Macro-Finance, Asset-Pricing, and the Role of Investor Beliefs

Zur Erlangung des akademischen Grades eines
Doktors der Wirtschaftswissenschaften

(Dr. rer. pol.)

von der Fakultät für Wirtschaftswissenschaften
des Karlsruher Instituts für Technologie (KIT)

genehmigte

DISSERTATION

von

M.Sc. Anian Roppel

Tag der mündlichen Prüfung:	05.05.2025
Hauptreferentin:	Prof. Dr. Marliese Uhrig-Homburg
Korreferent:	Prof. Dr. Fabian Krüger

Karlsruhe, 2025

Acknowledgements

I would like to express my deep gratitude to my supervisor, Prof. Dr. Marliese Uhrig-Homburg, for her continuous support, patience, and the faith she placed in my academic work. Our many interesting and inspiring discussions, her constructive feedback, and her expertise greatly enriched my research and this thesis.

I am deeply indebted to Dr. Nils Unger. He not only co-authored my first paper and helped kick-start my research journey, but also provided guidance, encouragement, and support throughout my doctoral studies. Nils's backing was vital to the successful completion of this work.

I extend my sincere thanks to Prof. Dr. Fabian Krüger, Prof. Dr. Christof Weinhardt, and Jun.-Prof. Dr. Julian Thimme for serving on the examination committee.

Additionally, I am grateful to the whole team at the Institute of Finance at KIT: Dr. Marcel Müller, Dr. Fabian Eska, Prof. Dr. Phillip Schuster, Tobias Kargus, Caroline Grauer, Viktoria Klaus, Philipp Cölsch, Dr. Michael Hofmann, Julian Böll, Annika Jung, Johannes Dinger, Matthias Molnar, Dr. Jelena Eberbach, Dr. Stefan Fiesel, Fanchen Meng, Dr. Michael Reichenbacher, Dr. Jan-Oliver Strych, and Dr. Richard Schubert. Our in-depth discussions, pleasant conversations, shared lunches, skiing adventures, and many other activities were not only a great asset to my academic work but also made this journey a memorable and enjoyable experience.

Lastly, special thanks go to my family and my partner for their unconditional love, patience, and endless support during the past years.

Contents

Acknowledgements i

List of Figures v

List of Tables vii

1 Introduction 1

2 Option-implied lower bound beliefs and the impact of negative interest rates 7
2.1 Introduction 7
2.2 Interest-rate distributions at the lower bound 13
2.2.1 Shape of interest-rate distributions close to the lower bound 13
2.2.2 Maximum lower bound 16
2.2.3 Level-volatility dependence 17
2.3 Option-implied lower bound beliefs 18
2.3.1 Data 18
2.3.2 Time variation in implied skewness, kurtosis and the level-volatility relation . 19
2.3.3 Option-implied limits for lower bound beliefs 21
2.3.4 Distribution characteristics and the level of interest rates 24
2.4 Structural break 28
2.5 Risk premia 40
2.6 Conclusion 43

3 Inflation cyclicality and stock-bond comovement: evidence from news media 45
3.1 Introduction 45
3.2 The role of stock-inflation dynamics in stock-bond comovement 51
3.3 A simple model of time-varying stock-bond correlation 54
3.3.1 Dynamics of consumption and inflation 54
3.3.2 Asset-pricing implications 56
3.4 Data 58
3.4.1 News measure of inflation cyclicality 58
3.4.2 Realized correlations 59
3.5 Consumption and inflation expectations in news data 60
3.5.1 News text sentiment and investor sentiment about (long-run) consumption . . 60
3.5.2 Content classification and investors' (long-term) inflation expecations 63
3.6 Inflation news sentiment and asset correlations 65
3.7 Conclusion 75

4 TIPS under FIRE - Formation of inflation expectations in bond markets . . . 77
4.1 Introduction . . . 77
4.2 Assessing the formation process of inflation expectations . . . 83
4.2.1 Error vs. revisions regression . . . 84
4.2.2 Error autocorrelation . . . 86
4.2.3 Revision autocorrelation . . . 89
4.2.4 Summary of model predictions . . . 90
4.3 Data . . . 93
4.3.1 Survey expectations . . . 93
4.3.2 Market expectations . . . 95
4.3.3 Risk premium proxies . . . 97
4.4 Inflation expectation formation across horizons . . . 99
4.4.1 Results for survey expectations . . . 99
4.4.2 Results for market expectations . . . 101
4.4.3 Horizon increasing overreaction . . . 102
4.5 A risk premia explanation? . . . 106
4.6 Empirical risk premium estimate . . . 108
4.7 High-frequency expectation shocks . . . 114
4.8 Conclusion . . . 116

5 Conclusion . . . 119

Appendix A Option-implied lower bound beliefs and the impact of negative interest rates . . . 123
A.1.1 Proof of Proposition 1 . . . 123
A.1.2 Proof of Proposition 2 . . . 123
A.1.3 Dataset . . . 125
A.1.4 Model-free estimation of option-implied interest rate moments . . . 127
A.1.5 Results for realized moments . . . 129

Appendix B Inflation cyclicality and stock-bond comovement: evidence from news media . . . 137
B.1.1 News reporting on inflation . . . 138
B.1.2 Model details . . . 139

Appendix C TIPS under FIRE - Formation of inflation expectations in bond markets . . . 149
C.1.1 Model derivations . . . 149
C.1.2 Additional results . . . 160

Bibliography . . . 167

List of Figures

2.1 Central-moment-floors near the lower bound . . . 16
2.2 Maximum lower bound . . . 17
2.3 Time-series of lower bound indicators . . . 20
2.4 Negative monetary policy interest rates . . . 22
2.5 Option-implied maximum lower bound . . . 23
2.6 Minimal distance to lower bound . . . 24
2.7 Break-date estimation . . . 30
2.8 Skewness and kurtosis before and after the break . . . 36
2.9 Break-date estimation: level-volatility dependence . . . 37
2.10 Time-series of realized moments . . . 41
3.1 News coverage . . . 47
3.2 Realized correlations . . . 53
3.3 $News_{NRC}$ and $MANews_{NRC}$ index . . . 66
3.4 $News_{NRC}$ and realized stock-inflation correlation . . . 68
4.1 Tests for overreaction and underreaction . . . 92
4.2 Regression results . . . 104
4.3 Cumulative revisions following CPI releases . . . 115
A.1 Break-date estimation realized moments . . . 131
C.1 Regression results: extended sample . . . 161
C.2 Regression results: financial participants . . . 162
C.3 Regression results: forecaster-by-forecaster regressions . . . 164

List of Tables

2.1 Conditional skewness . . . 25
2.2 Conditional kurtosis . . . 26
2.3 Conditional level-dependence of volatility . . . 27
2.4 Structural break regression: US skewness . . . 32
2.5 Structural break regression: EU skewness . . . 33
2.6 Structural break regression: US kurtosis . . . 34
2.7 Structural break regression: EU kurtosis . . . 35
2.8 Structural break regression: US level-dependence of volatility . . . 38
2.9 Structural break regression: EU level-dependence of volatility . . . 39
2.10 Predicting realized moments . . . 43
3.1 Yield volatilities and volatility ratios . . . 53
3.2 Comparison between aggregate text sentiment and other variables . . . 62
3.3 Predicting future consumption growth . . . 63
3.4 News coverage and changes in inflation rates . . . 65
3.5 Summary statistics . . . 67
3.6 Predicting realized stock-inflation correlations . . . 71
3.7 Predicting realized stock-yield correlations . . . 72
3.8 Predicting realized stock-inflation correlations at different horizons . . . 73
3.9 Predicting realized stock-yield correlations at different horizons . . . 74
4.1 Summary statistics of expectations data . . . 98
4.2 Horizon dependence . . . 105
4.3 Estimating the risk premium component . . . 110
4.4 Empirical properties of estimated risk premia . . . 111
4.5 Horizon dependence and an empirical estimate of the risk premium . . . 113
4.6 Horizon dependence of CPI-release-shock revisions . . . 116
A.1 Trading volume . . . 125
A.2 Sample composition . . . 126
A.3 Conditional realized skewness . . . 129
A.4 Conditional realized kurtosis . . . 130
A.5 Structural break regression: US realized skewness . . . 132
A.6 Structural break regression: EU realized skewness . . . 133
A.7 Structural break regression: US realized kurtosis . . . 134
A.8 Structural break regression: EU realized kurtosis . . . 135
B.1 Examples of inflation news articles . . . 138
C.1 Horizon dependence: extended sample . . . 160

C.2 Horizon dependence: financial participants 163
C.3 Horizon dependence: forecaster-by-forecaster regressions 165

Chapter 1

Introduction

Macro-Finance studies the interplay between macroeconomic dynamics and financial markets. Future cash flows of many assets are intimately related to future macroeconomic states. The prices for claims to these future cash flows are therefore determined by the collective beliefs of market participants about the distribution of possible macroeconomic outcomes. This outlines a rich set of challenging questions: What macroeconomic mechanisms explain the asset price dynamics we observe? Conversely, what can asset prices tell us about macroeconomic forces and associated collective investor beliefs? What is the role of *subjective* investor beliefs? How can potentially subjective belief formation help to reconcile observed macroeconomic data with financial market data? Can data on beliefs regarding unobserved macroeconomic dynamics help promote our understanding of economic mechanisms?

This dissertation aims to make contributions towards these questions. The guiding notion in this work is that investor beliefs are crucial to understand equilibrium asset prices and present a unique source of information to test macro-finance theories. The majority of the (macro-) finance literature assumes rational expectations (RE) and full information (FI) of the marginal investor. This posits that agents are perfectly rational, fully conceive the functional form of the true data generating process, know the values of its governing parameters and have seamless access to all relevant information, past and present. Under this paradigm there is no scope for separately studying investor beliefs. Beliefs under full information rational expectations are *objective* in the sense that they reflect the true law of motion of the relevant variables. Therefore, investor beliefs are conveniently implied by the models' assumed dynamics of the (macroeconomic) fundamentals. Empirically, beliefs under this paradigm can be recovered from the historical distribution of the fundamental variables. While providing a useful and handy benchmark, these are strong and in part implausible assumptions. Some aspects of these assumptions can find justification in highly competitive markets where investors are strongly incentivized to provide the best (and most rational) possible forecasts. Other aspects, however, such as investors possessing perfect knowledge of the true model of the world and facing no constraints in collecting and processing information, seem hard to justify theoretically. They can be regarded as a convenient simplification that circumvents the complexity of separately identifying subjective beliefs. To what extent these approximations provide a useful and sufficiently accurate description of reality to understand asset prices is ultimately an empirical question. Recent literature indicates that data on investor beliefs is in many cases hard to reconcile with the RE and/or the FI assumption. In contrast, allowing subjective investor beliefs to deviate from objective beliefs has proven quite successful in addressing various asset pricing puzzles. Born et al. (2023) and Adam and Nagel (2023) provide recent reviews of the existing evidence. Brunnermeier et al. (2021) emphasize that studying the subjective beliefs of investors is a promising avenue for asset pricing research. Following this avenue creates the need of finding ways to measure investor beliefs and to understand their formation. This dissertation contributes to this agenda by developing novel approaches to identify investor beliefs about the lower bound on nominal interest rates (Chapter 2), perceptions of the cyclicality of expected inflation (Chapter 3), and views regarding future inflation developments (Chapter 4). Building on the measurement of these beliefs, we shed light on their determinants and explore their asset pricing implications.

Even under the assumption of perfectly objective beliefs, actual data on beliefs can still be of great value to understand macro-finance relationships. Many intricate mechanisms regarding the dynamics of fundamental variables ingrained in modern finance theories are inherently very difficult to verify directly with fundamental data itself. This encompasses dynamics that operate at a very low frequency, concern very infrequent occurrences, or exclude the possibility of specific occurrences. Prime examples include the literature on long-run risks (Bansal and Yaron (2004), Drechsler and Yaron (2010)), rare disasters (Barro (2006), Wachter (2013), Gabaix (2012)), and the theory of a lower bound on nominal interest rates (Black (1995), McCallum (2000), Eggertsson and Woodford (2003a), Reifschneider and Williams (2000)). In many cases, these theories entertain asset pricing implications that are insufficiently sharp to provide convincing direct tests of the underlying mechanisms. A major criticism of many finance theories is their heavy reliance on such "dark matter" (Chen et al. (2024), Campbell (2017)) – essentially unverifiable elements of the model specification. This seriously questions the robustness of these theories, as it leads to a lack of internal refutability and external validity. It would be desirable to tie central model mechanisms to observable data allowing for direct empirical tests. While historical observations of fundamental data might often not suffice to verify proposed macroeconomic dynamics, investors' *beliefs* regarding these dynamics might be observable to some degree. In any case, as one allows for deviations from rational expectations, the subjective beliefs of investors about the economy replace the actual objective properties of macroeconomic dynamics as the central model ingredient. Therefore, measures of investor beliefs can help to evaluate otherwise hard to verify macro-finance theories and address the "dark matter" concern. We follow this strategy by extracting investor beliefs to evaluate theories regarding the dynamic interaction between long-run consumption risk and expected inflation and to approach questions regarding the lower bound on nominal interest rates.

We extract investor beliefs from a diverse set of data sources spanning surveys, asset prices - including derivatives, and news reporting. Surveys are an increasingly popular source of information on investor beliefs for asset pricing research. They provide a readily available and direct measure of expectations. A great effort mostly led by central banks has resulted in a plethora of available survey data. Survey expectation data exists for a wide range of economic indicators (such as inflation, GDP, interest rates), a spectrum of forecast horizons (from a few quarters up to a decade) and a variety of survey-respondent groups (e.g. individual households, firms, finance professionals). The main concern regarding surveys is that it is not entirely clear how the beliefs they measure connect to the beliefs of the marginal investor who is relevant for asset pricing. Surveys may contain substantial measurement errors due to weak incentives for respondents to provide accurate forecasts. Perhaps more importantly, it is unclear how the heterogenous beliefs of individual respondents should be aggregated to map into the beliefs of the marginal investor. When we address asset pricing questions with survey data in this dissertation, our approach to alleviate these concerns is to complement survey data evidence with evidence from beliefs embedded in asset prices.

Asset prices reflect the beliefs of the marginal investor. Market beliefs regarding several economic indicators can be extracted very sharply from (a combination of) asset prices (such as inflation expectations from the combination of nominal bond prices and inflation protected bond prices) or the prices of specific derivatives (such as expectations of short-term interest rates from the rates of futures on money-market rates). The prices of non-linear derivatives (e.g. options) additionally allow to focus on specific regions of interest of the forecasted distribution and to extract the higher moments of investor beliefs. Market beliefs can be extracted at high frequencies and for a fine grid of forecast horizons. The drawback of beliefs inferred from asset prices is that they reflect risk-neutral beliefs. Hence these beliefs are not purely based on perceptions of real-world probabilities but also contain an adjustment reflecting compensation for systematic risk. Separating between these components in measures of market-price implied beliefs poses a major difficulty. In this dissertation we develop several strategies to tackle this challenge. We also exploit the fact that for an important aspect risk-neutral distributions and real-world distributions coincide: the events in which they have measure zero. This simplifies our evaluation of option-implied beliefs on the lower bound on nominal interest rates, where we are specifically interested in the set of outcomes with zero probability.

While surveys and asset markets collectively present a rich source of information, many interesting aspects of investor beliefs are difficult to measure directly with existing survey and asset price data. An approach that has gained traction recently in the literature is to construct proxies of investor beliefs from textual analysis of print media or social media. Economic commentary covers opinions and views about the future and lays out causes and effects of economic events. Through the richness of narratives, text news can capture the complex interrelationships of current events and distill them into nuanced projections of the future. This gives researchers the opportunity to extract views on more subtle economic dynamics. But rather than purely reflecting public opinion, news media might also play an important role in shaping economic beliefs. Most people primarily consume economic information through news media. Therefore, the media act as information intermediaries between investors and the state of the economy. The theory of Narrative Economics (Shiller (2017)) emphasizes the potent role of media narratives for the transmission of economic beliefs affecting how facts are interpreted and contextualized by readership. While text news present a promising source of information of otherwise hard to measure investor beliefs, they also come with a set of empirical challenges. Any measure constructed from news text information is arguably a much less direct reflection of investor beliefs than surveys and asset prices. Next to the information of interest, media articles also contain a lot of information irrelevant for asset pricing. The route we take in this dissertation to face this challenge is to carefully examine several internal consistency conditions of measures constructed from news data. By relating specific aspects of the employed news measure to economic data and the beliefs observable from survey or asset price data, we increase our confidence in the aspects of the news measure that cannot be observed from alternative data sources.

Through the lens of subjective investor beliefs, this dissertation explores prominent asset pricing questions relating to interest-rate and inflation markets and their interplay with broad stock markets. In particular, the three research articles which this dissertation is built upon analyze (1) how investor beliefs about a lower bound on interest rates evolved over time, (2) how market views regarding the cyclical behavior of inflation drives stock-bond correlations, and (3) how bond market participants form expectations of future inflation across forecast horizons.

In Chapter 2, which is based on the paper *"Option-implied lower bound beliefs and the impact of negative interest rates"* (Roppel, Unger, and Uhrig-Homburg (2023)), we analyze market beliefs about a lower bound on nominal interest rates. Using forward looking information embedded in interest rate option prices both in the US and the Eurozone we track market participants' ex-ante views regarding the location of the lower bound over time. For this purpose, we derive model-free implications for interest rate distributions *close* to the lower bound. This allows us to assess the potential impact of a lower bound constraint on interest rate distributions and to infer a model-free upper limit for the market implied lower bound. We find that the evolution of option implied distributions is inconsistent with the idea of a fixed (zero) lower bound. In contrast, our analysis of short-term interest rate distribution moments during the past decade of ultra-low interest rates suggests a substantial downward revision of lower bound beliefs. Before the appearance of negative interest rates, derivatives markets both in the Eurozone and the US were consistent with a lower bound around zero: highly right-skewed distributions, high kurtosis, and a strong level-dependence of volatility at low interest rates. According to our structural break regressions, interest rate distributions show much weaker signs of being close to any lower bound precisely from the time when interest rates fell below zero in many economies in 2014/2015. Our model-free upper limit for the market implied lower bound shows that markets viewed substantially negative interest rates as possible not only in the Eurozone but also in the US starting from 2014/2015.

Chapter 3 builds on *"Inflation cyclicality and stock-bond comovement: evidence from news media"* (Roppel and Uhrig-Homburg (2024a)) and studies investors' views about the conditional cyclical properties of inflation. Using news article data, we show that media reporting on inflation captures a time-varying cyclical relationship between expected inflation and a persistent component of expected consumption growth. Our novel approach combines the content and the sentiment of news articles to extract investors' assessment of inflation news, i.e. whether inflation increases/decreases are currently being perceived as more favorable/unfavorable. This news-based measure allows us to directly test the central channel in prominent asset pricing models which attempt to explain the time variation in stock-bond correlations. We present a stylized model that provides directly testable asset pricing implications of this channel, given a measure of inflation-cyclicality. We find that our news measure significantly predicts realized correlations between stock market returns and market expectations of inflation as well as realized correlations between stock market returns and bond yields in line with theoretical predictions. These findings provide direct evidence for the (changing) cyclical properties of inflation as a major driver of stock-bond correlations.

Chapter 4 is based on the paper *"TIPS under FIRE - Formation of inflation expectations in bond markets"* (Roppel and Uhrig-Homburg (2024b)) and studies the formation of inflation expectations in treasury bond markets along the term structure. We analyze expectation data from inflation surveys as well as from bond market prices for a range of expectation horizons. Our results point to deviations from the full-information rational expectation (FIRE) model by the marginal investor. Market expectations overreact to inflation news and this overreaction becomes more pronounced at longer horizons. We provide comprehensive evidence that these results are both theoretically as well as empirically extremely difficult to reconcile with an alternative explanation by time-varying risk premia. Theoretically, we show that a risk premia explanation would require very implausible assumptions about the properties of risk premia dynamics. Empirically, we employ different identification strategies to separate expectations and risk premia. Our results uniformly support the conclusion that observed effects are driven by investor expectations and can be tied to deviations from FIRE. The non-FIRE inflation expectations we document offer an explanation for the puzzling excess-volatility and excess-news-sensitivity of long-term yields.

Chapter 5 concludes this dissertation and shares perspectives on possible future research and practical applications.

Chapter 2

Option-implied lower bound beliefs and the impact of negative interest rates

2.1 Introduction

It has long been a strong consensus among academics and practitioners alike that nominal interest rates cannot become negative granted by the opportunity of people to hold currency at presumably negligible cost. The zero lower bound was therefore widely believed to be a major driver for the shape and evolution of the yield curve and macroeconomic dynamics during past episodes of near-zero interest rates (see e.g. Bauer and Rudebusch (2016), King (2019), Gust et al. (2017)). Consequently, considerable effort has been made towards models with the ability to incorporate a lower bound constraint (Krippner (2012), Wu and Xia (2016), Filipović et al. (2017)). Extended periods of deeply negative interest rates in numerous major economies, however, proved the assumption of strictly positive nominal interest rates wrong, possibly calling the central economic role of a lower bound into serious question.

This chapter analyzes *ex ante* beliefs about a lower bound on interest rates through time. The hindsight knowledge that interest rates in fact can become negative tells little about how lower bound beliefs might have impacted economic expectations and interest-rate distributions. The crucial question we ask is how perceptions of a lower bound at the time shaped market expectations. Market participants may in fact have always correctly anticipated that interest rates can very well become (substantially) negative. In this case, the (zero) lower bound should have been largely irrelevant as it would not have entered market expectations. But it is also well conceivable that markets have been caught by surprise by this unprecedented event and that beliefs about a lower bound therefore changed over time. In this case, the zero lower bound may have had an important effect at the time investors were still convinced of the non-negativity constraint. In both cases, it is unclear how much a potentially negative lower bound impacts expectations in a negative interest rate environment. Finally, investors in the US may still hold on to their zero lower bound belief despite opposing international evidence. At least, this is a common implicit assumption in the literature focusing on US markets, where the breach of this bound in other economies has been largely ignored so far.

To address these questions, we study lower bound beliefs embedded in interest rate option prices in both the US and the Eurozone. Options reflect forward-looking distributions of market beliefs and hence allow us to directly estimate the conditional impact of a lower bound on this distribution. So far, the literature studying the impact of a lower bound on interest rate option prices finds conflicting evidence. Filipović et al. (2017) and Mertens and Williams (2021) emphasize the significant impact of the zero lower bound on interest rate expectations during low-rate periods. Filipović et al. (2017) focus on the period before the zero lower bound was breached, while Mertens and Williams (2021) analyze average option-implied distribution characteristics across the entire low-rate period in the US to distinguish between different equilibria. Our analysis complements their findings by exploring how conditional interest rate distributions evolved, particularly during periods when rates fell below zero in major economies. While Mertens and Williams (2021) highlight the zero lower bound's persistent influence on the unconditional distribution, we show that market perceptions shifted dramatically as negative rates became a reality. This shift may explain why Bauer and Chernov (2024) observed limited lower bound effects during certain low-rate periods. Before the breach of the zero lower bound, implied distributions clearly reflected its constraints. However, once rates turned negative, the lower bound's influence weakened significantly, suggesting that markets adapted to this new environment. Our findings provide a more nuanced understanding of the lower bound's role in shaping interest rate expectations and reconcile differing observations about the lower bound's effects during various low-rate phases.

We use three empirical measures to extract lower bound beliefs from option prices. While it is not possible to *directly* determine the location of the market anticipated lower bound, we show in a model-free theoretical analysis how the higher moments of (option-implied) interest-rate distributions can be used to infer whether interest rates are expected to be *close* to a lower bound. Particularly, we show that interest rate skewness and kurtosis are necessarily high and positive when distributions are tightly constrained by a lower bound. From this relationship, we are able to derive a fully model-free and measure-agnostic *upper limit* for the anticipated location of the lower bound. Furthermore, we follow Filipović et al. (2017) and Kim and Singleton (2012) and measure the conditional level-volatility dependence as an indicator for lower bound constrained interest rates. Close to the lower bound there should be a tight relation between the expected interest rate level and the volatility of interest-rate distributions. Together these measures provide us with a comprehensive description of interest rates distribution characteristics at the lower bound. We empirically analyze these distribution characteristics with model-free option-implied moments of short-term interest rates.

Consistent with the anticipation of a zero lower bound, the time-series of all three measures suggest tightly constrained interest-rate distributions during the near-zero interest rate periods both in the US and the Eurozone until 2014/2015. Interest rate skewness, kurtosis, and level-dependence of volatility are all highly elevated. Our model free upper-limit for the implied lower bound shows that implied moments are theoretically in line with a market-anticipated lower bound at zero. Around 2014/2015, however, all indications of lower-bound constrained distributions experience a sharp drop despite decreasing interest rates in the Eurozone and continued low interest rates in the US. The upper-limit for the implied lower bound documents that (substantially) negative interest rates were believed to be possible after 2015. Consequently, interest rate distributions were much less constrained. These results provide first evidence against the hypothesis that markets always anticipated the possibility of negative interest rates, but rather suggest a substantial change in lower-bound beliefs.

To test this mechanism more formally, we proceed in two steps. First, we examine whether interest-rate levels alone are able to explain time-variations in measures for lower-bound constrained distributions. Indeed this is not the case. Next, we formally investigate whether and when market lower bound beliefs may have changed over time. To this end we examine level-dependent distribution characteristics in a structural-break regression framework. Hence, we test whether a potentially lower bound induced level-dependent impact on distributions experienced a structural change because of a (downward) shift in lower bound beliefs. We find that low interest rates lead to strongly lower bound shaped distributions only in the former part of our sample, both in the Eurozone and the US. That is, we detect a large and significant break in the interest rate level effect on our indicators of lower bound constrained interest rates. Distributions of low interest rates after this break show much weaker signs of compression at a lower bound. We estimate this structural change to have happened in 2014/2015 over the exact same period when interest rates in several economies started to violate the zero lower bound. These results indicate that investors both in the US and the Eurozone were forced to revise their formerly held zero lower bound beliefs in light of opposing evidence.

Our results appear not to be driven by time-variation in risk premia. Firstly, we theoretically show that we can use option-implied skewness and kurtosis of the risk-neutral distribution $\mathbb{Q}$ to support many of our main conclusions. In particular, the upper limit for the lower bound on interest rates we derive is equally valid for the physical distribution $\mathbb{P}$ even when inferred from risk-neutral distribution moments. Secondly, we additionally analyze the characteristics of physical interest-rate distributions. To this end, we estimate realized skewness using the approach of Neuberger (2012) and realized kurtosis using the approach of Bae and Lee (2021). Realized moments quite closely resemble the low-frequency variation of implied moments. Running predictive regressions, we show that implied moments are highly predictive of future realized moments, hence providing a good description of the variation in conditional *physical* distribution moments. Finally, we show that the evaluation of physical distribution moments leads us to conclusions very similar to those we draw from risk-neutral distribution moments.

Understanding the impact of a lower bound on interest-rate distributions is important for analyzing and modeling the dynamics of the yield curve and therefore has implications for a broad range of related questions. Most generally, the estimation of dynamic term-structure models critically depends on reasonably accounting for a dynamic lower bound effect. Dynamic term-structure models are for example used by central banks to extract market expectations of future interest rates and inflation, to estimate risk premia, and to evaluate the success of their monetary policy. Furthermore, corporations rely on term-structure models to manage their interest-rate risk. Misspecification of the lower bound induced interest-rate asymmetry may materially distort risk-management decisions.

More particularly, shadow-rate models have been employed to generate alternative indicators summarizing monetary activities when the short rate is constrained by the lower bound. Most prominently, the shadow-short rate (SSR) is frequently used as generated regressor in macroeconomic studies to capture the effect of unconventional monetary policy. This practice relies on the assumption of a lower-bound constrained short rate and empirical results have been shown to be extremely sensitive to the lower bound specification (see Krippner (2015, 2017)). Moreover, the strength of a lower bound impact likely has important implications for the dynamics of risk premia embedded in interest rates. Within a shadow-rate framework for example, a tightly binding lower bound leads to the compression of risk premia and thereby induces a dependence between the level of interest rates and excess returns. Assessing to what degree current interest rates reflect market-participants' lower bound considerations also has direct practical consequences for the extraction of likely short-rate paths from futures prices: the distributional asymmetry at a lower bound drives a wedge between the expected value of future yields and their most likely realization.

Beyond interpretation and modeling of interest-rate *levels*, accounting for a lower bound impact is also crucial to understand interest-rate *volatility*. The large decline in interest-rate volatility during the period of low interest rates poses the question whether we experienced a corresponding decrease in underlying macroeconomic uncertainty or whether we simply observed compression at a lower bound. Distinguishing these effects is important when analyzing the relation between economic growth and uncertainty, where interest-rate volatility provides a natural candidate for measuring uncertainty (see e.g. Creal and Wu (2017) and Istrefi and Mouabbi (2018)). Additionally, the understanding of lower bound effects contributes to the literature on unspanned stochastic interest-rate volatility. Proximity to the lower bound compresses volatility and can be expected to decrease sensitivity to unspanned factors. Therefore, the portion of volatility variation spanned by the term structure should be increasing with interest rates approaching their lower bound. In fact, Filipović et al. (2017) document a substantial increase in volatility-level dependence during the near-zero interest-rate period in the US and attribute this phenomenon to the zero lower bound.

Related Literature Our research contributes to the large body of literature that studies the role of a lower bound on nominal interest rates. A vast macro literature with seminal contributions by Krugman et al. (1998) and Eggertsson and Woodford (2003b) has emphasized the economic importance of a lower bound. Consequently, a wide range of empirical facts during the last decades of ultra-low interest rates has been explained by the presence of a (zero) lower bound. Bauer and Rudebusch (2016) estimate a large impact on monetary policy expectations due to near-zero interest rates. King (2019) finds that the zero lower bound significantly weakened the duration channel of bond supply. Using counterfactual experiments, Gust et al. (2017) estimate that the zero lower bound constraint was responsible for about 30% of the US GDP contraction during the Great Recession. Chung et al. (2012) conclude that the ZLB had a first-order impact on macroeconomic outcomes in the United States. Christiano et al. (2011) measure a substantially increased government-spending multiplier near the zero lower bound. Analyzing high frequency data, Swanson and Williams (2014) show that long-term yields became less sensitive to macroeconomic announcements after 2008. Datta et al. (2021) argue that the zero lower bound induced stronger positive correlation between oil and equity returns and increased their sensitivity to macroeconomic news. Gourio and Ngo (2020) establish that the zero lower bound environment created positive correlation between inflation and stock returns. All of these studies crucially rely on assumptions about the location of the lower bound and its incorporation by market participants. We aim to provide missing empirical evidence for these (zero) lower bound beliefs.

Numerous studies explore the impact of a lower bound through lower bound consistent term-structure models fitted to the yield curve. These papers essentially rely on assumptions about the location of the lower bound to infer the model-implied impact on the yield-curve. For the US, Japan, and the UK, the lower bound is commonly assumed to be zero or slightly above. Examples include Krippner (2012, 2013), Christensen and Rudebusch (2015), Bauer and Rudebusch (2016), Wu and Xia (2016) for the US; Ichiue et al. (2013), Kim and Singleton (2012) for Japan; and Andreasen and Meldrum (2015), Carriero et al. (2018) for the UK. Some studies treat the lower bound as a free model-parameter (see Christensen and Rudebusch (2016), Kim and Priebsch (2013); the online implementation of Wu and Xia (2016) for the euro area), hence estimate its location from yield-curve data. None of these term-structure model applications, however, provides direct evidence for their respective lower bound assumptions. We fill this gap as we gather direct evidence from option-implied distributions without relying on model assumptions.

Evidence from option markets regarding a lower bound on interest rates has been studied before by Filipović et al. (2017), Mertens and Williams (2021), and Bauer and Chernov (2024). Our research contributes to this literature in several ways. Most importantly, this chapter explicitly addresses the violation of the zero lower bound and examines the implications for option-implied distributions. We provide first-time evidence for a change in lower bound beliefs related to this zero lower bound violation. For US markets, considering a departure from the zero lower bound conviction sets our research apart from prior literature. Moreover, our study extends existing option-based evidence for the US market to the European market, allowing us to analyze lower bound effects in a negative interest rate environment. Regarding the empirical methodology, Mertens and Williams (2021) and Bauer and Chernov (2024) also use skewness as indication for lower bound constrained yields. Their use of skewness is motivated by an interest rate model with a simple truncation mechanism, generating asymmetric distributions at the lower bound. We theoretically prove that these asymmetric distributions are not just the result of simplistic modelling assumptions but are - to a certain, quantifiable extent - a model-independent feature of lower bound constrained yields. Beyond focusing on the distribution asymmetry (skewness) we also consider the conditional level-volatility dependence as in Filipović et al. (2017) and additionally the fat tails (kurtosis). This provides us with a rich characterization of lower bound constrained distributions and hence helps us to identify lower bound effects more sharply than previous literature. Our research also differs from the literature in the estimation of interest-rate distribution moments. Filipović et al. (2017) use model-implied ATM volatilities and Mertens and Williams (2021) calculate the moments from a full probability density function fitted to option prices. In contrast, we directly estimate model-free moments from option prices which has shown to deliver robust results and hence became the standard approach in the literature. Bauer and Chernov (2024) use a similar model-free methodology but apply it to options on treasury futures instead of options on short-term interest rates as we do. For our purposes, options on short-term rates come with two advantages: (1) they capture lower bound effects more precisely, since short-term rates are more directly affected by a lower bound than long-term rates. (2) We do not need an approximate mapping from the futures price distribution to an interest-rate distribution as is required for options on treasury futures.

Few studies have addressed the violation of the zero lower bound. Lemke and Vladu (2017), Kortela (2016), Wu and Xia (2017), and Geiger and Schupp (2018) estimate lower-bound consistent term-structure models for the euro area. They all assume a time-varying lower bound to accommodate for the fact that euro area interest rates incrementally entered negative territory. A central input in these studies are estimates for the time-varying location of the lower-bound for which they rely on a variety of proxies. Our findings are consistent with the assumption of a time-variation of the lower-bound. However, contrary to assumptions in this literature, our results suggest that with the violation of the zero lower bound, markets revised their beliefs to such an extent that any lower-bound effects became negligible afterwards. In other markets, the violation of the zero lower bound has been largely ignored so far. Even after the general validity of the zero lower bound assumption has been proven wrong in multiple economies, numerous studies still rely on its unchanged validity in other markets where this bound has not been breached yet. Examples include Caballero and Simsek (2021), Johannsen and Mertens (2021), Andreasen and Meldrum (2019), Mertens and Williams (2021) for the US and Iiboshi et al. (2022) for Japan. We find this assumption to be inconsistent with evidence from option markets which suggests a revision of lower-bound beliefs not only in the Eurozone but also in the US where interest rates remained positive until now.

Regarding our model-free methodology to estimate option-implied interest-rate moments our research builds on Bakshi and Madan (2000), Carr and Madan (2001), and Bakshi et al. (2003). Recent studies applying similar methods to interest-rate distributions include Trolle and Schwartz (2014), Choi et al. (2017), Bauer et al. (2022), and Bauer and Chernov (2024). Our approach to estimate realized interest-rate moments builds on Neuberger (2012) and Bae and Lee (2021).

2.2 Interest-rate distributions at the lower bound

In this section, we explain how interest rate distributions behave as they approach the lower bound, based on observable data. This sets the foundation for our empirical analysis, where we compare these characteristics to option-implied distributions in order to infer market beliefs about the lower bound. It's important to note that we cannot directly pinpoint the exact location of the market-implied lower bound. This is because areas of the distribution assigned zero probability are indistinguishable from highly unlikely outcomes. This remains true even with forward-looking option-implied probabilities, as we don't observe contingent claims priced at zero for events considered impossible. Therefore, for our analysis, we focus on describing how interest rate distributions behave when nearing an implied lower bound. Our characterization relies on a model-free description of the interest-rate distribution *shape* and established implications for the dynamic relationship between the expectation and volatility of interest-rate distributions. While we cannot directly determine the exact location of the lower bound, our characterization of distribution shapes *close* to the lower bound allows us to derive model-free expressions for a *upper limit* on the lower bound, based on observed conditional distribution moments. The great advantage of this maximum lower bound is that it remains agnostic regarding the measure under which conditional moments are estimated.

2.2.1 Shape of interest-rate distributions close to the lower bound

We leverage the fact that interest-rate distributions close to the lower bound exhibit a very distinctive shape. They are highly asymmetric and fat tailed. Importantly, this is not just the prediction of a specific model, but a direct and general implication that follows immediately from the definition of a lower bound. We start by developing some intuition behind this theoretical result and then present formal proof.

Market participants anticipate a lower bound on interest rates by excluding the possibility of future interest-rate realizations below such a floor in their decision making. When market expectations of interest rates are well above their floor relative to the associated uncertainty, the impact on the shape of the interest-rate distribution is likely to be negligible. However, as interest rates approach their lower bound, the range of realizations around the expectation conceived possible by the market becomes increasingly constrained. When interest-rate expectations are close to the lower bound, a high level of uncertainty may only come from the probability of large upward deviations from the mean, as downward deviations are limited. With interest rate uncertainty necessarily driven mostly by extreme positive deviations at the lower bound, the distribution shapes are by definition asymmetric and fat tailed. This intuition holds for any distribution, whether it incorporates risk premia or not (i.e., for $\mathbb{P}$ and $\mathbb{Q}$ distributions alike). The shape characteristics of interest rates close to the lower bound can be concisely summarized in terms of the corresponding central moments: interest rates constrained at a lower bound exhibit high positive skewness and high (excess) kurtosis. This insight has immediate empirical implications: When market participants incorporate a lower bound in their investment decisions, this is directly reflected by the prices of interest rate contingent claims. Therefore, at the market anticipated lower bound, option-implied interest-rate distributions are characterized by high positive skewness and high kurtosis.

In the following we formalize our intuition about the relationship between proximity to the lower bound and the shape characteristics of interest-rate distributions. Our derivations are completely model-free, i.e. we make no assumptions about the interest-rate distribution other than the existence of its finite mean and variance. The distribution may be any continuous, discrete, or mixed distribution. Hence, we also do not need to rely on any a-priori assumption on the mechanism through which a lower bound might shape interest-rate distributions. In particular, note that our derivations equally apply to interest-rate distributions under the physical measure $\mathbb{P}$ as well as to option-implied distributions under the risk-neutral measure $\mathbb{Q}$. We will show that interest-rate distributions close to the lower bound are always asymmetric and that their central moments skewness and kurtosis are bounded from below. We are able to derive explicit expressions for these bounds and specify them in terms of the distance of interest-rate expectations to the lower bound relative to their volatility.

Consider L_T the LIBOR rate fixed at T for the period T,T+Δ with conditional date-t probability distribution $P_t \in \mathcal{P}$. $\mathcal{P}$ is the set of all probability measures with support on $[LB, \infty)$ and with finite and given mean μ and variance σ^2. Otherwise, the distribution P_t is allowed to be completely arbitrary. LB is the respective LIBOR-rate floor and $d = \frac{(\mu - LB)}{\sigma}, d \in (0, \infty)$ denotes the distance of the distribution mean to the floor in terms of its standard-deviations.
We start by noting that

PROPOSITION 1. *For $0 < d < 1$, the LIBOR-rate distribution P_t cannot be symmetric*[1].

Proof. See Appendix A.1.1

Next, we derive explicit lower bounds for skewness and kurtosis of the LIBOR-rate distribution. Our main result is:

PROPOSITION 2. *Let $\underline{\mathcal{S}}$ and $\underline{\mathcal{K}}$ denote lower bounds for skewness and kurtosis respectively over the set $\mathcal{P}$, i.e.*

$$\underline{\mathcal{S}} = \inf_{P \in \mathcal{P}} \mathbb{E}^P \left[\left(\frac{L_T - \mu}{\sigma} \right)^3 \right] \tag{2.1}$$

$$\underline{\mathcal{K}} = \inf_{P \in \mathcal{P}} \mathbb{E}^P \left[\left(\frac{L_T - \mu}{\sigma} \right)^4 \right]. \tag{2.2}$$

Then the solution to (2.1),(2.2) *exists and is given by*

$$\underline{\mathcal{S}} = \frac{1 - d^2}{d},$$

$$\underline{\mathcal{K}} = \begin{cases} 1 & \text{for } d \geq 1 \\ \frac{1+d^6}{d^4+d^2} & \text{for } 0 < d < 1. \end{cases}$$

Proof. See Appendix A.1.2

Figure 2.1 displays derived central moment bounds for a varying distance of the distribution-mean to the lower bound. Note that the bounds $\underline{\mathcal{S}}$ and $\underline{\mathcal{K}}$ are only feasible for rather pathological interest-rate distributions where all the mass is placed on only two points. Restricting interest-rate distributions to more realistic shapes would of course further tighten the central-moment-bounds at the interest-rate floor.

Derived relationships enable us to empirically identify interest rates close to their lower bound simply based on the central moments of their distribution. Interest-rate distributions tightly constrained by a lower bound necessarily exhibit both high positive skewness and kurtosis. At any point in time, we can therefore assess whether empirically observed (implied) distributions are theoretically consistent with markets anticipating a lower bound in close proximity to current

[1] P_t is said to be symmetric if and only if $\forall \delta \in \mathbb{R} : f(\mu - \delta) = f(\mu + \delta)$ where f is the probability density function or the probability mass function respectively.

Figure 2.1: Central-moment-floors near the lower bound

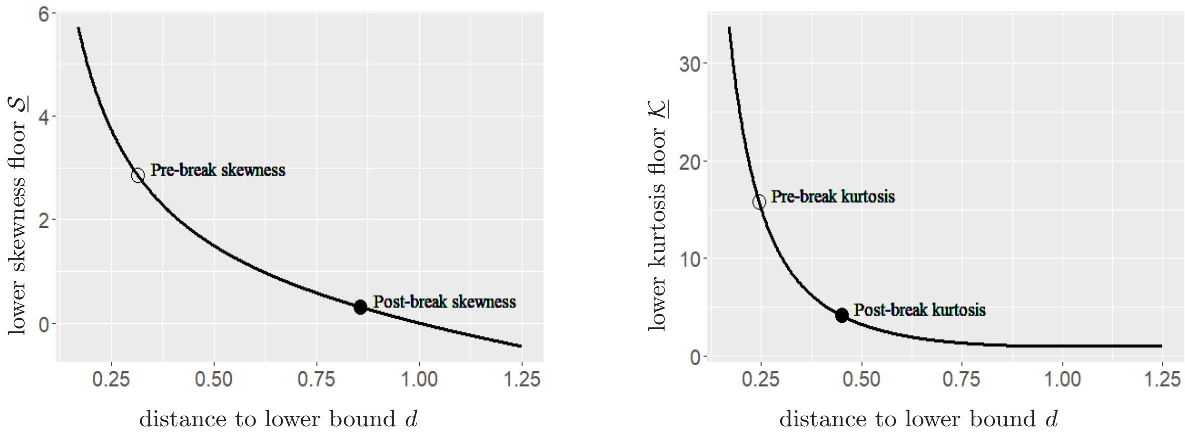

This figure shows the theoretical floors of skewness $\underline{\mathcal{S}}$ (left panel) and kurtosis $\underline{\mathcal{K}}$ (right panel), depending on the distance d of the distribution mean to the lower bound in terms of standard deviations.

expectations. For (implied) interest-rate distributions not exhibiting high levels of skewness and kurtosis we can exclude the possibility of a tightly binding lower bound. We will use these theoretical relationships to evaluate how lower-bound beliefs may have developed over time.

Our theoretical analysis of distribution shapes is useful to empirically distinguish in a model-agnostic way between interest rates stuck right at their lower bound - where skewness and kurtosis become explosive - and interest rates with still some space left to move downwards. At interest-rate levels more moderately away from the lower bound, however, the model-free impact on distributions becomes too weak as to expect strong further guidance from skewness and kurtosis regarding the exact distance to the lower bound.

2.2.2 Maximum lower bound

Derived bounds for skewness and kurtosis allow us to infer the maximum lower bound $\overline{LB}$. That is, a model-free upper limit for the location of the market-anticipated lower bound that is still consistent with observed values for conditional skewness and kurtosis.

COROLLARY 1. *A lower limit for d is given by*

$$d \geq \max\left[\left(-\frac{\mathcal{S}}{2}+\sqrt{\left(\frac{\mathcal{S}}{2}\right)^2+1}\right), \left(\sqrt{\frac{\mathcal{K}+1}{2}-\sqrt{\left(\frac{\mathcal{K}+1}{2}\right)^2-1}}\right)\right] \equiv \underline{d} \tag{2.3}$$

and the corresponding upper limit for LB is given by

$$LB \leq (\mu - \sigma \underline{d}) \equiv \overline{LB} \tag{2.4}$$

With this result we are able to exclude the possibility of a market anticipated lower bound higher than $\overline{LB}$ at any point in time just based on the moments of the distribution, without relying on any model assumptions. Note that $\overline{LB}$ is also measure-agnostic. Meaning, this limit holds under any alternative equivalent distribution measure, regardless of the specific measure used to observe the distribution moments. This directly follows from the fact that any two equivalent measures by definition have the same support. We can therefore empirically calculate a fully model-free upper limit for the market-implied lower bound: moments of the $\mathbb{Q}$-distribution, observed directly from option prices in a model-free manner can be translated into $\overline{LB}$, which also holds for the equivalent $\mathbb{P}$-distribution.

Figure 2.2: Maximum lower bound

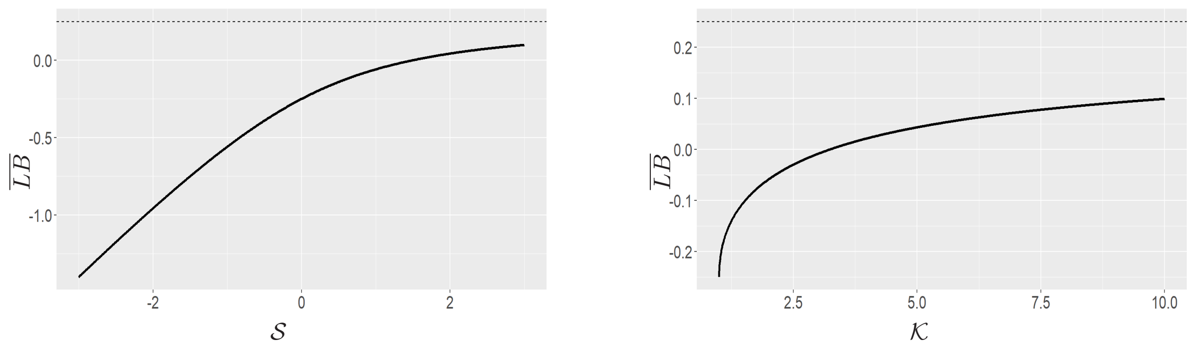

This figure shows the model-free maximum lower bound $\overline{LB}$ depending on conditional skewness (left panel) respectively kurtosis (right panel) with $\mu = 0.25$ and $\sigma = 0.5$.

2.2.3 Level-volatility dependence

Contrary to skewness and kurtosis, there is no similar model-free bound for the volatility of distributions close to the lower bound. Highly skewed distribution shapes could theoretically still produce arbitrarily high volatility even when expectations are very close to the lower bound. However, several studies have argued for a strong positive level-volatility dependence as a feature of lower-bound constrained yields (Kim and Singleton (2012), Filipović et al. (2017), King (2019), Kim and Priebsch (2013), Andreasen and Meldrum (2019), Christensen and Rudebusch (2015)). Intuitively, as interest-rate expectations increase, distributions become less constrained giving ceteris paribus rise to a higher volatility. And even though the level-volatility dependence of interest rates is not a fully model-independent feature of yields near the lower bound, it also does not depend on the exact way the lower bound is enforced in an interest-rate model. In fact, high level-volatility dependence close to the lower-bound has been shown to emerge from every popular lower-bound consistent dynamic term-structure model.[2] We will therefore consider the conditional level-volatility dependence of interest rates in addition to the derived model-free shape characterizations to identify lower-bound constrained yields.

2.3 Option-implied lower bound beliefs

2.3.1 Data

Our analysis uses futures and options on short-term interest rates for the US and the Eurozone market. We obtain daily settlement prices of futures on the three-month USD-LIBOR rate (Eurodollar futures) and the three-month Euribor rate, as well as options written on those futures. Data comes from Thomson Reuters Datastream and covers the period from June 2006 (May 2005) to December 2019 for the Eurodollar (Euribor) contracts. Eurodollar contracts traded at the CME and Euribor contracts traded at ICE Futures Europe constitute two of the largest markets for short-term interest rates in the world. Table A.1 in the Appendix reports trading statistics for December 2019 according to CME and ICE Futures Europe. To set these numbers into perspective we also list figures for US Treasury futures and options. Trading volumes of Eurodollar futures and options even exceeds those of the highly liquid Treasury note contracts. Naturally, trading in Euribor contracts is a bit more muted but average daily volume in these futures (options) still amounted to 537,316 (49,195) trades. These figures indicate that liquidity in LIBOR futures and options compares quite well with other segments of the interest rate market which ensures reliable market prices in our sample.

Payoffs are directly tied to the underlying reference rate (L_t) as futures contracts settle in cash based on $100 - L_T$. Futures prices are quoted as $100 - f_{t,T}$, where $f_{t,T}$ is the date-t future LIBOR rate with maturity T. In the remainder we will directly refer to $f_{t,T}$ as the futures rate. The underlying reference rate for the futures and option contracts is the three-month LIBOR (Euribor) rate, a daily benchmark rate representing unsecured interbank lending conditions in the US (Eurozone). It is the key cash-market benchmark within USD (EUR) money markets, serving as reference rate for forward rate agreements, interest rates swaps, futures, and options. Hence

[2] For a discussion of volatility compression at the lower bound see e.g. Kim and Priebsch (2013), Andreasen and Meldrum (2019), Kim and Singleton (2012) and Christensen and Rudebusch (2015) for shadow rate models; Kim and Singleton (2012) for affine term structure models with square-root processes as in Cox et al. (1985) and Dai and Singleton (2000); Andreasen and Meldrum (2019) and Kim and Singleton (2012) for quadratic term structure models as in Ahn et al. (2002b) and Leippold and Wu (2002); Filipović et al. (2017) for linear-rational models; Monfort et al. (2017) for their Autoregressive Gamma-zero (ARG_0) framework; and Feunou et al. (2022) for their tractable term structure framework.

our futures and options sample directly captures market expectations of the most important short-term interest rate in these two economies.

For each trading day we obtain estimates of the option-implied moments $Vol_{t,T}(L_T)$, $Skew_{t,T}(L_T)$ and $Kurt_{t,T}(L_T)$ for all available quarterly contract expirations. We limit our analysis to quarterly expiries as these contracts are most liquid and allow for straightforward estimation of implied moments. Quarterly Eurodollar (Euribor) options have maturities of up to four (two) years, providing us with 16 (8) points across the short-end of the yield curve at any given day. Appendix A.1.3 gives further details on our data selection procedure and the composition of our sample.

We estimate option-implied moments with a model-free methodology using the insights of Bakshi and Madan (2000), Carr and Madan (2001), Bakshi et al. (2003) and Jiang and Tian (2005). An application of their fundamental result allows us to derive formulas for volatility, skewness and kurtosis of the risk-neutral LIBOR rate distribution at given horizon T in terms of the corresponding futures rate and the cross-section of options with the same maturity written on those futures. Details of our model-free methodology are provided in Appendix A.1.4.

2.3.2 Time variation in implied skewness, kurtosis and the level-volatility relation

Figure 2.3 plots time-series of the interest futures rate in the US and the Eurozone, together with the three measures that we use to assess ex-ante beliefs about the lower bound: skewness, kurtosis and the level-volatility relationship. For illustration purposes we present time-series of option-implied skewness, kurtosis and the interest rate futures rate that refer to a constant maturity of one year. These constant maturity values are obtained by linearly interpolating between available adjacent maturities. To measure the level-volatility relationship, we follow Filipović et al. (2017) and regress changes in volatility on changes in the futures rate,

$$\Delta Vol_{i,t} = \alpha + \beta \Delta f_{i,t} + \epsilon_{i,t} \tag{2.5}$$

In the last row of Figure 2.3 we show 2-year rolling estimates of β.

Unconditionally, implied skewness and excess-kurtosis are slightly positive, confirming the results of Trolle and Schwartz (2014) for the swaption market. The level-volatility regression coefficient β is somewhat positive in the unconditional estimation. However, all three measures show substantial time-variation over our sample. Very notably, all three indicators are highly elevated during the period of near-zero interest rates - in the US 2008 – 2015 and in the Eurozone 2012 – 2015. Consistent with a tightly binding zero lower bound constraint, skewness, kurtosis and β all hover around 3-4 times their average values during this period. In contrast, prior to

Figure 2.3: Time-series of lower bound indicators

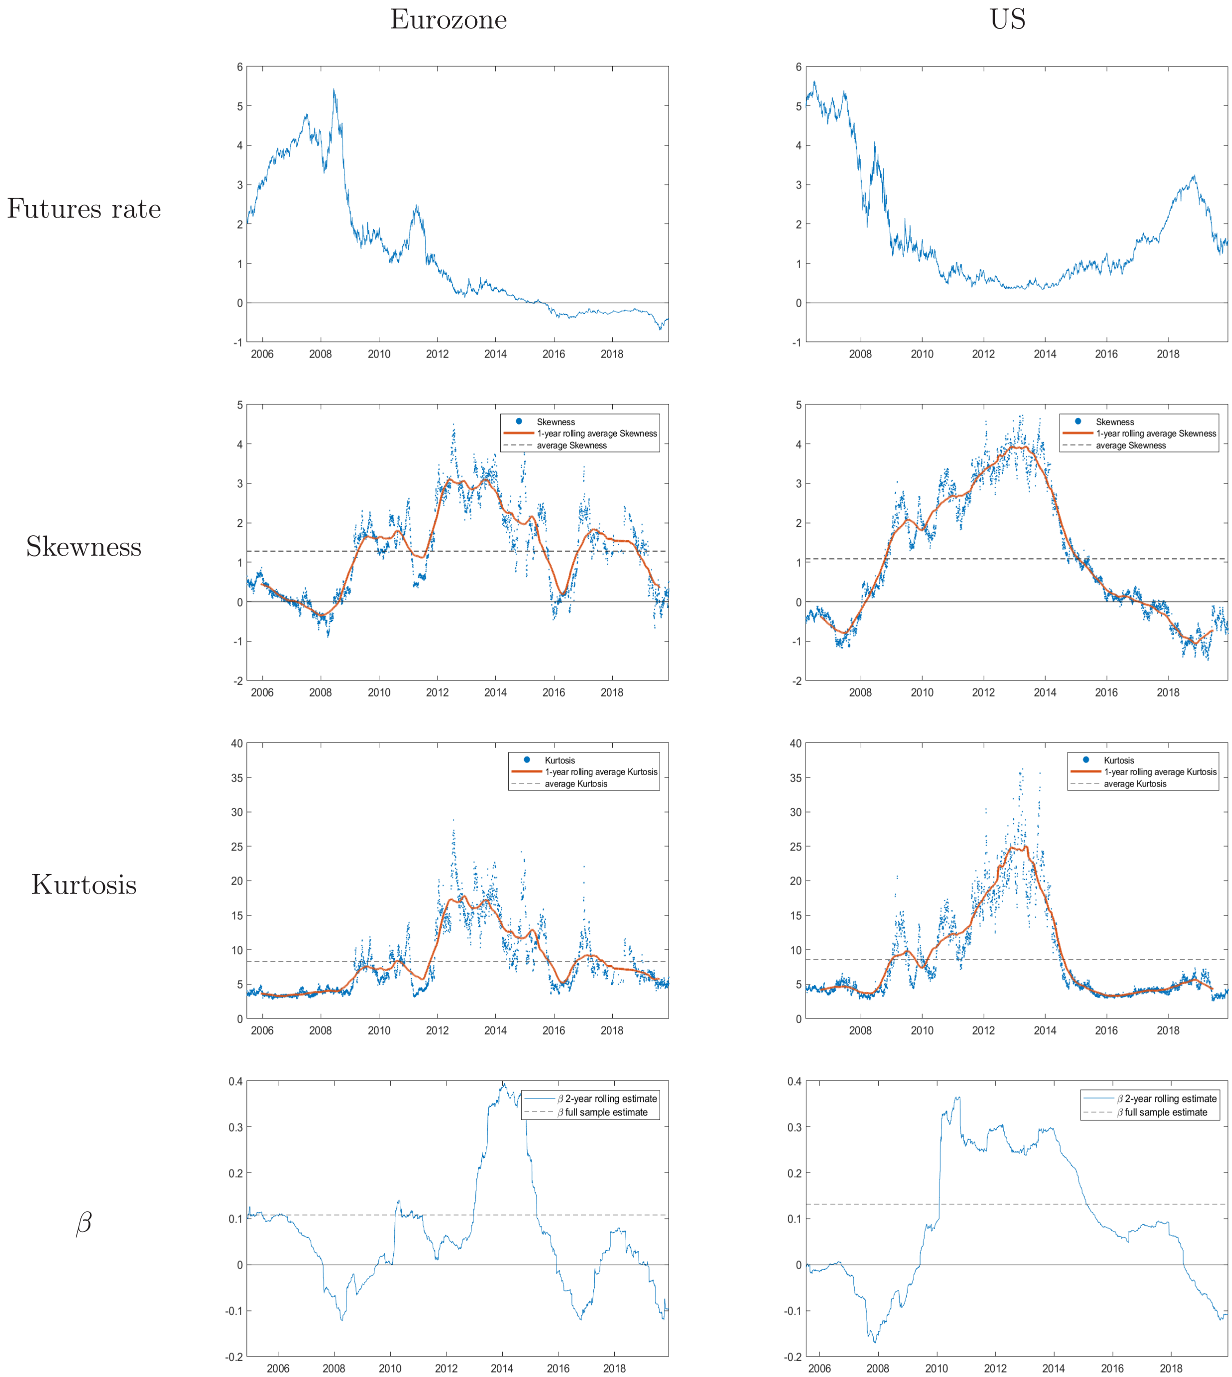

This figure shows time-series of skewness, kurtosis and β together with the level of the futures interest rate for both the US and the Eurozone market. The futures rate, skewness and kurtosis are each interpolated between adjacent maturities to a constant maturity of one year. Red lines show 252-day moving average values. The last row displays β from a rolling 2-year estimation of regression (2.5).

the respective zero-bound periods, skewness is close to zero or even negative in both markets, excess-kurtosis is still positive but much more muted and the level dependence of volatility fluctuates around zero. This combined evidence clearly speaks in favor of anticipation of a zero lower bound resulting in tightly constrained distributions for interest rates close to zero. However, all measures experience a sharp drop around 2014/2015 down to levels similar to those observed prior to the zero lower bound period. Strikingly, this drop coincides with numerous central banks subsequently lowering short-term rates into negative territory, hence breaching the zero lower bound (see Figure 2.4). This simultaneous decline in all indications of a constraining lower bound as interest rates continue to fall has revealing implications for markets lower bound beliefs: It is inconsistent with the idea that markets always anticipated a somewhat negative lower bound as this would result in more – not less – constrained distributions as interest rates become negative. It is also inconsistent with the idea that markets made only small subsequent readjustments to their zero lower-bound beliefs and repeatedly had to adapt convictions as they were again and again "surprised" by the experience of ever more negative interest rates. Instead, little indication for lower bound constrained interest rates after 2015 suggests that markets did not continue to anticipate a lower-bound close to observed interest rates. A more plausible explanation seems to be that market participants substantially revised their lower-bound beliefs as the previously held zero-lower bound conviction was contradicted.

In the US, the rapid decline of skewness, kurtosis and the level-volatility dependence *precedes* the lift-off of the Fed policy rate by more than a year. With US interest rates still about as close to zero as during the 2008 – 2015 period until the end of 2016, all lower-bound indicators fall to substantially below-average levels already by the end of 2014 when interest rates in several economies became negative. This means that the disappearance of lower-bound constrained distributions in the US was not merely a result of increasing interest rates. It seems that the breach of the zero lower bound in other economies also prompted US investors to materially revise their convictions. This interpretation is also supported by the developments towards the end of our sample: When US interest rates in 2019 returned to levels as low as in the end of 2008, skewness, kurtosis and β remained at about their lowest levels in our sample, with skewness and β being even negative.

2.3.3 Option-implied limits for lower bound beliefs

Expression (2.4) allows us to calculate a model-free, measure-agnostic *upper limit* for the market-implied lower bound. Figure 2.5 shows this option-implied *maximum* lower-bound $\overline{LB}$, which aligns with observed skewness and kurtosis at any given time. During the period 2008-2015, observed distributions are consistent with a lower-bound at or even slightly above zero. This provides model-free confirmation that the zero lower bound was likely incorporated

Figure 2.4: Negative monetary policy interest rates

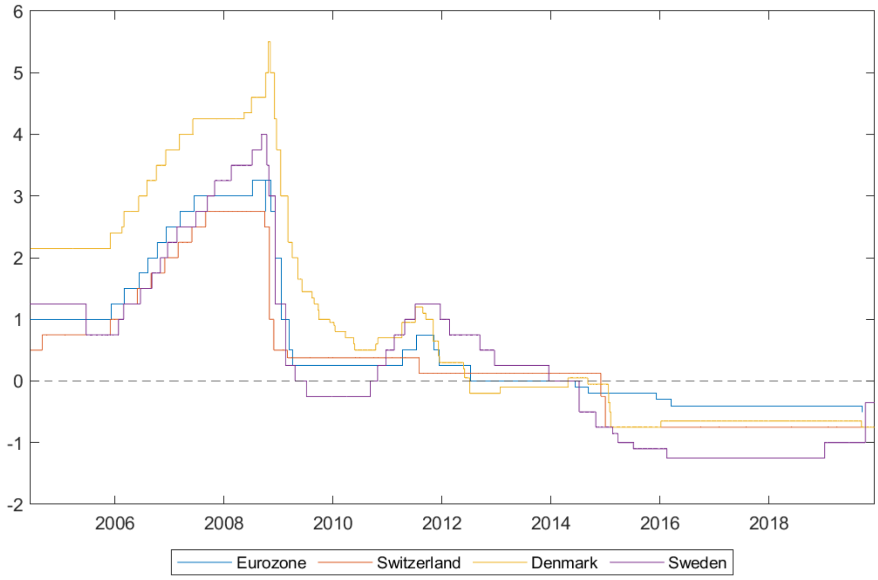

This figure shows policy rates for economies with negative interest rates. For the Eurozone, we show the ECB deposit facility rate, for Switzerland the midpoint of the SNB target range, for Denmark the Nationalbank certificates of deposit rate, and for Sweden the Riksbank deposit rate.

into market expectations at that time. Beginning in 2015, the maximum lower bound in the Eurozone drops into deeply negative territory, reaching less than -1% in early 2016. This implies that we can confidently exclude the possibility of a lower bound that is only slightly negative, as markets quickly adjusted their lower-bound beliefs to *at least* -1% following the breach of the zero lower bound. In the US, the maximum lower bound measure helps us to link changes in skewness and kurtosis to a corresponding change in lower-bound beliefs. While $\overline{LB}$ has been comfortably above 0 until 2015 in the US, it subsequently fell significantly, reaching negative levels in 2016. When interest rates in the US again fell after a short period of higher rates 2017/2018, $\overline{LB}$ displayed significant negative values of up to -0.3% at the end of 2019. This provides model-free evidence that investors no longer expect the zero lower bound to hold in the US, suggesting thath negative interest rates of *at least* -0.3% are now believed to be possible.

Another way to look at the data is to consider the minimal distance d of distribution means to the lower bound in terms of standard deviations. This implied minimal distance is plotted in Figure 2.6 for the potentially most constrained maturity at each point in time. This measure highlights the potential importance of the zero lower bound during the 2008-2015 period where

Figure 2.5: Option-implied maximum lower bound

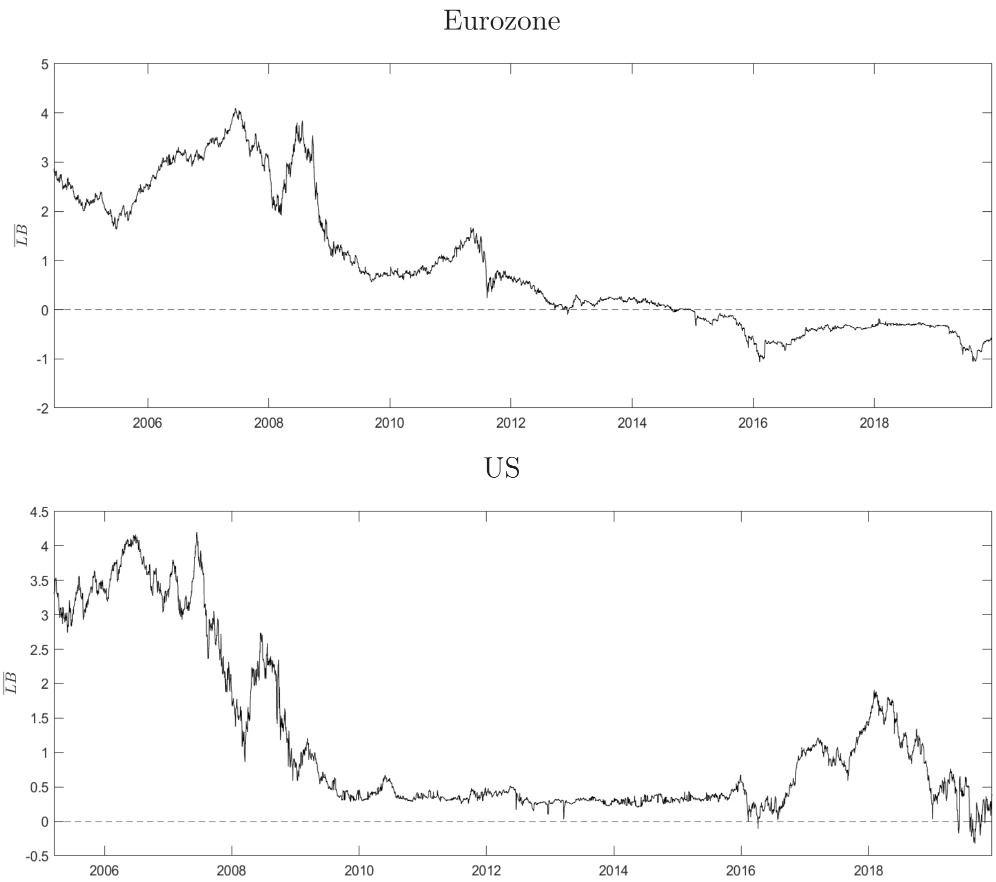

This figure shows the maximum possible market anticipated interest rate lower bound $\overline{LB_t}$ implied by conditional option skewness and kurtosis. At each point in time, we calculate for each maturity the highest value of a lower bound which would still be consistent with option-implied moments using expression 2.4. The lowest value for the maximum lower bound across all maturities on a given date gives us the sharpest upper limit for the lower bound on that day.

distributions are consistent with highly constrained interest rates of less than 0.2 standard deviations above the lower bound. After 2015, the importance of a lower bound quickly decreases as the minimal distance reverts to pre-2008 levels. Hence, distributions after 2015 imply a much less constraining lower bound with even the potentially most constrained maturities having at least 0.4 standard deviations room to move downward.

Figure 2.6: Minimal distance to lower bound

Eurozone

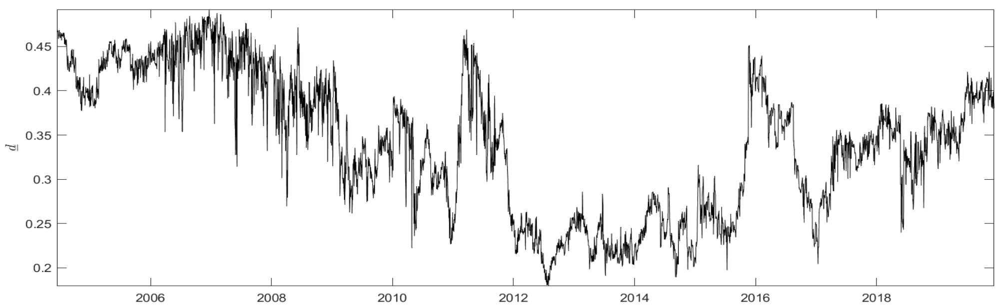

US

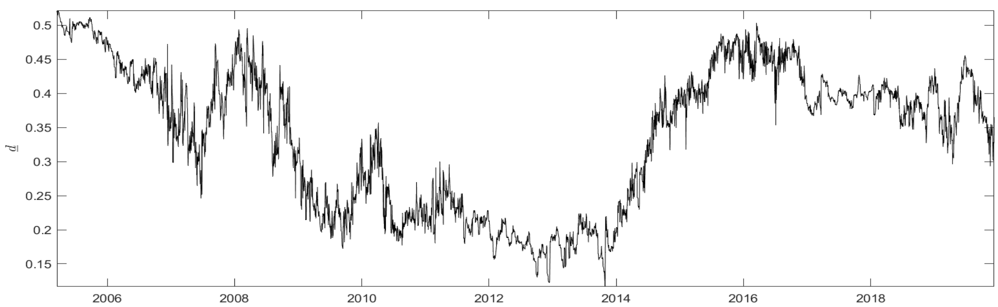

This figure shows the minimal possible distance $\underline{d}$ in terms of standard deviations between the futures rate and the lower bound implied by option skewness and kurtosis. At each point in time, we calculate for each maturity the smallest distance $\underline{d}$ which would still be consistent with option-implied moments. To do this, we invert relationships $\underline{\mathcal{S}}$ and $\underline{\mathcal{K}}$ to extract the minimal distance to the lower bound $\underline{d}$ from option-implied skewness and kurtosis. This directly gives us the minimal distance $\underline{d}$ consistent with skewness respectively kurtosis for any given maturity on a given date. On a given date, we plot the lowest value for the minimal distance $\underline{d}$ across all maturities, giving us the minimal possible distance $\underline{d}$ of the potentially most constrained maturity on that day.

2.3.4 Distribution characteristics and the level of interest rates

To further support our graphical evidence of a *change* in lower bound beliefs, we now look at how implied distribution characteristics depend on the level of interest rates. Table 2.1 and Table 2.2 show sample means of risk-neutral skewness and kurtosis conditional on the level of the futures rate across all maturities. They also report Newey-West t-statistics of the regressions $M_{t,T} = \alpha + \beta I_{\left(a \leq f_{t,T} < b\right)}$, testing the difference in means between the conditioning group and the remainder of the sample.

While unconditionally slightly positive, implied skewness substantially varies with the level of the interest rate. Conditional on interest rates being far above any potential lower bound, skewness in both markets is negative on average. This suggests that apart from the impact of a lower bound, market participants on average are willing to pay more for OTM options betting on falling

Table 2.1: Conditional skewness

	Total	<0%	0%-2%	2%-4%	>4%
Panel A: US					
Mean	0.94	-	1.65	0.14	-0.49
t-stat	-	-	9.29	-6.53	-9.78
SD	1.47	-	1.48	0.79	0.44
N	37646	0	21472	12414	3760
Panel B: EU					
Mean	1.22	0.99	2.04	0.51	-0.12
t-stat	-	-1.40	10.05	-6.53	-12.49
SD	1.09	0.82	0.86	0.52	0.39
N	25768	5665	11413	5781	2909

This table shows summary statistics of option-implied skewness for the US and the Eurozone across all available maturities. The first column displays unconditional results while the second to fifth column display results conditional on the futures rate being below 0%, in the interval 0%-2%, in the interval 2%-4% and above 4%, respectively. The second row within each panel shows t-statistics of the slope coefficient from the regression $Skew_{t,T} = \alpha + \beta I_{\left(a \leq f_{t,T} < b\right)}$. t-statistics are corrected for heteroscedasticity and serial correlation up to 252 lags (i.e. one year) using the approach of Newey and West (1987).

interest rates than for equivalent OTM options betting on rising interest rates. Intuitively, this negative implied interest-rate skewness matches with the experience that sharp short-rate cuts are much more likely than equivalent sharp increases.[3] However, with interest rates approaching zero, skewness increases and turns positive. Conditional on interest rates being in the interval 0%-2%, skewness is clearly positive and substantially higher than average in both markets. Also, the difference in conditional means is highly significant in both cases. This is in line with our earlier observation of the zero lower bound acting as a tight constraint during most of the near-zero interest rate period. However, average skewness for negative interest rates in the Eurozone reveals a discontinuity in this negative level-skewness relationship. When interest rates fall below 0% in the Eurozone, average skewness more than halves compared to its value for positive near-zero interest rates. This suggests that a static, level-dependent explanation based solely on the anticipation of a fixed lower bound is insufficient to explain observed empirical patterns in interest rate skewness. To better understand these changes, we need to consider that views on where the lower bound lies may shift over time, which means the relationship between interest rates and skewness could also change. We will consider this approach in the next section.

Level-conditional results for implied kurtosis show a very similar pattern. On average, and when interest rates are safely above a lower bound, kurtosis exceeds three. This indicates interest-rate

[3] Note that for this intuition of negative interest-rate skewness – namely negatively skewed short-rate increments – the effect should be expected to decrease with maturity. This is what we observe in Table 2.4 and 2.5

Table 2.2: Conditional kurtosis

	Total	<0%	0%-2%	2%-4%	>4%
Panel A: US					
Mean	7.61	-	10.00	4.49	4.27
t-stat	-	-	7.31	-6.78	-6.39
SD	6.61	-	7.90	1.12	1.16
N	37646	0	21472	12414	3760
Panel B: EU					
Mean	7.73	7.00	10.84	4.19	4.00
t-stat	-	-1.29	8.47	-8.87	-8.54
SD	4.53	2.65	4.72	0.97	0.88
N	25768	5665	11413	5781	2909

This table shows summary statistics of option-implied kurtosis for the US and the Eurozone across all available maturities. The first column displays unconditional results while the second to fifth column display results conditional on the futures rate being below 0%, in the interval 0%-2%, in the interval 2%-4% and above 4%, respectively. The second row within each panel shows t-statistics of the slope coefficient from the regression $Kurt_{t,T} = \alpha + \beta I_{\left(a \leq f_{t,T} < b\right)}$. t-statistics are corrected for heteroscedasticity and serial correlation up to 252 lags (i.e. one year) using the approach of Newey and West (1987).

distributions with heavier tails than the normal distribution independently from a lower-bound explanation. Again this finding is in line with the results of Trolle and Schwartz (2014). With falling interest rates, excess kurtosis becomes more pronounced. Conditional on interest rates being close to but above zero, kurtosis is significantly above average in both markets, consistent with (zero-)lower-bound-induced heavier right tails. However, as with skewness, average kurtosis for negative interest rates challenges the idea of a fixed lower bound: For negative interest rates, the sample mean of kurtosis is below average and substantially lower than for positive near-zero rates, indicating less heavy distribution tails despite lower interest rates. Hence our estimates for conditional kurtosis are also inconsistent with the notion of a negative lower bound being approached but point to a potential shift in lower-bound perceptions.

Table 2.3 presents our results for the level-conditional version of regression (2.5). Both unconditionally and when interest rates are far from a potential lower bound, volatility shows little relation with interest rate levels. The unconditional coefficients are significantly positive in both markets, but they are small in magnitude, and the R^2's are low. This is consistent with the large literature on unspanned stochastic interest-rate volatility which documents that the bulk of volatility variation in yields cannot be explained by variations in the term structure (see e.g. Collin-Dufresne and Goldstein (2002), Heidari and Wu (2003), Li and Zhao (2006, 2009)). However, when we condition on the futures rate being in the interval 0%-2%, the coefficients strongly increase, become highly significant, and R^2's rise substantially. This replicates the

Table 2.3: Conditional level-dependence of volatility

	Total	<0%	0%-2%	2%-4%	>4%
Panel A: US					
α	-0.001***	-	-0.001***	0.000	0.000
	(-3.124)	-	(-3.664)	(0.485)	(0.037)
β	0.132***	-	0.262***	0.088***	-0.129***
	(12.819)	-	(23.669)	(5.445)	(-6.253)
R^2	0.075	-	0.235	0.042	0.080
N	37580	0	21512	12327	3741
Panel B: EU					
α	0.000**	-0.001***	0.000	0.000	0.001
	(-2.468)	(-3.142)	(-1.640)	(-0.799)	(0.857)
β	0.109***	-0.050	0.282***	0.038**	0.008
	(8.738)	(-1.473)	(14.973)	(2.492)	(0.281)
R^2	0.033	0.005	0.154	0.006	0.000
N	25711	5770	11290	5764	2887

This table shows the results from regressing daily changes in volatility on changes in the futures rate. The first column displays unconditional results while the second to fifth column display results conditional on the futures rate being below 0%, in the interval 0%-2%, in the interval 2%-4% and above 4%, respectively. t-statistics of the coefficients are shown in parentheses and are corrected for heteroscedasticity and serial correlation up to two lags using the approach of Newey and West (1987). *, ** and *** denote significance at the 10%, 5%, and 1% levels, respectively.

finding of Filipović et al. (2017) and might be evidence for a zero-lower-bound-induced level-volatility dependence. In contrast, when we condition on negative interest rates in the Eurozone, the coefficients are small, statistically insignificant, and the R^2 values are near zero. In line with evidence for skewness and kurtosis, any potential lower-bound effect on level-volatility dependence (for positive near-zero rates) apparently did not carry into a negative interest-rate environment.

In summary, all three empirical measures indicate a substantial lower-bound impact on interest-rate distributions close to zero. However, the level-conditional distributions also reveal a strong decrease of lower-bound effects for negative interest rates. This result cannot be explained by the market's anticipation of a static lower bound. Instead, graphical investigation of the time-series of skewness, kurtosis and β suggests that the market's perception of the lower bound has shifted over time. In the next section we will formally investigate whether a shift in lower-bound views can explain the observed patterns in skewness, kurtosis, and level-volatility dependence.

2.4 Structural break

The previous analysis shows that the market's anticipation of a fixed lower bound is unable to fully explain observed implied distribution characteristics. In this section we formally analyze whether market participants might have revised their lower-bound beliefs during the sample period. To account for a potential shift of a lower bound over time, we extend the level-conditional analysis from the previous section by adopting a structural-break regression framework. Our goal is to analyze whether, and when, the observed lower-bound effects on interest rate distributions have changed over time. To this end, we regress skewness and kurtosis on a non-linear transformation of the futures rate, $\left(\frac{1}{1+f_{t,T}}\right)$. This transformation accounts for the explosive behavior of the derived lower-bound constraints, $\underline{\mathcal{S}}$ and $\underline{\mathcal{K}}$, which become increasingly relevant as interest rates approach the lower bound. We account for a potential shift in lower-bound beliefs with the following model specification:

$$M_{t,T} = \alpha + \beta_0 \left(\frac{1}{1+f_{t,T}}\right) I_{(t<T_b)} + \beta_1 \left(\frac{1}{1+f_{t,T}}\right) I_{(t \geq T_b)} + \epsilon_{t,T} \tag{2.6}$$

$M_{t,T}$ denotes the distribution moment (either skewness or kurtosis) in t for option maturity T. α, which we assume to be constant over the sample, is the mean of $M_{t,T}$ without the impact of a lower bound. β_0 and β_1 capture the impact of the lower bound on $M_{t,T}$ before and after the structural break, respectively. For a higher lower bound, the relationship between interest rates and $M_{t,T}$ is steeper, leading to a larger β on the convex function $\left(\frac{1}{1+f_{t,T}}\right)$. Conversely, when the lower bound is lower, this relationship flattens, resulting in a smaller estimate for β. The difference $\beta_1 - \beta_0$ quantifies the magnitude and direction of any shift in market beliefs about the lower bound. We remain agnostic about the timing of the structural break T_b, estimating it endogenously from the data. For each pooled regression (2.6) we estimate T_b as a common breakpoint across all maturities,[4] using the least-squares estimator of Bai and Perron (1998).[5]

Regarding the estimated timing of the occurred break, Figure 2.7 shows the structural-break regression R^2s for skewness and kurtosis in both markets as a function of T_b. The point estimates of T_b, which maximize the R^2 values, suggest break dates ranging from mid-2014 to September 2015. Strikingly, this period coincides with the fall of interest rates into deeply negative territory in many major economies for the first time in history (Figure 2.4). The European Central Bank (ECB) cut its deposit facility rate to -10 basis points in June 2014, followed by further rate cuts to -20 basis points in September 2014 and -30 basis points in December 2015. Denmarks Nationalbank, Swedens Riksbank, and the Swiss National Bank also all cut their policy rate to

[4] We obtain very similar break-date estimates when we estimate them individually for each maturity.

[5] The extension of this breakdate estimator to the panel data case relevant to our application was studied in Bai (2010) and Baltagi et al. (2016).

new below-zero lows during the same period.[6] Short-term market rates in the Eurozone such as 1-month Euribor started to drop below zero by the end of the first quarter 2015 and continued to fall to about -20 basis points by the end of the same year. Hence the impact of a lower bound on interest-rate distributions seems to have undergone a substantial structural change just when markets experienced the clear violation of the previously widely held zero-lower-bound assumption.

What's important to note is that while all the break dates fall within this narrow timeframe, they are unlikely to correspond to a single event. Instead of modeling a gradual change, we focused on detecting a clear structural break, which is a simple but effective way to pinpoint when market participants likely adjusted their expectations. That said, it's unlikely that markets revised their beliefs about the lower bound overnight. A sudden, one-day shift in such a deeply ingrained assumption is unrealistic. The concept of a lower bound is based on the idea that when interest rates fall below a certain level, investors would rather hold cash than lend at negative rates. However, the cost of holding cash varies for different participants and markets, meaning this shift likely occurred gradually as interest rates approached the lower bound. It likely took time for the market to assess whether the expected lower-bound mechanism was actually in play. Our approach simplifies this reality, but it effectively captures the period during which market participants gradually adapted their views on the lower bound and reflected this shift in option prices. In short, while we identify a clear break, we acknowledge that market beliefs evolved over time as participants learned that negative rates were possible, and the assumed 'zero-lower-bound' was no longer a hard limit.

Tables 2.4, 2.5, 2.6, and 2.7 display results of the structural-break regressions, evaluated separately for skewness and kurtosis in each market. The regressions are run for each maturity individually and the row "Total" shows the results from the regression pooling all maturities together. Regressions across all maturities are uniformly evaluated at the same break-date T_b as estimated previously on the full sample. For each regression, the tables report estimates for the coefficients α, β_0 and β_1, their Newey-West standard errors in parentheses, the regression-R^2, and the number of observations. Furthermore, column F shows the SupF-test statistic evaluating the significance of the estimated break for each individual maturity.
The estimated break for both skewness and kurtosis is (highly) significant for every single maturity in both markets. This indicates strong evidence for a structural break in distribution moments of low interest rates. We conclude that the impact of a lower bound on the distribution of interest rates has undergone significant change throughout the sample, implying a corresponding change in the market-implied location of the lower bound. Hence, to understand the strength

[6] Sweden and Denmark already experienced short periods of somewhat negative policy rates in 2009 and 2012, respectively.

Figure 2.7: Break-date estimation

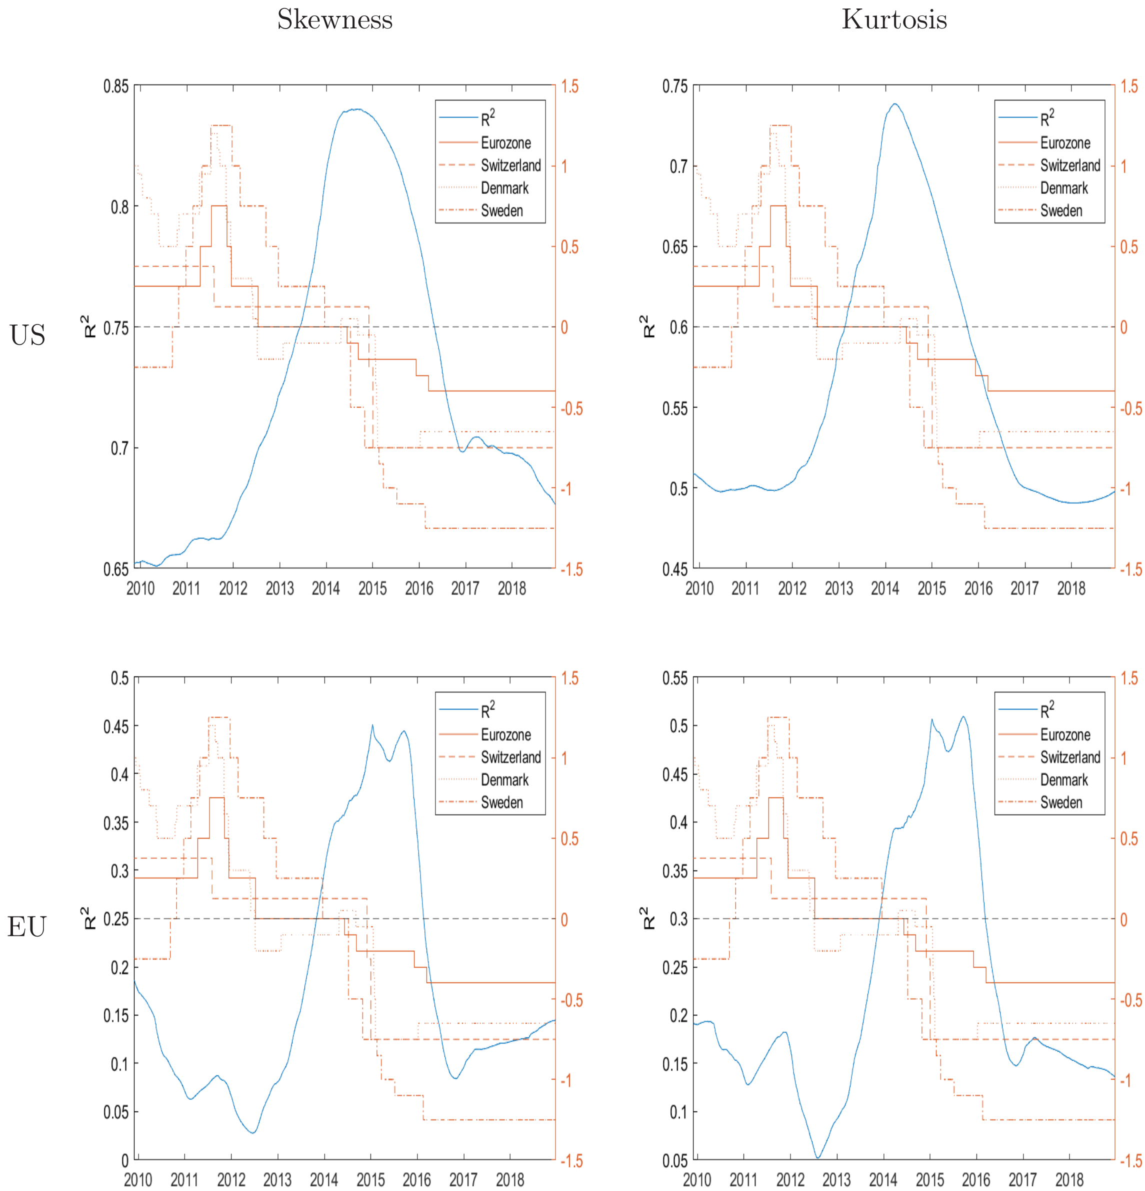

The figure plots R^2s of the structural-break-regressions (2.6) depending on the date of the structural break T_b. Our least-square estimates for the break-date parameter T_b maximize R^2. Additionally, policy rates for economies with negative interest rates are plotted.

of a lower-bound impact, and particularly the fixed lower-bound inconsistent conditional distributions of negative interest rates in the Eurozone, it will be crucial to account for this apparently occurred shift. Notably, the structural break tests are not only significant for the Eurozone data, where conditional distributions of negative rates already pointed to such a break, but also for the US data. This documents a completely novel finding, since a change in the location of the

lower bound in the US has not been considered yet in the literature.
The direction of change in model estimates indicates a substantial *downward revision* of lower bound beliefs in both markets. In the Eurozone, post-break estimates $\widehat{\beta_1}$ are only about a quarter of the pre-break estimates $\widehat{\beta_0}$, both for skewness and kurtosis. In the US, the break is less pronounced but still substantial, with post-break estimates $\widehat{\beta_1}$ of around half of the pre-break estimates $\widehat{\beta_0}$. This means that relative to the pre-break period, skewness and kurtosis are much less sensitive to the low levels of interest rates at which they exhibit much more moderate levels in the post-break period.
Consistent with a tightly binding lower bound, $\widehat{\beta_0}$ estimates for skewness in the pre-break period are all positive and highly significant for every single maturity in both markets. Particularly for the US, the impact of low interest rates on skewness is substantial. In the Eurozone, the absolute value of parameter estimates is somewhat more muted. While the post-break parameter estimates $\widehat{\beta_1}$ are in all cases much smaller, they are still positively significant in most cases. This suggests a much weaker but still visible impact from a lower bound constraint on interest rate distributions in the post-break period.
Results for kurtosis confirm these findings. $\widehat{\beta_0}$ estimates are all positive and highly significant. Estimates are also quantitatively large with the effect again more pronounced in US data. Hence distributions of low interest rates in the pre-break period have much heavier tails than normally. For the post-break period, estimates for $\widehat{\beta_1}$ are again much smaller for all maturities and in many cases insignificant. This confirms the strongly reduced importance of an impact of a lower bound on interest-rate distributions after the occurred break.
In summary, these results indicate tightly constrained interest-rate distributions leading to substantially above average levels of skewness and kurtosis in the pre-break period. In the post-break period, the impact on distributions is much weaker and only a fraction of the pre-break impact.

Figure 2.8 visualizes the differences in pre- and post-break distributions conditional on the level of interest rates. We plot skewness and kurtosis against the level of the futures rate, using different colors for pre- and post-break observations. We also show the fitted regression lines from our structural break regressions. The graphs clearly show a strongly convex relationship between skewness (and kurtosis) and the futures rate in the pre-break-period, consistent with the theoretical predictions of a lower bound close to zero. Post-break observations exhibit a much flatter relationship between the futures rate and conditional moments. Compared to the pre-break period, skewness and kurtosis in the post-break period are much smaller at similar values of the futures rate. Additionally, the convex region of the function that describes the explosive behavior of skewness and kurtosis at the lower bound has shifted to the left as market participants downward revised their lower bound beliefs.

Finally, we test whether a downward revision of the lower bound as indicated by skewness and kurtosis data is also reflected by a corresponding change in conditional level-volatility dependence.

Table 2.4: Structural break regression: US skewness

Maturity (in months)	α	β_0	β_1	F	R^2	Observations
0-3	-1.66***	7.10***	3.20***	156.49***	0.803	897
	(-4.58)	(11.19)	(3.93)			
3-6	-2.13***	8.16***	4.13***	128.28***	0.905	2672
	(-6.89)	(16.92)	(6.70)			
6-9	-1.91***	7.76***	4.03***	93.43***	0.908	3298
	(-8.51)	(24.83)	(7.92)			
9-12	-1.81***	7.74***	4.17***	57.25***	0.915	3516
	(-8.36)	(20.45)	(7.50)			
12-15	-1.79***	8.05***	4.42***	52.10***	0.908	3638
	(-8.44)	(20.28)	(7.40)			
15-18	-1.73***	8.31***	4.56***	58.02***	0.913	3702
	(-9.61)	(22.14)	(8.10)			
18-21	-1.71***	8.84***	4.79***	59.59***	0.916	3768
	(-9.03)	(22.05)	(7.63)			
21-24	-1.64***	8.92***	4.84***	55.70***	0.915	3686
	(-8.23)	(20.19)	(7.12)			
24-27	-1.61***	8.45***	5.03***	31.64***	0.904	2151
	(-3.65)	(9.87)	(3.94)			
27-30	-1.10**	7.16***	3.77***	38.11***	0.875	2093
	(-2.38)	(7.93)	(3.05)			
30-33	-1.00**	7.25***	3.63***	47.03***	0.850	2042
	(-2.31)	(7.81)	(3.06)			
33-36	-0.38	5.27***	2.04	33.67***	0.726	2011
	(-0.79)	(4.95)	(1.56)			
Total	-1.35	7.37	3.84		0.840	37646

This table shows the results from the structural-break regressions for skewness in the US. Regressions are run for each maturity individually and the row "Total" reports regression results from pooling all maturities together. The break-date parameter T_b is set to its least-square estimate uniformly for all regressions. The OLS estimator of T_b is given by $\hat{T}_b = argmin_{1<T_b<T-1} SSR(T_b)$, where $SSR(T_b) = \sum_{t=1}^{T}\sum_{T=1}^{N}\left(M_{t,T} - \hat{\alpha} - \widehat{\beta_0}\left(\frac{1}{1+f_{t,T}}\right) I_{(t<T_b)} + \widehat{\beta_1}\left(\frac{1}{1+f_{t,T}}\right) I_{(t\geq T_b)}\right)^2$. To assess significance of the estimated breaks we consider the SupF-test proposed by Andrews (1993), Andrews and Ploberger (1994) and Bai and Perron (1998). The table displays the regression coefficients with t-statistics in parentheses, R^2s and the number of observations. Column F reports the statistic of the F-test run under the null of no structural change against the alternative hypothesis of a break at T_b. Error-covariances for computation of the t-statistics and F-tests are estimated accounting for heteroscedasticity and serial correlation up to 252 lags (i.e. one year) using the approach of Newey and West (1987). *, ** and *** denote significance at the 10%, 5%, and 1% levels, respectively. Critical values used to assess significance of the SupF-Test are taken from Bai and Perron (1998).

To this end, we extend the level-volatility regression (2.5) to account for a relation with interest rates subject to a break:

$$\Delta Vol_{t,T} = \alpha + \beta \Delta f_{t,T} + \gamma_0 \Delta f_{t,T}\left(\frac{1}{1+f_{t,T}}\right) I_{(t<T_b)} + \gamma_1 \Delta f_{t,T}\left(\frac{1}{1+f_{t,T}}\right) I_{(t\geq T_b)} + \epsilon_{t,T} \quad (2.7)$$

Table 2.5: Structural break regression: EU skewness

Maturity (in months)	α	β_0	β_1	F	R^2	Observations
0-3	-0.54*	2.98***	1.27***	13.11**	0.325	1165
	(-1.80)	(4.75)	(5.55)			
3-6	-0.60	3.55***	0.76**	19.50***	0.461	2665
	(-1.55)	(5.07)	(2.33)			
6-9	-0.40	3.56***	0.78**	16.06***	0.445	3172
	(-0.92)	(4.10)	(2.33)			
9-12	-0.33	3.89***	0.86**	20.79***	0.484	3511
	(-0.76)	(4.52)	(2.05)			
12-15	-0.12	3.68***	0.79**	21.32***	0.486	3814
	(-0.29)	(4.34)	(1.98)			
15-18	-0.10	3.83***	0.80**	30.42***	0.529	3896
	(-0.26)	(4.99)	(2.02)			
18-21	-0.13	4.06***	0.97***	42.09***	0.576	3898
	(-0.38)	(5.97)	(2.60)			
21-24	-0.19	4.14***	1.25***	39.61***	0.659	3580
	(-0.78)	(7.21)	(4.39)			
Total	-0.19	3.58	0.84		0.450	25768

This table shows the results from the structural-break regressions for skewness in the Eurozone. Regressions are run for each maturity individually and the row "Total" reports regression results from pooling all maturities together. The break-date parameter T_b is set to its least-square estimate uniformly for all regressions. The OLS estimator of T_b is given by $\hat{T}_b = argmin_{1<T_b<T-1} SSR(T_b)$, where $SSR(T_b) = \sum_{t=1}^{T}\sum_{T=1}^{N}\left(M_{t,T} - \hat{\alpha} - \widehat{\beta_0}\left(\frac{1}{1+f_{t,T}}\right)I_{(t<T_b)} + \widehat{\beta_1}\left(\frac{1}{1+f_{t,T}}\right)I_{(t\geq T_b)}\right)^2$. To assess significance of the estimated breaks we consider the SupF-test proposed by Andrews (1993), Andrews and Ploberger (1994) and Bai and Perron (1998). The table displays the regression coefficients with t-statistics in parentheses, R^2s and the number of observations. Column F reports the statistic of the F-test run under the null of no structural change against the alternative hypothesis of a break at T_b. Error-covariances for computation of the t-statistics and F-tests are estimated accounting for heteroscedasticity and serial correlation up to 252 lags (i.e. one year) using the approach of Newey and West (1987). *, ** and *** denote significance at the 10%, 5%, and 1% levels, respectively. Critical values used to assess significance of the SupF-Test are taken from Bai and Perron (1998).

Here, γ_0 and γ_1 estimate the pre- and post-break impact of a lower bound on level-volatility dependence and β estimates the relationship without a lower-bound impact. Figure 2.9 shows that break-date estimates for this regression fall into the range of breaks identified for skewness and kurtosis in both markets. Column F in Tables 2.8 and 2.9 indicate that these breaks are (highly) significant for all except one maturity in the US-market. Hence the identified break in conditional skewness and kurtosis coincides with a significant change in the conditional level-volatility relation. $\widehat{\beta}$'s are small in magnitude and for many maturities not significantly different from zero, indicating little level-dependence of volatility apart from the impact of a lower bound. In the pre-break period, the relation of level-volatility dependence on the level of interest rates

Table 2.6: Structural break regression: US kurtosis

Maturity (in months)	α	β_0	β_1	F	R^2	Observations
0-3	0.94	22.42***	7.58***	70.34***	0.666	897
	(0.80)	(8.91)	(2.71)			
3-6	2.24	22.02***	4.12	174.71***	0.771	2672
	(1.61)	(9.98)	(1.50)			
6-9	-0.29	26.73***	10.49***	71.60***	0.715	3298
	(-0.18)	(6.84)	(2.99)			
9-12	-1.31	28.52***	12.93***	43.11***	0.760	3516
	(-0.87)	(7.14)	(3.98)			
12-15	-2.29	32.98***	15.76***	40.13***	0.716	3638
	(-1.54)	(7.47)	(4.65)			
15-18	-3.22**	36.15***	18.26***	43.04***	0.794	3702
	(-2.29)	(9.03)	(5.36)			
18-21	-4.85***	43.59***	22.99***	50.66***	0.827	3768
	(-3.32)	(10.60)	(5.94)			
21-24	-4.63***	42.94***	23.14***	39.91***	0.850	3686
	(-3.15)	(8.34)	(5.49)			
24-27	-1.66	34.67***	15.64	13.37**	0.744	2151
	(-0.34)	(3.33)	(1.17)			
27-30	1.36	23.29***	7.43	22.70***	0.791	2093
	(0.46)	(3.54)	(0.86)			
30-33	0.51	25.59***	10.03	20.89***	0.740	2042
	(0.16)	(3.45)	(1.03)			
33-36	3.11***	13.73***	2.71	54.31***	0.613	2011
	(2.76)	(5.20)	(0.73)			
Total	-1.01	30.06	13.44		0.740	37646

This table shows the results from the structural-break regressions for kurtosis in the US. Regressions are run for each maturity individually and the row "Total" reports regression results from pooling all maturities together. The break-date parameter T_b is set to its least-square estimate uniformly for all regressions. The OLS estimator of T_b is given by $\hat{T_b} = argmin_{1<T_b<T-1} SSR(T_b)$, where $SSR(T_b) = \sum_{t=1}^{T}\sum_{T=1}^{N}\left(M_{t,T} - \hat{\alpha} - \widehat{\beta_0}\left(\frac{1}{1+f_{t,T}}\right) I_{(t<T_b)} + \widehat{\beta_1}\left(\frac{1}{1+f_{t,T}}\right) I_{(t\geq T_b)}\right)^2$. To assess significance of the estimated breaks we consider the SupF-test proposed by Andrews (1993), Andrews and Ploberger (1994) and Bai and Perron (1998). The table displays the regression coefficients with t-statistics in parentheses, R^2s and the number of observations. Column F reports the statistic of the F-test run under the null of no structural change against the alternative hypothesis of a break at T_b. Error-covariances for computation of the t-statistics and F-tests are estimated accounting for heteroscedasticity and serial correlation up to 252 lags (i.e. one year) using the approach of Newey and West (1987). *, ** and *** denote significance at the 10%, 5%, and 1% levels, respectively. Critical values used to assess significance of the SupF-Test are taken from Bai and Perron (1998).

is highly significant and positive in both markets and for all maturities ($\widehat{\gamma_0}$). This confirms the finding of a tightly binding lower bound in the pre-break period. For the post-break period,

Table 2.7: Structural break regression: EU kurtosis

Maturity (in months)	α	β_0	β_1	F	R^2	Observations
0-3	3.77***	7.96***	1.31***	285.22***	0.272	1165
	(15.35)	(15.34)	(8.38)			
3-6	1.94***	12.28***	2.48***	47.31***	0.468	2665
	(2.84)	(7.05)	(5.27)			
6-9	1.71*	13.55***	3.03***	17.46***	0.459	3172
	(1.65)	(4.51)	(3.52)			
9-12	1.27	15.72***	3.57***	17.31***	0.526	3511
	(1.05)	(4.46)	(3.04)			
12-15	1.84	14.43***	3.22***	14.34***	0.531	3814
	(1.48)	(3.92)	(2.89)			
15-18	1.78	14.44***	3.25***	15.23***	0.55	3896
	(1.48)	(4.03)	(2.86)			
18-21	1.59	15.02***	3.72***	18.06***	0.581	3898
	(1.44)	(4.42)	(3.24)			
21-24	1.49	13.68***	4.37***	15.13***	0.583	3580
	(1.55)	(4.23)	(3.89)			
Total	1.78	13.85	3.30		0.510	25768

This table shows the results from the structural-break regressions for kurtosis in the Eurozone. Regressions are run for each maturity individually and the row "Total" reports regression results from pooling all maturities together. The break-date parameter T_b is set to its least-square estimate uniformly for all regressions. The OLS estimator of T_b is given by $\hat{T}_b = argmin_{1<T_b<T-1} SSR(T_b)$, where $SSR(T_b) = \sum_{t=1}^{T}\sum_{T=1}^{N}\left(M_{t,T} - \hat{\alpha} - \widehat{\beta_0}\left(\frac{1}{1+f_{t,T}}\right) I_{(t<T_b)} + \widehat{\beta_1}\left(\frac{1}{1+f_{t,T}}\right) I_{(t\geq T_b)}\right)^2$. To assess significance of the estimated breaks we consider the SupF-test proposed by Andrews (1993), Andrews and Ploberger (1994) and Bai and Perron (1998). The table displays the regression coefficients with t-statistics in parentheses, R^2s and the number of observations. Column F reports the statistic of the F-test run under the null of no structural change against the alternative hypothesis of a break at T_b. Error-covariances for computation of the t-statistics and F-tests are estimated accounting for heteroscedasticity and serial correlation up to 252 lags (i.e. one year) using the approach of Newey and West (1987). *, ** and *** denote significance at the 10%, 5%, and 1% levels, respectively. Critical values used to assess significance of the SupF-Test are taken from Bai and Perron (1998).

$\widehat{\gamma_1}$ estimates are substantially smaller on average as well as for all individual maturities, and in many cases insignificant. Post-break estimates $\widehat{\gamma_1}$ in the Eurozone are only about an eighth of the pre-break estimates $\widehat{\gamma_0}$. As for skewness and kurtosis, the break is less pronounced in the US. There, we see a similar reduction in parameter estimates to around a half. The structural change in conditional level-volatility is therefore also consistent with a strong downward revision of the lower bound around the introduction of negative interest rates.

Figure 2.8: Skewness and kurtosis before and after the break

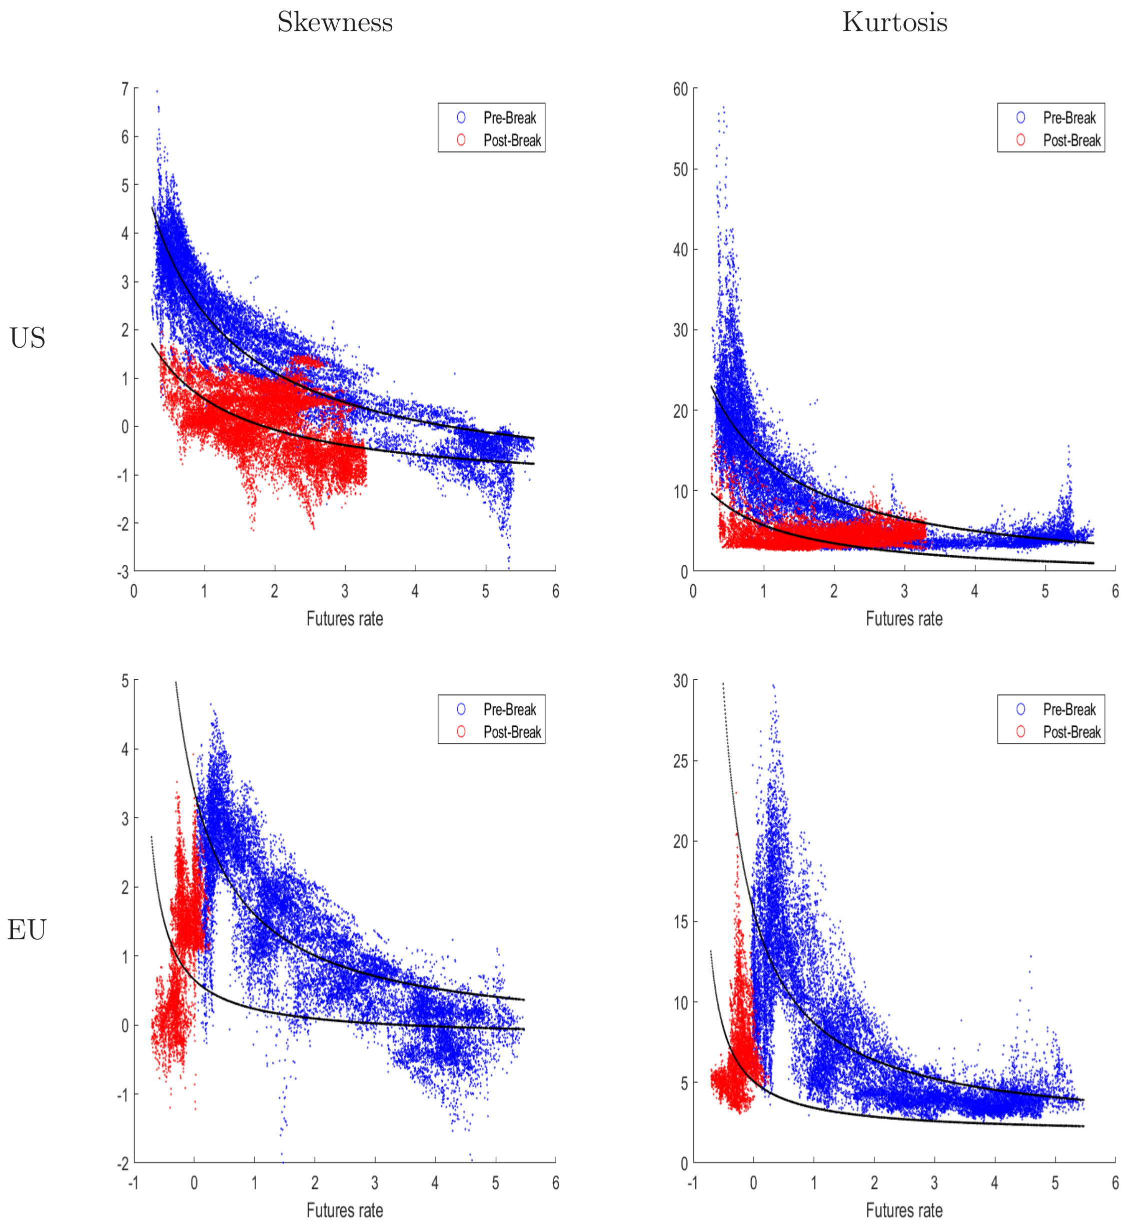

This figure plots skewness and kurtosis against the corresponding level of the futures rate. We use different colors for pre- and post-break observations, i.e. observations before and after the respective estimated break date T_b. Additionally, we plot the fitted regression lines from the individual structural-break regressions.

Figure 2.9: Break-date estimation: level-volatility dependence

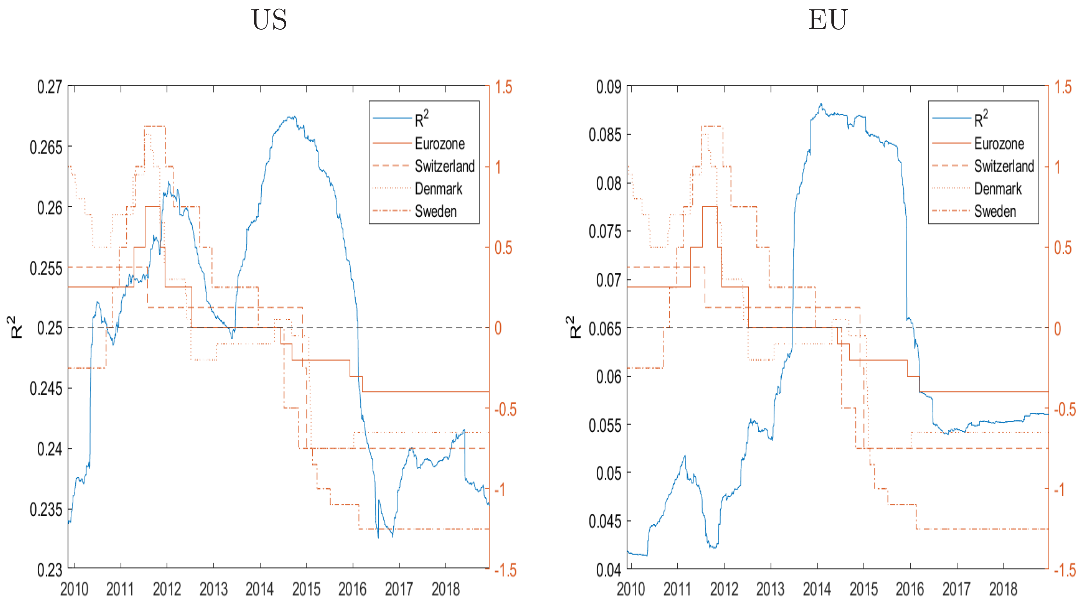

The figure plots R^2s of the structural-break regressions (2.7) depending on the date of the structural break T_b. Our least-square-estimates for the break-date parameter T_b maximize R^2. Additionally, policy rates for economies with negative interest rates are plotted.

Table 2.8: Structural break regression: US level-dependence of volatility

Maturity (in months)	α	β	γ_0	γ_1	F	R^2	Observations
0-3	0.0029	-0.54	3.96***	0.31	61.31***	0.251	897
	(1.03)	(-1.63)	(4.64)	(0.36)			
3-6	0.0006	-0.93***	3.78***	1.60***	277.13***	0.403	2672
	(0.84)	(-9.28)	(17.89)	(6.92)			
6-9	0.0002	-0.67***	2.82***	1.24***	290.08***	0.430	3298
	(0.38)	(-10.59)	(19.62)	(8.46)			
9-12	-0.0001	-0.50***	2.20***	0.99***	195.03***	0.413	3516
	(-0.21)	(-9.89)	(17.02)	(8.11)			
12-15	0.0000	-0.41***	1.86***	0.87***	157.35***	0.385	3638
	(-0.10)	(-9.98)	(17.77)	(8.58)			
15-18	-0.0002	-0.35***	1.59***	0.82***	217.55***	0.352	3702
	(-0.61)	(-9.54)	(19.55)	(8.75)			
18-21	-0.0001	-0.27***	1.41***	0.69***	201.12***	0.349	3768
	(-0.48)	(-8.22)	(19.44)	(7.72)			
21-24	0.0000	-0.17***	1.17***	0.49***	208.72***	0.376	3668
	(-0.11)	(-5.30)	(16.71)	(5.36)			
24-27	-0.0005**	-0.05	0.80***	0.17	192.00***	0.417	2141
	(-2.47)	(-0.97)	(6.57)	(1.22)			
27-30	-0.0005**	0.00	0.68***	0.04	182.72***	0.451	2092
	(-2.47)	(0.06)	(5.29)	(0.28)			
30-33	-0.0004*	-0.19*	1.10***	0.78**	6.70	0.520	2041
	(-1.74)	(-1.74)	(4.25)	(2.11)			
33-36	0.0000	0.11**	0.31**	-0.25*	135.52***	0.324	2002
	(-0.09)	(2.35)	(2.26)	(-1.83)			
Total	-0.0001	-0.37	1.84	1.00		0.270	37580

This table shows the results from the structural-break regressions for level-volatility dependence in the US. Regressions are run for each maturity individually and the row "Total" reports regression-results from pooling all maturities together. The break-date parameter T_b is set to its least-square estimate uniformly for all regressions. The OLS estimator of T_b is given by $\hat{T}_b = argmin_{1<T_b<T-1} SSR(T_b)$, where $SSR(T_b) = \sum_{t=1}^{T}\sum_{T=1}^{N}\left(\Delta Vol_{t,T} - \hat{\alpha} - \widehat{\beta}\Delta f_{t,T} - \widehat{\gamma_0}\Delta f_{t,T}\left(\frac{1}{1+f_{t,T}}\right) I_{(t<T_b)} - \widehat{\gamma_1}\Delta f_{t,T}\left(\frac{1}{1+f_{t,T}}\right) I_{(t\geq T_b)}\right)^2$. To assess significance of the estimated breaks we consider the SupF-test proposed by Andrews (1993), Andrews and Ploberger (1994) and Bai and Perron (1998). The table displays the regression coefficients with t-statistics in parentheses, R^2s and the number of observations. Column F reports the statistic of the F-test run under the null of no structural change against the alternative hypothesis of a break at T_b. Error-covariances for computation of the t-statistics and F-tests are estimated accounting for heteroscedasticity and serial correlation up to two lags using the approach of Newey and West (1987). *, ** and *** denote significance at the 10%, 5%, and 1% levels, respectively. Critical values used to assess significance of the SupF-Test are taken from Bai and Perron (1998).

Table 2.9: Structural break regression: EU level-dependence of volatility

Maturity (in months)	α	β	γ_0	γ_1	F	R^2	Observations
0-3	0.0007	-1.10***	3.54***	0.35*	28.74***	0.100	1165
	(0.38)	(-4.71)	(5.01)	(1.90)			
3-6	-0.0004	-0.58***	2.00***	0.41***	52.23***	0.125	2665
	(-0.62)	(-5.90)	(7.59)	(3.75)			
6-9	-0.0004	-0.34***	1.38***	0.16**	81.01***	0.131	3172
	(-0.96)	(-5.39)	(8.42)	(2.56)			
9-12	-0.0006*	-0.20***	0.96***	0.09*	62.83***	0.120	3512
	(-1.76)	(-3.54)	(7.10)	(1.75)			
12-15	-0.0004	-0.11**	0.76***	0.04	53.56***	0.130	3815
	(-1.28)	(-2.31)	(6.25)	(0.88)			
15-18	-0.0004	-0.04	0.56***	0.02	28.07***	0.135	3897
	(-1.60)	(-0.79)	(4.33)	(0.47)			
18-21	-0.0003	-0.03	0.57***	0.04	33.21***	0.164	3898
	(-1.05)	(-0.62)	(4.61)	(0.80)			
21-24	0.0000	0.00	0.44***	0.07	16.34***	0.131	3537
	(0.09)	(0.00)	(3.52)	(1.17)			
Total	-0.0003	-0.19	1.00	0.12		0.090	25711

This table shows the results from the structural-break regressions for level-volatility dependence in the Eurozone. Regressions are run for each maturity individually and the row “Total” reports regression-results from pooling all maturities together. The break-date parameter T_b is set to its least-square estimate uniformly for all regressions. The OLS estimator of T_b is given by $\hat{T}_b = argmin_{1<T_b\,<\,T-1} SSR\left(T_b\right)$, where $SSR\left(T_b\right) = \sum_{t=1}^{T}\sum_{T=1}^{N}\left(\Delta Vol_{t,T} - \hat{\alpha} - \widehat{\beta}\Delta f_{t,T} - \widehat{\gamma_0}\Delta f_{t,T}\left(\frac{1}{1+f_{t,T}}\right)I_{(t<T_b)} - \widehat{\gamma_1}\Delta f_{t,T}\left(\frac{1}{1+f_{t,T}}\right)I_{(t\geq T_b)}\right)^2$. To assess significance of the estimated breaks we consider the SupF-test proposed by Andrews (1993), Andrews and Ploberger (1994) and Bai and Perron (1998). The table displays the regression coefficients with t-statistics in parentheses, R^2s and the number of observations. Column F reports the statistic of the F-test run under the null of no structural change against the alternative hypothesis of a break at T_b. Error-covariances for computation of the t-statistics and F-tests are estimated accounting for heteroscedasticity and serial correlation up to two lags using the approach of Newey and West (1987). *, ** and *** denote significance at the 10%, 5%, and 1% levels, respectively. Critical values used to assess significance of the SupF-Test are taken from Bai and Perron (1998).

2.5 Risk premia

Our analysis up to here relies on option-implied distributions. Their forward-looking nature allows us to most directly capture markets' *ex ante* lower bound beliefs and provides us with reliable high frequency estimates. However, estimated distribution characteristics pertain to risk-neutral distributions. Therefore, one might be concerned as to what degree our conclusions could be driven by time variation in risk premia instead of physical distribution characteristics. Even though the physical ($\mathbb{P}$) distribution differs from the associated risk-neutral ($\mathbb{Q}$) distribution by risk premia, the two measures also have some common features. By definition, the two measures are equivalent and thus they coincide in which events have measure zero. Therefore, any lower bound of the physical interest-rate distribution (measure $\mathbb{P}$) is concomitantly the lower bound of the associated risk-neutral interest-rate distribution (measure $\mathbb{Q}$) and vice versa. Consequently, the maximum lower bound we derive from option-implied moments (see Figure 2.5) is generally valid, i.e., it applies to any measure, be it $\mathbb{Q}$ or $\mathbb{P}$. If we were able to determine the lower bound explicitly, we could safely rely on option-implied moments to make general statements about the lower bound. However, several of our statements rely on a relationship between the distance to the lower bound and the minimum skewness and kurtosis (see Figure 2.1). While this relationship also holds both under $\mathbb{P}$ and $\mathbb{Q}$, the "distance to the lower bound" d depends on the respective measure. Strictly speaking, option-implied moments thus only allow us to exclude a "tightly binding lower bound" as measured in terms of risk-neutral interest rate expectations and volatilities. Additionally, one might be worried that the elevated levels of implied skewness and kurtosis of near-zero interest rates in the former part of our sample might stem from risk premia instead of being induced by a zero lower bound.

To address these concerns, we additionally estimate the conditional moments of the physical distribution. We follow the literature on higher moments in stock returns and estimate realized moments using information from past realizations. Realized moments provide an unbiased measure for conditional moments of the objective, physical distribution. Their drawback is that they are backward-looking, very noisy and only available at low frequencies. Also, they might fail to adequately capture the *subjective* expectations of market participants if these differ from objective expectations. However, their great advantage is that they don't suffer from contamination with risk premia. This allows us to gauge the reliance of our main results on risk premia distortions.

Neuberger (2012) and Bae and Lee (2021) construct unbiased measures for realized skewness and kurtosis. We follow their approach to calculate realized skewness and kurtosis at a monthly frequency. Neuberger (2012) shows that skewness at longer horizons originates from two sources: skewness of high frequency price changes and the covariation between high frequency price changes and shocks to (implied) variance. Consequently, our estimation of realized skewness relies on daily changes in the futures rate as well as daily changes in implied variance (as estimated in Equation A.13) to account for both sources of skewness of the terminal interest-rate distribution. Similarly, Bae and Lee (2021) show that an unbiased estimate of the kurtosis of a longer horizon distribution needs to be constructed from the sub-period increments of not only prices but also the variance and the expected third raw moment. Hence, our estimates of realized kurtosis additionally require daily changes in the expected third raw moment, which we directly derive from option-implied skewness (Equation A.15) and variance (Equation A.13).

Figure 2.10 plots the time-series of realized moments. The realized moments are considerably more volatile than their option-implied counterparts and display some erratic behavior. However, although realized skewness and kurtosis appear much noisier, they closely resemble the low-frequency variation of implied moments. In particular, they are clearly elevated during the period of near-zero interest-rates and experience a sharp drop around 2014/2015. Taken together, the trajectories of the physical moments lead to similar conclusions about the importance and stability of market expectations about a lower bound on interest rates.

Figure 2.10: Time-series of realized moments

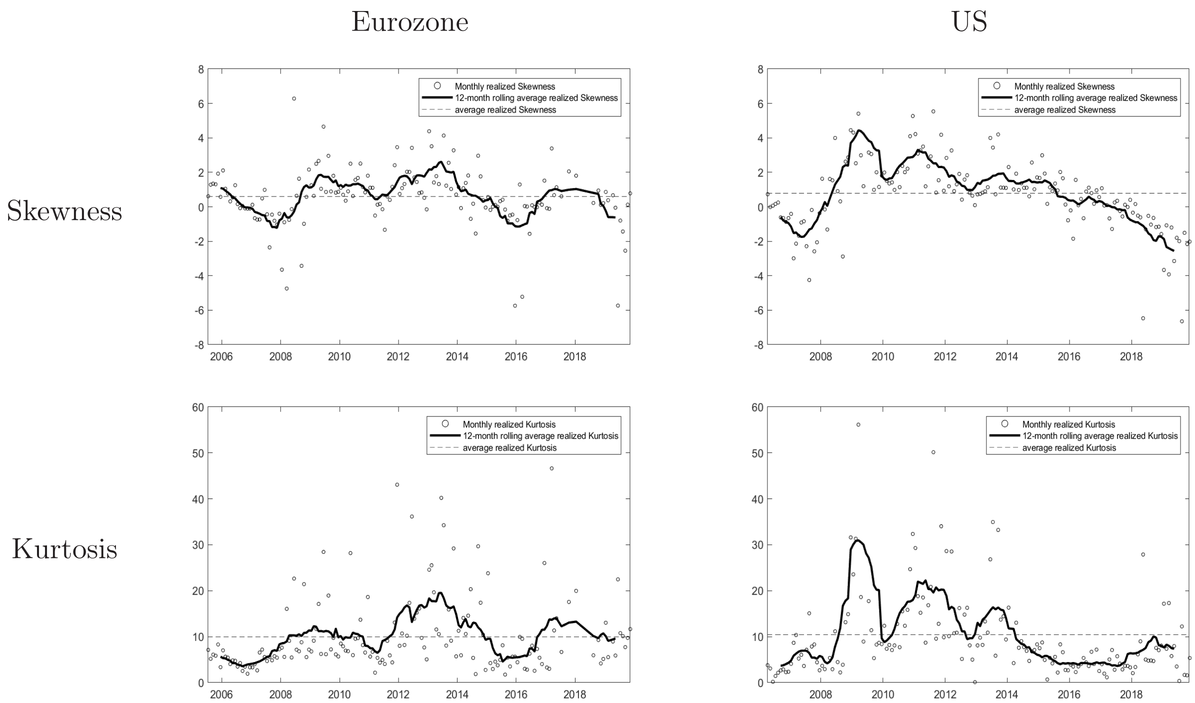

This figure shows time-series of monthly realized skewness and kurtosis, each interpolated between adjacent maturities to a constant maturity of one year. Red lines show 12-month moving average values.

Our time-series analysis of option-implied moments assumes that they provide a good description of the time-variation in conditional moments of the *physical* distribution. If risk premia are constant, risk-neutral moments will only differ from their physical counterparts by a shift in the *average* level. However, if risk premia vary considerably over time, the time-series variation of option-implied moments will fail to capture the *variation* in the true conditional moments of the physical distribution. We therefore attempt to validate whether risk-neutral interest rate moments provide a reasonable characterization of physical interest rate moments over time. If so, they should significantly predict future realized moments in a positive linear fashion. To assess this assertion, we run predictive regressions of realized moments on our option-implied measures:

$$\widetilde{M}_{t,T} = \alpha + \beta M_{t-1,T} + \epsilon_{t,T} \tag{2.8}$$

where $\widetilde{M}_{t,T}$ is the realized moment in month t and $M_{t-1,T}$ is the option implied moment sampled at the last day of month $t-1$. Table 2.10 shows the regression results. Option-implied moments are highly predictive of future realized moments. Coefficients are positive and strongly significant in every case. Intercepts are slightly negative for skewness and strongly positive for kurtosis, suggesting a risk-premia-induced bias that results in slightly higher average skewness and good bit lower average kurtosis in $\mathbb{Q}$-distributions. This suggests that risk-neutral market expectations somewhat underestimate the probability of large interest rate movements on average. R^2's are reasonably high. There is no reason to expect implied moments to fully explain realized moments even if they purely and perfectly reflect conditional expectations. Interestingly though, option-implied moments have similar - and in many cases even slightly higher - explanatory power for future realized moments compared to the lagged realized moments. Additionally, when we control for the own lag of the realized moments in a multivariate regression, option-implied moments still turn out to be strongly significant predictors. This suggests that option-implied measures capture variation in the *true* conditional $\mathbb{P}$-moments above and beyond the information contained in the $\mathbb{P}$-distribution *estimates.* Overall, these results suggest that our forward-looking option-implied estimates provide a very accurate assessment of the evolution of physical interest rate distribution characteristics.

Finally, we replicate our formal analysis replacing option-implied moments with realized moments. The results are summarized in Appendix A.1.5. Table A.3 and A.4 in the appendix replicate the results on level-dependent distribution characteristics from Table 2.1 and 2.2. Again, we see that the interest rate level alone is not able to explain the time variation in skewness and kurtosis. As a final robustness check, we also test our results on a potential shift in lower-bound views based on realized moments. Figure A.1 shows that break-date estimates based on realized moments quite closely match option-based estimates. Tables A.5 - A.8 replicate the structural break regression results. Again, our main results remain quite robust to the change in measure. Overall, our statements based on the distance to the lower bound also appear to be unaffected by time-varying risk premia.

Table 2.10: Predicting realized moments

	$\widetilde{Skew}_{t,T}$	$\widetilde{Skew}_{t,T}$	$\widetilde{Skew}_{t,T}$		$\widetilde{Kurt}_{t,T}$	$\widetilde{Kurt}_{t,T}$	$\widetilde{Kurt}_{t,T}$
Panel A: US							
α	0.015	0.347**	-0.021	α	4.546***	6.194***	3.317***
	(0.112)	(2.410)	(-0.181)		(5.747)	(6.855)	(4.644)
$Skew_{t-1,T}$	0.867***		0.576***	$Kurt_{t-1,T}$	0.710***		0.508***
	(9.528)		(5.606)		(6.056)		(4.624)
$\widetilde{Skew}_{t-1,T}$		0.531***	0.322***	$\widetilde{Kurt}_{t-1,T}$		0.365***	0.259***
		(6.792)	(4.587)			(5.243)	(3.991)
R^2	0.306	0.285	0.378	R^2	0.133	0.140	0.200
N	1703	1637	1637	N	1703	1637	1637
Panel B: EU							
α	-0.127	0.447***	-0.083	α	5.225***	7.776***	4.602***
	(-0.602)	(3.731)	(-0.432)		(6.637)	(9.217)	(5.524)
$Skew_{t-1,T}$	0.661***		0.510***	$Kurt_{t-1,T}$	0.608***		0.500***
	(5.520)		(4.250)		(5.151)		(4.619)
$\widetilde{Skew}_{t-1,T}$		0.324***	0.182***	$\widetilde{Kurt}_{t-1,T}$		0.199***	0.126*
		(5.885)	(4.206)			(2.917)	(1.960)
R^2	0.160	0.108	0.184	R^2	0.073	0.042	0.087
N	1157	1097	1097	N	1157	1097	1097

This table shows regression results for predicting monthly realized skewness ($\widetilde{Skew}_{t,T}$) and realized kurtosis ($\widetilde{Kurt}_{t,T}$) in the US and the Eurozone. We report regression coefficients with t-statistics in parentheses, R^2s and the number of observations. t-statistics are corrected for heteroscedasticity and serial correlation up to 12 lags (i.e. one year) using the approach of Newey and West (1987). *, ** and *** denote significance at the 10%, 5%, and 1% levels, respectively.

2.6 Conclusion

This chapter studies market participants' anticipation of a lower bound on interest rates. To this end we analyze the characteristics of option-implied interest-rate distributions in the US and the Eurozone.

We find that the lower bound temporarily had an important impact on the distributions of interest rates close to zero. Our estimated model free upper-limit for the market implied lower bound is consistent with the zero lower bound theory for several years. However, our results suggest that market participants significantly changed their lower-bound views over time. Using a structural change regression, we present evidence for a large structural break of the distribution characteristics of low interest rates, consistent with a marked downward revision of markets' lower-bound beliefs. Prior to the occurred break we document strong evidence of a substantial impact of a lower bound on interest-rate distributions. Pre-break distributions of low interest rates exhibit high positive skewness and kurtosis and strong positive level-dependence. After the break, the lower-bound impact is much weaker. The upper-limit for the lower bound in the post break period is clearly negative and comfortably below observed interest rates in both the Eurozone and the US. We consistently estimate this break to have occurred within the period when interest rates in many major economies fell into deeply negative territory for the first time in history in 2014/2015. These results point to the updating of investors' long-held zero lower bound beliefs in light of the negative interest rate experience.

Our results have important implications for modeling yield curves and the conduct of monetary policy. The estimation of dynamic term-structure models with historical data should account for a substantial shift in lower bound beliefs during the low-interest rate period. Prior to the occurrence of negative interest rates, it is important to enforce the zero lower bound constraint to capture the prevailing asymmetric interest rate distributions. After 2015, the model-implied lower bound should be substantially lower so as not to materially impact interest rate distributions. Practically, this could be done by estimating a lower bound consistent term structure model, such as a shadow rate model, and to exogenously provide a time-varying lower bound constraint. Properly accounting for this shift in lower-bound expectations is also essential for accurately estimating shadow rates, which heavily depend on assumptions regarding the location of the lower bound. Consequently, conclusions based on previously estimated shadow rates should be reassessed in light of this new evidence.

For future periods of low interest rates evidence from option implied distributions can be used to monitor the impact of a lower bound constraint on interest rates and asses the scope for additional monetary easing. Our model-free, measure agnostic upper limit for the lower bound provides a straightforward tool to track boundaries to market implied lower bound beliefs. It allows to assess in real-time whether current yields might already be constrained by a lower bound or whether there is still expected to be room for further rate cuts.

Chapter 3

Inflation cyclicality and stock-bond comovement: evidence from news media

3.1 Introduction

The recent drastic surge in inflation elevated market fears of stagflation to the dominant factor driving asset prices - reminiscent of the 1970s and 1980s. In contrast, the prolonged period of very low inflation preceding the pandemic gave much less rise to stagflation concerns in exchange for a much more prevalent threat of deflation. This "balance of risk" between stagflation and deflation potentially is a crucial determinant of the relationship between the two major asset classes – stocks and bonds. When markets perceive inflation news as *procyclical*, i.e. positively related to (long-run) consumption prospects, shocks to expected inflation move the prices of stocks and nominal bonds in opposite directions. On the other hand, if inflation is seen as countercyclical, expected inflation shocks cause stock and bond prices to move in the same direction. This mechanism may help explain the time-variation in the relation between stock- and bond-prices. In fact, the correlation between aggregate stock market returns and nominal bond yields fluctuates significantly over time, sometimes even changing sign. This dynamic is strongly tied to the correlation between stock market returns and expected inflation (see Figure 3.2). Inspired by this observation, a large theoretical literature connects time-varying stock-bond comovements to a changing relation between a persistent component of consumption growth and long-run expectations of inflation.[1] Such a relationship between persistent components - in combination with certain recursive preferences – permits already small shocks to expected inflation to generate large variations in stock prices.

Unfortunately, this inflation cyclicality channel is very difficult to verify empirically. Neither the persistent component of consumption growth, and even much less its relationship with expected inflation, can be directly observed in the data. This greatly amplifies the degree of "dark matter" (Chen et al. (2024)) – one of the main limitations of many asset pricing models. This chapter

[1] See e.g. Hasseltoft (2009), Burkhardt and Hasseltoft (2012), Song (2017), Dergunov et al. (2023).

aims to provide direct evidence for the changing cyclical behavior of inflation by evaluating data of media reporting on inflation.

To overcome the measurement problems in macroeconomic data, we attempt to extract the consumption/inflation relation directly from news media data. Our approach builds on the idea of Liu and Matthies (2022) that media reporting should capture investors' views on long-run growth prospects. We argue that if investors perceive inflation as strongly related to their long-run consumption outlook, this should be reflected in news as well. In fact, although the *inflation rate* itself remained very stable over most of our sample starting in 2000, Figure 3.1 reveals a remarkably high *press coverage of inflation* related topics in our news data set. Inflation even received more attention than other major economic topics like GDP, unemployment, or housing (Figure 3.1a). Some of this is due to the recent dramatic inflation surge in the post-covid period, when inflation news dominated the economic press with over 20% monthly coverage (Figure 3.1b). This is reminiscent of the attention inflation received in the 1980's (see e.g. Pfajfar and Santoro (2013)). But even aside these recent extremes, inflation-news-coverage throughout our sample was far from negligible and spiked above 10% on multiple occasions, documenting a steady substantial interest in inflation developments. These observations clearly indicate that inflation plays a significant role for readers of economic news. However, investors are not concerned with inflation dynamics per se but are ultimately focused on their long-term consumption outlook. The high level of inflation coverage suggests that investors perceive a strong link between inflation and their *real* consumption prospects. This is supported by anecdotal evidence, which shows that media articles often highlight the potentially drastic economic consequences of both high and low inflation scenarios, as illustrated in Table B.1. Therefore, media reporting on inflation is a crucial source for understanding whether investors view inflation as *procyclical*, meaning positively related to long-term consumption prospects, or *countercyclical*.

We construct such a measure of inflation cyclicality from news media data - the $News_{NRC}$ index. For this purpose we combine the content classification of news articles with their text sentiment. Our approach rests on the idea that economic articles containing "good news" ("bad news") for investors should on average be associated with positive (negative) text sentiment. "Good" respectively "bad" news for investors in the sense of a consumption-based (long-run risk) model are news about improvements/impairments of the long-term consumption outlook. Hence, we interpret the sentiment of economic news articles as indication for the consumption growth impact associated with the content of the article. We pair this sentiment assessment with the classification into articles about "rising inflation" (respectively "falling deflation") and "falling inflation" (respectively "rising deflation"). Accordingly, our measure compares the average text sentiment of "rising" vs. "falling" inflation articles. Thus, we measure whether inflation increases/decreases tend to be portrayed more positively or negatively in the news.

Figure 3.1: News coverage

(a) News coverage of major economic topics

(b) Monthly inflation news coverage

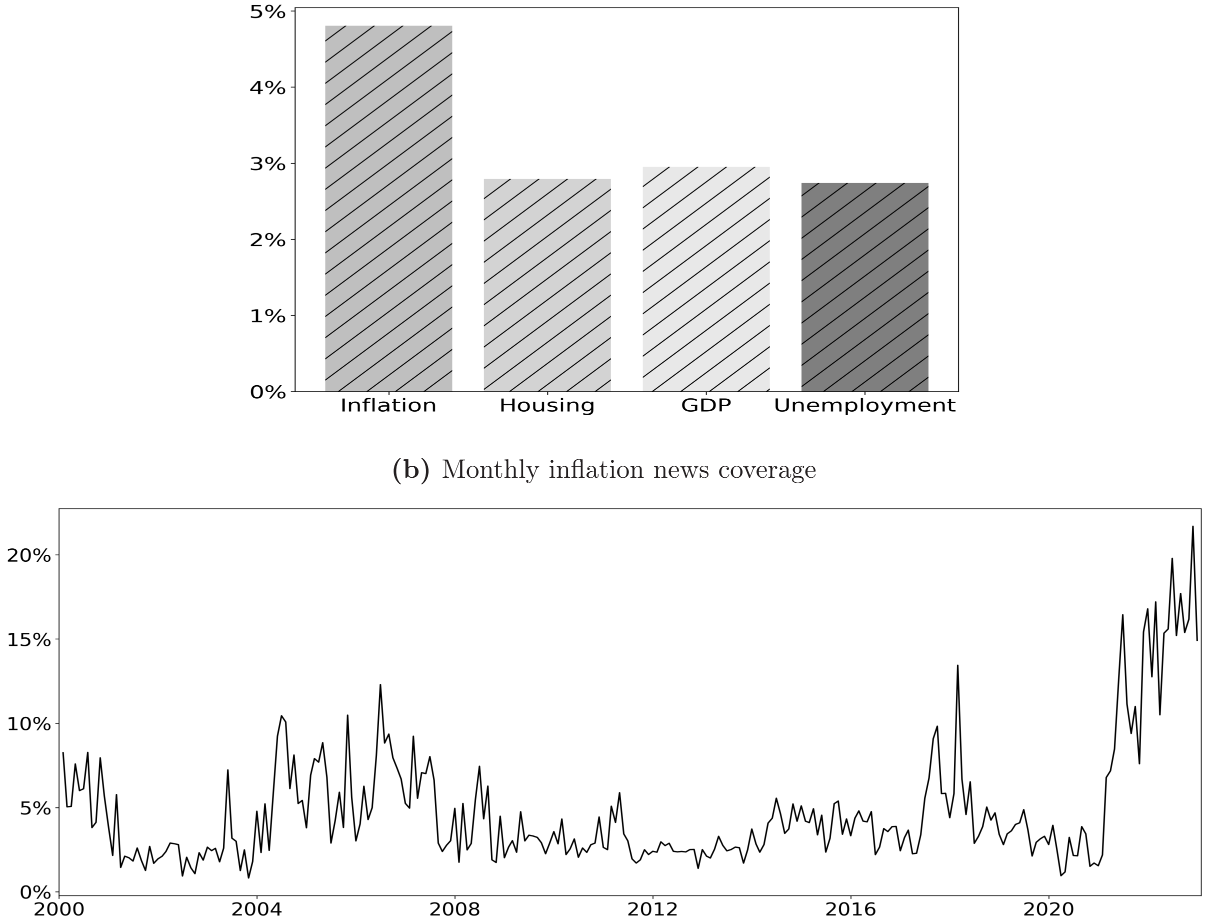

This figure depicts news coverage of inflation. Coverage is calculated as the monthly percentage share of inflation-related news articles to all economic news articles. Panel (a) compares average monthly coverage of inflation to coverage of other selected topics in our news-data sample from 01/2000-11/2022. Panel (b) shows the time-series of monthly inflation news coverage. News article data comes from Ravenpack and covers a broad range of textual news sources.

To highlight the relevance of our measure, we examine central predictions shared by models with a time-varying relationship between expected inflation and a persistent component of consumption growth. We present a stylized model where the state-dependent cyclical behavior of inflation is directly captured by a single exogenous variable. Given a measure for this state variable the asset-pricing implications of the inflation cyclicality channel allow for a straightforward empirical examination. Firstly, this state variable should forecast the realized correlations between stock-returns and inflation expectations. Secondly, it should directly determine the relationship between stock-returns and bond yields.

We employ our news measure for inflation cyclicality to empirically test these model predictions. We find that our cyclicality measure is in fact able to explain asset correlations in line with theoretical predictions. First, the news measure is *on average* positive. Hence, for most of our sample, news about inflation increases (decreases) have predominantly been associated with positive (negative) sentiment. This provides direct evidence that the mostly positive stock-return nominal bond yield correlation as well as the positive stock-return expected-inflation relation since 2000 can at least partly be rationalyzed by a procyclical inflation relation. Second, our measure is closely related to the *time-variation* in realized correlations. Predictive regressions reveal that our cyclicality news measure significantly predicts next-month realized correlations in a positive direction. This is true for both the correlations between stock returns and inflation expectations, and those between stock returns and bond yields. These results remain robust even after accounting for various variables previously suggested to predict stock-bond correlations. These variables likely also capture some cyclical aspects of inflation and may, therefore, operate through the same channel as our news measure. Our findings suggest that news carry valuable insights into cyclical relationships that are difficult to detect using fundamental data alone.

We further elaborate on this point by asking whether the conditional relationship between long-run consumption growth and inflation could also be identified with the empirical dynamics between inflation and consumption growth itself. To this end, we estimate rolling correlations between inflation and future cumulative long run consumption growth. We find that this rolling correlation exhibits a weakly positive yet insignificant relation with future asset correlations. The coefficient on our $News_{NRC}$ index remains significant when we include the rolling inflation/long-run consumption growth correlation as control variable. Moreover, we evaluate whether our results could be driven by a *contemporaneous* relation between inflation and consumption growth instead of a relation between *persistent* components. We proceed by estimating rolling correlations between contemporaneous consumption growth and inflation. Our results document that this measure of a relationship between *transitory* innovations has no significant predictive power for asset correlations. Controlling for this contemporaneous correlation hardly affects the coefficient on the $News_{NRC}$ index.

Finally, we test whether the data aligns with mean reversion in the inflation cyclicality state, a feature of both our model and alternative regime-switching models. The key implication is that the predictive power of the state variable ρ_t for future realized asset correlations should diminish with longer forecast horizons. Our regressions show that as the forecast horizon extends, both the regression coefficients and R^2 values decrease, confirming this model prediction.

We construct our news measure using pre-analyzed articles from Ravenpack, a leading provider in news analytics. This approach offers several benefits: it simplifies sentiment analysis and content classification by leveraging Ravenpack's expertise and extensive database. Most importantly, it eliminates the potential biases and variability in text sentiment and content classification that could arise from researcher judgment, by relying on Ravenpack's standardized methods. However, this approach depends on a "black-box" methodology, which can be challenging to evaluate. To address this, we conduct thorough preliminary checks to ensure the internal consistency of our measure.

Our approach to measuring inflation cyclicality from news data rests on two key assumptions: first, that text sentiment reflects investor views on economic prospects and captures persistent consumption growth; and second, that classifying articles into “falling inflation” and “rising inflation” accurately captures long-run inflation expectations. We empirically verify these assumptions using data from 2000 to 2022.

To verify the first key assumption regarding the relation between text sentiment and long-run consumption views, we examine how our sentiment scores align with long-term economic outlook indicators. We calculate an aggregate news sentiment measure by averaging Ravenpack’s sentiment scores for inflation-related articles. This measure shows a significant positive correlation with survey consumer confidence and the Liu and Matthies (2022) news indicator of persistent consumption growth. Additionally, aggregate news sentiment effectively predicts future consumption growth, with predictability increasing and peaking around a 60-month (5-year) horizon. This supports our assumption that text sentiment reflects long-term consumption growth views.

To assess the second key assumption about classification accuracy, we compare the relative coverage of rising versus falling inflation in our news data with actual changes in inflation rates.[2] Our analysis reveals a strong positive relationship, with an R^2 of 48% and a correlation of 0.69. Manual checks confirm accurate classification. Most importantly our measure of relative coverage effectively captures long-term inflation expectations, as evidenced by its significant correlation with market-implied and survey-based long-run inflation expectations.

Related Literature Our research contributes to the large literature that tries to connect asset prices to macroeconomic dynamics. Bansal and Yaron (2004) have argued that the key determinants of aggregate stock prices are expectations about long-run future real consumption growth rates. While successful in explaining many important stock-market features, one key ingredient of their long-run risk model - a small persistent component of expected consumption growth - is difficult to detect with univariate time-series methods. Bansal et al. (2012), Ortu et al. (2013) and Schorfheide et al. (2018) develop statistical methods to offer support for the existence of such a persistent component in consumption growth data. Closely related to our work, Liu and Matthies (2022) use economic news reporting to provide direct evidence for a persistent long-run risk factor. We extend this idea by investigating news-based evidence for a dynamic relationship between such a long-run risk factor and expected inflation.

[2] The plausible relationship between relative news coverage and inflation rate changes has already often been documented in the literature. See e.g. Lamla and Lein (2014, 2015), Lamla and Sarferaz (2012), Soroka (2006), Goidel and Langley (1995)

Starting with Hasseltoft (2009) and Burkhardt and Hasseltoft (2012), several papers tied time-variation in conditional stock-bond correlations to consumption/inflation dynamics. Burkhardt and Hasseltoft (2012) extend the long-run risk framework to jointly price stocks and nominal bonds. They match the empirical co-movements between these asset classes by modeling an exogenous, regime-dependent relation between persistent consumption growth and inflation expectations. Several authors followed their lead (see e.g. Dergunov et al. (2023), David and Veronesi (2013), Song (2017), Campbell et al. (2017), Seo (2023)). Campbell et al. (2020) rely on habit-preferences and exogenously changing *contemporaneous* correlations between inflation and consumption growth instead of a long-run risk mechanism. Some papers highlight the implications of time-varying inflation cyclicality for other asset classes (see Kang and Pflueger (2015), Bonelli et al. (2024), Boons et al. (2020), Burkhardt (2014)). Our contribution to this literature is to provide direct evidence for the inflation cyclicality channel which is the central model mechanism assumed in all these frameworks.

Our research is also related to the growing body of literature on textual analysis. Important contributions include Tetlock (2007), Baker et al. (2016), Manela and Moreira (2017). Our approach combines two recent strands of this literature. On the one hand, Liu and Matthies (2022) show that news articles contain views on long run consumption prospects predicting real consumption growth rates over long horizons. Similarly, van Binsbergen et al. (2024), Shapiro et al. (2022) and Doms and Morin (2004) report that media text sentiment is related to measures of consumer confidence and predicts future real consumption growth and output. On the other hand, starting with Carroll (2003), several papers document that news data capture information about inflation expectations. Carroll (2003) argues that news transmit inflation expectations of experts to households. He provides evidence that inflation news coverage induces households to update their expectations towards those of professional forecasters. Subsequent papers studied the interrelation between news and inflation expectations in further detail. Pfajfar and Santoro (2013) highlight that households may also adapt inflation views transmitted from news diverting from expert opinions. Lamla and Lein (2015) document that news coverage of falling vs. rising inflation importantly effects inflation perceptions. Lamla and Lein (2014) consider the tonality of inflation news and conclude that news media has a strong effect on households inflation expectations. Larsen et al. (2021) find that relative coverage of different news topics predicts inflation and inflation expectations. Picault et al. (2022) reveal that media coverage of monetary policy explains financial markets' implied expectations of long-term inflation. We merge these two thus far separate strands of the literature and use news data to extract information on the conditional relation between expected consumption growth and expected inflation.

Our theoretical model builds on the work of Bansal and Yaron (2004) and Bansal and Shaliastovich (2013). While Bansal and Shaliastovich (2013) add an *unconditional* relation between the persistent components of consumption growth and inflation, we introduce a *conditional* relation between expected consumption growth and inflation. This addition is crucial for our main objective: testing the asset-pricing implications of the inflation cyclicality channel using our news measure. By capturing inflation cyclicality with a single exogenous state variable, our model provides a straightforward and intuitive framework to guide this empirical analysis.

3.2 The role of stock-inflation dynamics in stock-bond comovement

Figure 3.2 shows realized correlations between daily US stock returns and daily changes in 10-year nominal Treasury yields for rolling samples of 253-days. The conditional stock/nominal bond yield correlation varies drastically over our sample, ranging from +0.7 to -0.35. Except for a few short periods, the correlation has been positive, implying a mostly negative relation between stock- and bond-returns. Figure 3.2 also displays the realized correlations between stock returns and changes in market expected inflation rates extracted from TIPS. Notably, the conditional stock/expected inflation correlation is very closely related to the stock/nominal bond correlation. To shed light on this relationship, note that the stock/nominal bond yield correlation can be decomposed into the correlation between stocks and real rates and stocks and market expected inflation:

$$cov_t(r_{m,t+1}, y^{\$}_{n,t+1}) = cov_t(r_{m,t+1}, y^{r}_{n,t+1}) + cov_t(r_{m,t+1}, y^{\pi}_{n,t+1}) \tag{3.1}$$

$$\rho^{\$,m}_{n,t} = \rho^{r,m}_{n,t}\frac{\sigma^{r}_{n,t}}{\sigma^{\$}_{n,t}} + \rho^{\pi,m}_{n,t}\frac{\sigma^{\pi}_{n,t}}{\sigma^{\$}_{n,t}} \tag{3.2}$$

where $r_{m,t}$ is the log stock market return at time t, $y^{\$}_{n,t}$ is the nominal bond yield of an n-period bond at time t, $y^{r}_{n,t}$ is the real bond yield of the same bond at time t and $y^{\pi}_{n,t}$ is the market implied inflation expectation at time t, calculated as $y^{\pi}_{n,t} = y^{\$}_{n,t} - y^{r}_{n,t}$. Hence the stock/nominal bond yield correlation $\rho^{\$,m}_{n,t}$ is given by the weighted sum of the stock/real yield correlation $\rho^{r,m}_{n,t}$ and the stock/expected inflation rate correlation $\rho^{\pi,m}_{n,t}$ with weights determined by the volatilities $\sigma^{\$}_{n,t}$, $\sigma^{r}_{n,t}$ and $\sigma^{\pi}_{n,t}$ of nominal yields, real yields and expected inflation. Therefore, the nominal stock-bond correlation can fundamentally be explained in theory with the stock-real relationship and the stock-expected inflation relationship. In this chapter we focus on the latter channel. Note that both $y^{r}_{n,t}$ and $y^{\pi}_{n,t}$ do not purely reflect expectations about future real rates respectively inflation but may also each contain a risk premium component. We focus on the decomposition into real- and inflation-components of stock/nominal bond yield correlations and do not further differentiate between expectation- and risk premia-contributions.

As (longer-term) inflation expectations have been quite stable, particularly during our sample since 2000, some researchers have argued that inflation should play only a minor role in explaining stock-bond correlations (Duffee (2023), Bansal et al. (2014)). Table 3.1 shows volatilities and volatility ratios of nominal yields, real yields, and market expected inflation rates in our sample. While long-term (market) inflation expectations in fact have been very stable (certainly as compared to the 1970s and 1980s), so have nominal yields. Volatility ratios indicate that the stock/expected inflation correlation has a significant weight of 0.64 in explaining the stock/nominal yield correlation. This is only slightly lower than the weight of 0.83 for the stock/real yield correlation, highlighting that inflation remains an important factor. Moreover, the conditional stock/expected inflation relation $\rho_t^{\pi,m}$ is very volatile and often very strong (see Figure 3.2). Together, these properties imply that the conditional stock/expected inflation correlation explains most of the time variation in the stock/nominal yield correlation. This conditional correlation has a time series correlation of 0.85 with the conditional stock/expected inflation correlation, as shown in Figure 3.2. Given that inflation expectations itself have been very stable raises the question what produces the occasional very strong positive stock/expected inflation relation and the large swings in this relationship over time. Any potential explanation must allow already small changes in expected inflation to have substantial impact on stock prices. In this chapter, we explore whether a time-varying relation between expected inflation and long-run consumption prospects can plausibly be made responsible for the observed patterns. To overcome the problem that long-run consumption expectations and their conditional relation with expected inflation cannot easily be observed from macroeconomic data, we turn to news media data to collect direct evidence for this channel.

Recent literature has increasingly focused on the real-rate channel to explain stock-bond correlations, with significant studies by Laarits (2020), Jones and Pyun (2024), Ermolov (2022), Kozak (2022), Chernov et al. (2021).[3] While our focus is on the inflation cyclicality channel, which we show significantly influences stock/expected inflation correlations, this does not diminish the importance of real-rate dynamics. Together, these approaches provide a more complete understanding of stock-bond comovement.

[3] Indeed, our results also support the real-rate channel: Real rates have been more volatile than market expected inflation ($\sigma^r = 0.76$ vs. $\sigma^\pi = 0.58$), the stock/real yield correlation varies significantly over time, though slightly less than the stock/expected inflation correlation ($\sigma(\rho_t^{r,m}) = 0.16$ vs. $\sigma(\rho_t^{\pi,m}) = 0.21$), and the stock/real yield correlation is almost as closely related to the stock/nominal yield correlation as the stock/expected inflation correlation ($Corr(\rho_t^{\$,m}, \rho_t^{r,m}) = 0.81$ vs. $Corr(\rho_t^{\$,m}, \rho_t^{\pi,m}) = 0.85$).

Figure 3.2: Realized correlations

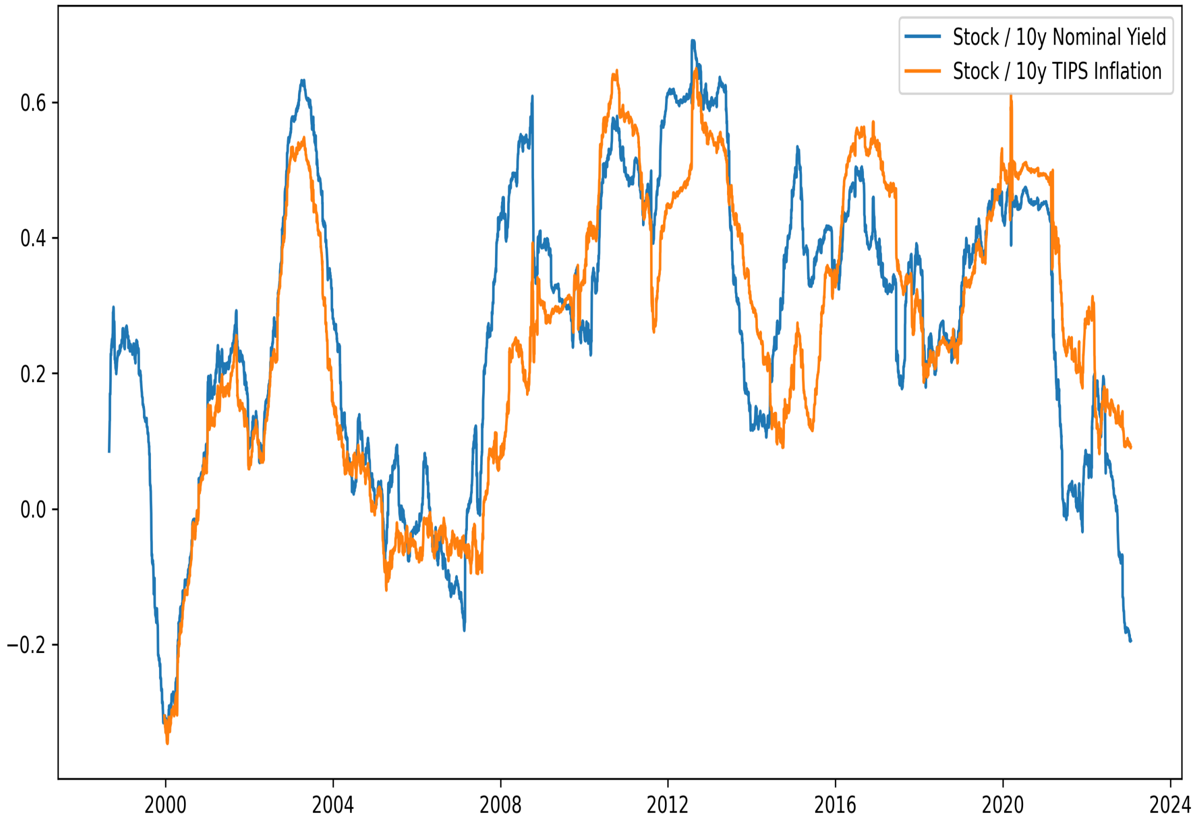

This figure shows conditional realized correlations between stock-market excess returns and changes in 10-year nominal treasury yields as well as changes in TIPS-implied breakeven inflation rates. Realized correlations are calculated over rolling samples of 253 days (1 year).

Table 3.1: Yield volatilities and volatility ratios

	Nominal	Real	Inflation
Volatility (σ^i)	0.912	0.757	0.583
Volatility ratio ($\frac{\sigma^i}{\sigma^{\$}}$)	1.000	0.831	0.640
Volatility of conditional correlation ($\sigma(\rho_t^{i,m}$)	0.220	0.164	0.210
Correlation of conditional correlations ($Corr(\rho_t^{\$,m}, \rho_t^{i,m})$)	1.000	0.811	0.849

This table shows unconditional volatilities and volatility ratios for 10-year nominal yields, real yields and market expected inflation for our sample from 2000 – 2022. It also shows the standard deviation and unconditional correlation-coefficients for the time-series of conditional correlations of nominal-, real- and market expected inflation-rates with stock-market returns.

3.3 A simple model of time-varying stock-bond correlation

To motivate our empirical analysis, we develop a tractable model in which the stock-bond correlation is driven by a time-varying relationship between expected consumption growth and inflation. We consider the discrete-time real endowment economy developed in Bansal and Yaron (2004) complemented by a process for inflation. Including inflation dynamics in the model allows us to jointly price nominal and real assets and to derive implications for the correlation between stock-returns and bond yields as well as the correlation between stock-returns and expected inflation. The model we present is intentionally similar in spirit to a range of models found in the literature.[4] Our goal is to provide a simple tractable framework that illustrates the inflation-cyclicality channel, which is the key mechanism shared across these models.

3.3.1 Dynamics of consumption and inflation

Our aim is to specify the most parsimonious dynamics that feature a time-varying relationship of persistent fluctuations in expectations of real growth and inflation. The exogenous processes for consumption growth and inflation are given by:

$$\Delta c_{t+1} = \mu_c + x_{c,t} + \sigma_c \eta_{c,t+1} \tag{3.3}$$

$$\pi_{t+1} = \mu_\pi + x_{\pi,t} + \sigma_\pi \eta_{\pi,t+1} \tag{3.4}$$

where $\eta_{c,t+1}$, $\eta_{\pi,t+1}$ are i.i.d. *transitory* shocks and $x_{c,t}$ and $x_{\pi,t}$ are the long-term components of consumption growth respectively inflation. Expected consumption growth and expected inflation follow the joint dynamics

$$x_{c,t+1} = \phi_{xc} x_{c,t} + \sigma_{xc} \eta_{xc,t+1} \tag{3.5}$$

$$x_{\pi,t+1} = \phi_{x\pi} x_{\pi,t} + \sigma_{x\pi} \zeta_{x\pi,t+1} \tag{3.6}$$

$$\zeta_{x\pi,t+1} = \rho_t \eta_{xc,t+1} + \sqrt{(1-\rho_t^2)}\eta_{x\pi,t+1} \tag{3.7}$$

where $\eta_{xc,t+1}$ and $\eta_{x\pi,t+1}$ are independent normal *persistent* shocks and ρ_t is the nominal-real correlation. ϕ_{xc} and $\phi_{x\pi}$ capture the persistence of expected consumption and expected inflation.

[4] See e.g. Burkhardt and Hasseltoft (2012), Dergunov et al. (2023), Song (2017)

Dividends have a leveraged exposure on expected consumption growth

$$\Delta d_{t+1} = \mu_d + \phi_d x_{c,t} + \sigma_d \eta_{d,t+1} \tag{3.8}$$

where $\eta_{d,t+1}$ is an independent normal shock and $\phi_d > 1$ can be interpreted as the leverage ratio on expected consumption growth of the companies issuing stocks. This captures the high volatility of dividends relative to consumption. For parsimony, we assume all conditional volatilities (σ_c, σ_π, σ_{xc}, $\sigma_{x\pi}$, σ_d) to be constant.

Equations (3.6) and (3.7) introduce a time-varying relation between expected inflation and expected consumption growth: shocks to expected inflation $\zeta_{x\pi,t+1}$ are composed of an independent inflation shock $\eta_{x\pi,t+1}$ and the expected consumption shock $\eta_{xc,t+1}$. The composition of shocks is determined by the variable ρ_t. If ρ_t is positive, shocks to expected consumption move inflation expectations in the same direction, therefore implying a *procyclical* inflation-consumption relation. Negative values of ρ_t lead to a *countercyclical* inflation-consumption relation with shocks to expected consumption moving expected inflation in the opposite direction. For values of ρ_t around zero, shocks to expected inflation are independent of shocks to expected consumption. This specification implies:

$$(\zeta_{x\pi,t+1}, \eta_{xc,t+1}) \sim N(0, \Sigma_t) \tag{3.9}$$

where the coviarance matrix Σ_t is:

$$\Sigma_t = \begin{pmatrix} 1 & \rho_t \\ \rho_t & 1 \end{pmatrix} \tag{3.10}$$

Hence, ρ_t captures the correlation between expected consumption growth and expected inflation, which we refer to as the nominal-real correlation in our model. The correlation ρ_t can vary between -1 and 1 over time. This is modelled through a Fisher-Transformation which follows a mean reverting process:

$$\rho_t = \frac{e^{2\tilde{\rho}_t} - 1}{e^{2\tilde{\rho}_t} + 1} \tag{3.11}$$

$$\tilde{\rho}_{t+1} = \mu_{\tilde{\rho}} + \phi_{\tilde{\rho}}(\tilde{\rho}_t - \mu_{\tilde{\rho}}) + \sigma_{\tilde{\rho}} \eta_{\tilde{\rho},t+1} \tag{3.12}$$

The transformation ensures that ρ_t stays within the -1 to 1 range and allows for changes in sign. In our model, the variable ρ_t is the key factor driving how correlations between stocks, inflation, and bonds change over time. This mechanism is the only one responsible for these time-dependent correlations. As discussed earlier, neither the long-term component of consumption growth ($x_{c,t}$) nor its connection with expected inflation (ρ_t) can be directly observed from macroeconomic data. Instead, we aim to test this mechanism directly by creating an empirical measure of ρ_t from news data.

3.3.2 Asset-pricing implications

We assume a representative investor with recursive preferences over consumption C_t described by the Kreps-Porteus, Epstein-Zin utility function (Kreps and Porteus (1978); Epstein and Zin (1989)). The key parameters are γ, which represents relative risk aversion, and ψ, denoting the elasticity of intertemporal substitution (EIS). Together, they determine the investor's preference for early resolution of uncertainty, with the condition $\gamma > \frac{1}{\psi}$ indicating a preference for resolving risk sooner. In order to derive analytical model solutions, we proceed as Bansal and Yaron (2004) and use a Campbell-Shiller (Campbell and Shiller (1988)) return log linearization. Details of the model solution are given in Appendix B.1.2. The real stochastic discount factor is given by

$$m_{t+1} = m_0 + m_{xc} x_{c,t} - \lambda_c \sigma_c \eta_{c,t+1} - \lambda_{xc} \sigma_{xc} \eta_{xc,t+1} \tag{3.13}$$

Its conditional mean is affine in expected consumption growth with loading $m_{xc} = -\frac{1}{\psi}$. Innovations to the real pricing kernel are a function of the fundamental consumption shocks and their respective market prices of risk λ_c and λ_{xc}.

Notably, expected inflation $x_{\pi,t}$ and the nominal-real correlation ρ_t do not independently enter the SDF, indicating that, once expected consumption growth $x_{c,t}$ is observed, expected inflation does not add separate risk. However, expected inflation is not riskless; its correlation with long-run consumption growth introduces a risk premium in asset pricing. The risk premium varies based on the sign of ρ_t: A negative relationship between expected growth and inflation ($\rho_t < 0$) leads to a positive inflation risk premium, as assets with low real payoffs in high inflation states are deemed risky. Conversely, a positive relationship ($\rho_t > 0$) implies a negative risk premium, as these assets provide a hedge during low inflation and low expected consumption growth. This mechanism makes expected inflation and the nominal-real correlation crucial for pricing nominal assets and their relation to real assets in our model.
We focus on the following key predictions of our model:

1. **Equity-inflation relation**
   The first and foremost prediction of our model is that the nominal-real correlation ρ_t directly determines the stock-inflation relation:

$$cov_t(r_{m,t+1}, x_{\pi,t+1}) = \beta_{xc} \sigma_{xc} \sigma_{x\pi} \rho_t \tag{3.14}$$

Given that the stock-market positively loads on long-run risk ($\beta_{xc} > 0$), a positive (negative) relationship between expected inflation and persistent consumption growth implies a positive (negative) stock-inflation correlation.[5]

2. **Equity-nominal bond yield relation**
The second central prediction of our model is that the nominal-real correlation ρ_t directly determines the stock-nominal bond yield relation:

$$cov_t(r_{m,t+1}, y^{\$}_{n,t+1}) = -\frac{1}{n}\left(\sigma^2_{xc}\beta_{xc}B_{xc,n} + \sigma_{x\pi}\sigma_{xc}\beta_{xc}B^{\$}_{x\pi,n}\rho_t\right) \tag{3.15}$$

Real yields are positively related to expected consumption growth ($B_{xc,n} < 0$), introducing a positive component in the stock-nominal yield relationship. The relation between expected consumption growth and expected inflation is determined by ρ_t, introducing a conditional component in the stock-nominal yield relationship.[6]

3. **Distant horizon asset correlations**
Lastly, our model predicts that the relationship between the current nominal-real correlation ρ_t and asset-correlations weakens for increasing prediction horizons:

$$\begin{aligned} cov_{t+h}(r_{m,t+h+1}, x_{\pi,t+h+1}) \approx & (\beta_{xc}\sigma_{xc}\sigma_{x\pi})\mu_\rho \sum_{i=1}^{h} \phi_\rho^{i-1} \\ & + (\beta_{xc}\sigma_{xc}\sigma_{x\pi})\phi_\rho^h \rho_t \\ & + (\beta_{xc}\sigma_{xc}\sigma_{x\pi})\sigma_\rho \sum_{i=1}^{h} \phi_\rho^{h-i}\eta_{\tilde{\rho},t+i} \end{aligned} \tag{3.16}$$

$$\begin{aligned} cov_{t+h}(r_{m,t+h+1}, y^{\$}_{n,t+h+1}) \approx -\frac{1}{n}\Bigg[& \sigma^2_{xc}\beta_{xc}B_{xc,n} + (\sigma_{x\pi}\sigma_{xc}\beta_{xc}B^{\$}_{x\pi,n})\mu_\rho \sum_{i=1}^{h} \phi_\rho^{i-1} \\ & + (\sigma_{x\pi}\sigma_{xc}\beta_{xc}B^{\$}_{x\pi,n})\phi_\rho^h \rho_t \\ & + (\sigma_{x\pi}\sigma_{xc}\beta_{xc}B^{\$}_{x\pi,n})\sigma_\rho \sum_{i=1}^{h} \phi_\rho^{h-i}\eta_{\tilde{\rho},t+i} \Bigg] \end{aligned} \tag{3.17}$$

The relationships weaken for increasing horizon h because inflation cyclicality ρ_t is expected to revert back to its unconditional mean. Moreover, innovations to ρ_t accumulate increasing uncertainty about future states of inflation cyclicality. Therefore, ρ_t increasingly loses explanatory power over the forecast horizon.

[5] When the EIS and the dividend leverage ratio are sufficiently high ($\phi_d > \frac{1}{\psi}$), stock market returns are positively related to expected consumption growth ($\beta_{xc} > 0$).

[6] If $\rho_t > -\frac{\sigma_{xc}B_{xc,n}}{\sigma_{x\pi}B^{\$}_{x\pi,n}}$, the real channel dominates and nominal yields increase with expected consumption growth. If $\rho_t < -\frac{\sigma_{xc}B_{xc,n}}{\sigma_{x\pi}B^{\$}_{x\pi,n}}$, the nominal channel dominates as the reaction to decreases in expected inflation associated with increases in expected growth surpass the real effect.

In this chapter we extract information about the nominal-real correlation ρ_t from the news. We refer to this news measure as the $News_{NRC}$ index. This measure allows us to empirically evaluate above model predictions for asset correlations and provide direct evidence for the inflation cyclicality channel.

3.4 Data

3.4.1 News measure of inflation cyclicality

Our news measure of inflation cyclicality links the text sentiment of news articles to their content. We use the text sentiment of economic news articles as a proxy for investor sentiment about future consumption. This is motivated by the literature showing that media text sentiment is related to measures of consumer confidence and predicts future real consumption and output (van Binsbergen et al. (2024), Shapiro et al. (2022), Doms and Morin (2004)). We exploit this relationship by employing text sentiment to identify whether news articles about rising respectively falling inflation convey good or bad long-run consumption news to investors.

We use news analytics data from Ravenpack. Ravenpack is the leading provider of news analytics and its dataset has been extensively used in the finance literature in recent years.[7] Our primary reason for relying on pre-analyzed news data is to eliminate the vast degrees of freedom that a custom procedure to calculate text sentiment and perform content classification would entail. Ravenpack tracks around 22.000 news sources including local, regional and national news (e.g. Dow Jones, Wall Street Journal, Barron's, MarketWatch, MT Newswires, Benzinga), press releases, and reputable blogs. Historical data is available from 2000, providing almost 22 years of data. The dataset is at the news level, with each observation representing a news article uniquely identified by its publication timestamp, the source, and the article headline. For each news article, Ravenpack provides a range of document-level news analytics, generated by automated proprietary text-processing algorithms. In constructing of our news measure, we mainly rely on two news analytics variables: the *text sentiment* of an article and its *topic classification*. The Ravenpack text sentiment score ("CSS" – "composite sentiment score") intends to capture the tonality of the news article by combining multiple textual analysis techniques. Generally, it evaluates emotionally charged words and phrases and matches articles to similar text with expert-rated sentiment. CSS scores range between -1 and 1, with negative (positive) values indicating negative (positive) text sentiment. The Ravenpack topic classification evaluates the content of the news article and identifies its topic or topics from a predefined list. This topic taxonomy comprises 6895 topics, hierarchically organized into topic-groups (such as "Economy"

[7] See e.g. Campbell and Shang (2022), Reed et al. (2020), Ben-Rephael et al. (2017), Bushman et al. (2017), von Beschwitz et al. (2020), Kolasinski and Yang (2018), Jiang et al. (2021).

or ”Politics”) and multiple sub-topic layers. For our news measure, we aggregate several Ravenpack topic categories related to inflation: “consumer price index”, “consumer price index rising”, “consumer price index falling”, “producer price index”, “producer price index rising”, “producer price index falling”, “inflation”, “inflation rising” and “inflation falling”. We then consolidate these categories into broader groups: “inflation”, “inflation rising”, and “inflation falling”.

We construct our time-series news index, the $News_{NRC}$ measure, by calculating the weighted average text sentiment of all “inflation rising” articles, net of “inflation falling” articles:

$$News_{NRC,t} = \frac{\sum\limits_{i \in N(t)} s_i \left(I_{\pi-rising}(i) - I_{\pi-falling}(i) \right)}{\sum\limits_{i \in N(t)} I_{\pi}(i)} \tag{3.18}$$

s_i is the Ravenpack text sentiment score of article i, and $I_\pi(i) \in (0,1)$, $I_{\pi-rising}(i) \in (0,1)$, $I_{\pi-falling}(i) \in (0,1)$ are indicator variables that are 1 if article i is classified in the respective content category (inflation "rising"/"falling") and 0 otherwise; and $N(t)$ denotes the set of articles published in month t. This allows us to measure how the average text sentiment associated with “inflation rising” articles compares to that of “inflation falling” articles. We aggregate this measure at monthly frequency resulting in a time-series of 275 observations from January 2000 to November 2022.

This monthly news measure could potentially be affected by measurement noise in the component capturing inflation cyclicality. To address this concern, we calculate an alternative, lower-frequency version of our index, the $MANews_{NRC}$ index, by applying a simple 3-month moving-average filter to the raw $News_{NRC}$ index.

For better interpretability, we standardize both the $News_{NRC}$ and the $MANews_{NRC}$ using their respective standard deviations.

3.4.2 Realized correlations

To evaluate the main asset pricing predictions of the inflation cyclicality channel in our model, we require an empirical measure of asset correlations. For this purpose, we rely on realized correlations between stock market returns, changes in treasury yields, and changes in market inflation expectations at a daily frequency. The monthly realized correlation between stock returns and yield changes (or inflation expectation changes) in month t is given by

$$\tilde{corr}_t(r_m, y_n) = \frac{\sum_{d=t_1}^{t_N} r_{m,d} \Delta y_{n,d}}{\sqrt{\sum_{d=t_1}^{t_N} r_{m,d}^2 \sum_{d=t_1}^{t_N} \Delta y_{n,d}^2}} \tag{3.19}$$

where t_1 and t_N denote the first and last available daily return observations' dates in month t, $r_{m,d}$ is the log stock market return on day d and $\Delta y_{n,d}$ are the yield changes of nominal yields or expected inflation rates. Andersen et al. (2001, 2003) and Barndorff-Nielsen and Shephard (2004), among others, show that this realized correlation estimator is an unbiased estimator of the actual (unobserved) correlation. It has the advantage to be a model free measure not depending on any particular parametric assumptions and is approximately free from measurement error given sufficient sampling frequency.

Stock market returns $r_{d,m}$ are the daily returns of the nearest-to-maturity S&P-500 futures traded at the Chicago Mercantile Exchange, which we obtain from Bloomberg. These contracts are highly liquid and serve as a popular equity benchmark to investors. Since futures returns are derived from the price quotes of a single security rather than an aggregation of multiple securities, they help mitigate potential microstructure-induced measurement noise (see Ahn et al. (2002a)). Furthermore, realized futures returns can conveniently be interpreted as excess returns, simplifying the calculation of second moments.

Nominal yield changes $\Delta y^{\$}_{n,d}$ are the daily changes of the continuously compounded 10-year Treasury zero coupon constant maturity yield from Gürkaynak et al. (2007). Expected inflation rate changes $\Delta y^{\pi}_{n,d}$ are the daily changes of the 10-year Break-even inflation rates from Gürkaynak et al. (2010). These are calculated as the difference between the 10-year nominal yields and the corresponding real yields of Treasury inflation protected securities (TIPS).

3.5 Consumption and inflation expectations in news data

In order for our news measure, $News_{NRC}$, to successfully capture views on inflation cyclicality and thus serve as an adequate proxy for the state variable ρ_t, two main conditions must be met: First, the text sentiment measure we rely on must reflect investors' views on (long-run) consumption prospects. Second, articles classified as "inflation rising" and "inflation falling" must be associated with corresponding views on (long-term) inflation. Before we proceed to use $News_{NRC}$ as an empricial proxy for ρ_t, we aim to validate these central internal consistency conditions.

3.5.1 News text sentiment and investor sentiment about (long-run) consumption

Our approach to measuring inflation cyclicality is based on the assumption that news text sentiment, as captured by Ravenpack's analysis of news articles on inflation, provides a meaningful proxy for investor sentiment regarding future long-run consumption growth. This assumption is supported by the findings of Liu and Matthies (2022), who show that news coverage of economic growth reflects investors' concerns about the economic outlook and serves as a reliable predictor of real consumption growth over long horizons. To assess whether this relationship holds in our data, we calculate aggregate text sentiment by averaging the sentiment scores of all inflation-related articles:

$$News_{s,t} = \frac{\sum\limits_{i \in N(t)} s_i I_\pi(i)}{\sum\limits_{i \in N(t)} I_\pi(i)} \tag{3.20}$$

This calculation allows us to assess the text sentiment variable in isolation, enabling its use in estimating inflation cyclicality with our main measure, the $News_{NRC}$ index. To enhance robustness, we compute $News_{s,t}$ as both a monthly index and a lower-frequency moving average, denoted as $MANews_{s,t}$, filtering out short-term noise.

First, we compare these news text sentiment measures to alternative measures of aggregate investor sentiment: the University of Michigan Consumer Confidence Survey ("UMC"), the N-Index ("LMN") and the HN-Index ("LMHN") from Liu and Matthies (2022), Option Implied Stock-Market Volatility ("VIX"), and the Economic Policy Uncertainty Index from Baker et al. (2016) ("EPU"). We expect $News_{s,t}$ to be related to these alternative measures of investor sentiment, albeit this relationship is likely far from perfect due to apparent conceptual differences. Table 3.2 shows the results from regressing $News_{s,t}$ on the alternative sentiment measures,

$$News_{s,t} = \alpha + \beta * sentiment_t^i + \epsilon_t \tag{3.21}$$

Both the raw and filtered versions of the aggregate text sentiment are negatively associated with the second moment measures, VIX and EPU, suggesting that increases in market volatility and economic policy uncertainty are linked to decreases in investor sentiment about future consumption growth. In contrast, the text sentiment measures load positively on UMC, the LMN-Index, and the LMHN-Index, indicating that improvements in consumer confidence and long-run consumption expectations are reflected in more optimistic news sentiment. The relationship between $News_{s,t}$ and UMC is particularly strong and significant, highlighting that news sentiment is closely aligned with investors' views on consumption outlook. Filtering out measurement noise further increases the relationship between text sentiment and survey consumer confidence. Moreover, $News_{s,t}$ is significantly related to the LMN-Index and its trend component, the LMHN-Index, from Liu and Matthies (2022), both of which serve as proxies for long-run consumption growth. The strong relationship between the low-frequency text sentiment index ($MANews_{s,t}$) and the LMHN-Index suggests that our news measure captures similar long-term consumption expectations as those reflected in the Liu and Matthies (2022) consumption growth indices.

Table 3.2: Comparison between aggregate text sentiment and other variables

	UMC	LMN	LMHN	VIX	EPU
Panel A: $News_s$ index					
Coefficient	0.017*	0.187*	0.472*	-0.015	-0.001
	(3.751)	(2.532)	(3.477)	(-1.244)	(-0.679)
Corr	0.227	0.187	0.252	-0.119	-0.051
R^2	0.048	0.029	0.058	0.010	-0.001
N	275	168	168	275	275
Panel B: $MANews_s$ index					
Coefficient	0.025*	0.285*	0.687*	-0.015	-0.001
	(4.920)	(3.862)	(4.640)	(-1.058)	(-0.646)
Corr	0.330	0.286	0.370	-0.119	-0.052
R^2	0.105	0.076	0.131	0.011	-0.001
N	275	168	168	275	275

This table documents the relationship between the text sentiment variable in our dataset and alternative measures of investors' views on consumption prospects and economic uncertainty. It shows the results from regressing our aggregate economic text sentiment measures $News_s$ (Panel A) and $MANews_s$ (Panel B) on the selected variables. We report coefficients, t-statistics (in parentheses), correlation coefficients and adjusted R^2s. Standard-errors are corrected for heteroscedasticity and serial correlation using the approach of Newey-West with one lag. The LMN and LMHN indices are only available until 2013, all other regressions are for the full sample period until 11/2022. * indicates significance at the 5%-level under the Newey-West adjusted standard errors.

Second, we evaluate the real consumption growth predictability as in Liu and Matthies (2022). If news text sentiment proxies for the persistent component $x_{c,t}$ in Equation (3.3) it should predict future real consumption growth. Predictability should increase over the cumulative horizon as the transitory shocks $\eta_{c,t}$ average out. Consequently, we run the following predictive regressions for cumulative future real consumption growth:

$$\sum_{i=1}^{K} \Delta c_{t+i} = \alpha + \beta News_{s,t} + e_{t+K} \quad K \in \{1, 2, \dots 60\} \tag{3.22}$$

Δc_{t+i} is per capita real consumption growth in month $t+i$. We measure monthly aggregate real consumption with personal consumption expenditures (PCE) of nondurable goods and services deflated by the respective PCE price indices. Monthly per capita real consumption growth is the log-difference between monthly real consumption levels over population numbers.

Table 3.3 reports the regression results. The coefficients are positive across all horizons from 1 year to 10 years. However, at short horizons, coefficients are insignificant and R^2's are close to

Table 3.3: Predicting future consumption growth

Horizon (years)	1	2	3	4	5	6	7	8	9	10
$News_s$	0.058	0.225	0.527*	0.674*	1.000*	0.757*	0.608*	0.498*	0.361*	0.387
	(0.214)	(0.839)	(2.091)	(2.117)	(3.691)	(2.641)	(2.369)	(2.510)	(2.581)	(1.663)
Adj. R^2	-0.003	0.002	0.029	0.046	0.089	0.047	0.029	0.018	0.007	0.012
N	268	256	244	232	220	208	196	184	172	160
$MANews_s$	0.280	0.480	0.744*	1.032*	1.404*	1.053*	0.735*	0.517	0.428*	0.338
	(1.233)	(1.735)	(2.089)	(2.863)	(4.120)	(2.904)	(2.215)	(1.560)	(2.466)	(1.458)
Adj. R^2	0.008	0.022	0.060	0.111	0.175	0.092	0.043	0.019	0.012	0.008
N	268	256	244	232	220	208	196	184	172	160

This table documents the predictive content of the news text sentiment variable in our dataset for real consumption growth. It shows the results from regressing future cumulative real per capita consumption growth on our aggregate economic text sentiment measures $News_s$ and $MANews_s$ for different horizons. We report coefficients, t-statistics (in parentheses) and adjusted R^2s. Standard-errors are corrected for heteroscedasticity and serial correlation using the approach of Newey-West with n lags, where n is the number of overlapping periods. Consumption growth rates are multiplied by 100. * indicates significance at the 5%-level under the Newey-West adjusted standard errors.

zero. T-statistics increase for longer horizons and indicate significance at the 5% level for horizons from 3 to 9 years. Adjusted R^2's also increase for longer horizons and peak at 5 years, where they reach 17.5% for the low-frequency $MANews_s$ measure. This aligns closely with the results of Liu and Matthies (2022), who find that predictability peaks at the 6-year horizon. While their study finds somewhat stronger predictability, potentially due to a longer sample period, our results are encouraging and broadly align with their conclusions. This finding provides further support of a predictive component in long-run consumption growth which can be captured with news data. It also justifies our interpretation of text sentiment as a proxy for *long-run* consumption growth and our main $News_{NRC}$ index as a measure between *persistent* components.

3.5.2 Content classification and investors' (long-term) inflation expecations

For our $News_{NRC}$ measure to successfully capture the relationship between consumption and expected inflation, the Ravenpack classification of inflation articles must be sufficiently accurate. Moreover, news coverage of inflation needs to reflect investors' inflation expectations. We verify that both conditions are met in our data.

First, we randomly select a sample of 100 articles classified as either "rising inflation" or "falling inflation" and manually review their content. We confirm that the vast majority (96) of these articles either include sections or fully focus on discussing rising or falling inflation developments, in line with the classification result. Second, as a more replicable plausibility check, we use the presumption that the relative news coverage of rising vs. falling inflation should closely align with realized developments in the inflation rate. Specifically, during periods of rising inflation rates, we expect news coverage to focus more on “rising inflation” compared to periods of falling inflation rates. This plausible relationship has already been documented in the literature (see e.g. Lamla and Lein (2014, 2015), Lamla and Sarferaz (2012), Soroka (2006), Goidel and Langley (1995)). To test this hypothesis in our data we construct a measure of relative directional news coverage on inflation. We proceed exactly as with our $News_{NRC}$ measure but instead of evaluating the sentiment, we only focus on the share of articles classified as “rising inflation” net of the share of articles classified as “falling inflation”:

$$News_{\pi,t,h} = \frac{\sum\limits_{i \in N(t,t-h)} (I_{\pi-rising}(i) - I_{\pi-falling}(i))}{\sum\limits_{i \in N(t,t-h)} I_{\pi}(i)} \tag{3.23}$$

We regress this measure on changes of realized inflation rates

$$News_{\pi,t,h} = \alpha + \beta \Delta\pi_{t-h,t} + \epsilon_{t,h}, \tag{3.24}$$

where both the change in the realized annual inflation rate $\Delta\pi_{t-h,t}$ and the relative directional news coverage $News_{\pi,t,h}$ are calculated over the past h months. Confirming our hypothesis, Table 3.4 shows that relative news coverage significantly depends on changes in the realized inflation rate. News coverage leans more towards “inflation rising” articles in times of rising inflation and vice versa. Annual relative news coverage is explained by realized changes in the inflation rate over the same period, with an R^2 of 48% and a correlation of 0.69. These results provide strong evidence of the reliable classification of articles into “inflation rising” and “inflation falling” in our dataset.

Next, we examine whether the classification into "rising inflation" and "falling inflation" articles not only accurately reflects inflation trends but, more importantly, captures information about investors’ inflation expectations. To evaluate this empirically, we analyze how our aggregate measure of relative directional news coverage relates to measures of aggregate inflation expectations. We expect a significant positive contemporaneous relation between news $News_{\pi,t,h}$ and changes in inflation expectations if news capture investors’ inflation views. Consequently, we run regression (3.24) for two measures of expected inflation: market-implied inflation expectations calculated from 10-year Treasury Inflation Protected Securities (y^{π}) and mean survey inflation expectations from the survey of professional forecasters (π_{SPF}). Table 3.4 documents a uniformly positive relationship between $News_{\pi,t,h}$ and changes in inflation expectation over all horizons. Coefficients are significant at the 5% level in all but one case. At the annual horizon, R^2 reaches 23%. Hence our news measure successfully captures information about investors’ inflation expectations.

Table 3.4: News coverage and changes in inflation rates

Period (months)		$\Delta\pi_{CPI}$	Δy^{π}	$\Delta\pi_{SPF}$
	Coefficient	15.197*	19.017*	0.387
		(3.684)	(2.554)	(1.313)
1	Corr	0.291	0.147	0.087
	R^2	0.081	0.018	0.004
	N	275	275	274
	Coefficient	9.608*	16.549*	0.378*
		(6.137)	(4.364)	(2.396)
3	Corr	0.555	0.320	0.203
	R^2	0.306	0.099	0.038
	N	275	275	272
	Coefficient	2.923*	6.166*	0.282*
		(9.981)	(3.893)	(5.168)
12	Corr	0.694	0.380	0.483
	R^2	0.480	0.141	0.230
	N	275	275	263

This table evaluates the inflation information contained in the content classification of articles into "rising inflation" and "falling inflation". It shows the relationship between our measure of relative directional news coverage, $News_{\pi}$, and changes in realized inflation, changes in market-implied inflation expectations, and changes in survey inflation expectations. For each variable, we run univariate regressions on $News_{\pi}$ for monthly, quarterly and annual (overlapping) periods. In each case, we report coefficients, t-statistics (in parentheses), correlation coefficients and adjusted R^2s. Standard-errors are corrected for heteroscedasticity and serial correlation using the approach of Newey-West with 12 lags. * indicates significance at the 5%-level under the Newey-West adjusted standard errors.

3.6 Inflation news sentiment and asset correlations

Given that the $News_{NRC}$ index proxies for the nominal-real correlation between persistent consumption growth and inflation, it should, according to our model hypotheses, explain the observed correlations between nominal and real assets (Prediction (1) and (2)). Figure 3.3 shows the monthly $News_{NRC}$ index together with its filtered version, the $MANews_{NRC}$ index. Table 3.5 provides summary statistics. On average, the $News_{NRC}$ index is positive with only brief negative periods, including late 2022. This suggests that news of rising inflation tend to be associated with a more positive sentiment compared to news of falling inflation. This aligns with the positive stock-inflation and negative stock-bond correlations observed since 2000. Studies, including Campbell et al. (2017) and Burkhardt and Hasseltoft (2012), link this shift around the turn of the century to a positive consumption-inflation relationship. Our findings offer direct evidence supporting this link, indicating that *long-run* consumption and inflation expectations likely drove these asset correlations since the early 2000s.

Figure 3.3: $News_{NRC}$ and $MANews_{NRC}$ index

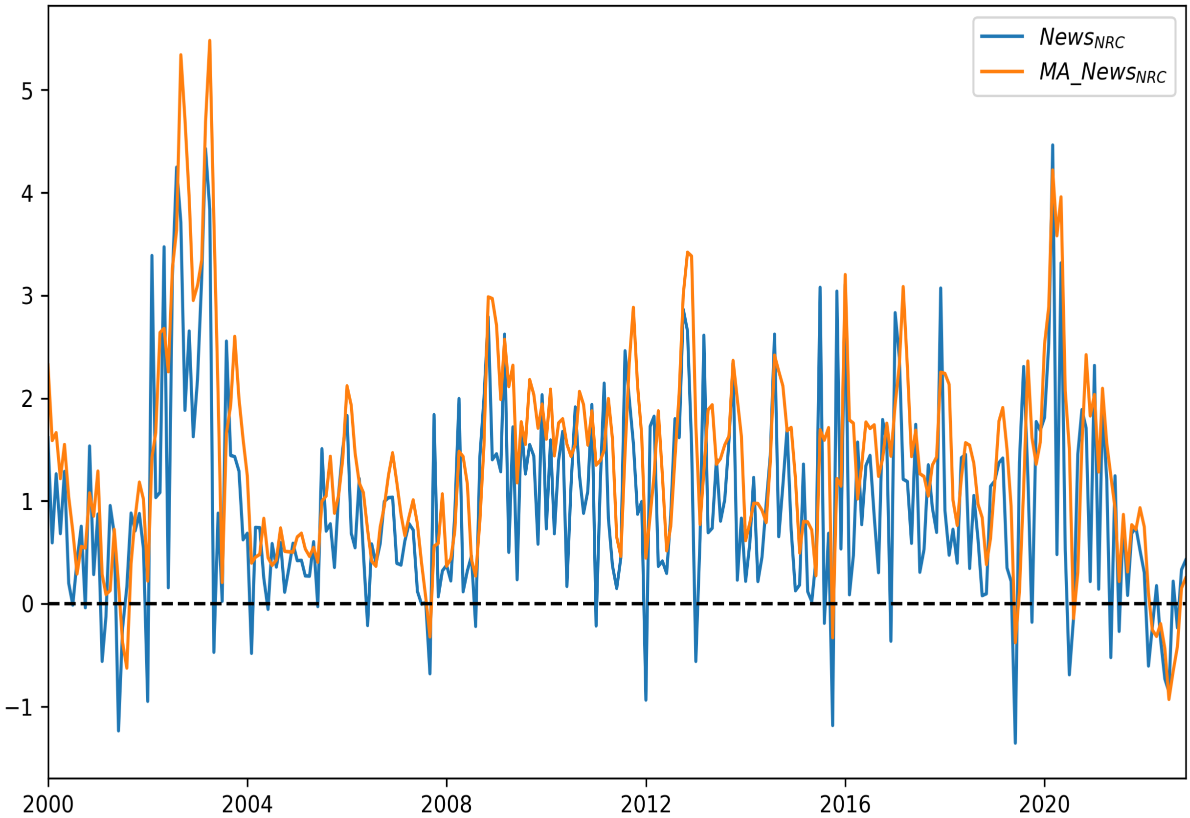

This figure shows the monthly timeseries of the $News_{NRC}$ index and the $MANews_{NRC}$ index from January 2000 to November 2022.

Figure 3.4 further illustrates that the $News_{NRC}$ index not only successfully captures the average stock-inflation relationship but is also closely aligned with the time-variation in asset correlations. To improve visual interpretability, Figure 3.4 plots a 12-month moving average version of our $News_{NRC}$ next to the realized correlation between stock-market returns and expected inflation estimated over rolling 1-year periods. The low-frequency components of our news index and asset-correlations seem strongly related. Our sample covers several episodes where stock-bond relations have been previously hard to connect to macroeconomic dynamics. Ermolov (2022) notes that the strong spike in stock-yield correlations around the dot-com bubble and the Enron scandal, as well as the persistent positive stock-yield correlations following the great recession, are difficult to understand with macroeconomic data. Our news measure offers a compelling explanation: these periods were marked by perceptions of procyclical inflation (and countercyclical deflation). Notably, the brief period of very high stock-inflation correlations in the early 2000s is strongly reflected in our news index. In contrast, the procyclical inflation implied by news in the post-recession period appears somewhat muted compared to the observed correlation data.

Table 3.5: Summary statistics

	$News_{NRC}$	$MANews_{NRC}$	$Corr(r_m, y^{\$})$	$Corr(r_m, y^{\pi})$
Mean	0.956	1.380	0.295	0.319
SD	1.000	1.000	0.342	0.406
p25	0.306	0.721	0.032	0.044
p50	0.742	1.349	0.278	0.352
p75	1.457	1.876	0.550	0.594
AR(1)	0.228	0.801	0.495	0.464

This table shows summary statistics for the variables used in our main analysis. We report the mean, standard deviation, quantiles and persistence of each variable.

To formally investigate the predictions 1 and 2 that the $News_{NRC}$ index captures time-variation in asset correlations, we run predictive regressions of monthly Fisher-transformed future realized correlations on the monthly $News_{NRC}$ index:

$$F\tilde{Corr}_{t+1} = \alpha + \beta News_{NRC,t} + \epsilon_{t+1} \tag{3.25}$$

where $F\tilde{Corr}_{t+1}$ denotes next-month Fisher transformed realized correlation. The Fisher transformation, given by

$$F\tilde{Corr}_t = \frac{1}{2} ln \left(\frac{1 + \tilde{Corr}_t}{1 - \tilde{Corr}_t} \right) \tag{3.26}$$

has the advantage to map the realized correlation coefficient from the range of $[-1, 1]$ to $[-\infty, \infty]$. The Fisher transformation of the realized correlation coefficient is known to quickly converge to a standard normal distribution thereby improving the finite sample properties of the realized correlation (see e.g. Anderson (1984) and Barndorff-Nielsen and Shephard (2004)).

Tables 3.6 and 3.7 document the regression results. Coefficients for $News_{NRC}$ in both the stock-inflation and the stock-yield correlations are positive and significant at the 5% level (column (1)). These findings provide direct evidence for an inflation cyclicality channel driving time variation in asset correlations.

Previous literature has identified several variables that may predict stock-bond correlations: the level of inflation (see e.g. Ilmanen (2003), Dergunov et al. (2023), David and Veronesi (2013)); inflation uncertainty (David and Veronesi (2008), Li et al. (2003)); stock-market uncertainty (Connolly et al. (2005, 2007)); and the correlation between stock-market uncertainty and stock-market returns (Bansal et al. (2014)). Some or all of these variables may also partially capture the cyclical properties of inflation, with the level of inflation likely beeing closely related to inflation cyclicality. To explore this, we run multivariate predictive regressions, including these

Figure 3.4: $News_{NRC}$ and realized stock-inflation correlation

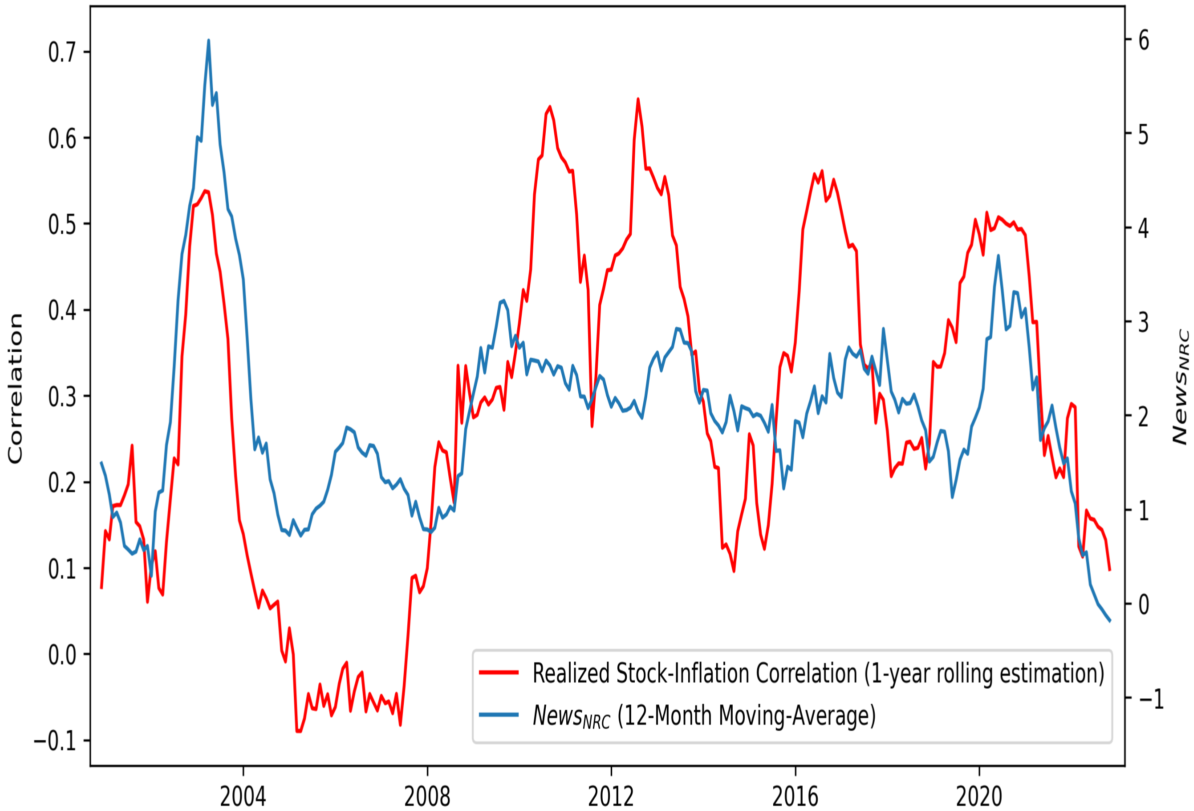

This figure plots the timeseries of the 12-month moving average of the $News_{NRC}$ index together with the 253-days (1 year) realized rolling correlation between stock-market returns and changes in expected inflation.

four variables as controls. Coefficients on the $News_{NRC}$ index remain positive and significant (column (6)). This indicates that news data can provide valuable insights into inflation cyclicality hard to capture with traditional macroeconomic data.

We further investigate to what extent the conditional relationship between persistent consumption growth and inflation could possibly also be measured with macroeconomic data. It is well known that econometrically detecting a small but highly persistent component in a time series is challenging. Nevertheless, we try to capture the conditional relation between inflation and long-run consumption growth from the realizations of the two variables. This approach to measure the nominal-real correlation has recently been applied by Boons et al. (2020). We estimate 60-month rolling correlations between future cumulative consumption growth and the level of inflation:

$$InfConsLT_{t,h} = Corr_{t-60,t}\left(\sum_{i=1}^{h} \Delta c_{t+i}, \pi_t\right) \tag{3.27}$$

In our model, this correlation is determined by the variable ρ_t. Therefore, $InfConsLT_{t,h}$ can be viewed as an alternative measure for ρ_t estimated from macroeconomic data instead of news data. We focus on the 60-month horizon (h=60) because our earlier results show that consumption growth predictability peaks around this horizon (see Table 3.3). Conclusions regarding our $News_{NRC}$ index do not depend on this choice. Column (2) of Tables 3.6 and 3.7 shows that there is a weakly positive but insignificant relationship between $InfConsLT_t$ and future realized correlations. Controlling for $InfConsLT_t$ only slightly reduces the coefficient estimates for our $News_{NRC}$ index, which remain positive and significant (column (3)). This provides further evidence that news data allows us to capture otherwise hard-to-detect conditional relations to a low-frequency component of consumption growth.

So far, we have focused on a conditional relation between the *long-run* components of inflation and consumption growth. Next, we analyze whether our results could also be driven by a contemporaneous relationship between inflation and consumption growth. To this end, we estimate rolling correlations between the level of inflation and contemporaneous consumption growth:

$$InfConsST_t = Corr_{t-60,t}\left(\Delta c_t, \pi_t\right) \tag{3.28}$$

The correlation between $InfConsST_t$ and $News_{NRC}$ is close to zero and even slightly negative at -0.08, suggesting that our news measure is mostly unrelated to correlations between *transitory* innovations of inflation and consumption growth. The relation between $InfConsST_t$ and future realized correlations is insignificant and even slightly negative for stock-yield correlations (column (4)). This documents that the contemporaneous inflation-consumption relation contributes little to the understanding of stock-inflation cyclicality and emphasizes the importance to focus on long-run dynamics. Including $InfConsST_t$ as control variable in a multivariate regression leaves the significance of the $News_{NRC}$ coefficients unchanged (column (5)).

Panel B of Tables 3.6 and 3.7 provides regression results for the low-frequency version of our news measure, the $MANews_{NRC}$ index. The results suggest that focusing on the lower-frequency component of our news measure successfully reduces measurement errors and improves explanatory power. The coefficients on the $MANews_{NRC}$ are positive and significant across all specifications. The improved measurement of news implied inflation cyclicality leads to an increase in R^2's in all cases. For the univariate regressions of stock-inflation (stock-yield) correlations, adjusted R^2's increase from 4.8% (5.4%) for the $News_{NRC}$ index to 9.6% for the $MANews_{NRC}$ index. The predictive power of the $MANews_{NRC}$ index is robust to the inclusion of the long-term consumption/inflation correlation, the contemporaneous consumption-inflation correlation, the level of inflation, inflation uncertainty, stock-market variance, and the stock-return/stock-uncertainty correlation.

Prediction 3 provides additional model implications relating to the forecast horizon of realized correlations. Specifically, we anticipate that the coefficient and predictive content of the state variable ρ_t will diminish for more distant realized correlations, as future states of ρ_t become increasingly uncertain and tend to revert to the mean over time. We empirically assess this prediction by regressing monthly Fisher transformed realized correlations for a range of prediction horizons on the $News_{NRC}$ index,

$$F\tilde{Corr}_{t+h} = \alpha + \beta News_{NRC,t} + \epsilon_{t+h} \quad (3.29)$$

We expect to see both coefficients and R^2's to decrease with horizon h if $News_{NRC}$ captures a stationary stance of inflation cyclicality. Table 3.8 and 3.9 display the results for prediction horizons from 1 to 6 months ahead. In all cases, coefficient estimates and R^2's clearly decrease with increasing prediction horizon. Notably, the coefficient estimates at the 6-month horizon are roughly 35%-50% of those at the 1-month horizon. For horizons of 4 to 5 months ahead, coefficients start to become insignificant.

Table 3.6: Predicting realized stock-inflation correlations

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: $News_{NRC}$ index							
Intercept	0.222*	0.310*	0.232*	0.304*	0.228*	0.040	0.054
	(5.521)	(5.149)	(3.485)	(8.083)	(5.725)	(0.316)	(0.330)
$News_{NRC}$	0.078*		0.075*		0.081*	0.046*	0.052*
	(4.336)		(3.343)		(4.233)	(2.733)	(2.731)
$InfConsLT$		0.148	0.126				-0.414
		(0.458)	(0.417)				(-1.425)
$InfConsST$				0.154	0.180		0.733*
				(1.157)	(1.502)		(2.143)
Other Controls						Yes	Yes
Adj. R^2	0.048	-0.002	0.036	0.008	0.060	0.122	0.159
N	275	232	232	275	275	275	232
Panel B: $MANews_{NRC}$ index							
Intercept	0.147*	0.310*	0.134	0.304*	0.152*	-0.056	-0.061
	(3.136)	(5.149)	(1.732)	(8.083)	(3.367)	(-0.427)	(-0.376)
$MANews_{NRC}$	0.108*		0.119*		0.111*	0.077*	0.093*
	(4.099)		(3.405)		(4.143)	(2.904)	(2.998)
$InfConsLT$		0.148	0.108				-0.422
		(0.458)	(0.374)				(-1.561)
$InfConsST$				0.154	0.185		0.750*
				(1.157)	(1.687)		(2.354)
Other Controls						Yes	Yes
Adj. R^2	0.096	-0.002	0.096	0.008	0.109	0.145	0.192
N	275	232	232	275	275	275	232

This table documents the predictive content of our inflation cyclicality news measures for future realized stock-market/expected inflation correlations. It shows the results from regressing next months Fisher transformed realized correlation between stock-market returns and changes in expected inflation on the $News_{NRC}$ index (Panel A) and the $MANews_{NRC}$ index (Panel B). Column (1) contains results for the univariate regression specification. Columns (3), (5), (6) and (7) report multivariate regression results for the specification $F\tilde{Corr}_{t+1} = \alpha + \beta News_{NRC,t} + Controls + \epsilon_{t+1}$. The control variables include the rolling correlation between inflation and future cumulative consumption growth ($InfConsLT$), the rolling correlation between contemporaneous consumption growth and inflation ($InfConsST$), and *other controls* comprises the level of inflation, inflation uncertainty, stock market uncertainty and the rolling correlation between stock market uncertainty and stock market returns. Columns (2) and (4) contains univariate regression results using $InfConsLT$ and $InfConsST$ as predictors for future realized correlations. We report intercepts, coefficients, t-statistics (in parentheses) and adjusted R^2s. Standard-errors are corrected for heteroscedasticity and serial correlation using the approach of Newey-West with one lag. * indicates significance at the 5%-level under the Newey-West adjusted standard errors.

Table 3.7: Predicting realized stock-yield correlations

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A: $News_{NRC}$ index							
Intercept	0.229*	0.377*	0.304*	0.307*	0.221*	-0.101	-0.221
	(4.396)	(5.251)	(3.654)	(7.178)	(4.497)	(-0.696)	(-0.950)
$News_{NRC}$	0.097*		0.071*		0.092*	0.072*	0.060*
	(3.831)		(2.547)		(3.846)	(3.594)	(2.635)
$InfConsLT$		0.187	0.166				-0.167
		(0.481)	(0.444)				(-0.452)
$InfConsST$				-0.279	-0.249		0.170
				(-1.655)	(-1.575)		(0.383)
Other Controls						Yes	Yes
Adj. R^2	0.054	-0.001	0.024	0.023	0.072	0.126	0.083
N	275	232	232	275	275	275	232
Panel B: $MANews_{NRC}$ index							
Intercept	0.146*	0.377*	0.211*	0.307*	0.138*	-0.232	-0.364
	(2.355)	(5.251)	(2.245)	(7.178)	(2.418)	(-1.639)	(-1.614)
$News_{NRC}$	0.127*		0.112*		0.124*	0.111*	0.112*
	(4.266)		(3.260)		(4.448)	(4.265)	(3.944)
$InfConsLT$		0.187	0.149				-0.177
		(0.481)	(0.413)				(-0.511)
$InfConsST$				-0.279	-0.245		0.193
				(-1.655)	(-1.661)		(0.449)
Other Controls						Yes	Yes
Adj. R^2	0.096	-0.001	0.066	0.023	0.114	0.156	0.123
N	275	232	232	275	275	275	232

This table documents the predictive content of our inflation cyclicality news measures for future realized stock-market/nominal yield correlations. It shows the results from regressing next months Fisher transformed realized correlation between stock-market returns and changes in nominal yields on the $News_{NRC}$ index (Panel A) and the $MANews_{NRC}$ index (Panel B). Column (1) contains results for the univariate regression specification. Columns (3), (5), (6) and (7) report multivariate regression results for the specification $F\tilde{Corr}_{t+1} = \alpha + \beta News_{NRC,t} + Controls + \epsilon_{t+1}$. The control variables include the rolling correlation between inflation and future cumulative consumption growth ($InfConsLT$), the rolling correlation between contemporaneous consumption growth and inflation ($InfConsST$), and *other controls* comprises the level of inflation, inflation uncertainty, stock market uncertainty and the rolling correlation between stock market uncertainty and stock market returns. Columns (2) and (4) contains univariate regression results using $InfConsLT$ and $InfConsST$ as predictors for future realized correlations. We report intercepts, coefficients, t-statistics (in parentheses) and adjusted R^2s. Standard-errors are corrected for heteroscedasticity and serial correlation using the approach of Newey-West with one lag. * indicates significance at the 5%-level under the Newey-West adjusted standard errors.

Table 3.8: Predicting realized stock-inflation correlations at different horizons

Horizon (months)	1	2	3	4	5	6
Panel A: $News_{NRC}$ index						
Intercept	0.222*	0.222*	0.227*	0.246*	0.268*	0.274*
	(5.521)	(5.315)	(5.303)	(5.858)	(5.808)	(6.307)
$News_{NRC}$	0.078*	0.078*	0.074*	0.054*	0.031	0.027
	(4.336)	(3.091)	(3.024)	(2.041)	(1.235)	(1.246)
Adj. R^2	0.048	0.048	0.044	0.021	0.005	0.003
N	275	275	274	273	272	271
Panel B: $MANews_{NRC}$ index						
Intercept	0.147*	0.162*	0.192*	0.223*	0.228*	0.226*
	(3.136)	(3.182)	(3.519)	(4.005)	(4.225)	(4.542)
$News_{NRC}$	0.108*	0.097*	0.077*	0.054	0.050	0.053
	(4.099)	(3.015)	(2.292)	(1.617)	(1.613)	(1.835)
Adj. R^2	0.096	0.078	0.047	0.021	0.018	0.019
N	275	275	274	273	272	271

This table documents the predictive content of our inflation cyclicality news measures for future realized stock-market/expected inflation correlations at prediction horizons from 1 to 6 months ahead. It shows the results from regressing future n-month ahead Fisher transformed realized correlation between stock-market returns and changes in expected inflation on the $News_{NRC}$ index (Panel A) and the $MANews_{NRC}$ index (Panel B). Each column reports the regression results for a specific prediction horizon. We report intercepts, coefficients, t-statistics (in parentheses) and adjusted R^2s. Standard-errors are corrected for heteroscedasticity and serial correlation using the approach of Newey-West with one lag. * indicates significance at the 5%-level under the Newey-West adjusted standard errors.

Table 3.9: Predicting realized stock-yield correlations at different horizons

Horizon (months)	1	2	3	4	5	6
Panel A: $News_{NRC}$ index						
Intercept	0.229*	0.237*	0.237*	0.286*	0.271*	0.285*
	(4.396)	(4.327)	(4.481)	(4.849)	(4.688)	(5.077)
$News_{NRC}$	0.097*	0.087*	0.088*	0.037	0.053	0.041
	(3.831)	(2.693)	(3.103)	(0.994)	(1.659)	(1.300)
Adj. R^2	0.054	0.042	0.044	0.004	0.014	0.006
N	275	275	274	273	272	271
Panel B: $MANews_{NRC}$ index						
Intercept	0.146*	0.182*	0.203*	0.235*	0.221*	0.249*
	(2.355)	(2.610)	(2.788)	(3.057)	(3.100)	(3.504)
$News_{NRC}$	0.127*	0.100*	0.086*	0.062	0.073	0.054
	(4.266)	(2.625)	(2.131)	(1.455)	(1.949)	(1.480)
Adj. R^2	0.096	0.057	0.041	0.020	0.028	0.013
N	275	275	274	273	272	271

This table documents the predictive content of our inflation cyclicality news measures for future realized stock-market/nominal yield correlations at prediction horizons from 1 to 6 months ahead. It shows the results from regressing future n-month ahead Fisher transformed realized correlation between stock-market returns and changes in nominal yields on the $News_{NRC}$ index (Panel A) and the $MANews_{NRC}$ index (Panel B). Each column reports the regression results for a specific prediction horizon. We report intercepts, coefficients, t-statistics (in parentheses) and adjusted R^2s. Standard-errors are corrected for heteroscedasticity and serial correlation using the approach of Newey-West with one lag. * indicates significance at the 5%-level under the Newey-West adjusted standard errors.

3.7 Conclusion

We show that news reporting on inflation captures a cyclical relationship between persistent consumption growth and long-run inflation expectations. Our results provide direct evidence that the changing cyclical properties of inflation drive the co-movement between stock returns, bond yields and expected inflation.

Our novel approach combines the text sentiment of news articles with the classification of their content into “rising” and “falling” inflation news. We find that news text sentiment proxies for long-run consumption growth. Aggregate text sentiment predicts future consumption growth at long horizons, with predictability increasing over the horizon. The content of news articles (“rising” vs. “falling” inflation) proxies for long-run inflation expectations. Aggregate directional news coverage (coverage of “rising” relative to “falling” inflation) is significantly positively related to measures of long-run inflation expectations. We combine these two elements into our measure of news implied inflation cyclicality – the $News_{NRC}$ index. The $News_{NRC}$ index is calculated by comparing the average text sentiment of “rising” vs. “falling” inflation articles. It measures how investors currently view changes in inflation expectations to be related to the long-run consumption outlook. We find that the $News_{NRC}$ index significantly positively predicts future realized asset correlations consistent with an equilibrium long-run risk model.

Overall, our results suggest that asset co-movements are in part driven by investors’ views about the *long-run* relationship between inflation and consumption growth. This chapter provides an approach to empirically link unobservable macroeconomic dynamics and asset prices using news data.

Chapter 4

TIPS under FIRE - Formation of inflation expectations in bond markets

4.1 Introduction

Market participants' expectations of future inflation developments are a crucial determinant of asset prices and shape the term structure of nominal interest rates. Understanding the nature of inflation belief formation embedded in asset prices therefore has central importance. Since the "rational-expectations revolution" (Muth (1960, 1961), Lucas (1972, 1973, 1976)), the bulk of the finance and economics literature relied on "model-consistent"[1] *full information rational expectations* (FIRE) as workhouse framework to understand expectation formation. Recently though, mounting evidence documents that actual expectations of market participants reflected in asset prices exhibit stark deviations from this benchmark model.[2] The literature on inflation expectation formation currently leaves open two crucial questions: 1) How does the *marginal* investor – as opposed to e.g. the average survey respondent - form inflation expectations? 2) How do investors form *long-term* inflation expectations as opposed to short-term expectations for a few quarters ahead? Addressing these questions is quintessential to assess how the formation of inflation expectations impacts the term structure of nominal interest rates. This chapter aims to fill this gap and empirically studies inflation expectation formation of the marginal bond market investor across a wide range of expectation horizons. Our results reveal a maturity pattern of horizon-increasing overreaction in violation of FIRE in bond market inflation expectations. This offers an inflation-expectations rooted explanation for the puzzling excess-volatility of long-term nominal interest rates and their overshooting reaction to macroeconomic news.

Existing evidence is limited to *short-horizon* inflation forecasts from *survey* data. This evidence insufficiently addresses the question of how the *marginal* bond market investor forms inflation expectations across the term structure up to *long-horizon* projections. In order to directly analyze

[1] "Model-consistent" refers to expectations being consistent with the assumed underlying inflation process.
[2] See e.g. Greenwood and Shleifer (2014), Cieslak (2018), Schmeling et al. (2022), d'Arienzo (2020), Wang (2021), Boeck and Feldkircher (2021), Bordalo et al. (2019), Bouchaud et al. (2019), Lochstoer and Muir (2022), Bordalo et al. (2018), Mokanov (2023), Valente et al. (2022).

this question, we extend the existing evidence towards two crucial dimensions: First, we consider expectations implied by market prices of inflation claims in addition to survey data. While survey data is useful as it essentially makes expectations directly observable, one must be very careful when drawing conclusions about asset prices. Expectational biases in surveys may not be reflected in market prices e.g. because more rational or better-informed agents actively trade until any impact of the biased average respondent is eliminated. In fact, the existing literature documents a discrepancy between *consensus* expectations, which have been shown to underreact to new information, and the expectations of *individual* forecasters, which have been shown to overreact to new information.[3] The implications for the expectation of the *marginal* investor remain completely uncharted from this survey evidence alone. Second, we evaluate inflation expectations for a wide range of forecast horizons including long-term projections of up to 10 years. Prior adjacent theoretical and empirical results indicate that there are distinct patterns of expectation formation depending on the forecast horizon.[4] To understand expectational effects along the full inflation term structure one therefore needs to gather empirical evidence for a spectrum of maturities and cannot just extrapolate from short-horizon observations. However, existing empirical evidence regarding horizon specific effects in inflation expectation formation is limited to the very short-end of expectation horizons (see Pang (2023)). Our work is the first to directly analyze expectation formation in both short-term and long-term inflation forecasts from both survey and market price data. We are therefore in the unique position to draw direct conclusions about bond market expectation formation embedded in the term structure of market inflation rates as well as the term structure of nominal interest rates.

We organize our empirical analysis around a simple reduced-form model of expectation formation. This framework nests FIRE and the leading departures from it: *underreaction* of agents to new information, as in models with information frictions such as (rational) inattention (Mankiw and Reis (2002), Woodford (2003), Sims (2003), Carroll (2003), Gabaix (2014)); and *overreaction*, as in models of diagnostic expectations (Bordalo et al. (2018, 2019, 2020)) or over-extrapolation (Angeletos et al. (2021), Barberis et al. (2018)). The model yields empirical testable implications for these theories, which we can compare against expectations data.

Our research objective brings along two main empirical challenges: First, while expectations extracted from market prices directly reflect the beliefs of the marginal investor, they may also contain a compensation for risk, severely complicating inference about expectational patterns. Second, prevalent tests of FIRE based on the predictability of forecast errors face restraining data limitations when it comes to long-horizon forecasts. We develop multiple strategies to address

[3] See e.g. Bordalo et al. (2020).

[4] See Pang (2023) and Angeletos et al. (2021) for inflation expectations, Bordalo et al. (2019) for earnings growth expectations, Wang (2021) and d'Arienzo (2020) for expectations of nominal interest rates, Valente et al. (2022) for exchange rate expectations and Lochstoer and Muir (2022) for expectations of stock-market volatility.

these challenges. Most importantly, we exploit multiple testable model predictions pertaining to deviations from FIRE by overreaction or underreaction. Previous literature has overwhelmingly focused on error-on-revision regressions. We show that our benchmark expectation model additionally yields implications that concern the autocorrelation of forecast errors and the autocorrelation of forecast revisions. Under plausible assumptions regarding the process governing the dynamics of risk premia these model implications each yield differentiating predictions for the impact of time-varying risk premia and specific deviations from FIRE. Together, these model implications jointly allow to empirically separate expectational effects from risk premia. Moreover, exploiting model implications for the autocorrelation of forecast *revisions* evades the severe data-limitation problems that practically prevent the evaluation of long-term forecasts based on forecast errors. Forecast revisions for long-term forecasts are available at the same frequency as for short-term forecasts. Independent observations of the forecast error on the other hand become increasingly sparse for longer horizons which is partly why the literature so far mainly focused on the evaluation of short-term forecast errors.

Our results point to clear deviations from FIRE in bond market inflation expectations, revealing distinctive horizon-dependent patterns along the term structure of expectations. Confirming previous results in the literature, we provide additional evidence for underreaction of consensus survey inflation expectations and overreaction of individual survey expectations to new information at very short horizons. Despite this discrepancy at short horizons, our examination of longer-term expectations reveals that both consensus and individual expectations each increasingly exhibit overreaction with extending forecast horizon. At medium-to-long-term horizons both consensus and individual expectations significantly overreact. This pattern of maturity increasing overreaction shows up remarkably similar across survey inflation expectations and market-implied inflation expectations. Maturity increasing overreaction hence seems to be a robust feature of bond market expectations.

Our empirical results are very difficult to reconcile with an alternative explanation that relies on time-varying risk premia. First, results for market-implied inflation expectations mirror results for survey expectations which do not contain a risk premium component. Second, under plausible assumptions regarding the evolution of risk premia, our model allows us to unambiguously link results to expectational effects instead of risk premia. Third, by relaxing our initial model assumptions, we demonstrate that an explanation via time-varying risk premia would require highly implausible time-series properties of inflation risk premia drastically deviating from what is commonly assumed and estimated in the literature. Essentially, to quantitatively match our results without resorting to deviations from FIRE one needs to assume an extremely high mean-reversion if not a negative autocorrelation of the risk premium component. Fourth, we additionally employ two independent empirical strategies to separate the expectations component from the risk premium component in market-implied inflation expectations. Both approaches suggest that our results are rooted in expectational patterns.

Our first approach to empirically isolate expectational biases in market-implied inflation is to study an estimate of the risk premium component. We construct a measure for the risk premium component by regressing the returns of the inflation contract on a wide range of variables suggested in the literature. We then use this measure to a) verify model assumptions regarding the evolution of risk premia and b) isolate expectational effects in our main regressions by including it as a control variable. Regarding a), we find the time-series properties of our risk premium estimate to be consistent with our modelling assumptions, but inconsistent with the properties theoretically required for a pure risk premium explanation of our main results. Mainly, our empirical risk premium estimate appears far too persistent as to explain our main regression results without resorting to overreaction. We confirm this conclusion as we assess b) and include our risk premium estimate as control variable in our main regressions. After isolating the expectations channel with this approach, horizon increasing overreaction becomes even more visible in the results. Additionally, the estimated impact of the risk premium component aligns encouragingly well with our theoretical predictions.

Our second approach to empirically single out expectational effects aims to directly identify shocks to market implied inflation (almost) exclusively driven by the expectation component. To this end, we rely on a high-frequency identification around scheduled releases of the consumer price index (CPI). We work with the assumption that short-term movements of market-implied inflation around CPI releases are primarily driven by changes in market expectations in reaction to new inflation information. Using these high-frequency expectation shocks, we exploit our model predictions concerning subsequent expectation revisions: underreaction implies continuing adjustment of expectations in the same direction as the initial shock while overreaction implies subsequent partial reversal. We observe that for medium-and long-term forecast horizons, a substantial fraction of around 40% of the initial expectation shock is subsequently reversed over three weeks following the CPI release. Short-horizon forecasts however display no significant pattern following CPI releases. These results indicate that markets overreact to CPI release information and increasingly do so for longer horizons.

Maturity increasing overreaction of bond market inflation expectations speaks to prominent puzzles in the bond market literature. Long-term nominal interest rates are known to be excessively sensitive to macroeconomic news and short-rate variations (Gürkaynak et al. (2005), Hanson et al. (2018), Brooks et al. (2018)). This high sensitivity is at odds with the observed weak persistence of shocks. Relatedly, long-term interest rates, as well as implied inflation rates, are excessively volatile relative to short-term interest rates. Giglio and Kelly (2018) show that relative variances along the term structure are inconsistent with a broad class of mainstream rational term structure models even after allowing for considerable discount rate variation. Our results reveal over- and underreaction along the term structure of inflation expectations as an explanation for excess sensitivity at long maturities.

This chapters approach to explore inflation expectation formation with both survey and market-price evidence is new to the literature and combines the strengths of both sources of expectation data. Inflation survey data has been very popular in this regard as it essentially makes expectations directly observable and is readily available. However, survey data has a few shortcomings when it comes to assessing the expectation of the marginal investor across the term structure. First off, survey respondents may not necessarily fully “act on their beliefs” and incentives to provide accurate forecasts can be weak. Market participants on the other hand have “skin in the game” and have a strong direct incentive to provide the most accurate forecast possible. Moreover, market trading allows more informed or more rational agents to capitalize on their competitive advantage. Averaging over survey responses is therefore potentially very different from the equilibrium result generated from market arbitrage mechanisms. Also, survey responses potentially somewhat depend on the exact interpretation of the survey questions (e.g. regarding compounding convention of growth rates) while the payoff structure of a financial inflation contract is completely unambiguous. Second, survey forecast data is only available at rather low frequencies and for only a few select horizons along the term structure. Market implied inflation on the other hand can be extracted at very high frequencies for a wide range of maturities. Market data is therefore extraordinarily well suited to analyze term structure effects of expectation formation. While market-implied expectations data addresses many of the shortcomings of survey data, it has the disadvantage of potential contamination with a time-varying investor compensation for risk. This chapter dedicates significant attention to addressing this issue to effectively complement survey evidence with insights from market prices.

Related Literature Our research contributes to the growing body of literature exploring the role of expectation formation and potential deviations from FIRE for the determination of asset prices. Greenwood and Shleifer (2014) document that investors unduly extrapolate recent past returns when forming expectations about future stock-market returns. Bouchaud et al. (2019) relate the profitability anomaly to investor underreaction to earnings news. Bordalo et al. (2019) provide evidence for diagnostic expectations in analysts’ and investors’ earnings-growth projections. Lochstoer and Muir (2022) find that market participants have slow moving beliefs about stock-market volatility, leading to initial underreaction and delayed overreaction to volatility shocks. Bordalo et al. (2018) show that boom-bust credit cycles can be explained by investor overreaction driven by diagnostic expectations. Analyzing expectations of future nominal short rates, Cieslak (2018) finds that investors overextrapolate the current short rate while ignoring predictive information from macroeconomic variables. Our work on bond market inflation expectations complements this literature by providing evidence on how market expectations are formed around one of the central macroeconomic factors.

Our research inquiry is motivated by an expansive macroeconomic literature documenting various deviations from FIRE in inflation survey expectations. Coibion and Gorodnichenko (2012, 2015) document that aggregate (short-horizon) survey forecasts of inflation underreact to new information, consistent with models of information rigidities. Kohlhas and Walther (2021) show that aggregate survey inflation forecasts simultaneously underreact to new information and overextrapolate from current conditions. Using the cross-sectional dimension of survey data, Bordalo et al. (2020) show that *individual* forecasters overreact to new information while *aggregate* forecasts underreact. Evaluating the impulse-response functions of aggregate inflation survey forecasts to inflation shocks, Angeletos et al. (2021) identify a pattern of initial underreaction followed by delayed over-shooting. Boeck and Feldkircher (2021) confirm this combination of under- and overreaction in surveys of nominal yields in response to monetary policy shocks. Our contribution to this literature is twofold: First, we convey these questions to an asset pricing context. Hence, we go beyond survey evidence and explore implications for market quantities. Second, we collect evidence across a wide range of forecast horizons

The horizon dimension of expectation formation gained some attention recently. d'Arienzo (2020) and Wang (2021) use short-term forecasts of *nominal* yields of different maturities to argue for maturity increasing overreaction in the nominal term structure. They attribute this to behavioral biases in expectations of future short rates or term premia dependent on the forecast horizon. Relative to these approaches, we directly look at expectations over different forecast horizons. Our results point to the peculiarities of bond market *inflation* expectations as a more fundamental cause. Lochstoer and Muir (2022) detect that short-horizon stock-market volatility expectations underreact to new information while long-horizon expectations overreact as they overestimate the persistence of volatility shocks. Similarly to our research, Pang (2023) also addresses the question of cross-horizon effects in inflation expectations. However, empirical evidence in Pang (2023) is limited to very short horizon (0-3 quarters) survey data while implications for longer horizons are extrapolated from a theoretical model. Moreover, while Pang (2023) finds strong discrepancies between individual and consensus survey forecasts, he leaves the unapparent implications for the *marginal* investor unexplored. This work is the first to explore expectation formation in medium- and long-horizon *market* inflation expectations empirically.

Our research also connects to the literature studying the role of risk premia, liquidity and market inefficiencies versus inflation expectations in TIPS markets. Examples include Grishchenko and Huang (2013), Grishchenko et al. (2016), Christensen et al. (2010), Chernov and Mueller (2012), Ang et al. (2008), Bahaj et al. (2023), Andreasen et al. (2021), Pflueger and Viceira (2016, 2011), Abrahams et al. (2016), Gürkaynak et al. (2010), Fleckenstein et al. (2014), d'Amico et al. (2018), Bauer (2015). Our work is the first to explicitly study the role of expectational biases in this market.

4.2 Assessing the formation process of inflation expectations

The workhouse model to understand expectation formation in macroeconomics and finance is the theory of full-information rational expectations (FIRE). It posits that agents fully know the true data-generating process including its governing parameters. Moreover, it is assumed that agents have all relevant (past and present) information when forming their forecasts. A key implication of FIRE is that future realized forecast errors should be unpredictable given the real-time available set of information. This implication can be used to empirically assess whether actual forecasts of market participants are consistent with FIRE. A major difficulty in empirical implementation is to identify the real-time information set without unduly imposing the ex-post knowledge of the econometrician. In this regard we adopt the popular approach in the literature (e.g. Coibion and Gorodnichenko (2015)) and rely on the information contained in the real-time forecaster's own forecasts. Any information contained in these forecasts are necessarily a sub-set of the full real-time information set.

To guide our empirical analysis, we rely on a simple reduced form model of expectation formation which we adopt from the literature:[5]

$$F_t(\pi_{t+h}) = (1-g)E_t(\pi_{t+h}) + gF_{t-1}(\pi_{t+h}) \tag{4.1}$$

Agents' forecasts $F_t(\pi_{t+h})$ in period t for realized inflation π_{t+h} in period $t+h$ are a weighted average of the current objective forecast $E_t(\pi_{t+h})$ comprising all available information, and the agents' own previous-period forecast $F_{t-1}(\pi_{t+h})$. FIRE forecasts are nested for the case $g = 0$ when agents' forecast equal the objective forecast $E_t(\pi_{t+h})$. The model allows us to capture two possible deviations from FIRE. For $g > 0$, agents are slow to react to new information, hence displaying *underreaction*. Coibion and Gorodnichenko (2015) show that this aggregate behavior could for example arise from information frictions. For $g < 0$, agents put too much weight on recent information, i.e. they overreact. *Overreaction* can result from a range of deviations from rationality such as diagnostic expectations (see e.g. Bordalo et al. (2018, 2019, 2020)) or over-extrapolation (see e.g. Angeletos et al. (2021) and Barberis et al. (2018)).
We will use this expectation formation model to derive empirically testable implications which we will bring to survey and market data. The advantage to evaluate a specific expectation model to test FIRE are threefold. First, it disciplines our empirical analysis to regressions directly specified by theory. Second, it allows us to not only evaluate the Null of the FIRE hypothesis but also reveals specific deviations from FIRE. Third, we can compare our results to the literature, which extensively studied models like (4.1).

[5] See e.g. Coibion and Gorodnichenko (2012, 2015), Bouchaud et al. (2019), Lochstoer and Muir (2022).

4.2.1 Error vs. revisions regression

A popular implication of model (4.1), as derived by Coibion and Gorodnichenko (2015), is a relationship between forecast revisions and forecast errors if forecasts deviate from FIRE. In the model, this relationship is given by

$$Err_{t+h,t} = \frac{g}{1-g}\Delta F_t(\pi_{t+h}) + \epsilon_{t+h,t} \tag{4.2}$$

where

$$\Delta F_t(\pi_{t+h}) = F_t(\pi_{t+h}) - F_{t-1}(\pi_{t+h}) \tag{4.3}$$

are forecast revisions,

$$Err_{t+h,t} = \pi_{t+h} - F_t(\pi_{t+h}) \tag{4.4}$$

are agents' forecast errors and

$$\epsilon_{t+h,t} = \pi_{t+h} - E_t(\pi_{t+h}) \tag{4.5}$$

is the unforecastable component of realized inflation. With survey data, both the error and the revision variable are observable so that one can estimate $g/(1-g)$ by running the regression

$$Err_{t+h,t} = \alpha + \beta * \Delta F_t(\pi_{t+h}) + u_{t+h,t} \tag{4.6}$$

Only $\beta = 0$ is consistent with FIRE. Significantly positive β signal underreaction, significantly negative β point to overreaction. This regression has been extensively applied to survey data in the literature. At short horizons, there is ample evidence for $\beta > 0$ (underreaction) for *aggregate* survey expectations and for $\beta < 0$ (overreaction) for *individual* survey expectations. Our goal is to extend this analysis along two dimensions: (1) Gather similar evidence across a range of forecast horizons to draw conclusions about term structure effects and (2) confront this analysis with bond market inflation expectations implied by market prices to learn about expectation formation of the *marginal* investor.

So far, we naturally referred to agents' expectations under real-world ("physical") probabilities. Working with market prices, we do not directly observe agents' physical inflation expectations. We only observe "risk-neutral" expectations, which may contain a risk premium:

$$F_t^{\mathbb{Q}}(\pi_{t+h}) = F_t^{\mathbb{P}}(\pi_{t+h}) + RP_t \tag{4.7}$$

$F_t^{\mathbb{Q}}(\pi_{t+h})$ and $F_t^{\mathbb{P}}(\pi_{t+h})$ indicate forecasts formed under the risk-neutral respectively the physical measure and RP_t is the risk premium component.

How do risk premia in our expectations data interfere with the theoretical relationship between forecast revisions and forecast errors? From Equation (4.1), which describes $\mathbb{P}$-expectations, adding RP_t gives the expressions for risk-neutral forecasts, forecast errors and revisions in our model:

$$F_t^{\mathbb{Q}}(\pi_{t+h}) = F_t^{\mathbb{P}}(\pi_{t+h}) + RP_t = (1-g)E_t(\pi_{t+h}) + gF_{t-1}^{\mathbb{P}}(\pi_{t+h}) + RP_t \tag{4.8}$$

$$Err_{t+h,t}^{\mathbb{Q}} \equiv \pi_{t+h} - F_t^{\mathbb{Q}}(\pi_{t+h}) = \pi_{t+h} - F_t^{\mathbb{P}}(\pi_{t+h}) - RP_t \tag{4.9}$$

$$\Delta F_t^{\mathbb{Q}}(\pi_{t+h}) = \Delta F_t^{\mathbb{P}}(\pi_{t+h}) + \Delta RP_t \tag{4.10}$$

If the risk premium component is constant, it does not affect the estimate of β of forecast errors on forecast revisions. To evaluate the impact of time-varying risk premia, let them evolve according to

$$RP_t = \theta RP_{t-1} + z_t \tag{4.11}$$

We can now calculate the theoretical β coefficient of risk neutral forecast errors on forecast revisions implied by the model. The regression coefficient β is defined as

$$\beta(Err_{t+h,t}^{\mathbb{Q}}, \Delta F_t^{\mathbb{Q}}(\pi_{t+h})) = \frac{Cov(Err_{t+h,t}^{\mathbb{Q}}, \Delta F_t^{\mathbb{Q}}(\pi_{t+h}))}{\sigma^2_{\Delta F^{\mathbb{Q}}}} \tag{4.12}$$

where $\sigma^2_{\Delta F^{\mathbb{Q}}}$ denotes the variance of $\Delta F_t^{\mathbb{Q}}(\pi_{t+h})$. We derive an expression for $Cov(Err_{t+h,t}^{\mathbb{Q}}, \Delta F_t^{\mathbb{Q}}(\pi_{t+h}))$ in detail in Appendix C.1.1. Central to this derivation is that risk neutral forecast errors and revisions can be represented as the sum of past underlying shocks to objective expectations ($\Delta E_t(\pi_{t+h})$) and the risk premia component (z_t) under our model assumptions. These shocks are by definition serially uncorrelated and unpredictable from past information. Therefore, we only have to evaluate the sum of the covariances and variances of the contemporaneous shocks to arrive at the following expression for β:

$$\begin{aligned}\beta(Err_{t+h,t}^{\mathbb{Q}}, \Delta F_t^{\mathbb{Q}}(\pi_{t+h})) &= \frac{1}{2}\left[\left(\frac{1+g}{1-g}\right)\frac{\sigma^2_{\Delta F^{\mathbb{P}}}}{\sigma^2_{\Delta F^{\mathbb{Q}}}}\right] - \frac{1}{2} \\ &= \frac{1}{2}\left[\left(\frac{1+g}{1-g}\right)\left(1 - \frac{\sigma^2_{\Delta RP}}{\sigma^2_{\Delta F^{\mathbb{Q}}}}\right)\right] - \frac{1}{2}\end{aligned} \tag{4.13}$$

where $\sigma^2_{\Delta F^{\mathbb{P}}}$ and $\sigma^2_{\Delta RP}$ denote the variance of $\Delta F_t^{\mathbb{P}}(\pi_{t+h})$ and ΔRP_t respectively. The relationship still depends on g, but variable risk premia add a negative effect to the link between forecast errors and forecast revisions. This happens because some of the changes in forecast revisions are caused by risk premia. When risk premia decrease, they lead to higher forecast errors, and when they increase, they lower forecast errors. If risk premia change very little compared to the objective forecast ($\sigma^2_{\Delta F^{\mathbb{Q}}} \approx \sigma^2_{\Delta F^{\mathbb{P}}}$), the sign of β mainly reflects the sign of g. When risk premia are constant ($\sigma^2_{\Delta F^{\mathbb{Q}}} = \sigma^2_{\Delta F^{\mathbb{P}}}$), the relationship simplifies to Equation (4.2). However, if risk premia vary significantly ($\sigma^2_{\Delta F^{\mathbb{Q}}} \gg \sigma^2_{\Delta F^{\mathbb{P}}}$), the negative impact of risk premia determines the sign of β. In this case, if we find a negative β in our empirical analysis, we cannot tell whether this is due to market overreaction or the high variability of risk premia. On the other hand, if we find a positive β, we can conclude that market expectations exhibit underreaction strong enough to outweigh the effect of variable risk premia. Therefore, the error-on-revision regression, when applied to market-implied forecasts, is only capable of identifying (strong) underreaction. To address this limitation, we develop strategies that reveal both overreaction and underreaction, even when the expectation data potentially includes time-varying risk premia. As we will demonstrate in Section 4.2.2 and Section 4.2.3, our expectation model provides additional testable predictions that make it possible to clearly distinguish between overreaction, underreaction, and time-varying risk premia.

An alternative strategy to the model-based distinction, laid out in Sections 4.2.1 –4.2.3, would be to account for an empirical estimate of RP_t in our regressions or trying to directly capture only the variation in forecasts driven by the expectation component $F_t^{\mathbb{P}}(\pi_{t+h})$. We will consider these approaches as a robustness check in Section 4.6 and Section 4.7.

4.2.2 Error autocorrelation

Next to the error-revision correlation discussed in the previous section, model (4.1) proposes another special case of the general test of FIRE concerning the predictability of forecast errors: it implies autocorrelation of agents' forecast errors if these forecasts deviate from FIRE. This is mainly because current period surprises to objective forecasts enter next period subjective forecast errors with the same sign (underreaction) respectively the opposite sign (overreaction) if the objective forecasts across forecast horizons are positively correlated ($Corr(\Delta E_t(\pi_{t+j}), \Delta E_t(\pi_{t+k})) > 0$).

To give an example: if there is a positive surprise in current quarter realized inflation relative to the objective forecast for the same quarter ($\epsilon_{t,t-1} > 0$ in Equation (4.5)), the subjective forecast error for this quarter will also tend to be positive ($Err^{\mathbb{P}}_{t,t-1} = \epsilon_{t,t-1} + E_{t-1}(\pi_t) - F^{\mathbb{P}}_{t-1}(\pi_t) > 0$). This upward surprise of the current quarter objective forecast ($E_t(\pi_t) - E_{t-1}(\pi_t) = \pi_t - E_{t-1}(\pi_t) > 0$) may be correlated with the objective forecast revision for next quarter inflation ($\Delta E_t(\pi_{t+1})$). If this cross-horizon correlation is positive, the revision of the objective forecast for next quarter inflation will also tend to be positive in this example ($\Delta E_t(\pi_{t+1}) > 0$). If agents underreact to this new information, the subjective forecast for next quarter inflation ($F_t^{\mathbb{P}}(\pi_{t+1})$) will only partially reflect the increase in objective expectations according to Equation

(4.1). The subjective forecast for next quarter inflation will then be below the objective forecast ($F_t^{\mathbb{P}}(\pi_{t+1}) < E_t(\pi_{t+1})$) and therefore likely also undershoot realization ($Err_{t+1,t}^{\mathbb{P}} = \epsilon_{t+1,t} + E_t(\pi_{t+1}) - F_t^{\mathbb{P}}(\pi_{t+1}) > 0$). This results in a positive relationship between the forecast errors of subsequent quarters.
To see this more generally, consider subsequent forecast errors over forecast horizon h

$$Err_{t+h,t}^{\mathbb{P}} = \frac{g}{1-g}\Delta F_t^{\mathbb{P}}(\pi_{t+h}) + \epsilon_{t+h,t} \tag{4.14}$$

$$Err_{t,t-h}^{\mathbb{P}} = \frac{g}{1-g}\Delta F_{t-h}^{\mathbb{P}}(\pi_t) + \epsilon_{t,t-h} \tag{4.15}$$

With these expressions, it is straightforward to calculate the theoretical β coefficient implied by the model as we show in Appendix C.1.1.2. Essentially, subjective forecast errors can be written as a sum of innovations to the objective forecasts for subsequent horizons. Evaluating the sum of covariances in terms of the cross-horizon correlation gives, for one-step-ahead[6] forecast errors, the following relationship:

$$\beta(Err_{t+1,t}^{\mathbb{P}}, Err_{t,t-1}^{\mathbb{P}}) = g * \rho_{\Delta E} \tag{4.16}$$

where $\rho_{\Delta E}$ denotes the correlation between one-step-ahead and two-step-ahead objective forecasts ($Corr(\Delta E_t(\pi_{t+1}), \Delta E_t(\pi_{t+2}))$. In practice, this correlation can safely be assumed to be clearly positive.[7]

We can therefore estimate the sign of g on survey data using the regression

$$Err_{t+h,t} = \alpha + \beta * Err_{t,t-h} + u_{t+h,t} \tag{4.17}$$

Again, under the Null of FIRE, $\beta = 0$. $\beta > 0$ indicates underreaction and $\beta < 0$ indicates overreaction. Hence, this implication allows us to gather additional evidence for model (4.1), complementing the results from error-vs-revision regressions found in the literature. More importantly, model implied error-autocorrelation yields new predictions regarding time-varying risk premia. If the expectations data contains a risk premium component, which is again assumed to follow (4.11), β is given by:[8]

[6] See Appendix C.1.1.2 for a general h-step-ahead expression.
[7] This is similar to assuming a dominant level effect in the term structure of inflation expectations.
[8] For the remainder of Section 4.2 and Section 4.4, we assume risk premium shocks and expectation shocks to be uncorrelated ($\rho(z, \Delta E) = 0$). We will relax this assumption in Section 4.5.

$$\begin{aligned}\beta(Err^{\mathbb{Q}}_{t+1,t}, Err^{\mathbb{Q}}_{t,t-1}) =& g\rho_{\Delta E}\left[1-\left(\frac{1}{2}\frac{1}{1-\theta}\right)\left(\frac{\sigma^2_{\Delta RP}}{\frac{1}{1-g^2}\sigma^2_{\Delta E}+0.5\frac{1}{1-\theta}\sigma^2_{\Delta RP}}\right)\right] \\ &+\theta\left[\left(\frac{1}{2}\frac{1}{1-\theta}\right)\left(\frac{\sigma^2_{\Delta RP}}{\frac{1}{1-g^2}\sigma^2_{\Delta E}+0.5\frac{1}{1-\theta}\sigma^2_{\Delta RP}}\right)\right]\end{aligned} \tag{4.18}$$

where $\sigma^2_{\Delta E}$ is the variance of objective expectations $\Delta E_t(\pi_{t+h}) = E_t(\pi_{t+h}) - E_{t-1}(\pi_{t+h})$. See Appendix C.1.1.2 for a detailed derivation of this result. The first term shows that g now has a dampened impact on β the more risk premia vary relative to the expectation component. This is because some of the variation in forecast errors is now caused by risk premia, blurring the effects from the expectation component. The direction of the impact that g exerts on β, however, is unchanged. The second term describes the additional effect risk premia have on β. The sign of this additional effect is determined by the persistence of risk premia θ. This term gains on importance with the variance of risk premia. Again, if risk premia are constant, the sign of β purely reflects the sign of g. The pure impact of the risk premium component on β under FIRE can be seen by setting $g = 0$:

$$\beta(Err^{\mathbb{Q}}_{t+1,t}, Err^{\mathbb{Q}}_{t,t-1}|g=0) = \theta\left[\left(\frac{1}{2}\frac{1}{1-\theta}\right)\left(\frac{\sigma^2_{\Delta RP}}{\sigma^2_{\Delta E}+0.5\frac{1}{1-\theta}\sigma^2_{\Delta RP}}\right)\right] \tag{4.19}$$

The effect of risk premia on β depends on the variability of risk premia relative to the variability of the expectation component and the persistence θ of risk premia. The higher the variance of risk premia, the more important the impact on β. The sign of the contribution solely depends on θ. Hence, persistent risk premia ($\theta > 0$) lead to a positive effect on error persistence. The intuition is straightforward: if the expected risk premium tomorrow is similar to the risk premium today, the expected forecast error tomorrow will be equally similar to the forecast error today. Estimates for the risk premium component in market implied inflation in the literature are usually highly persistent.[9] Therefore, under FIRE, we expect $\beta > 0$. This means, when we empirically estimate a positive β, we cannot distinguish between underreaction and the effect from persistent, variable risk premia. But this time, we get an *opposite* prediction for β for (strong) overreaction and risk premia. So while the error-on-revision regression allows us to reveal (strong) *underreaction*, the error-on-past-error regression complements our empirical testing set-up helping us to additionally uncover (strong) *overreaction*.

[9] See e.g. Buraschi and Jiltsov (2005), Chernov and Mueller (2012), Fleckenstein et al. (2016), Hördahl and Tristani (2014), Joyce et al. (2010), Christensen et al. (2010), Ang et al. (2008) for time series estimates of inflation and liquidity premia embedded in financial inflation contracts such as breakeven inflation.

4.2.3 Revision autocorrelation

While hitherto discussed approaches already allow to distinguish between overreaction, underreaction and time-varying risk premia, there is yet another testable implication of model (4.1) that will prove very useful for our empirical undertaking. It addresses two main shortcomings of the error-on-revision and the error-on-past-error regressions. They both rely on measurements of the forecast error. The unforecastable component of inflation realizations, even under the most informed fully rational forecast, is likely large relative to the effect in expectation formation we aim to measure. Moreover, estimating effects for a prediction horizon h requires observations over h periods for each independent forecast error. Hence, with increasing prediction horizon h, the number of independent observations decreases, making regressions for very long horizons impractical due to limited sample sizes. To address this, we now present a model prediction that can be analyzed using forecast revisions alone without requiring forecast errors.

We build on the general implication of the FIRE model that not only forecast errors, but also future forecast revisions, should be unpredictable with the real-time information set. This directly follows from the law of iterated expectations. Model (4.1) implies a correlation between subsequent forecast revisions if forecasts deviate from FIRE. This can be immediately seen from the fact that current periods revisions of agent's forecasts are a weighted average between revisions in the objective forecast and previous periods revisions of agent's forecasts. From Equation (4.1), forecast revisions are given by

$$\Delta F_{t+1}^{\mathbb{P}}(\pi_{t+h}) = F_{t+1}^{\mathbb{P}}(\pi_{t+h}) - F_{t}^{\mathbb{P}}(\pi_{t+h}) \tag{4.20}$$

$$= (1-g)\Delta E_{t+1}(\pi_{t+h}) + g\Delta F_{t}^{\mathbb{P}}(\pi_{t+h}) \tag{4.21}$$

Revisions in the objective forecasts, $\Delta E_{t+1}(\pi_{t+h})$, are unforecastable with information in period t. The relation between subsequent forecast revisions is therefore given by

$$\beta(\Delta F_{t+1}^{\mathbb{P}}(\pi_{t+h}), \Delta F_{t}^{\mathbb{P}}(\pi_{t+h})) = g \tag{4.22}$$

which can be empirically estimated with the regression

$$\Delta F_{t+1}(\pi_{t+h}) = \alpha + \beta * \Delta F_{t}(\pi_{t+h}) + u_{t+1,t} \tag{4.23}$$

Again, $\beta > 0$ ($\beta < 0$) indicates underreaction (overreaction) while $\beta = 0$ is consistent with FIRE.

What is the impact of risk premia on this relation? Revisions of risk neutral forecasts are composed of revisions to real world forecasts and changes in the risk premium component, as stated in Equation (4.10). Forecast revisions and changes in the risk premium component can again be written as the sum of unpredictable innovations to objective forecasts and the risk premium process, respectively. As we show in Appendix C.1.1.3, the relation between subsequent revisions of forecasts containing a risk premium component can then be calculated as[10]

$$\beta(\Delta F^{\mathbb{Q}}_{t+1}(\pi_{t+h}), \Delta F^{\mathbb{Q}}_{t}(\pi_{t+h})) = g\left(1 - \frac{\sigma^2_{\Delta RP}}{\sigma^2_{\Delta F^{\mathbb{Q}}}}\right) - \frac{1}{2}(1-\theta)\left(\frac{\sigma^2_{\Delta RP}}{\sigma^2_{\Delta F^{\mathbb{Q}}}}\right) \tag{4.24}$$

As the first term shows, g again has now a dampened impact on β as more of the variation in agents' risk neutral forecasts is driven by risk premia. The sign of the relation between g and β however is unchanged ($\frac{\sigma^2_{\Delta RP}}{\sigma^2_{\Delta F^{\mathbb{Q}}}} \leq 1$). The second component to β, introduced by risk premia, depends on the persistence of the risk premium process θ and again the variance of risk premia relative to the total variance of forecast revisions. For positively autocorrelated risk premia ($\theta > 0$), this component has a negative effect on β. This comes from the fact that changes in risk premia, which drive part of the agents' forecast revisions, will be followed by a (small) mean-reversion contribution to the next periods forecast revision. For highly persistent risk premia ($\theta \to 1$) this term vanishes as the mean-reversion effect in risk premia disappears. Therefore, the sign of the revision-on-past-revision coefficient will be approximately unaffected by a sufficiently persistent risk premium component. This means we can uncover both underreaction as well as overreaction with this empirical test even under the existence of time-varying risk premia. Empirical implementation only requires observing forecast revisions, making this test exceptionally suitable to analyze long-term inflation expectations.

4.2.4 Summary of model predictions

In summary, we rely on three empirical tests to analyze expectation formation. These tests are derived from a simple framework that describes specific deviations from FIRE. Figure 4.1 illustrates how these methods differentiate between underreaction, overreaction, and the impact of time-varying risk premia.

The first test, "Error vs. revision", runs a regression of forecast errors on forecast revisions. This well-established method, shown in Panel (a) of Figure 4.1, serves as a benchmark for identifying underreaction ($g > 0$) or overreaction ($g < 0$). The figure displays the theoretical coefficient β implied by the model depending on the true degree of overreaction or underreaction in agents‘ expectations. If the expectations measure contains no risk premium, as in survey data, the sign of β directly reflects the sign of g and therefore allows to clearly distinguish between overreaction and underreaction. If some of the variation in the expectations measure however comes from the variation of risk premia, as in market implied expectations, β no longer clearly indicates

[10] Again, we assume uncorrelated risk premium and expectation shocks for now.

underreaction and overreaction. In this case, a negative β may either result from overreaction or from risk premia variation. Only a positive β unambiguously indicates underreaction. Therefore, the "Error vs. revision" test allows only partial identification of deviations from FIRE when applied to market-implied expectations. The limitation of this test is most pronounced when the true deviation from FIRE, g, is small and the variability of risk premia is high.

The second test, "Error vs. past error", focuses on the autocorrelation of forecast errors. As shown in Panel (b) of Figure 4.1, this method is particularly valuable when expectations exhibit (strong) overreaction ($g < 0$). A negative β can be directly related to overreaction, even when the expectation measure contains a variable risk premium component. Conversely, a positive β in this case ambiguously points to either underreaction or a risk premium effect. On its own, the "Error vs. past error" test is therefore also only partially revealing about the true g. When combined with the "Error vs. revision" test however, their complementary predictions help to clearly identify both overreaction and underreaction.

The third test, "Revision vs. past revision", illustrated in Panel (c) of Figure 4.1, analyzes the autocorrelation of forecast revisions. This method effectively identifies both underreaction and overreaction even in the presence of time-varying risk premia. Risk premia variation only has a weak effect on this relation. Only for very slightly negative values of β there is ambiguity of whether this indicates overreaction, underreaction or risk premia variation. More pronounced negative values of β clearly indicate overreaction. Positive values of β always indicate underreaction. An additional advantage of this test is its empirical robustness, as it does not rely on noisy forecast errors which can be observed much less frequently with increasing forecast horizon. This makes this test particularly well-suited for studying long-term expectations.

In summary, our newly developed method of jointly studying tests on error-revision correlation, error autocorrelation, and revision autocorrelation allows to effectively analyze expectation formation in both survey and market implied expectation measures at short-term as well as long-term horizons.

Figure 4.1: Tests for overreaction and underreaction

(a) Error vs. revision

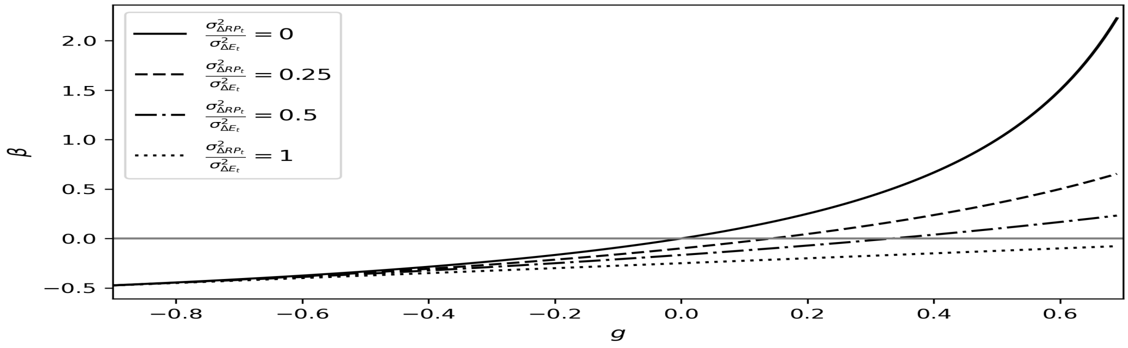

(b) Error vs. past error

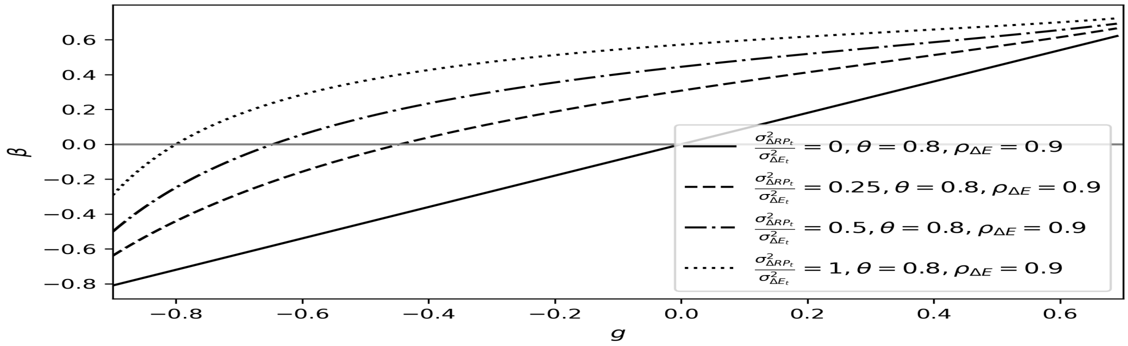

(c) Revision vs. past revision

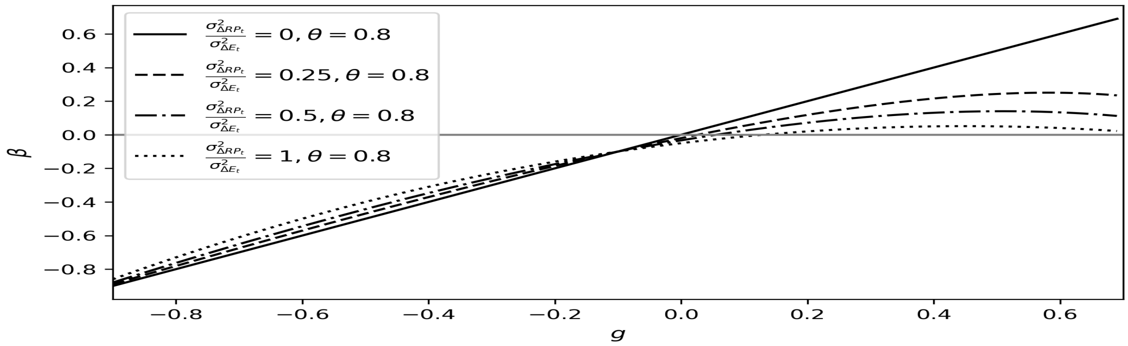

Model implications for regression coefficients of (a) forecast errors on forecast revisions, (b) forecast errors on previous period forecast errors and (c) forecast revisions on preceding forecast revisions, depending on the deviation from FIRE g. The x-axis shows the true deviation from FIRE in expectations. The y-axis shows the associated theoretical regression coefficient implied by the model. We show four versions of this relationship for different degrees of risk premia variability. The case $\frac{\sigma^2_{\Delta RP}}{\sigma^2_{\Delta E}} = 0$ refers to the case of a constant or non-existent risk premium component, as in survey expectations. The case $\frac{\sigma^2_{\Delta RP}}{\sigma^2_{\Delta E}} = 1$ refers to the other extreme case of a constant expectations component, with all variation in forecasts driven by risk premia.

4.3 Data

4.3.1 Survey expectations

We obtain survey inflation expectations from the Survey of Professional Forecasters (SPF). The SPF is run quarterly by the Federal Reserve Bank of Philadelphia and surveys expert forecasters, including agents who are likely active participants in inflation markets. It can therefore be considered to quite closely reflect the expectations of actual bond market participants. We focus on SPF forecasts of (seasonally adjusted) CPI inflation, which is directly traded in the financial markets. This allows us to link our results from survey data to the expectation formation of market participants in as unmediated a manner as possible. We work with both consensus expectations as well as the expectations of individual forecasters. Consensus expectations are given by the average forecast across individual forecasters.

The SPF survey is conducted at the beginning of each quarter, with the deadline for responses usually in the second week of the second month of the quarter, and the official survey release shortly after. The CPI index is usually released by the U.S. Bureau of Labor Statistics between the 10th and 15th of the month following the reference period. Hence, in most cases, the information set of forecasters does not yet include any official CPI release for the current quarter but does include all releases from prior quarters.

The survey asks respondents for forecasts of (annualised) changes in quarter-average CPI index levels. Therefore, we treat SPF forecasts as expectations

$$F_t^s(\pi_{t+j\to t+n}) = E_t^s\left(\frac{\overline{CPI}_{t+n}}{\overline{CPI}_{t+j}}\right)^{\frac{4}{n-j}} - 1 \tag{4.25}$$

for time-t forecasts of the realized rate of change of quarter-average CPI index levels between period $t+j$ and period $t+n$. $\overline{CPI}_t$ denotes the quarterly average of the three monthly CPI index levels in quarter t. To give an example: forecasts of the quarterly inflation rate in Q4 2025 consist of projections for the average CPI index level in October, November and December 2025, relative to projections for the average CPI index level in July, August and September 2025.

Forecast errors are consequently given by

$$Err^s_{t+j\to t+n,t} = \left(\frac{\overline{CPI}_{t+n}}{\overline{CPI}_{t+j}}\right)^{\frac{4}{n-j}} - 1 - F_t^s(\pi_{t+j\to t+n}) \tag{4.26}$$

Seasonally adjusted CPI inflation is subject to small revisions (revisions of the seasonal-adjustment factors). We follow Angeletos et al. (2021) and use final-release data. Nothing depends on this choice: when we alternatively consider first-vintage data, our results are almost

unchanged. Vintage CPI data, reflecting data releases at different points in time, is obtained from ALFRED.[11]
Quarter-over-quarter forecast revisions are calculated as

$$\Delta F_t^s(\pi_{t+j\to t+n}) = F_t^s(\pi_{t+j\to t+n}) - F_{t-1}^s(\pi_{t+j\to t+n}) \tag{4.27}$$

SPF CPI forecasts are available for single-quarter horizons (0-4 quarters ahead of the current quarter), as well as calendar-year horizons for 1, 2, 3, 5 and 10 years, each starting at the beginning of the current year. Previous literature has exclusively focused on the quarterly forecasts (0-4 quarters ahead). We extend the empirical evidence to the full term structure of expectations by using all available survey horizons. To this end, we construct forecasts for non-overlapping expectation horizons:

$$F_t^s(\pi_{t+j\to t+n}) = \frac{\left(F_t^s(\pi_{t+k\to t+n})\right)^{\frac{n-k}{n-j}}}{\left(F_t^s(\pi_{t+k\to t+j})\right)^{\frac{j-k}{n-j}}} - 1 \quad for \quad n > j > k \geq 0 \tag{4.28}$$

This reflects expectations of the CPI level growth rate between future periods $t+j$ and $t+n$, similar to a forward rate in interest rate markets. With this approach, we aim to isolate horizon-dependent effects and ensure that our results across horizons are driven by unique information. In the remainder of this chapter, we will simplify notation to $F_t^s(\pi_{t+j\to t+n}) = F_t^s(\pi_{t+h})$, where $h = (j+n)/2$.

The available data allows us to construct forecasts, forecast errors, and forecast revisions for forecast horizons referring to a single period (0, 1, 2 and 3 quarters ahead), as well as longer-term forecast horizons covering multiple periods (4-7, 5-18 and 17-38 quarters ahead). The specific horizons used in our empirical analysis vary slightly, as they depend on the regression specification which requires matching horizons for the dependent and independent variables. Panel (a) of Table 4.1 shows summary statistics of the survey data forecasts for the different horizons.

The various horizons have been included at different points in time in the SPF survey, resulting in varying historical data availability. Quarterly forecasts (0-5 quarters) have the longest history starting from Q3 1981, while 3- and 5-year horizons have only been included since Q3 2005. Our main analysis uses the sample Q1 1999 to Q4 2023, in accordance with market data availability. As Table 4.1 shows, this sample is largely consistent across survey horizons. Results for the available horizons over the longer sample from 1981 are in line with our main results (Appendix C.1.2).

[11] https://alfred.stlouisfed.org/

4.3.2 Market expectations

The most direct way to identify bond market inflation expectations is through the prices of traded assets which directly depend on future realized inflation. Consequently, we follow a broad literature and extract market-implied inflation expectations from the yields of nominal and real treasury bonds. Next to standard nominal bonds with fixed coupons, the U.S. department of the Treasury issues real bonds, known as TIPS, that compensate for realized inflation. The coupon and the principal amount of these TIPS bonds are not fixed but are indexed to the official (seasonally unadjusted) headline CPI index. Comparing the prices of nominal and TIPS bonds therefore allows to isolate market beliefs about future CPI index developments.

Although survey participants and markets forecast slightly different measures of inflation, these differences are negligible for our analysis. Specifically, SPF survey questions refer to *seasonally adjusted* CPI, while TIPS are based on *seasonally unadjusted* CPI. Since we focus on market forecasts of annual inflation rates, seasonality does not affect the analysis. Additionally, survey respondents predict the growth rate of the *quarterly average* of monthly CPI levels, whereas TIPS indexation relies on the raw CPI levels without averaging. This difference in averaging is unlikely to influence our results. Ang et al. (2007) demonstrate that the correlation between quarterly average CPI changes and end-of-quarter log CPI changes is almost perfect (0.994). Thus, we can confidently conclude that these discrepancies have a minimal impact on our estimation of expectation formation effects.

The price of a n-period real zero-coupon bond can be expressed as

$$P_t^n = E_t\left(\frac{CPI_{t+n-l}}{CPI_r}M_{t+n}^{\$}\right) = \left(\frac{CPI_{t-l}}{CPI_r}\right)E_t^{\mathbb{Q}(n)}\left(\frac{CPI_{t+n-l}}{CPI_{t-l}}\right)E_t\left(M_{t+n}^{\$}\right) \tag{4.29}$$

where $M_{t+n}^{\$}$ denotes the nominal pricing kernel applicable to valuing payoffs at date $t+n$, and CPI_t is the level of the CPI index. CPI_r refers to the reference CPI index value fixed at bond issuance, and l accounts for the 3-month indexation lag applied to adjust TIPS cashflows. This lag reflects the delay in the release of official CPI data. $E_t^{\mathbb{Q}(n)}$ is the expectation operator under the n-Forward measure, which takes a n-period zero coupon bond as it's numeraire. Specifically: $E_t^{\mathbb{Q}(n)}(X_{t+n}) = E_t\left(\frac{M_{t+n}^{\$}}{E_t(M_{t+n}^{\$})}X_{t+n}\right)$.

The yield of the TIPS bond is defined as

$$y_t^n = \frac{1}{n}\left(ln\left(\frac{CPI_{t-l}}{CPI_r}\right) - lnP_t^n\right) = y_t^{\$,n} - \frac{1}{n}lnE_t^{\mathbb{Q}(n)}\left(\frac{CPI_{t+n-l}}{CPI_{t-l}}\right) \tag{4.30}$$

which accounts for the inflation adjustment $\frac{CPI_{t-l}}{CPI_r}$ already applied to coupon and principal amount by time-t. The realized log return of this bond when held to maturity is given by

$$\begin{aligned}\frac{1}{n}r^n_{t+n} &= \frac{1}{n}\left(ln\left(\frac{CPI_{t+n-l}}{CPI_r}\right) - lnP^n_t\right)\\ &= y^{\$,n}_t - \frac{1}{n}lnE^{\mathbb{Q}(n)}_t\left(\frac{CPI_{t+n-l}}{CPI_{t-l}}\right) + \frac{1}{n}ln\left(\frac{CPI_{t+n-l}}{CPI_{t-l}}\right)\end{aligned} \tag{4.31}$$

where $y^{\$,n}_t$ is the yield of a nominal bond. To isolate inflation expectations, define

$$be^n_t = y^{\$,n}_t - y^n_t = \frac{1}{n}lnE^{\mathbb{Q}(n)}_t\left(\frac{CPI_{t+n-l}}{CPI_{t-l}}\right) \approx \frac{1}{n}E^{\mathbb{Q}(n)}_t\left(ln\left(\frac{CPI_{t+n-l}}{CPI_{t-l}}\right)\right) \tag{4.32}$$

which represents the yield of a portfolio that is long nominal bonds and short real bonds. be^n_t is commonly referred to as the break-even inflation rate and reflects bond market expectations about future CPI log growth rates. The last approximation neglects an inflation-convexity term.[12] The (annualised) forecast error for bond market expectations over period n can then be calculated as

$$Err^{\mathbb{Q}(n)}_{t+n,t} = \frac{1}{n}ln\left(\frac{CPI_{t+n-l}}{CPI_{t-l}}\right) - be^n_t \tag{4.33}$$

This forecast error can also be interpreted as the realized return of a long/short portfolio in nominal and real bonds. Consistent with our approach for survey expectations, we evaluate non-overlapping forecast horizons which can be constructed from forward yields of nominal and real bonds as

$$F^{be,j,n}_t = F^{\$,j,n}_t - F^{j,n}_t \approx \frac{1}{n-j}E^{\mathbb{Q}(n)}_t\left(ln\left(\frac{CPI_{t+n-l}}{CPI_{t+j-l}}\right)\right) \tag{4.34}$$

where $F^{\$,j,n}_t$ and $F^{j,n}_t$ are the nominal and real forward rates over periods $t+j$ to $t+n$. The associated forecast errors are

$$Err^{\mathbb{Q}(n)}_{t+j\rightarrow t+n,t} = \frac{1}{n-j}ln\left(\frac{CPI_{t+n-l}}{CPI_{t+j-l}}\right) - F^{be,j,n}_t \tag{4.35}$$

which equivalently can be interpreted as the realized return of a long/short portfolio in nominal and real forward contracts. Revisions of market inflation expectations over a period r are given by the change in break-even forward rates

$$\Delta F^{be,j,n}_t = F^{be,j,n}_t - F^{be,j,n}_{t-r} \tag{4.36}$$

[12] If $ln\left(\frac{CPI_{t+n-l}}{CPI_{t-l}}\right)$ is normally distributed under the $\mathbb{Q}(n)$-measure, then $lnE^{\mathbb{Q}(n)}_t\left(\frac{CPI_{t+n-l}}{CPI_{t-l}}\right) = E^{\mathbb{Q}(n)}_t\left(ln\left(\frac{CPI_{t+n-l}}{CPI_{t-l}}\right)\right) + \frac{1}{2}Var^{\mathbb{Q}(n)}_t\left(ln\left(\frac{CPI_{t+n-l}}{CPI_{t-l}}\right)\right)$.

We obtain nominal and real yields from the fitted yield curve model data from Gürkaynak, Sack, and Wright (2007, 2010). Data for real yields is available starting from 1999 when TIPS began trading. In our main analysis, we use market expectations at a monthly frequency with month-over-month revisions. We construct forecasts for expectation horizons ranging from 1 to 10 years. Panel (b) of Table 4.1 shows summary statistics for the market expectations data.

4.3.3 Risk premium proxies

We gather a number of variables that help us to capture the risk premium embedded in TIPS breakeven inflation. We will use them in Section 4.6 to construct an empirical estimate of the risk premium component in market inflation expectations in order to single out the effects from expectation formation in our analysis. We estimate inflation volatility with a GARCH(1,1) model on realized monthly CPI changes. The D1 variable (75th – 25th percentile) for 10-year inflation forecasts from the SPF data is taken as a measure for cross-sectional inflation forecaster dispersion. The "off-the-run/on-the-run" yield spread between the most recently issued (benchmark) bonds and previously issued old bonds at the 10-year maturity, taken from Bloomberg (NFCINMLI Index), captures the liquidity of the treasury bond market as a whole. Relative liquidity between TIPS bonds and nominal bonds is measured by the relative dealer transaction volume. Following Pflueger and Viceira (2016), we calculate relative transaction volume as $(Volume_t^{TIPS}/Volume_t^{\$})$ where $Volume_t^{TIPS}$ and $Volume_t^{\$}$ are the rolling 3-month-average transaction volumes in TIPS and nominal treasury bonds across all maturities. Primary dealer transaction volume is retrieved from the New York Federal Reserve website. Additionally, we use the synthetic-cash spread as an indicator for the cost of arbitraging between inflation-indexed and nominal bond markets for levered investors (see Pflueger and Viceira (2016)). The synthetic-cash spread is calculated as the 10-year inflation swap rate minus the 10-year cash bond breakeven inflation rate. The VIX index, which we download from Bloomberg, captures overall risk aversion in the market. Finally, we aim to directly measure the conditional covariance with the pricing kernel using the conditional stock-market beta of inflation contract returns. We estimate conditional stock-market betas by running rolling window regressions with 66-day windows[13] of inflation-returns on stock-market excess returns. Inflation returns are the realized forecast errors as described in the previous section. Stock-market excess returns are returns on S&P-500 futures from Bloomberg. Panel (c) of Table 4.1 shows summary statistics of the risk premium proxies.

[13] This reflects window lengths of one quarter, assuming 22 trading days per month.

Table 4.1: Summary statistics of expectations data

(a) Survey data

		Errors				Revisions			
Horizon (Quarters)	First observation	Mean	SD	Min	Max	Mean	SD	Min	Max
0	Q3 1981	0.09	1.28	-6.53	4.19	-0.05	1.09	-5.01	3.55
1	Q3 1981	0.03	2.07	-11.54	6.14	-0.08	0.42	-2.01	1.61
2	Q3 1981	-0.04	2.20	-11.34	6.93	-0.08	0.30	-1.39	0.86
3	Q3 1981	-0.10	2.21	-11.05	7.26	-0.08	0.30	-1.55	0.65
4-7	Q3 1981	-0.17	2.02	-11.00	7.36	-0.00	0.36	-1.51	1.26
5-18	Q3 2005	-0.30	1.05	-1.39	2.76	-0.02	0.16	-0.64	0.41
17-38	Q3 2005	-0.57	0.89	-1.34	1.50	-0.01	0.09	-0.38	0.24

(b) Market data

		Errors				Revisions			
Horizon (Quarters)	First observation	Mean	SD	Min	Max	Mean	SD	Min	Max
0-3	Q1 1999	0.69	1.74	-4.42	7.11	0.01	0.87	-4.31	4.82
4-7	Q1 1999	0.82	1.98	-4.17	7.69	0.00	0.56	-3.24	3.76
8-11	Q1 1999	0.71	2.02	-4.47	7.05	-0.01	0.38	-2.02	2.06
12-15	Q1 1999	0.57	1.94	-4.26	6.47	-0.01	0.28	-1.51	1.41
16-19	Q1 1999	0.39	2.04	-4.46	6.83	-0.01	0.22	-0.95	0.98
20-23	Q1 1999	0.23	2.16	-4.49	7.27	-0.01	0.19	-0.84	0.62
24-27	Q1 1999	0.03	2.19	-4.50	6.64	-0.00	0.18	-0.95	0.84
28-31	Q1 1999	-0.19	2.13	-4.61	6.24	-0.00	0.19	-0.93	0.98
32-35	Q1 1999	-0.33	2.08	-4.23	5.82	-0.00	0.19	-0.85	1.01
36-39	Q1 1999	-0.55	2.10	-5.12	6.07	-0.00	0.20	-0.73	0.97

(c) Risk premium proxies

Variable	First observation	Mean	SD	Min	Max
CPI Volatility (%)	Q1 1980	0.31	0.06	0.25	0.71
Survey Dispersion (%)	Q4 1991	0.48	0.18	0.10	1.00
Off-the-run spread (%)	Q1 1990	0.42	0.25	0.13	1.41
Rel. transaction volume (%)	Q2 1998	0.12	0.05	0.02	0.24
Synthetic-cash spread (%)	Q1 1980	0.24	0.09	0.02	0.98
VIX (%)	Q1 1990	19.70	7.80	9.58	69.95
Stock-market beta	Q3 1997	1.70	4.00	-13.01	11.62

This table shows summary statistics of the data used in our empirical analysis. Panel (a) describes the survey data, Panel (b) describes the market implied expectations data and Panel (c) describes the variables used to construct a risk premium estimate. Columns 3 to 6 of Panels (a) and (b) show statistics for forecast errors, columns 7 to 10 show statistics for forecast revisions. In each case, we report the first available observation, the mean, the standard deviation (SD), the minimum value, and the maximum value.

4.4 Inflation expectation formation across horizons

4.4.1 Results for survey expectations

Figure 4.2 shows results for running regressions (4.6), (4.17), and (4.23) separately for each horizon of survey and market expectations. Panel (a) of Figure 4.2 replicates a well-documented finding: short-horizon consensus survey inflation forecast errors are significantly *positively* related to consensus forecast revisions, indicating underreaction. Estimates are in the ballpark of results found in the literature. However, running the same regression for longer-horizon expectations reveals novel empirical evidence about the term-structure of inflation expectations. Specifically, as the expectation horizon increases, the relationship between consensus survey forecast errors and consensus forecast revisions diminishes rapidly. For medium horizons (around 2-5 quarters), estimates of β no longer differ significantly from zero. At even longer horizons, the relationship turns significantly negative, indicating *overreaction* in long-horizon consensus expectations.

The results for the error-on-past-error and revision-on-past revision regressions in Panel (b) and (c) of Figure 4.2, show similar patterns for the consensus survey data. Very short horizon estimates of β in both cases point to underreaction, further corroborating existing evidence. But again, β estimates decrease along the forecast horizon and eventually turn significantly negative.

Notably, the confidence intervals for the revision-on-past-revision regression coefficients are substantially narrower than those for the error-on-revision and error-on-past-error regressions. This highlights the practical advantages of the revision-on-past-revision regression, as it does not rely on the noisy inflation realization and ensures a constant number of observations independent of the forecast horizon. These properties faciliate a more efficient estimation of β.

For the error-on-past-error regressions, the absolute coefficient values are considerably smaller than those observed in the error-on-revision and revision-on-past-revision regressions. This is consistent with non-perfect cross-horizon correlations of expectation shocks ($\rho_{\Delta E} < 1$).

Regressions (4.6), (4.17), and (4.23) can also be estimated for individual forecasters instead of the consensus. Following Bordalo et al. (2020), the corresponding specifications are:

$$Err^i_{t+h,t} = \alpha + \beta * \Delta F^i_t(\pi_{t+h}) + u^i_{t+h,t} \tag{4.37}$$

$$Err^i_{t+h,t} = \alpha + \beta * Err^i_{t,t-h} + u^i_{t+h,t} \tag{4.38}$$

$$\Delta F^i_{t+1}(\pi_{t+h}) = \alpha + \beta * \Delta F^i_t(\pi_{t+h}) + u^i_{t+1,t} \tag{4.39}$$

where i indexes individual forecasters, and these equations directly parallel the consensus-based regressions (4.6), (4.17), and (4.23). This specification estimates whether, and to what extent, *individual* forecasters deviate from FIRE on average. The reaction of individual forecasters to new information might differ from the response of the consensus forecast. In fact, Coibion and Gorodnichenko (2015) show that underreaction at the consensus-level can arise from information frictions as a result of aggregation over heterogenous individual expectations. In this information-frictions-based explanation, individual forecasters maintain rational expectations and do not exhibit over- or underreaction to new information at the individual level ($\beta = 0$). Individual-level regressions more distinctly test deviations from the "rational expectations" assumptions of FIRE, as opposed to deviations from the "full information" assumption. For short-horizon expectations, the literature documents simultaneous underreaction of consensus forecasts and overreaction of individual forecasts,[14] which can be reconciled with a combination of information frictions and a form of irrational overreaction. Market expectations likely differ from the simple average across individual expectations. As the exact mechanism by which markets aggregate individual expectations into marginal expectations remains an open question, it is valuable to examine how individual expectations are formed across the term structure.

The regression coefficients shown in orange in Figure 4.2 displays individual-level results. Consistent with the findings in the literature, the individual-level error-vs-revision regressions reveal *overreaction* in short-term expectations, contrasting with the *underreaction* observed in consensus forecasts (Panel (a) of Figure 4.2). Again, estimates are quantitatively in a similar range as those in the literature. Interestingly, individual-level estimates share a similar cross-maturity pattern with consensus-level estimates: they decline with increasing maturity, suggesting that overreaction intensifies as the forecast horizon extends.

Results for the error-on-past-error and revision-on-past revision individual-level regressions (Panels (b) and (c) of Figure 4.2) confirm this finding: overreaction increases with the extended expectation horizon. At short horizons, the revision-on-past revision regression estimates show a similar discrepancy with consensus estimates as the error-on-revision regressions, indicating overreaction rather than underreaction (Panel (c)). The error-on-past error regression estimates deviate from this pattern, as they are almost identical between individual-level and consensus-level regressions (Panel (b)). Once again, the confidence intervals for revision-on-past-revision regression coefficients are much smaller, and the absolute coefficient values for the error-on-past-error regressions are more muted, in line with the results for consensus forecasts. Hence, while consensus-level and individual-level expectations show inconsistent deviations from FIRE at short horizons, they both consistently exhibit a term structure of maturity-increasing overreaction.

In summary, all three empirical tests on survey data portray a combination of underreaction and overreaction across inflation expectation horizons, with maturity-increasing overreaction in both consensus and individual expectations. These results are quite robust to alternative specifications, particularly for revision-on-past revision regressions. In the Appendix C.1.2, we

[14] See e.g. Bordalo et al. (2020), Pang (2023), Broer and Kohlhas (2024).

report results for several alternatives: A) we extend the sample back to 1981 using the longest available history for each forecast horizon; B) we only use the subsample of forecasts made by forecasters from financial services providers; C) we estimate forecaster-by-forecaster regressions, as in Bordalo et al. (2020), instead of the pooled panel specification in (4.37)-(4.38). Our main results, particularly maturity-increasing overreaction, are confirmed in each case.

4.4.2 Results for market expectations

Are similar patterns reflected in expectations implied by market prices, or are these results artifacts of survey data? The results for market data, shown in blue in Figure 4.2, indicate that error-on-revision regressions yield negative β estimates across all horizons, with a slight decrease for longer horizons. From the discussion in Section 4.2.1, we know that a negative β estimate in this case can be attributed to both overreaction and time-varying risk premia. We can therefore not clearly attribute this finding to overreaction. Much of the results is likely driven by time-varying risk premia. Just from this analysis we can only conclude that there is no underreaction across the term structure of market expectations strong enough to outweigh the risk premium effect.

Error-on-past-error regressions now allow us to distinguish between overreaction and the risk premium effect. Regression estimates display a very similar pattern to those for survey data: they decrease over the forecast horizon and are significantly negative at longer horizons. In this case, risk premia dynamics have a positive effect on β. Thus, negative β estimates provide evidence of strong overreaction, outweighing this positive risk premium effect. The results are consistent with maturity-increasing overreaction, confirming the survey data results.

Revision-on-past-revision regressions provide additional evidence pointing in the same direction. In this case, risk premia likely have only a small (negative) impact on β. The increasingly negative β estimates thus further support the finding of maturity-increasing overreaction.

Collectively, all evidence unanimously points to deviations from FIRE in agents' inflation expectations with maturity-increasing overreaction. Interestingly, results for market expectations align more closely with individual-level survey results than with consensus-level results. In particular, the lower β estimates of the individual-level error-on-revision regressions and revision-on-past-revision regressions are more consistent with market expectation results. This might suggest that the market's aggregation mechanism successfully avoids underreaction, despite information frictions which affect the survey consensus. On the other hand, the departures from rationality evident in individual-level expectations seem to carry over to the expectations of the marginal investor.

4.4.3 Horizon increasing overreaction

To confirm our visual results in Sections 4.4.1 and 4.4.2, we now formally test the horizon-dependence of overreaction. To this end, we run regressions (4.6), (4.17) and (4.23) pooled over all expectations horizons instead of independently for each horizon. The forecast horizon is included as an interaction term. Specifically, we estimate

$$Err_{t+h,t} = \alpha + \beta * \Delta F_t(\pi_{t+h}) + \gamma * h * \Delta F_t(\pi_{t+h}) + \phi * h + u_{t+h,t} \tag{4.40}$$

$$Err_{t+h,t} = \alpha + \beta * Err_{t,t-h} + \gamma * h * Err_{t,t-h} + \phi * h + u_{t+h,t} \tag{4.41}$$

$$\Delta F_t(\pi_{t+h}) = \alpha + \beta * \Delta F_{t-1}(\pi_{t+h}) + \gamma * h * \Delta F_{t-1}(\pi_{t+h}) + \phi * h + u_{t+1,t} \tag{4.42}$$

Here, γ measures how the degree of overreaction depends on the forecast horizon. For individual-level regressions we analogously estimate

$$Err^i_{t+h,t} = \alpha + \beta * \Delta F^i_t(\pi_{t+h}) + \gamma * h * \Delta F^i_t(\pi_{t+h}) + \phi * h + u^i_{t+h,t} \tag{4.43}$$

$$Err^i_{t+h,t} = \alpha + \beta * Err^i_{t,t-h} + \gamma * h * Err^i_{t,t-h} + \phi * h + u^i_{t+h,t} \tag{4.44}$$

$$\Delta F^i_t(\pi_{t+h}) = \alpha + \beta * \Delta F^i_{t-1}(\pi_{t+h}) + \gamma * h * \Delta F^i_{t-1}(\pi_{t+h}) + \phi * h + u^i_{t+1,t} \tag{4.45}$$

Table 4.2 presents the results. The average reaction to new information across horizons, reflected in the β estimates, varies strongly across the different regressions. For consensus-level survey regressions, β is positive in all cases, reflecting that the strong underreaction at short horizons dominates the cross-horizon average. In contrast, for individual level regressions, β is negative in most cases, highlighting the inconsistent reactions to new information between consensus and individual expectations at short horizons. Similarly, β for market expectations data is negative in most cases, confirming the previously observed similarity between individual-level survey expectations and market expectations.

While the maturity-average reaction to new information shows great heterogeneity, the maturity pattern of expectation formation is unambiguous across all specifications: γ is negative in every case, confirming a consistent positive relation between the forecast horizon and the degree of overreaction. In most cases, γ is statistically significant. It is insignificant for error-on-revision regressions of market expectations data, where the negative β's are partly driven by risk premia, blurring the effect from overreaction. γ is also insignificant for error-on-past-error regressions for surveys, likely due to the limited data availability for longer-term expectations. For revision-on-past-revision regressions, which are less affected by the two discussed issues, we find highly significant negative γ estimates across consensus survey expectations, individual survey expectations, and market expectations, further reinforcing the evidence of maturity-increasing overreaction.

Figure 4.2: Regression results

(a) Error vs. revision

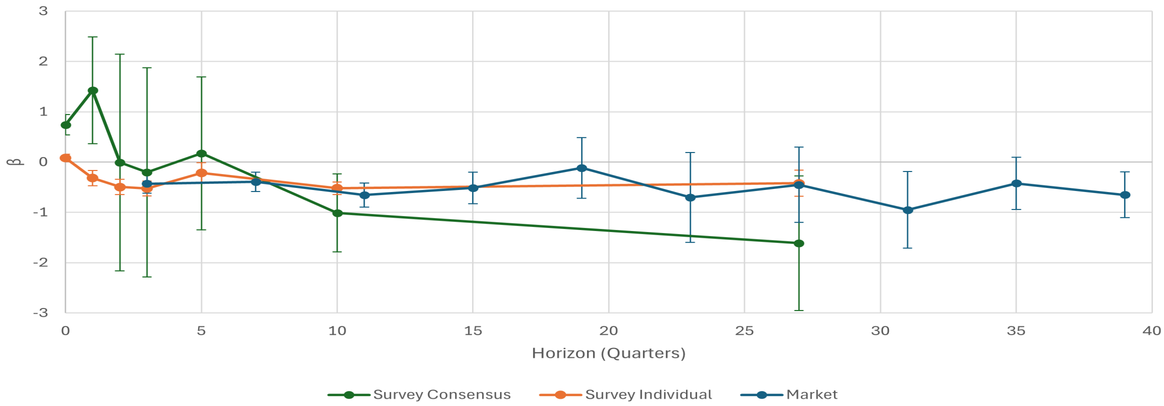

(b) Error vs. past error

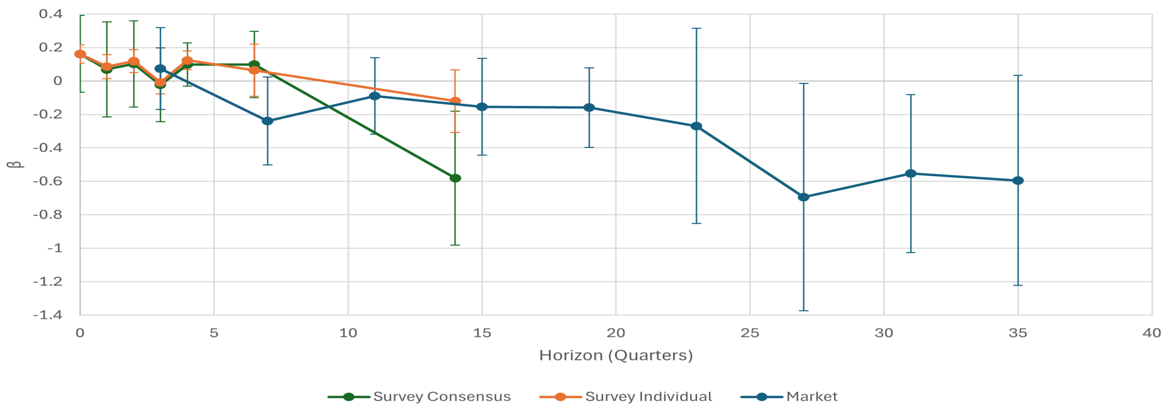

(c) Revision vs. past revision

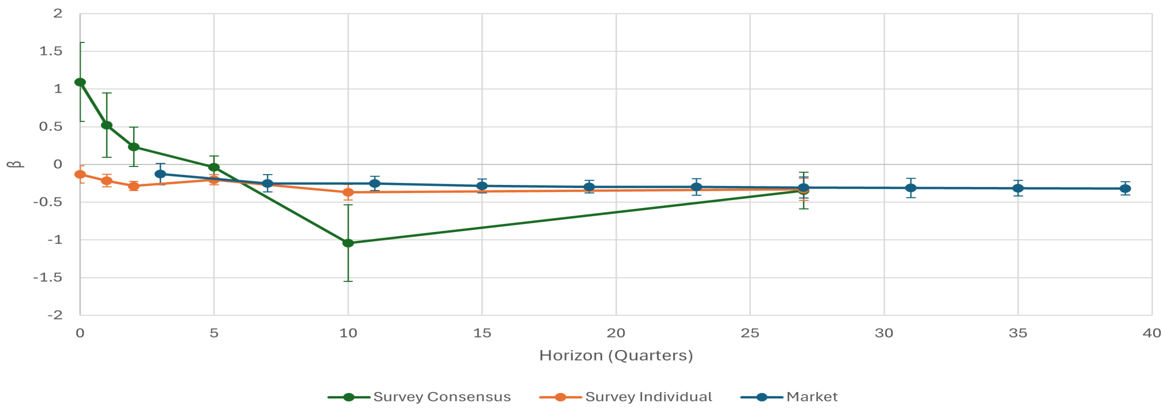

Empirical regression coefficients of (a) forecast errors on forecast revisions, (b) forecast errors on previous period forecast errors and (c) forecast revisions on preceding forecast revisions for regressions run for different forecast horizons seperately. Spikes report 90% confidence intervals calculated with HAC robust standard errors.

Table 4.2: Horizon dependence

(a) Error vs. revision

	Survey Consensus		Survey Individual		Market	
	$Err_{t+h,t}$		$Err_{t+h,t}$		$Err_{t+h,t}$	
α	0.0021**	0.0038***	0.0019***	0.0035***	0.0030***	0.0100***
$\Delta F_t(\pi_{t+h})$	0.727***	0.851***	-0.116***	-0.042	-0.455***	-0.416***
$h * \Delta F_t(\pi_{t+h})$		-0.458***		-0.151***		-0.024

(b) Error vs. past error

	Survey Consensus		Survey Individual		Market	
	$Err_{t+h,t}$		$Err_{t+h,t}$		$Err_{t+h,t}$	
α	0.0027***	0.0037***	0.0028***	0.0030***	0.0047***	0.0091***
$Err_{t,t-h}$	0.077	0.099	0.089***	0.118***	-0.154**	0.076
$h * Err_{t,t-h}$		-0.031		-0.039		-0.082***

(c) Revision vs. past revision

	Survey Consensus		Survey Individual		Market	
	$\Delta F_{t+1}(\pi_{t+h})$		$\Delta F_{t+1}(\pi_{t+h})$		$\Delta F_{t+1}(\pi_{t+h})$	
α	0.0001	0.0002	0.0000	0.0001	0.0000	0.0000
$\Delta F_t(\pi_{t+h})$	0.648***	0.917***	-0.205***	-0.178***	-0.203***	-0.148***
$h * \Delta F_t(\pi_{t+h})$		-0.510***		-0.038***		-0.025***

This table shows regression results for expectation data pooled over all horizons. We run regressions (4.40), (4.41), and (4.42) for consensus survey expectations and market expectations. For individual-level survey expectations we use the specifications (4.43), (4.44), and (4.45). Panel (a) shows the results for error-on-revision regressions, Panel (b) shows the results for error-on-past-error regressions, and Panel (c) shows the results for revision-on-past-revision regressions. Error-covariances for computation of the t-statistics are estimated accounting for heteroscedasticity and serial correlation using the approach of Newey and West (1987). *, ** and *** denote significance at the 10%, 5%, and 1% levels, respectively.

4.5 A risk premia explanation?

So far, we worked with some simplifying assumptions to gauge the impact of time-varying risk premia on regression estimates. Most importantly, we assumed a (highly) persistent process for risk premia ($\theta \to 1$) as well as approximately uncorrelated risk premium and expectation shocks ($\rho(z, \Delta E) \approx 0$). Under these assumptions, we find that market expectations overreact at longer horizons relative to FIRE. A natural follow up question is whether, and to what extent, we could alternatively explain previous empirical results with risk premia instead of overreactions, once we relax these assumptions.

As can be immediately seen from Equation (4.24), mean-reversion in risk premia could produce negative autocorrelations in subsequent revisions:

$$\beta(\Delta F^{\mathbb{Q}}_{t+1}(\pi_{t+h}), \Delta F^{\mathbb{Q}}_{t}(\pi_{t+h}) | g = 0) = -\frac{1}{2}(1-\theta)\left(\frac{\sigma^2_{\Delta RP}}{\sigma^2_{\Delta F^{\mathbb{Q}}}}\right) \tag{4.46}$$

$$= -\frac{1}{2}(1-\theta)\left(1 - \frac{\sigma^2_{\Delta E}}{\sigma^2_{\Delta F^{\mathbb{Q}}}}\right) \tag{4.47}$$

$$\geq -0.5 \quad for \quad \theta \geq 0 \tag{4.48}$$

For positively autocorrelated risk premia ($\theta \geq 0$), β under FIRE ($g = 0$) is bounded below by -0.5. This bound is achieved in the extreme case of maximal mean-reversion ($\theta = 0$) and constant inflation expectations ($\sigma^2_{\Delta E} = 0$), i.e. when market-implied forecasts are solely driven by variation in risk premia. As we show in Appendix C.1.1.3, additionally allowing for a non-zero correlation between risk premium shocks and expectation shocks ($\rho(z, \Delta E)$) does not impact Equation (4.47) and the lower bound on revision autocorrelation.

For error-on-past error β estimates, even strong mean-reversion of the risk premium by itself is insufficient to explain negative values, as can be directly seen from Equation (4.19). A risk premia explanation of negative values additionally requires a non-zero correlation between risk premium shocks and expectation shocks. If we allow $\rho(z, \Delta E)$ to be different from zero, Appendix C.1.1.2 shows that the error-on-past-error relation under FIRE now includes an additional term relative to Equation (4.19). The sign of this additional term is determined by $\rho(z, \Delta E)$. Specifically, the relationship is now given by:

$$\beta(Err^{\mathbb{Q}}_{t+1,t}, Err^{\mathbb{Q}}_{t,t-1}|g=0) = \theta\left[\left(\frac{1}{2}\frac{1}{1-\theta}\right)\left(\frac{\sigma^2_{\Delta RP}}{\sigma^2_{\Delta E}+0.5\frac{1}{1-\theta}\sigma^2_{\Delta RP}}\right)\right] - \rho(z,\Delta E)\left[\left(\frac{\sigma_{\Delta E}\sigma_{\Delta RP}\sqrt{0.5(1+\theta)}}{\sigma^2_{\Delta E}+0.5\frac{1}{1-\theta}\sigma^2_{\Delta RP}}\right)\right] \tag{4.49}$$

$$\geq -0.5 \quad for \quad \theta \geq 0 \tag{4.50}$$

Mean-reversion only reduces the *positive* impact of risk premia on β. To justify negative estimates with risk premia, one needs to additionally assume a *positive* correlation between risk premium shocks and expectation shocks ($\rho(z, \Delta E) \geq 0$). Positive premium-expectation correlation affects error-on-past error regression coefficients through the following channel: current period forecast-errors under FIRE are composed of current-period expectation shocks and previous-period risk premium

$$Err^{\mathbb{Q}}_{t,t-1}(g=0) = \epsilon_{t,t-1} - RP_{t-1} \tag{4.51}$$

as can be seen by setting $\pi_{t+h} - F^{\mathbb{P}}_t(\pi_{t+h}) = \pi_{t+h} - E_t(\pi_{t+h}) = \epsilon_{t+h,t}$ in Equation (4.9). Under $\rho(z, \Delta E) > 0$, high forecast errors $Err^{\mathbb{Q}}_{t,t-1}$ tend to coincide with increases in the risk premium RP_t, resulting in lower forecast errors $Err^{\mathbb{Q}}_{t+1,t} = \epsilon_{t+1,t} - RP_t$ in the next period.

The magnitude of β explainable purely by time-varying risk premia is again bounded by -0.5. This extreme case requires maximal mean-reversion ($\theta = 0$), perfectly correlated expectation- and risk-premium shocks ($\rho(z, \Delta E) = 1$), and equal variance of expectation shocks and risk premium shocks ($\sigma^2_{\Delta E} = \sigma^2_z = 0.5\sigma^2_{\Delta RP}$). Note that this last condition, which maximizes the impact of the covariance-channel, is inconsistent with the boundary condition for revision-on-past-revision estimates (which requires $\sigma^2_{\Delta E} = 0$ and $\sigma^2_{\Delta RP} > 0$). Therefore, these bounds cannot be achieved simultaneously.

Now, compare the derived relations under relaxed assumptions to our empirical results. The β estimates for error-on-past-error regressions range between -0.6 and -0.7 for long-horizon forecasts (Panel (b) of Figure 4.2). As demonstrated earlier, this magnitude is inconsistent with a pure risk premium explanation, given positively autocorrelated risk premia. Attributing this empirical result exclusively to risk premia variation, without allowing additional overreaction effects, would require the assumption of *negatively autocorrelated* risk premia. Under negatively autocorrelated risk premia ($\theta < 0$), β estimates under FIRE could be even more negative. However, such a risk premium process would be highly inconsistent with both empirical and theoretical findings in the literature and seems very difficult to justify economically.

β estimates for revision-on-past-revision regressions reach approximately -0.32 for long horizons

(Panel (c) of Figure 4.2). These values could be theoretically explained by positively autocorrelated risk premia. However, consider the parameter values required to match these results. $\beta = -0.3$ could be, for example, achieved with moderate variance of expectations and *maximal* mean-reversion

$$\theta = 0 \quad and \quad \frac{\sigma^2_{\Delta E}}{\sigma^2_{\Delta F^{\mathbb{Q}}}} \leq 0.4,$$

or with *constant* expectations and still very high mean-reversion

$$\theta \leq 0.4 \quad and \quad \frac{\sigma^2_{\Delta E}}{\sigma^2_{\Delta F^{\mathbb{Q}}}} = 0.$$

Note that our regressions are monthly. The corresponding quarterly persistence parameter, even for constant expectations, is close to zero ($0.4^3 = 0.064$). Therefore, any risk premium explanation of our results implies inplausibly high mean reversion of risk premia. This drastically deviates from previous literature, which commonly views inflation risk premia as highly persistent.

In conclusion, our results for error-autocorrelations and revision-autocorrelations appear to be extremely difficult to reconcile with a pure risk premia explanation, pointing to deviations from FIRE in bond market expectation formation.

4.6 Empirical risk premium estimate

Our previous regressions on market-implied inflation expectations may have been distorted by a time-varying risk premium component. The goal of this section is to empirically estimate the risk premium component in order to isolate the expectational bias in market-implied inflation. While the previous analysis shows that a pure risk premium-based explanation of our regression results would be theoretically challenging, this section aims to correct for the risk premium effect and isolate the pure expectation-driven effect.

We follow the literature estimating TIPS breakeven risk premia with a regression approach using a host of plausible proxies (see e.g. Söderlind (2011), Kajuth and Watzka (2011), Pflueger and Viceira (2016), Gürkaynak et al. (2010), Kupfer (2019)). To this end, we run the regression

$$Err^{\mathbb{Q}}_{t+h,t} = X_t\beta + u_{t+h} \tag{4.52}$$

where $Err^{\mathbb{Q}}_{t+h,t}$ represents the forecast error of market-implied inflation or the realized return of the inflation contract, as previously defined. The return $Err^{\mathbb{Q}}_{t+h,t} = \pi_{t+h} - F^{\mathbb{P}}_t(\pi_{t+h}) - RP_t$ (see Equation (4.9)) contains an expected return component RP_t that compensates for the systematic risk associated with holding the inflation contract. X_t collects the variables used to capture this systematic risk. According to the literature, the main sources of systematic risk embedded in TIPS breakeven inflation are *inflation* risk and *liquidity* risk. The compensation for liquidity risk enters market forecast errors with a positive sign. A positive compensation for bearing the risk of losses in high inflation states increases risk-neutral inflation expectations and therefore enters market forecast errors with a negative sign. One can therefore think of the market forecast error as being composed of an expectational error, a liquidity premium and an inflation risk premium:

$$Err^{\mathbb{Q}}_{t+h,t} = Err^{\mathbb{P}}_{t+h,t} + LIQ_t - INF_t \tag{4.53}$$

where LIQ_t is the liquidity premium, and INF_t is the inflation risk premium.

To capture inflation risk, we employ two measures of conditional inflation volatility: (1) GARCH(1,1) volatility estimated from historical CPI index changes (*"CPI Volatility"*), and (2) the cross-sectional dispersion of SPF forecasters 10 year CPI projections (*"Survey Dispersion"*). Liquidity risk arises from the fact that TIPS bonds are considerably less liquid than nominal treasury bonds. Following Pflueger and Viceira (2016), we capture liquidity risk using three variables that reflect complementary aspects of liquidity: The *"Off-the-run spread"* of 10-year nominal treasury bonds; the relative dealer transaction volume (*"Rel. transaction volume"*); and the *"Synthetic-cash spread"*. Additionally, we include the *"VIX"* index to capture overall risk aversion. Finally, we also incorporate the conditional *"Stock-market beta"* of the inflation contract. To estimate this conditional beta, we run a rolling window regression of $Err^{\mathbb{Q}}_{t+h,t}$ on the excess stock-market return. With this measure we intend to directly capture the conditional covariance of the return $Err^{\mathbb{Q}}_{t+h,t}$ with the pricing kernel and therefore pick up any potential embedded risk premium.

Table 4.3 shows the regression results. *"CPI Volatility"* and *"Survey Dispersion"* negatively predict forecast errors, indicating a positive relationship with inflation risk premia. This is consistent with findings in previous literature (see e.g. Söderlind (2011), Kajuth and Watzka (2011)). The loading on the *"VIX"* is negative. In the literature, both positive (e.g. Adrian and Wu (2009)) and negative (e.g. Söderlind (2011)) correlations between the VIX and TIPS risk premia have been documented. Kim et al. (2019) point out that the VIX might both be associated with inflation risk and liquidity risk and document that the correlation between the VIX and the total TIPS risk premium component depends strongly on the sample period. The positive coefficients for the liquidity variables *"Off-the-run spread"* and *"Synthetic-cash spread"* indicate a positive relationship with breakeven liquidity premia, confirming earlier literature (see e.g. Pflueger and Viceira (2011, 2016), Söderlind (2011), Kajuth and Watzka (2011)). However, the positive sign of the coefficient on the third liquidity variable, *"Rel. transaction volume"*, surprisingly suggests a positive effect on breakeven risk premia, at odds with the literature. Potentially, TIPS trading volume in our sample increased in times of heightened deflation fears, as

Table 4.3: Estimating the risk premium component

Intercept	0.0196
CPI Volatility	-9.4872***
Survey Dispersion	-0.0637***
Rel. transaction volume	0.1415***
Synthetic-cash spread	0.0358*
Off-the-run spread	5.0122***
VIX	-0.0005**
Stock-market beta	0.0015***
Adj. R^2	28.4%

This table shows regression results for regressing market forecast errors on a range of variables proxying for systematic risk in breakeven inflation. The variables include historical GARCH(1,1) volatility of CPI (*"CPI Volatility"*), the cross-sectional dispersion of SPF 10-year CPI forecasts (*"Survey Dispersion"*), the transaction volume of TIPS relative to nominal bonds (*"Rel. transaction volume"*), the spread between the 10-year inflation swap rate and the 10-year breakeven rate (*"Synthetic-cash spread"*), the spread between newly issued and older bonds (*"Off-the-run spread"*), the *"VIX"* and the conditional stock market beta of inflation contract returns (*"Stock-market beta"*). Error-covariances for computation of the t-statistics are estimated accounting for heteroscedasticity and serial correlation using the approach of Newey and West (1987). *, ** and *** denote significance at the 10%, 5%, and 1% levels, respectively.

investors tried to more actively hedge deflation risks in the TIPS market.[15] This might have led to a negative relation between *"Rel. transaction volume"* and INF_t, masking the negative effect on LIQ_t. Additionally, the relationship between *"Rel. transaction volume"* and breakeven liquidity risk might be less strong compared to older studies as the TIPS market matured and liquidity steadily increased.[16] Andreasen et al. (2021) also recently noted the inconsistent relationship between TIPS trading volume and other measures of TIPS liquidity. The coefficient on the *"Stock-market beta"* of $Err^{\mathbb{Q}}_{t+h,t}$ is significantly positive, in line with financial theory: positive exposure to broad stock-market returns positively predicts positive excess returns. Together the employed variables capture a significant share of the variation in forecast errors with an adjusted R^2 of 28%. This suggests that a substantial amount of the forecast error in market-implied expectations can be related to time-varying inflation and liquidity risk premia.

Having run regression (4.52), we then project X onto our empirical estimate of β, $\tilde{\beta}$, to obtain an estimate for the risk premium component:

$$\tilde{RP}_t = -X_t\tilde{\beta} \tag{4.54}$$

[15] Fleckenstein et al. (2017) show that market participants place substantial weight on deflation scenarios.
[16] See e.g. d'Amico et al. (2018).

The estimation results of regression (4.52) reveal the empirical properties of $\tilde{RP}_t$ summarized in Table 4.4. On average, the premium in our sample is negative, at -0.12% annualized (Panel (a) of Table 4.4). This can be rationalized with a positive compensation for liquidity risk as well as a positive deflation risk premium (i.e. a negative inflation risk premium) over our sample. Both potential explanations find strong support in the previous literature (see e.g. d'Amico, Kim, and Wei (2018), Pflueger and Viceira (2011), Shen (2006), Auckenthaler, Kupfer, and Sendlhofer (2015), Grishchenko and Huang (2013), Fleckenstein, Longstaff, and Lustig (2016, 2017), Campbell, Shiller, and Viceira (2009), Campbell, Sunderam, and Viceira (2017)). The high volatility indicates that $\tilde{RP}_t$ varies considerably over time, which is also consistent with previous observations. Most importantly, $AR(1)$ regressions in Panel (b) of Figure 4.4 show that $\tilde{RP}_t$ is highly persistent, with an empirical estimate $\tilde{\theta} = 0.83$. On the one hand, this is well in line with previous literature. On the other hand, as discussed in the previous sections, this estimated persistence is way too high to support a risk premium explanation in place of an overreaction explanation for any of our main regression results.

Table 4.4: Empirical properties of estimated risk premia

(a) Summary statistics

Mean	-0.12%
SD	1.19%
Max	2.72%
Min	-2.90%

(b) Autoregression results

	AR(1)	AR(2)	AR(3)
$\tilde{RP}_{t-1}$	0.8299***	0.8204***	0.8207***
$\tilde{RP}_{t-2}$		0.0114	0.0359
$\tilde{RP}_{t-2}$			-0.0300
$adj.R^2$	68.9%	68.8%	68.7%

This table shows empirical properties of the estimated $\tilde{RP}_t$ measure. Panel (a) shows summary statistics. We report the mean, the standard deviation (SD), the maximum value and the minimum value. Panel (b) shows results for the regression $\tilde{RP}_t = \alpha + \beta_1 \tilde{RP}_{t-1} + \beta_2 \tilde{RP}_{t-2} + \beta_3 \tilde{RP}_{t-3} + z_t$. Error-covariances for computation of the t-statistics are estimated accounting for heteroscedasticity and serial correlation using the approach of Newey and West (1987). *, ** and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Panel (b) of Table 4.4 also shows the results of estimating more flexible $AR(2)$ and $AR(3)$ specifications for $\tilde{RP}_t$. The coefficient on the first lag is roughly unchanged relative to the $AR(1)$

specification. Coefficients on lags two and three are an order of magnitude smaller and turn out to be statistically insignificant and don't improve the explanatory power. Hence the simple $AR(1)$ specification of the risk premium process already provides a quite satisfactory fit to the data and does not seem to be too much of an oversimplification. High regression R^2's support this conclusion.

In summary, our findings so far show that the time-series properties of estimated inflation risk premia (1) are *consistent* with the $AR(1)$ specification assumed to derive our theoretical results in Section 4.2, and (2) are *inconsistent* with strong mean reversion required to explain our main regression results with risk premia.

Next, we revisit regressions (4.40), (4.41) and (4.42) which assess the horizon dependence in expectation formation. We rerun these regressions for market expectations, this time controlling for $\tilde{RP}_t$ in order to isolate the effects from expectational biases. Table 4.5 shows the results. To facilitate comparison, the first two columns replicate the results from Table 4.2 (without controlling for risk premia), while columns 3 and 4 present the results with risk premia control. First, like with no risk premium control, the interaction term with the forecast horizon is negative in every case. However, after including $\tilde{RP}_t$ as control, the interaction term becomes statistically significant for the error-on-revision regression as well and remains significant in every other case. This confirms that the maturity-decreasing negative regression β's are not solely driven by risk premia but instead reflect maturity-dependent overreaction. Moreover, there are some notable differences when controlling for $\tilde{RP}_t$. For the error-on-revision regressions, including $\tilde{RP}_t$ as control variable results in considerably less negative estimates for β. This is consistent with the theoretically derived negative impact that time-varying risk premia exert on the error-on-revision relationship. However, the coefficients are still significantly negative, revealing an overreaction effect, which had been indistinguishable from a risk premium effect in the univariate regression specification. Additionally, as discussed earlier, the horizon effect of the expectational bias becomes much clearer after adjusting for contaminating risk premium effects. For the error-on-past-error regression, on the other hand, controlling for $\tilde{RP}_t$ results in markedly lower β's. Again, this aligns with our theoretical considerations regarding the positive impact that risk premia have on error autocorrelations. Finally, parameter estimates for the revision-on-past-revision regressions remain next to unchanged when we control for $\tilde{RP}_t$. This confirms the predictions of our simple framework, which suggests that persistent risk premia should have no strong impact on revision autocorrelation.

Table 4.5: Horizon dependence and an empirical estimate of the risk premium

(a) Error vs. revision

	$Err_{t+h,t}$			
α	0.0030***	0.0100***	-0.0001***	0.0000
$\Delta F_t(\pi_{t+h})$	-0.455***	-0.416***	-0.331***	-0.257***
$h * \Delta F_t(\pi_{t+h})$		-0.024		-0.038*
$\tilde{RP}_t$			-0.998***	-0.997***

(b) Error vs. past error

	$Err_{t+h,t}$			
α	0.0047***	0.0091***	-0.0002	-0.0005
$Err_{t,t-h}$	-0.154**	0.076	-0.238***	0.017
$h * Err_{t,t-h}$		-0.082***		-0.073***
$\tilde{RP}_t$			-1.176***	-1.166***

(c) Revision vs. past revision

	$\Delta F_{t+1}(\pi_{t+h})$			
α	0.0000	0.0000	0.0000	0.0002
$\Delta F_t(\pi_{t+h})$	-0.203***	-0.148***	-0.206***	-0.153***
$h * \Delta F_t(\pi_{t+h})$		-0.025***		-0.025***
$\tilde{RP}_t$			0.022***	0.026***

This table shows regression results for expectation data pooled over all horizons. We run regressions (4.40), (4.41) and (4.42) for market expectations and additionally include the empirical estimate of the risk premium component, $\tilde{RP}_t$, as a control variable. $\tilde{RP}_t$ is estimated with Equations (4.52) and (4.54). Panel (a) shows the results for error-on-revision regressions, Panel (b) shows the results for error-on-past-error regressions, and Panel (c) shows the results for revision-on-past-revision regressions. Error-covariances for computation of the t-statistics are estimated accounting for heteroscedasticity and serial correlation using the approach of Newey and West (1987). *, ** and *** denote significance at the 10%, 5%, and 1% levels, respectively.

4.7 High-frequency expectation shocks

In this section we pursue an alternative approach to empirically isolate the expectation channel from the dynamics of risk premia. Instead of constructing an empirical estimate for the risk premium component in market inflation expectations, we aim to directly identify "pure" expectation shocks. Our strategy focuses on high-frequency windows around points in time when new inflation information becomes available. Specifically, we exploit scheduled releases of official CPI numbers. The underlying assumption is that changes in market-implied inflation following CPI releases are primarily driven by the expectation component, as new inflation information gets incorporated into market prices.

We focus on the following hypothesis within our expectation framework: if market participants overreact to CPI releases, changes in implied inflation around CPI release dates ("CPI dates") will be partly reversed subsequently (see Equation (4.22)). To test this hypothesis, we run the regression

$$F^{\mathbb{Q}}_{t,i}(\pi_{t+h}) - F^{\mathbb{Q}}_{t,0}(\pi_{t+h}) = \alpha + \beta(F^{\mathbb{Q}}_{t,0}(\pi_{t+h}) - F^{\mathbb{Q}}_{t,-1}(\pi_{t+h})) + u_{t,i} \tag{4.55}$$

where $F^{\mathbb{Q}}_{t,i}(\pi_{t+h})$ denotes the end-of-day measurement of forecasts in period t on the i-th day relative to the CPI date. Our high-frequency identification assumption implies that

$$RP_{t,0} - RP_{t,-1} \approx 0$$

so that

$$F^{\mathbb{Q}}_{t,0}(\pi_{t+h}) - F^{\mathbb{Q}}_{t,-1}(\pi_{t+h}) \approx F^{\mathbb{P}}_{t,0}(\pi_{t+h}) - F^{\mathbb{P}}_{t,-1}(\pi_{t+h})$$

and

$$\begin{aligned} \beta(F^{\mathbb{Q}}_{t,i}(\pi_{t+h}) - F^{\mathbb{Q}}_{t,0}(\pi_{t+h}), F^{\mathbb{Q}}_{t,0}(\pi_{t+h}) - F^{\mathbb{Q}}_{t,-1}(\pi_{t+h})) \approx \\ \beta(F^{\mathbb{P}}_{t,i}(\pi_{t+h}) - F^{\mathbb{P}}_{t,0}(\pi_{t+h}), F^{\mathbb{P}}_{t,0}(\pi_{t+h}) - F^{\mathbb{P}}_{t,-1}(\pi_{t+h})) \end{aligned}$$

Figure 4.3 shows regression coefficients for cumulative revisions $F^{\mathbb{Q}}_{t,i}(\pi_{t+h}) - F^{\mathbb{Q}}_{t,0}(\pi_{t+h})$ over $i = 1 - 20$ days following CPI dates, for short-term (h = 1-year), medium-term (h = 5-years) and long-term (h = 10-years) maturities, respectively. The coefficients for medium- and long-term forecast horizons become increasingly negative, peaking around day 15 after the CPI release, at around -0.3 to -0.5. This indicates an average revision of about 30% to 50% of the initial shock in the following three weeks, which is consistent with strong overreaction by market participants to CPI releases. For short term maturities, coefficients are close to zero. This confirms our previous finding that overreaction tends to increase with the forecast horizon. The β estimates for cumulative 15-day revisions are also quite similar to our main revision-on-past revision regression results in Section 4.4 where we obtained estimates of about -0.3 for medium- to long-term horizons (Panel (c) of Figure 4.2).

Figure 4.3: Cumulative revisions following CPI releases

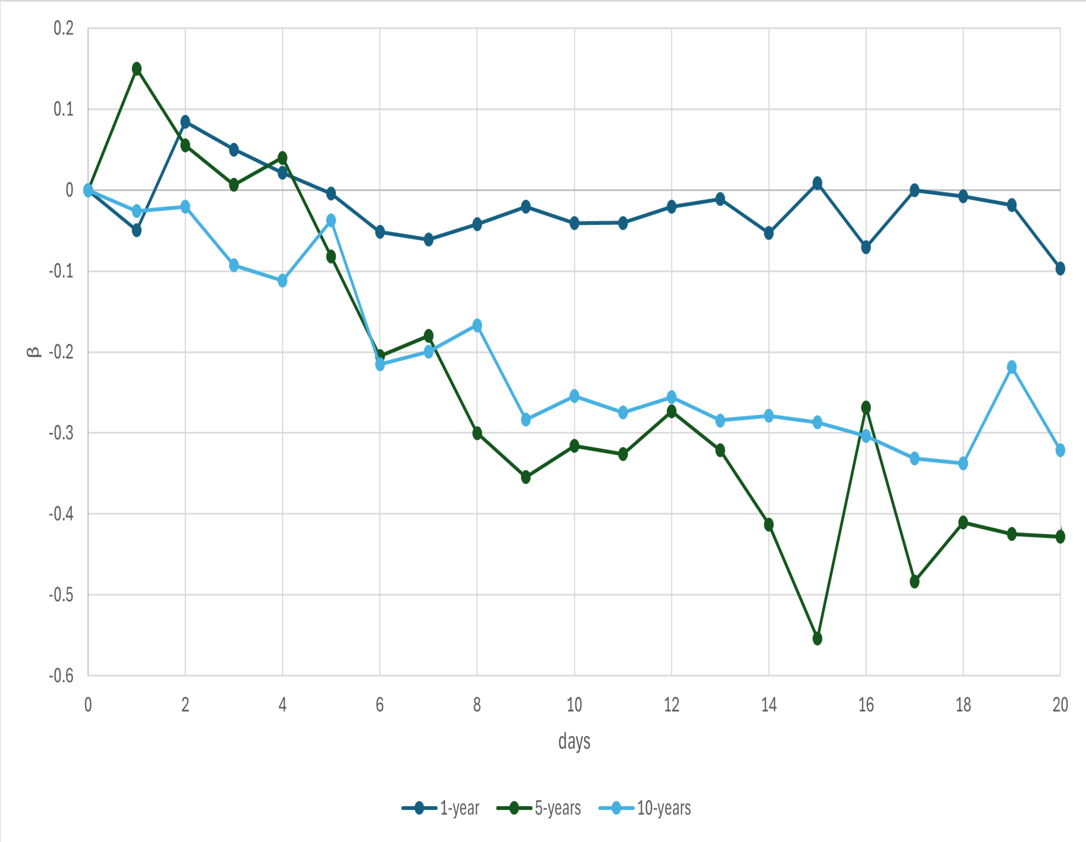

Empirical regression coefficients of cumulative post-CPI-release revisions on CPI-release-date shocks. We run the regression $F^{\mathbb{Q}}_{t,i}(\pi_{t+h}) - F^{\mathbb{Q}}_{t,0}(\pi_{t+h}) = \alpha + \beta(F^{\mathbb{Q}}_{t,0}(\pi_{t+h}) - F^{\mathbb{Q}}_{t,-1}(\pi_{t+h})) + u_{t,i}$ individually for $i = 1 - 20$ days, indicated on the x-axis. We show results for forecast horizons of h = 1-year, h = 5-years and h = 10-years.

Table 4.6 formally investigates the relation between forecast horizon and the degree of overreaction. We focus on the 15-day period following CPI releases and consider the interaction between the forecast horizon and the relationship between the initial shock and subsequent revisions:

$$\begin{aligned} F^{\mathbb{Q}}_{t,15}(\pi_{t+h}) - F^{\mathbb{Q}}_{t,0}(\pi_{t+h}) &= \alpha + \beta * (F^{\mathbb{Q}}_{t,0}(\pi_{t+h}) - F^{\mathbb{Q}}_{t,-1}(\pi_{t+h})) \\ &\quad + \gamma * h * (F^{\mathbb{Q}}_{t,0}(\pi_{t+h}) - F^{\mathbb{Q}}_{t,-1}(\pi_{t+h})) + \phi * h + u^{h}_{t,15} \end{aligned} \tag{4.56}$$

The interaction term with the forecast horizon turns out significantly negative, confirming the pattern of maturity-increasing overreaction. The estimate for γ closely resembles that from our main regression (-0.025 vs. -0.028) shown in Panel (c) of Table 4.2. Also, the estimates for β are not too far off. Controlling for our risk premium estimate leaves the results virtually unchanged. This confirms our conclusion that the horizon-dependent patterns we observe are rooted in agents' expectation formation. It also underscores once more the robustness of the revision-vs-past revision test, which remains largely unaffected by the presence of time-varying risk premia.

Table 4.6: Horizon dependence of CPI-release-shock revisions

	$\Delta F^{\mathbb{Q}}_{t,15}(\pi_{t+h})$			
α	0.0000	-0.0001*	0.0000	0.0000
$\Delta F^{\mathbb{Q}}_{t,0}(\pi_{t+h})$	-0.315***	-0.234***	-0.316***	-0.233***
$h * \Delta F^{\mathbb{Q}}_{t,0}(\pi_{t+h})$		-0.028***		-0.028***
$\tilde{RP}_t$			0.003*	0.003*

This table shows regression results for post-CPI-release revisions on CPI-release-date shocks pooled over all horizons. Error-covariances for computation of the t-statistics are estimated accounting for heteroscedasticity and serial correlation using the approach of Newey and West (1987). *, ** and *** denote significance at the 10%, 5%, and 1% levels, respectively.

4.8 Conclusion

We provide evidence that bond market inflation expectations deviate from FIRE. The strength of the deviation varies across the term structure of inflation expectations. Bond market participants overreact to new inflation information compared to the FIRE benchmark, and this overreaction becomes more pronounced at longer forecast horizons.

We find strong support that our empirical observations are neither driven by risk premia dynamics, nor by survey data artifacts not reflected in efficient equilibrium market prices. Instead, we argue that our results can be tied to non-FIRE expectation formation by marginal bond market participants, who direcly influence market prices.

First, we find similar patterns of horizon-increasing overreaction across inflation survey consensus data, individual survey forecast data, and market-implied inflation data. While survey data results exclude a risk premium explanation, market-implied data results affirm that expectation biases survive market mechanisms and affect equilibrium prices. Second, we show that an alternative explanation of our market-implied data results via risk premia would require rather unreasonable assumptions about risk premia dynamics. Mainly, any risk premia explanation would need to assume an implausible high mean-reversion, if not a negative autocorrelation, of inflation risk premia. Third, we empirically construct an estimate for a dynamic risk premium component in market-implied inflation, capturing a wide range of plausible sources of systematic risk, including inflation and liquidity. Accounting for this risk premium estimate in our main regressions supports the conclusion that observed results stem from the specifics of market participants expectation formation – not risk premium dynamics. Fourth and last, we zoom in on CPI release dates to identify shocks to market-implied inflation mostly driven by the expectation component. The pattern of subsequent revisions of these “pure-expectation-shocks” are in line with the previously identified patterns of over- and underreaction.

Our results contribute to the understanding of major bond market puzzles. The non-FIRE inflation expectations we document in this chapter help explain the high volatility and excess sensitivity to economic news of long-term nominal interest rates.

Chapter 5

Conclusion

This dissertation studies several aspects of the interconnection between the macroeconomy and financial markets and highlights the role of investor beliefs therein.

In Chapter 2 we study market beliefs about a lower bound on interest rates and how they evolved over time. We show that option-implied interest rate distributions reveal important information about market participants' anticipation of a lower bound and allow to infer an upper-limit for its location. Our findings reveal a distinctive evolution of lower bound beliefs: initially, markets were convinced of the "zero-lower-bound" - a widespread macroeconomic theory at the time. With the realization of negative interest rates in several major economies from 2014/2015, markets had to materially revise these prior lower bound beliefs. We document this pattern both in a market where the zero lower bound has been actually breached (the Eurozone), as well as one where it has not (the US). We provide evidence that the market anticipated lower bound turned negative after 2015 in both markets. As a result, the impact on interest rate distributions that arises from interest rates being close to the anticipated lower bound has been much lower after the downward revision of markets lower bound beliefs.
Our findings provide valuable guidance for interest rate modelling and monetary policy. When estimating dynamic term-structure models using historical data, researchers should account for a (zero) lower bound constraint subject to a substantial downward shift around 2015. For this purpose, established lower bound consistent modelling approaches, such as shadow rate models, could be provided with a suitable exogenous specification of the lower bound location. Correctly accounting for the asymmetric impact of a lower bound on interest rates in these models is crucial for accurate inference about interest rate expectations, risk premia and the effectiveness of monetary policy. The methodology we develop in this chapter helps to monitor markets lower bound beliefs in real time in a model-free manner. In future periods of low interest rates these tools can be used to assess markets views on remaining monetary policy margins for further rate cuts.

In Chapter 3 we focus on investors' views regarding the conditional cyclical relationship between persistent consumption growth and long-run inflation expectations. We construct a

measure from economic news reporting that quantifies how investors currently view changes in inflation expectations to be related to the long-run consumption outlook. We show that news data captures expectations about long-run consumption growth and about long-run inflation. Our cyclicality measure combines this news-embedded information to capture the stochastic relationship between both macroeconomic expectations. The measure empirically uncovers previously unobserved relationships between long-run components in macroeconomic expectations, which are crucial in an asset pricing context. We employ this cyclicality news measure to reconcile observed co-movements in asset prices within a macro-finance equilibrium asset pricing model. The news measure allows us to directly test the central channel common to many leading asset pricing models explaining the correlation between stocks, bonds and inflation with a macroeconomic mechanism. Our findings document strong consistency with model predictions. Most importantly, the news measure significantly positively predicts future realized asset correlations of stocks, bonds, and inflation contracts. These results provide direct evidence that asset co-movements depend on investors' views about the *long-run* relationship between inflation and consumption growth.
Our results in this chapter demonstrate that monitoring news data helps to get timely insights into the dynamic risk of inflation exposure and specifically the conditional relationship between the two major asset classes stocks and bonds. Investors may exploit this fact to improve their asset allocation decisions and get a sharper understanding of the macroeconomic drivers of their portfolios. The approach we put forward in this chapter to empirically link unobservable macroeconomic dynamics and asset prices using news data could be adopted in future research to address the critique of "dark-matter"[1] in asset pricing models.

Chapter 4 examines the formation of bond market inflation expectations. We analyze expectations data from surveys and bond prices at various expectation horizons. Our findings suggest that inflation expectation formation driving bond prices deviates from the established full information rational expectations (FIRE) model. Investors' inflation expectations overreact to new information relative to the FIRE benchmark with the strength of the overreaction increasing with extending forecast horizon. We carefully scrutinize whether our results could be driven by risk premia dynamics instead of expectation formation. Our results unambiguously point to the peculiarities of expectation formation as the most plausible explanation for our results. Theoretically, we show that an explanation of our results via risk premia would require very unusual and unreasonable assumptions about the dynamics of risk premia. Empirically, multiple independent pieces of evidence support an expectations-based explanation: (1) Survey expectations show similar patterns as market-implied expectations; (2) A regression-based estimate of the risk premium component is consistent with an expectation-driven interpretation of our results; (3) A high-frequency identification of expectation shocks around CPI releases

[1] Chen et al. (2024)

supports our conclusions regarding expectation formation.
Our findings in this chapter contribute evidence that deviations from FIRE are not only an important characteristic of survey expectations but are also reflected in the collective expectations embedded in equilibrium asset prices. Consequently, our results relate to prominent bond market puzzles. The deviations from FIRE we document help to understand the high volatility and excess sensitivity to economic news of long-term nominal interest rates. Further exploring empirically supported mechanisms of expectation formation to tackle asset pricing questions is a promising avenue for future research.

Overall, this dissertation highlights the importance of at times subjective, informational incomplete, and hard to measure beliefs of investors about macroeconomic dynamics. Understanding and measuring investor beliefs is essential to rationalize asset prices and their relation to the macroeconomy. Financial derivatives, surveys, and news articles are valuable sources of information in this regard.

Appendix A

Option-implied lower bound beliefs and the impact of negative interest rates

A.1.1 Proof of Proposition 1

Proof. Symmetry implies that P_t has support only on $[LB, 2\mu - LB]$. Popoviciu's inequality (Popoviciu 1935) for probability distributions bounded by a and b then immediately gives $\sigma^2 \leq \frac{(b-a)^2}{4} = (\mu - LB)^2$. Hence, the volatility for a symmetric distribution is bounded by $\sigma \leq (\mu - LB) = \overline{\sigma}$ and the distance to lower bound by $d \geq \frac{(\mu - LB)}{\overline{\sigma}} = 1$. □

A.1.2 Proof of Proposition 2

Proof. Problems (2.1),(2.2) can be formulated as the general moment problem

$$\inf_{P \in \mathcal{P}} \int_{LB}^{\infty} H(L_T) dP = \inf_{P \in \mathcal{P}} \int_{LB}^{\infty} \left(\frac{L_T - \mu}{\sigma} \right)^z dP, z \in \{3, 4\} \tag{A.1}$$

s.t.

$$\int_{LB}^{\infty} L_T dP = \mu \tag{A.2}$$

$$\int_{LB}^{\infty} L_T^2 dP = \mu^2 + \sigma^2 \tag{A.3}$$

Karr (1983) studies properties of generalized moment problems like (A.1). He shows that optimal solutions are obtained at extreme points of $\mathcal{P}$, which are probability measures with finite support on $L \leq N + 1 = 3$ distinct points (Theorem 2.1, Karr (1983)). If there exists the continuous 1^{st} derivative of H, H', and H' is strictly convex on $[LB, \infty)$, an insight from Krein and Nudelman (1977) (Chapter 4, Theorem 1.1) applies and further simplifies the problem: in this case, extremal solutions to (A.1) are obtained for distributions with only two distinct points of support $\{x_1, x_2\}$, where $x_1 = LB$ and x_2, as well as the probability weights $\{p_1, p_2\}$, are uniquely determined by conditions (A.2),(A.3) and $\int_{LB}^{\infty} dP = 1$.

For $\underline{\mathcal{S}}$ then immediately follows

$$\underline{\mathcal{S}} = \frac{1-d^2}{d}$$

which is obtained for the two-point-support distribution $\{\{x_1, x_2\}, \{p_1, p_2\}\}$ with

$$x_1 = LB$$
$$x_2 = \frac{\mu^2 + \sigma^2 - LB\mu}{\mu - LB}$$
$$p_1 = \frac{\sigma^2}{(\mu - LB)^2 + \sigma^2} = \frac{1}{d^2+1}$$
$$p_2 = \frac{d^2}{d^2+1}$$

For $\underline{\mathcal{K}}$ we consider the program

$$\min_{\{x_i, p_i\}} \sum_{i=1}^{3} \left(\frac{x_i - \mu}{\sigma}\right)^4 p_i \tag{A.4}$$

s.t.

$$\sum_{i=1}^{3} p_i = 1 \tag{A.5}$$
$$\sum_{i=1}^{3} x_i p_i = \mu \tag{A.6}$$
$$\sum_{i=1}^{3} x_i^2 p_i = \mu^2 + \sigma^2 \tag{A.7}$$
$$p_1, p_2, p_3 \geq 0 \tag{A.8}$$
$$LB \leq x_1 < x_2 < x_3 \tag{A.9}$$

The program can be solved with the Lagrange Method. Note that the problem is obviously only feasible if $x_1 \in [LB, \mu)$ and $x_3 > \mu$. For $\mu - \sigma \geq LB (d \geq 1)$ the optimal solution with inactive inequality constraint $x_1 \geq LB$ is given by the two-point-support distribution with $x_1 = \mu - \sigma, x_2 = \mu + \sigma, p_1 = p_2 = 0.5$ and $\underline{\mathcal{K}} = 1$.

For $\mu - \sigma < LB (d < 1)$, the optimal constrained solution is the same discrete distribution as derived for $\underline{\mathcal{S}}$. Kurtosis for this solution is given by $\underline{\mathcal{K}} = \frac{1+d^6}{d^4+d^2}$

□

A.1.3 Dataset

Our derivatives dataset consists of daily settlement prices of futures and options on the three-month USD-LIBOR rate (Eurodollar futures) and the three-month Euribor rate trading at CME and at ICE Futures Europe. Option maturities are available at monthly frequency, but we limit our analysis to quarterly expiries. These options are most liquid and much easier to value as their expirations coincide with those of the underlying futures. Conveniently, they are in effect options on the underlying reference rate itself. Quarterly Eurodollar (Euribor) options have maturities of up to four (two) years, providing us with 16 (8) points across the short-end of the yield curve at any given day. We further limit our data to out-of-the-money (OTM) options, which are typically more liquid than in-the-money options (see e.g. Puigvert-Gutiérrez and de Vincent-Humphreys (2012)). This gives us unique option prices for each strike-maturity combination as required by our empirical methodology.

Finally, we apply some additional standard filters to our options sample. We exclude option prices which are negative or below the minimum tick-size, and we exclude prices at the last trading day. Further, option prices are screened to meet the basic arbitrage considerations. After applying all filters, we estimate option-implied moments on a given day for all maturities with at least two put options and two call options. Table A.2 provides details on the composition of our option sample.

While Eurodollar and Euribor options are conceptually very similar, we point out one difference which affects their pricing and the calculation of implied moments: Unlike Eurodollar options, Euribor options are subject to futures-style margining. That is, the premium for Euribor options is not paid upfront at the time of purchase. Instead, option positions are marked-to-market giving rise to daily variation margin flows. This simplifies pricing of these options as it no longer involves discounting or an early-exercise premium (see Lieu (1990) and Chen and Scott (1993)).

Table A.1: Trading volume

	Eurodollar		Euribor		10-Yr US Treasury Note	
	Futures	Options	Futures	Options	Futures	Options
Open interest	10,940,505	36,266,803	3,821,640	4,653,484	3,623,839	3,116,891
Volume	36,492,894	15,285,898	11,283,635	1,033,095	28,427,186	15,066,820
Average daily volume	1,737,757	727,900	537,316	49,195	1,353,676	717,468

This table shows trading volume and open interest as of December 2019 for the future and option contracts used in our study. The last two columns additionaly provide trading statistics for 10-Year US treasury notes for comparison. Eurodollar contracts and 10-Year US treasury futures and options trade at CME, Euribor contracts trade at ICE Futures Europe. Data comes directly from CME and ICE Futures Europe.

Table A.2: Sample composition

Maturity (in months)	Moneyness					
	0.99-1	0.98-0.99	0.97-0.98	0.96-0.97	$\leq$ 0.96	Total
Panel A: US						
0-3	7,677	894	191	42	0	8,804
3-6	24,601	5,000	967	225	40	30,833
6-9	35,358	10,279	2,463	839	150	49,089
9-12	43,308	15,902	4,807	1,527	491	66,035
12-15	48,057	22,065	7,794	2,452	1,804	82,172
15-18	50,835	29,139	10,745	3,597	3,146	97,462
18-21	52,010	34,452	13,159	4,609	4,559	108,789
21-24	49,596	34,589	12,563	5,003	4,645	106,396
24-27	31,783	23,971	8,514	2,791	2,113	69,172
27-30	31,212	25,395	10,266	3,554	2,708	73,135
30-33	30,588	25,890	11,749	4,175	3,584	75,986
33-36	28,825	24,746	12,110	5,052	4,033	74,766
36-39	22,511	21,401	10,570	5,097	3,557	63,136
39-42	17,939	16,865	8,757	4,662	3,855	52,078
42-45	14,137	13,212	6,975	4,273	3,988	42,585
45-48	10,883	9,758	5,267	3,556	3,644	33,108
Total	499,632	313,817	127,045	51,557	42,443	1,034,494
Panel B: EU						
0-3	8,799	449	16	2	0	9,266
3-6	24,351	2,783	82	6	0	27,222
6-9	34,219	7,735	1,093	25	0	43,072
9-12	42,165	12,610	3,896	410	6	59,087
12-15	35,389	12,020	5,413	2,012	349	55,183
15-18	38,793	15,338	7,045	3,431	949	65,556
18-21	41,285	19,248	8,656	4,324	1,895	75,408
21-24	39,395	21,427	9,090	3,900	1,614	75,426
Total	264,396	91,610	35,291	14,110	4,813	410,220

This table provides details about the composition of our sample. For each combination of option maturity and moneyness, we count the number of available daily option prices in our sample running from June 2006 (May 2005) to December 2019 for the US (Eurozone) market. Moneyness is calculated for call (put) options as the ratio of futures (strike) price to strike (futures) price. Eurodollar contracts and 10-Year US treasury futures and options trade at CME, Euribor contracts trade at ICE Futures Europe. Data comes from Thomson Reuters Datastream.

A.1.4 Model-free estimation of option-implied interest rate moments

Consider a European-style call option with maturity T and strike K. Its value at time t is given by

$$C_t(K) = \mathbb{E}_t^{\mathbb{Q}}\left[e^{-\int_t^T r_t dt}\left(L_T - K\right)^+\right] \tag{A.10}$$

$$= B_{t,T}\mathbb{E}_t^{\mathbb{Q}_T}\left[\left(L_T - K\right)^+\right] \tag{A.11}$$

where $\mathbb{Q}$ denotes the risk-neutral measure, $\mathbb{Q}_T$ is the T-forward measure, and $B_{t,T}$ is the price of a default-free zero-coupon bond maturing in T.

For options on Euribor, valuation simplifies due to futures style margining: the discounting term drops from (A.10) and early exercise is never optimal, justifying the above treatment as European options. This result has been shown by Lieu (1990) and Chen and Scott (1993). Valuation for Euribor options therefore narrows down to just taking the expectation of the terminal payoff under the risk-neutral measure:

$$C_t(K) = \mathbb{E}_t^{\mathbb{Q}}\left[\left(L_T - K\right)^+\right]$$

For Eurodollar options, the use of the $\mathbb{Q}_T$ measure becomes necessary to account for stochastic interest rates (Equation (A.11)). The $\mathbb{Q}_T$ measure is closely related to the $\mathbb{Q}$ measure but uses a zero-coupon bond with maturity T as numeraire instead of the money-market cash-account. Furthermore, in case of Eurodollar options, the European style valuation in (A.11) constitutes an approximation, neglecting the value of early exercise. However, based on existing evidence for Eurodollar options, we expect this effect to be negligible (See Flesaker (1993b), Flesaker (1993a) and Cakici and Zhu (2001)).

Following Bakshi and Madan (2000), for any fixed Z, one can write any twice differentiable payoff function of L_T, $g(L_T)$ as:

$$g(L_T) = g(Z) + g'(Z)(L_T - Z) + \int_Z^\infty g''(K)[L_T - K]^+ dK + \int_0^Z g''(K)[K - L_T]^+ dK \tag{A.12}$$

Setting $Z = \mathbb{E}_t\left(L_T\right)$ and taking expectations under the appropriate measure one can obtain expressions for conditional moments in terms of out-of-the-money options.

Bakshi et al. (2003) use this result to derive formulas for the risk-neutral moments pertaining to the *log-return* distribution of the underlying. We are however interested in the distribution of the *level* L_T of future interest rates, which directly correspond to our derivations in Section 2.2.

Formulas for the central moments of the level of interest rates at the time horizon of the option expiry can be obtained based on (A.12) by setting $g(L_T)$ to the corresponding moment.

For Eurodollar options we obtain

$$\begin{aligned} Var_{t,T}^{\mathbb{Q}_T}(L_T) &= \mathbb{E}^{\mathbb{Q}_T}\left[(L_T - F_{t,T})^2\right] \\ &= \frac{2}{B_{t,T}}\left[\int_{F_{t,T}}^{\infty} C(K)dK + \int_0^{F_{t,T}} P(K)dK\right] \end{aligned} \quad (A.13)$$

$$Vol_{t,T}^{\mathbb{Q}_T}(L_T) = \sqrt{Var^{\mathbb{Q}_T}} \quad (A.14)$$

$$\begin{aligned} Skew_t^{\mathbb{Q}_T}(L_T) &= \mathbb{E}^{\mathbb{Q}_T}\left[\left(\frac{L_T - F_{t,T}}{\sigma}\right)^3\right] \\ &= \frac{6}{B_{t,T} Var_t^{\mathbb{Q}_T}(L_T)^{\frac{3}{2}}}\left[\int_{F_{t,T}}^{\infty} (K - F_{t,T})C(K)dK + \int_0^{F_{t,T}} (K - F_{t,T})P(K)dK\right] \end{aligned} \quad (A.15)$$

$$\begin{aligned} Kurt_{t,T}^{\mathbb{Q}_T}(L_T) &= \mathbb{E}^{\mathbb{Q}_T}\left[\left(\frac{L_T - F_{t,T}}{\sigma}\right)^4\right] \\ &= \frac{12}{B_{t,T} Var_t^{\mathbb{Q}_T}(L_T)^2}\left[\int_{F_{t,T}}^{\infty} (K - F_{t,T})^2 C(K)dK + \int_0^{F_{t,T}} (K - F_{t,T})^2 P(K)dK\right] \end{aligned} \quad (A.16)$$

where $F_{t,T}$ denotes the forward LIBOR rate. Formulas for Euribor options are analogous, only that moments are measured under $\mathbb{Q}$ and the discounting term $\frac{1}{B_{t,T}}$ drops from (A.13)-(A.16).

While evaluation of Equations (A.13)-(A.16) demands a continuum of option prices, traded LIBOR options are available only for a discrete finite set of strikes. Therefore empirical implementation requires inter- and extrapolation of option prices across strikes. We follow Jiang and Tian (2005) and Carr and Wu (2009) among others to fit an interpolating function in strike implied-volatility space, which is more reliable then interpolation directly in strike-price space. To calculate (A.13)-(A.16) for a given observation date and maturity we accordingly proceed as follows: First, we translate observed settlement prices of out-of-the-money put and call options into normal implied volatilities (IV). These volatilities equate theoretical option prices under the assumption of normally distributed future LIBOR rates to observed option prices. Second, we interpolate between obtained IVs by fitting a cubic spline and extrapolate outside the lowest and highest strikes using the IV at each boundary. Third, from the fitted IVs we obtain a fine grid of strike prices and evaluate the integrals in (A.13)-(A.16) through numerical integration using

Simpson's rule. This way, for each trading day we obtain estimates of $Vol_{t,T}(L_T)$, $Skew_{t,T}(L_T)$ and $Kurt_{t,T}(L_T)$ for all available quarterly expirations.

A.1.5 Results for realized moments

Table A.3: Conditional realized skewness

	Total	<0%	0%-2%	2%-4%	>4%
Panel A: US					
Mean	1.99	-	3.29	0.55	-0.68
t-stat	-	-	7.04	-5.00	-8.14
SD	3.70	-	4.12	2.17	1.01
N	1727	0	981	577	169
Panel B: EU					
Mean	0.69	-0.04	1.27	0.59	0.06
t-stat	-	-3.12	4.35	-0.49	-1.48
SD	1.78	1.85	1.54	1.51	2.24
N	1186	261	523	272	130

This table shows summary statistics of monthly realized skewness for the US and the Eurozone across all available maturities. The first column displays unconditional results while the second to fifth column display results conditional on the futures rate being below 0%, in the interval 0%-2%, in the interval 2%-4% and above 4%, respectively. The second row within each panel shows t-statistics of the slope coefficient from the regression $\widetilde{Skew}_{t,T} = \alpha + \beta I_{\left(a \leq f_{t,T} < b\right)}$. t-statistics are corrected for heteroscedasticity and serial correlation up to 12 lags (i.e. one year) using the approach of Newey and West (1987).

Table A.4: Conditional realized kurtosis

	Total	<0%	0%-2%	2%-4%	>4%
Panel A: US					
Mean	9.96	-	12.34	7.31	5.12
t-stat	-	-	4.42	-3.39	-4.55
SD	13.18	-	15.64	8.73	3.90
N	1727	0	981	577	169
Panel B: EU					
Mean	9.92	9.47	12.21	7.48	6.71
t-stat	-	-0.45	3.32	-3.04	-2.68
SD	10.30	9.92	12.35	5.66	6.87
N	1186	261	523	272	130

This table shows summary statistics of monthly realized kurtosis for the US and the Eurozone across all available maturities. The first column displays unconditional results while the second to fifth column display results conditional on the futures rate being below 0%, in the interval 0%-2%, in the interval 2%-4% and above 4%, respectively. The second row within each panel shows t-statistics of the slope coefficient from the regression $\widetilde{Kurt}_{t,T} = \alpha + \beta I_{\left(a \leq f_{t,T} < b\right)}$. t-statistics are corrected for heteroscedasticity and serial correlation up to 12 lags (i.e. one year) using the approach of Newey and West (1987).

Figure A.1: Break-date estimation realized moments

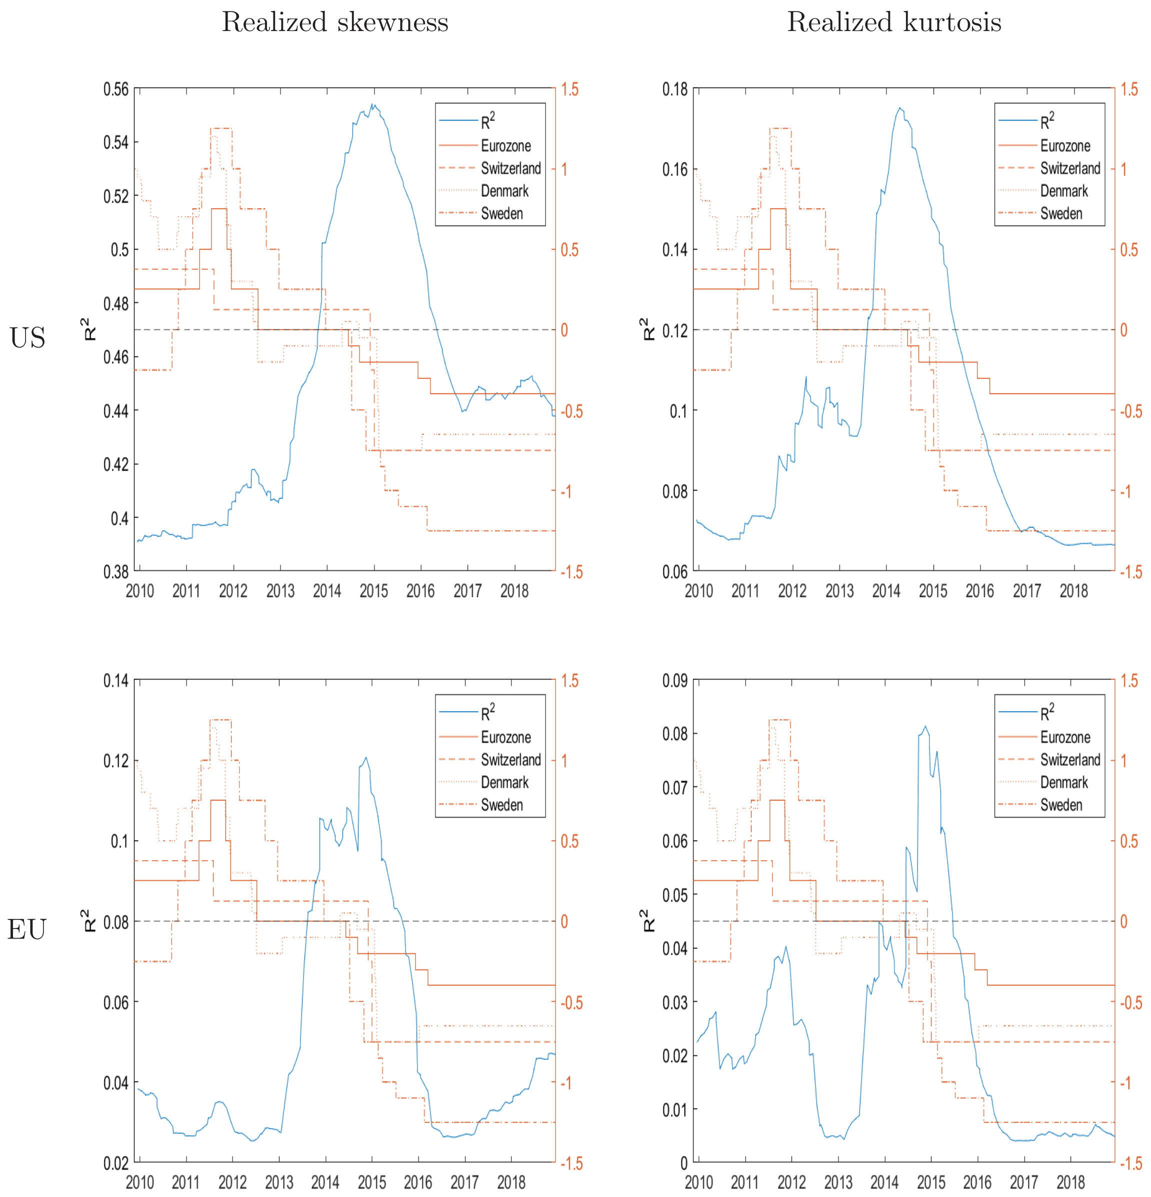

The figure plots R^2s of the structural-break-regressions (2.6) depending on the date of the structural break T_b. Our least-square estimates for the break-date parameter T_b maximize R^2. Additionally, the deposit facility rate set by the European central bank and the three-month Euribor rate are plotted.

Table A.5: Structural break regression: US realized skewness

Maturity (in months)	α	β_0	β_1	F	R^2	Observations
3-6	-4.45***	17.30***	8.59***	40.03***	0.474	127
	(-8.34)	(8.79)	(7.45)			
6-9	-3.68***	15.35***	7.77***	78.55***	0.548	148
	(-5.72)	(7.48)	(5.34)			
9-12	-3.50***	15.13***	7.90***	51.28***	0.582	168
	(-5.69)	(7.38)	(5.45)			
12-15	-3.29***	14.93***	7.93***	68.32***	0.751	165
	(-6.62)	(10.84)	(6.24)			
15-18	-3.40***	16.01***	8.62***	67.38***	0.797	173
	(-7.15)	(11.94)	(6.72)			
18-21	-3.44***	17.02***	9.25***	60.31***	0.757	172
	(-6.79)	(14.18)	(6.34)			
21-24	-2.75***	16.59***	7.77***	40.37***	0.650	174
	(-3.06)	(9.50)	(3.10)			
24-27	-3.04**	17.07***	8.86***	40.23***	0.671	102
	(-2.62)	(6.83)	(3.00)			
27-30	-1.86	16.01***	5.83	22.08***	0.392	96
	(-1.08)	(3.48)	(1.28)			
30-33	-2.13*	16.64***	6.92**	51.44***	0.627	92
	(-1.94)	(7.05)	(2.57)			
33-36	-0.98	13.23***	4.34	22.29***	0.555	90
	(-0.76)	(4.46)	(1.16)			
Total	-2.38	14.41	6.52		0.550	1727

This table shows the results from the structural-break regressions for monhtly realized skewness in the US. Regressions are run for each maturity individually and the row "Total" reports regression results from pooling all maturities together. The break-date parameter T_b is set to its least-square estimate uniformly for all regressions. The OLS estimator of T_b is given by $\hat{T_b} = argmin_{1<T_b < T-1} SSR(T_b)$, where $SSR(T_b) = \sum_{t=1}^{T}\sum_{T=1}^{N}\left(M_{t,T} - \hat{\alpha} - \widehat{\beta_0}\left(\frac{1}{1+f_{t,T}}\right) I_{(t<T_b)} + \widehat{\beta_1}\left(\frac{1}{1+f_{t,T}}\right) I_{(t\geq T_b)}\right)^2$. To assess significance of the estimated breaks we consider the SupF-test proposed by Andrews (1993), Andrews and Ploberger (1994) and Bai and Perron (1998). The table displays the regression coefficients with t-statistics in parentheses, R^2s and the number of observations. Column F reports the statistic of the F-test run under the null of no structural change against the alternative hypothesis of a break at T_b. Error-covariances for computation of the t-statistics and F-tests are estimated accounting for heteroscedasticity and serial correlation up to 12 lags (i.e. one year) using the approach of Newey and West (1987). *, ** and *** denote significance at the 10%, 5%, and 1% levels, respectively. Critical values used to assess significance of the SupF-Test are taken from Bai and Perron (1998).

Table A.6: Structural break regression: EU realized skewness

Maturity (in months)	α	β_0	β_1	F	R^2	Observations
3-6	-0.29	1.53**	-0.03	10.67**	0.115	122
	(-0.76)	(2.38)	(-0.12)			
6-9	-0.53	2.42***	0.16	19.66***	0.168	143
	(-1.10)	(3.16)	(0.49)			
9-12	-0.13	2.09**	-0.15	10.28**	0.142	165
	(-0.23)	(2.17)	(-0.32)			
12-15	0.32	1.87*	-0.32	8.11*	0.152	172
	(0.60)	(1.81)	(-0.73)			
15-18	0.45	2.28**	-0.40	9.79**	0.157	182
	(0.83)	(2.02)	(-0.79)			
18-21	0.33	3.11***	-0.24	34.34***	0.154	178
	(0.70)	(3.90)	(-0.49)			
21-24	0.11	3.28**	0.12	9.92**	0.184	174
	(0.23)	(2.50)	(0.22)			
Total	0.17	1.99	-0.21		0.120	1186

This table shows the results from the structural-break regressions for monthly realized skewness in the Eurozone. Regressions are run for each maturity individually and the row "Total" reports regression results from pooling all maturities together. The break-date parameter T_b is set to its least-square estimate uniformly for all regressions. The OLS estimator of T_b is given by $\hat{T}_b = argmin_{1<T_b<T-1} SSR(T_b)$, where $SSR(T_b) = \sum_{t=1}^{T}\sum_{T=1}^{N}\left(M_{t,T} - \hat{\alpha} - \widehat{\beta_0}\left(\frac{1}{1+f_{t,T}}\right)I_{(t<T_b)} + \widehat{\beta_1}\left(\frac{1}{1+f_{t,T}}\right)I_{(t\geq T_b)}\right)^2$. To assess significance of the estimated breaks we consider the SupF-test proposed by Andrews (1993), Andrews and Ploberger (1994) and Bai and Perron (1998). The table displays the regression coefficients with t-statistics in parentheses, R^2s and the number of observations. Column F reports the statistic of the F-test run under the null of no structural change against the alternative hypothesis of a break at T_b. Error-covariances for computation of the t-statistics and F-tests are estimated accounting for heteroscedasticity and serial correlation up to 12 lags (i.e. one year) using the approach of Newey and West (1987). *, ** and *** denote significance at the 10%, 5%, and 1% levels, respectively. Critical values used to assess significance of the SupF-Test are taken from Bai and Perron (1998).

Table A.7: Structural break regression: US realized kurtosis

Maturity (in months)	α	β_0	β_1	F	R^2	Observations
3-6	7.46***	4.69*	-5.67*	18.25***	0.119	127
	(4.50)	(1.93)	(-1.78)			
6-9	6.61***	9.83***	-2.53	11.75**	0.112	148
	(2.65)	(2.89)	(-0.49)			
9-12	6.49	13.55**	-2.54	10.79**	0.105	168
	(1.60)	(2.07)	(-0.29)			
12-15	2.21	26.32***	7.42	21.01***	0.262	165
	(0.81)	(3.43)	(1.20)			
15-18	-0.41	39.72***	13.85*	24.92***	0.253	173
	(-0.13)	(4.41)	(1.90)			
18-21	0.94	39.36***	11.56	21.30***	0.259	172
	(0.26)	(3.39)	(1.24)			
21-24	-1.52	47.64***	17.43*	20.12***	0.258	174
	(-0.42)	(3.44)	(1.75)			
24-27	19.78***	4.41	-41.71**	38.91***	0.349	102
	(3.07)	(0.37)	(-2.36)			
27-30	20.74**	7.26	-46.81	35.58***	0.350	96
	(2.03)	(0.32)	(-1.59)			
30-33	19.90*	11.37	-45.67	27.61***	0.409	92
	(1.78)	(0.45)	(-1.34)			
33-36	3.51	42.01**	6.35	10.59**	0.238	90
	(0.47)	(2.42)	(0.26)			
Total	5.50	21.98	0.03		0.180	1727

This table shows the results from the structural-break regressions for monthly realized kurtosis in the US. Regressions are run for each maturity individually and the row "Total" reports regression results from pooling all maturities together. The break-date parameter T_b is set to its least-square estimate uniformly for all regressions. The OLS estimator of T_b is given by $\hat{T_b} = argmin_{1<T_b<T-1} SSR(T_b)$, where $SSR(T_b) = \sum_{t=1}^{T}\sum_{T=1}^{N}\left(M_{t,T} - \hat{\alpha} - \widehat{\beta_0}\left(\frac{1}{1+f_{t,T}}\right) I_{(t<T_b)} + \widehat{\beta_1}\left(\frac{1}{1+f_{t,T}}\right) I_{(t\geq T_b)}\right)^2$. To assess significance of the estimated breaks we consider the SupF-test proposed by Andrews (1993), Andrews and Ploberger (1994) and Bai and Perron (1998). The table displays the regression coefficients with t-statistics in parentheses, R^2s and the number of observations. Column F reports the statistic of the F-test run under the null of no structural change against the alternative hypothesis of a break at T_b. Error-covariances for computation of the t-statistics and F-tests are estimated accounting for heteroscedasticity and serial correlation up to 12 lags (i.e. one year) using the approach of Newey and West (1987). *, ** and *** denote significance at the 10%, 5%, and 1% levels, respectively. Critical values used to assess significance of the SupF-Test are taken from Bai and Perron (1998).

Table A.8: Structural break regression: EU realized kurtosis

Maturity (in months)	α	β_0	β_1	F	R^2	Observations
3-6	4.50^{***}	3.42^{*}	0.52	5.03	0.044	122
	(4.38)	(1.88)	(0.76)			
6-9	3.53^{***}	9.94^{***}	1.19	53.88^{***}	0.145	143
	(2.62)	(4.91)	(1.26)			
9-12	4.2^{**}	11.91^{***}	2.71^{*}	15.44^{***}	0.122	165
	(2.17)	(3.72)	(1.88)			
12-15	3.99^{**}	17.33^{***}	3.43^{***}	34.50^{***}	0.159	172
	(2.19)	(5.40)	(2.61)			
15-18	3.18	24.02^{***}	4.71^{**}	19.81^{***}	0.156	182
	(1.31)	(4.32)	(2.42)			
18-21	0.48	34.48^{***}	8.21^{***}	54.35^{***}	0.154	178
	(0.18)	(6.94)	(3.55)			
21-24	1.14	31.07^{***}	9.10^{***}	11.17^{**}	0.20	174
	(0.47)	(3.63)	(3.36)			
Total	4.18	14.95	3.47		0.080	1186

This table shows the results from the structural-break regressions for monthly realized kurtosis in the Eurozone. Regressions are run for each maturity individually and the row "Total" reports regression results from pooling all maturities together. The break-date parameter T_b is set to its least-square estimate uniformly for all regressions. The OLS estimator of T_b is given by $\hat{T}_b = argmin_{1<T_b<T-1} SSR(T_b)$, where $SSR(T_b) = \sum_{t=1}^{T}\sum_{T=1}^{N}\left(M_{t,T} - \hat{\alpha} - \widehat{\beta_0}\left(\frac{1}{1+f_{t,T}}\right) I_{(t<T_b)} + \widehat{\beta_1}\left(\frac{1}{1+f_{t,T}}\right) I_{(t\geq T_b)}\right)^2$. To assess significance of the estimated breaks we consider the SupF-test proposed by Andrews (1993), Andrews and Ploberger (1994) and Bai and Perron (1998). The table displays the regression coefficients with t-statistics in parentheses, R^2s and the number of observations. Column F reports the statistic of the F-test run under the null of no structural change against the alternative hypothesis of a break at T_b. Error-covariances for computation of the t-statistics and F-tests are estimated accounting for heteroscedasticity and serial correlation up to 12 lags (i.e. one year) using the approach of Newey and West (1987). *, ** and *** denote significance at the 10%, 5%, and 1% levels, respectively. Critical values used to assess significance of the SupF-Test are taken from Bai and Perron (1998).

Appendix B

Inflation cyclicality and stock-bond comovement: evidence from news media

B.1.1 News reporting on inflation

Table B.1: Examples of inflation news articles

Publication Date	Outlet	Headline
21.02.2000	NYT	“Surge in oil prices is raising specter of inflation spike”
20.11.2002	NYT	“Threat of deflation is economists’ No. 1 worry”
11.05.2004	CNN	“Back to the 70’s”
06.02.2005	NYT	“Something in the Air: Inflation”
15.06.2006	NYT	“A Modest Rise Still Amplifies Inflation Fears”
07.06.2007	NYT	“More Fears on Inflation Hit Stocks”
18.12.2008	CNN	“The growing threat of deflation”
03.06.2009	WSJ	“Is Your Portfolio Ready for Hyperinflation?”
29.08.2010	WSJ	“Where to Find Shelter if Deflation Takes Hold”
16.02.2011	FT	“Investors rush to hedge against inflation threat”
25.02.2011	FT	“A toxic combination of inflation and deflation”
26.08.2011	Forbes	“Stagflation Is Knocking On The Door”
03.12.2013	FT	“We need to talk about deflation, again”
19.01.2014	FT	“Deflation ‘ogre’ is latest monster to scare global economy”
28.04.2014	Forbes	“Is U.S. Hyperinflation Imminent?”
16.10.2014	WSJ	“Risk of Deflation Feeds Global Fears”
28.05.2015	WSJ	“World War D – Deflation”
07.03.2016	WSJ	“Rising Global Debt and the Deflation Threat”
05.04.2018	WSJ	“The Big Risk of a Trade War: Inflation”
01.05.2019	WP	“Forget inflation. We should be concerned about deflation.”
05.04.2020	WP	“Opinion: Here comes the great deflation threat”
26.04.2020	FT	“The deflation threat from the virus will be long lasting”
12.05.2020	WSJ	“Specter of Deflation Rises Again”
19.05.2020	FT	“Deflation is the real killer of prosperity”
28.05.2020	Forbes	“Hyperinflation: Will America Dodge The Bullet?”
10.12.2020	WSJ	“Goodbye Covid, Hello Inflation?”

This table shows selected news headlines for inflation related articles in our dataset.

B.1.2 Model details

The dynamics of consumption growth, inflation, dividend growth and the nominal-real correlation are given by:

$$\begin{aligned}
\Delta c_{t+1} &= \mu_c + x_{c,t} + \sigma_c \eta_{c,t+1} \\
\Delta d_{t+1} &= \mu_d + \phi_d x_{c,t} + \sigma_d \eta_{d,t+1} \\
\pi_{t+1} &= \mu_\pi + x_{\pi,t} + \sigma_\pi \eta_{\pi,t+1} \\
x_{c,t+1} &= \phi_{xc} x_{c,t} + \sigma_{xc} \eta_{xc,t+1} \\
x_{\pi,t+1} &= \phi_{x\pi} x_{\pi,t} + \sigma_{x\pi} \left(\rho_t \eta_{xc,t+1} + \sqrt{(1-\rho_t^2)} \eta_{x\pi,t+1} \right) \\
\rho_t &= \frac{e^{2\tilde{\rho}_t} - 1}{e^{2\tilde{\rho}_t} + 1} \\
\tilde{\rho}_{t+1} &= \mu_{\tilde{\rho}} + \phi_{\tilde{\rho}} (\tilde{\rho}_t - \mu_{\tilde{\rho}}) + \sigma_{\tilde{\rho}} \eta_{\tilde{\rho},t+1} \\
\eta_{i,t+1} &\sim i.i.d.N(0,1) \quad \text{for} \quad i = \{c, d, \pi, xc, x\pi, \tilde{\rho}\}
\end{aligned} \tag{B.1}$$

We assume a representative investor with recursive preferences over consumption C_t described by the Kreps-Porteus, Epstein-Zin utility function (Kreps and Porteus (1978); Epstein and Zin (1989)):

$$U_t = \left[(1-\delta) C_t^{\frac{1-\gamma}{\theta}} + \delta \left(E_t \left(U_{t+1}^{1-\gamma} \right) \right)^{\frac{1}{\theta}} \right]^{\frac{\theta}{1-\gamma}} \tag{B.2}$$

where $\gamma \geq 0$ is the coefficient of relative risk aversion; $\theta = \frac{1-\gamma}{1-\frac{1}{\psi}}$; $\psi \geq 0$ is the intertemporal elasticity of substitution (EIS); and δ is the time discount factor. The agent prefers early (late) resolution of uncertainty when risk aversion γ is higher (lower) than the reciprocal of EIS ψ. A preference for early resolution and an EIS above one imply that $\theta < 1$. Note that when $\theta = 1$ (i.e., $\gamma = \frac{1}{\psi}$), above recursive preferences collapse to the standard case of time-separable power utility. In this case the agent is indifferent to when the uncertainty of the consumption path is resolved. Only the transitory innovations but not the persistent innovations are then priced. In the long-run risk model the agent prefers early resolution of uncertainty, that is $\gamma > \frac{1}{\psi}$. This permits also innovations to the persistent component $x_{c,t}$ to be priced.
The agent maximizes lifetime utility subject to the budged constraint

$$W_{t+1} = (W_t - C_t) R_{c,t+1} \tag{B.3}$$

where W_t denotes the total wealth of the agent and $R_{c,t+1}$ is the return on the total wealth portfolio. This portfolio pays aggregate consumption as dividends. As shown in Epstein and Zin (1989), the logarithm of the real stochastic discount factor (SDF) is then given by:

$$m_{t+1} = \theta log\delta - \frac{\theta}{\psi}\Delta c_{t+1} - (1-\theta) r_{c,t+1} \tag{B.4}$$

where $r_{c,t+1} = logR_{c,t+1}$ is the log return on the wealth portfolio, i.e. the claim on aggregate consumption. Hence, for $\theta < 1$, the SDF not only depends on consumption growth but also on the return on the aggregate consumption portfolio. This return is not observed in the data. It differs from the observed return of the market portfolio since aggregate consumption is much larger than aggregate dividends. Therefore, we use the Euler equation:

$$E_t\left[M_{t+1}R_{t+1}\right] = E_t\left[e^{(m_{t+1}+r_{t+1})}\right] = 1 \tag{B.5}$$

which holds for any continuous return $r_{t+1} = log(R_{t+1})$ including the return on the wealth portfolio. Together with the exogenous process we specified for consumption growth (Equation (3.3)) this allows us to solve for the unobserved price-to-consumption ratio and derive the equilibrium real prices of any asset in terms of the fundamental state variables.

To obtain an analytical model solution we proceed as Bansal and Yaron (2004) and use a Campbell-Shiller (Campbell and Shiller (1988)) log-linear Taylor expansion for returns of the total wealth portfolio $r_{c,t+1}$ and the aggregate market portfolio $r_{m,t+1}$:

$$r_{c,t+1} = k_0 + k_1 pc_{t+1} - pc_t + \Delta c_{t+1} \tag{B.6}$$

$$r_{m,t+1} = k_0^d + k_1^d pd_{t+1} - pd_t + \Delta d_{t+1} \tag{B.7}$$

where pc_t is the log price-consumption ratio and pd_t is the log price-dividend ratio. k_0, k_1, k_0^d and k_1^d are approximation constants determined by the unconditional means of the price-consumption ratio ($\bar{pc}$) respectively the price-dividend ratio ($\bar{pd}$). They are given by:

$$\begin{aligned} k_1 &= \frac{e^{\bar{pc}}}{e^{\bar{pc}}+1} \qquad k_0 = log(e^{\bar{pc}}+1) - k_1\bar{pc} \\ k_1^d &= \frac{e^{\bar{pd}}}{e^{\bar{pd}}+1} \qquad k_0^d = log(e^{\bar{pd}}+1) - k_1^d\bar{pd} \end{aligned} \tag{B.8}$$

B.1.2.1 Solution for the price-consumption ratio

We conjecture that the equilibrium log price-consumption ratio is affine in the state variables:

$$pc_t = A_0 + A_{xc} x_{c,t} \tag{B.9}$$

To solve for the coefficients, we substitute the expression for pc_t (B.9), the log-linear approximation of $r_{c,t+1}$ (B.6) and the real SDF (B.4) into the Euler Equation (B.5). Since $(m_{t+1} + r_{c,t+1})$ is normally distributed, we can express the Euler equation as

$$E_t\left[(m_{t+1} + r_{c,t+1})\right] + \frac{1}{2} Var_t\left[(m_{t+1} + r_{c,t+1})\right] = 0 \tag{B.10}$$

The coefficient and the constant are then given by:

$$\begin{aligned} A_{xc} &= \frac{1 - \frac{1}{\psi}}{1 - k_1 \phi_{xc}} \\ A_0 &= \frac{1}{1 - k_1}\left(k_0 + log(\delta) + \mu_c + \frac{1}{2}\theta(\sigma_c^2 + A_{xc}^2 k_1^2 \sigma_{xc}^2) - \frac{\mu_c + \theta\sigma_c^2}{\psi} + \frac{\theta\sigma_c^2}{2\psi^2}\right) \end{aligned} \tag{B.11}$$

With the solution for pc_t and the assumed process for Δc_t at hand, we can rewrite the real log SDF in terms of the fundamental state variables and the underlying shocks:

$$m_{t+1} = m_0 + m_{xc} x_{c,t} - \lambda_c \sigma_c \eta_{c,t+1} - \lambda_{xc} \sigma_{xc} \eta_{xc,t+1} \tag{B.12}$$

with

$$\begin{aligned} m_0 &= \theta log\delta - \frac{\theta\mu_c}{\psi} + (\theta - 1)(k_0 + A_0(k_1 - 1) + \mu_c) \\ m_{xc} &= -\frac{1}{\psi} \\ \lambda_c &= (1 - \theta) + \frac{\theta}{\psi} = \gamma \\ \lambda_{xc} &= k_1 \frac{\gamma - \frac{1}{\psi}}{1 - k_1 \phi_{xc}} \end{aligned} \tag{B.13}$$

The price of short-run consumption risk λ_c equals the risk-aversion coefficient γ. When the agent has a preference for early resolution, that is when $\gamma > \frac{1}{\psi}$, the market price for persistent consumption growth risk λ_{xc} is positive. States of low expected consumption growth are bad states and discounted less heavily. This is distinct from the standard power-utility case ($\gamma = \frac{1}{\psi}$):

here $\lambda_{xc} = 0$, hence there is no compensation for long-run consumption risk.
The nominal log SDF is the real one minus the inflation rate:

$$m^{\$}_{t+1} = m_{t+1} - \pi_{t+1} = m_0 - \mu_\pi + m_{xc}x_{c,t} - x_{\pi,t} - \lambda_c\sigma_c\eta_{c,t+1} - \lambda_{xc}\sigma_{xc}\eta_{xc,t+1} - \sigma_\pi\eta_{\pi,t+1} \quad \text{(B.14)}$$

B.1.2.2 The pricing of inflation risk

Expected inflation x_π and the nominal-real correlation ρ_t do not *separately* enter the real stochastic discount factor. Therefore, the inflation dynamics we assume do not alter the pricing of real assets. In fact, the model implications for real assets are exactly the same as in Bansal and Yaron (2004). However, this does by no means imply that expected inflation is riskless. The risk of expected inflation arises because it is correlated (with correlation ρ_t) with long-run consumption growth. Therefore expected inflation risk carries the same risk premium as long-run consumption risk, pre-multiplied by ρ_t:

$$- cov_t(-x_{\pi,t+1}, m_{t+1}) = -\lambda_{xc}\sigma_{xc}\sigma_{x\pi}\rho_t \quad \text{(B.15)}$$

When the price of long-run risk is positive ($\gamma > \frac{1}{\psi}$), a negative relation between expected growth and expected inflation ($\rho_t < 0$) implies a positive inflation risk premium. Assets with low real payoffs in states of high inflation (and low expected consumption) are risky and demand a compensation. For $\rho_t > 0$, these assets provide a hedge as they tend to have high payoffs in states of low inflation and low expected consumption growth. In this case, inflation will carry a negative risk premium. This mechanism makes expected inflation and the nominal-real correlation crucial for pricing nominal assets and their relation to real assets in our model. Transitory innovations to inflation ($\eta_{\pi,t+1}$) are not priced in our model because they are assumed to be uncorrelated with the real economy.

B.1.2.3 Solution for the stock market return

Analogously to the solution for the price-consumption ratio, we conjecture that the price-dividend ratio is also affine in the state variables:

$$pd_t = A^d_0 + A^d_{xc}x_{c,t} \quad \text{(B.16)}$$

Using the Euler equation for the aggregate market return $r_{m,t+1}$ we obtain:

$$
\begin{aligned}
A_{xc}^d &= \frac{\phi_d - \frac{1}{\psi}}{1 - k_1^d \phi_{xc}} \\
A_0^d &= \frac{1}{1 - k_1^d} \{k_0 + log(\delta) + \mu_d + \\
&\frac{1}{2}\left[(1-\theta)\sigma_c^2 + \sigma_d^2 + \left[\left(A_{xc}k_1 - A_{xc}^d k_1^d\right)^2 + \theta\left(A_{xc}^d k_1^d\right)^2\right]\sigma_{xc}^2\right] \\
&-\frac{\mu_c}{\psi} + \frac{\theta\sigma_c^2}{2\psi^2}\Bigg\}
\end{aligned}
\tag{B.17}
$$

Substituting this solution into (B.7), the return of the aggregate market portfolio is given by:

$$
r_{m,t+1} = r_0 + H_{xc}x_{c,t} + \sigma_{xc}\beta_{xc}\eta_{xc,t+1} + \sigma_d\beta_d\eta_{d,t+1} \tag{B.18}
$$

with

$$
r_0 = k_0^d - A_0^d(1 - k_1^d) + \mu_d \qquad H_{xc} = \frac{1}{\psi} \tag{B.19}
$$

$$
\beta_{xc} = k_1^d \frac{\phi_d - \frac{1}{\psi}}{1 - k_1^d \phi_{xc}} \qquad \beta_d = 1 \tag{B.20}
$$

When the EIS and the dividend leverage ratio are sufficiently high ($\phi_d > \frac{1}{\psi}$), stock market returns are positively related to expected consumption growth. With the expression for stock returns, we can calculate the covariance between stock returns and expected inflation:

$$
cov_t(r_{m,t+1}, x_{\pi,t+1}) = \beta_{xc}\sigma_{xc}\sigma_{x\pi}\rho_t \tag{B.21}
$$

Given that the stock-market positively loads on long-run risk ($\beta_{xc} > 0$), a positive (negative) relationship between expected inflation and persistent consumption growth implies a positive (negative) stock-inflation correlation.

B.1.2.4 Solution for bond returns

To obtain analytical solutions for n-period bond prices we consider a linear approximation of the nominal-real correlation process around the unconditional mean ($\mu_{\tilde{\rho}}$) of $\tilde{\rho}_t$:

$$\begin{aligned}\rho_{t+1} &\approx \frac{e^{2\mu_{\tilde{\rho}}}-1}{e^{2\mu_{\tilde{\rho}}}+1}+\frac{4e^{2\mu_{\tilde{\rho}}}}{(e^{2\mu_{\tilde{\rho}}}+1)^2}(\tilde{\rho}_{t+1}-\mu_{\tilde{\rho}})\\ &=\frac{e^{2\mu_{\tilde{\rho}}}-1}{e^{2\mu_{\tilde{\rho}}}+1}(1-\phi_{\tilde{\rho}})+\phi_{\tilde{\rho}}\rho_t+\frac{4e^{2\mu_{\tilde{\rho}}}}{(e^{2\mu_{\tilde{\rho}}}+1)^2}\sigma_{\tilde{\rho}}\eta_{\tilde{\rho},t+1}\\ &=\mu_\rho+\phi_\rho\rho_t+\sigma_\rho\eta_{\tilde{\rho},t+1}\end{aligned} \tag{B.22}$$

where $\mu_\rho = \frac{e^{2\mu_{\tilde{\rho}}}-1}{e^{2\mu_{\tilde{\rho}}}+1}(1-\phi_{\tilde{\rho}})$, $\sigma_\rho = \frac{4e^{2\mu_{\tilde{\rho}}}}{(e^{2\mu_{\tilde{\rho}}}+1)^2}\sigma_{\tilde{\rho}}$, and $\phi_\rho = \phi_{\tilde{\rho}}$
Real and nominal log bond prices then satisfy

$$p^r_{n,t} = E\left[e^{m_{t+1}+p_{n-1,t+1}}\right] = B_{0,n} + B_{xc,n}x_{c,t} + B_{x\pi,n}x_{\pi,t} + B_{\rho,n}\rho_t \tag{B.23}$$

$$p^\$_{n,t} = E\left[e^{m^\$_{t+1}+p^\$_{n-1,t+1}}\right] = B^\$_{0,n} + B^\$_{xc,n}x_{c,t} + B^\$_{x\pi,n}x_{\pi,t} + B^\$_{\rho,n}\rho_t \tag{B.24}$$

where $p_{n,t} = logP_{n,t}$ is the time-t log price of an n-period zero-coupon bond paying one unit of numeraire in n periods. The bond coefficients are given by

$$
\begin{aligned}
B_{0,n} &= B_{0,n-1} + m_0 + \frac{1}{2}\left[\lambda_c^2\sigma_c^2 + (B_{xc,n-1} - \lambda_{xc})^2\sigma_{xc}^2)\right] \\
B_{xc,n} &= -\frac{1}{\psi} + B_{xc,n-1}\phi_{xc} \\
&= -\frac{1}{\psi}\sum_{k=1}^{n}\phi_{xc}^{k-1} \\
B_{x\pi,n} &= 0 \\
B_{\rho,n} &= 0 \\
B_{0,n}^{\$} &= B_{0,n-1}^{\$} + m_0 - \mu_\pi + \mu_\rho B_{\rho,n-1}^{\$} \\
&\quad + \frac{1}{2}\left[\lambda_c^2\sigma_c^2 + (B_{xc,n-1}^{\$} - \lambda_{xc})^2\sigma_{xc}^2) + (B_{x\pi,n-1}^{\$})^2\sigma_{x\pi}^2 + (B_{\rho,n-1}^{\$})^2\sigma_\rho^2\right] \\
B_{xc,n}^{\$} &= -\frac{1}{\psi} + B_{xc,n-1}^{\$}\phi_{xc} \\
&= -\frac{1}{\psi}\sum_{k=1}^{n}\phi_{xc}^{k-1} \\
&= B_{xc,n} \\
B_{x\pi,n}^{\$} &= -1 + B_{x\pi,n-1}^{\$}\phi_{x\pi} \\
&= -\sum_{k=1}^{n}\phi_{x\pi}^{k-1} \\
B_{\rho,n}^{\$} &= \phi_\rho B_{\rho,n-1}^{\$} + \sigma_{xc}\sigma_{x\pi}B_{x\pi,n-1}^{\$}(B_{xc,n-1}^{\$} - \lambda_{xc}) \\
&= \sigma_{xc}\sigma_{x\pi}\frac{1}{\psi}\sum_{k=2}^{n}\left[\phi_\rho^{k-2}\left(\sum_{z=1}^{k}\phi_{x\pi}^{z-1}\left(\sum_{z=1}^{k}\phi_{xc}^{z-1} - \lambda_{xc}\right)\right)\right]
\end{aligned}
\tag{B.25}
$$

The real short rate increases in response to higher expected consumption growth. This results from the agents' desire to smooth consumption intertemporally. Faced with an increase in future consumption the agent wants to borrow to increase consumption today. In equilibrium, zero net supply of assets requires the real short rate to increase to discourage borrowing. The loading of longer dated bond yields on excpected consumption growth additionaly depends on the persistence of the growth process (ϕ_{xc}) and decreases with maturity.

The response of the nominal short rate to expected consumption growth depends on the variable ρ_t. Partly, it incorporates the same mechanism which connects real yields to expected consumption. Moreover, nominal yields increase with expected inflation according to a standard Fisher equation. As inflation expectations are correlated with expected consumption, the total reaction of nominal yields to innovations in expected consumption is a combination of both effects. If $\rho_t > -\frac{\sigma_{xc}B_{xc,n}}{\sigma_{x\pi}B_{x\pi,n}^{\$}}$, the real channel dominates and yields increase with expected consumption growth. If $\rho_t < -\frac{\sigma_{xc}B_{xc,n}}{\sigma_{x\pi}B_{x\pi,n}^{\$}}$, the nominal channel dominates as the reaction to decreases in expected inflation associated with increases in expected growth surpass the real effect. For longer dated bond yields the sensitivity to excpected inflation $B^{\$}_{x\pi,n}$ decreases with maturity depending on the persistence of the inflation process ($\phi_{x\pi}$). While the level of nominal short rates does not depend on the nominal-real correlation, longer term nominal yields negatively load on ρ_t ($B^{\$}_{\rho,n} > 0 \quad \text{for} \quad n > 1$). This arises because of two effects: Firstly, positive (negative) values of ρ_t increase (decrease) the variance of future short rates as expected inflation and expected consumption growth tend to move in the same (different) direction and both variables positively affect short rates. Because of Jensen's inequality, higher variance of future short rates decreases bond yields. Secondly, ρ_t affects the risk premium of long dated bonds. For positive (negative) values of ρ_t, long term bonds hedge (increase) consumption risk and therefore are priced at a premium (discount).

With the expression for log bond prices, we can calculate the covariance between stock returns and bond yields:

$$cov_t(r_{m,t+1}, y^r_{n,t+1}) = -\frac{1}{n}\left(\sigma^2_{xc}\beta_{xc}B_{xc,n}\right) \tag{B.26}$$

$$cov_t(r_{m,t+1}, y^{\$}_{n,t+1}) = -\frac{1}{n}\left(\sigma^2_{xc}\beta_{xc}B_{xc,n} + \sigma_{x\pi}\sigma_{xc}\beta_{xc}B^{\$}_{x\pi,n}\rho_t\right) \tag{B.27}$$

Given that the stock market loads positively on expected consumption growth ($\beta_{xc} > 0$), stock returns and real bond yields are positively correlated (as $B_{xc,n} < 0$). The correlation of nominal bond yields with the stock market additionally depends on the sign and magnitude of ρ_t.

B.1.2.5 Distant horizon asset correlations

From the linear approximation B.22 of ρ_t we can express ρ_{t+h} as

$$\begin{aligned}\rho_{t+h} &\approx \mu_\rho + \phi_\rho\rho_{t+h-1} + \sigma_\rho\eta_{\tilde{\rho},t+h} \\ &= \mu_\rho\sum_{i=1}^{h}\phi_\rho^{i-1} + \phi_\rho^h\rho_t + \sigma_\rho\sum_{i=1}^{h}\phi_\rho^{h-i}\eta_{\tilde{\rho},t+i}\end{aligned} \tag{B.28}$$

Plugging this expression into Equations B.21 and B.27, we obtain the relationships between ρ_t and asset correlations at distant horizons:

$$
\begin{aligned}
cov_{t+h}(r_{m,t+h+1}, x_{\pi,t+h+1}) \approx & (\beta_{xc}\sigma_{xc}\sigma_{x\pi})\mu_\rho \sum_{i=1}^{h} \phi_\rho^{i-1} \\
& + (\beta_{xc}\sigma_{xc}\sigma_{x\pi})\phi_\rho^h \rho_t \\
& + (\beta_{xc}\sigma_{xc}\sigma_{x\pi})\sigma_\rho \sum_{i=1}^{h} \phi_\rho^{h-i} \eta_{\tilde{\rho},t+i}
\end{aligned} \tag{B.29}
$$

$$
\begin{aligned}
cov_{t+h}(r_{m,t+h+1}, y^{\$}_{n,t+h+1}) \approx -\frac{1}{n} \Bigg[& \sigma_{xc}^2 \beta_{xc} B_{xc,n} + (\sigma_{x\pi}\sigma_{xc}\beta_{xc}B^{\$}_{x\pi,n})\mu_\rho \sum_{i=1}^{h} \phi_\rho^{i-1} \\
& + (\sigma_{x\pi}\sigma_{xc}\beta_{xc}B^{\$}_{x\pi,n})\phi_\rho^h \rho_t \\
& + (\sigma_{x\pi}\sigma_{xc}\beta_{xc}B^{\$}_{x\pi,n})\sigma_\rho \sum_{i=1}^{h} \phi_\rho^{h-i} \eta_{\tilde{\rho},t+i} \Bigg]
\end{aligned} \tag{B.30}
$$

Appendix C

TIPS under FIRE - Formation of inflation expectations in bond markets

C.1.1 Model derivations

From Equation (4.1) in the main text, we assume that agents form forecasts $F_t^{\mathbb{P}}(\pi_{t+h})$ in period t for realized inflation π_{t+h} in period $t+h$ according to

$$F_t^{\mathbb{P}}(\pi_{t+h}) = (1-g)E_t(\pi_{t+h}) + gF_{t-1}^{\mathbb{P}}(\pi_{t+h}) \tag{C.1}$$

where $E_t(\pi_{t+h})$ is the objective expectation of π_{t+h} based in information in t. The indexation by $\mathbb{P}$ indicates that these forecasts are formed under the real world measure. Expanding this recursive definition directly relates agents' subjective forecast to past and present objective forecasts:

$$F_t^{\mathbb{P}}(\pi_{t+h}) = (1-g)\sum_{i=0}^{\infty} g^i E_{t-i}(\pi_{t+h}) \tag{C.2}$$

Forecast revisions are defined as

$$\Delta F_t^{\mathbb{P}}(\pi_{t+h}) = F_t^{\mathbb{P}}(\pi_{t+h}) - F_{t-1}^{\mathbb{P}}(\pi_{t+h}) \tag{C.3}$$

Plugging in expression (C.2), they can be expressed in terms of past and present changes in the objective forecasts:

$$\Delta F_t^{\mathbb{P}}(\pi_{t+h}) = (1-g)\left[\sum_{i=0}^{\infty} g^i E_{t-i}(\pi_{t+h}) - \sum_{i=0}^{\infty} g^i E_{t-i-1}(\pi_{t+h})\right] \tag{C.4}$$

$$= (1-g)\sum_{i=0}^{\infty} g^i \Delta E_{t-i}(\pi_{t+h}) \tag{C.5}$$

Forecast errors are defined as

$$Err^{\mathbb{P}}_{t+h,t} = \pi_{t+h} - F^{\mathbb{P}}_t(\pi_{t+h}) \tag{C.6}$$
$$= E_t(\pi_{t+h}) + \epsilon_{t+h,t} - F^{\mathbb{P}}_t(\pi_{t+h}) \tag{C.7}$$

where $\epsilon_{t+h,t} = \pi_{t+h} - E_t(\pi_{t+h})$ is the unforecastable component of realized inflation. Plugging in Equation (C.1) relates agents' forecast errors to agents' forecast revisions:

$$Err^{\mathbb{P}}_{t+h,t} = \frac{g}{1-g}\Delta F^{\mathbb{P}}_t(\pi_{t+h}) + \epsilon_{t+h,t} \tag{C.8}$$

Under the risk neutral measure $\mathbb{Q}$, forecasts additionally contain a risk premium component and are therefore given by

$$F^{\mathbb{Q}}_t(\pi_{t+h}) = F^{\mathbb{P}}_t(\pi_{t+h}) + RP_t \tag{C.9}$$
$$= (1-g)E_t(\pi_{t+h}) + gF^{\mathbb{P}}_{t-1}(\pi_{t+h}) + RP_t \tag{C.10}$$
$$= (1-g)\sum_{i=0}^{\infty} g^i E_{t-i}(\pi_{t+h}) + RP_t \tag{C.11}$$

Equivalently, forecast revisions and forecast errors under $\mathbb{Q}$ are given by

$$\Delta F^{\mathbb{Q}}_t(\pi_{t+h}) = F^{\mathbb{Q}}_t(\pi_{t+h}) - F^{\mathbb{Q}}_{t-1}(\pi_{t+h}) \tag{C.12}$$
$$= \Delta F^{\mathbb{P}}_t(\pi_{t+h}) + \Delta RP_t \tag{C.13}$$
$$= (1-g)\sum_{i=0}^{\infty} g^i \Delta E_{t-i}(\pi_{t+h}) + \Delta RP_t \tag{C.14}$$
$$Err^{\mathbb{Q}}_{t+h,t} = \pi_{t+h} - F^{\mathbb{Q}}_t(\pi_{t+h}) \tag{C.15}$$
$$= \pi_{t+h} - F^{\mathbb{P}}_t(\pi_{t+h}) - RP_t \tag{C.16}$$
$$= E_t(\pi_{t+h}) + \epsilon_{t+h,t} - F^{\mathbb{P}}_t(\pi_{t+h}) - RP_t \tag{C.17}$$
$$= \frac{g}{1-g}\Delta F^{\mathbb{P}}_t(\pi_{t+h}) + \epsilon_{t+h,t} - RP_t \tag{C.18}$$

The risk premium component is assumed to evolve according to

$$RP_t = \theta RP_{t-1} + z_t \tag{C.19}$$

which can be expressed in terms of the shocks z_t:

$$RP_t = \sum_{i=0}^{\infty} \theta^i z_{t-i} \tag{C.20}$$

The change in the risk premium component can then be written as

$$\begin{aligned}
\Delta RP_t &= RP_t - RP_{t-1} && \text{(C.21)}\\
&= (\theta - 1)RP_{t-1} + z_t && \text{(C.22)}\\
&= (\theta - 1)\sum_{i=0}^{\infty} \theta^i z_{t-i-1} + z_t && \text{(C.23)}\\
&= \theta \Delta RP_{t-1} + z_t - z_{t-1} && \text{(C.24)}
\end{aligned}$$

All shocks are assumed to have zero mean and to be serially independent and identically distributed, i.e

$$\begin{aligned}
Var(z_{t+i}) &= \sigma_z^2 \quad \forall i\\
Var(\Delta E_{t+i}(\pi_{t+h})) &= \sigma_{\Delta E}^2 \quad \forall i, h\\
Cov(z_{t+i}, \Delta E_{t+j}(\pi_{t+h}))) &= 0 \quad \forall h, i \neq j\\
Cov(z_{t+i}, z_{t+j}) &= 0 \quad \forall i \neq j\\
Cov(\Delta E_{t+i}(\pi_{t+h})), \Delta E_{t+j}(\pi_{t+k}))) &= 0 \quad \forall h, k, i \neq j
\end{aligned} \tag{C.25}$$

but may be contemporaneously correlated with constant covariance

$$Cov(z_{t+i}, \Delta E_{t+i}(\pi_{t+h})) = Cov(z, \Delta E) = \rho(z, \Delta E)\sigma_z \sigma_{\Delta E} \quad \forall i, h \tag{C.26}$$

where $\rho(z, \Delta E)$ denotes the correlation between z_t and $\Delta E_t(\pi_{t+h})$.

Note that the unforecastable component of realized inflation, $\epsilon_{t+h,t}$, can be expressed as

$$\epsilon_{t+h,t} = \pi_{t+h} - E_t(\pi_{t+h}) = E_{t+h}(\pi_{t+h}) - E_t(\pi_{t+h}) = \sum_{i=0}^{h-1} \Delta E_{t+h-i}(\pi_{t+h}) \tag{C.27}$$

C.1.1.1 Error-revision relation

Using Equations (C.18) and (C.13), the covariance between risk neutral forecast errors and forecast revisions can be expressed as

$$Cov(Err^{\mathbb{Q}}_{t+h,t}, \Delta F^{\mathbb{Q}}_t(\pi_{t+h})) = Cov\left(\left[\frac{g}{1-g}\Delta F^{\mathbb{P}}_t(\pi_{t+h}) + \epsilon_{t+h,t} - RP_t\right], \left[\Delta F^{\mathbb{P}}_t(\pi_{t+h}) + \Delta RP_t\right]\right) \tag{C.28}$$

$$\begin{aligned} =&\frac{g}{1-g}\sigma^2_{\Delta F^{\mathbb{P}}} - Cov(\Delta F^{\mathbb{P}}_t(\pi_{t+h}), RP_t) \\ &+ \frac{g}{1-g}Cov(\Delta F^{\mathbb{P}}_t(\pi_{t+h}), \Delta RP_t) - Cov(RP_t, \Delta RP_t) \end{aligned} \tag{C.29}$$

where $\sigma^2_{\Delta F^{\mathbb{P}}} = Var(\Delta F^{\mathbb{P}}_t(\pi_{t+h}))$ denotes the variance of risk premia changes.

We use Equations (C.5) and (C.23) to calculate the term $Cov(\Delta F^{\mathbb{P}}_t(\pi_{t+h}), \Delta RP_t)$ as:

$$\begin{aligned} Cov(\Delta F^{\mathbb{P}}_t(\pi_{t+h}), \Delta RP_t) &= Cov\left(\left[(1-g)\sum_{i=0}^{\infty} g^i \Delta E_{t-i}(\pi_{t+h})\right], \left[(\theta-1)\sum_{i=0}^{\infty}\theta^i z_{t-i-1} + z_t\right]\right) \\ &= (1-g)\left(Cov(z_t, \Delta E_t(\pi_{t+h})) + (\theta-1)\sum_{i=1}^{\infty}\theta^{i-1}g^i Cov(z_{t-i}, \Delta E_{t-i}(\pi_{t+h}))\right) \\ &= \frac{(1-g)^2}{(1-\theta g)}Cov(z, \Delta E) \end{aligned} \tag{C.30}$$

The term $Cov(\Delta F^{\mathbb{P}}_t(\pi_{t+h}), RP_t)$ can be calculated with Equations (C.5) and (C.20) as:

$$\begin{aligned} Cov(\Delta F^{\mathbb{P}}_t(\pi_{t+h}), RP_t) &= Cov\left(\left[(1-g)\sum_{i=0}^{\infty} g^i \Delta E_{t-i}(\pi_{t+h})\right], \left[\sum_{i=0}^{\infty}\theta^i z_{t-i}\right]\right) \\ &= (1-g)\sum_{i=0}^{\infty}\theta^i g^i Cov(z_{t-i}, \Delta E_{t-i}(\pi_{t+h})) \\ &= \frac{(1-g)}{(1-\theta g)}Cov(z, \Delta E) \end{aligned} \tag{C.31}$$

And the term $Cov(RP_t, \Delta RP_t)$ can be expressed, using Equations (C.20) and (C.23), as:

$$\begin{aligned} Cov(RP_t, \Delta RP_t) &= Cov\left(\left[\sum_{i=0}^{\infty}\theta^i z_{t-i}\right], \left[(\theta-1)\sum_{i=0}^{\infty}\theta^i z_{t-i-1} + z_t\right]\right) \\ &= \left(1 + (\theta-1)\sum_{i=1}^{\infty}\theta^i\theta^{i-1}\right)\sigma_z^2 \\ &= \frac{1}{(1+\theta)}\sigma_z^2 \\ &= \frac{1}{2}\sigma_{\Delta RP}^2 \end{aligned} \tag{C.32}$$

where $\sigma_{\Delta RP}^2$ denotes the variance of risk premia changes. The last row follows from the fact that the variance of risk premia changes can be expressed as

$$\begin{aligned} \sigma_{\Delta RP}^2 &= Var(\Delta RP_t) \\ &= Var\left((\theta-1)\sum_{i=0}^{\infty}\theta^i z_{t-i-1} + z_t\right) \\ &= \left(\frac{(\theta-1)^2}{(1-\theta)} + 1\right)\sigma_z^2 \\ &= 2\frac{1}{(1+\theta)}\sigma_z^2 \end{aligned} \tag{C.33}$$

Plugging Equations (C.30), (C.31) and (C.32) into Equation (C.29) yields:

$$Cov(Err_{t+h,t}^{\mathbb{Q}}, \Delta F_t^{\mathbb{Q}}(\pi_{t+h})) = \frac{g}{1-g}\sigma_{\Delta F^{\mathbb{P}}}^2 - \frac{(1-g)^2}{(1-\theta g)}Cov(z, \Delta E) - \frac{1}{2}\sigma_{\Delta RP}^2 \tag{C.34}$$

The variances of forecast revisions under $\mathbb{P}$ and $\mathbb{Q}$ can be calculated from Equations (C.5) and (C.13) as

$$\begin{aligned} \sigma_{\Delta F^{\mathbb{P}}}^2 &= Var(\Delta F_t^{\mathbb{P}}(\pi_{t+h})) \\ &= Var\left((1-g)\sum_{i=0}^{\infty}g^i\Delta E_{t-i}(\pi_{t+h})\right) \\ &= \frac{(1-g)^2}{(1-g^2)}\sigma_{\Delta E}^2 \end{aligned} \tag{C.35}$$

and

$$\begin{aligned}\sigma^2_{\Delta F^{\mathbb{Q}}} &= Var(\Delta F^{\mathbb{Q}}_t(\pi_{t+h})) \\ &= Var(\Delta F^{\mathbb{P}}_t(\pi_{t+h}) + \Delta RP_t) \\ &= Var(\Delta F^{\mathbb{P}}_t(\pi_{t+h})) + Var(\Delta RP_t) + 2Cov(\Delta F^{\mathbb{P}}_t(\pi_{t+h}), \Delta RP_t) \\ &= \sigma^2_{\Delta F^{\mathbb{P}}} + \sigma^2_{\Delta RP} + 2Cov(\Delta F^{\mathbb{P}}_t(\pi_{t+h}), \Delta RP_t) \\ &= \frac{(1-g)^2}{(1-g^2)}\sigma^2_{\Delta E} + \sigma^2_{\Delta RP} + 2\frac{(1-g)^2}{(1-\theta g)}Cov(z, \Delta E)\end{aligned} \tag{C.36}$$

where the last row plugs in Equation (C.35) and(C.30). This allows to express $Cov(Err^{\mathbb{Q}}_{t+h,t}, \Delta F^{\mathbb{Q}}_t(\pi_{t+h}))$ in terms of the variances of the forecast revisions under $\mathbb{P}$ and $\mathbb{Q}$ by substituting (C.35) and (C.36) into (C.34):

$$Cov(Err^{\mathbb{Q}}_{t+h,t}, \Delta F^{\mathbb{Q}}_t(\pi_{t+h})) = \frac{1}{2}\left[\left(\frac{1+g}{1-g}\right)\sigma^2_{\Delta F^{\mathbb{P}}} - \sigma^2_{\Delta F^{\mathbb{Q}}}\right] \tag{C.37}$$

Finally, the beta of risk neutral forecast errors on forecast revisions is given by

$$\begin{aligned}\beta(Err^{\mathbb{Q}}_{t+h,t}, \Delta F^{\mathbb{Q}}_t(\pi_{t+h})) =& \frac{Cov(Err^{\mathbb{Q}}_{t+h,t}, \Delta F^{\mathbb{Q}}_t(\pi_{t+h}))}{\sigma^2_{\Delta F^{\mathbb{Q}}}} \\ =& \frac{1}{2}\left[\left(\frac{1+g}{1-g}\right)\frac{\sigma^2_{\Delta F^{\mathbb{P}}}}{\sigma^2_{\Delta F^{\mathbb{Q}}}}\right] - \frac{1}{2}\end{aligned} \tag{C.38}$$

where the operator $\beta(y, x)$ denotes the regression coefficient of y onto x. Under the assumption of uncorrelated risk premium and expectation shocks ($\rho(z, \Delta E) = 0$), this simplifies to

$$\sigma^2_{\Delta F^{\mathbb{P}}} = \sigma^2_{\Delta F^{\mathbb{Q}}} - \sigma^2_{\Delta RP} \tag{C.39}$$

$$\beta(Err^{\mathbb{Q}}_{t+h,t}, \Delta F^{\mathbb{Q}}_t(\pi_{t+h})|\rho(z, \Delta E) = 0) = \frac{1}{2}\left[\left(\frac{1+g}{1-g}\right)\left(1 - \frac{\sigma^2_{\Delta RP_t}}{\sigma^2_{\Delta F^{\mathbb{Q}}}}\right)\right] - \frac{1}{2} \tag{C.40}$$

C.1.1.2 Error autocorrelation

Subsequent forecast errors over the same forecast horizon are given by

$$Err^{\mathbb{Q}}_{t+h,t} = \frac{g}{1-g}\Delta F^{\mathbb{P}}_t(\pi_{t+h}) + \epsilon_{t+h,t} - RP_t \tag{C.41}$$

$$= g\sum_{i=0}^{\infty} g^i \Delta E_{t-i}(\pi_{t+h}) + \sum_{i=0}^{h-1} \Delta E_{t+h-i}(\pi_{t+h}) - \sum_{i=0}^{\infty} \theta^i z_{t-i} \tag{C.42}$$

$$Err^{\mathbb{Q}}_{t,t-h} = \frac{g}{1-g}\Delta F^{\mathbb{P}}_{t-h}(\pi_t) + \epsilon_{t,t-h} - RP_{t-h} \tag{C.43}$$

$$= g\sum_{i=0}^{\infty} g^i \Delta E_{t-h-i}(\pi_t) + \sum_{i=0}^{h-1} \Delta E_{t-i}(\pi_t) - \sum_{i=0}^{\infty} \theta^i z_{t-h-i} \tag{C.44}$$

For simplicity, we additionally assume a constant correlation between revisions of objective forecasts for different forecast horizons:

$$Cov(\Delta E_{t-i}(\pi_t), \Delta E_{t-i}(\pi_{t+h})) = Cov(\Delta E_t(\pi_{t+h}), \Delta E_t(\pi_{t+2h})) \quad \forall i, h \tag{C.45}$$

$$= \rho_{\Delta E}\sigma^2_{\Delta E} \tag{C.46}$$

With these expressions, the covariance between subsequent risk neutral forecast errors can then be calculated as

$$
\begin{aligned}
Cov(Err^{\mathbb{Q}}_{t+h,t}, Err^{\mathbb{Q}}_{t,t-h}) = g^2 \sum_{i=0}^{\infty} g^i g^{i+h} Cov(\Delta E_{t-h-i}(\pi_t), \Delta E_{t-h-i}(\pi_{t+h})) \\
+ g \sum_{i=0}^{h-1} g^i Cov(\Delta E_{t-i}(\pi_t), \Delta E_{t-i}(\pi_{t+h})) \\
- g \sum_{i=0}^{\infty} g^i \theta^{i+h} Cov(z_{t-h-i}, \Delta E_{t-h-i}(\pi_t)) \\
- \sum_{i=0}^{h-1} \theta^i Cov(z_{t-i}, \Delta E_{t-i}(\pi_t)) \\
- g \sum_{i=0}^{\infty} g^{i+h} \theta^i Cov(z_{t-h-i}, \Delta E_{t-h-i}(\pi_{t+h})) \\
+ \sum_{i=0}^{\infty} \theta^i \theta^{i+h} \sigma_z^2
\end{aligned}
\tag{C.47}
$$

$$
\begin{aligned}
= \left[\frac{g^{h+2}}{1-g^2} + \frac{g(1-g^h)}{1-g} \right] \rho_{\Delta E} \sigma^2_{\Delta E} \\
- \left[\frac{g(g^h + \theta^h)}{1-\theta g} + \frac{1-\theta^h}{1-\theta} \right] Cov(z, \Delta E) \\
+ \frac{1}{2} \frac{\theta^h}{1-\theta} \sigma^2_{\Delta RP}
\end{aligned}
\tag{C.48}
$$

and the variance of risk neutral forecast errors as

$$
Var(Err^{\mathbb{Q}}_{t,t-h}) = \left[\frac{g^2}{1-g^2} + h \right] \sigma^2_{\Delta E} + \frac{1}{2} \frac{1}{1-\theta} \sigma^2_{\Delta RP} - 2 \frac{g}{1-\theta g} Cov(z, \Delta E) \tag{C.49}
$$

The beta coefficient on forecast errors on past forecast errors is given by

$$
\beta(Err^{\mathbb{Q}}_{t+h,t}, Err^{\mathbb{Q}}_{t,t-h}) = \frac{Cov(Err^{\mathbb{Q}}_{t+h,t}, Err^{\mathbb{Q}}_{t,t-h})}{Var(Err^{\mathbb{Q}}_{t,t-h})} \tag{C.50}
$$

Under FIRE ($g = 0$), these expressions simplify to

$$Cov(Err^{\mathbb{Q}}_{t+h,t}, Err^{\mathbb{Q}}_{t,t-h}|g=0) = -\frac{1-\theta^h}{1-\theta}Cov(z,\Delta E) + \frac{1}{2}\frac{\theta^h}{1-\theta}\sigma^2_{\Delta RP} \tag{C.51}$$

$$Var(Err^{\mathbb{Q}}_{t,t-h}|g=0) = h\sigma^2_{\Delta E} + \frac{1}{2}\frac{1}{1-\theta}\sigma^2_{\Delta RP} \tag{C.52}$$

$$\begin{aligned}\beta(Err^{\mathbb{Q}}_{t+h,t}, Err^{\mathbb{Q}}_{t,t-h}|g=0) = \theta^h &\left[\left(\frac{1}{2}\frac{1}{1-\theta}\right)\left(\frac{\sigma^2_{\Delta RP}}{h\sigma^2_{\Delta E} + 0.5\frac{1}{1-\theta}\sigma^2_{\Delta RP}}\right)\right] \\ - \rho(z,\Delta E) &\left[\left(\frac{\sigma_{\Delta E}\sigma_{\Delta RP}\sqrt{0.5(1+\theta)}}{h\sigma^2_{\Delta E} + 0.5\frac{1}{1-\theta}\sigma^2_{\Delta RP}}\right)\left(\frac{1-\theta^h}{1-\theta}\right)\right]\end{aligned} \tag{C.53}$$

Setting $h = 1$ in the last equation gives Equation (4.49) in the main text text for subsequent one-step-ahead forecast errors.

Using the simplifying assumption that risk premia are uncorrelated with objective expectations ($\rho(z,\Delta E) = 0$) and focusing on $h = 1$ we obtain

$$Var(Err^{\mathbb{Q}}_{t,t-h}|\rho(z,\Delta E)=0) = \left[\frac{1}{1-g^2}\right]\sigma^2_{\Delta E} + \frac{1}{2}\frac{1}{1-\theta}\sigma^2_{\Delta RP} \tag{C.54}$$

$$\begin{aligned}Cov(Err^{\mathbb{Q}}_{t+1,t}, Err^{\mathbb{Q}}_{t,t-1}|\rho(z,\Delta E)=0) =& \frac{g}{1-g^2}\rho_{\Delta E}\sigma^2_{\Delta E} + \frac{1}{2}\frac{\theta}{1-\theta}\sigma^2_{\Delta RP} \\ =& g\rho_{\Delta E}Var(Err^{\mathbb{Q}}_{t,t-h}|\rho(z,\Delta E)=0) \\ &- g\rho_{\Delta E}\frac{1}{2}\frac{\theta}{1-\theta}\sigma^2_{\Delta RP} + \frac{1}{2}\frac{\theta}{1-\theta}\sigma^2_{\Delta RP}\end{aligned} \tag{C.55}$$

$$\begin{aligned}\beta(Err^{\mathbb{Q}}_{t+1,t}, Err^{\mathbb{Q}}_{t,t-1}|\rho(z,\Delta E)=0) =& \frac{Cov(Err^{\mathbb{Q}}_{t+1,t}, Err^{\mathbb{Q}}_{t,t-1}|\rho(z,\Delta E)=0)}{Var(Err^{\mathbb{Q}}_{t,t-h}|\rho(z,\Delta E)=0)} \\ =& g\rho_{\Delta E}\left[1 - \left(\frac{1}{2}\frac{1}{1-\theta}\right)\left(\frac{\sigma^2_{\Delta RP}}{\frac{1}{1-g^2}\sigma^2_{\Delta E} + 0.5\frac{1}{1-\theta}\sigma^2_{\Delta RP}}\right)\right] \\ &+ \theta\left[\left(\frac{1}{2}\frac{1}{1-\theta}\right)\left(\frac{\sigma^2_{\Delta RP}}{\frac{1}{1-g^2}\sigma^2_{\Delta E} + 0.5\frac{1}{1-\theta}\sigma^2_{\Delta RP}}\right)\right]\end{aligned} \tag{C.56}$$

Setting $\sigma^2_{\Delta RP} = 0$ in the last equation we obtain Equation (4.16) in the main text applicable for forecast errors under the physical measure. Setting $g = 0$ in the last equation retrieves Equation (4.19) in the main text.

C.1.1.3 Revision autocorrelation

Subsequent forecast revisions are given by

$$\Delta F_t^{\mathbb{Q}}(\pi_{t+h}) = \Delta F_t^{\mathbb{P}}(\pi_{t+h}) + \Delta RP_t \tag{C.57}$$

$$\Delta F_{t+1}^{\mathbb{Q}}(\pi_{t+h}) = (1-g)\Delta E_{t+1}(\pi_{t+h}) + g\Delta F_t^{\mathbb{Q}}(\pi_{t+h}) - g\Delta RP_t + \Delta RP_{t+1} \tag{C.58}$$

With these equations, the covariance between subsequent risk neutral forecast revisions is given by

$$\begin{aligned} Cov(\Delta F_{t+1}^{\mathbb{Q}}(\pi_{t+h}), \Delta F_t^{\mathbb{Q}}(\pi_{t+h})) =& g\left(\sigma^2_{\Delta F^{\mathbb{Q}}} - cov(\Delta F_t^{\mathbb{P}}(\pi_{t+h}), \Delta RP_t) - \sigma^2_{\Delta RP}\right) \\ &+ cov(\Delta F_t^{\mathbb{P}}(\pi_{t+h}), \Delta RP_{t+1}) + cov(\Delta RP_{t+1}, \Delta RP_t) \end{aligned} \tag{C.59}$$

The term $Cov(\Delta F_t^{\mathbb{P}}(\pi_{t+h}), \Delta RP_{t+1})$ can be calculated from Equations (C.1) and (C.24) as

$$\begin{aligned} Cov(\Delta F_t^{\mathbb{P}}(\pi_{t+h}), \Delta RP_{t+1}) &= Cov((1-g)\Delta E_t(\pi_{t+h}) + g\Delta F_{t-1}^{\mathbb{P}}(\pi_{t+h}), \theta\Delta RP_t + z_{t+1} - z_t) \\ &= \theta Cov(\Delta F_t^{\mathbb{P}}(\pi_{t+h}), \Delta RP_t) - (1-g)Cov(z, \Delta E) \end{aligned} \tag{C.60}$$

with $Cov(\Delta F_t^{\mathbb{P}}(\pi_{t+h}), \Delta RP_t)$ given in Equation (C.30).
The term $Cov(\Delta RP_{t+1}, \Delta RP_t)$ derives from Equation (C.23) as

$$\begin{aligned} Cov(\Delta RP_{t+1}, \Delta RP_t) &= \left[(\theta-1)^2 \sum_{i=0}^{\infty} \theta^i \theta^{i+1} + (\theta - 1)\right]\sigma_z^2 \\ &= (\theta-1)\left[\frac{(\theta-1)\theta}{1-\theta^2} + 1\right]\sigma_z^2 \\ &= -\frac{(1-\theta)^2}{(1-\theta^2)}\sigma_z^2 \\ &= -\frac{1}{2}(1-\theta)\sigma^2_{\Delta RP} \end{aligned} \tag{C.61}$$

Substituting (C.60), (C.61) and (C.30) into (C.59) gives

$$
\begin{aligned}
Cov(\Delta F^{\mathbb{Q}}_{t+1}(\pi_{t+h}), \Delta F^{\mathbb{Q}}_{t}(\pi_{t+h})) =& g\left(\sigma^2_{\Delta F^{\mathbb{Q}}} - \frac{(1-g)^2}{1-\theta g}Cov(z,\Delta E) - \sigma^2_{\Delta RP}\right) \\
&+ \theta Cov(\Delta F^{\mathbb{P}}_{t}(\pi_{t+h}), \Delta RP_t) - (1-g)Cov(z,\Delta E) \qquad (C.62) \\
&- \frac{1}{2}(1-\theta)\sigma^2_{\Delta RP} \\
=& g\left[\sigma^2_{\Delta F^{\mathbb{Q}}} - \sigma^2_{\Delta RP} - \left(\frac{(1-g)^2}{1-\theta g} - 1\right)Cov(z,\Delta E)\right] \\
&- \frac{1}{2}(1-\theta)\sigma^2_{\Delta RP} - \left(1 - \theta\frac{(1-g)^2}{1-\theta g}\right)Cov(z,\Delta E) \qquad (C.63)
\end{aligned}
$$

The beta coefficient of subsequent risk neutral forecast revisions is then given by

$$
\begin{aligned}
\beta(\Delta F^{\mathbb{Q}}_{t+1}(\pi_{t+h}), \Delta F^{\mathbb{Q}}_{t}(\pi_{t+h})) =& \frac{Cov(\Delta F^{\mathbb{Q}}_{t+1}(\pi_{t+h}), \Delta F^{\mathbb{Q}}_{t}(\pi_{t+h}))}{\sigma^2_{\Delta F^{\mathbb{Q}}}} \\
=& g\left[1 - \frac{\sigma^2_{\Delta RP}}{\sigma^2_{\Delta F^{\mathbb{Q}}}} - \left(\frac{(1-g)^2}{1-\theta g} - 1\right)\frac{Cov(z,\Delta E)}{\sigma^2_{\Delta F^{\mathbb{Q}}}}\right] \qquad (C.64) \\
&- \frac{1}{2}(1-\theta)\frac{\sigma^2_{\Delta RP}}{\sigma^2_{\Delta F^{\mathbb{Q}}}} - \left(1 - \theta\frac{(1-g)^2}{1-\theta g}\right)\frac{Cov(z,\Delta E)}{\sigma^2_{\Delta F^{\mathbb{Q}}}}
\end{aligned}
$$

And under FIRE ($g = 0$), expressions simplify to

$$
\sigma^2_{\Delta F^{\mathbb{Q}}} = \sigma^2_{\Delta E} + \sigma^2_{\Delta RP} + 2Cov(z,\Delta E) \qquad (C.65)
$$

$$
\begin{aligned}
\beta(\Delta F^{\mathbb{Q}}_{t+1}(\pi_{t+h}), \Delta F^{\mathbb{Q}}_{t}(\pi_{t+h})|g=0) &= -\frac{1}{2}(1-\theta)\frac{\sigma^2_{\Delta RP}}{\sigma^2_{\Delta F^{\mathbb{Q}}}} - (1-\theta)\frac{Cov(z,\Delta E)}{\sigma^2_{\Delta F^{\mathbb{Q}}}} \qquad (C.66) \\
&= -\frac{1}{2}(1-\theta)\left(1 - \frac{\sigma^2_{\Delta E}}{\sigma^2_{\Delta F^{\mathbb{Q}}}}\right) \qquad (C.67)
\end{aligned}
$$

Under the simplifying assumption $\rho(z,\Delta E) = 0$, Equation (C.64) reduces to

$$
\beta(\Delta F^{\mathbb{Q}}_{t+1}(\pi_{t+h}), \Delta F^{\mathbb{Q}}_{t}(\pi_{t+h})|Cov(z,\Delta E)=0) = g\left(1 - \frac{\sigma^2_{\Delta RP}}{\sigma^2_{\Delta F^{\mathbb{Q}}}}\right) - \frac{1}{2}(1-\theta)\left(\frac{\sigma^2_{\Delta RP}}{\sigma^2_{\Delta F^{\mathbb{Q}}}}\right) \qquad (C.68)
$$

C.1.2 Additional results

Table C.1: Horizon dependence: extended sample

(a) Error vs. revision

	Survey Consensus		Survey Individual		Market	
	$Err_{t+h,t}$		$Err_{t+h,t}$		$Err_{t+h,t}$	
α	-0.0002	0.0008	-0.0005*	0.0006*	0.0030***	0.0100***
$\Delta F_t(\pi_{t+h})$	0.599***	0.738***	-0.142***	-0.065**	-0.455***	-0.416***
$h * \Delta F_t(\pi_{t+h})$		-0.361*		-0.137***		-0.024

(b) Error vs. past error

	Survey Consensus		Survey Individual		Market	
	$Err_{t+h,t}$		$Err_{t+h,t}$		$Err_{t+h,t}$	
α	0.0000	0.0005	0.0004*	0.0008**	0.0047***	0.0091***
$Err_{t,t-h}$	0.085*	0.108	0.108***	0.128***	-0.154**	0.076
$h * Err_{t,t-h}$		-0.029		-0.027		-0.082***

(c) Revision vs. past revision

	Survey Consensus		Survey Individual		Market	
	$\Delta F_{t+1}(\pi_{t+h})$		$\Delta F_{t+1}(\pi_{t+h})$		$\Delta F_{t+1}(\pi_{t+h})$	
α	-0.0005**	-0.0005*	-0.0006***	-0.0006***	0.0000	0.0000
$\Delta F_t(\pi_{t+h})$	0.331***	0.63***	-0.187***	-0.162***	-0.203***	-0.148***
$h * \Delta F_t(\pi_{t+h})$		-0.536***		-0.037***		-0.025***

This table shows regression results for expectation data pooled over all horizons. We replicate our main regression results using the longest historical data sample available for each horizon. The earliest observations start in 1981. Error-covariances for computation of the t-statistics are estimated accounting for heteroscedasticity and serial correlation using the approach of Newey and West (1987). *, ** and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Figure C.1: Regression results: extended sample

(a) Error vs. revision

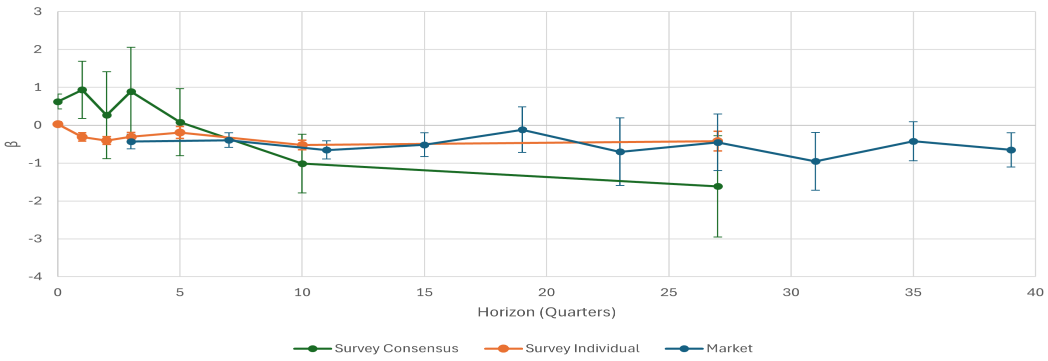

(b) Error vs. past error

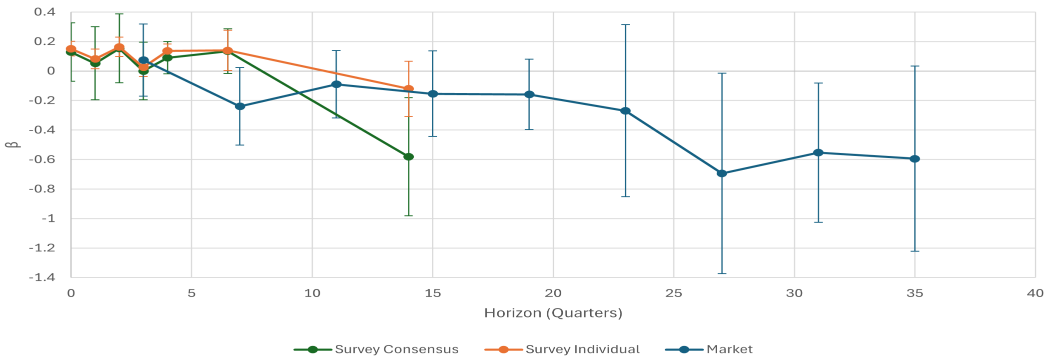

(c) Revision vs. past revision

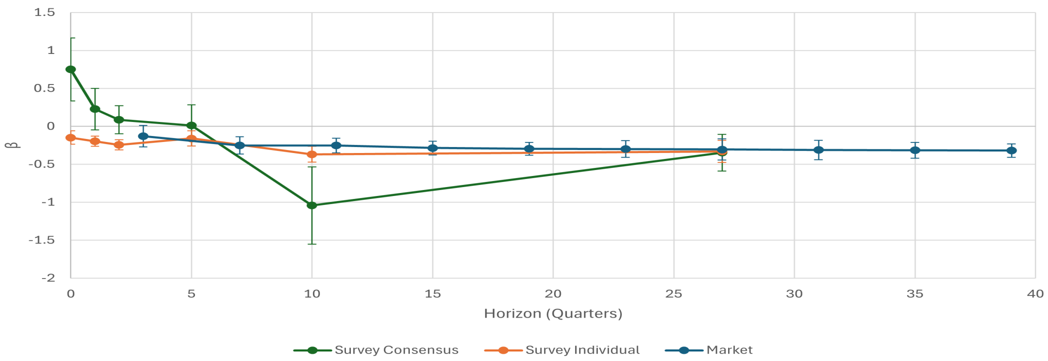

We replicate our main regression results using the longest historical data sample available for each horizon. The earliest observations start in 1981. We show empirical regression coefficients of (a) forecast errors on forecast revisions, (b) forecast errors on previous period forecast errors and (c) forecast revisions on preceding forecast revisions for regressions run for different forecast horizons seperately. Spikes report 90% confidence intervals calculated with HAC robust standard errors.

Figure C.2: Regression results: financial participants

(a) Error vs. revision

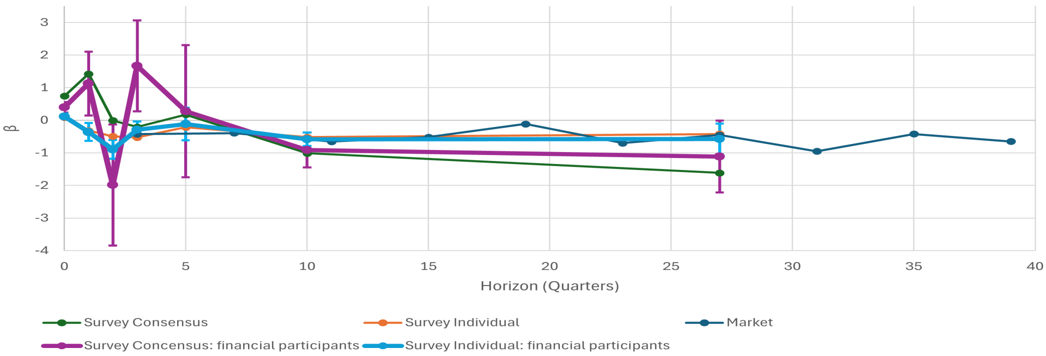

(b) Error vs. past error

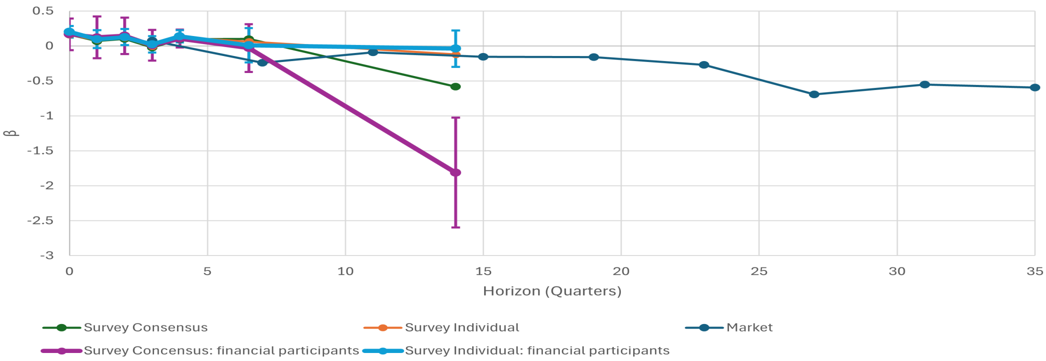

(c) Revision vs. past revision

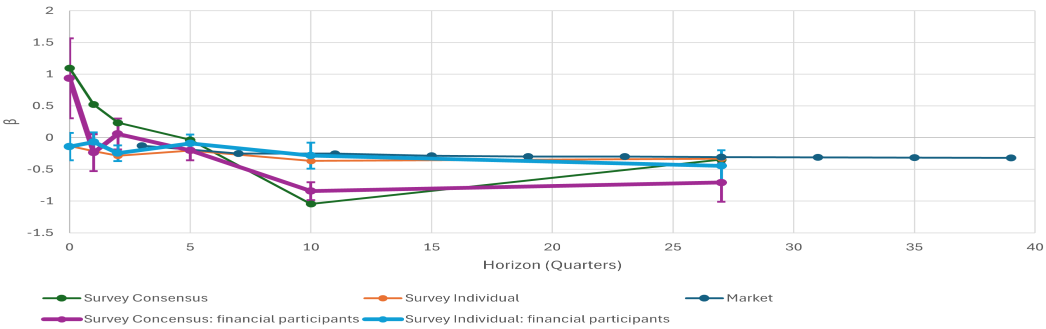

This Figure superimposes our main regression results with the results for the survey forecasts of only financial participants. Survey forecasts of financial participants are the survey responses from the SPF subsample of forecasters with industry classification "financial service providers". The consensus forecast of financial participants is calculated as the mean forecast of this subsample. We show empirical regression coefficients of (a) forecast errors on forecast revisions, (b) forecast errors on previous period forecast errors and (c) forecast revisions on preceding forecast revisions for regressions run for different forecast horizons seperately. Spikes report 90% confidence intervals calculated with HAC robust standard errors.

Table C.2: Horizon dependence: financial participants

(a) Error vs. revision

	Survey Consensus		Survey Individual	
	$Err_{t+h,t}$		$Err_{t+h,t}$	
α	0.0035***	0.0053***	0.0029***	0.0049***
$\Delta F_t(\pi_{t+h})$	0.389***	0.442***	-0.084*	-0.013
$h * \Delta F_t(\pi_{t+h})$		-0.210		-0.157***

(b) Error vs. past error

	Survey Consensus		Survey Individual	
	$Err_{t+h,t}$		$Err_{t+h,t}$	
α	0.0044***	0.0047***	0.0041***	0.0044***
$Err_{t,t-h}$	0.101*	0.154	0.106***	0.130**
$h * Err_{t,t-h}$		-0.069		-0.033

(c) Revision vs. past revision

	Survey Consensus		Survey Individual	
	$\Delta F_{t+1}(\pi_{t+h})$		$\Delta F_{t+1}(\pi_{t+h})$	
α	0.0003	0.0004	0.0000	0.0001
$\Delta F_t(\pi_{t+h})$	0.313	0.555**	-0.151***	-0.114*
$h * \Delta F_t(\pi_{t+h})$		-0.340***		-0.047**

This table shows regression results for expectation data pooled over all horizons. We replicate our main regression results with the results for the survey forecasts of only financial participants. Survey forecasts of financial participants are the survey responses from the SPF subsample of forecasters with industry classification "financial service providers". The consensus forecast of financial participants is calculated as the mean forecast of this subsample. Error-covariances for computation of the t-statistics are estimated accounting for heteroscedasticity and serial correlation using the approach of Newey and West (1987). *, ** and *** denote significance at the 10%, 5%, and 1% levels, respectively.

Figure C.3: Regression results: forecaster-by-forecaster regressions

(a) Error vs. revision

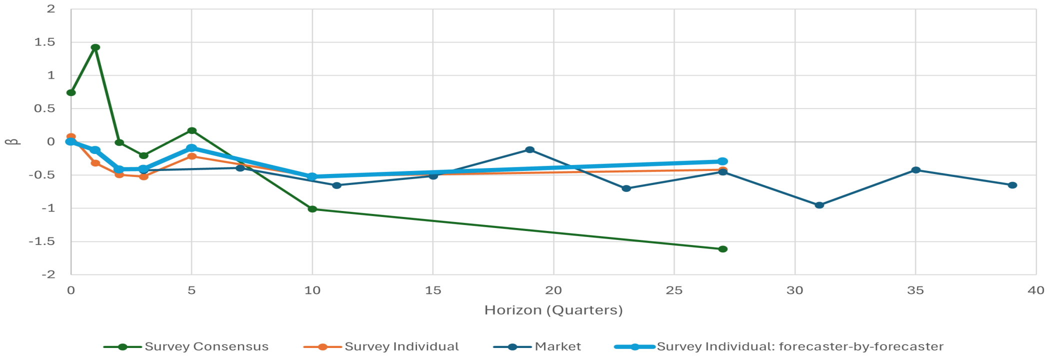

(b) Error vs. past error

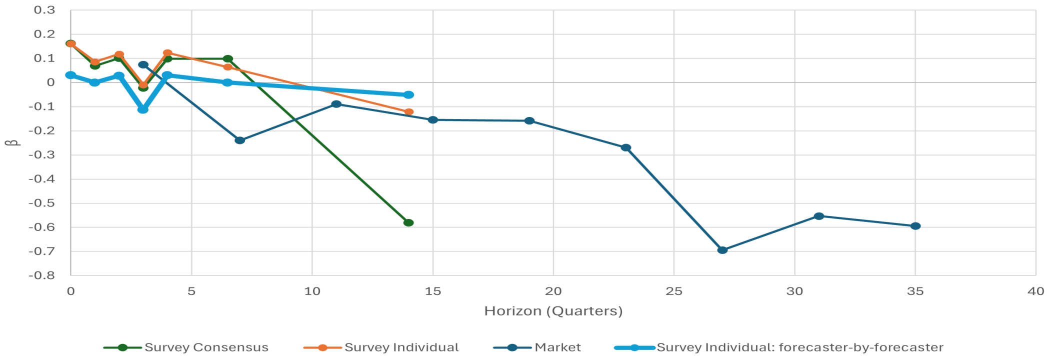

(c) Revision vs. past revision

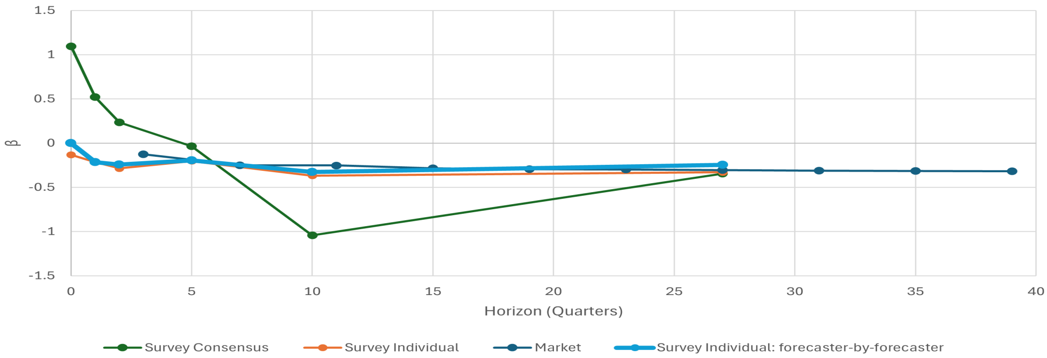

This Figure superimposes our main regression results with the results for forecaster-by-forecaster regressions. In (a) we estimate individual coefficients of forecast errors on forecast revisions with the specification $Err^i_{t+h,t} = \alpha^i + \beta^i * \Delta F^i_t(\pi_{t+h}) + u^i_{t+h}$. In (b) we estimate individual coefficients of forecast errors on previous period forecast errors with the specification $Err^i_{t+h,t} = \alpha^i + \beta^i * Err^i_{t,t-h} + u^i_{t+h}$. In (c) we estimate individual coefficients of forecast revisions on preceding forecast revisions with the specification $\Delta F^i_{t+1}(\pi_{t+h}) = \alpha^i + \beta^i * \Delta F^i_t(\pi_{t+h}) + u^i_{t+1}$. In each case, we show the median of the individual regression estimates across individual forecasters. All regressions are run for different forecast horizons seperately.

Table C.3: Horizon dependence: forecaster-by-forecaster regressions

(a) Error vs. revision

	$Err^i_{t+h,t}$	$Err^i_{t+h,t}$
α^i	0.0013	0.0024
$\Delta F^i_t(\pi_{t+h})$	-0.156	-0.028
$h * \Delta F^i_t(\pi_{t+h})$		-0.173

(b) Error vs. past error

	$Err^i_{t+h,t}$	$Err^i_{t+h,t}$
α^i	0.0029	0.0013
$Err^i_{t,t-h}$	-0.007	0.000
$h * Err^i_{t,t-h}$		-0.031

(c) Revision vs. past revision

	$\Delta F^i_{t+1}(\pi_{t+h})$	$\Delta F^i_{t+1}(\pi_{t+h})$
α^i	0.0001	0.0003
$\Delta F^i_t(\pi_{t+h})$	-0.183	-0.056
$h * \Delta F^i_t(\pi_{t+h})$		-0.084

This table shows regression results for expectation data pooled over all horizons. We replicate our main regression results with the results for individual forecaster-by-forecaster regressions. In (a) we estimate individual coefficients of forecast errors on forecast revisions with the specification $Err^i_{t+h,t} = \alpha^i + \beta^i * \Delta F^i_t(\pi_{t+h}) + \gamma^i * h * \Delta F^i_t(\pi_{t+h}) + \phi^i * h + u^i_{t+h}$. In (b) we estimate individual coefficients of forecast errors on previous period forecast errors with the specification $Err^i_{t+h,t} = \alpha^i + \beta^i * Err^i_{t,t-h} + \gamma^i * h * Err^i_{t,t-h} + \phi^i * h + u^i_{t+h}$. In (c) we estimate individual coefficients of forecast revisions on preceding forecast revisions with the specification $\Delta F^i_{t+1}(\pi_{t+h}) = \alpha^i + \beta^i * \Delta F^i_t(\pi_{t+h}) + \gamma^i * h * \Delta F^i_t(\pi_{t+h}) + \phi^i * h + u^i_{t+1}$. In each case, we show the median of the individual regression estimates across individual forecasters.

Bibliography

Abrahams, Michael, Tobias Adrian, Richard K Crump, Emanuel Moench, and Rui Yu, 2016, Decomposing real and nominal yield curves, *Journal of Monetary Economics* 84, 182–200.

Adam, Klaus, and Stefan Nagel, 2023, Expectations data in asset pricing, in *Handbook of economic expectations*, 477–506 (Elsevier).

Adrian, Tobias, and Hao Z Wu, 2009, The term structure of inflation expectations, *FRB of New York Staff Report* .

Ahn, Dong-Hyun, Jacob Boudoukh, Matthew Richardson, and Robert F Whitelaw, 2002a, Partial adjustment or stale prices? implications from stock index and futures return autocorrelations, *The Review of Financial Studies* 15, 655–689.

Ahn, Dong-Hyun, Robert F Dittmar, and A Ronald Gallant, 2002b, Quadratic term structure models: Theory and evidence, *The Review of Financial Studies* 15, 243–288.

Andersen, Torben G, Tim Bollerslev, Francis X Diebold, and Paul Labys, 2001, The distribution of realized exchange rate volatility, *Journal of the American statistical association* 96, 42–55.

Andersen, Torben G, Tim Bollerslev, Francis X Diebold, and Paul Labys, 2003, Modeling and forecasting realized volatility, *Econometrica* 71, 579–625.

Anderson, T. W., 1984, *An Introduction to Multivariate Statistical Analysis*, second edition (John Wiley and Sons).

Andreasen, Martin M, Jens HE Christensen, and Simon Riddell, 2021, The tips liquidity premium, *Review of Finance* 25, 1639–1675.

Andreasen, Martin M, and Andrew Meldrum, 2015, Market beliefs about the uk monetary policy life-off horizon: a no-arbitrage shadow rate term structure model approach, Bank of England working papers 541, Bank of England.

Andreasen, Martin M, and Andrew Meldrum, 2019, A shadow rate or a quadratic policy rule? the best way to enforce the zero lower bound in the united states, *Journal of Financial and Quantitative Analysis* 54, 2261–2292.

Andrews, Donald WK, 1993, Tests for parameter instability and structural change with unknown change point, *Econometrica: Journal of the Econometric Society* 62, 821–856.

Andrews, Donald WK, and Werner Ploberger, 1994, Optimal tests when a nuisance parameter is present only under the alternative, *Econometrica: Journal of the Econometric Society* 62, 1383–1414.

Ang, Andrew, Geert Bekaert, and Min Wei, 2007, Do macro variables, asset markets, or surveys forecast inflation better?, *Journal of monetary Economics* 54, 1163–1212.

Ang, Andrew, Geert Bekaert, and Min Wei, 2008, The term structure of real rates and expected inflation, *The Journal of Finance* 63, 797–849.

Angeletos, George-Marios, Zhen Huo, and Karthik A Sastry, 2021, Imperfect macroeconomic expectations: Evidence and theory, *NBER Macroeconomics Annual* 35, 1–86.

Auckenthaler, Julia, Alexander Kupfer, and Rupert Sendlhofer, 2015, The impact of liquidity on inflation-linked bonds: A hypothetical indexed bonds approach, *The North American Journal of Economics and Finance* 32, 139–154.

Bae, Kwangil, and Soonhee Lee, 2021, Realized higher-order comoments, *Quantitative Finance* 21, 421–429.

Bahaj, Saleem, Robert Czech, Sitong Ding, and Ricardo Reis, 2023, The market for inflation risk, Working paper, Bank of England.

Bai, Jushan, 2010, Common breaks in means and variances for panel data, *Journal of Econometrics* 157, 78–92.

Bai, Jushan, and Pierre Perron, 1998, Estimating and testing linear models with multiple structural changes, *Econometrica* 66, 47–78.

Baker, Scott R, Nicholas Bloom, and Steven J Davis, 2016, Measuring economic policy uncertainty, *The Quarterly Journal of Economics* 131, 1593–1636.

Bakshi, Gurdip, Nikunj Kapadia, and Dilip Madan, 2003, Stock return characteristics, skew laws, and the differential pricing of individual equity options, *The Review of Financial Studies* 16, 101–143.

Bakshi, Gurdip, and Dilip Madan, 2000, Spanning and derivative-security valuation, *Journal of Financial Economics* 55, 205–238.

Baltagi, Badi H, Qu Feng, and Chihwa Kao, 2016, Estimation of heterogeneous panels with structural breaks, *Journal of Econometrics* 191, 176–195.

Bansal, Naresh, Robert A Connolly, and Chris Stivers, 2014, The stock-bond return relation, the term structure's slope, and asset-class risk dynamics, *Journal of Financial and Quantitative Analysis* 49, 699–724.

Bansal, Ravi, Dana Kiku, and Amir Yaron, 2012, An empirical evaluation of the long-run risks model for asset prices, *Critical Finance Review* 1, 183–221.

Bansal, Ravi, and Ivan Shaliastovich, 2013, A long-run risks explanation of predictability puzzles in bond and currency markets, *The Review of Financial Studies* 26, 1–33.

Bansal, Ravi, and Amir Yaron, 2004, Risks for the long run: A potential resolution of asset pricing puzzles, *The Journal of Finance* 59, 1481–1509.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer, 2018, Extrapolation and bubbles, *Journal of Financial Economics* 129, 203–227.

Barndorff-Nielsen, Ole E, and Neil Shephard, 2004, Econometric analysis of realized covariation: High frequency based covariance, regression, and correlation in financial economics, *Econometrica* 72, 885–925.

Barro, Robert J, 2006, Rare disasters and asset markets in the twentieth century, *The Quarterly Journal of Economics* 121, 823–866.

Bauer, Michael, 2015, Inflation expectations and the news, *International Journal of Central Banking* 11, 1–40.

Bauer, Michael D., and Mikhail Chernov, 2024, Interest Rate Skewness and Biased Beliefs, *The Journal of Finance* 79, 173–217.

Bauer, Michael D, Aeimit Lakdawala, and Philippe Mueller, 2022, Market-based monetary policy uncertainty, *The Economic Journal* 132, 1290–1308.

Bauer, Michael D, and Glenn D Rudebusch, 2016, Monetary policy expectations at the zero lower bound, *Journal of Money, Credit and Banking* 48, 1439–1465.

Ben-Rephael, Azi, Zhi Da, and Ryan D Israelsen, 2017, It depends on where you search: Institutional investor attention and underreaction to news, *The Review of Financial Studies* 30, 3009–3047.

Black, Fischer, 1995, Interest rates as options, *The Journal of Finance* 50, 1371–1376.

Boeck, Maximilian, and Martin Feldkircher, 2021, The impact of monetary policy on yield curve expectations, *Journal of Economic Behavior & Organization* 191, 887–901.

Bonelli, Diego, Berardino Palazzo, and Ram Yamarthy, 2024, Good inflation, bad inflation: Implications for risky asset prices, Working paper.

Boons, Martijn, Fernando Duarte, Frans De Roon, and Marta Szymanowska, 2020, Time-varying inflation risk and stock returns, *Journal of Financial Economics* 136, 444–470.

Bordalo, Pedro, Nicola Gennaioli, Yueran Ma, and Andrei Shleifer, 2020, Overreaction in macroeconomic expectations, *American Economic Review* 110, 2748–2782.

Bordalo, Pedro, Nicola Gennaioli, Rafael La Porta, and Andrei Shleifer, 2019, Diagnostic expectations and stock returns, *The Journal of Finance* 74, 2839–2874.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer, 2018, Diagnostic expectations and credit cycles, *The Journal of Finance* 73, 199–227.

Born, Benjamin, Zeno Enders, and Gernot J Müller, 2023, On fire, news, and expectations, Working paper.

Bouchaud, Jean-Philippe, Philipp Krueger, Augustin Landier, and David Thesmar, 2019, Sticky expectations and the profitability anomaly, *The Journal of Finance* 74, 639–674.

Broer, Tobias, and Alexandre N Kohlhas, 2024, Forecaster (mis-) behavior, *Review of Economics and Statistics* 106, 1334–1351.

Brooks, Jordan, Michael Katz, and Hanno Lustig, 2018, Post-fomc announcement drift in us bond markets, Working paper, National Bureau of Economic Research.

Brunnermeier, Markus, Emmanuel Farhi, Ralph SJ Koijen, Arvind Krishnamurthy, Sydney C Ludvigson, Hanno Lustig, Stefan Nagel, and Monika Piazzesi, 2021, Perspectives on the future of asset pricing, *The Review of Financial Studies* 34, 2126–2160.

Buraschi, Andrea, and Alexei Jiltsov, 2005, Inflation risk premia and the expectations hypothesis, *Journal of Financial Economics* 75, 429–490.

Burkhardt, Dominic, 2014, Time-varying inflation risk and the cross-section of stock returns, Working paper.

Burkhardt, Dominic, and Henrik Hasseltoft, 2012, Understanding asset correlations, Working paper.

Bushman, Robert M, Christopher D Williams, and Regina Wittenberg-Moerman, 2017, The informational role of the media in private lending, *Journal of Accounting Research* 55, 115–152.

Caballero, Ricardo J, and Alp Simsek, 2021, A model of endogenous risk intolerance and lsaps: Asset prices and aggregate demand in a "covid-19" shock, *The Review of Financial Studies* 34, 5522–5580.

Cakici, Nusret, and Jintao Zhu, 2001, Pricing eurodollar futures options with the heath—jarrow—morton model, *Journal of Futures Markets: Futures, Options, and Other Derivative Products* 21, 655–680.

Campbell, Dennis W, and Ruidi Shang, 2022, Tone at the bottom: Measuring corporate misconduct risk from the text of employee reviews, *Management Science* 68, 7034–7053.

Campbell, John Y, 2017, *Financial decisions and markets: a course in asset pricing* (Princeton University Press).

Campbell, John Y, Carolin Pflueger, and Luis M Viceira, 2020, Macroeconomic drivers of bond and equity risks, *Journal of Political Economy* 128, 3148–3185.

Campbell, John Y, and Robert J Shiller, 1988, The dividend-price ratio and expectations of future dividends and discount factors, *The Review of Financial Studies* 1, 195–228.

Campbell, John Y, Robert J Shiller, and Luis M Viceira, 2009, Understanding inflation-indexed bond markets, *Brookings Papers on Economic Activity* 2009, 79–120.

Campbell, John Y, Adi Sunderam, and Luis M Viceira, 2017, Inflation bets or deflation hedges? the changing risks of nominal bonds, *Critical Finance Review* 6, 263–301.

Carr, P., and D. Madan, 2001, Optimal positioning in derivative securities, *Quantitative Finance* 1, 19–37.

Carr, Peter, and Liuren Wu, 2009, Variance risk premiums, *The Review of Financial Studies* 22, 1311–1341.

Carriero, Andrea, Sarah Mouabbi, and Elisabetta Vangelista, 2018, Uk term structure decompositions at the zero lower bound, *Journal of Applied Econometrics* 33, 643–661.

Carroll, Christopher D, 2003, Macroeconomic expectations of households and professional forecasters, *the Quarterly Journal of Economics* 118, 269–298.

Chen, Hui, Winston Wei Dou, and Leonid Kogan, 2024, Measuring “dark matter” in asset pricing models, *The Journal of Finance* 79, 843–902.

Chen, Ren-Raw, and Louis Scott, 1993, Pricing interest rate futures options with futures-style margining, *The Journal of Futures Markets* 13, 15.

Chernov, Mikhail, Lars A Lochstoer, and Dongho Song, 2021, The real explanation of nominal bond-stock puzzles, Working paper.

Chernov, Mikhail, and Philippe Mueller, 2012, The term structure of inflation expectations, *Journal of Financial Economics* 106, 367–394.

Choi, Hoyong, Philippe Mueller, and Andrea Vedolin, 2017, Bond variance risk premiums, *Review of Finance* 21, 987–1022.

Christensen, Jens HE, Jose A Lopez, and Glenn D Rudebusch, 2010, Inflation expectations and risk premiums in an arbitrage-free model of nominal and real bond yields, *Journal of Money, Credit and Banking* 42, 143–178.

Christensen, Jens HE, and Glenn D Rudebusch, 2015, Estimating shadow-rate term structure models with near-zero yields, *Journal of Financial Econometrics* 13, 226–259.

Christensen, Jens HE, and Glenn D Rudebusch, 2016, Modeling yields at the zero lower bound: Are shadow rates the solution?, *Advances in Econometrics* 35, 75–125.

Christiano, Lawrence, Martin Eichenbaum, and Sergio Rebelo, 2011, When is the government spending multiplier large?, *Journal of Political Economy* 119, 78–121.

Chung, Hess, Jean-Philippe Laforte, David Reifschneider, and John C Williams, 2012, Have we underestimated the probability of hitting the zero lower bound?, *Journal of Money, Credit and Banking* 44, 47–48.

Cieslak, Anna, 2018, Short-rate expectations and unexpected returns in treasury bonds, *The Review of Financial Studies* 31, 3265–3306.

Coibion, Olivier, and Yuriy Gorodnichenko, 2012, What can survey forecasts tell us about information rigidities?, *Journal of Political Economy* 120, 116–159.

Coibion, Olivier, and Yuriy Gorodnichenko, 2015, Information rigidity and the expectations formation process: A simple framework and new facts, *American Economic Review* 105, 2644–2678.

Collin-Dufresne, Pierre, and Robert S Goldstein, 2002, Do bonds span the fixed income markets? theory and evidence for unspanned stochastic volatility, *The Journal of Finance* 57, 1685–1730.

Connolly, Robert, Chris Stivers, and Licheng Sun, 2005, Stock market uncertainty and the stock-bond return relation, *Journal of Financial and Quantitative Analysis* 40, 161–194.

Connolly, Robert A, Chris Stivers, and Licheng Sun, 2007, Commonality in the time-variation of stock–stock and stock–bond return comovements, *Journal of Financial Markets* 10, 192–218.

Cox, John C, Jonathan Ingersoll, and Stephen Ross, 1985, A theory of the term structure of interest rates, *Econometrica* 53, 385–407.

Creal, Drew D, and Jing Cynthia Wu, 2017, Monetary policy uncertainty and economic fluctuations, *International Economic Review* 58, 1317–1354.

Dai, Qiang, and Kenneth J. Singleton, 2000, Specification analysis of affine term structure models, *The Journal of Finance* 55, 1943–1978.

d'Arienzo, Daniele, 2020, Maturity increasing overreaction and bond market puzzles, Working paper.

Datta, Deepa D, Benjamin K Johannsen, Hannah Kwon, and Robert J Vigfusson, 2021, Oil, equities, and the zero lower bound, *American Economic Journal: Macroeconomics* 13, 214–53.

David, Alexander, and Pietro Veronesi, 2008, Inflation and earnings uncertainty and volatility forecasts: A structural form approach, Chicago gsb research paper, university of calgary haskyane school of business working paper.

David, Alexander, and Pietro Veronesi, 2013, What ties return volatilities to price valuations and fundamentals?, *Journal of Political Economy* 121, 682–746.

Dergunov, Ilya, Christoph Meinerding, and Christian Schlag, 2023, Extreme inflation and time-varying expected consumption growth, *Management Science* 69, 2972–3002.

Doms, Mark E, and Norman J Morin, 2004, Consumer sentiment, the economy, and the news media, FRB of San Francisco Working Paper 2004-09.

Drechsler, Itamar, and Amir Yaron, 2010, What's vol got to do with it, *The Review of Financial Studies* 24, 1–45.

Duffee, Gregory R, 2023, Macroeconomic news and stock–bond comovement, *Review of Finance* 27, 1859–1882.

d'Amico, Stefania, Don H Kim, and Min Wei, 2018, Tips from tips: the informational content of treasury inflation-protected security prices, *Journal of Financial and Quantitative Analysis* 53, 395–436.

Eggertsson, Gauti, and Michael Woodford, 2003a, The zero bound on interest rates and optimal monetary policy, *Brookings Papers on Economic Activity* 34, 139–235.

Eggertsson, Gauti B., and Michael Woodford, 2003b, Optimal Monetary Policy in a Liquidity Trap, NBER Working Papers 9968, National Bureau of Economic Research, Inc.

Epstein, Larry G., and Stanley E. Zin, 1989, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937–969.

Ermolov, Andrey, 2022, Time-varying risk of nominal bonds: How important are macroeconomic shocks?, *Journal of Financial Economics* 145, 1–28.

Feunou, Bruno, Jean-Sébastien Fontaine, Anh Le, and Christian Lundblad, 2022, Tractable term structure models, *Management Science* 68, 8411–8429.

Filipović, Damir, Martin Larsson, and Anders B Trolle, 2017, Linear-rational term structure models, *The Journal of Finance* 72, 655–704.

Fleckenstein, Matthias, Francis A Longstaff, and Hanno Lustig, 2014, The tips-treasury bond puzzle, *The Journal of Finance* 69, 2151–2197.

Fleckenstein, Matthias, Francis A Longstaff, and Hanno Lustig, 2016, Inflation-adjusted bonds and the inflation risk premium, *Handbook of Fixed-Income Securities* 41–52.

Fleckenstein, Matthias, Francis A Longstaff, and Hanno Lustig, 2017, Deflation risk, *The Review of Financial Studies* 30, 2719–2760.

Flesaker, Bjorn, 1993a, Arbitrage free pricing of interest rate futures and forward contracts, *The Journal of Futures Markets* 13, 77.

Flesaker, Bjorn, 1993b, Testing the heath-jarrow-morton/ho-lee model of interest rate contingent claims pricing, *Journal of Financial and Quantitative Analysis* 28, 483–495.

Gabaix, Xavier, 2012, Variable rare disasters: An exactly solved framework for ten puzzles in macro-finance, *The Quarterly Journal of Rconomics* 127, 645–700.

Gabaix, Xavier, 2014, A sparsity-based model of bounded rationality, *The Quarterly Journal of Economics* 129, 1661–1710.

Geiger, Felix, and Fabian Schupp, 2018, With a little help from my friends: Survey-based derivation of euro area short rate expectations at the effective lower bound, Discussion Papers 27/2018, Deutsche Bundesbank.

Giglio, Stefano, and Bryan Kelly, 2018, Excess volatility: Beyond discount rates, *The Quarterly Journal of Economics* 133, 71–127.

Goidel, Robert K, and Ronald E Langley, 1995, Media coverage of the economy and aggregate economic evaluations: Uncovering evidence of indirect media effects, *Political Research Quarterly* 48, 313–328.

Gourio, François, and Phuong Ngo, 2020, Risk Premia at the ZLB: A Macroeconomic Interpretation, Working Paper Series WP 2020-01, Federal Reserve Bank of Chicago.

Greenwood, Robin, and Andrei Shleifer, 2014, Expectations of returns and expected returns, *The Review of Financial Studies* 27, 714–746.

Grishchenko, Olesya V, and Jing-Zhi Huang, 2013, The inflation risk premium: Evidence from the tips market, *The Journal of Fixed Income* 22, 5–30.

Grishchenko, Olesya V, Joel M Vanden, and Jianing Zhang, 2016, The informational content of the embedded deflation option in tips, *Journal of Banking & Finance* 65, 1–26.

Gürkaynak, Refet S, Brian Sack, and Eric Swanson, 2005, The sensitivity of long-term interest rates to economic news: Evidence and implications for macroeconomic models, *American Economic Review* 95, 425–436.

Gürkaynak, Refet S, Brian Sack, and Jonathan H Wright, 2007, The us treasury yield curve: 1961 to the present, *Journal of Monetary Economics* 54, 2291–2304.

Gürkaynak, Refet S, Brian Sack, and Jonathan H Wright, 2010, The tips yield curve and inflation compensation, *American Economic Journal: Macroeconomics* 2, 70–92.

Gust, Christopher, Edward Herbst, David López-Salido, and Matthew E Smith, 2017, The empirical implications of the interest-rate lower bound, *American Economic Review* 107, 1971–2006.

Hanson, Samuel G, David O Lucca, and Jonathan H Wright, 2018, The excess sensitivity of long-term rates: A tale of two frequencies, *Federal Reserve Bank of New York Staff Report* 810.

Hasseltoft, Henrik, 2009, The 'fed model' and the changing correlation of stock and bond returns: An equilibrium approach, Working paper.

Heidari, Massoud, and Liuren Wu, 2003, Are interest rate derivatives spanned by the term structure of interest rates?, *The Journal of Fixed Income* 13, 75–86.

Hördahl, Peter, and Oreste Tristani, 2014, Inflation risk premia in the euro area and the united states, *International Journal of Central Banking* 10, 1–47.

Ichiue, Hibiki, Yoichi Ueno, et al., 2013, Estimating term premia at the zero bound: An analysis of japanese, us, and uk yields, Working paper, Bank of Japan.

Iiboshi, Hirokuni, Mototsugu Shintani, and Kozo Ueda, 2022, Estimating a nonlinear new keynesian model with the zero lower bound for japan, *Journal of Money, Credit and Banking* 54, 1637–1671.

Ilmanen, Antti, 2003, Stock-bond correlations, *The Journal of Fixed Income* 13, 55.

Istrefi, Klodiana, and Sarah Mouabbi, 2018, Subjective interest rate uncertainty and the macroeconomy: A cross-country analysis, *Journal of International Money and Finance* 88, 296–313.

Jiang, George J, and Yisong S Tian, 2005, The model-free implied volatility and its information content, *The Review of Financial Studies* 18, 1305–1342.

Jiang, Hao, Sophia Zhengzi Li, and Hao Wang, 2021, Pervasive underreaction: Evidence from high-frequency data, *Journal of Financial Economics* 141, 573–599.

Johannsen, Benjamin K, and Elmar Mertens, 2021, A time-series model of interest rates with the effective lower bound, *Journal of Money, Credit and Banking* 53, 1005–1046.

Jones, Christopher S, and Sungjune Pyun, 2024, Consumption growth persistence and the stock-bond correlation, *Journal of Financial and Quantitative Analysis* 1–29.

Joyce, Michael AS, Peter Lildholdt, and Steffen Sorensen, 2010, Extracting inflation expectations and inflation risk premia from the term structure: a joint model of the uk nominal and real yield curves, *Journal of Banking & Finance* 34, 281–294.

Kajuth, Florian, and Sebastian Watzka, 2011, Inflation expectations from index-linked bonds: Correcting for liquidity and inflation risk premia, *The Quarterly Review of Economics and Finance* 51, 225–235.

Kang, Johnny, and Carolin E Pflueger, 2015, Inflation risk in corporate bonds, *The Journal of Finance* 70, 115–162.

Karr, Alan F, 1983, Extreme points of certain sets of probability measures, with applications, *Mathematics of Operations Research* 8, 74–85.

Kim, Don, Cait Walsh, and Min Wei, 2019, Tips from tips: Update and discussions, Feds notes.

Kim, Don H, and Marcel Priebsch, 2013, Estimation of multi-factor shadow-rate term structure models, Federal reserve board discussion paper series, Federal Reserve Board.

Kim, Don H, and Kenneth J Singleton, 2012, Term structure models and the zero bound: an empirical investigation of japanese yields, *Journal of Econometrics* 170, 32–49.

King, Thomas B, 2019, Expectation and duration at the effective lower bound, *Journal of Financial Economics* 134, 736–760.

Kohlhas, Alexandre N, and Ansgar Walther, 2021, Asymmetric attention, *American Economic Review* 111, 2879–2925.

Kolasinski, Adam C, and Nan Yang, 2018, Managerial myopia and the mortgage meltdown, *Journal of Financial Economics* 128, 466–485.

Kortela, Tomi, 2016, A shadow rate model with time-varying lower bound of interest rates, Bank of Finland Research Discussion Paper 19.

Kozak, Serhiy, 2022, Dynamics of bond and stock returns, *Journal of Monetary Economics* 126, 188–209.

Krein, MG, and AA Nudelman, 1977, The markov moment problem and extremal problems, transl. math, *Monographs, American Math. Soc., Providence* 50.

Kreps, David M, and Evan L Porteus, 1978, Temporal resolution of uncertainty and dynamic choice theory, *Econometrica: Journal of the Econometric Society* 185–200.

Krippner, Leo, 2012, Modifying gaussian term structure models when interest rates are near the zero lower bound, Reserve Bank of New Zealand Discussion Paper 2012/02.

Krippner, Leo, 2013, Measuring the stance of monetary policy in zero lower bound environments, *Economics Letters* 118, 135–138.

Krippner, Leo, 2015, A comment on wu and xia (2015), and the case for two-factor shadow short rates, Cama working paper.

Krippner, Leo, 2017, A comment on wu and xia (2016) from a macroeconomic perspective, Cama working paper.

Krugman, Paul R, Kathryn M Dominquez, and Kenneth Rogoff, 1998, It's baaack: Japan's slump and the return of the liquidity trap, *Brookings papers on economic activity* 1998, 137–205.

Kupfer, Alexander, 2019, Estimating inflation risk premia using inflation-linked bonds: A review, *Contemporary Topics in Finance: A Collection of Literature Surveys* 117–149.

Laarits, Toomas, 2020, Precautionary savings and the stock-bond covariance, Working paper, NYU Stern School of Business.

Lamla, Michael J, and Sarah M Lein, 2014, The role of media for consumers' inflation expectation formation, *Journal of Economic Behavior & Organization* 106, 62–77.

Lamla, Michael J, and Sarah M Lein, 2015, Information rigidities, inflation perceptions, and the media: lessons from the euro cash changeover, *Economic Inquiry* 53, 9–22.

Lamla, Michael J, and Samad Sarferaz, 2012, Updating inflation expectations, Kof working papers.

Larsen, Vegard H, Leif Anders Thorsrud, and Julia Zhulanova, 2021, News-driven inflation expectations and information rigidities, *Journal of Monetary Economics* 117, 507–520.

Leippold, Markus, and Liuren Wu, 2002, Asset pricing under the quadratic class, *Journal of Financial and Quantitative Analysis* 37, 271–295.

Lemke, Wolfgang, and Andreea Liliana Vladu, 2017, Below the zero lower bound: a shadow-rate term structure model for the euro area, Working Paper Series 1991, European Central Bank.

Li, Haitao, and Feng Zhao, 2006, Unspanned stochastic volatility: Evidence from hedging interest rate derivatives, *The Journal of Finance* 61, 341–378.

Li, Haitao, and Feng Zhao, 2009, Nonparametric estimation of state-price densities implicit in interest rate cap prices, *The Review of Financial Studies* 22, 4335–4376.

Li, Lingfeng, et al., 2003, Macroeconomic factors and the correlation of stock and bond returns, Working paper, Yale School of Management.

Lieu, Derming, 1990, Option pricing with futures-style margining, *The Journal of Futures Markets* 10, 327.

Liu, Yukun, and Ben Matthies, 2022, Long-run risk: Is it there?, *The Journal of Finance* 77, 1587–1633.

Lochstoer, Lars A, and Tyler Muir, 2022, Volatility expectations and returns, *The Journal of Finance* 77, 1055–1096.

Lucas, Robert E, 1972, Expectations and the neutrality of money, *Journal of Economic Theory* 4, 103–124.

Lucas, Robert E., 1973, Some international evidence on output-inflation tradeoffs, *The American Economic Review* 63, 326–334.

Lucas, Robert E., 1976, Econometric policy evaluation: A critique, *Carnegie-Rochester Conference Series on Public Policy* 1, 19–46.

Manela, Asaf, and Alan Moreira, 2017, News implied volatility and disaster concerns, *Journal of Financial Economics* 123, 137–162.

Mankiw, N Gregory, and Ricardo Reis, 2002, Sticky information versus sticky prices: a proposal to replace the new keynesian phillips curve, *The Quarterly Journal of Economics* 117, 1295–1328.

McCallum, Bennett T, 2000, Theoretical analysis regarding a zero lower bound on nominal interest rates, *Journal of Money, Credit and Banking* 870–904.

Mertens, Thomas M., and John C. Williams, 2021, What to expect from the lower bound on interest rates: Evidence from derivatives prices, *American Economic Review* 111, 2473–2505.

Mokanov, Denis, 2023, Deviations from rational expectations and the uncovered interest rate parity puzzle, Working paper.

Monfort, Alain, Fulvio Pegoraro, Jean-Paul Renne, and Guillaume Roussellet, 2017, Staying at zero with affine processes: An application to term structure modelling, *Journal of Econometrics* 201, 348–366.

Muth, John F., 1960, Optimal properties of exponentially weighted forecasts, *Journal of the American Statistical Association* 55, 299–306.

Muth, John F., 1961, Rational expectations and the theory of price movements, *Econometrica* 29, 315–335.

Neuberger, Anthony, 2012, Realized skewness, *The Review of Financial Studies* 25, 3423–3455.

Newey, Whitney K., and Kenneth D. West, 1987, A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix, *Econometrica* 55, 703–708.

Ortu, Fulvio, Andrea Tamoni, and Claudio Tebaldi, 2013, Long-run risk and the persistence of consumption shocks, *The Review of Financial Studies* 26, 2876–2915.

Pang, Hao, 2023, The term structure of inflation expectations and treasury yields, Working paper.

Pfajfar, Damjan, and Emiliano Santoro, 2013, News on inflation and the epidemiology of inflation expectations, *Journal of Money, Credit and Banking* 45, 1045–1067.

Pflueger, Carolin E, and Luis M Viceira, 2011, An empirical decomposition of risk and liquidity in nominal and inflation-indexed government bonds, Working paper.

Pflueger, Carolin E, and Luis M Viceira, 2016, Return predictability in the treasury market: real rates, inflation, and liquidity, *Handbook of Fixed-Income Securities* 191–209.

Picault, Matthieu, Julien Pinter, and Thomas Renault, 2022, Media sentiment on monetary policy: Determinants and relevance for inflation expectations, *Journal of International Money and Finance* 124, 102626.

Popoviciu, Tiberiu, 1935, Sur les équations algébriques ayant toutes leurs racines réelles, *Mathematica* 9, 129–145.

Puigvert-Gutiérrez, Josep Maria, and Rupert de Vincent-Humphreys, 2012, A quantitative mirror on the euribor market using implied probability density functions, *Eurasian Economic Review* 2, 1–31.

Reed, Adam V, Mehrdad Samadi, and Jonathan S Sokobin, 2020, Shorting in broad daylight: Short sales and venue choice, *Journal of Financial and Quantitative Analysis* 55, 2246–2269.

Reifschneider, David, and John C Williams, 2000, Three lessons for monetary policy in a low-inflation era, *Journal of Money, Credit and Banking* 936–966.

Roppel, Anian, and Marliese Uhrig-Homburg, 2024a, Inflation cyclicality and stock-bond comovement: evidence from news media, Working paper.

Roppel, Anian, and Marliese Uhrig-Homburg, 2024b, Tips under fire: Formation of inflation expectations in bond markets, Working paper.

Roppel, Anian, Nils Unger, and Marliese Uhrig-Homburg, 2023, Option-implied lower-bound beliefs and the impact of negative interest rates, Working paper.

Schmeling, Maik, Andreas Schrimpf, and Sigurd AM Steffensen, 2022, Monetary policy expectation errors, *Journal of Financial Economics* 146, 841–858.

Schorfheide, Frank, Dongho Song, and Amir Yaron, 2018, Identifying long-run risks: A bayesian mixed-frequency approach, *Econometrica* 86, 617–654.

Seo, Jinyoung, 2023, The determinants of bond-stock correlation: the role of trend inflation and monetary policy, Working paper.

Shapiro, Adam Hale, Moritz Sudhof, and Daniel J Wilson, 2022, Measuring news sentiment, *Journal of Econometrics* 228, 221–243.

Shen, Pu, 2006, Liquidity risk premia and breakeven inflation rates., *Economic Review* 91.

Shiller, Robert J, 2017, Narrative economics, *American economic review* 107, 967–1004.

Sims, Christopher A, 2003, Implications of rational inattention, *Journal of Monetary Economics* 50, 665–690.

Söderlind, Paul, 2011, Inflation risk premia and survey evidence on macroeconomic uncertainty, *International Journal of Central Banking* 7, 113–133.

Song, Dongho, 2017, Bond market exposures to macroeconomic and monetary policy risks, *The Review of Financial Studies* 30, 2761–2817.

Soroka, Stuart N, 2006, Good news and bad news: Asymmetric responses to economic information, *The Journal of Politics* 68, 372–385.

Swanson, Eric T, and John C Williams, 2014, Measuring the effect of the zero lower bound on medium-and longer-term interest rates, *American Economic Review* 104, 3154–85.

Tetlock, Paul C, 2007, Giving content to investor sentiment: The role of media in the stock market, *The Journal of Finance* 62, 1139–1168.

Trolle, Anders B, and Eduardo S Schwartz, 2014, The swaption cube, *The Review of Financial Studies* 27, 2307–2353.

Valente, Joao Paulo, Kaushik Vasudevan, and Tianhao Wu, 2022, The role of beliefs in asset prices: Evidence from exchange rates, Working paper.

van Binsbergen, Jules H, Svetlana Bryzgalova, Mayukh Mukhopadhyay, and Varun Sharma, 2024, (almost) 200 years of news-based economic sentiment, Working paper, National Bureau of Economic Research.

von Beschwitz, Bastian, Donald B Keim, and Massimo Massa, 2020, First to “read” the news: News analytics and algorithmic trading, *The Review of Asset Pricing Studies* 10, 122–178.

Wachter, Jessica A, 2013, Can time-varying risk of rare disasters explain aggregate stock market volatility?, *The Journal of Finance* 68, 987–1035.

Wang, Chen, 2021, Under-and overreaction in yield curve expectations, Working paper.

Woodford, Michael, 2003, Imperfect common knowledge and the effects of monetary policy, *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps* 25.

Wu, Jing Cynthia, and Fan Dora Xia, 2016, Measuring the macroeconomic impact of monetary policy at the zero lower bound, *Journal of Money, Credit and Banking* 48, 253–291.

Wu, Jing Cynthia, and Fan Dora Xia, 2017, Time-varying lower bound of interest rates in europe, Chicago Booth Research Paper 17-06.