wu, L., 2009. Variance risk premiums. Review of Financial Studies 22, 1311–1341.

Cheng, I.-H., 2019. The VIX premium. Review of Financial Studies 32 (1), 180–227.

Cheng, M., Madhavan, A., 2009. The dynamics of leveraged and inverse exchange-traded funds. Journal of Investment Management 7 (4), 43–62.

Christensen, J. H. E., Gillan, J. M., 2018. Does quantitative easing affect market liquidity? Working paper 2013-26. Federal Reserve Bank of San Francisco.

Copeland, A., Martin, A., Walker, M., 2010. The tri-party repo market before the 2010 reforms. Tech. rep., Staff Report, Federal Reserve Bank of New York.

Copeland, A., Martin, A., Walker, M., 2014. Repo runs: Evidence from the tri-party repo market. Journal of Finance 69 (6), 2343–2380.

Dang, T. V., Gorton, G., Holmström, B., 2011. Haircuts and repo chains. Working paper, Columbia University.

Dang, T. V., Gorton, G., Holmström, B., 2013. The information sensitivity of a security. Working paper, Columbia University.

De Fiore, F., Uhlig, H., 2015. Corporate debt structure and the financial crisis. Journal of Money, Credit and Banking 47 (8), 1571–1598.

Demertzis, M., Wolff, G., 2016. The effectiveness of the European Central Bank's asset purchase programme. European Parliament, Policy Department A: Economic and Scientific Policy.

Denbee, E., Julliard, C., Li, Y., Yuan, K., 2014. Network risk and key players: A structural analysis of interbank liquidity. Working paper.

Dew-Becker, I., Giglio, S., Le, A., Rodriguez, M., 2017. The price of variance risk. Journal of Financial Economics 123, 225–250.

Diamond, D. W., Dybvig, P. H., 1983. Bank runs, deposit insurance, and liquidity. Journal of Political Economy 91 (3), 401–419.

Dickey, D., Fuller, W., 1981. Likelihood ratio statistics for autoregressive time series with a unit root. Econometrica 49, 1057–1072.

Dong, X. S., 2016. Price impact of ETP demand on underliers. Working paper, Princeton University.

Duffie, D., 1996. Special repo rates. The Journal of Finance 51 (2), 493–526.

Eraker, B., Wu, Y., 2017. Explaining the negative returns to volatility claims: An equilibrium approach. Journal of Financial Economics 125, 72–98.

Fama, E. F., French, K. R., 1987. Commodity futures prices: Some evidence on forecast power, premiums, and the theory of storage. Journal of Business 60, 55–73.

Financial Stability Board, 2013. Global shadow banking monitoring report. Tech. rep.

Fostel, A., Geanakoplos, J., 2012. Tranching, CDS, and asset prices: How financial innovation can cause bubbles and crashes. American Economic Journal: Macroeconomics 4 (1), 190–225.

FT, 2018. https://www.ft.com/content/be68aac6-3d13-11e8-b9f9-de94fa33a81e. Last accessed: August 2019.

FT, 2020. https://www.ft.com/content/a8a7dc16-1a9a-48f9-9bfa-89269f94218a. Last accessed: April 2020.

Garleanu, N., Pedersen, L. H., 2011. Margin-based asset pricing and deviations from the law of one price. Review of Financial Studies 24 (6), 1980–2022.

Garleanu, N., Pedersen, L. H., Poteshman, A. M., 2009. Demand-based option pricing. Review of Financial Studies 22 (10), 4259–4299.

Geanakoplos, J., 1997. Promises, promises. In: The economy as an evolving complex system II. Addison-Wesley Reading, MA, pp. 285–320.

Geanakoplos, J., 2003. Liquidity, default, and crashes: Endogenous contracts in general equilibrium. In: Advances in economics and econometrics: theory and applications: Eighth World Congress. Vol. 170.

Geanakoplos, J., 2010. The leverage cycle. In: Acemoglu, D., Rogoff, K., Woodford, M. (Eds.), NBER Macroeconomics Annual. MIT Press, pp. 1–65.

Geanakoplos, J., Zame, W., 2002. Collateral and the enforcement of intertemporal contracts. Working paper, Yale University.

Gerba, E., Macchiarelli, C., 2016. Policy options and risks of an extension of the ECB's quantitative easing programme: An analysis. European Parliament, Policy Department A: Economic and Scientific Policy.

Goetzmann, W. N., Rouwenhorst, K. G., 2005. The Origins of Value. Oxford University Press.

Gorton, G., Hayashi, F., Rouwenhorst, K., 2013. The fundamentals of commodity futures returns. Review of Finance 17 (1), 35–105.

Gorton, G., Metrick, A., 2012. Securitized banking and the run on repo. Journal of Financial Economics 104 (3), 425–451.

Gorton, G., Ordonez, G., 2014. Collateral crises. American Economic Review 104 (2), 343–378.

Gottardi, P., Maurin, V., Monnet, C., 2017. A theory of repurchase agreements, collateral re-use, and repo intermediation. CESifo Working Paper Series No. 6579.

Gromb, D., Vayanos, D., 2002. Equilibrium and welfare in markets with financially constrained arbitrageurs. Journal of Financial Economics 66, 361–407.

Gromb, D., Vayanos, D., 2018. The dynamics of financially constrained arbitrage. Journal of Finance 73 (4), 1713–1750.

He, Z., Kelly, B., Manela, A., 2017. Intermediary asset pricing: New evidence from many asset classes. Journal of Financial Economics 126 (1), 1–35.

Hu, X., Pan, J., Wang, J., 2019. Tri-party repo pricing. Journal of Financial and Quantitative Analysis forthcoming.

Huang, J., Wang, J., 2009. Liquidity and market crashes. Review of Financial Studies 22, 2607–2643.

ICMA, 2013. Frequently asked questions on repo. http://www.icmagroup.org/Regulatory-Policy-and-Market-Practice/short-term-markets/Repo-Markets/frequently-asked-questions-on-repo/.

ICMA, 2016. ECB Corporate Sector Purchase Programme. The potential impact on European corporate bond market liquidity. A briefing note by the ICMA IG corporate bond secondary market practices committee. Working paper. International Capital Market Association.

Ivanov, I., Lenkey, S., 2018. Do leveraged ETFs really amplify late-day returns and volatility? Journal of Financial Markets 41, 36–56.

Kadan, O., Tang, X., 2020. A bound on expected stock returns. Review of Financial Studies 33, 1565–1617.

Kashyap, A. K., Stein, J. C., Wilcox, D. W., 1993. Monetary policy and credit conditions: Evidence from the composition of external finance. American Economic Review 83 (1), 78–98.

Keynes, J. M., 1930. Treatise on Money: The Applied Theory of Money. Volume 2. London: Macmillan.

Koijen, R. S., Koulischer, F., Nguyen, B., Yogo, M., 2017. Euro-area quantitative easing and portfolio rebalancing. American Economic Review 107 (5), 621–627.

Koijen, R. S., Moskowitz, T. J., Pedersen, L. H., Vrugt, E. B., 2018. Carry. Journal of Financial Economics 127 (2), 197–225.

Krishnamurthy, A., Nagel, S., Orlov, D., 2014. Sizing up repo. The Journal of Finance 69 (6), 2381–2417.

Krishnamurthy, A., Vissing-Jorgensen, A., 2011. The effects of quantitative easing on interest rates: Channels and implications for policy. Brookings Papers on Economic Activity.

Kwiatkowski, D., Phillips, P. C., Schmidt, P., Shin, Y., 1992. Testing the null hypothesis of stationarity against the alternative of a unit root: How sure are we that economic time series have a unit root? Journal of Econometrics 54, 159–178.

Li, D., Schürhoff, N., 2018. Dealer networks. Journal of Finance 74 (1), 91–144.

Lou, D., 2012. A flow-based explanation for return predictability. Review of Financial Studies 25 (12), 3457–3489.

Malamud, S., 2015. A dynamic equilibrium model of ETFs. Working Paper, Swiss Finance Institute.

Mancini, L., Ranaldo, A., Wrampelmeyer, J., 2016. The euro interbank repo market. Review of Financial Studies 29 (7), 1747–1779.

Martin, A., Skeie, D., Von Thadden, E.-L., 2014. Repo runs. Review of Financial Studies 27 (4), 957–989.

Martin, I., 2013. Consumption-based asset pricing with higher cumulants. Review of Economic Studies 80, 745–773.

Martin, I., 2015. Simple variance swaps. Working paper, London School of Economics.

Martin, I., 2017. What is the expected return on the market? Quarterly Journal of Economics, Oxford University Press 132 (1), 367–433.

Mixon, S., Onur, E., 2015. Volatility derivatives in practice: activity and impact. Working paper.

Mou, Y., 2011. Limits to arbitrage and commodity index investment: front-running the Goldman roll. Working paper, Columbia University.

Nagel, S., 2012. Evaporating liquidity. The Review of Financial Studies 25 (7), 2005–2039.

Newey, W., West, K., 1987. A simple, positive semi-definite heteroscedasticity and autocorrelation-consistent covariance matrix. Econometrica 55, 703–708.

Ozdenoren, E., Yuan, K., Zhang, S., 2018. Dynamic asset-backed security design. Working paper.

Pagano, M., Sanchez Serrano, A., Zechner, J., 2019. Can ETFs contribute to systemic risk? Report of the Advisory Scientific Committee 9, European Systemic Risk Board.

ProShares, 2019. Prospectus of a leveraged VIX ETF. https://www.proshares.com/funds/uvxy.html. Last accessed: August 2019.

Rauh, J. D., Sufi, A., 2010. Capital structure and debt structure. Review of Financial Studies 23 (12), 4242–4280.

Rehlon, A., Nixon, D., 2013. Central counterparties: What are they, why do they matter and how does the bank supervise them? Bank of England Quarterly Bulletin 53 (2), 147–156.

Reuters, 2018. https://www.reuters.com/article/us-global-markets-volatility/explainer-investors-burned-as-bets-on-low-market-volatility-implode-idUSKBN1FQ2GL. Last accessed: August 2019.

Shleifer, A., Vishny, R. W., 1997. The limits of arbitrage. Journal of Finance 52 (1), 35–55.

Simsek, A., 2013. Belief disagreements and collateral constraints. Econometrica 81 (1), 1–53.

Singleton, K. J., 2013. Investor flows and the 2008 boom/bust in oil prices. Management Science 60 (2), 300–318.

Song, Z., Zhu, H., 2018. Quantitative easing auctions of treasury bonds. Journal of Financial Economics 128 (1), 103–124.

S&P500, 2019. https://us.spindices.com/indices/strategy/sp-500-vix-short-term-index-mcap. Last accessed: August 2019.

Sparks, A., 2017. What is the future of the ECB's corporate bond programme. MSCI research. https://www.msci.com/www/blog-posts/what-is-the-future-of-the-ecb-s/0473342167. Last accessed: May 2018.

Sushko, V., Turner, G., 2019. The implications of passive investing for securities markets. Working paper, Bank for International Settlements.

TABB Group, 2016. Bond liquidity metrics: Reading between the lines. Focus note. Markit.

Tang, K., Xiong, W., 2012. Index investment and the financialization of commodities. Financial Analysts Journal 68 (5), 54–74.

Tuzun, T., 2012. Are leveraged and inverse ETFs the new portfolio insurers? Working Paper, Board of Governors, Federal Reserve System.

Vayanos, D., Vila, J.-L., 2009. A preferred-habitat model of the term structure of interest rates. Working paper No. 15487. National Bureau of Economic Research.

Warther, V. A., 1995. Aggregate mutual fund flows and security returns. Journal of Financial Economics 39, 209–35.

Woodford, M., 2012. Methods of policy accommodation at the interest-rate lower bound. Working paper. Columbia University Academic Commons.