



Essays on Asset Prices and Macroeconomic News Announcements

Citation

Zhou, John Cong. 2016. Essays on Asset Prices and Macroeconomic News Announcements. Doctoral dissertation, Harvard University, Graduate School of Arts & Sciences.

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:33493485>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Essays on Asset Prices and Macroeconomic News Announcements

A dissertation presented

by

John Cong Zhou

to

The Committee for the PhD in Business Economics

in partial fulfillment of the requirements

for the degree of

Doctor of Philosophy

in the subject of

Business Economics

Harvard University

Cambridge, Massachusetts

April 2016

© 2016 John Cong Zhou
All rights reserved.

Dissertation Advisors:
Professor John Campbell
Professor Jeremy Stein

Author:
John Cong Zhou

Essays on Asset Prices and Macroeconomic News Announcements

Abstract

My dissertation is composed of three chapters that are unified by their exploration of asset prices and macroeconomic news announcements. With respect to asset prices, my main focus is on the price discovery process: how do asset prices reveal information relevant for asset fundamentals? Through my research, I provide new answers to this question. My work gets at core issues in asset pricing: whether financial markets are informationally efficient; why some assets earn unconditionally high premia; and how the sensitivity of prices to information varies over time and across assets. Specifically, chapter one shows evidence that sophisticated traders with an informational advantage inefficiently impound their edge into the aggregate U.S. stock market and U.S. Treasury bonds. In chapter two, I explore a model in which investors are averse to ambiguity (Knightian uncertainty) to explain why the equity premium is concentrated around specific events. Finally, chapter three investigates how the Federal Reserve's zero lower bound affects the response of asset prices, in particular interest rates, to information.

Each of the three chapters explores the price discovery process using the unique setting of U.S. macroeconomic news announcements, which are made by government agencies and private-sector organizations and cover macroeconomic data on inflation, output, and unemployment. Analyzing financial markets in this setting deepens our understanding of how asset prices reflect information about macroeconomic fundamentals. At the same time, the results have macroeconomic implications; for example, the assumptions of monetary policy models in theory and the effectiveness of unconventional monetary policy in practice.

Contents

Abstract	iii
Acknowledgments	xii
Introduction	1
1 Sophisticated Trading and Market Efficiency: Evidence from Macroeconomic News Announcements	4
1.1 Introduction	4
1.2 Empirical Setting	11
1.2.1 MNAs	12
1.2.2 Stocks and Bonds	14
1.3 Sophisticated Trading around MNAs	18
1.3.1 Pre-Announcement Conditional Drift	18
1.3.2 Order Imbalances and Sophisticated Trading Profits	25
1.4 Source of Informational Advantage	26
1.4.1 (Illegal) Information Leakage Channel	26
1.4.2 Homework Channel	28
1.4.2.1 Doing Homework: Public Information Predicts Macroeconomic News Reports	28
1.4.2.2 Trading on the Homework: Explaining the Pre-Announcement Conditional Drift	31
1.5 Market Efficiency around MNAs	34
1.5.1 Over Time	34
1.5.2 Across Assets	35
1.6 Model	37
1.6.1 Model with One Risky Asset	38
1.6.1.1 Efficient-Market Specification: All Traders are Sophisticated, and $i = 1$	40
1.6.1.2 Inefficient-Market Specification: Some Traders are Unsophisticated, and $i \in (0, 1)$	42
1.6.2 Model with Two Risky Assets	43

1.6.2.1	Efficient-Market Specification: All Traders are Sophisticated, and $i = 1$	44
1.6.2.2	Inefficient-Market Specification: Some Traders are Unsophisticated, and $i \in (0,1)$	45
1.6.3	Data are Consistent with the Model	47
1.7	Trading on Market Inefficiency	50
1.7.1	Trading on the Predictable Component of Macroeconomic News Reports	50
1.7.1.1	Optimal Trade Timing	52
1.7.2	Trading on Overreaction and Return Reversal	54
1.8	Conclusion	57
2	The Good, the Bad, and the Ambiguous: The Aggregate Stock Market Dynamics around Macroeconomic News	59
2.1	Introduction	59
2.2	Model with Ambiguity	67
2.2.1	Utility Function	67
2.2.2	Three Period Model	68
2.2.3	Equilibrium Prices at $t = 0$ and $t = 1$ with Ambiguity-Aversion	69
2.2.3.1	Price at $t = 1$ with Ambiguity-Aversion	70
2.2.3.2	Price at $t = 0$ with Ambiguity-Aversion	72
2.2.3.3	Volatility of the Risky Asset with Ambiguity-Aversion	75
2.2.4	Equilibrium Prices at $t = 1$ and $t = 0$ under an Unambiguous Benchmark	77
2.2.5	Summary of Model Predictions	78
2.3	The Stock Market Response to Macroeconomic News	82
2.3.1	Data on MNAs	82
2.3.2	Data on High-Frequency Stock Prices	85
2.4	Empirical Results: Whole Sample	86
2.4.1	The Asymmetry Effect and No News is Good News Effect	86
2.4.2	The Ambiguity Premium	91
2.5	State Variable for Ambiguity	96
2.5.1	Variance Risk Premium	96
2.5.2	Does the Variance Risk Premium Predict Stock Returns?	105
2.5.3	Alternate State Variables for Ambiguity	109
2.6	Empirical Results: Time-Varying Ambiguity	113
2.6.1	The Asymmetry Effect and No News is Good News Effect	114
2.6.2	The Ambiguity Premium	118
2.6.3	Risky Asset Volatility	122
2.7	Alternative Explanations	124

2.7.1	Good Times versus Bad Times	125
2.7.2	Volatility Feedback	130
2.8	Conclusion	130
3	Asset Price Reactions to News at the Zero Lower Bound	132
3.1	Introduction	132
3.2	Shadow Rate Term Structure Model	138
3.2.1	Setup and Forward Rate	139
3.2.2	Expected Duration of the ZLB	140
3.2.3	Predictions for Forward Rate Sensitivity to MNAs	141
3.3	Data and Basic Regressions	143
3.3.1	MNAs	143
3.3.2	High-Frequency Asset Prices	145
3.3.3	Basic Regressions for All Assets	146
3.4	Empirical Test of Model Predictions for Interest Rates	148
3.4.1	Methodology	149
3.4.2	Survey Proxy for Expected Duration of the ZLB	151
3.4.3	Fed Funds Futures Proxy for Expected Duration of the ZLB	154
3.5	Stock Market's Sensitivity to MNAs	158
3.5.1	Methodology and Results	158
3.5.2	Decomposition: Interest Rate News vs. Cash Flow and Risk Premium News	158
3.6	Conclusion	165
	References	166
	Appendix A Appendix to Chapter 1	173
A.1	Intraday Stock and Bond Dollar Volume	173
A.2	Order Imbalances	176
A.3	Alternative Tests for Anchoring Bias	183
A.4	Alternative Constructions of the Predictable and Unpredictable Components of News Reports	184
A.5	Model with One Risky Asset	186
A.5.1	$t = 4$	188
A.5.2	$t = 3$	188
A.5.3	$t = 2$	188
A.5.4	$t = 1$	189
A.5.5	$t = 0$	190
A.6	Model with Two Risky Assets	190

A.7	Notes for Figures and Tables	191
A.7.1	Figure 1.6	191
A.7.2	Figure 1.7	191
A.7.3	Table A.2	192
Appendix B	Appendix to Chapter 2	193
B.1	Conditional Distribution	193
B.2	Conditional Means of Normal and Squared Normal Random Variables	193
B.3	Notes for Figures and Tables	194
B.3.1	Figure 2.4	194
B.3.2	Figure 2.7	195
B.3.3	Table 2.9	195
Appendix C	Appendix to Chapter 3	197
C.1	Forward Rate	197
C.2	Comparative Statics	199
C.2.1	Proof of Proposition 1	200
C.2.2	Proof of Proposition 2	200
C.2.3	Proof of Proposition 3	201
C.2.3.1	Proof of Claim 1	201
C.2.4	Proof of Proposition 4	202
C.2.4.1	Proof of Claim 2	204
C.2.4.2	Proof of Claim 3	205
C.2.4.3	Proof of Claim 4	205
C.3	Notes for Figures	206
C.3.1	Figure 3.2	206
C.3.2	Figure 3.3	207
C.3.3	Figure 3.4	207

List of Tables

1.1	Full Sample of MNAs	13
1.2	Summary Statistics of the Pre-Announcement Conditional Drift in Stocks and Bonds	20
1.3	Regressions of the Pre-Announcement Conditional Drift in Stocks and Bonds	23
1.4	Predicting Macroeconomic News Reports	31
1.5	Regressions of Stock and Bond Returns on the Predictable and Unpredictable Components of Macroeconomic News Reports	33
1.6	Regressions of Return Momentum in Stocks and Bonds	36
1.7	Regressions of Return Reversal in Stocks and Bonds	55
2.1	Full Sample of MNAs	83
2.2	The Whole Sample Stock Market Reaction to MNAs	88
2.3	Summary Statistics of MNA Surprises	91
2.4	The Ambiguity Premium and its Decomposition	92
2.5	Summary Statistics of 10-minute Stock Market Returns for Periods with and without MNAs	94
2.6	Summary Statistics of Variance Risk Premium and Components	100
2.7	Predictability of Stock Returns	107
2.8	Stock Market Reaction to MNAs Bucketed into Variance Risk Premium Quintiles	115
2.9	Decomposing the Predictability of Stock Returns	120
2.10	Stock Market Volatility around MNAs Bucketed into Variance Risk Premium Quintiles	124
2.11	Replication of Results Using the Relative Price-to-Earnings Ratio to Bucket the Data	128
3.1	Full sample of MNAs from 12/17/08 to 3/6/14	144
3.2	Reactions to news during the ZLB	147
3.3	Attenuation in the sensitivity of ED futures due to the ZLB using the survey proxy	154
3.4	Attenuation in the sensitivity of ED futures due to the ZLB using the FF24 proxy	157

A.1	Summary Statistics of Pre-Announcement and Announcement Order Imbalances in Stocks and Bonds	178
A.2	Regressions of Pre-Announcement and Announcement Order Imbalances in Stocks and Bonds	180
A.3	Alternative Tests for Anchoring Bias	184
A.4	Predicting Macroeconomic News Reports Grouping All MNAs Together . . .	185
A.5	Regressions of Stock and Bond Returns on the Predictable and Unpredictable Components of Macroeconomic News Reports	187

List of Figures

1.1	Time Series of Macroeconomic News Reports	15
1.2	Intraday Stock and Bond Dollar Volumes on Days with MNAs	17
1.3	The Pre-Announcement Conditional Drift in Stocks and Bonds	19
1.4	Timelines of Events in the Model	39
1.5	Price(s) of the Risky Asset(s) in the Model	41
1.6	Trading on Perfect Foresight of Macroeconomic News Reports	48
1.7	Trading on the Predictable Component of Macroeconomic News Reports with Transaction Costs	51
1.8	Trading on Return Reversal after MNAs with Transaction Costs	56
2.1	Graphical Representation of Model Predictions	79
2.2	Comparison of Intraday versus Daily Stock Market Data	87
2.3	Time Series of Standardized News for the Whole Sample of MNAs	90
2.4	Time Series of Variance Risk Premium and Components	101
2.5	Monthly Variance Risk Premium with Annotations	103
2.6	Time Series of Macroeconomic Forecast Dispersions	110
2.7	Stock Market Reaction to MNAs Bucketed into Variance Risk Premium Quintiles	116
2.8	Time Series of Intraday and Daily Returns around MNAs	123
2.9	Time Series of Absolute and Relative Price-to-Earnings Ratios	127
3.1	Proxies for the expected duration of the ZLB	133
3.2	Time-varying sensitivities of interest rates to MNAs and the impact of the ZLB using the survey proxy	152
3.3	Time-varying sensitivities of interest rates to MNAs and the impact of the ZLB using the FF24 proxy	155
3.4	Time-varying sensitivities of stocks to MNAs and the impact of the ZLB . . .	159
3.5	Decomposition of the stock market reaction to MNAs into interest rate news versus cash flow and risk premium news	163
A.1	Intraday Stock and Bond Dollar Volumes on Days without MNAs	174

A.2	Differences in Intraday Stock and Bond Dollar Volumes between Days with and without MNAs	175
A.3	Pre-Announcement Order Imbalances in Stocks and Bonds	182

Acknowledgments

First, I would like to thank my advisers. I am particularly grateful to my dissertation committee: John Campbell, Jeremy Stein, Robin Greenwood, and Owen Lamont. I have been fortunate to have also received excellent advising throughout graduate school from Lauren Cohen, Sam Hanson, Chris Malloy, Neil Shephard, Erik Stafford, Andrei Shleifer, Adi Sunderam, and Luis Viceira. Thank you not only for your academic and research guidance but also for your career and life advice, without which I would not be where I am today.

Second, I would like to give a big shout-out to my fellow graduate students and friends. In particular, I would like to thank my seatmate Tom Powers, Jon Rhinesmith, and Filippo Mezzanotti for all the good times (and for getting through all the not-so-good times together). I learned so much from you. Thank you also to my friends in the other doctoral programs/cohorts and the MBA Class of 2015 (especially, Section J). You have made this journey memorable.

Finally, I would like to thank my parents: Suwei and Xiaofeng Zhou. You have been with me every step of the way.

To my parents

Introduction

My dissertation is composed of three chapters that are unified by their exploration of asset prices and macroeconomic news announcements (MNAs). With respect to asset prices, my main focus is on the price discovery process: how do asset prices reveal information relevant for asset fundamentals? Through my research, I provide new answers to this question. My work gets at core issues in asset pricing: whether financial markets are informationally efficient; why some assets earn unconditionally high premia; and how the sensitivity of prices to information varies over time and across assets.

Each of the three chapters explores the price discovery process using the unique setting of U.S. MNAs, which are made by government agencies and private-sector organizations and cover macroeconomic data on inflation, output, and unemployment. Analyzing financial markets in this setting deepens our understanding of how asset prices reflect information about macroeconomic fundamentals. At the same time, the results have macroeconomic implications; for example, the assumptions of monetary policy models in theory and the effectiveness of unconventional monetary policy in practice.

Chapter one studies how the views of sophisticated traders are impounded into stocks and bonds around MNAs. I find evidence that sophisticated traders trade on predictions of macroeconomic news reports before announcements and obtain their informational advantage using public information. Specifically, consensus forecasts of upcoming data releases suffer from anchoring bias and overweight past data releases. By correcting this bias, sophisticated traders can predict news reports. The results suggest that stock and bond markets are inefficient in this setting. Over time, there is a “late trading puzzle”:

sophisticated traders can predict news reports days before announcements but appear to trade these predictions into stock and bond prices just hours before announcements. Across assets, there is a related puzzle: the predictable component of news reports is eventually fully impounded into bonds but only partially impounded into stocks. Stocks but not bonds react to announcements of the predictable component and display return momentum. I develop a model showing that market inefficiency can arise when unsophisticated traders neglect public information that predicts news reports, and risk management concerns deter sophisticated traders from acting on their informational edge. Trading earlier and trading riskier assets such as stocks exposes sophisticated traders to greater risk. As a result, sophisticated traders wait to trade and trade safer assets such as bonds.

In chapter two, I argue that the behavior of ambiguity-averse investors in response to ambiguous macroeconomic data can account for stock market dynamics around MNAs. Using a representative agent model in which the investor is averse to ambiguity (Knightian uncertainty) and sees an ambiguous piece of news about the fundamental value of a risky asset, I show a number of predictions for the dynamics of stocks around news: stocks respond more strongly to bad news than to good news, respond positively to neutral news, and increase on average through news. In times of high ambiguity, the magnitudes of each effect is larger, and the volatility of stocks around news changes in a predictable manner as well. I provide empirical evidence consistent with the model by analyzing the high-frequency behavior of the aggregate stock market around MNAs from November, 1997 to March, 2014. The model helps to understand features of the data that challenge existing frameworks; e.g., the findings that the stock market reacts especially strongly to bad news versus good news during crisis periods and that about 1/3 of equity returns in the 17 year sample accrues in the 10 minutes around the release of macroeconomic data. In addition to providing evidence for the role of ambiguity in financial markets generally and in how financial assets reflect macroeconomic shocks specifically, the empirical results also have implications for the behavior of investors. Investors treat bad news as more relevant in bad times than in good times but treat good news the same in good and bad times.

The third chapter analyzes the reaction of interest rates and the stock market to MNAs at the zero lower bound (ZLB). I start by using a shadow rate term structure model to formulate three predictions for the sensitivity of interest rates to MNAs. First, “better”-than-expected macroeconomic data increases interest rates. Second, as the expected duration of the ZLB increases, whether because economic conditions are worse or because monetary policy changes, interest rates become less sensitive to macroeconomic data. Third, this attenuation in the sensitivity of interest rates is largest for intermediate-maturity rates. I verify these predictions by using a broad sample of MNAs and high-frequency intraday futures data on interest rates. Turning to stocks, I show that the stock market’s reaction to MNAs can be decomposed into an interest rate news term that is directly related to interest rates’ reaction to MNAs and a cash flow plus risk premium news term. Using the same sample of MNAs and high-frequency intraday futures data on the stock market, I empirically estimate the stock market’s sensitivity to macroeconomic data as well as that of the constituent news terms. Based on the interest rate news term alone, the expected duration of the ZLB should increase the sensitivity of stocks to macroeconomic news. The data furthermore suggests that the expected duration of the ZLB decreases the magnitude of the cash flow plus risk premium news term.

Chapter 1

Sophisticated Trading and Market Efficiency: Evidence from Macroeconomic News Announcements

1.1 Introduction

Are markets efficient or inefficient? This question is central to economics. In efficient markets, asset prices quickly and completely reflect public information relevant for asset fundamentals. Conversely, asset prices slowly and partially incorporate such information in inefficient markets. In this paper, I study sophisticated trading around U.S. macroeconomic news announcements (MNAs) and show how and why the aggregate U.S. stock and bond markets are inefficient in this setting.

Government agencies and private-sector organizations regularly make MNAs and release macroeconomic data on inflation, output, and unemployment. To assess the informational content of an announcement, observers look at the difference between the announced value of the data release and economists' forecasts. Consistent with common practice, I

define this difference as the “news report” such that a positive news report corresponds to higher-than-expected inflation and output and lower-than-expected unemployment. A negative news report is the opposite. As documented in an extensive and active literature, MNAs substantially impact the prices of many assets; in particular, the U.S. aggregate stock market, or “stocks,” and U.S. Treasury bonds, or “bonds.”¹ Stocks increase {decrease} and bonds decrease {increase} on announcements of positive {negative} news reports about the economy.²

In the first part of this paper, I present evidence that sophisticated traders predict macroeconomic news reports and trade on this informational advantage in stocks and bonds in the hours prior to the official release times. For a comprehensive sample of 18 MNAs from July 2003 to March 2014, stocks increase {decrease} and bonds decrease {increase} *before* positive {negative} news reports. More specifically, in the hours prior to positive {negative} news reports of greater than 1 standard deviation in magnitude, the E-mini S&P 500 futures contract increases up to 7.8 basis points {decreases up to 3.1 basis points}, and the 10-Year U.S. Treasury Note futures contract decreases up to 2.8 basis points {increases up to 3.8 basis points}. This pre-announcement conditional drift is sizable relative to the announcement return: up to 57% and 26% of the price adjustment of stocks and bonds, respectively, to news reports occurs in the pre-announcement window.

What is the source of sophisticated traders’ informational advantage? In the second part

¹Many papers have explored how asset prices reflect macroeconomic risk; for example, Jones, Lamont, and Lumsdaine (1998), Andersen, Bollerslev, Diebold, and Vega (2003, 2007), Boyd, Hu, and Jagannathan (2005), Faust, Rogers, Wang, and Wright (2007), Faust and Wright (2012), and Savor and Wilson (2013, 2014). These studies have investigated an assortment of assets (such as stocks, bonds, and currencies) in various sample periods (for instance, expansionary and contractionary states). How asset prices vary in response to macroeconomic data can give insight into other areas of economics as well. Swanson and Williams (2013, 2014) and Zhou (2014), for example, show that the Federal Reserve’s zero lower bound attenuates interest rate sensitivity to news about the economy. Zhou (2015) argues that the behavior of ambiguity-averse investors in response to ambiguous macroeconomic data can explain stock market dynamics around MNAs.

²Throughout this paper, I use curly braces to distinguish between the cases of positive news reports and negative news reports. A sentence without the braced words corresponds to the case of positive news reports; in this instance, “Stocks increase and bonds decrease on announcements of positive news reports about the economy.” A sentence that replaces the words before braces with the words inside braces corresponds to the case of negative news reports; in this instance, “Stocks decrease and bonds increase on announcements of negative news reports about the economy.”

of this paper, I provide an answer to this natural follow-up question. The data suggest that sophisticated traders are able to obtain at least some of their informational advantage via the “homework channel.” That is, sophisticated traders skillfully process public information to predict macroeconomic news reports and trade on these predictions. I first demonstrate that homework is doable by showing an example in which public information predicts news reports. In my sample of MNAs, economists’ forecasts display anchoring bias and are systematically biased to the values of past releases. When forecasts are higher than past data realizations, forecasts are not high enough due to overweighting of prior data, and conversely for forecasts lower than past data realizations. As such, differences between forecasts and past data realizations positively predict news reports. Since this predictor variable is common knowledge prior to the release of macroeconomic data, public information predicts news reports. Second, I find support that the homework is being acted on. There is a strong relationship between the predictability of news reports based on anchoring bias and the pre-announcement conditional drift. Stocks increase {decrease} and bonds decrease {increase} before news reports predicted to be positive {negative}. It is striking that public information based on the straightforward example of anchoring bias predicts news reports and explains so much of the pre-announcement conditional drift. Market participants have the incentives and resources to analyze a more comprehensive set of public information using more advanced techniques. Thus, sophisticated traders in practice likely obtain an even greater informational advantage via the homework channel than shown in my research.

Based on the above results, I show in the third part of this paper that stock and bond markets are inefficient both over time and across assets. Over time, there is a “late trading puzzle”: the public information that predicts macroeconomic news reports is impounded into stocks and bonds with a significant lag. More precisely, the public information that predicts news reports is available days prior to the official release times of MNAs. Contrary to the logic of market efficiency, however, the predictable component of news reports is impounded into asset prices just hours prior to the official release times.

Across assets, there is a related puzzle: by the announcement times, the public in-

formation that predicts news reports is fully incorporated into bonds but only partially incorporated into stocks. That is, the bond market eventually becomes efficient, while the stock market remains inefficient. Stocks but not bonds react to announcements of the predictable component of news reports based on anchoring bias. A 1 standard deviation positive {negative} value of the predictable component elicits a statistically significant 1.5 basis points increase {decrease} in stocks. Bonds are more efficient than stocks in another way: stocks but not bonds display return momentum between pre-announcement and announcement returns. For stocks, a 10 basis points higher pre-announcement return is statistically significant in predicting up to 0.6 basis point higher announcement return.

In the fourth part of this paper, I provide a model-based explanation for why stocks and bonds are inefficient around MNAs. In the model, all traders see a public announcement of a news report that provides information about risky asset values. Sophisticated traders obtain advance knowledge of this news report, possibly through the skillful processing of public information. The model makes three key assumptions. First, unsophisticated traders do not know beforehand the news report and thus neglect public information that predicts the news report. Second, sophisticated traders have risk management concerns, as modeled through risk aversion. Third, traders do not extract each others' information from prices.

In the baseline version of the model, there is one risky asset. Sophisticated traders have two opportunities to trade on their informational advantage before the announcement: during an early period and during a late period. In each period, holders of the risky asset face fundamental risk. If the fundamental risk is large in the early period relative to the late period, risk-averse sophisticated traders recognize that they can achieve a higher risk-adjusted return by trading little on their information early on and instead waiting to trade a lot later on. As a result, the advance knowledge of sophisticated traders is inefficiently impounded into the risky asset with a delay.

Just as sophisticated traders avoid risk in the time series by trading later, the model predicts that sophisticated traders avoid risk in the cross section by trading less risky assets. In an extension of the model, the public announcement of a news report provides information

about the values of two risky assets. Prior to the announcement, the fundamental risk associated with holding one of the assets versus the other is higher (in both early and late periods). Risk-averse sophisticated traders recognize that they can achieve a higher risk-adjusted return by trading the less risky asset than the riskier one. In doing so, sophisticated traders impound more of their informational edge in the former than the latter, which makes the less risky asset more efficiently priced than the riskier asset. Compared to the riskier asset, the less risky asset reacts less strongly to the public announcement of the news report already known to sophisticated traders and displays less return momentum between pre-announcement and announcement returns.

The data are consistent with the model's explanation of market inefficiency. Sophisticated traders with advance knowledge of macroeconomic news reports find it less attractive to trade stocks and bonds in the days versus hours before announcements due to higher risk. These sophisticated traders also find it less attractive to trade stocks versus bonds before announcements because of greater risk. As evidence, I consider *hypothetical* trading strategies that go long {short} stocks and short {long} bonds before positive {negative} news reports above {below} a given threshold. For both assets, the Sharpe ratios are lower when the strategies are implemented in the days before announcements than the hours before announcements. The reason is that the return standard deviations are higher and not that the return means are lower. Similarly, before announcements, the Sharpe ratios are lower for strategies implemented in stocks versus bonds. Once again, the reason is that the return standard deviations are higher and not that the return means are lower.

In the fifth and final part of this paper, I investigate whether the inefficiency of stock and bond markets creates attractive trading opportunities in practice. To do so, I construct *implementable* trading strategies using the predictable component of macroeconomic news reports derived from anchoring bias. These strategies go long {short} stocks and short {long} bonds before news reports predicted to be above {below} a given threshold. Even after transaction costs, the Sharpe ratios are sizable: the Sharpe ratio is as high as 0.82 for stocks and 1.12 for bonds. These risk-adjusted returns are consistent with the interpretation

that sophisticated traders are deterred by risk from trading the predictable component into prices and consequently leave money on the table.

This paper builds on the broader informed trading literature in which some, sophisticated traders have an informational advantage.³ A number of studies empirically document settings in which there is or appears to be informed trading; for example, Meulbroek (1992) around *ex-post* insider trading events based on U.S. Securities and Exchange Commission (SEC) charges, Acharya and Johnson (2007) in the credit default swap market, Irvine, Lipson, and Puckett (2007) around sell-side analysts' recommendations, and Campbell, Ramadorai, and Schwartz (2009) around earnings announcements. Many of these papers focus on traders acting upon an informational edge relevant for the values of individual securities. Recently, a few researchers have presented evidence of traders with advance knowledge of macroeconomic events who trade in aggregate markets. Cieslak, Morse, and Vissing-Jorgensen (2014) and Bernile, Hu, and Tang (2015) find patterns in the U.S. stock market consistent with sophisticated traders acting on information concerning the Federal Reserve. Bernile, Hu, and Tang (2015) and Kurov, Sancetta, Strasser, and Wolfe (2015) also look for evidence of informed trading around MNAs.⁴ Much of the aforementioned literature either attributes informed trading to (illegal) information leakage or is unable to specify the mechanism by which sophisticated traders obtain their information.

My research contributes to a separate strand of the informed trading literature that is concerned with how traders optimally trade and impound their informational advantage into asset prices. A number of theories provide insights into optimal trade timing. Kyle (1985) and Back (1992) show in multi-period microstructure models that a single, risk-neutral informed trader spreads his trading out over time such that private information is gradually

³I use the terms "informed" and "sophisticated" interchangeably.

⁴Bernile, Hu, and Tang (2015) were unable to find evidence of informed trading. This no result is likely due to their small sample of MNAs (3 types of announcements versus 18 in my analysis) and their focus on the last 30 minutes prior to the official release times. Using an expanded set of MNAs, Kurov, Sancetta, Strasser, and Wolfe (2015) do find evidence of informed trading in the last 30 minutes prior to the official release times. This result is a subset of my finding in the first part of this paper that sophisticated trading appears to take place in the hours prior to the official release times.

incorporated into asset prices. Holden and Subrahmanyam (1992, 1994) demonstrate that when there are multiple informed traders, even if they are risk-averse, competition ensues, which precipitates a rush to trade. Traders aggressively front-run each other, and private information is quickly incorporated into asset prices. On the other hand, Foster and Viswanathan (1996) and Back, Cao, and Willard (2000) argue that informed traders with heterogeneous private information wait to trade. Each informed trader is incentivized to wait in order to conserve his information monopoly and prevent others from using prices to infer his private information. Under this scenario, private information is slowly incorporated into asset prices. While there is a sizable theoretical literature on optimal trade timing, there is little complementary empirical work on whether sophisticated traders actually rush to trade, wait to trade, or spread out their trading. One example of such a study is Di Mascio, Lines, and Naik (2015) who document the gradual alpha decay of stocks purchased by institutional investors and argue that this phenomenon is indicative of sophisticated traders waiting to trade on their information. Other frameworks analyze aspects of optimal trade aside from timing. In De Long, Shleifer, Summers, and Waldmann (1990) and Brunnermeier (2005), for example, sophisticated traders take advantage of unsophisticated traders. The sophisticated front-run the unsophisticated before public announcements and then flip their positions after public announcements. Kadan, Michaely, and Moulton (2014) find supportive empirical evidence in the trading of institutional investors around sell-side analysts' recommendations. Some of the predictions of the above theories are outside of my model. At the end of this paper, I show features of the data that offer new evidence for the existing theories in the literature and motivate a more comprehensive extension of my model.

Finally, my research is related to the substantial body of work concerning news about the economy. As discussed in the beginning of this paper, MNAs are frequently used to measure the impact of macroeconomic shocks on financial assets. The standard approach is to regress returns in narrow announcement windows on macroeconomic news reports, which proxy for the informational content of announcements. If traders informed about

news reports act on this information prior to announcements, however, the aforementioned regression inaccurately estimates the extent to which asset prices reflect macroeconomic risk. The presence of sophisticated traders around MNAs has monetary policy implications as well. Romer and Romer (2000) show that Federal Reserve forecasts of the economy are more accurate than those of private-sector economists. The authors interpret the result as evidence that the Federal Reserve has an informational advantage about the economy over market participants. Yet my research suggests that some market participants, the sophisticated traders, also have superior information about the economy versus private-sector economists. When properly accounting for sophisticated traders, the Federal Reserve may have a smaller or possibly no informational advantage about the economy over market participants. Such a conclusion calls into question the asymmetric information assumption of monetary policy models that central banks have better knowledge of the economy than other agents.

The organization of this paper is as follows. Section 1.2 introduces the sample of MNAs as well as the data on stocks and bonds. Section 1.3 presents evidence of sophisticated trading in the hours ahead of official release times, as shown by the pre-announcement conditional drift of stocks and bonds. The results in Section 1.4 suggest that sophisticated traders obtain their informational advantage via the homework channel. In Section 1.5, I outline how stock and bond markets are inefficient over time and across assets. Section 1.6 provides a model-based explanation of the documented market inefficiency. Section 1.7 shows that trading on market inefficiency yields sizable risk-adjusted returns. Section 1.8 concludes.

1.2 Empirical Setting

The empirical setting in which I analyze sophisticated trading and market efficiency is the stock and bond markets around MNAs. The first subsection introduces the data on MNAs, and the second subsection introduces the data on stocks and bonds.

1.2.1 MNAs

Table 1.1 shows the sample of MNAs considered in this paper. For each data series, the table presents the name, units, number of observations, start date, end date, frequency, government agency or private-sector organization responsible, and intraday announcement timestamp. In all, I analyze 18 announcements, which cover the lion's share of important MNAs. Data for the majority of announcements begin in July 2003, when the data on stocks and bonds begin, as discussed further down, and extend to March 2014. The Existing Home Sales and Pending Home Sales data series start in 2005. All of the announcements occur once a month excluding Initial Jobless Claims, which is a weekly data release. The data in my sample are released at either 8:30 AM, 9:15 AM, or 10:00 AM ET. Macroeconomic news releases are pre-scheduled and occur precisely at those pre-scheduled times. The former feature ensures that if sophisticated trading were to occur, it would start prior to well-defined release times. The latter feature means that evidence suggestive of sophisticated trading is unlikely to be contaminated by early data releases. For days on which multiple announcements occur, I keep in my sample the announcements that occur first and drop the other announcements. By ensuring no MNAs occur prior to the MNAs in my sample, the selection criteria removes a source of noise from the analysis of asset returns prior to macroeconomic releases.⁵

Investigating sophisticated trading around MNAs requires defining the informational content of MNAs that traders attempt to predict. I quantify the informational content of a given economic indicator by calculating the difference between the announced value of the data release and economists' forecasts. Specifically, I construct the news report variable

$$NR_t = \frac{A_t - F_t}{\hat{\sigma}}, \quad (1.1)$$

with A_t the announced value of the data release and F_t the median forecasted value of the data release from a survey of economists. $\hat{\sigma}$ is the sample standard deviation of $A_t - F_t$. I obtain both A_t and F_t from Bloomberg, which is the standard source of such data for

⁵Extending my sample to the full set of announcements does not substantively alter the results.

Table 1.1: Full Sample of MNAs

Event Name	Units	N	Start	End	Freq.	Source	Time (ET)	Sign
Change in Nonfarm Payrolls	Thousands	129	07/03/2003	03/07/2014	M	BLS	8:30 AM	1
Consumer Confidence	Index	110	07/29/2003	03/25/2014	M	CB	10:00 AM	1
CPI	% ch. mom	129	07/16/2003	03/18/2014	M	BLS	8:30 AM	1
Durable Orders	% ch. mom	129	07/25/2003	03/26/2014	M	Census	8:30 AM	1
Existing Home Sales	Millions	73	04/25/2005	03/20/2014	M	NAR	10:00 AM	1
Factory Orders	% ch. mom	69	07/02/2003	03/06/2014	M	Census	10:00 AM	1
Housing Starts	Millions	127	07/17/2003	03/18/2014	M	Census	8:30 AM	1
Industrial Production	% ch. mom	35	09/15/2003	03/17/2014	M	Fed	9:15 AM	1
Initial Jobless Claims	Thousands	561	07/03/2003	03/27/2014	W	ETA	8:30 AM	-1
ISM Manufacturing	Index	84	09/02/2003	03/03/2014	M	ISM	10:00 AM	1
ISM Non-Manufacturing	Index	81	08/05/2003	03/05/2014	M	ISM	10:00 AM	1
Leading Index	Index	28	07/21/2003	03/20/2014	M	CB	10:00 AM	1
New Home Sales	Thousands	51	10/27/2003	03/25/2014	M	Census	10:00 AM	1
Pending Home Sales	% ch. mom	58	06/01/2005	03/27/2014	M	NAR	10:00 AM	1
Personal Income	% ch. mom	128	08/01/2003	03/28/2014	M	BEA	8:30 AM	1
PPI	% ch. mom	129	07/11/2003	03/14/2014	M	BLS	8:30 AM	1
Retail Sales	% ch. mom	129	07/15/2003	03/13/2014	M	Census	8:30 AM	1
Unemployment Rate	% rate	129	07/03/2003	03/07/2014	M	BLS	8:30 AM	-1

Notes: Freq. refers to monthly (M) or weekly (W). Source uses the following acronyms: BEA, Bureau of Economic Analysis; BLS, Bureau of Labor Statistics; CB, Conference Board; Census, Census Bureau; ETA, Employment and Training Administration; Fed, Federal Reserve Board of Governors; ISM, Institute for Supply Management; and NAR, National Association of Realtors.

both academics and practitioners. The news report variable is identical to the measure of the informational content of MNAs employed by Balduzzi, Elton, and Green (2001) and subsequent papers in the literature of asset price reactions to MNAs. Consistent with the literature, I define the news report variable such that a positive value corresponds to higher-than-expected inflation and output and lower-than-expected unemployment. For this reason, I multiply the Initial Jobless Claims and Unemployment Rate data by -1 while keeping the signs of the other data unchanged, as shown by the sign column of Table 1.1.

As constructed, the news report variable provides a single metric that is standardized across different types of news about macroeconomic fundamentals and allows for comparability. This is important because different types of news are released with different units. Figure 1.1 plots the time series of all the news reports in my sample. NR_t is well-behaved in the sense that it has a mean around zero and is symmetric. On average, macroeconomic news reports are neutral, and positive news reports are as likely as negative news reports. Moreover, the distribution of news reports appears stable in the time series.

1.2.2 Stocks and Bonds

I study sophisticated trading and market efficiency in the U.S. aggregate stock market and U.S. Treasury bonds. I proxy for stock and bond markets using the front E-mini S&P 500 (ES) futures contract and 10-Year U.S. Treasury Note (TY) futures contract, respectively. There are several reasons for my choice of assets and securities. First, as discussed in the beginning of this paper, stocks and bonds reflect market-wide risk, and, in particular, macroeconomic risk. Traders with an informational advantage concerning MNAs can profitably express their views in both assets. Second, stock and bonds futures are available for trading around the clock.⁶ Importantly, futures are available for trading in the hours around MNAs, which is when I find the strongest evidence of sophisticated trading. Third, ES and TY futures are among the most liquid securities. As such, these securities are attractive to traders seeking

⁶For weekday t , the ES futures contract is available for trading from 5:00 PM CT of day $t - 1$ to 4:15 PM CT of day t with a halt from 3:15 PM CT to 3:30 PM CT. The TY futures contract is available for trading from 5:00 PM CT of day $t - 1$ to 4:00 PM CT of day t .




Figure 1.1: Time Series of Macroeconomic News Reports

Notes: Figure plots the news report variable NR_t for the full sample of MNAs. For a given MNA, the news report variable is the announced data less the forecasted data based on a Bloomberg survey of economists, all normalized by the sample standard deviation of the announced less the forecasted difference for that MNA: $NR_t = (A_t - F_t) / \hat{\sigma}$.

to minimize trading costs and price impact. Other proxies for stocks and bonds such as ETFs are much less liquid and unavailable for trading for most of the overnight trading session in the U.S. Fourth, ES and TY have low margins and thus offer easy leverage for sophisticated traders wanting to maximize the value of their information.


Investigating trading in the hours around MNAs necessitates high-frequency, intraday data. I obtain minute-level price and volume data on ES and TY futures from Tick Data, a data vendor. The data are available from July 2003 to March 2014.⁷

In Figure 1.2, I use the volume data to substantiate my claims about ES and TY liquidity. Panels A and B plot the intraday, minute-by-minute and cumulative dollar volumes for stocks and bonds, respectively. The values are averaged across all days in the sample that have MNAs. The three dashed vertical lines correspond to the 8:30 AM, 9:15 AM, and 10:00 AM ET announcement times. For both assets, there are substantial spikes in trading activity corresponding to these timestamps.⁸ For stocks, the jumps in volume are comparable in magnitude to the well-known surges in trading associated with the 9:30 AM ET open and 4:00 PM ET close of the U.S. stock market. It is clear that when they are announced, macroeconomic data command the attention of market participants who react by trading ES and TY. Figure 1.2 also highlights the abundant liquidity of stock and bond futures. The average daily cumulative dollar volumes are on the order of \$100 billion. There is even liquidity in the early morning hours around MNAs: by 8:30 AM ET, the average cumulative dollar volume is over \$7 billion for ES futures and over \$17 billion for TY futures. Appendix A.1 compares the stock and bond dollar volumes on days with MNAs to those on days without MNAs. ES and TY are more liquid in the former case than the latter case, in particular on and immediately after timestamps associated with MNAs.

⁷Data prior to July 2003 are missing information on overnight trading in the U.S. Since MNAs occur during morning trading in the U.S., the overnight data are important for studying trading in the hours prior to announcements.

⁸There is no noticeable spike associated with the middle dashed vertical line. This is because only one of the MNAs in my sample, Industrial Production, is announced at 9:15 AM ET. All the others are announced at 8:30 AM or 10:00 AM ET.

Panel A: Stocks



Panel B: Bonds




Figure 1.2: Intraday Stock and Bond Dollar Volumes on Days with MNAs

Notes: Panels A and B plot the intraday minute-by-minute and cumulative dollar volumes for stocks (front ES futures contract) and bonds (front TY futures contract), respectively. The values are averaged across all days in the sample that have MNAs. The three dashed vertical lines correspond to 8:30 AM, 9:15 AM, and 10:00 AM ET.

1.3 Sophisticated Trading around MNAs

1.3.1 Pre-Announcement Conditional Drift


The first set of results shows evidence consistent with sophisticated traders predicting macroeconomic news reports and trading on this informational advantage in stocks and bonds in the hours prior to official release times. In Panels A and B of Figure 1.3, I plot the mean cumulative return (in basis points or bps) starting 360 minutes before MNAs for stocks and bonds, respectively. I do so for three samples of MNAs: the “Full Sample” of announcements, those announcements with “News Report > 1 ” or greater than 1 standard deviation positive news reports, and those announcements with “News Report < -1 ” or less than 1 standard deviation negative news reports. The news report variable is constructed based on Eq. (1.1). First, consider the results for stocks in Panel A. Upon announcements of positive {negative} macroeconomic news reports, stocks increase {decrease}. In the hours before announcements, stocks also increase {decrease}. Parallel results hold for bonds in Panel B. Upon announcements of positive {negative} macroeconomic news reports, bonds decrease {increase}. In the hours before announcements, bonds also decrease {increase}. This pre-announcement conditional drift is the key piece of support for sophisticated trading in stocks and bonds around MNAs.

To quantify the conditional drift, I calculate summary statistics in Table 1.2. Row $-x$ corresponds to the cumulative pre-announcement return (in bps) from x minutes to 5 minutes before MNAs, and row 0 corresponds to the ± 5 minute announcement return (in bps) around MNAs. For Table 1.2 and subsequent results in this paper, I define the end of the pre-announcement window and the start of the announcement window to be 5 minutes before MNAs. This demarcation minimizes the risk that early releases of macroeconomic data are contributing to pre-announcement findings.^{9,10} Symmetrically, I define the end of

⁹Some MNAs in my sample have been accidentally released earlier than scheduled, but, to the best of my knowledge, only at most a few seconds before the official release times.

¹⁰Stipulating that the pre-announcement window ends 5 minutes before MNAs may understate the magnitude of sophisticated trading if sophisticated traders literally wait until the last moment to act on their

Panel A: Stocks



Panel B: Bonds




Figure 1.3: *The Pre-Announcement Conditional Drift in Stocks and Bonds*

Notes: Panels A and B plot the mean cumulative return (in bps) starting 360 minutes before MNAs for stocks (front ES futures contract) and bonds (front TY futures contract), respectively. “Full Sample” includes the full sample of MNAs. “News Report > 1” includes only the MNAs with greater than 1 standard deviation positive news reports; that is, $NR_t > 1$. “News Report < -1” includes only the MNAs with less than 1 standard deviation negative news reports; that is, $NR_t < -1$. NR_t measures the informational content of MNAs and is defined in Figure 1.1.

Table 1.2: Summary Statistics of the Pre-Announcement Conditional Drift in Stocks and Bonds

Panel A: Stocks

	Full Sample ($n = 2179$)			News Report > 1 ($n = 291$)			News Report < -1 ($n = 286$)		
	Mean	Std. Dev.	t -stat	Mean	Std. Dev.	t -stat	Mean	Std. Dev.	t -stat
-720	4.5	52.0	4.1	10.4	60.3	2.9	2.3	57.0	0.7
-660	4.3	51.2	3.9	9.1	60.0	2.6	2.4	53.8	0.8
-600	4.3	48.9	4.1	9.5	57.2	2.8	2.8	51.2	0.9
-540	4.7	48.4	4.5	10.4	57.3	3.1	2.9	50.4	1.0
-480	4.5	48.3	4.4	10.9	55.7	3.4	1.8	48.6	0.6
-420	3.7	48.4	3.5	10.8	55.2	3.3	0.2	45.2	0.1
-360	2.9	46.3	2.9	7.7	54.8	2.4	-0.5	41.9	-0.2
-300	1.8	42.7	1.9	8.3	50.2	2.8	-1.6	40.5	-0.7
-240	1.2	40.0	1.4	6.7	47.6	2.4	-1.4	38.0	-0.6
-180	0.5	34.6	0.7	5.1	39.5	2.2	-2.2	33.6	-1.1
-120	0.9	30.4	1.4	3.6	38.8	1.6	-0.6	30.8	-0.3
-60	0.3	23.3	0.6	3.0	29.1	1.8	0.7	24.2	0.5
0	1.1	26.7	1.9	8.7	29.3	5.1	-15.5	34.2	-7.7

Panel B: Bonds

	Full Sample ($n = 2179$)			News Report > 1 ($n = 291$)			News Report < -1 ($n = 286$)		
	Mean	Std. Dev.	t -stat	Mean	Std. Dev.	t -stat	Mean	Std. Dev.	t -stat
-720	-0.5	19.8	-1.2	-3.1	19.4	-2.8	2.6	20.4	2.1
-660	-0.7	19.3	-1.6	-3.4	19.3	-3.0	2.6	19.7	2.2
-600	-0.9	19.0	-2.2	-3.7	19.1	-3.3	2.1	19.3	1.9
-540	-1.0	18.9	-2.4	-3.7	19.3	-3.3	2.3	19.0	2.1
-480	-1.2	18.4	-3.0	-4.0	19.2	-3.6	2.1	18.5	1.9
-420	-1.0	17.9	-2.7	-3.9	18.4	-3.6	2.0	18.4	1.9
-360	-0.8	16.8	-2.2	-3.0	16.7	-3.1	2.8	17.3	2.7
-300	-0.8	16.0	-2.2	-3.4	16.6	-3.5	2.5	17.6	2.6
-240	-0.8	15.0	-2.5	-3.0	15.8	-3.3	2.0	17.2	2.0
-180	-0.8	14.0	-2.6	-3.2	15.9	-3.4	1.3	15.8	1.4
-120	-0.7	12.5	-2.7	-2.3	14.0	-2.8	1.5	14.5	1.8
-60	-0.7	9.2	-3.7	-2.0	10.4	-3.3	0.3	10.3	0.6
0	0.2	19.3	0.5	-8.5	21.0	-6.9	11.4	23.6	8.2

Notes: Panels A and B present summary statistics for stocks (front ES futures contract) and bonds (front TY futures contract), respectively. Row $-x$ corresponds to the cumulative pre-announcement return (in bps) from x minutes to 5 minutes before MNAs. Row 0 corresponds to the ± 5 minute announcement return (in bps) around MNAs. The “Full Sample” columns include the full sample of MNAs. The “News Report > 1 ” columns include only the MNAs with greater than 1 standard deviation positive news reports; that is, $NR_t > 1$. The “News Report < -1 ” columns include only the MNAs with less than 1 standard deviation negative news reports; that is, $NR_t < -1$. NR_t measures the informational content of MNAs and is defined in Figure 1.1.

the announcement window and start of the post-announcement window to be 5 minutes after MNAs. After 5 minutes, the vast majority of activity in response to announcements is completed, as reflected in the renormalization of trading volumes and return volatilities. The ± 5 minute announcement window is an interval comparable to announcement windows used in the previously cited literature.

Panel A presents summary statistics for stocks. In the “Full Sample” of MNAs, stocks unconditionally increase both in the 10-minute window around announcements (mean of 1.1 bps) and in the pre-announcement period (mean of up to 4.7 bps).^{11,12} While this unconditional pattern is interesting, I concentrate on the conditional results. As shown in the “News Report > 1 ” and “News Report < -1 ” columns, the announcement return of stocks is positive (mean of 8.7 bps) for greater than 1 standard deviation positive news reports and negative (mean of -15.5 bps) for less than 1 standard deviation negative news reports. The announcement return is preceded by the pre-announcement return in the same direction. Stocks increase (up to 10.9 bps on average) in the hours prior to greater than 1 standard deviation positive news reports and decrease (up to -2.2 bps on average) in the hours prior to less than 1 standard deviation negative news reports.

Panel B presents analogous summary statistics for bonds. In the “Full Sample” of MNAs in Panel B of Table 1.2, bonds are flat through announcements (mean of 0.2 bps) and unconditionally decrease before announcements (mean of up to -1.2 bps).¹³ Opposite of stocks, bonds decrease (mean of -8.5 bps) on announcements of greater than 1 standard

informational advantage.

¹¹The announcement premium earned by stocks around MNAs is related to Savor and Wilson (2013), who show the same result using daily data, and to Frazzini and Lamont (2007), who document an earnings announcement premium for individual stocks. Zhou (2015) explains the announcement premium around MNAs as partly resulting from ambiguity-averse investors facing ambiguous macroeconomic news.

¹²The pre-announcement drift of stocks before MNAs is related to the pre-FOMC announcement drift documented by Lucca and Moench (2015). They show that the U.S. aggregate stock market unconditionally increases in the hours prior to FOMC meetings. My finding of the pre-announcement drift of stocks before MNAs, which the authors were unable to find, suggests that the unconditional increase of stocks prior to macroeconomic events may be a more widespread phenomenon.

¹³In other words, there is no announcement premium for bonds around MNAs, but there is a pre-announcement drift of bonds before MNAs.

deviation positive news reports and increase (mean of 11.4 bps) on announcements of less than 1 standard deviation negative news reports. Identical to the case of stocks, however, the announcement return of bonds is preceded by the pre-announcement return of bonds in the same direction. Bonds decrease (up to -4.0 bps on average) in the hours prior to greater than 1 standard deviation positive news reports and increase (up to 2.8 bps on average) in the hours prior to less than 1 standard deviation negative news reports.

Table 1.3 shows that the pre-announcement conditional drift in stocks and bonds is statistically and economically significant. Panel A runs regressions of pre-announcement and announcement returns on the news report variable NR_t :

$$R_t = \alpha + \beta NR_t + \epsilon_t.$$

In row $-x$, the left-hand-side variable R_t is the cumulative pre-announcement return (in bps) from x minutes to 5 minutes before MNAs. In row 0, R_t is the ± 5 minute announcement return (in bps) around MNAs. The coefficient β measures the return sensitivity to a 1 standard deviation news report. For both assets, $\hat{\beta}$ in row 0 is statistically significant: stocks increase {decrease} 6.7 bps more and bonds decrease {increase} 5.3 bps more in the 10 minutes around announcements of 1 standard deviation positive {negative} news reports. For rows $-x$, the majority of the $\hat{\beta}$ for stocks and all the $\hat{\beta}$ for bonds are statistically significant: stocks increase {decrease} up to 2.3 bps more and bonds decrease {increase} up to 1.5 bps more before announcements of 1 standard deviation positive {negative} news reports. The ratio of the pre-announcement $\hat{\beta}$ to the sum of the pre-announcement $\hat{\beta}$ and announcement $\hat{\beta}$ demonstrates the economic magnitude of the results. For stocks and bonds, up to 26% and 23%, respectively, of the impact of macroeconomic news reports takes place in the pre-announcement window.

In Panel B, I evaluate how news reports impact returns in a different regression specification:

$$R_t = \alpha + \beta^+ D_t^+ + \beta^- D_t^- + \epsilon_t.$$

Table 1.3: Regressions of the Pre-Announcement Conditional Drift in Stocks and Bonds

Panel A: Regressions of the Form $R_t = \alpha + \beta NR_t + \epsilon_t$

	Stocks			Bonds		
	$\hat{\beta}$ (bps)	$\hat{\alpha}$ (bps)	R_{adj}^2 (%)	$\hat{\beta}$ (bps)	$\hat{\alpha}$ (bps)	R_{adj}^2 (%)
-720	1.8	4.5***	0.07	-1.4***	-0.5	0.50
-660	1.5	4.3***	0.04	-1.5***	-0.7	0.60
-600	1.4	4.3***	0.04	-1.5***	-0.9**	0.58
-540	1.5	4.7***	0.05	-1.5***	-1.0**	0.63
-480	2.0*	4.5***	0.13	-1.5***	-1.2***	0.67
-420	2.3**	3.6***	0.18	-1.5***	-1.0***	0.71
-360	1.9*	2.8***	0.12	-1.5***	-0.8**	0.74
-300	2.3***	1.7*	0.26	-1.5***	-0.7**	0.84
-240	1.9**	1.1	0.19	-1.3***	-0.8**	0.70
-180	1.8**	0.5	0.24	-1.2***	-0.8**	0.75
-120	1.0	0.9	0.06	-1.1***	-0.7***	0.75
-60	0.8*	0.3	0.08	-0.7***	-0.7***	0.60
0	6.7***	1.0*	6.34	-5.3***	0.3	7.66

Panel B: Regressions of the Form $R_t = \alpha + \beta^+ D_t^+ + \beta^- D_t^- + \epsilon_t$

	Stocks				Bonds			
	$\hat{\beta}^+$ (bps)	$\hat{\beta}^-$ (bps)	$\hat{\alpha}$ (bps)	R_{adj}^2 (%)	$\hat{\beta}^+$ (bps)	$\hat{\beta}^-$ (bps)	$\hat{\alpha}$ (bps)	R_{adj}^2 (%)
-720	6.6**	-1.5	3.9***	0.12	-2.6**	3.1**	-0.6	0.47
-660	5.4*	-1.3	3.7***	0.05	-2.6**	3.3***	-0.8	0.54
-600	5.8*	-0.9	3.7***	0.08	-2.8**	3.0**	-0.9*	0.54
-540	6.4**	-1.1	4.0***	0.13	-2.7**	3.4***	-1.1**	0.61
-480	7.1**	-2.1	3.9***	0.20	-2.8**	3.3***	-1.3***	0.64
-420	7.8**	-2.7	3.0**	0.28	-2.8**	3.1***	-1.1**	0.63
-360	5.1*	-3.1	2.6**	0.13	-2.0*	3.8***	-1.0**	0.77
-300	7.2***	-2.8	1.2	0.32	-2.6**	3.3***	-0.8**	0.81
-240	6.1**	-2.0	0.6	0.23	-2.1**	2.9***	-0.9**	0.67
-180	4.9**	-2.4	0.2	0.23	-2.5***	2.1**	-0.7**	0.60
-120	3.0	-1.2	0.7	0.05	-1.5*	2.4***	-0.8***	0.56
-60	3.3**	1.0	-0.3	0.13	-1.3**	1.0*	-0.7***	0.35
0	6.0***	-18.2***	2.7***	6.38	-8.2***	11.6***	-0.2	7.08

Notes: The table shows results from regressions of pre-announcement and announcement stock (front ES futures contract) and bond (front TY futures contract) returns on the news report variable NR_t . Regressions are run on the full sample of MNAs, and $n = 2179$. Panels A and B represent different regression specifications: $R_t = \alpha + \beta NR_t + \epsilon_t$, and $R_t = \alpha + \beta^+ D_t^+ + \beta^- D_t^- + \epsilon_t$, respectively. In row $-x$, the left-hand-side variable R_t is the cumulative pre-announcement return (in bps) from x minutes to 5 minutes before MNAs. In row 0, the left-hand-side variable R_t is the ± 5 minute announcement return (in bps) around MNAs. NR_t measures the informational content of MNAs and is defined in Figure 1.1. D_t^+ and D_t^- are dummy variables equal to 1 if $NR_t > 1$ and $NR_t < -1$, respectively, and zero otherwise. t -statistics (not shown) are based on heteroskedasticity-consistent standard errors. *** denotes significance at the 1% level, ** denotes significance at the 5% level, and * denotes significance at the 10% level.

D_t^+ and D_t^- are dummy variables equal to 1 if $NR_t > 1$ and $NR_t < -1$, respectively, and R_t is as defined above. The coefficients β^+ and β^- measure the return sensitivity to large magnitude news reports that are positive and negative, respectively. For both assets, $\hat{\beta}^+$ and $\hat{\beta}^-$ in row 0 are statistically significant: stocks increase 6.0 bps {decrease 18.2 bps} more and bonds decrease 8.2 bps {increase 11.6 bps} more for announcements of greater than 1 standard deviation positive {negative} news reports.¹⁴ For rows $-x$, the majority of the $\hat{\beta}^+$ but none of the $\hat{\beta}^-$ for stocks and all the $\hat{\beta}^+$ and $\hat{\beta}^-$ for bonds are statistically significant: stocks increase up to 7.8 bps {decrease up to 3.1 bps} more and bonds decrease up to 2.8 bps {increase up to 3.8 bps} more before announcements of greater than 1 standard deviation positive {negative} news reports.¹⁵ For each of $\hat{\beta}^+$ and $\hat{\beta}^-$, I compute the ratio of the pre-announcement coefficient to the sum of the pre-announcement and announcement coefficients. In the stock market, up to 57% {15%} of the impact of large, positive {negative} macroeconomic news reports takes place in the pre-announcement window. In the bond market, the analogous number is 26% {25%}.¹⁶

Insofar as the results in Figure 1.3 and Tables 1.2 and 1.3 group all types of MNAs together, the magnitude of sophisticated trading is likely to be underestimated. The reason is that the intensity of sophisticated trading is likely to be greater around some types of announcements than others. For example, traders may only be able to predict macroeconomic news reports for some types of announcements. Jointly analyzing announcements with and without sophisticated trading thus downward biases the magnitude of the pre-announcement conditional drift.

¹⁴Stocks react more strongly to negative macroeconomic news reports than to positive ones. Zhou (2015) discusses this asymmetry. Bonds appear to react more similarly to positive and negative news reports.

¹⁵It is unclear the reason for the stronger pre-announcement reaction of stocks to positive news reports than negative news reports.

¹⁶Though I do not focus on the $\hat{\alpha}$ results in Table 1.3, the regression intercepts have economic content. For stocks in Panels A and B, the row 0 $\hat{\alpha}$ is positive and statistically significant. This is the announcement premium for stocks around MNAs referenced in Fn. (11). The row $-x$ $\hat{\alpha}$ is also positive and generally statistically significant. This is the pre-announcement drift of stocks before MNAs referenced in Fn. (12). For bonds, the row 0 $\hat{\alpha}$ is statistically insignificant, but the row $-x$ $\hat{\alpha}$ is negative and generally statistically significant. As discussed in Fn. (13), there is no announcement premium for bonds around MNAs, but there is a pre-announcement drift of bonds before MNAs.

The pre-announcement conditional drift is difficult to rationalize with explanations other than that of sophisticated trading. One alternative explanation, for example, is that of underreaction. Suppose that positive {negative} macroeconomic news reports are preceded by public announcements that elicit positive {negative} stock returns and negative {positive} bond returns. If stocks and bonds underreact to such public announcements, the pre-announcement conditional drift could result from the gradual correction of this underreaction. Such an explanation is improbable: MNAs in my sample occur from 8:30 AM to 10:00 AM ET, and the conditional drift begins in the hours before during overnight trading in the U.S. It is unlikely that important public announcements occur at those times.

1.3.2 Order Imbalances and Sophisticated Trading Profits

Analysis of order imbalances provides further support for sophisticated trading around MNAs. In Appendix A.2, I use the bulk volume classification algorithm of Easley, López de Prado, and O'Hara (2013) to impute order flow and split trade volume into buys and sells.¹⁷ On announcements of positive {negative} macroeconomic news reports, there are order imbalances of buys {sells} in stocks and sells {buys} in bonds. In the hours before announcements, there are also imbalances of buys {sells} in stocks and sells {buys} in bonds. This pre-announcement conditional pattern in the order imbalance is analogous to the pre-announcement conditional drift.

The data on order imbalances make possible a back of the envelope calculation that reveals the magnitude of sophisticated trading profits. Consider the case of stocks first. As indicated in Table 1.2, there are 291 announcements with greater than 1 standard deviation positive news reports and 286 announcements with less than 1 standard deviation negative news reports. Thus, there are about 50 such announcements per year from July 2003 to March 2014. Based on Panel B of Table 1.3, a greater than 1 standard deviation magnitude news report is associated with a combined pre-announcement and announcement return

¹⁷Based on papers such as Easley, López de Prado, and O'Hara (2013) and Panayides, Shohfi, and Smith (2014), the bulk volume classification approach often outperforms other trade classification algorithms such as that of Lee and Ready (1991). See Appendix A.2 for details.

on the order of 10 bps in magnitude. From Figure 1.2, there is an average of about \$10 billion in dollar volume traded by the time of the first announcements at 8:30 AM ET. Finally, Appendix A.2 shows that a greater than 1 standard deviation magnitude news report leads to a change in the order imbalance on the order of 50 bps in magnitude. There is thus $(\$10 \text{ billion}) \times (50 \text{ bps}) = \50 million of excess buy or sell dollar volume. This value represents a conservative lower bound of capital used for sophisticated trading. Per announcement, the profits to sophisticated trading are estimated to be $(\$50 \text{ million}) \times (10 \text{ bps}) = \$50,000$. Over the course of a year, the profits are around $\$50,000 \times 50 = \2.5 million . A calculation using the same numbers applies to bonds.

The overall profits to sophisticated trading in stocks and bonds are higher than indicated by the above calculation. Two reasons are that I only consider announcements with greater than 1 standard deviation magnitude news reports and that I only attribute the order imbalance portion of dollar volume to sophisticated traders' capital. A reasonable estimate of the profits to sophisticated trading in stock and bond futures is on the order of \$10 million a year in each security. Sophisticated traders likely also profit from trading other securities and assets.

1.4 Source of Informational Advantage

The previous section's results lead to a natural follow-up question: what is the source of sophisticated traders' informational advantage? The first subsection below discusses (illegal) information leakage as one possible source. The second subsection considers the homework done by skilled traders as another possible source. The empirical evidence supports the latter channel of information flows.

1.4.1 (Illegal) Information Leakage Channel

Regulators have expressed concerns that some market participants may obtain an informational advantage about upcoming macroeconomic news reports through an (illegal)

information leakage channel.¹⁸ One concern of regulators centers on the role of the media. For many types of MNAs, data are initially released to media organizations during lockup periods prior to public releases. While locked up, these organizations are embargoed from communicating with the outside. Yet there have been accounts of media circumventing the lockup process both unintentionally and intentionally. Another concern of regulators is more standard and centers on the role of insiders working at data-releasing institutions who leak macroeconomic data. Security procedures around the release of macroeconomic data have been scrutinized by the SEC, the Federal Bureau of Investigation, the U.S. Commodity Futures Trading Commission, and even Sandia National Laboratories, which is responsible for the safety of U.S. nuclear weapons.

Thus, it is plausible that those who I call sophisticated traders actually obtain their informational edge via an (illegal) information leakage channel. While there is concrete evidence of idiosyncratic leaks, however, there is little evidence of systematic leaks. I also do not find evidence of systematic leaks that can explain the pre-announcement conditional drift. I do find that information leakage by the media during lockups cannot account for all of the sophisticated trading around MNAs. Based on results from Section 1.3, stocks and bonds begin conditionally drifting at least hours prior to announcements, while media lockups begin only 30 minutes to 60 minutes prior to announcements. Relatedly, the pre-announcement conditional drift cannot be explained by certain media organizations' business strategy of releasing macroeconomic data early to paying customers. This pay to play practice takes place anywhere from a few seconds to a few minutes before official release times and thus occurs within the ± 5 minute announcement window, not the pre-announcement period.

¹⁸See, for example, the following press reports in order of publication date: the March 7, 2012 CNBC article "Labor Dept. Asks Nuclear Guardians for Help Keeping Jobs Data Secret"; the April 10, 2013 *The Wall Street Journal* article "A Probe on Data Releases is Revived"; the August 5, 2013 *The Wall Street Journal* article "FBI Finds Holes in System Protecting Economic Data"; the January, 2, 2014 *The Wall Street Journal* article "Labor Department Panel Calls for Ending Lockup for Jobs Data"; and the December 1, 2014 Bloomberg News article "Fed Leak Handed Traders Profitable Tip, Prompted Secret Inquiry."

1.4.2 Homework Channel

Though absence of evidence does not rule out (illegal) information leakage, I show that sophisticated traders can obtain their informational advantage via a more innocent mechanism: the homework channel. It is not altogether surprising that skilled traders are able to do their homework and use public information to predict macroeconomic news reports. These traders have a strong incentive: money. They moreover have the necessary resources to predict news reports: traders can combine an array of predictor variables including proprietary data with cutting-edge forecasting algorithms.¹⁹

Romer and Romer (2000) in effect appeal to the homework channel in order to explain their finding that the Federal Reserve is able to produce better forecasts about the economy than private-sector economists. By committing resources to the forecasting process, the central bank is able to do so using only “publicly available information.” Given that there is a well-trodden path from the Federal Reserve to the financial services industry, perhaps the sophisticated traders analyzed in my paper even get homework help from former central bank officials or are former central bank officials themselves.

Section 1.4.2.1 provides an example of how skilled traders might do their homework and use public information to predict macroeconomic news reports. Section 1.4.2.2 shows that the predictable component of news reports is able to explain a significant portion of the pre-announcement conditional drift.

1.4.2.1 Doing Homework: Public Information Predicts Macroeconomic News Reports

To illustrate the broader idea that sophisticated traders can use public information to predict macroeconomic news reports, I show a simple example. Traders are able to do their homework and predict the news report of an upcoming MNA using the difference between economists’ forecasts of and past realizations of the data to be released.

¹⁹PriceStats and State Street provide examples of such proprietary data. PriceStats uses online prices to track inflation daily in multiple countries including the U.S. Per its official website, the State Street Investor Confidence Index “measures investor confidence or risk appetite quantitatively by analyzing the actual buying and selling patterns of institutional investors.” These proprietary data are potentially useful for forecasting news reports of CPI, Consumer Confidence, and other types of MNAs.

The success of this method rests on the systematic bias of economists' forecasts of macroeconomic data. Utilizing a different sample of MNAs, Campbell and Sharpe (2009) find that forecasts are "anchored toward the recent past values of the series being forecasted."²⁰ Viewed through the lens of Tversky and Kahneman (1974) and related studies, this anchoring phenomenon manifests a behavioral bias: the human tendency to overweight a reference point or anchor when making decisions. Alternatively, economists may strategically anchor their macroeconomic forecasts as an optimal response to incentives.

If economists' forecasts also display anchoring bias in my sample of MNAs, it is straightforward to use public information to predict news reports. To see this, consider the following unstandardized version of the news report variable defined in Eq. (1.1):

$$\tilde{NR}_t = \hat{\sigma}NR_t = A_t - F_t. \quad (1.2)$$

As before, A_t is the announced value of the data release, and F_t is the median forecasted value of the data release from a survey of economists. If there is anchoring bias, the forecast F_t puts weight $\lambda < 1$ on the unbiased predicted value of the data release $\mathbb{E}[A_t]$ and weight $1 - \lambda > 0$ on some anchor \bar{A}_h :

$$F_t = \lambda \mathbb{E}[A_t] + (1 - \lambda) \bar{A}_h. \quad (1.3)$$

I define \bar{A}_h as the historical average of the past h actual data releases of the economic indicator being forecasted. From the definition of the unstandardized news report variable in Eq. (1.2), $\mathbb{E}[A_t] = \mathbb{E}[\tilde{NR}_t] + F_t$. Plugging this expression into that of the biased forecast in Eq. (1.3) yields $\mathbb{E}[\tilde{NR}_t] = \gamma (F_t - \bar{A}_h)$, with

$$\gamma = \frac{1 - \lambda}{\lambda}. \quad (1.4)$$

When there is anchoring bias, $\lambda < 1 \implies \gamma > 0$. In this case, $F_t - \bar{A}_h$ positively predicts

²⁰Their sample runs from 1991 to 2006, and they collect economists' forecasts from surveys by Money Market Services.

$\tilde{N}R_t$ in the regression

$$\tilde{N}R_t = \kappa + \gamma (F_t - \bar{A}_h) + \epsilon_t. \quad (1.5)$$

Since both the forecast F_t and the anchor \bar{A}_h are in the public domain before a MNA, public information predicts the news report variable. The intuition here is that if the forecast is higher than the anchor, it is not high enough due to overweighting of the lower value of the anchor, and conversely if the forecast is lower than the anchor.

Whether there is anchoring bias in my sample is an empirical question. In each row of Table 1.4, I run the regression in Eq. (1.5) for each type of MNA and for the two specifications $h = 1$ (anchor is the last actual data release) and $h = 3$ (anchor is the average of the last 3 actual data releases). To avoid look-ahead bias, I estimate the regression in a dynamic manner over a changing window.²¹ The coefficient of interest is γ . For each type of MNA, the output is two series of $\hat{\gamma}$, which correspond to the $h = 1$ and $h = 3$ specifications. As indicated in the “Anchor” column, Table 1.4 displays the results corresponding to the anchor specification that more frequently has a higher adjusted R^2 . Table 1.4 shows the mean, standard deviation, maximum, and minimum of $\hat{\gamma}$. In most cases, the mean of $\hat{\gamma}$ is positive and large relative to the standard deviation. Economists’ forecasts exhibit anchoring bias. Equivalently and more importantly, differences between economists’ forecasts and past data realizations predict macroeconomic news reports. The results in Table 1.4 are economically significant as well. Using Eq. (1.4) to translate between $\hat{\gamma}$ and the implied λ , I find that forecasts are substantially weighted to historical data realizations. For example, with a mean $\hat{\gamma}$ of 0.69 in the case of the Consumer Confidence data series with $h = 1$, the implied $\lambda = 0.59$. On average, forecasts place 59% weight on the previous month’s data and only 41% weight on the unbiased predicted value of the upcoming data release. As shown in Appendix A.3, the anchoring bias results are robust to a battery of tests. In particular, the results hold when I estimate Eq. (1.5) using different anchors, different windows, and

²¹My sample begins in 2003 for 16 out of 18 types of MNAs and 2005 for the other 2 types of MNAs (see Table 1.1). For the portion of my sample beginning in 2003, data are available extending back to 1997, 1998, or 1999. I make use of this data to predict news reports starting from 2003 and thus do not lose any observations. For the portion of my sample beginning in 2005, I use the first 4 years to predict news reports starting from 2009 and thus lose 4 years of observations.

Table 1.4: *Predicting Macroeconomic News Reports*

Event Name	Anchor	Mean $\hat{\gamma}$	Std. Dev. $\hat{\gamma}$	Max. $\hat{\gamma}$	Min. $\hat{\gamma}$
Change in Nonfarm Payrolls	1	-0.12	0.01	-0.09	-0.16
Consumer Confidence	1	0.69	0.06	0.87	0.57
CPI	3	0.17	0.02	0.24	0.14
Durable Orders	3	0.41	0.03	0.48	0.35
Existing Home Sales	1	0.31	0.08	0.54	0.13
Factory Orders	3	0.05	0.02	0.08	0.03
Housing Starts	1	0.31	0.05	0.38	0.16
Industrial Production	3	0.31	0.06	0.47	0.20
Initial Jobless Claims	1	-0.12	0.02	-0.08	-0.16
ISM Manufacturing	3	0.17	0.07	0.37	0.09
ISM Non-Manufacturing	1	0.47	0.08	0.61	0.34
Leading Index	3	0.21	0.02	0.27	0.18
New Home Sales	1	0.40	0.13	0.88	0.29
Pending Home Sales	3	0.28	0.09	0.46	-0.08
Personal Income	3	0.04	0.05	0.17	0.00
PPI	3	0.33	0.03	0.42	0.28
Retail Sales	1	0.18	0.02	0.22	0.16
Unemployment Rate	3	-0.01	0.09	0.11	-0.17

Notes: For each type of MNA, the corresponding row runs the regression $\tilde{NR}_t = \kappa + \gamma (F_t - \bar{A}_h) + \epsilon_t$. The unstandardized news report variable $\tilde{NR}_t = \hat{\sigma} NR_t$ measures the informational content of a given data release, with $\hat{\sigma}$ and NR_t both defined in Figure 1.1. For the same data release, F_t is the median forecasted value from a survey of economists, and \bar{A}_h is the historical average of the past h realizations. Regressions are run for the specifications $h = 1$ and $h = 3$. To avoid look-ahead bias, I estimate the regression in a dynamic manner over a changing window (see Fn. (21) for details). For each type of MNA, the output is two series of $\hat{\gamma}$, which correspond to the $h = 1$ and $h = 3$ specifications. As indicated in the “Anchor” column, the table displays the results corresponding to the specification that more frequently has a higher adjusted R^2 .

grouping all types of MNAs together.

1.4.2.2 Trading on the Homework: Explaining the Pre-Announcement Conditional Drift

The previous subsection shows that market participants can do their homework to predict macroeconomic news reports. This subsection goes a step further and provides evidence that market participants trade on this homework, which can explain the documented asset price patterns around MNAs. To do so, I first use the results of Table 1.4 to split the news report variable NR_t into a predictable component $NR_{p,t}$ and an unpredictable component

$NR_{u,t}$:

$$NR_t = NR_{p,t} + NR_{u,t}.$$

For each type of MNA, $NR_{p,t}$ is the one-step-ahead forecast of the corresponding regression in Table 1.4 divided by the $\hat{\sigma}$ defined in Eq. (1.1). $NR_{u,t}$ is the regression residual divided by $\hat{\sigma}$.

Table 1.5 shows that the predictable component of news reports explains much of the pre-announcement conditional drift in stocks and bonds. Consider the results corresponding to rows $-x$. Each row runs regressions of pre-announcement returns on the predictable and unpredictable components of news reports:

$$R_t = \alpha + \beta^p NR_{p,t} + \beta^u NR_{u,t} + \epsilon_t. \quad (1.6)$$

As in previous tables, R_t is the cumulative pre-announcement return (in bps) from x minutes to 5 minutes before MNAs. Focusing on $\hat{\beta}^p$, all the estimated coefficients are positive for stocks (with some statistically significant) and negative for bonds (with all statistically significant). Stocks increase {decrease} and bonds decrease {increase} in the hours before news reports predicted to be positive {negative}. To quantify the results, I multiply the standard deviation of $NR_{p,t}$ (0.29) with $\hat{\beta}^p$ to find that before a 1 standard deviation positive {negative} value of the predictable component of the news report, stocks increase {decrease} up to 1.7 bps and bonds decrease {increase} up to 1.8 bps. In Appendix A.4, I construct the predictable and unpredictable components of news reports using alternative estimates of Eq. (1.5). These different constructions of the predictable component continue to explain the pre-announcement conditional drift in stocks and bonds via Eq. (1.6).

The evidence is consistent with the interpretation that sophisticated traders skillfully predict macroeconomic news reports with public information and then trade on these predictions. In doing so, these traders contribute to the pre-announcement conditional drift. Of course, the results are based on a specific prediction method that relies on anchoring bias. Practitioners probably have more advanced ways of predicting news reports using more comprehensive public information. This likelihood can explain why $\hat{\beta}^u$ in Table

Table 1.5: Regressions of Stock and Bond Returns on the Predictable and Unpredictable Components of Macroeconomic News Reports

	Stocks				Bonds			
	$\hat{\beta}^p$ (bps)	$\hat{\beta}^u$ (bps)	$\hat{\alpha}$ (bps)	R_{adj}^2 (%)	$\hat{\beta}^p$ (bps)	$\hat{\beta}^u$ (bps)	$\hat{\alpha}$ (bps)	R_{adj}^2 (%)
-720	4.7	1.6	4.3***	0.06	-6.0***	-1.2***	-0.3	0.95
-660	4.2	1.4	4.1***	0.03	-5.7***	-1.3***	-0.5	1.00
-600	3.5	1.3	4.2***	0.01	-5.6***	-1.3***	-0.7	0.98
-540	2.7	1.2	4.5***	-0.01	-4.9***	-1.3***	-0.8*	0.88
-480	5.0	1.5	4.1***	0.07	-5.1***	-1.3***	-1.0**	0.94
-420	5.5	1.7	3.2***	0.12	-5.4***	-1.2***	-0.8**	1.07
-360	4.1	1.4	2.4**	0.05	-5.0***	-1.2***	-0.5	1.09
-300	5.2*	1.7*	1.3	0.17	-4.5***	-1.3***	-0.6	1.11
-240	3.7	1.3	0.9	0.08	-3.7***	-1.1***	-0.7**	0.86
-180	5.7**	1.4*	0.2	0.26	-4.1***	-1.0***	-0.6**	1.07
-120	3.6	1.0	1.0	0.11	-3.3***	-1.0***	-0.6**	1.03
-60	3.0*	0.3	0.1	0.06	-2.0***	-0.6***	-0.7***	0.68
0	5.0***	6.4***	0.9	5.81	-2.3	-5.5***	0.2	7.73

Notes: The table shows results from regressions of pre-announcement and announcement stock (front ES futures contract) and bond (front TY futures contract) returns on the predictable and unpredictable components of news reports $NR_{p,t}$ and $NR_{u,t}$, respectively: $R_t = \alpha + \beta^p NR_{p,t} + \beta^u NR_{u,t} + \epsilon_t$. Regressions are run on the full sample of MNAs, and $n = 2108$. In row $-x$, the left-hand-side variable R_t is the cumulative pre-announcement return (in bps) from x minutes to 5 minutes before MNAs. In row 0, the left-hand-side variable R_t is the ± 5 minute announcement return (in bps) around MNAs. For each type of MNA, $NR_{p,t}$ is the one-step-ahead forecast of the corresponding regression in Table 1.4 divided by the $\hat{\sigma}$ defined in Figure 1.1. $NR_{u,t}$ is the regression residual divided by $\hat{\sigma}$. t -statistics (not shown) are based on heteroskedasticity-consistent standard errors. *** denotes significance at the 1% level, ** denotes significance at the 5% level, and * denotes significance at the 10% level.

1.5 is positive for stocks (and sometimes statistically significant) and negative for bonds (and always statistically significant). Based on the results, stocks and bonds conditionally drift before the release of the unpredictable component of news reports. What I call the unpredictable component of news reports, however, is only unpredictable when news reports are predicted with a technique utilizing anchoring bias. The unpredictable component of news reports may actually be predictable if news reports are predicted using superior methods that practitioners likely employ. Alternatively, if the unpredictable component of news reports is truly unpredictable with public information, $\hat{\beta}^u$ could be picking up signs of (illegal) information leakage. The bottom line is that even though using anchoring bias to predict news reports is simple, this method nonetheless makes successful predictions and is able to account for a significant portion of the pre-announcement conditional drift.

1.5 Market Efficiency around MNAs

Based on the above results, stock and bond markets are inefficient around MNAs. In the two subsections below, I elaborate on how markets are inefficient over time and across assets, respectively.

1.5.1 Over Time

Over time, there is a late trading puzzle: the public information that predicts macroeconomic news reports is impounded into stocks and bonds with a significant lag. Section 1.4, supplemented by Appendix A.3, argues that sophisticated traders' informational advantage arises from the homework channel. Traders are able to do their homework days before MNAs: most types of public information predictive of news reports are available at that time. For example, predicting news reports based on anchoring bias requires only economists' forecasts and historical data realizations. Sections 1.3 and 1.4 and Appendix A.4, however, provide evidence of sophisticated trading hours before MNAs. At least some of this trading appears to be on the predictable component of news reports obtainable from public information. Contrary to the logic of market efficiency, the public information that predicts

news reports days before announcements is impounded into stocks and bonds just hours before announcements.

1.5.2 Across Assets

Across assets, there is a related puzzle: by the official release times of MNAs, the public information that predicts news reports is fully incorporated into bonds but only partially incorporated into stocks. As a result, the bond market eventually becomes efficient, while the stock market remains inefficient. Table 1.5 shows that the predictable component of news reports based on anchoring bias is completely impounded into bonds but only partially impounded into stocks. Stocks but not bonds react to announcements of the predictable component. Row 0 of Table 1.5 runs regressions of the ± 5 minute stock or bond announcement return around MNAs (in bps) on the predictable and unpredictable components of news reports:

$$R_t = \alpha + \beta^p NR_{p,t} + \beta^u NR_{u,t} + \epsilon_t. \quad (1.7)$$

$\hat{\beta}^p$ is positive and statistically significant for stocks but smaller in magnitude and statistically insignificant for bonds. Multiplying the standard deviation of $NR_{p,t}$ (0.29) with $\hat{\beta}^p$, I find that, on the announcement of a 1 standard deviation positive {negative} value of the predictable component of the news report, stocks increase {decrease} 1.5 bps while bonds decrease {increase} 0.7 bp.²² In Appendix A.4, I run regressions based on Eq. (1.7) using alternative constructions of $NR_{p,t}$ and $NR_{u,t}$ and find similar results. The two components are derived from varying the specifications of the anchoring bias regression in Eq. (1.5).

Table 1.6 provides additional evidence that bonds are more efficient than stocks: stocks but not bonds display return momentum between pre-announcement and announcement returns. The table runs regressions of return momentum in both assets:

²²Note that, as expected, stocks and bonds react to announcements of the unpredictable component of news reports: $\hat{\beta}^u$ is positive for stocks, negative for bonds, and statistically significant for both assets.

Table 1.6: Regressions of Return Momentum in Stocks and Bonds

	Stocks			Bonds		
	$\hat{\beta}$ (bps)	$\hat{\alpha}$ (bps)	R_{adj}^2 (%)	$\hat{\beta}$ (bps)	$\hat{\alpha}$ (bps)	R_{adj}^2 (%)
-720	0.03**	1.06*	0.26	0.00	-0.04	-0.06
-660	0.03**	1.05*	0.30	0.01	-0.03	-0.05
-600	0.02**	1.07*	0.20	0.00	-0.03	-0.06
-540	0.03**	1.04*	0.28	0.01	-0.03	-0.04
-480	0.03***	1.02*	0.42	0.02	-0.02	-0.02
-420	0.05***	1.00*	0.94	0.02	-0.02	-0.02
-360	0.06***	1.00*	1.34	0.03	-0.02	0.01
-300	0.06***	1.06*	1.10	0.03	-0.02	0.05
-240	0.06***	1.11*	0.94	0.03	-0.02	0.00
-180	0.05***	1.17*	0.38	0.03	-0.02	-0.01
-120	0.02	1.17*	-0.01	0.03	-0.02	0.01
-60	0.05**	1.17*	0.18	0.01	-0.03	-0.05

Notes: The table shows results from regressions of announcement returns on pre-announcement returns for stocks (front ES futures contract) and bonds (front TY futures contract): $evtR_t = \alpha + \beta preR_t + \epsilon_t$. Regressions are run on the full sample of MNAs, and $n = 1705$. The left-hand-side variable $evtR_t$ is the ± 5 minute announcement return (in bps) around MNAs. The right-hand-side variable $preR_t$ in row $-x$ is the cumulative pre-announcement return (in bps) from x minutes to 5 minutes before MNAs. t -statistics (not shown) are based on heteroskedasticity-consistent standard errors. *** denotes significance at the 1% level, ** denotes significance at the 5% level, and * denotes significance at the 10% level.

$$evtR_t = \alpha + \beta preR_t + \epsilon_t.$$

The left-hand-side variable $evtR_t$ is the ± 5 minute announcement return (in bps) around MNAs. The right-hand-side variable $preR_t$ in row $-x$ is the cumulative pre-announcement return (in bps) from x minutes to 5 minutes before MNAs. $\hat{\beta}$ is positive and statistically significant for stocks. A 10 bps higher pre-announcement return predicts up to 0.6 bp higher announcement return. In contrast, the estimated coefficient is smaller in magnitude and statistically insignificant for bonds.

1.6 Model

To understand why stock and bond markets are inefficient around MNAs, I construct a model. In the model, all traders see a public announcement of a news report, which can be thought of as a MNA, that provides information about the values of risky assets, which can be thought of as stocks and bonds. Consistent with empirical observations, a minority of traders in the model are sophisticated and obtain advance knowledge of the news report, possibly through the skillful processing of public information.

The model has three key assumptions that are justified by the literature. The first assumption of the model is that the majority of traders are unsophisticated and do not have advance knowledge of the news report. Unsophisticated traders neglect public information that predicts the news report. In the parlance of the inattention literature, we can think of unsophisticated traders as being inattentive to such public information; in contrast, sophisticated traders are attentive. Two related examples of the inattention literature are Tetlock (2011) and Gilbert, Kogan, Lochstoer, and Ozyildirim (2012), which show that asset prices react to stale information. The difference in attention between the two types of traders may reflect a difference in incentives and resources. For example, suppose that sophisticated traders work at funds that employ trading strategies based on predicting news reports. These traders then have the incentives (such as money) and institutional resources to pay close attention to public information potentially useful for making predictions. Other, unsophisticated traders may lack such incentives and resources.

The second assumption of the model is that sophisticated traders have risk management concerns and dislike risk. Risk impedes sophisticated traders from acting on their informational advantage and, consequently, impedes the extent to which their informational advantage is incorporated into asset prices. I model risk management concerns with risk aversion, which is sufficient but not necessary: some other limit to arbitrage could work just as well. For example, sophisticated traders could be risk-neutral but capital constrained as in Shleifer and Vishny (1997).

The third assumption of the model is that traders do not extract each others' information

from asset prices. In particular, no trader, sophisticated or unsophisticated, extracts from pre-announcement prices the advance knowledge of (other) sophisticated traders. As a result, there is no strategic front-running that could unravel the equilibrium. Other papers in the literature have similar assumptions. Each newswatcher in Hong and Stein (1999), for example, obtains information but fails to infer other newswatchers' information from prices. Eyster, Rabin, and Vayanos (2015) model the equilibrium outcome in a financial market in which traders are differentially informed and are cursed in the sense of failing to fully account for what prices signal about each others' information. In a laboratory experiment, Carrillo and Palfrey (2011) show evidence that traders neglect asymmetric information.

The first subsection below outlines the baseline model with one risky asset. The second subsection extends the model to include two risky assets.

1.6.1 Model with One Risky Asset

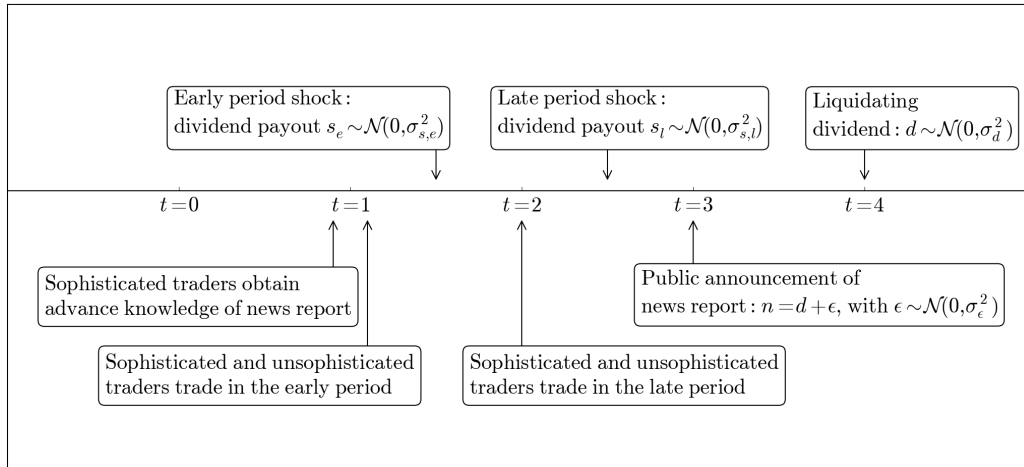
The baseline model has one risky asset that corresponds to stocks or bonds and one risk-free asset. By assumption, the risk-free rate is zero, and the risky asset is in zero net supply. Agents in the model are split into two types of traders, with measure i sophisticated and measure $1 - i$ unsophisticated. I assume that there is no time discounting and that all traders have constant absolute risk aversion (CARA) utility. The coefficients of absolute risk aversion for sophisticated and unsophisticated traders are A_S and A_U , respectively. Given the assumptions, both types of traders seek the maximum risk-adjusted return.

Panel A of Figure 1.4 shows the timeline of the model. There are five dates: $t = 0, 1, 2, 3,$ and 4 . At $t = 4$, the end of the model, the risky asset pays a liquidating dividend $d \sim \mathcal{N}(0, \sigma_d^2)$. At $t = 3$, there is a public announcement of a news report that is seen by all traders and that is interpreted as a noisy signal of d :

$$n = d + \epsilon,$$

with $\epsilon \sim \mathcal{N}(0, \sigma_\epsilon^2)$ and independent from d . Unsophisticated traders see the news report for the first time, but sophisticated traders already know the news report. Specifically, at

Panel A: Model with One Risky Asset



Panel B: Model with Two Risky Assets




Figure 1.4: *Timelines of Events in the Model*

Notes: Panels A and B plot the timelines of events in the model with one risky asset and two risky assets, respectively.

$t = 1$, sophisticated traders obtain advance knowledge of the news report and infer n .²³ Immediately after, both types of traders have the first opportunity to trade in the “early period” before the public announcement. At $t = 2$, the sophisticated traders still have an informational advantage over the unsophisticated traders. Both types of traders have the second opportunity to trade in the “late period” before the public announcement. Trading in both periods is risky due to information flows about the fundamental value of the risky asset. To model this risk, I assume that between $t = 1$ and $t = 2$, there is a shock to holding the risky asset in the form of dividend payout $s_e \sim \mathcal{N}(0, \sigma_{s,e}^2)$. Similarly, in between $t = 2$ and $t = 3$, there is another shock in the form of dividend payout $s_l \sim \mathcal{N}(0, \sigma_{s,l}^2)$. By assumption, s_e and s_l are independent from each other and from ϵ and d .

Sections 1.6.1.1 and 1.6.1.2 below describe the average *cum*-dividend price of the risky asset in two specifications of the model. I first consider the efficient-market specification in which all traders are sophisticated; that is, $i = 1$. I then consider the inefficient-market specification in which some traders are unsophisticated; that is, $i \in (0, 1)$. Without loss of generality, the discussion below takes the case of a news report that is good for the risky asset: $n > 0$. Appendix A.5 provides details including derivations.


1.6.1.1 Efficient-Market Specification: All Traders are Sophisticated, and $i = 1$

The solid $i = 1$ line in Panel A of Figure 1.5 plots the average *cum*-dividend price of the risky asset when all traders are sophisticated. The price of the risky asset starts at 0. All traders obtain advance knowledge of the news report at $t = 1$, so the price increases to the rational expectation of the liquidating dividend. The price remains at this efficient-market level at $t = 2, 3$, and 4:

$$p_1 = p_2 = p_3 = p_4 = \alpha n, \tag{1.8}$$

²³I assume that sophisticated traders have perfect foresight of the public announcement. The source of the informational advantage is left unmodeled and could be due to the homework channel, as the evidence in this paper suggests, or the (illegal) information leakage channel.

Panel A: Model with One Risky Asset



Panel B: Model with Two Risky Assets




Figure 1.5: Price(s) of the Risky Asset(s) in the Model

Notes: Panels A and B plot the average cum-dividend price(s) of the risky asset(s) in the model with one risky asset and the model with two risky assets (a and b), respectively. I assume that the news report is good for the risky asset(s): $n > 0$ in Panel A, and $n_a > 0$ and $n_b > 0$ in Panel B; otherwise, flip the figures. In both panels, the $i = 1$ and $i \in (0, 1)$ lines correspond to the case when all traders are sophisticated and the case when measure i of traders are sophisticated, respectively. In Panel A, the risky asset price starts at 0 and ends at the efficient-market level αn . In Panel B, the prices of a and b start at 0 and end at the efficient-market levels $\alpha_a n_a$ and $\alpha_b n_b$, respectively. I scale the parameters such that $\alpha_a n_a = \alpha_b n_b$ (see Section 1.6.2). α , α_a , and α_b are defined in Appendices A.5 and A.6 as functions of model parameters.

with α defined in Appendix A.5 as a function of model parameters. In particular, the price does not change on the public announcement, as no new information arrives. When $i = 1$, the market is efficient: once the news report is knowable to traders, the news report is immediately and completely impounded into the price of the risky asset.

1.6.1.2 Inefficient-Market Specification: Some Traders are Unsophisticated, and $i \in (0, 1)$

The dashed $i \in (0, 1)$ line in Panel A of Figure 1.5 plots the average *cum*-dividend price of the risky asset when only measure i of traders are sophisticated. The price of the risky asset starts at 0. It increases at $t = 1$ and $t = 2$ as sophisticated traders impound their advance knowledge of a good news report. At $t = 3$, the unsophisticated also see the news report. The price increases to that in Eq. (1.8) and stays there at $t = 4$.

As the figure clearly shows, the pre-announcement price of the risky asset does not increase all the way to the post-announcement efficient-market level: the market is inefficient. The price of the risky asset at $t = 2$ is a fraction of that at $t = 3$:

$$p_2 = \frac{1}{1 + \left(\frac{1-i}{i}\right) \left(\frac{A_S}{A_U}\right) \frac{\sigma_{s,l}^2}{(\alpha^2(\sigma_d^2 + \sigma_\epsilon^2) + \sigma_{s,l}^2)}} p_3. \quad (1.9)$$

p_2 is at a discount to p_3 because sophisticated traders are risk-averse (positive A_S) and face fundamental risk in the late trading period (positive $\sigma_{s,l}^2$). As risk aversion or fundamental risk increase, the discount increases. Similarly, the price of the risky asset at $t = 1$ is a fraction of that at $t = 2$ because risk-averse sophisticated traders hold back their trading due to fundamental risk in the early trading period:

$$p_1 = \frac{1}{1 + \left(\frac{1-i}{i}\right) \left(\frac{A_S}{A_U}\right) \frac{\sigma_{s,e}^2}{(k_2^2 \alpha^2 (\sigma_d^2 + \sigma_\epsilon^2) + \sigma_{s,e}^2)}} p_2, \quad (1.10)$$

with k_2 defined in Appendix A.5 as a function of model parameters.

The above analysis leads to the following proposition that describes how the model accounts for the late trading puzzle.

Proposition 1 *Sophisticated traders obtain advance knowledge of the news report at $t = 1$. They can trade immediately at $t = 1$ and/or wait to trade at $t = 2$ right before the announcement. Suppose that the risk faced by sophisticated traders in the early trading period is large relative to that faced by sophisticated traders in the late trading period: $\sigma_{s,e}^2 > \sigma_{s,l}^2$. If so, sophisticated traders recognize that they can achieve a higher risk-adjusted return by optimally timing their trading and trading little immediately at $t = 1$ and instead waiting to trade a lot at $t = 2$. As a result, the price of the risky asset changes in the direction of the news report before the public announcement at both $t = 1$ and $t = 2$, but the change is small at $t = 1$ and large at $t = 2$. The market is inefficient over time: the advance knowledge of sophisticated traders is inefficiently impounded into the risky asset with a delay.*

It is a simplistic assumption that sophisticated traders, after acquiring their informational advantage, have only two pre-announcement trading opportunities and face greater risk trading early than trading late. Yet this assumption is also realistic. In practice, traders are boundedly rational. Cognitive and time limitations likely prevent traders from trading too frequently and instead compel traders to use heuristics such as trading in early and late periods. Whether there actually is greater risk from trading early versus trading late is ultimately an empirical question, which I tackle in Section 1.6.3.

1.6.2 Model with Two Risky Assets

In the model above, risk-averse sophisticated traders dislike risk in the time series and consequently vary their trading intensity over time. I introduce a second risky asset into the model and show that risk-averse sophisticated traders also dislike risk in the cross section and consequently vary their trading intensity across assets.

The timeline of the model with two risky assets is shown in Panel B of Figure 1.4. At $t = 4$, risky assets a and b pay liquidating dividends $d_a \sim \mathcal{N}(0, \sigma_{d,a}^2)$ and $d_b \sim \mathcal{N}(0, \sigma_{d,b}^2)$, respectively. At $t = 3$, the public announcement of a news report is separately interpreted

as noisy signals of d_a and d_b :

$$n_a = d_a + \epsilon_a, \text{ and}$$

$$n_b = d_b + \epsilon_b.$$

$\epsilon_a \sim \mathcal{N}(0, \sigma_{\epsilon,a}^2)$, and $\epsilon_b \sim \mathcal{N}(0, \sigma_{\epsilon,b}^2)$. In the early trading period between $t = 1$ and $t = 2$, there are two shocks to holding risky assets a and b in the form of dividend payouts $s_{e,a} \sim \mathcal{N}(0, \sigma_{s,e,a}^2)$ and $s_{e,b} \sim \mathcal{N}(0, \sigma_{s,e,b}^2)$, respectively. Similarly, in between $t = 2$ and $t = 3$ in the late trading period, there are two shocks in the form of dividend payouts for assets a and b : $s_{l,a} \sim \mathcal{N}(0, \sigma_{s,l,a}^2)$ and $s_{l,b} \sim \mathcal{N}(0, \sigma_{s,l,b}^2)$, respectively. All the above random variables are independent.

Assuming independence simplifies the analysis. Sophisticated traders operate in completely segmented asset markets, so the results of the model with one risky asset apply to each of assets a and b separately. Analogous to Panel A of Figure 1.5, Panel B shows the average *cum*-dividend price of both risky assets in two specifications of the model. I consider the case in which the news report is good for both risky assets: $n_a > 0$, and $n_b > 0$. The efficient-market levels for assets a and b are given in Eq. (1.8) after substituting n_a and n_b , respectively, for n and α_a and α_b , respectively, for α . α_a and α_b are defined in Appendix A.6 as functions of model parameters. For illustrative purposes, I scale the parameters such that the efficient-market levels of assets a and b are the same. I do so by setting $\sigma_{\epsilon,a}^2 = \sigma_{\epsilon,b}^2$, $\sigma_{d,a}^2 = \sigma_{d,b}^2$, and $n_a = n_b$.

1.6.2.1 Efficient-Market Specification: All Traders are Sophisticated, and $i = 1$

The “ $i = 1$, Assets a and b ” solid line corresponds to the baseline version of the two risky assets model in which all traders are sophisticated. As in the case with one risky asset described in Section 1.6.1.1, the prices of risky assets a and b start at 0 before jumping to the efficient-market levels upon traders obtaining advance knowledge of the news report.

1.6.2.2 Inefficient-Market Specification: Some Traders are Unsophisticated, and $i \in (0, 1)$

The “ $i \in (0, 1)$, Asset a ” dotted line and “ $i \in (0, 1)$, Asset b ” dashed line correspond to the version of the two risky assets model in which measure i of traders are sophisticated. As in the case with one risky asset described in Section 1.6.1.2, the prices of risky assets a and b start at 0. Trading based on advance knowledge of the good news report pushes prices up at $t = 1$ and $t = 2$, but not all the way due to fundamental risk: $\sigma_{s,e,a}^2, \sigma_{s,e,b}^2, \sigma_{s,l,a}^2, \sigma_{s,l,b}^2$ are all positive. Eqs. (1.9) and (1.10) still hold after substituting for the appropriate asset-specific parameters. For each asset, the price at $t = 1$ is at a discount to that at $t = 2$, which is in turn at a discount to that at $t = 3$. Only after the public announcement of the news report do both assets obtain the (same) efficient-market levels.

To capture the notion that traders face different levels of risk over time, I assume that it is riskier to trade early than late: $\sigma_{s,e,a}^2, \sigma_{s,e,b}^2 > \sigma_{s,l,a}^2, \sigma_{s,l,b}^2$. There is little trading in the early period, so the prices of a and b increase modestly at $t = 1$. Most of the trading occurs in the late period, so the prices of a and b increase more sharply at $t = 2$. These time series results follow directly from the one risky asset version of the model.

To capture the notion that traders face different levels of risk trading different assets, I assume that asset a is riskier than asset b . Holding a position in the former entails weathering more fundamental risk in both early and late trading periods than holding a position in the latter: $\sigma_{s,e,a}^2 > \sigma_{s,e,b}^2$, and $\sigma_{s,l,a}^2 > \sigma_{s,l,b}^2$. Risk-averse sophisticated traders recognize that they can achieve a higher risk-adjusted return by trading the safer asset than the riskier one. As a result, sophisticated traders trade asset b more intensely than asset a , impound more of their informational advantage into b than a , and make b more efficient than a before the public announcement. Plugging in $\sigma_{s,l,a}^2$ and $\sigma_{s,l,b}^2$ for $\sigma_{s,l}^2$ in Eq. (1.9), $p_{2,b} > p_{2,a}$. Similarly, plugging in $\sigma_{s,e,a}^2$ and $\sigma_{s,e,b}^2$ for $\sigma_{s,e}^2$ in Eq. (1.10), $p_{1,b} > p_{1,a}$. The safer asset increases more than the riskier asset at $t = 1$ and $t = 2$ and reflects more of the good news report. Since most sophisticated trading in both assets takes place in the late period, the discrepancy between a and b is particularly noticeable then.

Two propositions emerge from the above analysis. First, both assets react to the public announcement of the news report: $p_{3,a} - p_{2,a} > 0$, and $p_{3,b} - p_{2,b} > 0$. This reaction occurs despite the fact that sophisticated traders can perfectly predict the news report in advance. The nature of the reaction differs for assets a and b , however.

Proposition 2 *The safer asset b reacts less strongly than the riskier asset a to the public announcement of the news report: $p_{3,b} - p_{2,b} < p_{3,a} - p_{2,a}$.*

Risk prevents sophisticated traders from bidding up the pre-announcement prices of a and b to the efficient-market levels. The greater risk of a impedes market efficiency more than the lower risk of b .

Second, there is return momentum in each asset between the respective pre-announcement and announcement returns: $p_{2,a} - p_{1,a} > 0$ is followed by $p_{3,a} - p_{2,a} > 0$, and $p_{2,b} - p_{1,b} > 0$ is followed by $p_{3,b} - p_{2,b} > 0$.²⁴ The response of risky assets to the public announcement is preceded by pre-announcement returns in the same direction due to sophisticated traders impounding their advance knowledge of the good news report. The return momentum differs for assets a and b , however.

Proposition 3 *Under a general condition, the safer asset b has less return momentum between pre-announcement and announcement returns than the riskier asset a .*²⁵

Proposition 3 is directly related to Proposition 2. The pre-announcement return in assets a and b are followed by larger and smaller announcement returns, respectively.

²⁴Pre-announcement here refers specifically to the late trading period immediately prior to the public announcement.

²⁵The condition is that $\sigma_{s,l,a}^2$ cannot be too large. Otherwise the price of a changes little at $t = 2$ due to the high risk of trading a during the late period, and the return momentum of a is small. Since the model assumes that the risk of trading any asset in the late period is low (and consequently incentivizes sophisticated traders to wait to trade), this condition seems reasonable.

1.6.3 Data are Consistent with the Model

To summarize, Proposition 1 of the model says that the news report is inefficiently impounded into the price of the risky asset with a delay. The empirical analog from Section 1.5.1 is that the public information that predicts macroeconomic news reports is available days before MNAs but impounded into stocks and bonds just hours before MNAs. Propositions 2 and 3 of the model say that the news report is more efficiently impounded into the safer asset than the riskier asset. The empirical counterpart from Section 1.5.2 is that the public information that predicts news reports is fully incorporated into bonds but only partially incorporated into stocks before MNAs.

A key idea underlying the model propositions is that sophisticated traders recognize that they can achieve a higher risk-adjusted return due to lower risk by trading late instead of early and by trading the safer asset instead of the riskier asset. This idea holds in the data as well: there is evidence that sophisticated traders with advance knowledge of macroeconomic news reports find it more attractive to trade stocks and bonds closer to announcements and stocks versus bonds before announcements. One piece of evidence is based on the following thought experiment in which sophisticated traders have perfect foresight of macroeconomic news reports. Nobody in practice has perfect foresight, so the real-world data on stocks and bonds does not reflect trading on perfect foresight. How should sophisticated traders exploit perfect foresight given that this informational advantage has not been traded in stocks and bonds?

To answer this question, I construct hypothetical trading strategies that go long {short} stocks and short {long} bonds before positive {negative} news reports above {below} a given threshold.²⁶ Panel A of Figure 1.6 plots the annualized Sharpe ratios from the strategies starting some number of days or minutes before MNAs, as indicated on the x-axis. In both stocks and bonds, positions are unwound 5 minutes after MNAs. The risk-adjusted returns from trading stocks and bonds are low in the days prior to MNAs, much higher

²⁶The results shown are for news reports of greater than 0.7 standard deviation in magnitude. Changing the threshold does not substantively alter the results.

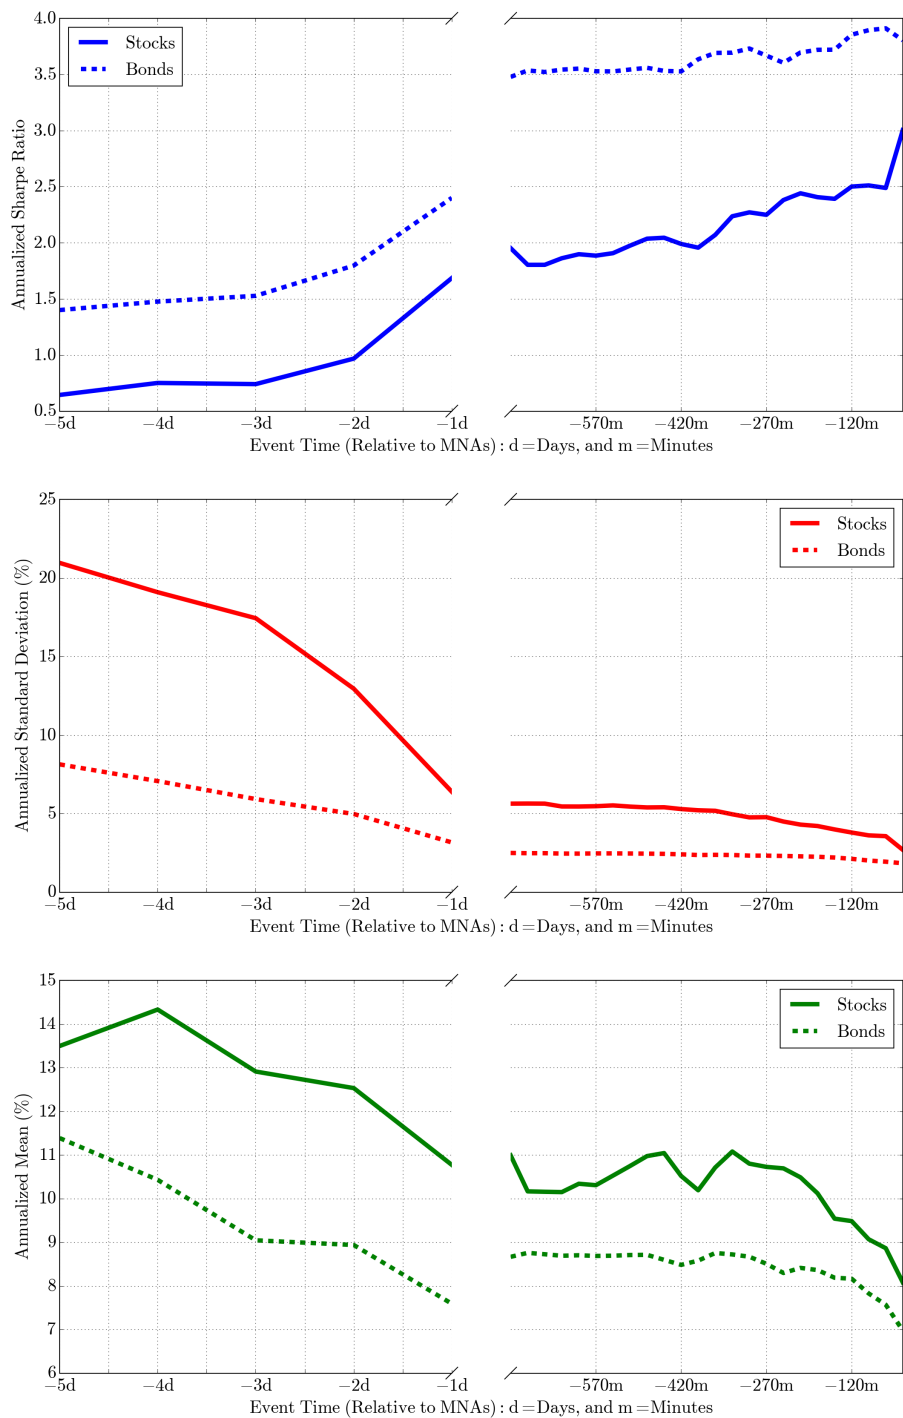


Figure 1.6: Trading on Perfect Foresight of Macroeconomic News Reports

Notes: See Appendix A.7.1. Panels A, B, and C from top to bottom correspond to annualized Sharpe ratios, annualized standard deviations, and annualized means, respectively.

in the hours prior to MNAs, and highest in the moments just prior to MNAs. In Panels B and C of Figure 1.6, I decompose the annualized Sharpe ratios into their two components: the annualized standard deviations in the denominators and the annualized means in the numerators, respectively. The increases in the Sharpe ratios from trading closer to MNAs are completely due to the decreases in the return standard deviations. The return means actually decrease but not as quickly as the return standard deviations. On some level, this finding is obvious: trading earlier entails holding risky positions for longer. Since risk scales with time, there has to be more risk from trading earlier. Trading in the hours before MNAs, however, is especially low risk because announcements occur in the morning Eastern Time, and there is little volatility in stocks and bonds during the overnight trading session. Based on Figure 1.6, sophisticated traders are incentivized to trade closer to MNAs and obtain higher risk-adjusted returns by avoiding risk. Viewed through the lens of the model, markets are inefficient over time in the manner described in Section 1.5.1.

Panel A of Figure 1.6 also presents evidence that the risk-adjusted returns from trading bonds are higher than those from trading stocks and particularly so in the hours prior to MNAs. Panels B and C of Figure 1.6 show that the higher Sharpe ratios of bonds versus stocks are completely driven by the lower return standard deviations of bonds versus stocks, which offset the lower return means of the former asset. The data suggest that bonds are a safer asset than stocks for the purpose of expressing an informational edge on macroeconomic news reports. Such a conclusion is consistent with the intuition that orthogonal sources of volatility play a more important role in stocks than bonds. For example, around the time MNAs are released in the morning Eastern Time, earnings announcements are often also released. Earnings announcements are unlikely to be related to the MNAs explored in this paper. Earnings announcements are, however, much more likely to impact stocks than bonds. Sophisticated traders acting on predictions of macroeconomic news reports are plausibly more wary of trading stocks than bonds and being exposed to volatility from earnings surprises. Based on Figure 1.6, sophisticated traders are incentivized to trade bonds more intensely than stocks before MNAs and obtain higher risk-adjusted returns by

avoiding risk. Viewed through the lens of the model, markets are inefficient across assets in the manner described in Section 1.5.2.

1.7 Trading on Market Inefficiency

1.7.1 Trading on the Predictable Component of Macroeconomic News Reports


Market inefficiency is commonly associated with money being left on the table. How much money do sophisticated traders leave on the table as a result of being deterred by risk from trading their predictions of macroeconomic news reports into prices? Figure 1.7 plots the annualized Sharpe ratios from implementable trading strategies that use the predictable component of macroeconomic news reports derived from anchoring bias. The strategies take positions in stocks (Panels A and B) and bonds (Panels C and D) starting some number of days (Panels A and C) or minutes (Panels B and D) before MNAs. In stocks, the strategies go long {short} before announcements of the top {bottom} x percentile of the predictable component, with x indicated in the “Trade Threshold” axis. In bonds, the strategies go the opposite direction. Positions are unwound 5 minutes after MNAs for stocks and 5 minutes before MNAs for bonds.²⁷ To account for transaction costs, I include the round-trip bid-ask spread into calculations of the Sharpe ratios.²⁸

The Sharpe ratios in Figure 1.7 are sizable. The maximum Sharpe ratio in stocks is 0.82 for trading strategies initiated in the days before MNAs and 0.65 for strategies initiated in the hours before MNAs. For bonds, the maximum Sharpe ratio is even higher: 1.12 and 0.82 for strategies starting in the days and hours, respectively, before MNAs. It is notable that a simple algorithm based on anchoring bias can present viable trading strategies. For market participants with more powerful ways to predict macroeconomic news reports, the Sharpe


²⁷Based on Table 1.5, stocks but not bonds react to announcements of the predictable component. Thus, optimal trading strategies based on the predictable component should hold through announcements for stocks but not for bonds.

²⁸For stocks (ES futures contract), the round-trip bid-ask spread is 0.25 point. For bonds (TY futures contract), the round-trip bid-ask spread is $1/64 = 0.015625$ point.


Panel A: Stocks in the Days before MNAs



Panel B: Stocks in the Hours before MNAs



Panel C: Bonds in the Days before MNAs



Panel D: Bonds in the Hours before MNAs




Figure 1.7: Trading on the Predictable Component of Macroeconomic News Reports with Transaction Costs

Notes: See Appendix A.7.2.

ratios in Figure 1.7 represent a lower bound of achievable risk-adjusted returns.

The results pose the question of why more traders do not go from being unsophisticated to being sophisticated. Increasing the measure i of sophisticated traders in the model makes asset prices more efficient and reduces the amount of money left on the table. A more comprehensive extension of the model in which traders endogenously choose to become sophisticated may shed light on this question. Alternatively, it is also possible that the seemingly attractive risk-adjusted returns in Figure 1.7 are in fact misleading. First, the plotted Sharpe ratios fluctuate considerably depending on the trade threshold and time of entry. In that sense, the trading strategies are not particularly robust. Second, there are uncalculated transaction costs such as price impact and trading commissions. Third, the Sharpe ratio may not adequately capture the risk involved in trading around MNAs; for example, higher return moments may reveal hidden risks.

1.7.1.1 Optimal Trade Timing

In addition to showing the profitability of trading on market inefficiency, Figure 1.7 illustrates an important point about optimal trade timing. Front running is central to all the strategies that have high Sharpe ratios: strategies that take positions right before MNAs have low or negative risk-adjusted returns. In other words, it is better to trade earlier than to wait until the last moment to trade the predictable component of macroeconomic news reports.

This result may seem at odds with a mantra of this paper that sophisticated traders with an informational advantage about news reports are incentivized to trade later. In the model, sophisticated traders recognize that they can obtain a higher risk-adjusted return by trading late *on information not yet impounded into prices*. To this point, Figure 1.6 shows that perfect foresight trading strategies yield the highest Sharpe ratios when implemented immediately before MNAs. These Sharpe ratio calculations are based on real-world stock and bond data in which perfect foresight has not been traded into prices.

In the model, after sophisticated traders trade, they make prices more efficient in the late period than the early period. To an econometrician observing this equilibrium *in*

which information is already impounded into prices, the difference between the risk-adjusted returns from trading late and trading early is diminished. Trading early is still riskier than trading late. Trading early, however, also gives a higher return mean than trading late, since sophisticated traders drive the return mean down more in the late period. The Sharpe ratios calculated in Figure 1.7 are from trading the predictable component of news reports. The predictable component is already (partly) impounded into stock and bond prices and more so closer to MNAs. Though there is less risk from trading the predictable component later, there is even less reward. Hence, the most attractive risk-adjusted returns arise from front running.

The above discussion gives insight to models of optimal trade timing in the informed trading literature. One factor that affects timing is the number of informed traders. As the number increases, strategic considerations come into play: each informed trader worries about being front-run and is consequently incentivized to trade earlier (Holden and Subrahmanyam (1992, 1994)). Another factor is the heterogeneity of private information. As heterogeneity increases, each informed trader benefits from delaying trades to prevent others from using prices to infer that part of the private information known only to him (Foster and Viswanathan (1996) and Back, Cao, and Willard (2000)).

Consider the Sharpe ratios based on perfect foresight in Figure 1.6. These Sharpe ratios are available to an informed trader who is the only possessor of his informational advantage. The reason is that nobody in practice knows the exact values of unreleased news reports. Thus, this scenario describes the case in which there is only the single informed trader, or the informed trader has private information unrelated to that of other informed traders. Existing models predict that, qualitatively, there is little rush to trade, which is borne out in the data. Not rushing to trade is different from waiting to trade, however: these models predict that trading is spread out over time, which is at odds with the data. My model suggests that risk compels traders to wait to trade.

Now consider the Sharpe ratios based on the predictable component of news reports in Figure 1.7. These Sharpe ratios are available to an informed trader who is one of many who

possess the same informational advantage. The reason is that other market participants can and do trade the predictable component. Thus, this scenario describes the case in which there are multiple informed traders and informed traders have homogeneous (or at least not very heterogeneous) private information. The aforementioned literature predicts that, qualitatively, there is a rush to trade, which is evidenced by the data. The rush to trade is so strong in these models, however, that informed traders immediately and fully impound their private information into prices. In contrast, the data reveal that even if trading the predictable component immediately before announcements is unprofitable, doing so a few hours earlier is still attractive. Once again, my model's explanation is that traders are influenced by risk to wait to trade.

1.7.2 Trading on Overreaction and Return Reversal

There is intriguing evidence of another form of market inefficiency in the data. Table 1.7 tests for overreaction around MNAs using regressions of return reversal in stocks and bonds:

$$\text{post}R_t = \alpha + \beta \text{evt}R_t + \epsilon_t.$$

In row x , the left-hand-side variable $\text{post}R_t$ is the cumulative post-announcement return (in bps) from 5 minutes to x minutes after MNAs. The right-hand-side variable $\text{evt}R_t$ is the ± 5 minute announcement return (in bps) around MNAs. $\hat{\beta}$ is negative and statistically significant for stocks but close to zero and generally insignificant for bonds. Stocks decrease {increase} up to 4.0 bps after a 10 bps positive {negative} announcement return. Bonds, however, are essentially flat in the hours following a 10 bps positive {negative} announcement return. The data show that stocks are less efficient than bonds in another important manner: stocks display overreaction and subsequent return reversal, but bonds do not. Overreaction creates attractive trading opportunities. Analogous to Figure 1.7, Figure 1.8 plots the annualized Sharpe ratios inclusive of transaction costs from trading strategies that go short {long} each asset after announcement returns above {below} positive {negative} thresholds. The strategies initiate positions 5 minutes after MNAs and unwind positions some number

Table 1.7: Regressions of Return Reversal in Stocks and Bonds

	Stocks			Bonds		
	$\hat{\beta}$ (bps)	$\hat{\alpha}$ (bps)	R_{adj}^2 (%)	$\hat{\beta}$ (bps)	$\hat{\alpha}$ (bps)	R_{adj}^2 (%)
60	-0.04	-0.32	0.04	0.03*	0.41	0.12
120	-0.20***	-1.24	0.89	0.00	0.85*	-0.06
180	-0.29***	-1.87	1.39	0.04	1.11**	0.01
240	-0.27***	-1.06	1.00	0.06	0.96	0.09
300	-0.30***	0.35	1.01	0.06	0.63	0.08
360	-0.40***	0.58	1.18	0.04	1.06	-0.01
420	-0.37***	0.60	0.84	0.05	1.26	0.02
480	-0.35***	1.38	0.62	0.05	1.49*	0.01
540	-0.35***	0.30	0.60	0.04	1.29	-0.01
600	-0.34***	0.22	0.56	0.04	1.33*	-0.02
660	-0.34***	0.15	0.55	0.03	1.51*	-0.03
720	-0.32***	0.25	0.50	0.03	1.58**	-0.04


Notes: The table shows results from regressions of post-announcement returns on announcement returns for stocks (front ES futures contract) and bonds (front TY futures contract): $postR_t = \alpha + \beta evtR_t + \epsilon_t$. Regressions are run on the full sample of MNAs, and $n = 1705$. In row x , the left-hand-side variable $postR_t$ is the cumulative post-announcement return (in bps) from 5 minutes to x minutes after MNAs. The right-hand-side variable $evtR_t$ is the ± 5 minute announcement return (in bps) around MNAs. t -statistics (not shown) are based on heteroskedasticity-consistent standard errors. *** denotes significance at the 1% level, ** denotes significance at the 5% level, and * denotes significance at the 10% level.

of minutes later. The Sharpe ratio is as high as 1.21 for stocks but zero or negative for bonds.

The above finding motivates a more comprehensive model. Brunnermeier (2005) provides intuition for how such a model might look. In his work, sophisticated traders impound their informational advantage into prices prior to a public announcement. Other, unsophisticated market participants attempt to learn from pre-announcement prices what the sophisticated traders know but are not able to do so perfectly, which leads to an overreaction to the announcement. Sophisticated traders take advantage by front-running pre-announcement and flipping post-announcement; that is, they follow the well-known Wall Street adage and “buy the rumor and sell the news.”²⁹ Kadan, Michaely, and Moulton (2014) find support in the data for the theory: institutional investors buy before sell-side analyst upgrades and

²⁹If sophisticated traders’ advance knowledge indicates the value of the risky asset should be higher, they buy the rumor and sell the news. Otherwise, they “sell the rumor and buy the news.”

Panel A: Stocks



Panel B: Bonds




Figure 1.8: Trading on Return Reversal after MNAs with Transaction Costs

Notes: The figure plots the annualized Sharpe ratios from trading strategies based on return reversal after MNAs. The strategies take positions in stocks (front ES futures contract) and bonds (front TY futures contract) starting 5 minutes after MNAs. In both assets, the strategies go short {long} after the top {bottom} “Trade Threshold” percentile of ± 5 minute announcement returns. Positions are unwound the indicated number of minutes after MNAs.

sell when the upgrades are publicly announced. If, in my model, unsophisticated traders attempt to extract from pre-announcement prices the information of sophisticated traders, there is also overreaction and return reversal. The intuition is similar to that of Brunnermeier (2005). The challenge, however, is to construct a model that simultaneously incorporates overreaction and return reversal while preserving the other results shown in my research.

1.8 Conclusion

In this paper, I find novel evidence of sophisticated trading in a unique setting: stocks and bonds around MNAs. The data suggest that sophisticated traders acquire their informational advantage via the homework channel and skillfully process public information to predict macroeconomic news reports. In contrast, the literature generally attributes sophisticated trading in various settings to (illegal) information leakage or is unable to specify the mechanism by which sophisticated traders obtain their information.

Using these results, I show that stock and bond markets are inefficient around MNAs. Over time, there is a late trading puzzle: the public information that predicts macroeconomic news reports is available days before announcements but inefficiently impounded into stocks and bonds just hours before announcements. Across assets, there is a related puzzle: by the announcement times, the public information that predicts news reports is fully impounded into bonds but only partially impounded into stocks. Bonds eventually become efficient, while stocks remain inefficient: stocks but not bonds react to announcements of the predictable component of news reports and display return momentum between pre-announcement and announcement returns.

I construct a model to explain the documented market inefficiency. The model provides new insights into how sophisticated traders optimally trade and impound their informational advantage into asset prices. In the model, sophisticated traders have risk management concerns. If trading earlier is riskier than trading later, sophisticated traders wait to trade. If trading one asset is riskier than trading another asset, sophisticated traders trade the safer asset more intensely than the riskier asset. Consistent with the data, the model implies that,

in the time series, the information of sophisticated traders is impounded into asset prices later when it is safer. Similarly, in the cross section, the information is impounded more fully into safer assets such as bonds than riskier assets such as stocks.

My results are directly relevant for practitioners as well. As a result of being deterred by risk from trading their predictions of macroeconomic news reports into prices, sophisticated traders leave money on the table. Implementable trading strategies that exploit market inefficiency yield sizable Sharpe ratios in bonds and stocks.

This paper has three other implications. First, any research that uses MNAs as the empirical setting should take into account the presence of sophisticated traders. For example, looking at the price adjustment of stocks and bonds in response to announcements of macroeconomic news reports is misleading when a sizable fraction of the price adjustment takes place before announcements. Second, market participants may be more knowledgeable about the state of the economy than thought otherwise. How realistic, then, is the asymmetric information assumption of monetary policy models that central banks have an informational advantage about the economy over other agents? Finally, if markets are inefficient in a controlled setting such as around MNAs, markets may be similarly inefficient in many other settings. The two ingredients for market inefficiency in my research are that unsophisticated traders fail to recognize market inefficiency, and sophisticated traders recognize market inefficiency but fail to act due to risk. It is plausible that both ingredients are present in other settings.

There are a number of directions for future work. Most intriguingly, features of the data suggest that factors aside from risk influence informed trading: the number of informed traders, the heterogeneity of private information, and the strategic behavior of the informed who take advantage of the uninformed. Incorporating all these factors into a more sophisticated theory is challenging but promises to deepen our understanding of the price discovery process.

Chapter 2

The Good, the Bad, and the Ambiguous: The Aggregate Stock Market Dynamics around Macroeconomic News¹

2.1 Introduction

Uncertainty plays a central role in determining asset prices. How investors interpret news and impound that information into financial assets is a likewise important topic in asset pricing. These two themes are very much related, for investors are routinely exposed to news of uncertain value and must pass judgement on the relevance of said news for financial markets. The first focus of this paper is the use of a representative agent model based off Epstein and Schneider (2008) to investigate how investors process news that has a specific type of uncertainty, Knightian uncertainty or ambiguity, and the resulting implications for asset pricing. In the model, the investor sees a piece of news about the fundamental

¹I thank Geert Bekaert, Itamar Drechsler, and Marie Hoerova for helpful discussions and suggestions.

value of a risky asset, but this news is ambiguous in the sense that the investor knows the variance of the noise component only within a range. If the investor believes that the variance is low, the news is precise or relevant; otherwise, the news is imprecise or irrelevant. The investor is furthermore ambiguity-averse and picks the worst-case belief for the noise variance that minimizes expected utility. This setup yields four predictions for the behavior of stocks around ambiguous news, the first three of which are (i) the asymmetry effect or asymmetrically strong response of stocks to bad news versus good news, (ii) the no news is good news effect or the positive response of stocks to no (neutral) news, and (iii) the ambiguity premium or the average positive return of stocks around news. As the amount of ambiguity, equivalent to the range of the noise variance, increases, the model further conjectures that the magnitude of each of these three phenomenon increases. There is an additional (iv) risky asset volatility effect or corresponding increase or decrease in the volatility of stocks around news depending on the source of the increase in ambiguity.

Much of the work involving Knightian uncertainty is theoretical in nature.² The second focus of this paper is thus empirical and shows that the aforementioned model with ambiguity accurately captures the real-world behavior of stocks in a way that existing frameworks in the literature cannot. To arrive at this conclusion, I investigate the behavior of the aggregate stock market around macroeconomic news announcements (MNAs) concerning inflation, output, and the labor market. The results demonstrate that ambiguity can contribute to our understanding of how asset markets price macroeconomic shocks. Using high-frequency intraday data on the S&P 500 futures contract in the ± 5 minutes around a comprehensive sample of macroeconomic data releases from November, 1997 to March, 2014, I show whole sample evidence consistent with the asymmetry effect, the no news is good news effect, and the ambiguity premium. In the 10-minute window around MNAs, stocks increase 2.416 bps to a unit of good news but decrease 8.557 bps to a unit of bad news, which results in a meaningful 6.141 bps asymmetry effect.³ At the same time, stocks increase 3.204 bps to

²It is not without reason that another name for Knightian uncertainty or ambiguity is unmeasurable uncertainty.

³News is a standardized variable such that good news corresponds to “better”-than-expected data; i.e.,

no (neutral) news. All of the referenced numbers are statistically significant. Strikingly, the stock market increases a significant 1 bp on average through the 3,194 10-minute return intervals that contain macroeconomic news releases yet only an insignificant 0.014 bps on average through the other 588,671 10-minute return intervals during the sample period. In other words, stocks compounded 36.250% around MNAs and 66.267% otherwise such that about 1/3 of the stock market return over 17 years accrues in that tiny fraction of time around macroeconomic news releases. The standard deviation of returns is, however, only modestly higher in intervals containing MNAs relative to all other intervals, so this set of empirical observations is hard to reconcile with standard risk-based stories. The ambiguity premium provides an explanation for this puzzle.

The model makes several sharper predictions for how the aggregate stock market dynamics around macroeconomic data releases depends on the amount of ambiguity. Motivated by earlier academic studies such as Drechsler (2013), I use the variance risk premium, the difference between risk-neutral and physical expectations of stock market return variance, as a state variable for ambiguity in empirical tests. Existing measures of the variance risk premium, however, are problematic as proxies for ambiguity. To overcome these issues, I construct my own time series which behaves as a more reasonable proxy. The variance risk premium is elevated during well-known times of market stress such as the 2008 financial crisis when investors are plausibly less sure about the relevance of a piece of macroeconomic news for the fundamental value of the stock market; i.e., ambiguity is high. In less turbulent times, such as the mid-2000s, the variance risk premium is lower, consistent with the intuition that investors are more certain about the salience of macroeconomic data for stocks; i.e., ambiguity is low. As a consequence of how the variable is constructed, the variance risk premium predicts stock returns at long horizons but not short horizons in contrast to results in the literature.

Using the variance risk premium time series as a proxy for ambiguity, I find that, consistent with the model, the stock market reacts most asymmetrically to bad news versus

higher-than-expected inflation or output and lower-than-expected unemployment. See Section 2.3.1 for details.

good news during times of high ambiguity and progressively less so during times of lower ambiguity. The reaction of stocks to signed news is essentially symmetric when ambiguity is at its lowest. Intriguingly, the stronger reaction of stocks to bad news relative to good news in times of high ambiguity comes from the increased sensitivity of stocks to bad news and not from the decreased sensitivity of stocks to good news. This result, interpreted through the model, says that ambiguity, defined as the range of news noise variance, changes due to changes in the lower bound of the range instead of the upper bound of the range. The real-world implication for the behavior of investors is that during times of high ambiguity, which also tend to be times of panic in financial markets, investors think that bad macroeconomic news is more relevant for the stock market but good macroeconomic news is not any less relevant for the stock market. Phrased simply, investors focus more on bad news in bad times than in good times but treat good news the same in good and bad times.

The data also corroborate the no news is good news effect: Stocks react more strongly to neutral MNAs during times of high ambiguity than during times of low ambiguity. The evidence for the ambiguity premium is less clear-cut. True to the model, stocks earn a positive average return around macroeconomic data releases, but this average return is not increasing in ambiguity. Indeed, there is little time-variation in the average return of stocks around MNAs. The implication is that the ambiguity premium exists but is not time-varying. Based on the asymmetry effect, we know that changes in ambiguity occur in a specific way, and how ambiguity changes helps to reconcile the lack of time-variation in the ambiguity premium with the model. Moreover, the way in which ambiguity changes implies, through the risky asset volatility effect, that the volatility of stocks around MNAs should be higher in times of high ambiguity. This result shows up in the data.

How does the model generate the predictions that it does? When the investor sees a piece of news with noise variance known only within a range, ambiguity aversion dictates that the investor believes the actual noise variance is the one that minimizes his expected utility. If the news is positive for the fundamental value of the risky asset, the investor believes that the noise variance is high; hence, the news has low relevance and the risky

asset responds weakly. If the news is negative, the investor believes that the noise variance is low; hence, the news has high relevance and the risky asset responds strongly. This is the asymmetry effect. Prior to the arrival of news, the investor knows that ambiguous news is coming. Because the investor dislikes ambiguity, he must be compensated for holding the risky asset. Neutral news provides no information on the value of the risky asset but does resolve the ambiguity, so the rise of the risky asset in response to no news measures this compensation. This is the no news is good news effect. The average change in the price of the risky asset around news is equal to the sum of the no news is good news and asymmetry effects. This sum is positive and is equal to the ambiguity premium. Finally, if ambiguity or the range of news noise variance increases because the lower (upper) bound of the range decreases (increases), the risky asset responds more (less) strongly to negative (positive) news. As such, the volatility of the risky asset around news increases (decreases). This is the risky asset volatility effect.

While we must always exercise caution in interpreting the implications of a representative agent model for the real-world, there are several reasons why the behavior of the aggregate stock market (the risky asset) around news about macroeconomic conditions (the news) is an appropriate environment to test the model. Investors care about the fundamental value of the stock market, which is ultimately a stream of dividends, but can only infer future dividends from various news reports. One important type of news is that concerning the macroeconomy, for it is well-established that market participants pay close attention to and actively trade such reports. As a result, the aggregate stock market reacts in a systematic manner to MNAs: Stocks react positively to news of higher-than-expected inflation or output and lower-than-expected unemployment. MNAs are not, however, precise pieces of news about stock market fundamentals. A piece of news about macroeconomic conditions does not lay out the future stream of dividends with certainty, and there is considerable freedom in interpreting the relevance of such news. In other words, MNAs are ambiguous pieces of news about the stock market. Taking the model literally, investors have a set of beliefs encoded in the set of precisions of a given macroeconomic news release and its

relevance for the fundamental value of the stock market. The amount of ambiguity can vary over time with high levels of ambiguity corresponding to environments in which investors have a great deal of uncertainty in interpreting the salience of macroeconomic news for stocks; that is, the precision of the news in investors' minds could be anywhere from very low to very high. When ambiguity is low, however, investors are more sure of how informative macroeconomic news is for the stock market, so investors have a narrower range of precisions in mind. *Ex ante*, we might expect that ambiguity is highest at identifiable times such as during the 2008 financial crisis or the 2011 height of the European sovereign-debt crisis and concurrent U.S. debt-ceiling crisis. At other times, such as the few years before the financial crisis when economic and financial conditions were relatively stable, we might expect lower levels of ambiguity. Introspection suggests that such a narrative is reasonable.

This paper builds on the literature at the intersection of ambiguity and finance. In particular, the model with ambiguity follows from Epstein and Schneider (2008) who use the max-min expected utility theory of Gilboa and Schmeidler (1989) to capture the idea of ambiguity aversion. Hansen and Sargent (2010) among other papers by the two co-authors represents an alternative but related way to formulate aversion to ambiguity using robust control theory. A few studies have looked at how ambiguity can aid in explaining various phenomena in financial markets. For example, Illeditsch (2011) shows that ambiguity can account for portfolio inertia and excess volatility, Boyarchenko (2012) attributes spikes in CDS spreads during the 2008 financial crisis to increases in ambiguity, and Drechsler (2013) calibrates a model with ambiguity that matches the variance risk premium and implied volatility surface of stock index options as well as other properties of stocks. Related to my findings, Williams (2014) provides empirical evidence that individual stocks respond more strongly to bad earnings news than to good earnings news and especially so when the VIX, used as a proxy for ambiguity, is elevated. To my knowledge, however, no existing research explores the broad set of equity characteristics addressed in this paper and how these features in the data provide evidence for a model with ambiguity.

The behavior of the aggregate stock market around macroeconomic data releases is a

novel setting in which to evaluate the impact of ambiguity. Many papers in the literature have used this setting to explore topics related to how financial assets reflect macroeconomic risk; e.g., Jones, Lamont, and Lumsdaine (1998), Andersen, Bollerslev, Diebold, and Vega (2003, 2007), Boyd, Hu, and Jagannathan (2005), Faust, Rogers, Wang, and Wright (2007), Faust and Wright (2012), and Savor and Wilson (2013, 2014). These studies have investigated an assortment of assets (e.g., stocks, bonds, and currencies) in various sample periods (e.g., expansionary and contractionary states). How asset prices vary in response to macroeconomic news can give insight into other areas of economics as well. Swanson and Williams (2013, 2014) and Zhou (2014), for example, show that the Federal Reserve's zero lower bound attenuates interest rate sensitivity to news about the macroeconomy.

An integral part of my research is the construction of an appropriate proxy for ambiguity. The state variable for ambiguity that I choose, the variance risk premium, has been analyzed from a number of different angles. Evidence from Bollerslev, Tauchen, and Zhou (2009), Faust and Wright (2012), Drechsler (2013), and Bekaert and Hoerova (2014) shows that the variance risk premium is a significant predictor of stock returns in the time series. Bekaert, Hoerova, and Lo Duca (2013) also find interesting lead-lag relationships between the variance risk premium and the stance of monetary policy. Importantly, academic work such as Bekaert, Hoerova, and Lo Duca (2013) and, especially, Drechsler (2013), have discussed and provided justification for the variance risk premium as a measure of ambiguity. There is no data series in the literature that serves as the gold standard for the variance risk premium, however, so much work on the subject is necessarily methodological in nature; e.g., Carr and Wu (2009) in addition to some of the aforementioned studies. One of the contributions of my paper is the construction of a variance risk premium series that behaves as a more reasonable metric of ambiguity in the time series. This data series differs meaningfully from those in the literature in other ways, for example in the predictive power for stock returns.

The predictions that the model with ambiguity makes and that I subsequently verify empirically are related to several strands of research. Consistent with the asymmetry effect, Barberis, Shleifer, and Vishny (1998) and Veronesi (1999) have models that imply that

stocks react more strongly to bad news than to good news. In their models, however, the asymmetry occurs in good times and decreases (or even flips) in bad times. Conrad, Cornell, and Landsman (2002) find that individual stocks do indeed react asymmetrically stronger to bad earnings announcements versus good earnings announcements in good times, as measured by the equity market valuation, but not in bad times. Brown, Harlow, and Tinic (1988) and Campbell and Hentschel (1992) are alternative explanations for the asymmetric response of stocks to signed news. None of these frameworks involve ambiguity, however, and I show in Section 2.7 of this paper that the behavior of stocks around macroeconomic news supports an ambiguity-based explanation for the asymmetry effect over some of these existing explanations. Note that, as mentioned previously, Williams (2014) shows additional evidence that ambiguity drives the asymmetry effect using the response of individual stocks to earnings announcements.

My finding of a large, positive average return of stocks around MNAs is most closely related to work by Savor and Wilson (2013) who show a similar result using daily data instead of intraday data. Instead of interpreting this positive return as evidence of an ambiguity premium, the co-authors construct a model in which macroeconomic data releases reveal important information about the state of the economy. Stocks do poorly if the state of the economy is bad, which creates a large risk premium on days with MNAs. Faust and Wright (2012) find that standard return predictability regressions cannot predict the amount stocks earn around MNAs, which suggests that the premium earned is not time-varying. The increase of the aggregate stock market around macroeconomic announcements also parallels the earnings announcement premium for individual stocks for which there are various explanations. Frazzini and Lamont (2007), for example, reason that earnings announcements grab the attention of investors who are short-sale constrained or otherwise have a predisposition to be long.

Finally, the risky asset volatility effect is related to the well-documented time-variation of stock return volatility. Campbell and Hentschel (1992) and Veronesi (1999) are two of the papers that propose mechanisms generating the volatility fluctuations seen in the data.

The organization of this paper is as follows. Section 2.2 introduces the representative agent model with ambiguity and highlights the predictions for the behavior of the risky asset around ambiguous news. Section 2.3 describes the data on macroeconomic news releases and high-frequency stock prices used to empirically evaluate the model. Through analysis of the aggregate stock market around MNAs, Section 2.4 presents whole sample evidence supportive of the asymmetry effect, no news is good news effect, and ambiguity premium. Section 2.5 discusses the variance risk premium, its properties, and its suitability as a state variable for ambiguity. Using the variance risk premium to proxy for ambiguity, Section 2.6 shows how the asymmetry effect, no news is good news effect, ambiguity premium, and risky asset volatility effect vary with ambiguity. Section 2.7 address alternative explanations for the model predictions. Section 2.8 concludes.

2.2 Model with Ambiguity

2.2.1 Utility Function

To illustrate the impact of ambiguity on financial markets, I consider a model with a representative agent based off Epstein and Schneider (2008). The agent is ambiguity-averse in the sense of having recursive multiple-priors utility with utility function U_t at time t :

$$U_t = \min_{m_t \in \mathcal{M}_t} \mathbb{E}_{m_t} [u(C_t) + \beta U_{t+1}]. \quad (2.1)$$

$u(\cdot)$, C_t , and β are the familiar notations for the Bernoulli utility function, consumption at time t , and the discount factor, respectively. I introduce the notation \mathcal{M}_t and m_t to denote the set of models considered by the investor at time t and a specific model within that set, respectively. Finally, $\mathbb{E}_{m_t}[\cdot]$ is the expectation given the beliefs generated by model m_t .

The utility function in Eq. (2.1) captures both risk-aversion and ambiguity-aversion. The former shows up through the Bernoulli utility function $u(\cdot)$. To see the latter, note that if \mathcal{M}_t consists of a single model, we are back in a standard framework, and hence there is no ambiguity-aversion. As the size of \mathcal{M}_t increases, ambiguity-aversion plays an increasingly

important role. The reason is that the investor contemplates a range of models within \mathcal{M}_t and picks the model m_t that generates worst-case beliefs and leads to the lowest expected utility.

2.2.2 Three Period Model

The stylized model has three dates: $t = 0, 1$, and 2 . The investor can invest in either a risky asset, the aggregate stock market, or a risk-free asset. The risky asset is a claim on a dividend d that is revealed at $t = 2$, while the risk-free rate is zero by assumption. As such, wealth at $t = 2$ is $W_2 = W_0 + (d - p)\theta$, in which W_0 is starting wealth at $t = 0$, p is the price of the risky asset, and θ is the amount invested in the risky asset. The investor cares only about consumption at $t = 2$ and consumes all of his wealth $C_2 = W_2$. I assume that there is no time discounting and that period utility is exponential; that is, $\beta = 1$, and $u(C_t) = -\exp(-AC_t)$, with A the coefficient of absolute risk aversion.

Working backward in time, at $t = 2$, the dividend

$$d = \bar{d} + \epsilon_d$$

is revealed. \bar{d} is the mean dividend, and $\epsilon_d \sim \mathcal{N}(0, \sigma_d^2)$ is a normal, mean-zero shock.

At $t = 1$, a piece of news arrives about the dividend:

$$n = \alpha\epsilon_d + \epsilon_n. \tag{2.2}$$

$\alpha \in [0, 1]$ measures the relevance of the news. For example, if $\alpha = 1$, the news is just a noisy estimate of ϵ_d . On the other hand, if $\alpha = 0$, the news provides no information on the dividend to be paid in one period. $\epsilon_n \sim \mathcal{N}(0, \sigma_n^2)$ is a normal, mean-zero shock that is independent from ϵ_d . The news is ambiguous in the sense that the agent only knows that the variance of ϵ_n is within an interval: $\sigma_n^2 \in [\underline{\sigma}_n^2, \bar{\sigma}_n^2]$. As the range $\bar{\sigma}_n^2 - \underline{\sigma}_n^2$ of this interval increases, there is a greater amount of ambiguity and the agent is less confident about the news' precision. Referring to prior notation, $[\underline{\sigma}_n^2, \bar{\sigma}_n^2]$ maps to \mathcal{M}_1 , and σ_n^2 maps to m_1 . From

Eq. (2.1), the utility function at $t = 1$ is

$$U_1 = \min_{\sigma_n^2 \in [\underline{\sigma}_n^2, \bar{\sigma}_n^2]} \mathbb{E}_{\sigma_n^2} [-\exp(-A(W_0 + (d-p)\theta)) | n]. \quad (2.3)$$

I ignore the $u(C_1)$ term because the investor only cares about consumption at $t = 2$, and I substitute $U_2 = -\exp(-A(W_0 + (d-p)\theta))$ to reflect that final period utility is simply the utility from consuming all available wealth W_2 . Finally, the utility function in Eq. (2.3) is conditional on observing the news n in Eq. (2.2).

At $t = 0$, the agent knows that ambiguous news will arrive in one period. His utility function at $t = 0$ is thus

$$U_0 = \min_{\sigma_n^2 \in [\underline{\sigma}_n^2, \bar{\sigma}_n^2]} \mathbb{E}_{\sigma_n^2} \left[\min_{\sigma_n^2 \in [\underline{\sigma}_n^2, \bar{\sigma}_n^2]} \mathbb{E}_{\sigma_n^2} [-\exp(-A(W_0 + (d-p)\theta)) | n] \right]. \quad (2.4)$$

This expression comes from straightforward substitution of Eq. (2.3) into Eq. (2.1). $\mathcal{M}_0 = [\underline{\sigma}_n^2, \bar{\sigma}_n^2]$ is the range of the noise variance from the first period perspective and is equal to \mathcal{M}_1 . From \mathcal{M}_0 , the agent at $t = 0$ selects the σ_n^2 that results in the worst-case expected utility.

2.2.3 Equilibrium Prices at $t = 0$ and $t = 1$ with Ambiguity-Aversion

We are interested in finding the price of the risky asset at $t = 0$, before the arrival of the news, and at $t = 1$, after the arrival of the news, in the case when the news is ambiguous, as well as in the case when the news is not ambiguous. To focus on the effect of ambiguity in the model, I assume that the investor is risk-neutral ($A = 0$). Consider the case of ambiguous news first, with $\bar{\sigma}_n^2 > \underline{\sigma}_n^2$.

2.2.3.1 Price at $t = 1$ with Ambiguity-Aversion

At $t = 1$, the investor's optimization problem based on Eq. (2.3) is

$$\begin{aligned} \max_{\theta} U_1 &= \max_{\theta} \min_{\sigma_n^2 \in [\underline{\sigma}_n^2, \bar{\sigma}_n^2]} \mathbb{E}_{\sigma_n^2} [-\exp(-A(W_0 + (d-p)\theta)) | n] \\ &= \max_{\theta} \min_{\sigma_n^2 \in [\underline{\sigma}_n^2, \bar{\sigma}_n^2]} -\log \left(\mathbb{E}_{\sigma_n^2} [\exp(-A(W_0 + (d-p)\theta)) | n] \right). \end{aligned} \quad (2.5)$$

The second equality comes from the fact that the $\log(\cdot)$ function is monotonic. Since the joint distribution of d and n is given by

$$\begin{pmatrix} n \\ d \end{pmatrix} = \begin{pmatrix} \alpha \epsilon_d + \epsilon_n \\ \bar{d} + \epsilon_d \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} 0 \\ \bar{d} \end{pmatrix}, \begin{pmatrix} \alpha^2 \sigma_d^2 + \sigma_n^2 & \alpha \sigma_d^2 \\ \alpha \sigma_d^2 & \sigma_d^2 \end{pmatrix} \right), \quad (2.6)$$

following the formula in Appendix B.1 gives the result that dividend d given news n is conditionally normally distributed:

$$d|n \sim \mathcal{N} \left(\bar{d} + \phi n, \sigma_d^2 (1 - \alpha \phi) \right), \quad (2.7)$$

with

$$\phi = \frac{\alpha \sigma_d^2}{\alpha^2 \sigma_d^2 + \sigma_n^2}. \quad (2.8)$$

Note that the conditional distribution in Eq. (2.7) depends on σ_n^2 through ϕ . In particular, $\phi \in [\underline{\phi}, \bar{\phi}] \subset [0, 1]$, with

$$\underline{\phi} = \frac{\alpha \sigma_d^2}{\alpha^2 \sigma_d^2 + \bar{\sigma}_n^2}, \text{ and} \quad (2.9)$$

$$\bar{\phi} = \frac{\alpha \sigma_d^2}{\alpha^2 \sigma_d^2 + \underline{\sigma}_n^2}. \quad (2.10)$$

Since I assume that $\bar{\sigma}_n^2 > \underline{\sigma}_n^2$, $\bar{\phi} > \underline{\phi}$. Using Eq. (2.7), I rewrite the optimization problem in Eq. (2.5):

$$\begin{aligned}
\max_{\theta} U_1 &= \max_{\theta} \min_{\phi \in [\underline{\phi}, \bar{\phi}]} -\mathbb{E}_{\phi} [-A (W_0 + (d - p) \theta) | n] - \frac{1}{2} \mathbb{V}_{\phi} [-A (W_0 + (d - p) \theta) | n] \\
&= \max_{\theta} \min_{\phi \in [\underline{\phi}, \bar{\phi}]} \theta \mathbb{E}_{\phi} [d - p | n] - \frac{1}{2} A \theta^2 \mathbb{V}_{\phi} [d - p | n] \\
&= \max_{\theta} \min_{\phi \in [\underline{\phi}, \bar{\phi}]} \theta \mathbb{E}_{\phi} [d - p | n]. \tag{2.11}
\end{aligned}$$

The first equality uses the fact that for a log-normal variable z , $\log(\mathbb{E}[z]) = \mathbb{E}[\log(z)] + \frac{1}{2} \mathbb{V}[\log(z)]$. I also replace σ_n^2 with ϕ and $[\underline{\sigma}_n^2, \bar{\sigma}_n^2]$ with $[\underline{\phi}, \bar{\phi}]$ due to the one-to-one mapping between the variables in Eqs. (2.8), (2.9), and (2.10). $\bar{\phi} - \underline{\phi}$ thus proxies for the amount of ambiguity in the same way that $\bar{\sigma}_n^2 - \underline{\sigma}_n^2$ does. The second step simplifies the optimization problem by dropping additive and multiplicative constants. The final step sets $A = 0$ to reflect that the agent is risk-neutral.

Solving Eq. (2.11) for the equilibrium price at $t = 1$ yields

$$p_1(n) = \min_{\phi \in [\underline{\phi}, \bar{\phi}]} \mathbb{E}_{\phi} [d | n] = \min_{\phi \in [\underline{\phi}, \bar{\phi}]} \bar{d} + \phi n = \begin{cases} \bar{d} + \underline{\phi} n & \text{if } n > 0 \\ \bar{d} + \bar{\phi} n & \text{if } n \leq 0 \end{cases}. \tag{2.12}$$

Since the investor evaluates actions under the worst-case scenario, the price at $t = 1$ minimizes the conditional mean dividend and depends on the news value. Because $\underline{\phi} < \bar{\phi}$, the result is that the price of the risky asset is asymmetric in the value of the news. Moreover, the extent of this asymmetry is larger as the difference between $\bar{\phi}$ and $\underline{\phi}$ grows and there is greater ambiguity. If the news is good and $n > 0$, $p_1(n)$ has sensitivity $\underline{\phi}$ to n . On the other hand, if the news is bad and $n \leq 0$, $p_1(n)$ has sensitivity $\bar{\phi}$ to n . When the investor sees positive news, he thinks that it has high variance $\bar{\sigma}_n^2$ (which corresponds to $\bar{\phi}$) and is therefore not particularly informative about the dividend in the next period. Upon seeing a negative piece of news, however, the investor thinks that it has low variance $\underline{\sigma}_n^2$ (which corresponds to $\underline{\phi}$) and contains precise information about the dividend.

While Eq. (2.12) gives the price conditional on the news value, we are also interested in

the average price of the risky asset after the investor sees the news. To derive this quantity, I use the distribution of n from Eq. (2.6), and make the key assumption that the true variance of ϵ_n is $\bar{\sigma}_n^2 \in [\underline{\sigma}_n^2, \bar{\sigma}_n^2]$. The average price \bar{p}_1 is then

$$\begin{aligned}
\bar{p}_1 &= \mathbb{E}_{\bar{\sigma}_n^2} \left[\bar{d} + \underline{\phi} n | n > 0 \right] \times \mathbb{P}_{\bar{\sigma}_n^2} [n > 0] + \mathbb{E}_{\bar{\sigma}_n^2} \left[\bar{d} + \bar{\phi} n | n \leq 0 \right] \times \mathbb{P}_{\bar{\sigma}_n^2} [n \leq 0] \\
&= \bar{d} + \frac{1}{2} \underline{\phi} \mathbb{E}_{\bar{\sigma}_n^2} [n | n > 0] + \frac{1}{2} \bar{\phi} \mathbb{E}_{\bar{\sigma}_n^2} [n | n \leq 0] \\
&= \bar{d} - \frac{1}{2} (\bar{\phi} - \underline{\phi}) \mathbb{E}_{\bar{\sigma}_n^2} [n | n > 0] \\
&= \bar{d} - \frac{1}{\sqrt{2\pi}} (\bar{\phi} - \underline{\phi}) \sqrt{\alpha^2 \sigma_d^2 + \bar{\sigma}_n^2}. \tag{2.13}
\end{aligned}$$

The second and third steps use the fact that the news is symmetric around zero, and the last step uses the property of the normal distribution shown in Appendix B.2 that for $z \sim N(0, \sigma_z^2)$, $\mathbb{E}[z | z > 0] = \sigma_z \sqrt{2}/\sqrt{\pi}$. If the news value equals zero, the price at $t = 1$ equals $p_1(0) = \bar{d}$ from Eq. (2.12). The average price, however, is below \bar{d} due to the asymmetric reaction of the price to the news. A visual way to see this result is that the price is concave in the news, so averaging over the price creates a Jensen's inequality effect. With greater ambiguity and a larger magnitude of $\bar{\phi} - \underline{\phi}$, the average price is discounted more relative to \bar{d} due to the stronger price reaction to a negative piece of news versus a positive piece of news. The size of this asymmetry discount is

$$\bar{p}_1 - p_1(0) = -\frac{1}{\sqrt{2\pi}} (\bar{\phi} - \underline{\phi}) \sqrt{\alpha^2 \sigma_d^2 + \bar{\sigma}_n^2}. \tag{2.14}$$

2.2.3.2 Price at $t = 0$ with Ambiguity-Aversion

Moving back to $t = 0$, the agent's optimization problem based on Eq. (2.4) is

$$\begin{aligned}
\max_{\theta} U_0 &= \max_{\theta} \min_{\sigma_n^2 \in [\underline{\sigma}_n^2, \bar{\sigma}_n^2]} \mathbb{E}_{\sigma_n^2} \left[\min_{\sigma_n^2 \in [\underline{\sigma}_n^2, \bar{\sigma}_n^2]} \mathbb{E}_{\sigma_n^2} [-\exp(-A(W_0 + (d-p)\theta)) | n] \right] \\
&= \max_{\theta} \min_{\sigma_n^2 \in [\underline{\sigma}_n^2, \bar{\sigma}_n^2]} \mathbb{E}_{\sigma_n^2} \left[\min_{\sigma_n^2 \in [\underline{\sigma}_n^2, \bar{\sigma}_n^2]} -\log \mathbb{E}_{\sigma_n^2} [\exp(-A(W_0 + (d-p)\theta)) | n] \right]. \tag{2.15}
\end{aligned}$$

I once again utilize the fact that the $\log(\cdot)$ function is monotonic. Parallel to Eq. (2.11), the steps below simplify the optimization problem in Eq. (2.15) using the property of the log-normal distribution and the assumed risk-neutral nature of the investor:

$$\begin{aligned}
\max_{\theta} U_0 &= \max_{\theta} \min_{\sigma_n^2 \in [\underline{\sigma}_n^2, \bar{\sigma}_n^2]} \mathbb{E}_{\sigma_n^2} \left[\min_{\phi \in [\underline{\phi}, \bar{\phi}]} -\mathbb{E}_{\phi} [-A(W_0 + (d-p)\theta) | n] - \frac{1}{2} \mathbb{V}_{\phi} [-A(W_0 + (d-p)\theta) | n] \right] \\
&= \max_{\theta} \min_{\sigma_n^2 \in [\underline{\sigma}_n^2, \bar{\sigma}_n^2]} \mathbb{E}_{\sigma_n^2} \left[\min_{\phi \in [\underline{\phi}, \bar{\phi}]} \theta \mathbb{E}_{\phi} [d-p|n] - \frac{1}{2} A \theta^2 \mathbb{V}_{\phi} [d-p|n] \right] \\
&= \max_{\theta} \min_{\sigma_n^2 \in [\underline{\sigma}_n^2, \bar{\sigma}_n^2]} \mathbb{E}_{\sigma_n^2} \left[\min_{\phi \in [\underline{\phi}, \bar{\phi}]} \theta \mathbb{E}_{\phi} [d-p|n] \right]. \tag{2.16}
\end{aligned}$$

For notational purposes, I replace the σ_n^2 variables with the ϕ counterparts only in the inner minimization procedure. Solving Eq. (2.16) for the price at $t = 0$, we see that

$$\begin{aligned}
p_0 &= \min_{\sigma_n^2 \in [\underline{\sigma}_n^2, \bar{\sigma}_n^2]} \mathbb{E}_{\sigma_n^2} \left[\min_{\phi \in [\underline{\phi}, \bar{\phi}]} \mathbb{E}_{\phi} [d|n] \right] \\
&= \min_{\sigma_n^2 \in [\underline{\sigma}_n^2, \bar{\sigma}_n^2]} \mathbb{E}_{\sigma_n^2} [p_1(n)].
\end{aligned}$$

p_0 minimizes the average price at $t = 1$ with the appropriate choice of σ_n^2 . Plugging in for $p_1(n)$ using Eq. (2.12), we see that

$$\begin{aligned}
p_0 &= \min_{\sigma_n^2 \in [\underline{\sigma}_n^2, \bar{\sigma}_n^2]} \mathbb{E}_{\sigma_n^2} \left[\bar{d} + \underline{\phi} n | n > 0 \right] \times \mathbb{P}_{\sigma_n^2} [n > 0] + \mathbb{E}_{\sigma_n^2} \left[\bar{d} + \bar{\phi} n | n \leq 0 \right] \times \mathbb{P}_{\sigma_n^2} [n \leq 0] \\
&= \min_{\sigma_n^2 \in [\underline{\sigma}_n^2, \bar{\sigma}_n^2]} \bar{d} + \frac{1}{2} \underline{\phi} \mathbb{E}_{\sigma_n^2} [n | n > 0] + \frac{1}{2} \bar{\phi} \mathbb{E}_{\sigma_n^2} [n | n \leq 0] \\
&= \min_{\sigma_n^2 \in [\underline{\sigma}_n^2, \bar{\sigma}_n^2]} \bar{d} - \frac{1}{2} (\bar{\phi} - \underline{\phi}) \mathbb{E}_{\sigma_n^2} [n | n > 0] \\
&= \min_{\sigma_n^2 \in [\underline{\sigma}_n^2, \bar{\sigma}_n^2]} \bar{d} - \frac{1}{\sqrt{2\pi}} (\bar{\phi} - \underline{\phi}) \sqrt{\alpha^2 \sigma_d^2 + \sigma_n^2} \\
&= \bar{d} - \frac{1}{\sqrt{2\pi}} (\bar{\phi} - \underline{\phi}) \sqrt{\alpha^2 \sigma_d^2 + \bar{\sigma}_n^2}. \tag{2.17}
\end{aligned}$$

The first four steps in the derivation correspond exactly the four steps leading to the derivation of \bar{p}_1 in Eq. (2.13), with the only difference the $\min_{\sigma_n^2 \in [\underline{\sigma}_n^2, \bar{\sigma}_n^2]}$ operation. The final step substitutes in $\bar{\sigma}_n^2$ as the ϵ_n variance that minimizes the expression.

p_0 is equal to the mean dividend \bar{d} less a positive term that represents a no news is good news effect. The investor earns exactly this term from before the arrival of the news to after the arrival of the news when the news is equal to zero, or there is no news. In other words, this term measures the compensation that the investor receives for weathering ambiguity in the case when the news itself is not informative for the fundamental value of the risky asset. More precisely, the $t = 0$ investor knows that an ambiguous piece of news will arrive in one period and that he will interpret the news in an asymmetric fashion. As noted, this asymmetry implies that the price of the risky asset at $t = 1$ is a concave function of the news. By Jensen's inequality, the worst-case belief for the investor at $t = 0$ that minimizes the expected price of the risky asset at $t = 1$ is that the news yet to arrive has high variance. That is, the variance of ϵ_n is at the upper end of the range $\bar{\sigma}_n^2$. As the amount of ambiguity or $\bar{\phi} - \underline{\phi}$ increases, the no news is good news effect increases, and p_0 correspondingly decreases.

Having established the price p_0 before the news and the average price \bar{p}_1 after the news, I show that the price of the risky asset increases on average through the news; that is, the ambiguity premium $\bar{p}_1 - p_0 > 0$. The result is immediate from Eqs. (2.13) and (2.17):

$$\begin{aligned}\bar{p}_1 - p_0 &= \left(\bar{d} - \frac{1}{\sqrt{2\pi}} (\bar{\phi} - \underline{\phi}) \sqrt{\alpha^2 \sigma_d^2 + \bar{\sigma}_n^2} \right) - \left(\bar{d} - \frac{1}{\sqrt{2\pi}} (\bar{\phi} - \underline{\phi}) \sqrt{\alpha^2 \sigma_d^2 + \sigma_n^2} \right) \\ &= \frac{1}{\sqrt{2\pi}} (\bar{\phi} - \underline{\phi}) \left(\sqrt{\alpha^2 \sigma_d^2 + \bar{\sigma}_n^2} - \sqrt{\alpha^2 \sigma_d^2 + \sigma_n^2} \right).\end{aligned}\quad (2.18)$$

Since I have already assumed that $\bar{\phi} - \underline{\phi} > 0$, $\bar{p}_1 - p_0 > 0$ as long as $\bar{\sigma}_n^2 > \sigma_n^2$; that is, the true variance of ϵ_n is less than the upper range of the variance considered by the investor. The ambiguity premium depends on $\bar{\phi} - \underline{\phi}$ and is consequently increasing in the amount of ambiguity. There are two competing effects at work in $\bar{p}_1 - p_0$: the no news is good news effect and the asymmetry effect. As we have seen, the no news is good news effect decreases p_0 to compensate the investor for the news ambiguity. Upon the arrival of the news, however, the price of the risky asset decreases more for a negative piece of news than the price increases for a positive piece of news, and this asymmetric reaction has a negative effect on \bar{p}_1 . It turns out that the impact of the no news is good news effect is greater (in an

absolute sense) than that of the asymmetry effect, so the price of the risky asset rises on net through the arrival of the news. The reason is that the magnitude of the asymmetry effect is related to a Jensen's inequality effect for the average price of the risky asset at $t = 1$ using $\bar{\sigma}_n^2$ as the actual variance of ϵ_n . The magnitude of the no news is good news effect, however, is related to a Jensen's inequality effect using $\bar{\sigma}_n^2$ as the investor's belief for the variance of ϵ_n . Since $\bar{\sigma}_n^2 > \sigma_n^2$ by assumption, Jensen's inequality has a stronger influence on the no news is good news effect than on the asymmetry effect, so the ambiguity premium is positive.

2.2.3.3 Volatility of the Risky Asset with Ambiguity-Aversion

Using the expressions for the equilibrium prices of the risky asset, I show that the volatility of the change in the price of the risky asset through the news is increasing or decreasing in news ambiguity depending on the driver of the change in the amount of ambiguity. The variance of the change in the risky asset price from $t = 0$ to $t = 1$ is

$$\mathbb{V}_{\bar{\sigma}_n^2} [p_1(n) - p_0] = \mathbb{V}_{\bar{\sigma}_n^2} [p_1(n)] = \mathbb{E}_{\bar{\sigma}_n^2} [p_1^2(n)] - \bar{p}_1^2. \quad (2.19)$$

The first component of Eq. (2.19) is

$$\begin{aligned} \mathbb{E}_{\bar{\sigma}_n^2} [p_1^2(n)] &= \mathbb{E}_{\bar{\sigma}_n^2} \left[\left(\bar{d} + \underline{\phi}n \right)^2 | n > 0 \right] \times \mathbb{P}_{\bar{\sigma}_n^2} [n > 0] + \mathbb{E}_{\bar{\sigma}_n^2} \left[\left(\bar{d} + \bar{\phi}n \right)^2 | n \leq 0 \right] \times \mathbb{P}_{\bar{\sigma}_n^2} [n \leq 0] \\ &= \frac{1}{2} \times \mathbb{E}_{\bar{\sigma}_n^2} \left[\bar{d}^2 + \underline{\phi}^2 n^2 + 2\bar{d}\underline{\phi}n | n > 0 \right] + \frac{1}{2} \times \mathbb{E}_{\bar{\sigma}_n^2} \left[\bar{d}^2 + \bar{\phi}^2 n^2 + 2\bar{d}\bar{\phi}n | n \leq 0 \right] \\ &= \bar{d}^2 + \frac{1}{2} \left(\bar{\phi}^2 + \underline{\phi}^2 \right) \times \mathbb{E}_{\bar{\sigma}_n^2} [n^2 | n > 0] - \bar{d} \left(\bar{\phi} - \underline{\phi} \right) \times \mathbb{E}_{\bar{\sigma}_n^2} [n | n > 0] \\ &= \bar{d}^2 + \frac{1}{2} \left(\bar{\phi}^2 + \underline{\phi}^2 \right) \left(\alpha^2 \sigma_d^2 + \bar{\sigma}_n^2 \right) - \bar{d} \left(\bar{\phi} - \underline{\phi} \right) \frac{\sqrt{2}}{\sqrt{\pi}} \sqrt{\alpha^2 \sigma_d^2 + \bar{\sigma}_n^2}. \end{aligned}$$

The first equality substitutes in for $p_1(n)$ from Eq. (2.12), and the final equality uses the math formulas in Appendix B.2 that, for $z \sim N(0, \sigma_z^2)$, $\mathbb{E}[z | z > 0] = \sigma_z \sqrt{2} / \sqrt{\pi}$ and $\mathbb{E}[z^2 | z > 0] = \sigma_z^2$. The second component of Eq. (2.19) simply squares \bar{p}_1 , which we already know from Eq. (2.13):

$$\begin{aligned}\bar{p}_1^2 &= \left(\bar{d} - \frac{1}{\sqrt{2\pi}} (\bar{\phi} - \underline{\phi}) \sqrt{\alpha^2 \sigma_d^2 + \bar{\sigma}_n^2} \right)^2 \\ &= \bar{d}^2 + \frac{1}{2\pi} (\bar{\phi} - \underline{\phi})^2 (\alpha^2 \sigma_d^2 + \bar{\sigma}_n^2) - \bar{d} \frac{\sqrt{2}}{\sqrt{\pi}} (\bar{\phi} - \underline{\phi}) \sqrt{\alpha^2 \sigma_d^2 + \bar{\sigma}_n^2}.\end{aligned}$$

The variance of the risky asset price change is thus

$$\begin{aligned}\mathbb{V}_{\bar{\sigma}_n^2} [p_1(n) - p_0] &= \mathbb{E}_{\bar{\sigma}_n^2} [p_1^2(n)] - \bar{p}_1^2 \\ &= \frac{1}{2} (\alpha^2 \sigma_d^2 + \bar{\sigma}_n^2) \left(\bar{\phi}^2 + \underline{\phi}^2 - \frac{1}{\pi} (\bar{\phi} - \underline{\phi})^2 \right) \\ &= \frac{1}{2} (\alpha^2 \sigma_d^2 + \bar{\sigma}_n^2) \left(\left(1 - \frac{1}{\pi}\right) \bar{\phi}^2 + \left(1 - \frac{1}{\pi}\right) \underline{\phi}^2 + \frac{2}{\pi} \bar{\phi} \underline{\phi} \right).\end{aligned}\quad (2.20)$$

Simple comparative statics on the variance show that the risky asset volatility is increasing in $\bar{\phi}$ and $\underline{\phi}$:

$$\begin{aligned}\frac{\partial}{\partial \bar{\phi}} \mathbb{V}_{\bar{\sigma}_n^2} [p_1(n) - p_0] &= \frac{1}{2} (\alpha^2 \sigma_d^2 + \bar{\sigma}_n^2) \left(2 \left(1 - \frac{1}{\pi}\right) \bar{\phi} + \frac{2}{\pi} \underline{\phi} \right) > 0 \\ \frac{\partial}{\partial \underline{\phi}} \mathbb{V}_{\bar{\sigma}_n^2} [p_1(n) - p_0] &= \frac{1}{2} (\alpha^2 \sigma_d^2 + \bar{\sigma}_n^2) \left(2 \left(1 - \frac{1}{\pi}\right) \underline{\phi} + \frac{2}{\pi} \bar{\phi} \right) > 0\end{aligned}$$

If, on the one hand, an increase in $\bar{\phi}$ (decrease in $\bar{\sigma}_n^2$) drives an increase in ambiguity, the volatility of the risky asset increases. If, on the other hand, a decrease in $\underline{\phi}$ (increase in $\bar{\sigma}_n^2$) drives an increase in ambiguity, the volatility of the risky asset decreases. The reason for this result is that the volatility of the risky asset depends on how the risky asset responds to the news, and $\bar{\phi}$ and $\underline{\phi}$ control the sensitivity of the risky asset price to negative and positive news, respectively. A larger $\bar{\phi}$ means the asset responds more strongly to bad news, which increases volatility, whereas a smaller $\underline{\phi}$ means the asset responds less strongly to good news, which decreases volatility. Note that the variance in Eq. (2.20) is more sensitive to $\bar{\phi}$ than $\underline{\phi}$:

$$\frac{\partial}{\partial \bar{\phi}} \mathbb{V}_{\bar{\sigma}_n^2} [p_1(n) - p_0] > \frac{\partial}{\partial \underline{\phi}} \mathbb{V}_{\bar{\sigma}_n^2} [p_1(n) - p_0].$$

To see this, I simplify the left- and right-hand sides in the following steps and arrive at an obviously true inequality:

$$\begin{aligned} & 2 \left(1 - \frac{1}{\pi}\right) \bar{\phi} + \frac{2}{\pi} \underline{\phi} > 2 \left(1 - \frac{1}{\pi}\right) \underline{\phi} + \frac{2}{\pi} \bar{\phi} \\ \iff & \bar{\phi} - \frac{1}{\pi} (\bar{\phi} - \underline{\phi}) > \underline{\phi} - \frac{1}{\pi} (\underline{\phi} - \bar{\phi}) \\ \iff & \bar{\phi} - \underline{\phi} > \frac{2}{\pi} (\bar{\phi} - \underline{\phi}). \end{aligned}$$

Thus, even if an increase in ambiguity, as proxied by $\bar{\phi} - \underline{\phi}$, is driven equally by increases in $\bar{\phi}$ and decreases in $\underline{\phi}$, the volatility of the risky asset still increases.

2.2.4 Equilibrium Prices at $t = 1$ and $t = 0$ under an Unambiguous Benchmark

To arrive at the unambiguous benchmark, all we have to consider is the case of $\sigma_n^2 = \underline{\sigma}_n^2 = \bar{\sigma}_n^2$; or, equivalently, $\phi = \underline{\phi} = \bar{\phi}$. With an unambiguous piece of news, the investor has only one model in mind with which to interpret any given piece of news. Based on Eq. (2.12), the price at $t = 1$ in response to a piece of news n is

$$p_1(n) = \bar{d} + \phi n.$$

There is no more asymmetry in the reaction of the risky asset to the news: The price increases as much for a positive piece of news as it decreases for a negative piece of news. It follows that the average price of the risky asset after the news arrival is just the mean of the dividend:

$$\bar{p}_1 = \bar{d}.$$

From Eq. (2.17), we see that the price at $t = 0$ before the news arrival is also equal to the mean of the dividend:

$$p_0 = \bar{d}.$$

There is no longer a no news is good news effect, because the investor at $t = 0$ knows that the news to arrive in one period is unambiguous. Finally, it is obvious that the risky asset price on average does not change in response to the news, so there is no ambiguity premium:

$$\bar{p}_1 - p_0 = 0.$$

2.2.5 Summary of Model Predictions

To summarize, the model leads to the following four propositions, three of which are illustrated in Figure 2.1.

Proposition 1: The asymmetry effect *The price of the risky asset reacts asymmetrically to a piece of news about asset fundamentals (dividend payment). That is, the price decreases more for a given negative piece of news than the price increases for a positive piece of news of the same magnitude. Time-variation in the amount of news ambiguity drives time-variation in the extent of this asymmetry, with times of high ambiguity corresponding to times of large asymmetry.*

The investor is uncertain of how to interpret the news; in particular, how precise the news is about the fundamental value of the asset. Since the investor conducts worst-case scenario evaluation, he views positive news as being imprecise and negative news as being precise, which leads to the asymmetry effect. News ambiguity is defined as the range of precisions about the news that the investor considers, so greater ambiguity is equivalent to a greater range of precisions and a more pronounced asymmetry effect. Note that greater asymmetry could be due to either positive news being viewed as less relevant, negative news being viewed as more relevant, or both.

Figure 2.1 illustrates the asymmetry effect. Under ambiguity in Panel B, the solid blue line $p_1(n) - p_0$ has a gentler slope $\underline{\phi} < \bar{\phi}$ to the right of the y-axis (corresponding to positive news) and a steeper slope $\bar{\phi} > \underline{\phi}$ to the left of the y-axis (corresponding to negative news). The point of intersection of the solid blue line and the y-axis is equal to $p_1(0) - p_0$, the




Figure 2.1: Graphical Representation of Model Predictions

Notes: Each panel plots in the solid blue line $p_1(n) - p_0$, the change in price of the risky asset after the arrival of a piece of news n at $t = 1$ versus the prevailing price at $t = 0$, against the value of the news n . Panel A corresponds to the benchmark in which there is no news ambiguity. The solid blue line has slope $\bar{\phi} = \underline{\phi}$ to the left and right of the y-axis and intersects the origin. The asymmetry effect, no news is good news effect, and ambiguity premium are all equal to zero. Panel B corresponds to the case in which the news is ambiguous. The solid blue line has a steeper slope $\bar{\phi} > \underline{\phi}$ to the left of the y-axis and a gentler slope $\underline{\phi} < \bar{\phi}$ to the right of the y-axis. This line intersects the y-axis at $p_1(0) - p_0 > 0$, which measures the size of the no news is good news effect. A dashed blue line connects the solid blue line at symmetric points around the y-axis at $\bar{p}_1 - p_0 > 0$, which measures the size of the ambiguity premium. The distance between the intersections of the solid and dashed blue lines, $\bar{p}_1 - p_1(0) < 0$, measures the asymmetry effect.

price change of the risky asset for a neutral piece of news. Due to the concavity of the solid blue line, the dashed blue line, which connects the ends of the solid blue lines at symmetric points around the y-axis, intersects the y-axis at a lower value equal to $\bar{p}_1 - p_0$, the average price change of the risky asset through a piece of news. The gap between the two points of intersection is thus equal to $\bar{p}_1 - p_1(0)$, the asymmetry discount from Eq. (2.14). As a result of the risky asset reacting more strongly to negative pieces of news than positive ones, the average price of the risky asset decreases below the price obtained under neutral news. In contrast, Panel A of Figure 2.1 shows the unambiguous benchmark in which the solid blue line has equal slopes $\bar{\phi} = \underline{\phi}$ to the left and right of the y-axis. Without asymmetry, a dashed blue line (not shown) that connects the ends of the solid blue lines at symmetric points around the y-axis must intersect the y-axis at the same point as the solid blue line. There is thus no asymmetry discount.

Proposition 2: The no news is good news effect *The investor in the risky asset earns a positive return when no news (neutral or value zero) is released. This no news is good news effect is exactly equal to the increase in the price of the risky asset in response to a neutral piece of news. Time-variation in the amount of news ambiguity drives time-variation in the magnitude of the no news is good news effect, with times of high ambiguity corresponding to times of a large no news is good news effect.*

Before the arrival of the news, the investor knows that ambiguous information about the risky asset will arrive in the future. The arrival of a neutral piece of news contains no information for inferring the fundamental value of the asset but does resolve the ambiguity associated with the news. As a result, the price of the risky asset still rises in response to a neutral piece of news and rises more when there is more ambiguity.

Turning to Figure 2.1, each solid blue line intersects the y-axis at $p_1(0) - p_0$, which is the change in the price of the risky asset if the news is neutral. The distance between the point of intersection and the origin thus measures the size of the no news is good news

effect. In the case of ambiguity in Panel B, $p_1(0) - p_0 > 0$, so the no news is good news effect is positive, whereas with no ambiguity in Panel A, $p_1(0) - p_0 = 0$, and there is no no news is good news effect.

Proposition 3: The ambiguity premium *On average, the price of the risky asset increases from before the arrival of the news to after the arrival of the news. That is, the ambiguity premium is positive.*

The ambiguity premium has two components: the no news is good news effect and the asymmetry effect. The former leads to an increase in the price of the risky asset with the arrival of the news, but the latter on average leads to a decrease in the price of the asset after the news is revealed. On balance, the magnitude of the asymmetry effect is larger than that of the no news is good news effect, which leads to a positive ambiguity premium.

In Panel B of Figure 2.1, we have already graphically identified the no news is good news effect as the distance between the intersection of the solid blue line and the origin (equal to $p_1(0) - p_0$) and the asymmetry effect as the distance between the intersections of the solid and dashed blue lines (equal to $\bar{p}_1 - p_1(0)$). The ambiguity premium should then just be the distance between the intersection of the dashed blue line and the origin: $(p_1(0) - p_0) + (\bar{p}_1 - p_1(0)) = \bar{p}_1 - p_0$, the average price of the risky asset after the arrival of the news less the price before the arrival of the news. In the unambiguous benchmark of Panel A, the ambiguity premium is necessarily zero because the no news is good news effect and asymmetry effect are both zero.

Proposition 4: The risky asset volatility effect *The volatility of the change in the price of the risky asset from before the news to after the news is increasing in news ambiguity if the increase in ambiguity is due to the investor allowing for a smaller lower bound for the variance of the news noise. The volatility is decreasing in news ambiguity if the increase in ambiguity is due to the investor allowing for a larger upper bound for the variance of the news noise.*

A smaller lower bound for the variance of the news noise leads to a stronger reaction of the asset to bad news (since a smaller $\underline{\sigma}_s^2$ leads to a larger $\bar{\phi}$, which determines the sensitivity of the risky asset to negative news). A larger upper bound for the variance of the news noise leads to a weaker reaction of the asset to good news (since a larger $\bar{\sigma}_s^2$ leads to a smaller $\underline{\phi}$, which determines the sensitivity of the risky asset to positive news).

2.3 The Stock Market Response to Macroeconomic News

To demonstrate how the model with ambiguity-aversion can accurately describe real-world financial markets, I explore the dynamics of the aggregate stock market around macroeconomic news releases; i.e., news about inflation, output, or the labor market. The first section below introduces the sample of MNAs, and the second section discusses the data on the aggregate stock market.

2.3.1 Data on MNAs

Table 2.1 shows the full sample of MNAs considered in this paper. For each data release, the table presents the name, units, number of observations, start date, end date, frequency, government agency or private-sector firm responsible, and intraday announcement timestamp. In all, I analyze 18 economic announcements, which cover the lion's share of important MNAs. Data for the majority of the announcements begin in late-1997 and extend to early-2014, with the primary expectations being Existing Home Sales and Pending Home Sales, which start in 2005. All of the announcements occur once a month excluding Initial Jobless Claims, which is a weekly data release. The data in my sample are released at either 8:30 AM, 9:15 AM, or 10:00 AM ET.

While the macroeconomic announcements are officially produced and released by various government agencies (e.g., the Bureau of Labor Statistics), and private-sector firms (e.g., the National Association of Realtors), I collect the data from Bloomberg. Specifically, for each event, I download both the actual data release and the expected data release, the latter of

Table 2.1: Full Sample of MNAs

Event Name	Units	N	Start	End	Freq.	Source	Time (ET)	Sign
Change in Nonfarm Payrolls	Thousands	196	11/07/1997	03/07/2014	M	BLS	8:30 AM	1
Consumer Confidence	Index	197	11/25/1997	03/25/2014	M	CB	10:00 AM ^a	1
CPI	% ch. mom	197	11/18/1997	03/18/2014	M	BLS	8:30 AM	1
Durable Orders	% ch. mom	197	11/26/1997	03/26/2014	M	Census	8:30 AM	1
Existing Home Sales	Millions	109	03/23/2005	03/20/2014	M	NAR	10:00 AM	1
Factory Orders	% ch. mom	196	11/05/1997	03/06/2014	M	Census	10:00 AM	1
Housing Starts	Millions	189	03/17/1998	03/18/2014	M	Census	8:30 AM	1
Industrial Production	% ch. mom	194	11/17/1997	03/17/2014	M	Fed	9:15 AM	1
Initial Jobless Claims	Thousands	852	11/06/1997	03/27/2014	W	ETA	8:30 AM	-1
ISM Manufacturing	Index	197	11/03/1997	03/03/2014	M	ISM	10:00 AM ^b	1
ISM Non-Manufacturing	Index	182	01/06/1999	03/05/2014	M	ISM	10:00 AM	1
Leading Index	Index	197	11/04/1997	03/20/2014	M	CB	10:00 AM	1
New Home Sales	Thousands	195	12/03/1997	03/25/2014	M	Census	10:00 AM	1
Pending Home Sales	% ch. mom	107	06/01/2005	03/27/2014	M	NAR	10:00 AM	1
Personal Income	% ch. mom	196	11/03/1997	03/28/2014	M	BEA	8:30 AM	1
PPI	% ch. mom	193	12/12/1997	03/14/2014	M	BLS	8:30 AM	1
Retail Sales	% ch. mom	195	11/14/1997	03/13/2014	M	Census	8:30 AM	1
Unemployment Rate	% rate	196	11/07/1997	03/07/2014	M	BLS	8:30 AM	-1

Notes: Freq. refers to monthly (M) or weekly (W). Source uses the following acronyms: BEA, Bureau of Economic Analysis; BLS, Bureau of Labor Statistics; CB, Conference Board; Census, Census Bureau; ETA, Employment and Training Administration; Fed, Federal Reserve Board of Governors; ISM, Institute for Supply Management; and NAR, National Association of Realtors. Table above covers the whole period from November, 1997 to March, 2014.

^aOn 7/31/2001, data was released at 9:36 AM.

^bOn 8/1/2000, data was released at 9:26 AM.

which comes from a survey of economists. The difference between the actual number and the expected number constitutes the news that impacts the stock market.

The metric of the news surprise that I use in this paper is the conventional standardized news variable employed by Balduzzi, Elton, and Green (2001) as well as many subsequent papers in the literature of asset price reactions to MNAs. For a given economic indicator, the standardized news variable at time t is

$$S_t = \frac{A_t - \mathbb{E}_{t-} [A_t]}{\hat{\sigma}}. \quad (2.21)$$

A_t is the actual data, $\mathbb{E}_{t-} [A_t]$ is the expectation of the data from the Bloomberg survey, and $\hat{\sigma}$ is the sample standard deviation of $A_t - \mathbb{E}_{t-} [A_t]$. Consistent with the model, I define the variable such that a positive value corresponds to “better”-than-expected data; i.e., higher-than-expected inflation or output and lower-than-expected unemployment. For this reason, I multiply the Initial Jobless Claims and the Unemployment Rate data by -1 while keeping the sign of the other data unchanged, as shown by the sign column of Table 2.1. As the name suggests, the standardized news variable provides a single metric that is standardized across different types of news about macroeconomic fundamentals and allows for comparability. This is important because different types of news are released with different units.

The standardized news variable S_t in Eq. (2.21) is the empirical counterpart to the theoretical news n in Eq. (2.2) of the model. A positive (negative) news surprise with $S_t > 0$ ($S_t < 0$), like positive (negative) news $n > 0$ ($n < 0$), suggests a higher (lower) fundamental value of the stock market in the form of higher (lower) dividends. A neutral news release with $S_t = 0$, like neutral news $n = 0$, provides no new information on stock market fundamentals, since investors’ expectations are exactly met.

The use of the standardized news variable in conjunction with data from Bloomberg is common in testing the effect of macroeconomic news surprises on asset prices. Many financial market participants use Bloomberg to get a sense of the consensus forecast for any given MNA and to see how the actual data compares to the forecasted data at the time of

the release. The survey expectation is furthermore unbiased and unlikely to be stale, since economists can adjust their forecasts until the very last moment.

2.3.2 Data on High-Frequency Stock Prices

I use the front E-mini S&P 500 (ES) futures contract as the primary source for stock market price data and obtain intraday tick data from Tick Data, a data vendor. The data are available from late-1997 to mid-2014. High-frequency data are important because they allow me to consider the reaction of the stock market in a narrow window (± 5 minutes) around each macroeconomic news event. Stocks likely vary over the window only due to any surprise embedded in the news announcement. By extension, lower frequency data can be problematic because events unrelated to MNAs are more likely to influence prices. As an example of the importance of high-frequency data, the scatterplot in Figure 2.2 compares the 10-minute return of the stock market around a MNA (based on ES data) and the daily return of the stock market for the day of the MNA (based on S&P 500 index data). The correlation between the intraday return and the daily return is a positive 0.159, which suggests that the intraday response of stocks to macroeconomic news tends to influence the daily return of stocks in the expected manner. That is, when stocks do well (poorly) intraday in response to good (bad) news, the daily return also tends to be higher (lower). Nonetheless, the correlation is far from one, and the scatterplot clearly shows many instances in the 2nd and 4th quadrants in which stocks react positively (negatively) to macroeconomic news in the short interval around the news release but fall (rise) for the day. The annotated arrow in Figure 2.2 points to the example of the 10/16/2008 9:15 AM ET release of industrial production data. A worse-than-expected industrial production data release led to a -1.024% return of ES between 9:10 AM ET and 9:20 AM ET. For the day, however, the S&P 500 rallied 4.251% as a bounceback from the pummeling that stocks had taken earlier in the month as the 2008 financial crisis spread. The use of daily data in this instance and many others would lead to the erroneous observation that good (bad) macroeconomic news counterintuitively results in negative (positive) stock performance. I show in unreported robustness tests that

the use of daily price data on the stock market in the form of the S&P 500 index weakens the empirical results in this paper due to the lower resolution of the data.

While high-frequency data are available for instruments aside from futures, the use of futures data is important because many MNAs are released early in the morning before equity markets officially open at 9:30 AM ET. Whereas some markets (most relevantly, the traditional stock market) are less liquid at that early hour, futures markets are already active. Not surprisingly, financial market participants also tend to react to data surprises by trading in these futures directly.

2.4 Empirical Results: Whole Sample

The following two sections present empirical results in support of Propositions 1-3 using the whole sample. That is, I show that the asymmetry effect, no news is good news effect, and ambiguity premium exist using the entire sample period from November, 1997 to March, 2014.

2.4.1 The Asymmetry Effect and No News is Good News Effect

To assess the reaction of the stock market to macroeconomic news, I rely on the regression specification below:

$$R_t = \alpha + \beta^+ D_t^+ S_t + \beta^- (1 - D_t^+) S_t + \epsilon_t, \quad (2.22)$$

in which R_t is the % change in the price level (measured in bps) associated with the front ES futures contract in a ± 5 minute window around a given MNA, S_t is the standardized news for that data release defined in Eq. (2.21), and D_t is a dummy variable equal to one if $S_t > 0$ and zero otherwise. β^+ measures the reaction of the stock market to positive news surprises, β^- measures the reaction of the stock market to negative news surprises, and α measures the reaction of the stock market in the case of neutral news.

Table 2.2 shows the results from running Eq. (2.22) over the whole sample as defined in




Figure 2.2: Comparison of Intraday versus Daily Stock Market Data

Notes: The scatterplot compares the 10-minute return of the stock market in the ± 5 minutes around a MNA (horizontal axis) and the daily return of the stock market for the day of the MNA (vertical axis). Intraday returns are calculated as the % change in the price level associated with the front ES futures contract, and daily returns are calculated as the % change in the S&P 500 index. The sample plotted includes all MNAs from Table 2.1 excluding those that occur on days in which the S&P 500 was closed. I do not double count MNAs that occur at exactly the same time (e.g., 8:30 AM ET on the same day), but I treat MNAs that occur at different times on the same day as separate data points (associated with the same daily stock market return). The annotated arrow points to the 10/16/2008 9:15 AM ET release of industrial production data associated with a -1.024% 10-minute ES return and a 4.251% daily S&P 500 return. Figure above covers the whole period from November, 1997 to March, 2014. Lighter colors indicate earlier dates.

Table 2.2: The Whole Sample Stock Market Reaction to MNAs

Event Name	$\hat{\alpha}$ (bps)	$\hat{\beta}^+$ (bps)	$\hat{\beta}^-$ (bps)	Adj. R^2	N	$\hat{\beta}^+ - \hat{\beta}^-$ (bps)
Change in NFP	10.093**	-7.618	8.079	0.042	196	-15.697
Consumer Confidence	3.984	7.106*	23.105***	30.727	197	-15.999***
CPI	3.654	-11.889***	-5.830	11.397	197	-6.059
Durable Orders	1.475	4.335	10.287***	14.404	197	-5.952
Existing Home Sales	0.692	14.196**	11.820***	16.529	109	2.376
Factory Orders	1.294	3.519	5.249**	3.140	196	-1.730
Housing Starts	1.974	1.401	2.883	0.918	189	-1.481
Industrial Production	1.774	4.521	7.497*	21.998	194	-2.976
Initial Jobless Claims	2.487***	2.879***	8.715***	11.101	852	-5.836***
ISM Manufacturing	5.946*	9.593**	17.262***	15.582	197	-7.670
ISM Non-Manufacturing	0.379	10.209***	11.388***	14.783	182	-1.179
Leading Index	-4.025*	7.051*	-0.505	2.337	197	7.556
New Home Sales	6.783**	0.439	9.555**	3.578	195	-9.116*
Pending Home Sales	0.664	13.393***	9.535**	19.740	107	3.858
Personal Income	-0.016	0.955	0.361	-0.641	196	0.594
PPI	4.093*	-9.431***	5.177	6.463	193	-14.608***
Retail Sales	9.040***	1.597	18.229***	17.559	195	-16.632***
Unemployment Rate	5.206	2.767	10.411	0.604	196	-7.644
All	3.204***	2.416***	8.557***	4.701	3985	-6.141***

Notes: The table shows results from regressions of ES returns on standardized news for each event and then all events (last row): $R_t = \alpha + \beta^+ D_t^+ S_t + \beta^- (1 - D_t^+) S_t + e_t$. The left-hand-side variable R_t is the % change in the price level (bps) associated with the front ES futures contract in a ± 5 minute window around a MNA. The right-hand side variable S_t for the same MNA is the actual data less the expected data based on a Bloomberg survey of economists, all normalized by the sample standard deviation of the actual less the expected difference for that MNA: $S_t = (A_t - \mathbb{E}_{t-}[A_t]) / \hat{\sigma}$. D_t is a dummy variable equal to one if $S_t > 0$ and zero otherwise. t -statistics (not shown) are based on heteroskedasticity-consistent standard errors. Rightmost column presents $\hat{\beta}^+ - \hat{\beta}^-$ and performs a Wald test that the difference is not zero (F -statistics not shown). *** denotes significance at the 1% level, ** denotes significance at the 5% level, and * denotes significance at the 10% level. Table above covers the whole period from November, 1997 to March, 2014.

Table 2.1. Each row of Table 2.2 corresponds to a different regression with the specified event, and the last row aggregates all the MNAs together. If there is news ambiguity and investors are averse to this ambiguity, Propositions 1 and 2 imply the existence of the asymmetry effect and the no news is good news effect.

The asymmetry effect is that $\hat{\beta}^+ - \hat{\beta}^- < 0$, so that a unit of bad macroeconomic news lowers the stock market more than a unit of good macroeconomic news increases the stock market. We see from the last column of Table 2.2 that most of the $\hat{\beta}^+ - \hat{\beta}^-$ are negative, and of the numbers that are significantly different from zero based on a Wald test, all are negative. The last row says that $\hat{\beta}^+ - \hat{\beta}^- = -6.141$ bps, which has the economic interpretation that a one unit positive MNA surprise increases the stock market by 6.141 bps less than a one unit negative MNA surprise decreases the stock market. In particular, the positive surprise increases stocks by $\hat{\beta}^+ = 2.416$ bps, but the negative surprise decreases stocks by a larger $\hat{\beta}^- = 8.557$ bps.

The no news is good news effect is that $\hat{\alpha} > 0$, so that the stock market rises in response to neutral news as the arrival of the news resolves ambiguity. We see that most of the $\hat{\alpha}$ in the corresponding column of Table 2.2 are greater than zero, and all of the significant parameter estimates are greater than zero, with the exception of the weakly significant $\hat{\alpha}$ corresponding to Leading Index data releases. From the last row, $\hat{\alpha} = 3.204$ bps, so on average the stock market increases by 3.204 bps due to the resolution of ambiguous but neutral macroeconomic news.

To ensure that the covariate S_t used in the regressions is well-behaved, I plot the time series of standardized news in Figure 2.3 and calculate accompanying summary statistics in Panel A of Table 2.3. S_t behaves in the manner one would expect a news variable to behave in that it is unbiased and symmetric. In fact, MNA surprises look like normal shocks. Importantly, Figure 2.3 shows visual evidence that the distribution of standardized news is stable in the time series.




Figure 2.3: Time Series of Standardized News for the Whole Sample of MNAs

Notes: Figure plots standardized news S_t for all the MNAs, as defined in Table 2.2. Figure above covers the whole period from November, 1997 to March, 2014.

Table 2.3: Summary Statistics of MNA Surprises

Panel A: Whole Sample

	Mean	Median	Std. Dev.	Max.	Min.	Skewness	Kurtosis	<i>N</i>
Whole Sample	0.030	0.000	1.010	8.018	-5.514	0.181	2.973	3985

Panel B: Sample Bucketed into Variance Risk Premium Quintiles

VRP_t Quintile	Mean	Median	Std. Dev.	Max.	Min.	Skewness	Kurtosis	<i>N</i>
1	0.042	0.000	1.005	6.257	-4.171	0.254	2.786	796
2	0.003	0.000	0.962	4.866	-3.452	0.042	2.331	780
3	0.027	0.000	0.920	3.907	-4.315	0.137	1.616	815
4	0.034	0.000	0.966	5.214	-4.006	-0.113	1.872	781
5	0.044	0.000	1.176	8.018	-5.514	0.368	3.818	813

Notes: Panel A shows summary statistics for the standardized news variable S_t for the whole sample of MNAs. Panel B shows summary statistics for S_t bucketed according to VRP_t quintiles, with the 1st quintile corresponding to the smallest values and the 5th quintile corresponding to the largest values. S_t aggregates the news surprises of all MNAs, as defined in Table 2.2. VRP_t is the “Monthly Averages” variance risk premium series, which corresponds to the monthly time series in Panel B of Figure 2.4. Tables above cover the whole period from November, 1997 to March, 2014.

2.4.2 The Ambiguity Premium

Moving on to Proposition 3, I evaluate whether there is an ambiguity premium over the whole sample from November, 1997 to March, 2014. Under risk-neutrality, the model defines the ambiguity premium as the average return of the stock market around news. In reality, this average return also reflects various types of risk premia. Since I do not focus on these risk premia, I equate the ambiguity premium with the average return of the stock market around news and show that this quantity is positive and too large to be explained by conventional stories.

Panel A of Table 2.4 calculates the arithmetic sum of returns in the 10-minute window around MNAs in the first row (expressed in %) and the mean of returns in the bottom row (expressed in bps). Stocks increase a cumulative sum of 43.903% around MNAs or 1.102 bps on average per macroeconomic news release, which is suggestive evidence for a positive ambiguity premium. Panel A of Table 2.4 further decomposes the ambiguity premium

Table 2.4: The Ambiguity Premium and its Decomposition

Panel A: Whole Sample

		Return	No News	+ News	- News	Asymm.	N
Sum (%)	Whole Sample	43.903	127.688	37.011	-120.795	-83.784	3985
Mean (bps)	Whole Sample	1.102	3.204	0.929	-3.031	-2.102	3985

Panel B: Sample Bucketed into Variance Risk Premium Quintiles

	VRP _t Quintile	Return	No News	+ News	- News	Asymm.	N
Sum (%)	1	20.026	22.516	1.877	-4.368	-2.491	796
	2	8.237	11.848	12.024	-15.635	-3.612	780
	3	2.695	19.470	5.726	-22.502	-16.776	815
	4	7.256	36.758	-5.389	-24.113	-29.502	781
	5	5.690	38.489	18.186	-50.985	-32.798	813
Mean (bps)	1	2.516	2.829	0.236	-0.549	-0.313	796
	2	1.056	1.519	1.541	-2.005	-0.463	780
	3	0.331	2.389	0.703	-2.761	-2.058	815
	4	0.929	4.706	-0.690	-3.087	-3.777	781
	5	0.700	4.734	2.237	-6.271	-4.034	813

Notes: Panel A considers the whole sample of MNAs, and Panel B buckets the whole sample into variance risk premium quintiles as in Table 2.3. The “Sum (%)” row(s) of each panel present(s) arithmetic sum of returns in %. The “Mean (bps)” row(s) of each panel present(s) mean of returns in bps. “Return” is the actual return of the front ES contract around MNAs and is equal to the ambiguity premium. Other columns depend on the estimated $\hat{\alpha}$, $\hat{\beta}^+$, and $\hat{\beta}^-$ from Table 2.2 (Panel A) and Table 2.8 (Panel B). “No News” is that portion of the ambiguity premium attributed to the no news is good news effect and is equal to $\hat{\alpha} \times N$ for the “Sum (%)” row(s) and just $\hat{\alpha}$ for the “Mean (bps)” row(s). “+ News” is that portion of the ambiguity premium attributed to the reaction of the stock market to positive news and is equal to the sum of $\hat{\beta}^+ \times S_t$ for $S_t > 0$ for the “Sum (%)” row(s) and the same quantity divided by N for the “Mean (bps)” row(s). “- News” is that portion of the ambiguity premium attributed to the reaction of the stock market to negative news and is equal to the sum of $\hat{\beta}^- \times S_t$ for $S_t < 0$ for the “Sum (%)” row(s) and the same quantity divided by N for the “Mean (bps)” row(s). “Asymm.” is that portion of the ambiguity premium attributed to the asymmetric reaction of the stock market to positive news versus negative news and simply sums the “+ News” and “- News” columns. “Return” is equal to “No News” plus “Asymm.” N is the number of MNAs, and S_t aggregates the news surprises of all MNAs, as defined in Table 2.2. Tables above cover the whole period from November, 1997 to March, 2014.

into two components using the estimated $\hat{\alpha}$, $\hat{\beta}^+$, and $\hat{\beta}^-$ from the last row of Table 2.2: the positive no news is good news effect and the negative asymmetry effect. The “No News” column is that portion of the ambiguity premium attributed to the no news is good news effect and is equal to $\hat{\alpha} \times N$ for the top row, in which N is the number of MNAs, and just $\hat{\alpha}$ for the bottom row. The “+ News” column is that portion of the ambiguity premium attributed to the reaction of the stock market to positive news and is equal to the sum of $\hat{\beta}^+ \times S_t$ for $S_t > 0$ for the top row and the same quantity divided by N for the bottom row. The “- News” column is that portion of the ambiguity premium attributed to the reaction of the stock market to negative news and is equal to the sum of $\hat{\beta}^- \times S_t$ for $S_t < 0$ for the top row and the same quantity divided by N for the bottom row. The “Asymm.” column is that portion of the ambiguity premium attributed to the asymmetric reaction of the stock market to positive news versus negative news and simply sums the “+ News” and “- News” columns. As would be expected in a decomposition, the “Return” column is equal to the “No News” column plus the “Asymm.” column. In sum (on average), the 43.903% (1.102 bps) ambiguity premium can be attributed to a positive no news is good news effect of 127.688% (3.204 bps) which is greater in magnitude than the -83.784% (-2.102 bps) asymmetry effect.

While Panel A of Table 2.4 is useful for understanding how the ambiguity premium can be decomposed into its two components, the results double count some return intervals when multiple MNAs occur at the exact same time. I address this issue by presenting additional evidence for the ambiguity premium in Table 2.5, which calculates unique return intervals that include MNAs.⁴ In particular, I compute 10-minute return intervals for the whole November, 1997 to March, 2014 sample and split these intervals into those that contain MNAs (“MNA”) and those that do not (“nMNA”). The table presents summary statistics on the 3,194 “MNA” return intervals, 588,671 “nMNA” return intervals, and the difference between these two types of intervals.

Methodology-wise, returns are calculated as the % change in the price level (measured in bps) associated with the front ES futures contract. I calculate return intervals only for those

⁴All results in the paper are robust to dropping MNA observations that occur at the exact same time.

Table 2.5: Summary Statistics of 10-minute Stock Market Returns for Periods with and without MNAs

	MNA	nMNA	Diff.	MNA	nMNA	Diff.
	All Observations			Winsorized (1% and 99%)		
Total	3624.966	6626.671	-3001.705	3510.359	6757.771	-3247.412
Mean	1.000**	0.014	0.986**	0.983***	0.012	0.972***
<i>t</i> -statistic	(2.245)	(1.035)	(2.212)	(2.632)	(1.208)	(2.600)
1st Percentile	-76.543	-31.627	-44.915	-60.426	-23.354	-37.072
25th Percentile	-9.025	-2.208	-6.817	-8.819	-2.163	-6.657
Median	0.000	0.000	0.000	0.000	0.000	0.000
75th Percentile	10.549	2.246	8.303	10.271	2.210	8.061
99th Percentile	78.233	30.665	47.568	61.183	22.779	38.404
Std. Dev.	25.179	10.460		20.899	7.300	
Skewness	0.143	0.568		0.101	-0.071	
Kurtosis	5.826	68.635		2.107	3.625	
N	3194	588671		3130	576899	

Notes: Table splits 10-minute ES return intervals into those that contain MNAs (“MNA”) and those that do not (“nMNA”) and presents summary statistics. Returns are calculated as the % change in the price level (bps) associated with the front ES futures contract. I calculate ES return intervals only on those days that the S&P 500 was open as well as the following dates on which the S&P 500 was closed but futures markets were open: 4/2/1999, 4/6/2007, 4/2/2010, and 4/6/2012 (all corresponding to Good Friday); and 10/29/2012 (corresponding to Hurricane Sandy). For each day, I forward fill in price data for all 1440 minutes and attempt to split the day into 144 10-minute increments. For any MNA, I ensure that there is a corresponding 10-minute interval formed from the ± 5 minutes window around the MNA, consistent with the event window used elsewhere in the paper. If doing so prevents the formation of 144 evenly spaced 10-minute intervals for that day, I allow the intervals surrounding the interval containing the MNA to be slightly longer. For example, if the MNA occurs at 8:30 AM, the interval that contains the MNA is from 8:25 AM to 8:35 AM. The preceding interval is from 8:10 AM to 8:25 AM, and the following interval is from 8:35 AM to 8:50 AM; both surrounding intervals cover 15 minutes instead of 10 minutes. The rest of the intervals for that day are then exactly 10 minutes. The “Total” row shows returns compounded over all intervals, expressed in bps. Leftmost three columns include the whole sample, and rightmost three columns winsorize returns at the 1% and 99%. *** denotes significance at the 1% level, ** denotes significance at the 5% level, and * denotes significance at the 10% level. Table above covers the whole period from November, 1997 to March, 2014.

days that the S&P 500 was open as well as on the following dates on which the S&P 500 was closed but futures markets were open: 4/2/1999, 4/6/2007, 4/2/2010, and 4/6/2012 (all corresponding to Good Friday); and 10/29/2012 (corresponding to Hurricane Sandy). For each day, I forward fill in price data for all 1440 minutes and attempt to split the day into 144 10-minute increments. For any MNA, I ensure that there is a corresponding 10-minute interval formed from the ± 5 window around the MNA, consistent with the event window used elsewhere in the paper. If doing so prevents the formation of 144 evenly spaced 10-minute intervals for that day, I allow the intervals surrounding the interval containing the MNA to be slightly longer. For example, if the MNA occurs at 8:30 AM, the interval that contains the MNA is from 8:25 AM to 8:35 AM. The preceding interval is from 8:10 AM to 8:25 AM, and the following interval is from 8:35 AM to 8:50 AM; both surrounding intervals cover 15 minutes instead of 10 minutes. The rest of the intervals for that day are then exactly 10 minutes.

Focusing on the first three columns of Table 2.5 and the “Mean” row, we see that the stock market rises on average a statistically significant 1 bp in the 10-minute window around MNAs, which is consistent with Proposition 3. Interestingly, however, the stock market, rises a statistically insignificant 0.014 bps in all other 10-minute windows. The difference in means of 0.986 bps is statistically significant. The “Total” row shows returns compounded over all intervals, expressed in bps. Over the whole sample period, the stock market rises a compounded 36.250% over the 10-minute windows containing MNAs and 66.267% in all other 10-minute windows.⁵ That is, approximately 1/3 of the stock market return from November, 1997 to March, 2014 accrues in the tiny fraction of time around macroeconomic news releases. Despite the much higher mean of “MNA” return intervals versus “nMNA” return intervals, the standard deviation is only about 2.5 times greater in the former (25.179 bps) than the latter (10.460 bps). The Sharpe ratio of “MNA” intervals is consequently much higher than that of “nMNA” intervals. As the rightmost three columns demonstrate, the

⁵This total return differs from that reported in Panel A of Table 2.4 for two reasons. First, as already noted, Panel A of Table 2.4 double counts some return intervals. Second, Panel A of Table 2.4 presents the arithmetic sum of returns as opposed to the compounded total return.

results are robust to winsorizing “MNA” and “nMNA” return intervals separately at the 1% and 99%.

Note that the above methodology ignores certain days in which there are ES futures price data (for example, Sundays) and thus understates the number of “nMNA” return intervals. Since these days tend to have more stale prices, forward filling the price data to compute return intervals would create additional “nMNA” return intervals with zero returns. Including these extra days would thus lower both the mean return and standard deviation of “nMNA” intervals. The impact on results, however, is unlikely to be significant.

2.5 State Variable for Ambiguity

Aside from stating the existence of the asymmetry effect, no news is good news effect, and ambiguity premium, Propositions 1-3 have the more nuanced prediction that the magnitudes of these behaviors should be larger at times of greater ambiguity. Similarly, Proposition 4 predicts that the volatility of the stock market around MNAs depends on the amount of ambiguity. It is consequently vital to have a state variable for ambiguity in order to fully test the model. Below, I construct such a proxy for ambiguity, the variance risk premium, explore some of its properties for the predictability of stock returns, and discuss alternate state variables for ambiguity.

2.5.1 Variance Risk Premium

The variance risk premium $VRP_{t,T}$ is the difference between risk-neutral (Q-measure) and physical (P-measure) expectations of stock market return variance over a given horizon $RV_{t+1,T}$:

$$VRP_{t,T} = \mathbb{E}_t^Q [RV_{t+1,T}] - \mathbb{E}_t^P [RV_{t+1,T}]. \quad (2.23)$$

Expectations are taken at time t for realized variance from the period $t + 1$ to T .

One way to interpret the variance risk premium is that it measures the premium

embedded in an option on the equity market. The premium exists because investors not only care about uncertainty of the stock market return, as embodied by the return variance, but also uncertainty about the return variance itself. An option on the stock market hedges against high realized variance and thus demands a positive variance risk premium for this insurance.

Alternatively, consider a variance swap, an over-the-counter derivative. An investor who is long the swap receives the realized variance over the maturity of the swap less a fixed quantity, the variance swap rate. It is costless to enter such a swap, so by no arbitrage, the variance swap rate has to equal $\mathbb{E}_t^Q [RV_{t+1,T}]$, the first component of the variance risk premium. By being long the variance swap, an investor protects himself against high realized variance. The price of this protection is simply the variance swap rate less the true/statistical expectation of variance or the variance risk premium.

As previewed, I use the variance risk premium as a proxy for ambiguity. The motivation for this choice of state variable is Drechsler (2013). In that paper, a representative investor has a range of models in mind about the dynamics of economic fundamentals (e.g., the frequency and magnitude of jump shocks to expected growth and growth volatility processes for consumption and dividends). That is, the investor is uncertain about the true model governing economic fundamentals. Because the investor is ambiguity-averse, he evaluates decisions under the worst-case model, a model in which realized variance is high and thus options have high payoffs. Options hedge the investor's model uncertainty, which results in a price premium, the variance risk premium. Drechsler (2013) shows that the size of the variance risk premium is directly linked to ambiguity in the model.

I construct a daily empirical measure of the variance risk premium in the time series for a one-month horizon (the interval between $t + 1$ and T is one month). For the sake of notation, I set $T = t + 22$ and assume that there are 22 trading days in a month. Britten-Jones and Neuberger (2000) and Carr and Wu (2009), among others, have shown that, in continuous time, the risk-neutral expectation of stock market realized variance is equal to the value of a portfolio of European options on the stock market. A natural proxy for $\mathbb{E}_t^Q [RV_{t+1,t+22}]$ is

then the daily Chicago Board Options Exchange (CBOE) VIX index which is a risk-neutral expectation of S&P 500 variance over 30 calendar days or about 22 trading days, consistent with my assumption. The VIX produces this measure of implied volatility in a model-free manner (without relying on an option-pricing model) by calculating a weighted average of a portfolio of S&P 500 calls and puts. Since the VIX is quoted as an annual volatility variable, I create a squared VIX variable $VIX_t^2/12$ that simply squares the day-end VIX and divides by 12. This squared VIX quantity is a monthly variance variable that I use to proxy for $\mathbb{E}_t^Q [RV_{t+1,t+22}]$.

To create a daily series for $\mathbb{E}_t^P [RV_{t+1,t+22}]$, I first calculate a daily time series of monthly realized variances $RV_{t+1,t+22}$. For a given t , $RV_{t+1,t+22}$ is the sum of the 22 daily realized variances between the two dates (inclusive). Daily realized variance is the sum of squared five-minute log returns on ES futures from 9:30 AM ET to 4:00 PM ET and the squared close-to-open log return. Calculating realized variance by summing high-frequency squared returns is a well-established procedure. Papers as early as French, Schwert, and Stambaugh (1987) and Schwert (1989) estimate monthly realized variance by summing daily squared returns. The literature has subsequently provided formal justification for this practice; e.g., Andersen, Bollerslev, Diebold, and Ebens (2001). Sampling returns at a higher frequency increases the accuracy of estimated realized variance at a lower frequency. Sampling returns at too-high of a frequency, however, can be problematic due to market microstructure issues such as the presence of the bid-ask spread, so a five-minute frequency is a reasonable compromise.

With a daily series for $RV_{t+1,t+22}$ in hand, I construct $\mathbb{E}_t^P [RV_{t+1,t+22}]$ using the standard variance forecasting technique of projecting realized variance onto variables in a lagged information set. That is, I run a daily time series regression of $RV_{t+1,t+22}$ on one-month lagged realized variance $R_{t-21,t}$ and one-month lagged squared VIX, $VIX_{t-22}^2/12$. The regression covers the period from late-1997 (when my sample of intraday ES futures data for constructing $RV_{t+1,t+22}$ begins) to mid-2014. The regression with heteroskedasticity-

consistent standard errors is below:

$$RV_{t+1,t+22} = \underset{(1.145)}{2.130} + \underset{(0.047)}{0.310} \times R_{t-21,t} + \underset{(0.050)}{0.464} \times VIX_{t-22}^2/12 + e_t. \quad (2.24)$$

I subsequently construct $\mathbb{E}_t^{\mathbb{P}} [RV_{t+1,t+22}]$ from the one-step-ahead forecasts of this regression. Several considerations motivate the specification of Eq. (2.24). First, the use of $R_{t-21,t}$ and $VIX_{t-22}^2/12$ as covariates stems from the persistence of realized variance and evidence in the literature that implied volatility has bearing on future stock market variance. These covariates do a good job of forecasting future realized variance as evidenced by the high adjusted $R^2 = 0.519$. Second, Bekaert, Hoerova, and Lo Duca (2013) and Bekaert and Hoerova (2014) have tested a number of different regression specifications for forecasting realized variance. They find that the simple model in Eq. (2.24) does well on a number of criteria such as subsample coefficient stability and out-of-sample root-mean-squared error.

Differencing $VIX_t^2/12$ and $\mathbb{E}_t^{\mathbb{P}} [RV_{t+1,t+22}]$ yields a daily time series for the variance risk premium $VRP_{t,t+22}$. The first set of rows indexed by “Daily” in Table 2.6 and Panel A of Figure 2.4 show summary statistics and plot the three series, respectively, from November, 1997 to March, 2014. $VRP_{t,t+22}$ is generally positive, persistent, right-skewed, and has significant excess kurtosis. We expect the variance risk premium to be positive because it represents a premium paid to hedge realized variance, and it appears that approximately a quarter of the risk-neutral expectation of realized variance is attributed to the premium. Panel A of Figure 2.4 provides more color on the non-normality of the variance risk premium as well as that of the squared VIX and the physical expectation of realized variance. We see that all three time series spike dramatically at events such as the 2008 financial crisis.

The daily time series of the variance risk premium has two drawbacks as a proxy for ambiguity. First, while the variance risk premium is mostly positive, it clearly goes negative at times; e.g., the premium obtains a minimum value of -64.363 on 11/4/2008. These negative values violate the intuition and theoretical restriction that the premium should be positive. Relatedly, while the time series fluctuates in a well-behaved manner most of the time, the variance risk premium can fluctuate wildly at times; e.g., just a few

Table 2.6: Summary Statistics of Variance Risk Premium and Components

	Mean	Median	Std. Dev.	Max.	Min.	Skewness	Kurtosis	AC(1)
Daily								
$VIX_t^2/12$	46.163	34.987	45.872	544.862	8.151	4.245	27.277	0.969
$\mathbb{E}_t^{\mathbb{P}} [RV_{t+1,t+22}]$	34.261	25.048	35.756	400.182	7.351	4.582	29.392	0.988
$VRP_{t+1,t+22}$	11.902	8.721	12.665	179.842	-64.363	3.630	24.995	0.883
Monthly Averages								
$VIX_t^2/12$	46.204	35.780	43.431	332.292	9.781	3.655	18.578	0.838
$\mathbb{E}_t^{\mathbb{P}} [RV_{t+1,t+22}]$	34.264	25.248	34.526	281.605	8.474	4.261	24.549	0.822
$VRP_{t+1,t+22}$	11.941	9.631	10.076	56.722	0.757	2.105	5.796	0.803
Monthly End-of-Month								
$VIX_t^2/12$	45.241	37.066	39.436	298.901	9.048	3.069	13.220	0.795
$\mathbb{E}_t^{\mathbb{P}} [RV_{t+1,t+22}]$	32.562	24.514	32.931	304.136	6.471	4.243	26.996	0.771
$VRP_{t+1,t+22}$	12.679	9.740	10.223	59.564	-5.235	2.152	6.145	0.437

Notes: $VIX_t^2/12$ is the squared VIX , $\mathbb{E}_t^{\mathbb{P}} [RV_{t+1,t+22}]$ is the physical expectation of realized variance, and $VRP_{t+1,t+22}$ is the variance risk premium. The rows indexed by "Daily" correspond to the daily time series in Panel A of Figure 2.4, the rows indexed by "Monthly Averages" correspond to the monthly time series in Panel B of Figure 2.4, and the rows indexed by "Monthly End-of-Month" correspond to the monthly time series in Panel C of Figure 2.4. Refer to the Figure 2.4 caption for details. Table above covers the whole period from November, 1997 to March, 2014.






Figure 2.4: Time Series of Variance Risk Premium and Components

Notes: See Appendix B.3.1. Panels A, B, and C from top to bottom correspond to "Daily," "Monthly Averages," and "Monthly End-of-Month" time series, respectively.

days after hitting its minimum value, the premium obtains a maximum value of 179.842 on 11/20/2008. It is unreasonable to infer that, during a period of high ambiguity such as the 2008 financial crisis, investors faced the lowest amount of ambiguity on record on 11/4/2008 and subsequently faced the highest amount of ambiguity on record just two weeks later on 11/20/2008. Instead, the fact that both drawbacks of the daily time series present themselves at times when the variance risk premium is elevated suggests a more straightforward explanation. As surmised by Bekaert and Hoerova (2014), crisis periods lead to spikes in realized variance which in turn affect the physical expectation of realized variance through Eq. (2.24). Fluctuations in the physical expectation then drive fluctuations in the variance risk premium and result in the premium sometimes turning negative. The time series model for forecasting variance in Eq. (2.24) is likely too simple to capture the reality that realized variance has different components, some of which should be allowed to mean-revert more quickly and therefore not affect the physical expectation of realized variance as much. To my knowledge, the literature has not established more sophisticated econometric models that adequately address the aforementioned drawbacks.

Instead of attempting to fix the daily time series, I instead construct monthly time series for the variance risk premium as well as its two components by taking a simple average of corresponding daily data in each month. Smoothing the data in this manner results in a well-behaved estimate of the variance risk premium at the monthly frequency. The “Monthly Averages” index in Table 2.6 and Panel B of Figure 2.4 show summary statistics and plot the three series, respectively, from November, 1997 to March, 2014. The monthly $VRP_{t,t+22}$ has similar properties to the daily $VRP_{t,t+22}$, as both are persistent and display non-normal properties due to large moves during times of crisis. The key difference and major improvement is that the monthly variance risk premium series is never negative and fluctuates in a reasonable manner; that is, the two drawbacks associated with the daily series no longer present problems. As such, I use this monthly time series of the variance risk premium as my main proxy for ambiguity.

Figure 2.5 magnifies the monthly variance risk premium series from Panel B of Figure




Figure 2.5: Monthly Variance Risk Premium with Annotations

Notes: The figure graphs the variance risk premium from Panel B of Figure 2.4 and annotates major peaks and troughs. Figure above covers the whole period from November, 1997 to March, 2014.

2.4 and annotates the plot based on identifiable events. Peaks in the time series correspond to various crises: LTCM/Russian financial crisis, September 11, the corporate scandals of 2002, the 2008 financial crisis, and the European sovereign debt and U.S. debt ceiling crises. Troughs correspond to times of relative economic and financial stability such as during the mid-2000s. This time series behavior of the variance risk premium suggests that the premium is an appropriate proxy for ambiguity. That is, it is reasonable to think that during crisis periods, investors are highly unsure of how informative macroeconomic news is for the stock market; the news could be very relevant, very irrelevant, or somewhere in the middle. On the flipside, during placid times, investors are more certain of how to interpret the significance of macroeconomic news for stocks.

My monthly variance risk premium series is a more reasonable proxy for ambiguity than estimates of the premium on the stock market in the literature; e.g., Carr and Wu (2009), Bollerslev, Tauchen, and Zhou (2009), Drechsler (2013), and Bekaert and Hoerova (2014). The reason is that the estimates in the literature suffer from the drawbacks of my daily variance risk premium metric discussed above: The premium is at times highly volatile and liable to turn negative, especially during crisis periods. Consider the following result representative of variance risk premium estimates in the literature. I directly estimate a monthly time series by first using end-of-month squared VIX VIX_t^2 to proxy for $\mathbb{E}_t^Q [RV_{t+1,t+22}]$, with t the last trading day of a given month. Note that I keep all notation the same as before, but obviously some months have more trading days than other months. Monthly realized variance $RV_{t+1,t+22}$ is the sum of daily realized variances for the month that includes days $t + 1$ to $t + 22$. Analogous to Eq. (2.24), I estimate $\mathbb{E}_t^P [RV_{t+1,t+22}]$ from one-step-ahead forecasts of the monthly regression of realized variance on one-month lagged realized variance and one-month lagged squared VIX:

$$RV_{t+1,t+22} = \underset{(2.762)}{0.195} + \underset{(0.144)}{0.282} \times R_{t-21,t} + \underset{(0.098)}{0.511} \times VIX_{t-22}^2 / 12 + e_t \quad (2.25)$$

Though the adjusted $R^2 = 0.437$ of Eq. (2.25) is high, it is substantially lower than that of the daily regression in Eq. (2.24), which suggests that the daily regression provides substantially

more statistical power than the monthly regression. The monthly variance risk premium $VRP_{t,t+22}$, for the month with t as its last trading day, is still $VIX_t^2/12 - \mathbb{E}_t^{\mathbb{P}} [RV_{t+1,t+22}]$. The third set of rows indexed by “Monthly End-of-Month” in Table 2.6 and Panel C of Figure 2.4 show summary statistics and plot the three series, respectively, from November, 1997 to March, 2014. We see that the variance risk premium once again obtains a negative value some of the time; e.g., the premium is at its lowest value of -5.235 for 10/2008. The monthly estimate is also unreasonably volatile: Just one month later, the the premium is near its maximum value at 48.858 for 11/2008. As a proxy for ambiguity, this time series says that ambiguity was at an all-time low in one month of the 2008 financial crisis and at an all-time high in the next month of the crisis. This behavior is problematic and justifies my construction of a different monthly series for the variance risk premium.

2.5.2 Does the Variance Risk Premium Predict Stock Returns?

To illustrate the impact of methodology choice in constructing the variance risk premium, I explore how the premium predicts stock returns in the time series. A number of papers in the literature including Bollerslev, Tauchen, and Zhou (2009), Faust and Wright (2012), Drechsler (2013), and Bekaert and Hoerova (2014) have found evidence suggesting that the variance risk premium predicts stock returns even when controlling for traditional predictor variables such as the price-to-earnings or dividend-to-price ratios. Does the “Monthly Averages” variance risk premium that I construct and argue is a better proxy for ambiguity than existing variance risk premium measures in the literature also predict stock returns? To address this question, I run basic predictability regressions using the “Monthly Averages” variance risk premium as a predictor variable and compare the results to the same regressions using the “Monthly End-of-Month” variance risk premium, which suffers similar drawbacks to extant series in the literature.

The return predictability regressions, using monthly data from November, 1997 to March, 2014, are of the form

$$r_{t,t+h} = a + bX_t + \epsilon_{t,t+h}, \quad (2.26)$$

in which $r_{t,t+h}$ is the continuously compounded return of the stock market over a h -month horizon from the end of month t to the end of month $t + h$, and X_t is the set of predictor covariates known at month t . I calculate $r_{t,t+h}$ as the annualized log return of the front ES futures contract from 4:00 PM ET of the last trading day (in which the S&P 500 is open) of month t to the same time of the last trading day of month $t + h$. Most predictability regressions in the literature use excess returns of spot stock prices over a measure of the risk-free rate on the left-hand-side of Eq. (2.26). Assuming no-arbitrage between spot and future prices and ignoring the distinction between futures and forwards, the return calculated using futures data is equivalent to the excess return calculated using spot data.⁶ To adjust for serial correlation from overlapping stock returns, I use $\max\{3, 2h\}$ Newey-West lags as in Bekaert and Hoerova (2014).

Table 2.7 shows \hat{b} , associated t -statistics, and adjusted R^2 for varying horizons h and 5 different combinations of covariates X_t . The first set of rows corresponds to $X_t = \{\log(P/E)_t\}$, the log of the cyclically adjusted price-to-earnings ratio obtained from Robert Shiller's website. In this benchmark regression, the price-to-earnings ratio negatively predicts stock returns, though only starting at $h = 12$ months are the coefficient estimates statistically significant. Note that the adjusted R^2 and magnitude of the parameter estimate increase with h . Boudoukh, Richardson, and Whitelaw (2008) point out that in return predictability regressions with persistent regressors and overlapping returns, even under the null hypothesis of no predictability, it is possible to observe this increasing pattern for the adjusted R^2 and magnitude of the parameter estimate. Thus it is important to be careful in interpreting the results of these and subsequent predictability regressions.

In the second set of rows, I set $X_t = \{VRP_t\}$, the "Monthly Averages" variance risk premium. From $h = 6$ on, the "Monthly Averages" variance risk premium is significant in positively predicting stock returns. In the third set of rows, I include both the "Monthly

⁶Suppose the stock market is worth S_t at t and S_{t+h} at $t + h$, and r is the continuously compounded risk-free rate. Assuming no dividends, the respective forward prices to be paid at time T are $S_t e^{r(T-t)}$ and $S_{t+h} e^{r(T-t-h)}$. The log return calculated from the two forward prices is $\ln(S_{t+h} e^{r(T-t-h)} / S_t e^{r(T-t)}) = \ln(S_{t+h}/S_t) - rh$, which is identical to the excess return.

Table 2.7: Predictability of Stock Returns

Horizon	1	3	6	9	12	15	18	21	24
$\log(P/E)_t$	-12.980 (-0.656)	-14.716 (-0.823)	-20.076 (-1.362)	-22.229 (-1.568)	-22.729* (-1.698)	-24.159** (-2.023)	-25.391** (-2.429)	-26.465*** (-2.947)	-26.897*** (-3.585)
Adj. R^2	0.156	0.731	3.572	6.783	9.482	13.452	17.953	22.866	27.506
VRP_t	0.183 (0.299)	0.385 (0.861)	0.526** (2.298)	0.490** (2.202)	0.440** (2.248)	0.396** (2.358)	0.409*** (2.848)	0.450*** (3.453)	0.503*** (4.040)
Adj. R^2	0.404	0.785	3.736	4.807	5.075	5.034	6.530	9.398	13.882
$\log(P/E)_t$	-11.905 (-0.686)	-11.316 (-0.675)	-15.418 (-0.971)	-18.210 (-1.147)	-19.360 (-1.265)	-21.497 (-1.547)	-22.686* (-1.866)	-23.300** (-2.238)	-22.969** (-2.562)
VRP_t	0.095 (0.164)	0.301 (0.712)	0.411 (1.549)	0.351 (1.313)	0.289 (1.160)	0.223 (0.983)	0.219 (1.110)	0.247 (1.444)	0.295* (1.814)
Adj. R^2	0.645	0.951	5.453	8.804	11.206	14.586	19.332	25.108	31.512
VRP_t	0.915** (2.237)	0.819*** (3.289)	0.701*** (3.415)	0.528** (2.600)	0.461** (2.531)	0.406** (2.520)	0.390*** (2.911)	0.412*** (3.096)	0.440*** (3.462)
Adj. R^2	2.266	5.530	7.251	5.880	5.836	5.557	6.162	8.205	11.091
$\log(P/E)_t$	-6.655 (-0.344)	-9.156 (-0.533)	-15.539 (-1.086)	-19.031 (-1.333)	-20.018 (-1.466)	-21.881* (-1.746)	-23.239** (-2.125)	-24.110** (-2.558)	-24.262*** (-3.052)
VRP_t	0.885** (2.278)	0.778*** (3.307)	0.629*** (2.991)	0.439** (2.016)	0.365* (1.919)	0.298 (1.617)	0.269* (1.762)	0.281** (2.037)	0.305** (2.341)
Adj. R^2	1.853	5.511	9.148	10.584	12.862	16.147	20.582	26.352	32.452

Monthly Averages
Monthly
End-of-Month
Monthly End-of-Month
Monthly End-of-Month

Notes: The table shows results from return predictability regressions using monthly data from November, 1997 to March, 2014: $r_{t,t+h} = a + bX_t + \epsilon_{t,t+h}$. $r_{t,t+h}$ is the annualized log return of the front ES futures contract from 4:00 PM ET of the last trading day (in which the S&P 500 is open) of month t to the same time of the last trading day of month $t+h$. Table displays results for horizons $h = 1, 3, 6, 9, 12, 15, 18, 21, 24$ and five sets of covariates X_t drawn from $\log(P/E)_t$, the log of the cyclically adjusted price-to-earnings ratio from Robert Shiller's website, the "Monthly Averages" variance risk premium series from Panel B of Figure 2.4, and the "Monthly End-of-Month" variance risk premium series from Panel C of Figure 2.4. Table reports estimated parameters \hat{b} , associated t -statistics, and adjusted R^2 . t -statistics (shown in parentheses) are based on Newey-West standard errors with $\max\{3, 2h\}$ lags. *** denotes significance at the 1% level, ** denotes significance at the 5% level, and * denotes significance at the 10% level.

Averages” variance risk premium and the price-to-earnings ratio; $X_t = \{\log(P/E)_t, VRP_t\}$. Neither variable is significant at short horizons, with $\log(P/E)_t$ first significant at $h = 18$ and VRP_t first significant at $h = 24$. Comparing the adjusted R^2 of the regression with both predictors to those with an individual predictor, we see that there are significant improvements in the adjusted R^2 for nearly all horizons. For example, at $h = 6$, the regression with $X_t = \{\log(P/E)_t, VRP_t\}$ has an adjusted R^2 of 5.453 versus an adjusted R^2 of 3.572 for $X_t = \{\log(P/E)_t\}$ and 3.736 for $X_t = \{VRP_t\}$. The “Monthly Averages” variance risk premium that I argue is a state variable for ambiguity thus has some predictive power for stock returns, especially at longer horizons.

How does the “Monthly End-of-Month” variance risk premium compare in predicting stock returns? The fourth and fifth set of rows in Table 2.7 are analogous to the second and third set of rows, respectively, with the exception of using the “Monthly End-of-Month” variance risk premium instead of the “Monthly Averages” one. With only $X_t = \{VRP_t\}$ in the fourth set of rows, we see that the “Monthly End-of-Month” variance risk premium is a strong predictor of stock returns at all horizons. The fifth set of rows includes the price-to-earnings ratio for the regression with $X_t = \{\log(P/E)_t, VRP_t\}$. $\log(P/E)_t$ is significant only at long horizons, but VRP_t is highly significant at short horizons, less significant at intermediate horizons, and significant again at long horizons. From an adjusted R^2 perspective, the regressions with both covariates are significantly better than the corresponding regressions with only one of the covariates.

The strong predictive power of the “Monthly End-of-Month” variance risk premium at short horizons is consistent with findings in the literature; e.g., Bollerslev, Tauchen, and Zhou (2009), Faust and Wright (2012), and Drechsler (2013). This is unsurprising given that I construct the “Monthly End-of-Month” variance risk premium to mimic the properties of the series in the literature. Why is it that the “Monthly End-of-Month” variance risk premium predicts returns so well at short horizons, but the “Monthly Averages” variance risk premium does not? The reason is that the former series is volatile and sometimes negative during crisis periods, the two justifications I gave for ruling the series out as a


proxy for ambiguity. For example, as discussed, the “Monthly End-of-Month” variance risk premium is at its lowest value of -5.235 for 10/2008. Stocks dropped for the remainder of 2008 and the early part of 2009, which leads to a high degree of short-run predictability. In contrast, the “Monthly Averages” variance risk premium is near its highest value for 10/2008, which, while sensible for the interpretation of the variable as a metric of ambiguity, reduces the short-run predictive power for stock returns.

2.5.3 Alternate State Variables for Ambiguity

I consider several alternate state variables for ambiguity but ultimately reject these proxies in favor of the “Monthly Averages” variance risk premium. From an empirical perspective, both the squared VIX and the physical expectation of realized variance constructed in the previous section seem like plausible proxies for ambiguity. Williams (2014), for example, explicitly uses the VIX to measure ambiguity. Based on Figure 2.4, the time series of these two variables spike at the “right” times; that is, during crisis episodes when the variance risk premium also spikes. Whereas there is theoretical justification for why the variance risk premium measures ambiguity, however, there is no such justification for the true/statistical expectation of variance. By extension, the VIX, which incorporates the physical expectation of realized variance as one of its two components, is a less-ideal metric of ambiguity than the variance risk premium.

Another measure of ambiguity that feels plausible is the dispersion of forecasts about macroeconomic conditions. This intuition seems especially relevant since this paper assesses how the stock market responds to news about macroeconomic conditions. To explore this idea further, I obtain dispersion data from Bloomberg on the standard deviation and the range (maximum minus minimum) of the professional forecasts for each of the MNAs in the sample. In order to allow for comparison across different events with different units, I separately standardize the two types of dispersion data for each event to have standard deviations equal to one over the entire sample. A daily simple moving average of the standardized dispersion data for all MNAs in a one year window yields the two

Panel A: Short-Horizon Forecast Dispersion Time Series



Panel B: Long-Horizon Forecast Dispersion Time Series




Figure 2.6: Time Series of Macroeconomic Forecast Dispersions

Notes: Panel A plots two forecast dispersion series at a short forecast horizon. I obtain dispersion data from Bloomberg on the standard deviation and the range (maximum minus minimum) of forecasts for all MNAs in the sample (see Table 2.1). To allow for comparison across events with different units, I separately standardize the two types of dispersion data for each event to have standard deviations equal to one over the entire sample. For each day t in the plotted series, I average the dispersion data for all MNAs contained in the 1-year window centered at t to produce dispersion series based on the standard deviation of forecasts (“stdev”) and the range of forecasts “range.” The centered window ranges from 3/12/1998 to 12/19/2013. Panel B plots forecast dispersion series at long forecast horizons ranging from 1 quarter (“1q”) to 4 quarters (“4q”). Data are quarterly from 1997Q1 to 2014Q3 for q/q real GDP growth and come from the Philadelphia Fed’s SPF. Dispersion is measured as the 75th percentile minus the 25th percentile.

candidate daily time series in Panel A of Figure 2.6: one based on the standard deviation of forecasts (“stdev”) and one based on the range of forecasts (“range”). As evident from the figure, the dispersion series behave quite differently from the variance risk premium series and do not seem to be good empirical proxies for ambiguity. On the one hand, the dispersion series increase significantly during the 2008 financial crisis consistent with the rise of the variance risk premium during a time of heightened stress. On the other hand, the dispersion series also spike in the fall of 2005 (due to Hurricane Katrina), whereas the variance risk premium remains low, and the dispersion series remain low in the fall of 2011, whereas the variance risk premium spikes. Patton and Timmermann (2010) provide some intuition for why the dispersion series do not proxy for ambiguity. They suggest that at short forecast horizons, dispersion in forecasts about some macroeconomic fundamental is driven primarily by heterogeneity in private information related to that fundamental as opposed to heterogeneity in models about the evolution of that fundamental, the latter of which corresponds more closely with the idea of ambiguity. At long horizons, however, the opposite is true. The rationale is that forecasters have access to and use different information to form short-horizon forecasts. At longer horizons, these informational differences are likely to be less important. More important are modeling differences such as the econometric structure or the choice of sample period for model estimation. Since my entire sample of MNAs concerns forecasts at very short horizons (on the order of one week to one month), the resulting dispersion series may reflect forecasters’ different information signals instead of any notion of ambiguity. That the dispersion series diverge from the variance risk premium series in the fall of 2005 is evidence in support of short-horizon forecast dispersion representing informational differences instead of ambiguity. As a result of Hurricane Katrina, forecasters may have disagreed on an informational basis such as whether unemployment insurance initial claims would be filed in a timely manner. It is less likely that a one-off event could significantly increase modeling disagreement among forecasters.

The Patton and Timmermann (2010) reasoning implies that dispersion in longer-horizon forecasts on macroeconomic conditions do represent model uncertainty and hence could

better serve as a proxy for ambiguity. To explore this implication, I look at the dispersion (75th percentile minus 25th percentile) in forecasts of a number of macroeconomic variables from the Federal Reserve Bank of Philadelphia's Survey of Professional Forecasters (SPF). In particular, I look at the quarterly time series of dispersion for q/q real GDP growth for forecast horizons ranging from 1 quarter to 4 quarters. Panel B of Figure 2.6 plots this series for all four forecast horizons from 1997Q1 to 2014Q3. We see that the series spike during the 2008 financial crisis, consistent with the time series behavior of short-horizon forecast dispersions and the variance risk premium. Moreover, in line with the Patton and Timmermann (2010) intuition, the long-horizon forecast dispersions appear to proxy for model uncertainty better than the short-horizon series; e.g., the former do not increase in the fall of 2005 as a result of Hurricane Katrina. Other papers in the literature also interpret long-horizon dispersions as proxying for model uncertainty. Drechsler (2013), for example, shows that the 1-quarter forecast dispersion for q/q real GDP growth from the SPF correlates well with his measure of the variance risk premium.

Nonetheless, I find the long-horizon forecast dispersions in Panel B of Figure 2.6 to be lacking as measures of ambiguity. Take the European sovereign debt and U.S. debt ceiling crises episode during the later part of 2011. As noted, the variance risk premium rose significantly during that period, consistent with the intuition that ambiguity must also have been high at that time. It is clear, however, that the long-horizon forecast dispersions do not rise during that time (neither do the short-horizon dispersions). I find broadly similar results when plotting (unreported) q/q growth time series of the GDP price index and nominal GDP from the SPF. Note that Drechsler (2013) does not address this divergence between the long-horizon dispersions and his variance risk premium because his time series end in 2009.

What explains this divergence? While long-horizon forecast dispersions capture model uncertainty, the uncertainty is that experienced by economic forecasters who do not know the right model for predicting the evolution of macroeconomic conditions. In contrast, I focus on the model uncertainty that investors experience when interpreting news about

macroeconomic conditions for the fundamental value of the stock market. These two types of ambiguity, while plausibly correlated, do not have to move in lockstep. In 2011, the European sovereign debt and U.S. debt ceiling crises may have led investors to become less sure of their interpretation of macroeconomic news for stocks without forcing forecasters to become less sure of how macroeconomic variables would evolve in the short- and long-horizons.

The final proxy for ambiguity that I consider is the Baker, Bloom, and Davis (2013) index for economic policy uncertainty. Their monthly index has three underlying components: newspaper article mentions of words related to economic policy uncertainty, upcoming expirations of federal tax code provisions, and dispersion of macroeconomic forecasts based on the SPF. In unreported results, I look at the time series behavior of the overall index as well as its components. Though the empirical behavior of the various series suggest some ability to track ambiguity (e.g., increases during the 2008 financial crisis), two concerns prevent serious consideration of the series. First, the overall index includes dispersion of macroeconomic forecasts, which, as discussed, is a problematic proxy for the ambiguity that I model. Second, there is no good theoretical justification for why newspaper references to economic policy uncertainty or expirations of federal tax code provisions should measure investors' uncertainty about the impact of macroeconomic news on the stock market. As such, I stick with the variance risk premium as my measure of ambiguity.

2.6 Empirical Results: Time-Varying Ambiguity

Using the “Monthly Averages” variance risk premium as a state variable for ambiguity, I test the sharper predictions of the model that the behavior of stocks in response to macroeconomic news depends on the amount of news ambiguity. The first section shows that the asymmetry effect and the no news is good news effect are greater in magnitude when there is greater ambiguity. The increased strength of the asymmetry effect is due to a stronger response of the stock market to bad news as opposed to a weaker response to good news. When viewed through the lens of the model, this finding implies that increases in ambiguity come from investors thinking that the lower bound for the noise of

macroeconomic news in determining the fundamental value of the stock market is smaller as opposed to the upper bound being larger. In the second section, I show that the ambiguity premium exists independent of the amount of ambiguity, a result which can be reconciled with the model. Finally, the third section finds that the volatility of stocks around MNAs is greater in times of higher ambiguity.

2.6.1 The Asymmetry Effect and No News is Good News Effect

If Propositions 1 and 2 are true, the magnitudes of the asymmetry effect and no news is good news effect should be increasing in the amount of ambiguity. Running Eq. (2.22) over the sample of MNAs that occur at times of high ambiguity should lead to a larger $\hat{\alpha}$ and a more negative $\hat{\beta}^+ - \hat{\beta}^-$ compared to the same regression over a sample of MNAs that occur at times of lower ambiguity. To this end, I use the variance risk premium time series to proxy for ambiguity and associate each MNA with the corresponding value of the variance risk premium in the month of the data release. I then split the full sample of MNAs into five variance risk premium quintiles, with the first quintile consisting of macroeconomic news announced when the variance risk premium is at its lowest and the fifth quintile consisting of news announced when the premium is at its highest. Panel B of Table 2.3 provides summary statistics on the standardized news variable S_t across the quintiles. We see that for each quintile, the distribution of news is unbiased and symmetric. Moreover, the distribution of S_t is similar across the five quintiles. One might expect that at certain times, such as during the 2008 financial crisis, the frequency or magnitude of negative news about the macroeconomy should increase. Since the variance risk premium time series also spikes at these times based on Figure 2.5, the above logic suggests that S_t may be negatively-biased in the higher variance risk premium quintiles. The reason that this effect is not present in Panel B of Table 2.3 is that S_t is calculated as the actual data release relative to expectations. While times of crisis do result in more negative data releases of macroeconomic variables, economic forecasters correspondingly revise lower expectations for data releases. This downward revision of expectations offsets the poorer data, so S_t is

Table 2.8: Stock Market Reaction to MNAs Bucketed into Variance Risk Premium Quintiles

VRP_t Quintile	$\hat{\alpha}$ (bps)	$\hat{\beta}^+$ (bps)	$\hat{\beta}^-$ (bps)	Adj. R^2	N	$\hat{\beta}^+ - \hat{\beta}^-$ (bps)
1	2.829*** (3.748)	0.596 (0.589)	1.551 (1.614)	0.229	796	-0.955 (0.322)
2	1.519 (1.568)	4.407*** (3.181)	5.784*** (2.912)	4.606	780	-1.377 (0.341)
3	2.389** (2.085)	1.980 (0.962)	8.431*** (4.576)	4.188	815	-6.451** (5.858)
4	4.706*** (3.418)	-1.837 (-1.017)	9.032*** (3.776)	2.953	781	-10.869*** (12.759)
5	4.734*** (2.756)	5.030*** (2.602)	15.648*** (7.718)	10.497	813	-10.618*** (11.715)

Notes: The table shows results from regressions bucketed by variance risk premium VRP_t quintiles, with the 1st quintile corresponding to the smallest values and the 5th quintile corresponding to the largest values. The regression in each quintile is of ES returns on standardized news grouping all events together: $R_t = \alpha + \beta^+ D_t^+ S_t + \beta^- (1 - D_t^+) S_t + \epsilon_t$. R_t is the stock market return, S_t aggregates the news surprises of all MNAs, and D_t^+ is a dummy variable. All regression variables are as defined in Table 2.2. VRP_t is the “Monthly Averages” variance risk premium series, which corresponds to the monthly time series in Panel B of Figure 2.4. t -statistics (shown in parentheses) are based on heteroskedasticity-consistent standard errors. Rightmost column presents $\hat{\beta}^+ - \hat{\beta}^-$ and performs a Wald test that the difference is not zero (F-statistics shown in parentheses). *** denotes significance at the 1% level, ** denotes significance at the 5% level, and * denotes significance at the 10% level. Table above covers the whole period from November, 1997 to March, 2014.

distributed symmetrically in each quintile and comparably across quintiles.

Table 2.8 shows results from the regression in Eq. (2.22) for each quintile of the variance risk premium, and Figure 2.7 plots the regression parameters against the quintiles. Together, Table 2.8 and Figure 2.7 provide the main evidence in support of Propositions 1 and 2. Looking at the last column of Table 2.8 and Panel B of Figure 2.7, we see that, consistent with Proposition 1, the stock market exhibits little asymmetry to macroeconomic news in times of low ambiguity but large asymmetry in times of high ambiguity. $\hat{\beta}^+ - \hat{\beta}^- = -0.955$ bps is statistically indistinguishable from zero in the first variance risk premium quintile based on a Wald test, so the stock market reacts by a similar amount to a unit of positive surprise versus a unit of negative surprise. Moving to higher variance risk premium quintiles, $\hat{\beta}^+ - \hat{\beta}^-$ decreases consistently and becomes statistically different from zero. In the fifth




Figure 2.7: Stock Market Reaction to MNAs Bucketed into Variance Risk Premium Quintiles

Notes: See Appendix B.3.2. Panels A, B, and C from top to bottom correspond to $\hat{\beta}^+$ and $\hat{\beta}^-$ versus VRP_t , $\hat{\beta}^+ - \hat{\beta}^-$ versus VRP_t , and $\hat{\alpha}$ versus VRP_t plots, respectively.

quintile, $\hat{\beta}^+ - \hat{\beta}^- = -10.618$ bps; that is, the stock market falls 10.618 bps more for a unit of negative macroeconomic news than it rises for a unit of positive news.

The increase in the asymmetry effect comes entirely from the increased sensitivity of the stock market to negative news as opposed to the decreased sensitivity of the stock market to positive news. From the $\hat{\beta}^+$ and $\hat{\beta}^-$ columns of Table 2.8 and Panel A of Figure 2.7, we see that $\hat{\beta}^+$ shows no clear pattern across the variance risk premium quintiles, whereas $\hat{\beta}^-$ increases consistently from 1.551 bps in the first quintile to 15.648 bps in the fifth quintile. In the context of the model and Eq. (2.12), the results imply that while greater news ambiguity increases $\bar{\phi} - \underline{\phi}$ (equivalently, $\bar{\sigma}_s^2 - \underline{\sigma}_s^2$), this increase comes primarily through an increase in $\bar{\phi}$ (decrease in $\underline{\sigma}_s^2$) and not through a decrease in $\underline{\phi}$ (increase in $\bar{\sigma}_s^2$). The representative investor has a range of variances in mind for the noisiness of a given piece of news about the fundamental value of the risky asset. With greater ambiguity, this range of variances expands, but it expands in a skewed manner such that the lower end of the range decreases, but the upper end of the range stays approximately the same. When the investor sees a positive piece of news, he believes that the variance of the noise component is at the upper end of the range, which reduces the impact of the good news on the risky asset. Since the upper end of the range is similar for times of low and high ambiguity, the risky asset has a similarly subdued response to good news in both states of the world. When the investor sees a negative piece of news, he believes that the variance of the noise component is at the lower end of the range, which amplifies the impact of the bad news on the risky asset. The lower end of the range is lower for times of high ambiguity versus low ambiguity, however, so the risky asset responds more strongly to the negative news in the former case versus the latter case. Since ambiguity, as proxied by the variance risk premium in Figure 2.5, is generally high during times of market stress, another way of stating the above result is that investors believe that good news is equally relevant (or irrelevant) for stocks in good and bad times. Investors believe that bad news, however, is more relevant for stocks in bad states than in good states.

From the $\hat{\alpha}$ column of Table 2.8 and Panel C of Figure 2.7, we see that $\hat{\alpha}$ increases from

the first variance risk premium quintile to the fifth quintile, which supports Proposition 2. The magnitude of $\hat{\alpha}$ averages to 2.174 bps in the first two quintiles and 4.720 bps in the last two quintiles, an increase of 2.546 bps. That is, in times of high news ambiguity, the stock market increases 2.546 bps more in response to no news than in times of low news ambiguity.

2.6.2 The Ambiguity Premium

Based on the model, the ambiguity premium from Eq. (2.18) should, all else equal, be increasing in the news ambiguity $\bar{\phi} - \underline{\phi}$, which is the first term in parentheses. Taking the data from Panel A of Table 2.4, Panel B of the same table breaks the stock market return through MNAs into quintiles by the variance risk premium. The top set of rows presents the arithmetic sum of returns (expressed in %), and the bottom set of rows presents the mean of returns (expressed in bps). Focusing on the “Return” column, which is the actual return of the front ES contract around MNAs, we see that returns are positive in each of the variance risk premium quintiles. The ambiguity premium is thus not only positive over the whole sample, but also when the sample is split in five by the variance risk premium. It is clear from Table 2.4, however, that there is not much of a pattern in the “Return” column across variance risk premium quintiles. To better understand this result, I decompose the ambiguity premium in each quintile of Panel B into two components, as in Panel A, using the estimated $\hat{\alpha}$, $\hat{\beta}^+$, $\hat{\beta}^-$ from Table 2.8: the positive no news is good news effect and the negative asymmetry effect. Consistent with earlier results, the asymmetry effect exerts a more negative effect on the average return of the stock market around MNAs as ambiguity increases. The no news is good news effect has the counterbalancing role of exerting a more positive effect on the average return of the stock market around MNAs as ambiguity increases. These two results depend on the verification of Propositions 1 and 2 in Table 2.8 as well as the similar distribution of MNA surprises across variance risk premium quintiles documented in Table 2.3. The asymmetry effect and the no news is good news effect both change in such a way across variance risk premium quintiles that their sum, the ambiguity

premium, exhibits no discernable pattern across quintiles.

An alternate method to assess whether the ambiguity premium is time-varying is to conduct predictability regressions as in Eq. (2.26). Instead of forecasting $r_{t,t+h}$, the stock return from month t to month $t+h$, I forecast $r_{t,t+h}^{MNA}$, that portion of the stock return attributed to the release of macroeconomic news. If the ambiguity premium varies over time, $r_{t,t+h}^{MNA}$ should be predictable. More specifically, if the ambiguity premium is higher at times of greater ambiguity, the variance risk premium should be a positive, significant predictor of $r_{t,t+h}^{MNA}$. My regression specification is

$$r_{t,t+h}^{MNA} = a_{MNA} + b_{MNA}X_t + \epsilon_{t,t+h}^{MNA}, \quad (2.27)$$

in which $r_{t,t+h}^{MNA}$ is the annualized sum of the continuously compounded return of the stock market in the ± 5 minutes around MNAs from the end of month t to the end of month $t+h$. As discussed previously, $r_{t,t+h}^{MNA}$ can be thought of as an excess return due to the use of futures data in calculating returns. Panel A of Table 2.9 shows \hat{b}_{MNA} , associated t -statistics, and adjusted R^2 for a range of horizons and three sets of covariates X_t . When $X_t = \{\log(P/E)_t\}$, there is little evidence of predictability except at the longest horizons. Setting $X_t = \{VRP_t\}$, the “Monthly Averages” variance risk premium, no coefficients are significant. Similarly, incorporating both the price-to-earnings ratio and variance risk premium for $X_t = \{\log(P/E)_t, VRP_t\}$ yields no significant coefficients. There is some suggestive evidence that the variance risk premium positively predicts stock returns at short horizons. For example, at $h = 3$ months, the coefficient for the $X_t = \{VRP_t\}$ regression has a t -statistic of 1.457, which is significant at the 15% level, and an adjusted R^2 of 1.324. The presence of the variance risk premium also improves the adjusted R^2 of the $X_t = \{\log(P/E)_t, VRP_t\}$ regression specification to 1.581 versus the -0.326 with only $X_t = \{\log(P/E)_t\}$. On the whole, however, the return of stocks around macroeconomic news does not appear to vary much over time, which suggests that the ambiguity premium is not increasing in the amount of ambiguity as measured by the variance risk premium. Since we know from Table 2.7 and the extensive literature on return predictability that stock

Table 2.9: Decomposing the Predictability of Stock Returns

Panel A: Predictability of Monthly Stock Returns around MNAs

Horizon	1	3	6	9	12	15	18	21	24
$\log(P/E)_t$	1.508 (0.327)	1.220 (0.242)	2.020 (0.397)	2.583 (0.593)	3.275 (0.972)	3.420 (1.286)	3.488 (1.595)	3.298* (1.746)	3.001* (1.802)
Adj. R^2	-0.407	-0.326	0.504	1.806	4.434	6.289	8.390	9.527	10.024
VRP_t	0.043 (0.427)	0.097 (1.457)	0.079 (1.191)	0.046 (0.948)	0.006 (0.174)	0.000 (-0.009)	-0.013 (-0.482)	-0.017 (-0.635)	-0.016 (-0.599)
Adj. R^2	-0.383	1.324	1.842	0.587	-0.525	-0.556	-0.374	-0.177	-0.139
$\log(P/E)_t$	2.172 (0.480)	2.536 (0.530)	3.212 (0.658)	3.450 (0.803)	3.706 (1.058)	3.826 (1.338)	3.743 (1.478)	3.499 (1.552)	3.186 (1.567)
VRP_t	0.059 (0.564)	0.116 (1.496)	0.104 (1.222)	0.074 (1.062)	0.036 (0.666)	0.032 (0.715)	0.020 (0.435)	0.015 (0.342)	0.013 (0.316)
Adj. R^2	-0.696	1.581	3.715	3.860	4.707	6.575	8.242	9.273	9.775

120

Panel B: Predictability of Monthly Stock Returns Excluding MNAs

Horizon	1	3	6	9	12	15	18	21	24
$\log(P/E)_t$	-14.477 (-0.754)	-15.783 (-0.888)	-21.932 (-1.490)	-24.424* (-1.788)	-25.597** (-2.063)	-27.030** (-2.540)	-28.127*** (-3.070)	-29.112*** (-3.751)	-29.233*** (-4.596)
Adj. R^2	-0.053	0.956	4.286	7.903	11.262	15.768	20.718	25.980	30.591
VRP_t	0.143 (0.242)	0.307 (0.685)	0.468** (2.141)	0.474** (2.430)	0.465*** (2.665)	0.432*** (2.836)	0.452 *** (3.290)	0.495*** (3.845)	0.531*** (4.391)
Adj. R^2	-0.447	0.323	2.760	4.185	5.208	5.572	7.556	10.767	14.783
$\log(P/E)_t$	-14.045 (-0.820)	-13.442 (-0.788)	-18.149 (-1.102)	-20.785 (-1.320)	-22.174 (-1.520)	-24.203* (-1.903)	-25.125** (-2.313)	-25.586*** (-2.767)	-25.107*** (-3.166)
VRP_t	0.038 (0.067)	0.206 (0.471)	0.329 (1.226)	0.309 (1.264)	0.284 (1.243)	0.226 (1.088)	0.232 (1.300)	0.261 (1.639)	0.299** (1.982)
Adj. R^2	-0.567	0.786	5.265	9.242	12.722	16.799	22.164	28.342	34.477

Notes: See Appendix B.3.3.

returns in general are predictable yet find stock returns around MNAs are not predictable, a natural follow-up is to check if stock returns in times excluding MNAs are predictable. In Panel B of Table 2.9, I run regressions of the form

$$r_{t,t+h}^{nMNA} = a_{nMNA} + b_{nMNA}X_t + \epsilon_{t,t+h}^{nMNA}, \quad (2.28)$$

in which $r_{t,t+h}^{nMNA} = r_{t,t+h} - r_{t,t+h}^{MNA}$, the annualized sum of the continuously compounded return of the stock market outside of the ± 5 minutes around MNAs from the end of month t to the end of month $t+h$. We see evidence for predictability at various horizons and specifications of X_t . Thus, the predictability of stock returns comes almost entirely from returns earned in times excluding macroeconomic data releases, a result consistent with Faust and Wright (2012).

Can the lack of time-variation in the ambiguity premium be reconciled with the model? Consider the expression for the ambiguity premium in Eq. (2.18), specifically the second term in parentheses $\left(\sqrt{\alpha^2\sigma_d^2 + \bar{\sigma}_s^2} - \sqrt{\alpha^2\sigma_d^2 + \tilde{\sigma}_s^2}\right)$. Based on the fact that the asymmetry effect is due primarily to the increased sensitivity of the stock market to negative news as opposed to the decreased sensitivity of the stock market to positive news, we have inferred in the previous section that increases in $\bar{\phi}$ (decreases in $\underline{\sigma}_s^2$) drive increases in news ambiguity $\bar{\phi} - \underline{\phi}$, not decreases in $\underline{\phi}$ (increases in $\bar{\sigma}_s^2$). As such, the second term in parentheses in the expression for the ambiguity premium does not increase with news ambiguity unlike in the case if news ambiguity were driven by decreases in $\underline{\phi}$ (increases in $\bar{\sigma}_s^2$). This result dampens the positive relationship between the ambiguity premium and ambiguity. A second explanation for the stability of the ambiguity premium centers on $\tilde{\sigma}_s^2$, the true variance of the ϵ_s component of the news. If $\tilde{\sigma}_s^2$ were close to $\bar{\sigma}_s^2$, the second term in parentheses of the ambiguity premium would be close to zero, which would make it difficult to see an increasing relationship between the ambiguity premium and ambiguity. The investor thinks that the news could be a very noisy indicator of the fundamental value of the risky asset ($\bar{\sigma}_s^2$), and it turns out that the news is in fact a quite noisy indicator ($\tilde{\sigma}_s^2$). Alternatively, if $\tilde{\sigma}_s^2$ were increasing in news ambiguity for some reason, the second term in

parentheses would be decreasing in ambiguity, which would offset the rise in $\bar{\phi} - \underline{\phi}$ and sever the link between the ambiguity premium and ambiguity.


2.6.3 Risky Asset Volatility

The final prediction of the model, Proposition 4, says that the volatility of the stock market around macroeconomic news releases can be either increasing or decreasing in the amount of ambiguity depending on what drives increases in ambiguity. If, as implied by the results corroborating the asymmetry effect, increases in ambiguity in the model are due to investors' allowing for smaller lower bounds for the variance of the news noise, stock market volatility around MNAs should be increasing in ambiguity. Panel A of Figure 2.8 plots the time series of the 10-minute return of the stock market in the ± 5 minutes around MNAs. We see that this return is heteroscedastic and is highly variable at predictable episodes such as the 2008 financial crisis. Table 2.10 splits this return time series into variance risk premium quintiles and calculates the standard deviation of returns in each quintile in the "10-Minute Std." column. As suggested by Panel A of Figure 2.8 and consistent with Proposition 4, the standard deviation of stock returns around macroeconomic news increases substantially as the amount of ambiguity increases. This standard deviation is 14.791 bps in the lowest quintile versus 36.587 bps in the highest quintile.

The observation that the volatility of stock returns around MNAs increases with ambiguity is related to the well-studied phenomenon that the volatility of stock returns in general fluctuates substantially in the time series. Panel B of Figure 2.8 plots the time series of daily stock market returns for days with MNAs.⁷ There is substantial heteroscedasticity, and the volatility of daily returns appears to increase at the same time that the volatility of intraday returns increases. The "1-Day Std." column of Table 2.10 calculates the standard deviation of daily returns when separated into variance risk premium quintiles. Similar to the result for the volatility of intraday returns, the volatility of daily returns also increases

⁷Since there are macroeconomic data releases on most days, this time series is essentially equivalent to the time series of daily stock returns.

Panel A: Intraday Returns around MNAs



Panel B: Daily Returns around MNAs




Figure 2.8: Time Series of Intraday and Daily Returns around MNAs

Notes: Panel A plots the 10-minute return of the stock market in the ± 5 minutes around a MNA. Panel B plots the daily return of the stock market for the day of the MNA. The data in Panels A and B are the same as those plotted on the horizontal and vertical axes of Figure 2.2, respectively. Refer to the Figure 2.2 caption for details. Figures above cover the whole period from November, 1997 to March, 2014.

Table 2.10: *Stock Market Volatility around MNAs Bucketed into Variance Risk Premium Quintiles*

VRP_t Quintile	10-Minute Std.	1-Day Std.	Fraction of Std.	N
1	14.791	64.703	0.229	632
2	20.407	89.368	0.228	636
3	21.994	108.013	0.204	636
4	26.375	137.742	0.191	643
5	36.587	205.910	0.178	642

Notes: Table takes the 10-minute return of the stock market in the ± 5 minutes around a MNA and the daily return of the stock market for the day of the MNA and calculates standard deviations (bps) in the “10-Minute Std.” and “1-Day Std.” columns, respectively. Standard deviations are calculated for each VRP_t quintile, with the 1st quintile corresponding to the smallest values and the 5th quintile corresponding to the largest values. The “Fraction of Std.” column divides the “10-Minute Std.” column by the “1-Day Std.” column. VRP_t is the “Monthly Averages” variance risk premium series, which corresponds to the monthly time series in Panel B of Figure 2.4. The intraday and daily return data are the same as those plotted on the horizontal and vertical axes of Figure 2.2, respectively. Refer to the Figure 2.2 caption for details. Table above covers the whole period from November, 1997 to March, 2014.

in ambiguity from a low of 64.703 bps in the first quintile to a high of 205.910 bps in the fifth quintile. Together, Figure 2.8 and Table 2.10 provide suggestive evidence that some of the time variation in the volatility of stock returns more generally can be attributed to the impact of ambiguity on the volatility of stock returns around MNAs.

2.7 Alternative Explanations

To my knowledge, no explanation in the literature provides a unified story for the full set of empirical patterns exhibited by the aggregate stock market around MNAs. There are, however, stories not related to ambiguity that could plausibly explain a subset of the findings. In the two sections below, I investigate different narratives that have the prediction that the stock market reacts asymmetrically to good news versus bad news. The data support the model based on ambiguity over these alternative explanations.

2.7.1 Good Times versus Bad Times

Veronesi (1999) presents a rational, regime-switching framework in which the aggregate stock market can react asymmetrically to news depending on whether the news is good or bad. In his model, investors are uncertain about the state of the market and must infer whether times are good or bad. During good times, bad news has a large impact on stocks by both decreasing the probability investors place on the state of the world being good and increasing uncertainty about the state of the world. In contrast, good news has a small impact on stocks because investors already place a high probability on being in the good state of the world. During bad times, bad news has a small impact on stocks because investors believe they are already in the bad state of the world. The impact of good news on stocks is also muted, however, because while good news increases the probability investors place on the state of the world being good, good news also increases uncertainty about the state of the world. Altogether, one implication of Veronesi (1999) is that the aggregate stock market reacts more strongly to bad news than to good news during good times, and this asymmetry diminishes or is nonexistent during bad times.

The model based on investor sentiment in Barberis, Shleifer, and Vishny (1998) has a related prediction to that of Veronesi (1999) for the response of individual stocks to earnings news. In their model, investors extrapolate good (bad) times, characterized by strings of good (bad) earnings news, and expect future good (bad) earnings news. During good times, bad earnings news is unexpected and generates a large effect, but good earnings news is expected and generates a small effect. In contrast, during bad times, bad earnings is expected and generates a small effect, but good earnings news is unexpected and generates a large effect. The Barberis, Shleifer, and Vishny (1998) framework thus suggests that individual stocks react more strongly to bad earnings news than to good earnings news during good times, but this asymmetry flips during bad times.

Conrad, Cornell, and Landsman (2002) document empirical evidence in support of Veronesi (1999) and Barberis, Shleifer, and Vishny (1998) through investigating the response of individual stocks to earnings announcements. By using a relative price-to-earnings ratio


as a proxy for good times versus bad times, Conrad, Cornell, and Landsman (2002) find that stocks indeed react most strongly to bad earnings news versus good earnings news when the relative value of the market is high (good times), with the asymmetry diminishing in magnitude when the market valuation is lower (bad times).⁸

The results of the aforementioned three papers suggest that the state variable that drives asymmetry is some measure of how good or bad the state of the world is. I assess whether this alternative explanation can account for the stock market behavior around macroeconomic news releases. To do so, I roughly follow Conrad, Cornell, and Landsman (2002) and use the cyclically-adjusted price-to-earnings ratio ($CAPE_t$) from Robert Shiller's website to create a proxy for how good or bad times are. I then replicate my results using this state variable instead of the variance risk premium. Panel A of Figure 2.9 shows the time series of $CAPE_t$ from November, 1997 to March, 2014, and Panel A of Table 2.11 provides summary statistics. Conrad, Cornell, and Landsman (2002) argue that using the absolute price-to-earnings ratio is problematic. People may not care, for example, about how good or bad times are on an absolute scale but relative to recent experience. Moreover, as the figure shows, using $CAPE_t$ as the state variable is essentially the same as just using a time dummy. The market valuation is highest in the early part of the sample around the dot-com bubble, intermediate during the mid-2000s, and lowest during and subsequent to the 2008 financial crisis. Categorizing the whole sample into good times, mediocre times, and bad times in that chronological order seems both simplistic and inaccurate. Similar to Conrad, Cornell, and Landsman (2002), I address this issue by creating a relative price-to-earnings ratio ($DCAPE_t$) in Panel B of Figure 2.9 by taking $CAPE_t$ and subtracting its own trailing 12-month average. The time series of $DCAPE_t$ exhibits more reasonable dispersion that is not solely based on time. Summary statistics for $DCAPE_t$ appear in Panel A of Table 2.11.

To assess how relative market valuation drives the differential response of the stock market to macroeconomic news, I run regressions as before based on Eq. (2.22). I use

⁸On a related note, papers such as Boyd, Hu, and Jagannathan (2005) suggest that the relationship between stocks and macroeconomic news depends on the phase of the business cycle; i.e., whether the news arrives during NBER recessions or expansions.

Panel A: Cyclically-Adjusted Price-to-Earnings Ratio



Panel B: Relative Cyclically-Adjusted Price-to-Earnings Ratio




Figure 2.9: Time Series of Absolute and Relative Price-to-Earnings Ratios

Notes: Panel A plots the time series of the cyclically adjusted price-to-earnings ratio $CAPE_t$ from Robert Shiller's website. Panel B plots the time series of the relative cyclically adjusted price-to-earnings ratio $DCAPE_t$ calculated as $CAPE_t$ less its own trailing 12-month average. Figures above cover the whole period from November, 1997 to March, 2014.

Table 2.11: Replication of Results Using the Relative Price-to-Earnings Ratio to Bucket the DataPanel A: Summary Statistics of $CAPE_t$ and $DCAPE_t$

	Mean	Median	Std. Dev.	Max.	Min.	Skewness	Kurtosis	AC(1)
$CAPE_t$	27.016	25.650	7.300	44.190	13.320	0.871	0.009	0.990
$DCAPE_t$	-0.225	0.241	2.808	5.158	-8.282	-0.671	0.136	0.927

Panel B: Summary Statistics of MNA Surprises, with Sample Bucketed into $DCAPE_t$ Quintiles

$DCAPE_t$ Quintile	Mean	Median	Std. Dev.	Max.	Min.	Skewness	Kurtosis	N
1	-0.019	0.000	1.227	8.018	-5.514	0.391	3.495	779
2	0.009	0.000	1.016	4.866	-4.315	0.311	2.537	802
3	0.044	0.000	0.970	6.257	-3.527	0.533	3.198	805
4	0.014	0.052	0.912	3.303	-4.171	-0.495	1.437	788
5	0.100	0.017	0.895	3.332	-3.452	-0.141	0.749	811

Panel C: Stock Market Reaction to MNAs Bucketed into $DCAPE_t$ Quintiles

$DCAPE_t$ Quintile	$\hat{\alpha}$ (bps)	$\hat{\beta}^+$ (bps)	$\hat{\beta}^-$ (bps)	Adj. R^2	N	$\hat{\beta}^+ - \hat{\beta}^-$ (bps)
1	4.443** (2.533)	3.763* (1.941)	14.080*** (7.293)	9.958	779	-10.317*** (12.473)
2	2.346* (1.848)	4.217** (2.473)	7.445*** (4.079)	4.272	802	-3.228 (1.387)
3	3.983*** (3.875)	1.356 (1.120)	5.318** (2.136)	1.765	805	-3.962 (2.570)
4	1.651 (1.572)	3.876** (2.024)	5.079*** (3.268)	3.270	788	-1.203 (0.212)
5	3.314*** (3.200)	-1.836 (-1.141)	7.032*** (3.883)	2.108	811	-8.869*** (11.759)

Notes: Panel A shows summary statistics for $CAPE_t$ and $DCAPE_t$, the two variables plotted in Panels A and B of Figure 2.9, respectively. Panel B replicates Panel B of Table 2.3 and provides summary statistics for the standardized news variable S_t bucketed according to $DCAPE_t$ quintiles, with the 1st quintile corresponding to the smallest values and the 5th quintile corresponding to the largest values. Panel C replicates Table 2.8 and shows the stock market reaction to MNAs bucketed into $DCAPE_t$ quintiles. Tables above cover the whole period from November, 1997 to March, 2014.

$DCAPE_t$ to split my sample of MNAs into five quintiles with the first quintile consisting of macroeconomic news occurring when the relative market valuation is lowest and the fifth quintile consisting of macroeconomic news occurring when the relative market valuation is highest. Analogous to Panel B of Table 2.3, Panel B of Table 2.11 shows summary statistics for standardized news S_t across the $DCAPE_t$ quintiles. S_t is unbiased and symmetric in each quintile and seems to be distributed similarly across quintiles. Analogous to Table 2.8, Panel C of Table 2.11 shows regression results for each quintile. From the rightmost column, $\hat{\beta}^+ - \hat{\beta}^-$ is most negative in the first quintile at -10.317 bps, statistically indistinguishable from zero in the middle three quintiles, and second-most negative in the fifth quintile at -8.869 bps. This pattern is inconsistent with the story that how good or bad the state of the world is, as measured by the relative market valuation, drives the differential response of the stock market to good news versus bad news. If that state variable were the driving force, we would expect $\hat{\beta}^+ - \hat{\beta}^-$ to be most negative in the fifth quintile and progressively less negative moving down to the first quintile. The response of the aggregate stock market to macroeconomic news provides conflicting evidence for the implications of Veronesi (1999) and Barberis, Shleifer, and Vishny (1998) in contrast with the support that I find in the data for the framework based on ambiguity. This outcome is not particularly surprising given that times of high ambiguity, as measured by the variance risk premium, tend to be bad times, yet these are also times in which the stock market responds more strongly to bad news than to good news.

There are several reasons why my empirical results differ so significantly from those of Conrad, Cornell, and Landsman (2002). First, I look at the aggregate market response to macroeconomic news instead of the individual stock response to earnings news.⁹ A second explanation for the discrepancy is my use of high-frequency intraday data instead of daily data, which I have stressed makes a significant difference for the empirical results in this paper.

⁹Note that the Veronesi (1999) model applies to the aggregate market, the Barberis, Shleifer, and Vishny (1998) model applies to individual stocks, and the model based on ambiguity in my paper can apply to either the stock market as a whole or individual stocks (though I focus on the former).

2.7.2 Volatility Feedback

Volatility feedback can also lead to the asymmetric response of the stock market to good news versus bad news. Campbell and Hentschel (1992) have such a model in which large pieces of news tend to be followed by large pieces of news, which reflects the persistence of volatility. This phenomenon dampens the effect of good news on stocks and amplifies the effect of bad news, thus providing a potential explanation for why the aggregate stock market responds more strongly to bad news than to good news. Volatility feedback is unlikely, however, to explain the asymmetry effect exhibited by stocks around MNAs. The reason is that MNAs are anticipated. If there is a large surprise in a given piece of macroeconomic news, investors take this information into account in forecasting the next piece of macroeconomic news. This self-corrective mechanism means that there is little persistence in the magnitude of the surprise of MNAs. Consistent with this notion, the distribution of MNA surprises is time-invariant as shown in Figure 2.3, Table 2.3, and Panel B of Table 2.11.

2.8 Conclusion

Using a theoretical framework with ambiguity as its centerpiece, this paper has shown a number of testable predictions for the behavior of stocks around ambiguous news. Stocks react asymmetrically stronger to bad news than to good news, increase in response to neutral news, and have a positive average return through news. Times of higher ambiguity coincide with larger magnitudes for the aforementioned effects and moreover coincide with either higher or lower volatility of stocks around news depending on the driver of the increase in ambiguity. I find that these predictions receive considerable empirical support based on the dynamics of the aggregate stock market around news about the macroeconomy. To fully test the model, I construct a variance risk premium time series that is a more reasonable proxy for ambiguity than existing data series. Among the observations that corroborate the model with ambiguity, I find that stocks react more strongly to bad macroeconomic

news than good macroeconomic news and especially so during crisis periods. I also show that about 1/3 of the equity return from 1997 to 2014 accrues in just 10 minutes around MNAs. Viewed through the model, these results suggest that investors behave in such a way that they treat bad news as more salient in bad times than in good times but good news as equally salient in both times. The empirical results are hard to reconcile with existing models and suggest that ambiguity is an important factor in asset pricing and, in this specific setting, how financial assets reflect macroeconomic risks. As such, ambiguity warrants further consideration in finance and macrofinance models.

Chapter 3

Asset Price Reactions to News at the Zero Lower Bound

3.1 Introduction

A fundamental question in economics is how asset prices impound news about macroeconomic fundamentals. When news arrives of higher-than-expected inflation or output and lower-than-expected unemployment, do asset prices appreciate or depreciate and by how much? I tackle this question by analyzing the sensitivity of interest rates and the stock market in the U.S. to a broad sample of macroeconomic news announcements (MNAs). By considering the time period from December 17th, 2008 to March 6th, 2014 during which the federal funds (FF) rate was at the zero lower bound (ZLB), I find that the extent to which the ZLB binds drives the reaction of interest rates and stocks to data surprises.

We can think of the extent to which the ZLB binds as the expected duration of the ZLB, which varies significantly from the end of 2008 to 2014 even though the nominal short rate is fixed at the lower bound throughout. Figure 3.1 shows two proxies for the expected duration of the ZLB. The right-hand graph plots the median anticipated number of quarters until the first rate hike based on primary dealer surveys by the Federal Reserve Bank of New York (FRBNY). We see large variations in the anticipated time until the first rate hike from




Figure 3.1: Proxies for the expected duration of the ZLB

Notes: The left-hand plot is of the implied FF rate (%) from FF futures with maturities 1 month to 24 months out. Higher-value lines correspond to longer-maturity FF futures. The right-hand plot is the median number of quarters until the first rate hike based on primary dealer surveys by the FRBNY. The two vertical lines correspond to the dates of the August, 2011 FOMC meeting (when calendar guidance was introduced) and the December, 2012 FOMC meeting (when data thresholds was introduced), respectively.

a low of 4 quarters in mid-2011 to a high of 12 quarters in late 2012. The left-hand graph plots implied FF rates from FF futures contracts with 1 month to 24 months maturities. As the expected duration of the ZLB increases, we expect FF rates in the future to decrease, and it indeed appears that implied FF rates are negatively correlated with the surveyed time until the first rate hike. The expected duration of the ZLB can change either because economic conditions change or because the monetary policy reaction function that sets the nominal short rate changes. For example, the first vertical line in both graphs of Figure 3.1 corresponds to the August, 2011 Federal Open Market Committee (FOMC) meeting during which the Federal Reserve decided to implement “*exceptionally low levels for the federal funds rate at least through mid-2013.*” This decision was clearly a surprise as evidenced by the sharp drop in implied FF rates and the sharp increase in the surveyed time until the first rate hike. This increase in the expected duration of the ZLB could have been due to pessimism about the state of the economy or beliefs that monetary policy changed. This paper shows that regardless of whether the expected duration of the ZLB changes due to economic conditions or to monetary policy, asset price sensitivities to news change in a certain manner.

Focusing on interest rates first, I employ a shadow rate term structure model that incorporates the ZLB to make three predictions about the sensitivity of interest rates to MNAs. The first prediction is that “better”-than-expected or “positive” macroeconomic news (e.g., higher-than-expected inflation or output and lower-than-expected unemployment) increases interest rates at all maturities. The second prediction is that as the expected duration of the ZLB increases, whether because economic conditions are worse or because monetary policy is less responsive to economic conditions, the reaction of interest rates to MNAs decreases at all maturities. The third and final prediction is that this attenuation in the sensitivity of interest rates to news is initially increasing in maturity and then decreasing in maturity. In other words, the attenuation in the sensitivity of interest rates is greatest for intermediate maturity rates. I verify each of these predictions in the data using a representative sample of 18 MNAs and high-frequency data on interest rate futures at various maturities. To do so, I measure the reactions of various interest rates in a tight

± 5 minute window around MNAs and then compare time variation in these estimated reactions to the aforementioned proxies for the expected duration of the ZLB. As expected, “positive” MNAs increase interest rates, but greater expected duration of the ZLB (e.g., after the August, 2011 FOMC meeting) moderates this reaction. The extent of this moderation is moreover greatest for interest rates with maturities that are not too short and not too long.

For stocks, I initially measure the time-varying sensitivity of the stock market to data using the same sample of MNAs and high-frequency data on the S&P 500 futures contract. Based off the ideas of Campbell (1991) and Campbell and Ammer (1993), we know that when the stock market reacts to macroeconomic news, it must be due to a combination of a change in expected future dividends (cash flow news), a change in expected future interest rates (interest rate news), and a change in expected future excess returns (risk premium news). Using a simple Gordon growth model as in Boyd, Hu, and Jagannathan (2005), I decompose the reaction of the stock market to MNAs into an interest rate news term and a cash flow plus risk premium news term. The interest rate news term is exactly the sensitivity of interest rates to MNAs (adjusted by a factor). Based on this news term alone, I show that as the expected duration of the ZLB increases, we expect the sensitivity of stocks to macroeconomic data to increase as well. The empirical relationship between the expected duration of the ZLB and the sensitivity of stocks is less clear-cut because of time-variation in the offsetting cash flow and risk premium news terms. Intriguingly, the expected duration of the ZLB is negatively correlated with the magnitude of the cash flow and risk premium news terms.

The ideas presented above connect to several research areas in the literature. Many papers have looked at how the release of macroeconomic news affects asset prices; e.g., Andersen, Bollerslev, Diebold, and Vega (2003, 2007), Boyd, Hu, and Jagannathan (2005), and Faust, Rogers, Wang, and Wright (2007). These papers have asked nuanced questions about various aspects of asset price reactions to MNAs. How are the reactions different for stocks versus bonds versus foreign exchange? Do asset prices react differently to “better”-than-expected news compared to “worse”-than-expected news? And do asset reactions

depend on whether the economy is in an expansionary state or a contractionary state?

My research agenda explores asset price reactions to news from a different angle by considering the effect of the ZLB and monetary policy. Prior work taking this approach is necessarily limited due to the recentness of the ZLB. Swanson and Williams (2013, 2014), as the first papers to study how the ZLB influences interest rate sensitivities to MNAs, focus attention to this research area. They find that the sensitivity of interest rates to macroeconomic data during the ZLB was lower than in normal times, though longer maturity interest rates were still surprisingly responsive to data. The results presented in this paper differ from those of Swanson and Williams (2013, 2014) in a few ways. First, I focus on how variation throughout the ZLB, specifically the expected duration of the ZLB, drives asset sensitivities to MNAs as opposed to comparing sensitivities during the ZLB to that during normal times. Second, I utilize a shadow rate term structure model to analytically derive comparative statics on interest rate reactions to macroeconomic news and how these reactions should depend on the expected duration of the ZLB. As such, I am able to relate the findings to the large literature on term structure models and test precise predictions for how interest rates should react to MNAs. Third, I use high-frequency intraday data on asset prices to measure the effect of data surprises on assets instead of daily data. By focusing on a tight ± 5 minute window around news announcements, I am able to more cleanly measure any asset reactions to data without worrying about events throughout the day that are unrelated to macroeconomic news. Finally, I analyze not only interest rates but also the stock market. I show through a decomposition how the reaction of interest rates to MNAs has direct implications for the reaction of the stock market to MNAs.

In related work, Raskin (2013) uses the reaction of interest rate options to macroeconomic data at the ZLB to assess how actions of the Federal Reserve altered perceptions of the monetary policy reaction function. Instead of focusing on specific Federal Reserve events such as the FOMC meeting in August, 2011 and testing how asset sensitivities change before and after these events, I choose to exploit the richer variation in asset sensitivities throughout the ZLB period.

The work of Swanson and Williams (2013, 2014) and Raskin (2013) suggest several reasons for why we might care about how the ZLB affects the sensitivity of assets to MNAs. One reason is that the sensitivity of interest rates to macroeconomic news measures the effectiveness of monetary policy. The existence of the ZLB prevents the Federal Reserve from using its traditional instrument of monetary policy, the federal funds rate, which is locked at the lower bound. A number of papers such as Eggertsson and Woodford (2003) argue that a central bank can still conduct monetary policy at the ZLB by influencing expectations of short rates in the future, which in turn can affect other asset prices and the broader economy. The Federal Reserve has indeed employed this strategy of forward guidance extensively such as with the previously mentioned action taken at the August, 2011 FOMC meeting. If interest rates are reasonably sensitive to MNAs despite the ZLB, then monetary policy can still be effective because it is still possible for monetary policy to meaningful move around expectations of short rates as well as longer-maturity rates. On the other hand, if interest rates are insensitive to MNAs at the ZLB, then the ZLB constrains the ability of monetary policy to influence interest rates.

The sensitivity of interest rates to MNAs is also informative about the effectiveness of fiscal policy. Work by Christiano, Eichenbaum, and Rebelo (2011) among others shows that the fiscal multiplier may be larger at the ZLB because interest rates do not “crowd out” fiscal spending as much. If interest rates are less sensitive to MNAs at the ZLB, then it may be true that a government spending shock does not increase interest rates much, so the fiscal multiplier is indeed larger. On the other hand, if interest rates are still reasonably sensitive to MNAs, the “crowding out” effect of rising interest rates in response to government spending may diminish the fiscal multiplier.

Finally, my results on the stock market’s reaction to macroeconomic news at the ZLB continue an avenue of inquiry on the relationship between monetary policy and the stock market. Papers that address this area include Rigobon and Sack (2003), Bernanke and Kuttner (2005), and Campbell, Pflueger, and Viceira (2014). Better understanding this relationship between stocks and monetary policy is important because monetary policy

often has to work through financial assets such as the stock market in order to ultimately influence the macroeconomy.

The organization of the paper is as follows. Section 3.2 introduces the shadow rate term structure model and links the model to three predictions about the sensitivity of interest rates to macroeconomic news surprises. Section 3.3 first describes the data for MNAs and high-frequency asset prices before discussing the basic regression that tests the effect of MNAs on asset prices. Section 3.4 takes the model predictions about the sensitivity of interest rates to the data and presents supporting evidence. Section 3.5 analyzes the sensitivity of stocks to macroeconomic data and performs a decomposition of this sensitivity. Section 3.6 concludes.

3.2 Shadow Rate Term Structure Model

I start by building a basic term structure model and analytically exploring its implications for the sensitivity of interest rates to MNAs. The presence of the ZLB prevents the use of the canonical Gaussian affine term structure model to establish relationships among interest rates at various maturities. As such, I follow the literature in utilizing a shadow rate term structure model that incorporates the ZLB. Introduced by Black (1995), the idea of the shadow rate has become particularly relevant since the 2008 financial crisis as the nominal short rate has remained stuck at the ZLB in the U.S. and other countries. Shadow rate term structure models have been approximately solved in both continuous time by Ichiue and Ueno (2013) and Krippner (2013) and in discrete time by Wu and Xia (2014).

The model I develop and analyze is a one-factor version of Wu and Xia (2014). I first set up the model and solve for the forward rate, which is the interest rate of interest. The focus of the analysis is not on the forward rate itself but rather on how the forward rate changes in response to news about macroeconomic fundamentals and the role of the ZLB therein. To clarify the role of the ZLB, I next show in the model that the expected duration of the ZLB increases when economic conditions deteriorate and/or when monetary policy is less responsive to the economy. Finally, I highlight three predictions of the model for the

sensitivity of interest rates to MNAs.

3.2.1 Setup and Forward Rate

Let the nominal short rate i_t equal the maximum of the shadow rate s_t and some lower bound \underline{i} :

$$i_t = \max \{s_t, \underline{i}\} \quad (3.1)$$

with $0 < \underline{i} \leq 0.25$ reflecting the reality that the Federal Reserve still pays interest on excess reserves. If the shadow rate is higher than the lower bound, the short rate equals the shadow rate; otherwise, Eq. (3.1) restricts the short rate from venturing below the lower bound.

Analogous to the eponymous Taylor rule discussed in Taylor (1993), the shadow rate is linear in a general state variable x_t :

$$s_t = \delta_0 + \delta_1 x_t. \quad (3.2)$$

The magnitude of δ_1 determines how strongly policy makers respond to the state variable in setting the short rate. For tractability purposes, I assume that x_t is a one-dimensional proxy for economic conditions. In a Taylor rule, there are two state variables corresponding to the output gap and inflation relative to some benchmark, so we can also think of x_t as reflecting some combination of the output gap and inflation. Low-dimensionality is justified by Bernanke and Boivin (2003) who show that a few factors (e.g., corresponding to principal components) can summarize the large number of data series in a data-rich environment.

The state variable x_t follows an AR(1) process under the physical measure (\mathbb{P}):

$$x_{t+1} = \mu + \rho x_t + \sigma_\epsilon \epsilon_{t+1}$$

with $\epsilon_{t+1} \sim \mathcal{N}(0, 1)$. Assuming that the price of risk λ_t is linear in the state variable (such that there is a constant risk premium term and a time-varying risk premium term)

$$\lambda_t = \lambda_0 + \lambda_1 x_t,$$

and the log stochastic discount factor is essentially affine

$$M_{t+1} = \exp\left(-i_t - \frac{1}{2}\lambda_t^2 - \lambda_t\epsilon_{t+1}\right),$$

the state variable x_t also follows an AR(1) process under the risk neutral measure (\mathbb{Q}):

$$x_{t+1} = \mu^{\mathbb{Q}} + \rho^{\mathbb{Q}}x_t + \sigma_\epsilon\epsilon_{t+1}^{\mathbb{Q}} \quad (3.3)$$

with $\epsilon_{t+1}^{\mathbb{Q}} \stackrel{\mathbb{Q}}{\sim} \mathcal{N}(0, 1)$.

Given this setup, the forward rate $f_{n,t}$, defined as the interest rate on a 1-period investment to be made at $t + n$, is approximately

$$f_{n,t} \approx \underline{i} + \sigma_n^{\mathbb{Q}} \left(\left(\frac{a_n + b_n x_t - \underline{i}}{\sigma_n^{\mathbb{Q}}} \right) \Phi \left(\frac{a_n + b_n x_t - \underline{i}}{\sigma_n^{\mathbb{Q}}} \right) + \phi \left(\frac{a_n + b_n x_t - \underline{i}}{\sigma_n^{\mathbb{Q}}} \right) \right). \quad (3.4)$$

Φ and ϕ are the standard normal CDF and pdf, respectively. Appendix C.1 sketches the derivation for the forward rate and provides expressions for a_n , b_n , and $\sigma_n^{\mathbb{Q}}$ in terms of model parameters.

3.2.2 Expected Duration of the ZLB

In addition to observing the forward rate, we can see how the model illustrates the extent to which the ZLB binds, loosely termed the expected duration of the ZLB. In particular, the probability the ZLB still binds at time $t + n$ is simply measured by the probability that the shadow rate at that time is still below the lower bound:

$$\Pr_t^{\mathbb{Q}} [s_{t+n} < \underline{i}] = \Phi \left(\frac{\underline{i} - \bar{a}_n - b_n x_t}{\sigma_n^{\mathbb{Q}}} \right). \quad (3.5)$$

Appendix C.1 expresses \bar{a}_n in terms of model parameters and shows the conditional distribution of s_{t+n} , which directly implies the above probability. The expected duration of the ZLB in the model can change either because economic conditions change or because the responsiveness of monetary policy changes.

Proposition 1 *As economic conditions worsen (x_t is more negative), the probability that the ZLB*

is binding in the future increases: $\partial/\partial x_t \left(Pr_t^Q [s_{t+n} < \underline{i}] \right) < 0$. Symmetrically, as monetary policy becomes less responsive to economic conditions (δ_1 is smaller), the probability that the ZLB is binding in the future increases: $\partial/\partial \delta_1 \left(Pr_t^Q [s_{t+n} < \underline{i}] \right) < 0$.

Appendix C.2 works out the comparative statics in the above proposition, but the intuition is fairly straightforward. First, as economic conditions worsen at t , the shadow rate is set lower in Eq. (3.2), which increases the probability that the shadow rate in the future remains below the lower bound. Second, consider a small Taylor rule coefficient in Eq. (3.2), which implies that the shadow rate is set close to δ_0 and largely independent of economic conditions. Assuming that $\underline{i} > \delta_0$ such that it is still possible for the ZLB to bind in this scenario, an unresponsive policy rule means that the shadow rate in the future is also set close to δ_0 and relatively less influenced by economic conditions then. This insight in turn explains the increased probability that the ZLB still binds.

3.2.3 Predictions for Forward Rate Sensitivity to MNAs

Finally, the model delivers three predictions on the sensitivity of interest rates to news about macroeconomic fundamentals and how this sensitivity interacts with the expected duration of the ZLB. This sensitivity is defined by $\partial f_{n,t}/\partial x_t$, which quantifies how a change in our proxy for economic conditions affects the forward rate defined in Eq. (3.4). Though I focus on the forward rate, the predictions below hold more generally. For example, the predictions hold for the yield to maturity of a zero-coupon bond, which we can mechanically see by expressing the yield to maturity in terms of a series of forward rates. Proofs for each of the three propositions below are in Appendix C.2.

Proposition 2 *News that economic conditions are “better”-than-expected (x_t is more positive) increases forward rates at all maturities: $\partial f_{n,t}/\partial x_t > 0$.*

Examples of “better”-than-expected macroeconomic news include higher-than-expected

inflation or output and lower-than-expected unemployment. Such a surprise increases the shadow rate today via Eq. (3.2) and, through the persistent AR(1) process for x_t , leads to a more positive state variable in the future that increases the shadow rate at that time as well. The resulting effect is higher interest rates.

Proposition 3 *As the ZLB is expected to bind for longer (probability that the shadow rate will be below the lower bound increases), whether because economic conditions are worse (x_t is more negative) or because monetary policy is less responsive to economic conditions (δ_1 is smaller), the sensitivity of forward rates to news about macroeconomic fundamentals decreases at all maturities: $\partial/\partial x_t (\partial f_{n,t}/\partial x_t) > 0$ and $\partial/\partial \delta_1 (\partial f_{n,t}/\partial x_t) > 0$.*

We know from Proposition 1 that poorer economic conditions and a less responsive monetary policy reaction function both lead to an increase in the expected duration of the ZLB. Because the shadow rate is more likely to remain below the lower bound, Eq. (3.1) says that the nominal short rate is more likely to stay anchored at \underline{i} . This anchoring of the nominal short rate prevents interest rates at all maturities from reacting as strongly to MNA surprises. Consider an edifying example in which the expected duration of the ZLB binds completely for a short time such that the nominal short rate is guaranteed to equal the lower bound within this time period. Then a small macroeconomic news surprise has no effect on short-dated interest rates, which are set as a constant. As I show in the data, this example corresponds to reality. Over the ZLB period, the Federal Reserve has credibly committed to setting the federal funds rate close to zero for at least a quarter or two. The near-perfect certainty of short maturity interest rates has resulted in essentially no sensitivity of these rates to data surprises.

Proposition 4 *The attenuation in the sensitivity of forward rates to news in Proposition 3 is greater for longer maturity forward rates (n is larger) provided that the ZLB is sufficiently binding (probability that the shadow rate will be below the lower bound is sufficiently large):*

$\partial/\partial n (\partial/\partial x_t (\partial f_{n,t}/\partial x_t)) > 0$ and $\partial/\partial n (\partial/\partial \delta_1 (\partial f_{n,t}/\partial x_t)) > 0$. Otherwise, if the ZLB does not bind sufficiently, the attenuation in the sensitivity of rates to news is decreasing in maturity n .

In other words, if the expected duration of the ZLB is sufficiently long, shorter maturity interest rates are essentially set equal to the lower bound and are thus insensitive to macroeconomic news. As the expected duration of the ZLB increases, the sensitivity of short rates to data has little room to decrease. The sensitivity of longer maturity interest rates, however, has more room to fall, since long rates are less bound by the ZLB to begin with.

Conversely, if the expected duration of the ZLB is relatively short, neither short-dated nor long-dated interest rates are significantly constrained by the ZLB. As the expected duration of the ZLB increases, the sensitivity of both types of interest rates has room to decrease. Since short rates represent rates of interest over shorter time periods while long rates represent rates of interest over longer time periods, an increase in the expected duration of the ZLB has a more noticeable impact on the former than on the latter. For example, the sensitivity of a 100-year bond to MNAs should not change much because the nominal short rate is expected to stay at zero for six months instead of three months. Thus, the magnitude of the sensitivity decrease is larger in short rates than in long rates.

3.3 Data and Basic Regressions

3.3.1 MNAs

Model predictions in hand, I turn to the data to assess how various assets react to macroeconomic news at the ZLB. Table 3.1 shows the full sample of MNAs considered in this paper from December, 2008 to March, 2014. For each data release, the table presents the name, units, number of observations, frequency, government agency or private-sector firm responsible, and timestamp. In all, I analyze 18 economic announcements, which cover the lion's share of important MNAs. All of the announcements occur once a month with the exception of Initial Jobless Claims, which is a weekly data release. The data in my sample

Table 3.1: Full sample of MNAs from 12/17/08 to 3/6/14

Event Name	Units	N	Freq.	Source	Time (ET)	Sign
Change in Nonfarm Payrolls	Thousands	62	M	BLS	8:30 AM	1
Consumer Confidence	Index	63	M	CB	10:00 AM	1
CPI	% ch. mom	62	M	BLS	8:30 AM	1
Durable Orders	% ch. mom	63	M	Census	8:30 AM	1
Existing Home Sales	Millions	63	M	NAR	10:00 AM	1
Factory Orders	% ch. mom	62	M	Census	10:00 AM	1
Housing Starts	Millions	60	M	Census	8:30 AM	1
Industrial Production	% ch. mom	62	M	Fed	9:15 AM	1
Initial Jobless Claims	Thousands	273	W	ETA	8:30 AM	-1
ISM Manufacturing	Index	63	M	ISM	10:00 AM	1
ISM Non-Manufacturing	Index	63	M	ISM	10:00 AM	1
Leading Index	Index	63	M	CB	10:00 AM	1
New Home Sales	Thousands	62	M	Census	10:00 AM	1
Pending Home Sales	% ch. mom	63	M	NAR	10:00 AM	1
Personal Income	% ch. mom	63	M	BEA	8:30 AM	1
PPI	% ch. mom	61	M	BLS	8:30 AM	1
Retail Sales	% ch. mom	62	M	Census	8:30 AM	1
Unemployment Rate	% rate	62	M	BLS	8:30 AM	-1

Notes: *Freq.* refers to monthly (M) or weekly (W). *Source* uses the following acronyms: BEA, Bureau of Economic Analysis; BLS, Bureau of Labor Statistics; CB, Conference Board; Census, Census Bureau; ETA, Employment and Training Administration; Fed, Federal Reserve Board of Governors; ISM, Institute for Supply Management; and NAR, National Association of Realtors.

are released at either 8:30 AM, 9:15 AM, or 10:00 AM ET.

While the economic announcements are officially produced and released by various government agencies (e.g., the Bureau of Labor Statistics), and private-sector firms (e.g., the National Association of Realtors), I collect the data from Bloomberg. Specifically, for each event, I download both the actual data release and the expected data release, the latter of which comes from a survey of economists. The difference between the actual number and the expected number constitutes the news that impacts asset prices.

The metric of the news surprise that I use in this paper is the conventional standardized news variable employed by Balduzzi, Elton, and Green (2001) as well as many subsequent papers in the literature of asset price reactions to MNAs. For economic indicator i , the

standardized news variable at time t is

$$S_{i,t} = \frac{A_{i,t} - \mathbb{E}_{t-} [A_{i,t}]}{\hat{\sigma}_i}. \quad (3.6)$$

$A_{i,t}$ is the actual data, $\mathbb{E}_{t-} [A_{i,t}]$ is the expectation of the data from the Bloomberg survey, and $\hat{\sigma}_i$ is the sample standard deviation of $A_{i,t} - \mathbb{E}_{t-} [A_{i,t}]$. Consistent with the model, I define the variable such that a positive value corresponds to “better”-than-expected data; i.e., higher-than-expected inflation or output and lower-than-expected unemployment. For this reason, I multiply the Initial Jobless Claims and the Unemployment Rate data by -1 while keeping the sign of the other data unchanged, as shown by the sign column of Table 3.1. As the name suggests, the standardized news variable provides a single metric that is standardized across different types of news about macroeconomic fundamentals and allows for comparability. This is important because different types of news are released with different units.

The use of the standardized news variable in conjunction with data from Bloomberg is common in testing the effect of macroeconomic news surprises on asset prices. Many financial market participants use Bloomberg to get a sense of the consensus forecast for any given MNA and to see how the actual data compares to the forecasted data at the time of the release. The survey expectation is furthermore unbiased and unlikely to be stale, since economists can adjust their forecasts until the very last moment.

3.3.2 High-Frequency Asset Prices

I consider three different futures contracts for the asset prices in my analysis: Eurodollar Futures (ED), 10-Year U.S. Treasury Note Futures (TY), and E-mini S&P 500 Futures (ES). I obtain intraday tick data from Tick Data, a data vendor. The use of futures data is important because many MNAs are released early in the morning before equity markets officially open at 9:30 AM ET. Whereas some markets are less liquid at that early hour, futures markets are already active. Not surprisingly, financial market participants also tend to react to data surprises by trading in these futures directly.

I use ED and TY futures data to evaluate the impact of MNAs on interest rates. ED futures settle at maturity based on the spot 3-month LIBOR rate, so the ED future with a certain maturity corresponds roughly to the 3-month forward rate at that maturity. I construct series of generic ED contracts with different maturities in a straightforward manner. For ED_n , the n -quarter(s) out ED future is simply the n th contract with $n = 1, 2, \dots, 16$. Each contract is rolled over a few days before maturity. The ED1 contract thus corresponds roughly to the 3-month forward rate that is 1 quarter out, while the ED16 contract is the 3-month forward rate that is 16 quarters out. By considering how the full range of ED1 to ED16 futures responds to MNA surprises, I am able to distinguish interest rate reactions at different maturities. Since the longest maturity I consider for ED futures is 16 quarters or approximately 4 years, I also analyze the front TY futures, which corresponds to a longer-dated interest rate.

Assessing the reaction of the stock market to macroeconomic news surprises is more straightforward: I directly use the front ES futures contract.

3.3.3 Basic Regressions for All Assets

The motivating regression in evaluating the impact of a MNA surprise on asset prices is simply to regress asset returns R_t in a window around event i on the $S_{i,t}$ standardized news variable defined in Eq. (3.6):

$$R_t = \alpha_i + \beta_i S_{i,t} + \epsilon_t. \quad (3.7)$$

β_i measures the reaction of the asset to a news surprise and is the object of focus. The availability of high-frequency asset prices allows me to set the event window to ± 5 minutes around the event, which is an interval comparable to other studies that have used intraday return data. It is thus likely that asset prices vary over the window only due to any surprise embedded in the news announcement.

Table 3.2 documents the $\hat{\beta}_i$ from estimating Eq. (3.7) for various asset-MNA combinations over the ZLB date range. For example, consider Housing Starts and the ED16 futures contract. A one unit positive standardized news surprise in Housing Starts increases the

Table 3.2: Reactions to news during the ZLB

Event Name	ED1	ED4	ED8	ED16	TY	ES
	$\hat{\beta}$ (bps)	$\hat{\beta}$ (bps)	$\hat{\beta}$ (bps)	$\hat{\beta}$ (bps)	$\hat{\beta}$ (%)	$\hat{\beta}$ (%)
Change in NFP	-0.340**	-1.58	0.110	2.511	-0.116	0.158
Consumer Confidence	0.059	0.486***	0.855***	0.868***	-0.053***	0.196***
CPI	0.017	0.316*	0.494**	0.657**	-0.037***	-0.046
Durable Orders	-0.004	0.332**	0.834***	1.383***	-0.058***	0.123***
Existing Home Sales	0.027	0.264**	0.579***	0.704***	-0.036***	0.143***
Factory Orders	0.037	0.148	0.239	0.347*	-0.016	0.035*
Housing Starts	-0.066	0.393***	1.110***	2.153***	-0.080***	0.095***
Industrial Production	0.032	0.132	0.288**	0.323**	-0.017**	0.052***
Initial Jobless Claims	0.007	0.293***	0.563***	0.561**	-0.037***	0.088***
ISM Manufacturing	0.017	0.608***	1.451***	2.065***	-0.110***	0.228***
ISM Non-Manufacturing	0.061	0.893***	1.628***	2.004***	-0.098***	0.162***
Leading Index	0.088*	0.148	0.305*	0.187	-0.014	0.074**
New Home Sales	-0.051	0.835***	1.410***	1.523***	-0.084***	0.237***
Pending Home Sales	-0.015	0.327***	0.690***	0.822***	-0.037***	0.101***
Personal Income	0.034	0.090	0.054	0.118	-0.006	0.019**
PPI	0.052	0.552***	0.662**	0.694*	-0.039**	0.041
Retail Sales	0.222***	0.992***	1.672***	1.659***	-0.096***	0.215***
Unemployment Rate	0.014	0.294	0.560	0.980	-0.027	0.037
All	0.017	0.268***	0.617***	0.831***	-0.043***	0.097***

Notes: Regressions of asset returns on standardized news for each event i and then grouping all events together (last row): $R_t = \alpha_i + \beta_i S_{i,t} + \epsilon_t$. The left-hand side variable R_t is the asset return in a ± 5 minute window around a MNA (basis points change in the forward rate associated with ED futures or % change in the price level associated with TY and ES futures). The right-hand side variable $S_{i,t}$ is defined as the actual data less the expected data based on a Bloomberg survey of economists, all normalized by the sample standard deviation of the actual less the expected difference: $S_{i,t} = (A_{i,t} - \mathbb{E}_{t-}[A_{i,t}]) / \hat{\sigma}_i$. t -statistics (not shown) are based on heteroskedasticity-consistent standard errors. *** denotes significance at 1%, ** denotes significance at 5%, and * denotes significance at 10%.

3-month forward rate 4 years out approximately 2.15 bps, which is both statistically and economically significant. Examining Table 3.2, we clearly see that the reaction of ED1 to a news surprise is muted, as $\hat{\beta}_i$ is generally small and statistically insignificant. This finding is not surprising given that interest rates are almost certain to remain fixed near zero for the next several months due to the ZLB. Results for ED4, ED8, and ED16 futures contracts show that the reactions of longer-dated interest rates to a news surprise is both statistically and economically significant and in the predicted direction: a positive standardized news surprise increases interest rates. The column corresponding to the TY futures tells the same story: the TY futures data is left in terms of prices instead of yields, so a positive news surprise that increases yields decreases prices. The final column of Table 3.2 shows that a positive data surprise also has a positive, significant impact on the stock market across a broad range of economic indicators.

Finally, the last row of Table 3.2 displays the results of running Eq. (3.7) grouping together all 18 MNAs. A one unit positive standardized news surprise in any macroeconomic announcement leads to a 0.017 bps, 0.268 bps, 0.617 bps, and 0.831 bps increase in the forward rate implied from the ED1, ED4, ED8, and ED16 contracts, respectively. This same surprise leads the TY futures contract to decrease by 0.043% and the ES futures contract to increase by 0.097%.

3.4 Empirical Test of Model Predictions for Interest Rates

While Table 3.2 shows the responses of interest rates to macroeconomic news surprises, the table is insufficient to formally demonstrate the three predictions of the shadow rate term structure model. In particular, Table 3.2 does not make use of time variation in the sensitivity of interest rates to MNAs; rather, it presents one number to summarize the estimated sensitivity of a given interest rate to economic data over the entire ZLB period. I formally test the model predictions for interest rates by focusing on ED futures and using TY futures as a robustness check.

3.4.1 Methodology

For an ED future with maturity n , I perform a daily 1-year rolling regression over the ZLB period analogous to Eq. (3.7):

$$\Delta f_{n,t} = \alpha_n + \beta_n S_t + \epsilon_{n,t}, \quad (3.8)$$

with the regression grouping together all 18 MNAs identical to the last row of Table 3.2. The result of the rolling regression is a time series of $\hat{\beta}_{n,t}$, the estimated sensitivity of the 3-month forward rate n -quarter(s) out over a 1-year window centered at t . By grouping together all the MNAs together, Eq. (3.8) artificially imposes the constraint that $f_{n,t}$ reacts in the same manner to surprises in different data releases. This constraint is clearly an oversimplification; e.g., it is well known that the Change in NFP (nonfarm payrolls) data release is the most closely watched economic indicator with the greatest impact on financial markets. The upshot of forcing interest rates to react the same to different macroeconomic series is the greater statistical power from estimating Eq. (3.8) over more observations. Furthermore, the time series of $\hat{\beta}_{n,t}$ from estimating the regression over any specific data release without grouping has a similar shape to the time series of $\hat{\beta}_{n,t}$ from estimating the regression over all 18 MNAs. Swanson and Williams (2014) suggest a nonlinear least squares regression to measure the time-varying sensitivity of interest rates in a way that does allow for differential reactions to different macroeconomic news. Their methodology imposes that the relative magnitudes of the sensitivity of interest rates to different data releases are constant over time. I choose to not employ their technique for two reasons. The first is that their method requires choosing a benchmark time period when interest rates react “normally” to macroeconomic data. The choice of this benchmark period can affect the sensitivity of interest rates to data outside of the benchmark period. More importantly, the time-varying sensitivity of interest rates from simply grouping together all the MNAs as I have done is similar in shape to that produced by the Swanson and Williams (2014) methodology.

Using the time series of $\hat{\beta}_{n,t}$, I regress $\hat{\beta}_{n,t}$ on a proxy for the expected duration of the

ZLB:

$$\hat{\beta}_{n,t} = \gamma_n + \delta_n ZLB_t + \eta_{n,t}. \quad (3.9)$$

δ_n measures how the estimated sensitivity of $f_{n,t}$ to MNAs changes as the expected duration of the ZLB changes. The three model predictions then map neatly to the following testable implications.

Proposition 2 implies that $\hat{\beta}_{n,t} > 0$ for all n and for all 1-year windows centered at t : positive news surprises increase interest rates. This result is well-established empirically, and we have already seen some evidence in Table 3.2. In the table, forward rates implied from ED contracts clearly jump higher in reaction to positive macroeconomic news. At the same time, prices fall (yields rise) of TY futures. The rolling regression methodology in this section confirms the prediction for all 1-year windows.

Proposition 3 implies that $\hat{\delta}_n < 0$ for all n assuming that ZLB_t is a positive proxy for the expected duration of the ZLB. As the ZLB is expected to bind for longer, which could be because economic conditions are worse or because monetary policy is less responsive to economic conditions, the sensitivity of interest rates to MNAs decreases. If ZLB_t is a negative proxy for the expected duration of the ZLB, then the prediction mechanically switches to $\hat{\delta}_n > 0$.

Proposition 4, the last of the three predictions, implies that $|\hat{\delta}_n|$, the magnitude of $\hat{\delta}_n$, is hump-shaped in n (first increasing, then decreasing). Consider four interest rates with the following maturities: $n_1 < n_2 \ll n_3 < n_4$. The attenuation in the sensitivity of interest rates to data is greater for the n_2 -maturity rate than for the n_1 -maturity rate because the sensitivity of the n_1 -maturity rate is more constrained by the ZLB and has little room to decrease. In contrast, the sensitivity of the n_2 -maturity rate has more room to fall as the ZLB becomes more binding. The attenuation in the sensitivity of interest rates to data is smaller for the n_4 -maturity rate than for the n_3 -maturity rate because neither rates are greatly constrained by the ZLB to begin with. Both the n_3 - and n_4 -maturity rates encompass a time period over which the ZLB is binding. As the expected duration of the ZLB increases incrementally, the ZLB is binding for more of the period encompassed by the n_3 -maturity rate than for

that encompassed by the n_4 -maturity rate, which explains why the sensitivity of the former is attenuated more than that of the latter. For example, if the expected duration of the ZLB increased s.t. all rates out to n_3 flattened to the lower bound, the sensitivity of the n_3 -maturity rate would drop to zero, while the sensitivity of the n_4 -maturity rate would still be non-zero.

3.4.2 Survey Proxy for Expected Duration of the ZLB

Figure 3.2 illustrates the results of applying the abovementioned methodology. In particular, the figure shows the time-varying sensitivities of various interest rates to MNAs and how the expected duration of the ZLB affects these sensitivities. Focusing first on Panels A, B, C, and D, which correspond to the interest rates associated with the ED1, ED4, ED8, and ED16 futures contracts, respectively, the right-hand plots shows two time series. “fcsts,” plotted on the right y-axis, is the same variable shown in the right-hand plot of Figure 3.1: the median anticipated number of quarters until the first rate hike based on primary dealer surveys by the FRBNY. Since the FRBNY surveys primary dealers only once every month or every other month, I create a daily variable with a simple forward fill. The result is the “fcsts” variable that I use as the survey proxy for the expected duration of the ZLB; i.e., ZLB_t in Eq. (3.9). The second time series plotted in each right-hand plot is the time-varying sensitivity of a given interest rate to MNAs $\hat{\beta}_{n,t}$ in Eq. (3.9). The FRBNY started publishing primary dealer surveys in January, 2011, so I necessarily restrict my analysis such that the earliest t in Eq. (3.9) is 1/1/11. For each panel, the left-hand plot is simply a scatterplot of “beta” on “fcsts” that graphically demonstrates the relationship in Eq. (3.9). Each point in the scatterplot is color-coded based on the timestamp: red for 12/17/08 to 8/9/11 corresponding to the period of general forward guidance (“fwd”); green for 8/10/11 to 12/12/12 corresponding to the period of calendar guidance (“cdr”); and blue for 12/13/12 to 3/6/14 corresponding to the period of data thresholds (“dat”).

Verifying Proposition 2, we see in Panels A to D of Figure 3.2 that $\hat{\beta}_{n,t} > 0$: the dots in the left-hand scatterplots, or equivalently the time-series of “beta” in the right-hand plots,




Figure 3.2: Time-varying sensitivities of interest rates to MNAs and the impact of the ZLB using the survey proxy

Notes: See Appendix C.3.1. Panels A and B (1st row), C and D (2nd row), and E (3rd row) correspond to ED1, ED4, ED8, ED16, and TY, respectively.

are generally positive. Positive MNA surprises increase interest rates at all maturities and at all points in time. That some of the $\hat{\beta}_{n,t}$ are essentially at zero is entirely plausible and corresponds with the second prediction that a binding ZLB reduces the sensitivity of interest rates to MNAs. In the few cases that $\hat{\beta}_{n,t}$ is negative (Panel A of Figure 3.2), the magnitudes are economically small and statistically indistinguishable from zero.

Verifying Proposition 3, we see in the scatterplots that there is a negative relationship between $\hat{\beta}_{n,t}$ and ZLB_t ; i.e., $\hat{\delta}_n < 0$ in Eq. (3.9). Table 3.3 presents the $\hat{\delta}_n$ for $n = 1, \dots, 16$: all the $\hat{\delta}_n < 0$ and are statistically significant. As primary dealers expect the ZLB to bind for longer and the first rate hike to be later, the sensitivity of interest rates at all maturities to MNAs becomes smaller. This link is particularly apparent in Panels B, C, and D for longer-maturity interest rates. In Panel A, on the other hand, the negative relationship is less clear for a simple reason: the ZLB binds to such an extent that the 1 quarter out 3-month forward rate is essentially fixed at the lower bound, and $\hat{\beta}_{1,t}$ is therefore approximately zero. Nonetheless, $\hat{\delta}_1$ is negative and significantly different from zero in Table 3.3. The economic magnitudes of $\hat{\delta}_n$ are significant as well. For example, $\hat{\delta}_8 = -0.116$ in Table 3.3 implies that a one quarter increase in the expected duration of the ZLB lowers the sensitivity of the 8 quarters out 3-month forward rate 0.116 bps. As the scatterplot in Panel C shows, increasing the expected duration of the ZLB from 0 quarters to 10 quarters thus flattens the sensitivity of this forward rate from around 1.2 bps to a unit of standardized news to essentially nothing. Figure 3.2 also provides color on the effect of a monetary policy move such as the introduction of calendar guidance in August, 2011, a date marked by the first vertical line in the right-hand plots. By declaring that the federal funds rate would be held at exceptionally low levels until mid-2013, the Federal Reserve increased primary dealers' expected duration of the ZLB approximately 4 quarters, from around 5 quarters to around 9 quarters. As predicted, the sensitivity of interest rates to MNAs subsequently plummeted.

Finally, verifying Proposition 3, we see in Table 3.3 that the magnitude of $\hat{\delta}_n$ is first increasing in n and then decreasing. The attenuation in the sensitivity of interest rates to macroeconomic news is greatest for intermediate maturity rates, as the ZLB has already

Table 3.3: Attenuation in the sensitivity of ED futures due to the ZLB using the survey proxy

Rate	$\hat{\delta}_n$	Rate	$\hat{\delta}_n$	Rate	$\hat{\delta}_n$	Rate	$\hat{\delta}_n$
ED1	-0.007***	ED5	-0.079***	ED9	-0.124***	ED13	-0.093***
ED2	-0.003***	ED6	-0.107***	ED10	-0.122***	ED14	-0.094***
ED3	-0.027***	ED7	-0.112***	ED11	-0.117***	ED15	-0.078***
ED4	-0.050***	ED8	-0.116***	ED12	-0.091***	ED16	-0.082***

Notes: First-stage regression is a daily 1-year rolling regression over the ZLB period that aggregates all 18 MNAs: $\Delta f_{n,t} = \alpha_n + \beta_n S_t + \epsilon_{n,t}$. $\Delta f_{n,t}$ is the basis points change in the forward rate implied from the EDn contract in a ± 5 minute window around a MNA. The right-hand side variable S_t aggregates $S_{i,t}$ for MNA i . $S_{i,t}$ is defined as the actual data less the expected data based on a Bloomberg survey of economists, all normalized by the sample standard deviation of the actual less the expected difference: $S_{i,t} = (A_{i,t} - \mathbb{E}_{t-}[A_{i,t}]) / \hat{\sigma}_i$. Second-stage regression takes the estimated $\hat{\beta}_{n,t}$ corresponding to the first-stage regression over a 1-year window centered at t and regresses $\hat{\beta}_{n,t}$ on ZLB_t : $\hat{\beta}_{n,t} = \gamma_n + \delta_n ZLB_t + \eta_{n,t}$. ZLB_t is the median number of quarters until the first rate hike based on primary dealer surveys by the FRBNY. t -statistics (not shown) are based on heteroskedasticity-consistent standard errors. *** denotes significance at 1%, ** denotes significance at 5%, and * denotes significance at 10%.

constrained the sensitivity of short-term rates yet is not particularly relevant for long-term rates.

Panel E of Figure 3.2 provides a sanity check of the results using the TY futures contract instead of ED futures. The main difference is that the left-hand-side variable $\Delta f_{n,t}$ in Eq. (3.8) is replaced by R_t , the % change in the price level of the TY contract in the ± 5 minute window around a MNA. This switch reverses the signs of Propositions 2 and 3, as we now expect $\hat{\beta}_{n,t} < 0$ and $\hat{\delta}_n > 0$. This is indeed what Panel E demonstrates: positive MNA surprises decrease TY prices (increase TY yields), and the extent of this reaction declines as the expected duration of the ZLB increases.

3.4.3 Fed Funds Futures Proxy for Expected Duration of the ZLB

Figure 3.3 provides additional support for the model predictions using a different proxy for the ZLB: the implied FF rate from the FF24 futures contract, which is also one of the variables shown in the left-hand plot of Figure 3.1. The only difference between Figures 3.2 and 3.3 is that “FF24” replaces “fcsts” as a plotted variable in the right-hand plots and as the independent variable in the scatterplots. As the expected duration of the ZLB increases,




Figure 3.3: Time-varying sensitivities of interest rates to MNAs and the impact of the ZLB using the FF24 proxy

Notes: See Appendix C.3.2. Panels A and B (1st row), C and D (2nd row), and E (3rd row) correspond to ED1, ED4, ED8, ED16, and TY, respectively.

the FF rate 24 months out should decrease. For example, after the introduction of calendar guidance in August, 2011, “FF24” decreases significantly consistent with the increase in the surveyed time until the first rate hike. Though not as direct a measure of the expected duration of the ZLB, the implied FF rate is never a stale proxy in the way that the survey proxy for the ZLB could be due to the infrequency of surveys. The implied FF rate is also available before January, 2011, which allows the analysis to begin further back in time, though Figure 3.3 covers the same period as Figure 3.2 for the sake of comparability.

Since the $\hat{\beta}_{n,t}$ are the same in Figures 3.2 and 3.3, Panels A to D of Figure 3.3 support Proposition 2 in the same way: $\hat{\beta}_{n,t} > 0$ for most (n, t) pairs and is otherwise indistinguishable from zero. The relationship between $\hat{\beta}_{n,t}$ and ZLB_t in Figure 3.3 is positive in the scatterplots, which corroborates Proposition 3. The reason that Proposition 3 implies $\hat{\delta}_n > 0$ in Figure 3.3 is that the implied FF rate from the FF24 futures contract is a negative proxy for the expected duration of the ZLB: as the latter increases, the former decreases. Table 3.4 produces the quantitative estimates of $\hat{\delta}_n$ from performing the regression in Eq. (3.9). I drop the constant (restrict $\gamma_n = 0$) based on the economic intuition that if “FF24” were to equal zero, then the ZLB must be acutely binding and the 3-month forward rate n -quarter(s) out must be fixed at the lower bound. Consequently, these interest rates must be insensitive to MNAs. Though not explicitly stated, this logic is also at play in Figure 3.2. If primary dealers expect the first rate hike to be sufficiently far in the future, the predicted sensitivity of a given interest rate to macroeconomic news is zero, because the downward sloping regression line intersects the x-axis. As we see in Table 3.4, the $\hat{\delta}_n$ are greater than zero with statistically and economically significant magnitudes. A 100 bps drop in the implied FF rate decreases the reaction of the 16 quarters out 3-month forward rate to a unit of standardized news by 1.285 bps.

A seemingly incongruent feature of Table 3.4 is that the magnitude of $\hat{\delta}_n$ is monotonically increasing in n , which is at odds with the hump-shaped pattern implied by Proposition 4. That is, Table 3.4 suggests that the ZLB attenuates the sensitivity to MNAs of long-dated rates more than that of shorter-dated rates. This result is in fact supportive of Proposition 4.

Table 3.4: Attenuation in the sensitivity of ED futures due to the ZLB using the FF24 proxy

Rate	$\hat{\delta}_n$	Rate	$\hat{\delta}_n$	Rate	$\hat{\delta}_n$	Rate	$\hat{\delta}_n$
ED1	0.019***	ED5	0.463***	ED9	0.988***	ED13	1.231***
ED2	0.056***	ED6	0.624***	ED10	1.067***	ED14	1.251***
ED3	0.155***	ED7	0.750***	ED11	1.124***	ED15	1.263***
ED4	0.304***	ED8	0.883***	ED12	1.185***	ED16	1.285***

Notes: First-stage regression is a daily 1-year rolling regression over the ZLB period that aggregates all 18 MNAs: $\Delta f_{n,t} = \alpha_n + \beta_n S_t + \epsilon_{n,t}$. $\Delta f_{n,t}$ is the basis points change in the forward rate implied from the EDn contract in a ± 5 minute window around a MNA. The right-hand side variable S_t aggregates $S_{i,t}$ for MNA i . $S_{i,t}$ is defined as the actual data less the expected data based on a Bloomberg survey of economists, all normalized by the sample standard deviation of the actual less the expected difference: $S_{i,t} = (A_{i,t} - \mathbb{E}_{t-}[A_{i,t}]) / \hat{\sigma}_i$. Second-stage regression takes the estimated $\hat{\beta}_{n,t}$ corresponding to the first-stage regression over a 1-year window centered at t and regresses $\hat{\beta}_{n,t}$ on ZLB $_t$: $\hat{\beta}_{n,t} = \gamma_n + \delta_n \text{ZLB}_t + \eta_{n,t}$ with the constraint that $\gamma_n = 0$. ZLB $_t$ is the implied FF rate (%) from the FF24 futures contract. t -statistics (not shown) are based on heteroskedasticity-consistent standard errors. *** denotes significance at 1%, ** denotes significance at 5%, and * denotes significance at 10%.

The reason is that by setting $\gamma_n = 0$ in Eq. (3.9), I anchor each scatterplot with a point at the origin that corresponds to the ZLB binding “forever” such that the sensitivity of interest rates to data is zero. As “FF24” decreases, the sensitivity of short-dated rates falls quickly to zero and remains there, which leads to a small magnitude of $\hat{\delta}_n$. As “FF24” decreases, the sensitivity of long-dated rates initially remains above zero. As “FF24” nears zero and the ZLB binds for an increasingly longer time, the sensitivity of these long-dated rates must eventually collapse to zero by construction, and the suddenness of this attenuation is reflected in the larger magnitude of $\hat{\delta}_n$. The pattern in $\hat{\delta}_n$ thus comes from using the linear specification in Eq. (3.9) to analyze the nonlinear relationship between “beta” and “FF24”. In results not shown, estimating an unconstrained Eq. (3.9) without restricting $\gamma_n = 0$ does lead to the magnitude of $\hat{\delta}_n$ having a hump-shaped pattern in n .

As in Figure 3.2, Panel E of Figure 3.3 provides a sanity check of the results using the TY futures contract instead of ED futures. We see that $\hat{\beta}_{n,t} < 0$ and $\hat{\delta}_n < 0$ as expected: positive MNA surprises decrease TY prices (increase TY yields), and the extent of this reaction declines as the implied FF rate decreases.

3.5 Stock Market's Sensitivity to MNAs

3.5.1 Methodology and Results

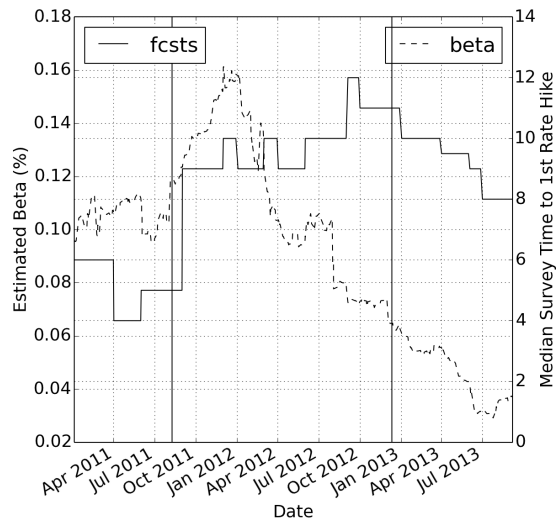
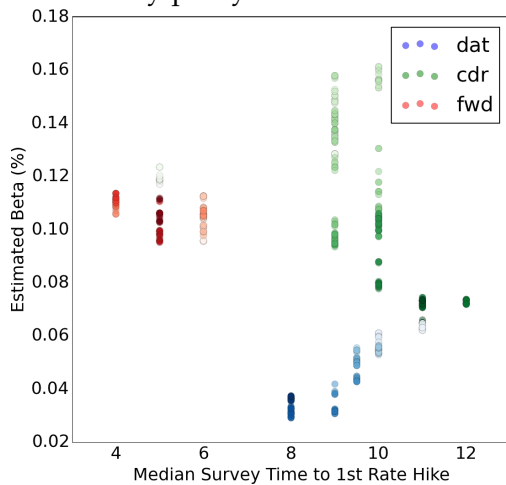
Having looked at the sensitivity of interest rates to MNAs and how the ZLB may have affected this sensitivity, I turn the analysis to another important asset: the stock market. I first use the empirical framework already presented for interest rates and let the data speak for itself. That is, I run a daily 1-year rolling regression over the ZLB period replacing the left-hand variable of Eq. (3.8) with R_t , the % change in the price level of the ES contract in the ± 5 minute window around a MNA. Using the resulting time series of $\hat{\beta}_{s,t}$, I run Eq. (3.9) using the two proxies for the expected duration of the ZLB: the median surveyed time until the first rate hike and the implied FF rate from the FF24 futures contract.

Analogous to Figures 3.2 and 3.3, Figure 3.4 shows the results of this methodology for stocks. Panel A presents findings using the survey proxy “fcsts” for ZLB_t , and Panel B uses the future FF rate “FF24” for ZLB_t . Focusing on the “beta” in the right-hand plots, we see that from early 2011 to mid-2012, stocks reacted on the order of 10 to 15 bps for every unit of standardized news. Starting in mid-2012, the sensitivity of stocks started dropping significantly, reaching a level of around 2 bps for every unit of standardized news in 2014. The scatterplots do not show a clear relationship between “beta” and the expected duration of the ZLB. Up until mid-2012, the sensitivity of stocks to MNAs seemed to increase as the ZLB became more binding. The subsequent sharp drop in the sensitivity of stocks to data does not have a clear explanation, however, for the expected duration of the ZLB stayed stable before moderately decreasing.

3.5.2 Decomposition: Interest Rate News vs. Cash Flow and Risk Premium News

To better understand the time-varying sensitivity $\hat{\beta}_{s,t}$ of stocks to economic data in Figure 3.4, I utilize a standard decomposition of stock returns. When stocks jump in response to economic data, we know based on the work of Campbell (1991) and Campbell and

Panel A: Survey proxy for the ZLB



Panel B: FF24 proxy for the ZLB





Figure 3.4: Time-varying sensitivities of stocks to MNAs and the impact of the ZLB

Notes: See Appendix C.3.3.

Ammer (1993) that the reaction could be due to any combination of a change in expected future dividends (cash flow news), a change in expected future interest rates (interest rate news), and a change in expected future excess returns (risk premium news). To see this decomposition algebraically, I follow Boyd, Hu, and Jagannathan (2005) in using a simple Gordon growth model:

$$P = \frac{D(1+G)}{R+\Pi-G} \quad (3.10)$$

with P the price level of the stock market, D the current dividends of constituent stocks, G a constant growth rate in the dividends, R a long-dated risk-free rate, and Π the equity risk premium. Defining S as the surprise in a given MNA, $(dP/P)/dS$ is the percentage change in the stock market in response to a data surprise and maps to $\hat{\beta}_{s,t}$. Using the Gordon growth model, we see that

$$\begin{aligned} \frac{1}{P} \frac{dP}{dS} &= \frac{1}{P} \frac{D \frac{dG}{dS} (R+\Pi-G) - \left(\frac{dR}{dS} + \frac{d\Pi}{dS} - \frac{dG}{dS} \right) D(1+G)}{(R+\Pi-G)^2} \\ &= -\frac{P}{D} \frac{1}{1+G} \left(\frac{dR}{dS} + \frac{d\Pi}{dS} - \left(1 + \frac{D}{P}\right) \frac{dG}{dS} \right) \\ &\approx -\frac{P}{D} \left(\frac{dR}{dS} + \frac{d\Pi}{dS} - \frac{dG}{dS} \right) \end{aligned} \quad (3.11)$$

with the approximation in the last step holding for small G and small D/P . As qualitatively stated previously, the reaction of stocks to macroeconomic news is due to a combination of cash flow news (dG/dS term), interest rate news (dR/dS term), and risk premium news ($d\Pi/dS$ term).

Focusing on the interest rate channel, the interest rate news term

$$N_R \equiv -\left(\frac{P}{D}\right) \left(\frac{dR}{dS}\right) \quad (3.12)$$

has two components: P/D , the price-to-dividend ratio, and dR/dS , the sensitivity of interest rates to MNAs. In general, the interest rate channel attenuates the sensitivity of stocks to economic data such that the absence of this channel would increase the magnitude of the stock market reaction to data. To see this, consider a positive MNA surprise. Based on

Figure 3.4 and Table 3.2 in addition to the literature on the reaction of stocks to MNAs, we know that positive MNA surprises tend to increase the stock market, so the left-hand side of Eq. (3.11) is positive. We also know from Proposition 2 that positive MNA surprises increase interest rates, so $dR/dS > 0$. Thus the interest rate news term $N_R < 0$, which reduces the magnitude of the stock market reaction to positive MNA surprises. The intuition is simply that higher interest rates discount dividends more heavily, which reduces stock prices. The exact same reasoning works for negative MNA surprises: stocks tend to decrease ($(dP/P)/dS < 0$) as do interest rates ($dR/dS < 0$), so the interest rate news term ($N_R > 0$) has an offsetting effect on the stock market due to the discounting of dividends at lower interest rates.

If the reaction of interest rates to MNAs offsets the reaction of stocks to MNAs, logic dictates that, all else equal, when the sensitivity of interest rates to MNAs is lower, the sensitivity of stocks to MNAs is higher. Since we have already seen from Proposition 3 that the sensitivity of interest rates to macroeconomic news decreases as the ZLB is expected to bind for longer, we might expect to see that longer expected duration of the ZLB increases the sensitivity of stocks to macroeconomic news. Turning to the data, the right-hand plots of Figure 3.4 suggest that this positive relationship between the expected duration of the ZLB and the sensitivity of stocks holds only from early 2011 to mid-2012 (“early period”) but not from mid-2012 to the end of the sample (“late period”). In the early period, the surveyed time until the first rate hike rose and the implied FF rate 24 months out fell, so the sensitivity of interest rates to data decreased. At the same time, the sensitivity of stocks to data increased, as we would expect if the interest rate news term were dominant in Eq. (3.11). In the late period, the expected duration of the ZLB was stable before decreasing modestly, which resulted in a stable and then modestly increasing sensitivity of interest rates to data. The sensitivity of the stock market to data, however, decreased continuously and precipitously throughout this period.

We can actually calculate the magnitude of the interest rate news term N_R in the data. To proxy for P/D in Eq. (3.12), I use the price-to-dividend ratio of the S&P 500 shown as

“dp” in the right-hand plot of Figure 3.5 (right y-axis). dR/dS in Eq. (3.12) is the number of basis points that R changes by due to a data surprise. Since R is the interest rate on a long-term risk-free bond, I measure dR/dS using the sensitivity of the forward rate from the ED16 contract; i.e., “beta” $\hat{\beta}_{16,t}$ in Panel D of Figures 3.2 and 3.3. The right-hand plot of Figure 3.5 plots $\hat{\beta}_{16,t}$ as “betas_ED16” (left y-axis). I use the ED16 contract because it is the longest-maturity ED futures contract in my data, and $\hat{\beta}_{16,t}$ is conveniently already measured in basis points. Note that while the TY futures contract corresponds to a longer maturity interest rate than the ED16 contract, converting TY prices to yields is a complicating step. Figure 3.5 plots the interest rate news term in the right y-axis of the left-hand plot as “only_R.” We see a hump-shaped pattern in the interest rate news term, consistent with our intuition. The expected duration of the ZLB increased in the early period before decreasing in the late period, which resulted in the interest rate news term decreasing in magnitude (from 75 bps to 25 bps per unit of standardized news) and then increasing (back to 75 bps). It is clear that if only the interest rate news term were time-varying in Eq. (3.11), we would expect a positive relationship between the expected duration of the ZLB and the reaction of stocks to MNAs.

We cannot, however, ignore the reality that the change in the stock market in response to a data surprise also depends on cash flow and risk premium news. Time variation in the two former terms must offset time variation in interest rate news to produce the final time-varying sensitivity of stocks to data. I back out the cash flow and risk premium news terms by subtracting the interest rate news term from the overall reaction of stocks to data:

$$N_{CF+RP} \equiv \frac{1}{P} \frac{dP}{dS} - N_R \approx -\frac{P}{D} \left(\frac{d\Pi}{dS} - \frac{dG}{dS} \right).$$

In the data, N_{CF+RP} is just the $\hat{\beta}_{s,t}$ of stocks, plotted as “beta” in the right-hand graph of Figure 3.5, less the “only_R” interest rate news term. The left-hand graph of Figure 3.5 plots the resulting cash flow and risk premium news terms as “PI_G” on the right y-axis. As expected, these two news terms demonstrate time-varying sensitivity to MNAs; in particular, a hump-shaped pattern that is almost a reflection of the hump-shaped pattern in the interest




Figure 3.5: Decomposition of the stock market reaction to MNAs into interest rate news versus cash flow and risk premium news

Notes: The right-hand plot shows two time series. “dp,” plotted on the right y-axis, is the dividend-to-price ratio of the S&P 500. “beta_ED16,” plotted on the left y-axis, is the sensitivity of the forward rate from the ED16 contract; i.e., “beta” $\hat{\beta}_{16,t}$ in Panel D of Figures 3.2 and 3.3. The left-hand plot shows three time series. “beta,” plotted on the left y-axis, is the sensitivity of the stock market from the ES contract; i.e., “beta” $\hat{\beta}_{s,t}$ in Figure 3.4. “only_R,” plotted on the right y-axis, is the portion of the stock market reaction to MNAs attributed to interest rate news: $N_R \equiv - (P/D) (dR/dS)$. I compute “only_R” by using “dp” to proxy for P/D and $\hat{\beta}_{16,t}$ to proxy for dR/dS. “PI_G,” also plotted on the right y-axis, is the portion of the stock market reaction to MNAs attributed to cash flow and risk premium news: $N_{CF+RP} \equiv - (P/D) (d\Pi/dS - dG/dS)$. I compute “PI_G” simply by subtracting “only_R” from $\hat{\beta}_{s,t}$, as it is an accounting identity that a change in the stock market must come from either interest rate news or cash flow and risk premium news (or a combination). The two vertical lines correspond to the dates of the August, 2011 FOMC meeting (when calendar guidance was introduced) and the December, 2012 FOMC meeting (when data thresholds was introduced), respectively.

rate news term. At the beginning of the early period, a standardized unit of news changed the combined expectations of future dividends and excess returns by 90 bps. This reaction decreased steadily to less than 50 bps by mid-2012 before increasing back to 75 bps toward the end of the sample.

The decomposition in Figure 3.5 yields a few interesting observations. First, the magnitudes of the interest rate news term and the combined cash flow and risk premium news terms are an order larger than the magnitude of $\hat{\beta}_{s,t}$. Second, while N_{CF+RP} appears to be a reflection of N_R , the symmetry is clearly not perfect. For example, at the beginning of 2011, cash flow and risk premium news reacted more strongly to data than interest rate news, so stocks on net increased by 10 to 15 bps to a positive unit of surprise. Toward the beginning of 2014, the magnitude of N_{CF+RP} and N_R terms were nearly identical, so stocks on net increased by only 2 bps to a positive unit of surprise. Finally, relating the expected duration of the ZLB to time-variation in the cash flow and risk premium news terms is intriguing. As the expected duration of the ZLB increased over the early period, the magnitude of N_{CF+RP} decreased. Then, as the expected duration of the ZLB stabilized before decreasing over the late period, the magnitude of N_{CF+RP} increased again. The result is a negative correlation between the expected duration of the ZLB and magnitude of the combined change in expected future dividends and expected future excess returns in response to macroeconomic news.

An obvious extension is to break N_{CF+RP} into its constituent cash flow news and risk premium news terms. This step would allow us to separately analyze time variation in how expectations of dividends react to macroeconomic news and how expectations of excess returns react to macroeconomic news. The challenge is finding reliable empirical measures of changes in expected dividends and expected excess returns over short time periods around MNAs. Boyd, Hu, and Jagannathan (2005) attempt to decompose the reaction of stocks to unemployment news into each of the three news components by using monthly proxies for cash flow news and risk premium news. A primary concern with this strategy is that unemployment news only arrives once a month, so any monthly proxy necessarily captures

information aside from unemployment releases. As I consider many MNAs throughout a month, it would be challenging to map the impact of a given MNA to monthly proxies for cash flow and risk premium news. Bernanke and Kuttner (2005) attempt a different decomposition of stocks to surprises in the FF rate. They use a monthly VAR based off the methodology of Campbell (1991) and Campbell and Ammer (1993). A monthly VAR for my purposes, however, would run into the same problems just mentioned.

3.6 Conclusion

This paper has shown how the reactions of interest rates and stocks to macroeconomic news have changed over the ZLB period. A key reason for these time-varying reactions is the expected duration of the ZLB. In a shadow rate term structure model, I analytically derive predictions for how interest rates should react to data surprises. First, “positive” data surprises increases interest rates. Second, as the expected duration of the ZLB increases, whether because economic conditions deteriorate or because monetary policy shifts, interest rates become less sensitive to data surprises. Third and finally, this attenuation in the sensitivity of interest rates is greater for medium-maturity rates. Each of these three predictions is statistically and economically significant in the data. Insofar as the extent to which interest rates move in response to MNAs contains information on the effectiveness of monetary policy and fiscal policy, my results provide policy insights.

The sensitivity of interest rates to news further affects the sensitivity of the stock market to news. If the former were the only factor influencing the latter, greater expected duration of the ZLB would increase the responsiveness of stocks to MNAs. This result does not appear to hold in the data because the reaction of the stock market to news also depends on changes in expected future dividends and changes in expected future excess returns. The expected duration of the ZLB does, however, correlate negatively with the magnitude of this cash flow plus risk premium news term. An interesting idea going forward is to separate cash flow news from risk premium news and see how the magnitude of each term varies over time.

References

- Acharya, Viral V. and Timothy C. Johnson, 2007, "Insider Trading in Credit Derivatives," *Journal of Financial Economics* 84, 110-141.
- Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Heiko Ebens, 2001, "The Distribution of Stock Return Volatility," *Journal of Financial Economics* 61, 43-76.
- Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Clara Vega, 2003, "Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange," *American Economic Review* 93, 38-62.
- Andersen, Torben G., Tim Bollerslev, Francis X. Diebold, and Clara Vega, 2007, "Real-Time Price Discovery in Global Stock, Bond and Foreign Exchange Markets," *Journal of International Economics* 73, 251-277.
- Back, Kerry, 1992, "Insider Trading in Continuous Time," *Review of Financial Studies* 5, 387-409.
- Back, Kerry, C. Henry Cao, and Gregory A. Willard, 2000, "Imperfect Competition among Informed Traders," *Journal of Finance* 55, 2117-2155.
- Baker, Scott R., Nicholas Bloom, and Steven J. Davis, 2013, "Measuring Economic Policy Uncertainty," unpublished paper, Stanford University, Stanford University, and the University of Chicago.
- Balduzzi, Pierluigi, Edwin J. Elton, and Clifton T. Green, 2001, "Economic News and Bond Prices: Evidence from the U.S. Treasury Market," *Journal of Financial and Quantitative Analysis* 36, 523-43.
- Barberis, Nicholas, Andrei Shleifer, and Robert Vishny, 1998, "A Model of Investor Sentiment," *Journal of Financial Economics* 49 307-343.
- Bekaert, Geert and Marie Hoerova, 2014, "The VIX, the Variance Premium and Stock Market Volatility," forthcoming *Journal of Econometrics*.
- Bekaert, Geert, Marie Hoerova, and Marco Lo Doca, 2013, "Risk, Uncertainty and Monetary Policy," *Journal of Monetary Economics* 60 771-788.

Bernanke, Ben S. and Jean Boivin, 2003, "Monetary Policy in a Data-Rich Environment," *Journal of Monetary Economics* 50, 525-546.

Bernanke, Ben S. and Kenneth N. Kuttner, 2005, "What Explains the Stock Market's Reaction to Federal Reserve Policy?" *The Journal of Finance* 60, 1221-1257.

Bernile, Gennaro, Jianfeng Hu, and Yuehua Tang, 2015, "Can Information Be Locked Up? Informed Trading Ahead of Macro-News Announcements," unpublished paper, Singapore Management University, Singapore Management University, and Singapore Management University.

Black, Fischer, 1995, "Interest Rates as Options," *The Journal of Finance* 50, 1371-1376.

Bollerslev, Tim, George Tauchen, and Hao Zhou, 2009, "Expected Stock Returns and Variance Risk Premia," *Review of Financial Studies* 22, 4463-4492.

Boudoukh, Jacob, Matthew Richardson, and Robert F. Whitelaw, 2008, "The Myth of Long-Horizon Predictability," *Review of Financial Studies* 21, 1577-1605.

Boyarchenko, Nina, 2012, "Ambiguity Shifts and the 2007-2008 Financial Crisis," *Journal of Monetary Economics* 59, 493-507.

Boyd, John H., Jian Hu, and Ravi Jagannathan, 2005, "The Stock Market's Reaction to Unemployment News: Why Bad News Is Usually Good for Stocks," *Journal of Finance* 60, 649-672.

Britten-Jones, Mark and Anthony Neuberger, 2000, "Option Prices, Implied Price Processes, and Stochastic Volatility," *Journal of Finance* 55, 839-866.

Brown, Keith C., W.V. Harlow, and Seha M. Tinic, 1988, "Risk Aversion, Uncertain Information, and Market Efficiency," *Journal of Financial Economics* 22, 355-385.

Brunnermeier, Markus K., 2005, "Information Leakage and Market Efficiency," *Review of Financial Studies* 18, 417-457.

Carr, Peter and Luiren Wu, 2009, "Variance Risk Premiums," *Review of Financial Studies* 22, 1311-1341.

Campbell, John Y., 1991, "A Variance Decomposition for Stock Returns," *The Economic Journal* 101, 157-179.

Campbell, John Y. and John Ammer, 1993, "What Moves the Stock and Bond Markets? A Variance Decomposition for Long-Term Asset Returns," *The Journal of Finance*

48, 3-37.

Campbell, John Y. and Ludger Hentschel, 1992, "No News is Good News," *Journal of Financial Economics* 31, 281-318.

Campbell, Sean D. and Steven A. Sharpe, 2009, "Anchoring Bias in Consensus Forecasts and Its Effect on Market Prices," *Journal of Financial and Quantitative Analysis* 44 369-390.

Campbell, John Y., Carolin E. Pflueger, and Luis M. Viceira, 2014, "Monetary Policy Drivers of Bond and Equity Risks," unpublished paper, Harvard University, The University of British Columbia, and Harvard University.

Campbell, John Y., Tarun Ramadorai, and Allie Schwartz, 2009, "Caught on Tape: Institutional Trading, Stock Returns, and Earnings Announcements," *Journal of Financial Economics* 92 66-91.

Carrillo, Juan D. and Thomas R. Palfrey, 2011, "No Trade," *Games and Economic Behavior* 71, 66-87.

Christiano, Lawrence, Martin Eichenbaum, and Sergio Rebelo, 2011, "When is the Government Spending Multiplier Large?" *Journal of Political Economy* 119, 78-121.

Cieslak, Anna, Adair Morse, and Annette Vissing-Jorgensen, 2014, "Stock Returns over the FOMC Cycle," unpublished paper, Northwestern University, University of California, Berkeley, and University of California, Berkeley.

Conrad, Jennifer, Bradford Cornell, and Wayne R. Landsman, 2002, "When Is Bad News Really Bad News?," *Journal of Finance* 57, 2507-2532.

De Long, J. Bradford, Andrei Shleifer, Lawrence H. Summers, and Robert J. Waldmann, 1990, "Positive Feedback Investment Strategies and Destabilizing Rational Speculation," *Journal of Finance* 45, 379-395.

Di Mascio, Rick, Anton Lines, and Narayan Y. Naik, 2015, "Alpha Decay," unpublished paper, Inalytics, London Business School, and London Business School.

Drechsler, Itamar, 2013, "Uncertainty, Time-Varying Fear, and Asset Prices," *Journal of Finance* 68, 1843-1889.

Easley, David, Marcos López de Prado, and Maureen O'Hara, 2013, "Bulk Classification of Trading Activity," unpublished paper, Cornell University, Hess Energy Trading Company, and Cornell University.

Eggertsson, Gauti B., and Michael Woodford, 2003, "The Zero Interest-Rate Bound and

- Optimal Monetary Policy," *Brookings Papers on Economic Activity* 1, 139-211.
- Epstein, Larry G. and Martin Schneider, 2008, "Ambiguity, Information Quality, and Asset Pricing," *Journal of Finance* 63, 197-228.
- Eyster, Erik, Matthew Rabin, and Dimitri Vayanos, 2015, "Financial Markets where Traders Neglect the Informational Content of Prices," unpublished paper, The London School of Economics, Harvard University, and The London School of Economics.
- Faust, Jon, John H. Rogers, Shing-Yi B. Wang, and Jonathan H. Wright, 2007, "The High-Frequency Response of Exchange Rates and Interest Rates to Macroeconomic Announcements," *Journal of Monetary Economics* 54, 1051-1068.
- Faust, Jon and Jonathan H. Wright, 2012, "Risk Premia in the 8:30 Economy," unpublished paper, The Johns Hopkins University and The Johns Hopkins University.
- Foster, F. Douglas and S. Viswanathan, 1996, "Strategic Trading When Agents Forecast the Forecasts of Others," *Journal of Finance* 51, 1437-1478.
- Frazzini, Andrea and Owen Lamont, 2007, "The Earnings Announcement Premium and Trading Volume," unpublished paper, Yale University and the University of Chicago.
- French, Kenneth R., G. William Schwert, and Robert F. Stambaugh, 1987, "Expected Stock Returns and Volatility," *Journal of Financial Economics* 19, 3-30.
- Gilbert, Thomas, Shimon Kogan, Lars Lochstoer, and Ataman Ozyildirim, 2012, "Investor Inattention and the Market Impact of Summary Statistics," *Management Science* 58, 336-350.
- Gilboa, Itzhak and David Schmeidler, 1989, "Maxmin Expected Utility with Non-Unique Prior," *Journal of Mathematical Economics* 18, 141-153.
- Hansen, Lars Peter and Thomas J. Sargent, 2010, "Fragile Beliefs and the Price of Uncertainty," *Quantitative Economics* 1, 129-162.
- Holden, Craig W. and Avanidhar Subrahmanyam, 1992, "Long-Lived Private Information and Imperfect Competition," *Journal of Finance* 47, 247-270.
- Holden, Craig W. and Avanidhar Subrahmanyam, 1994, "Risk Aversion, Imperfect Competition, and Long-Lived Information," *Economic Letters* 44, 181-190.
- Hong, Harrison and Jeremy C. Stein, 1999, "A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets," *Journal of Finance* 54, 2143-2184.
- Ichiue, Hibiki and Yoichi Ueno, 2013, "Estimating Term Premia at the Zero Bound: An

Analysis of Japanese, US, and UK Yields," unpublished paper, Bank of Japan and Bank of Japan.

Illeditsch, Philipp K., 2011, "Ambiguous Information, Portfolio Inertia, and Excess Volatility," *Journal of Finance* 66, 2213-2247.

Irvine, Paul, Marc Lipson, and Andy Puckett, 2007, "Tipping," *Review of Financial Studies* 20, 741-768.

Jones, Charles M., Owen Lamont, and Robin L. Lumsdaine, 1998, "Macroeconomic News and Bond Market Volatility," *Journal of Financial Economics* 47 315-337.

Kadan, Ohad, Roni Michaely, and Pamela C. Moulton, 2014, "Speculating on Private Information: Buy the Rumor, Sell the News," unpublished paper, Washington University in St. Louis, Cornell University, and Cornell University.

Krippner, Leo, 2013, "A Tractable Framework for Zero Lower Bound Gaussian Term Structure Models," unpublished paper, Reserve Bank of New Zealand.

Kurov, Alexander, Alessio Sancetta, Georg Strasser, and Marketa Halova Wolfe, 2015, "Price Drift before U.S. Macroeconomic News: Private Information about Public Announcements?," unpublished paper, West Virginia University, University of London, Boston College, and Washington State University.

Kyle, Albert S., 1985, "Continuous Auctions and Insider Trading," *Econometrica* 53, 1315-1335.

Lee, Charles M. C. and Mark J. Ready, 1991, "Inferring Trade Direction from Intraday Data," *Journal of Finance* 46, 733-746.

Lucca, David O. and Emanuel Moench, 2015, "The Pre-FOMC Announcement Drift," *Journal of Finance* 70, 329-371.

Mehra, Rajnish and Edward C. Prescott, 1985, "The Equity Premium: A Puzzle," *Journal of Monetary Economics* 15, 145-161.

Meulbroek, Lisa K., 1992, "An Empirical Analysis of Illegal Insider Trading," *Journal of Finance* 47, 1661-1699.

Panayides, Marios, Thomas Shohfi, and Jared Smith, 2014, "Comparing Trade Flow Classification Algorithms in the Electronic Era: The Good, the Bad, and the Uninformative," unpublished paper, University of Pittsburgh, University of Pittsburgh, and Clemson University.

Patton, Andrew J. and Allan Timmermann, 2010, "Why Do Forecasters Disagree? Lessons

from the Term Structure of Cross-Sectional Dispersion," *Journal of Monetary Economics* 57 803-820.

Raskin, Matthew D., 2013, "The Effects of the Federal Reserve's Date-Based Forward Guidance," unpublished paper, Federal Reserve Board

Rigobon, Roberto and Brian Sack, 2003, "Measuring the Reaction of Monetary Policy to the Stock Market," *Quarterly Journal of Economics* 118, 639-670.

Romer, Christina D. and David H. Romer, 2000, "Federal Reserve Information and the Behavior of Interest Rates," *American Economic Review* 90 429-457.

Savor, Pavel and Mungo Wilson, 2013, "How Much Do Investors Care About Macroeconomic Risk? Evidence from Scheduled Economic Announcements," *Journal of Financial and Quantitative Analysis* 48 343-375.

Savor, Pavel and Mungo Wilson, 2014, "Asset Pricing: A Tale of Two Days," *Journal of Financial Economics* 113 171-201.

Schwert, G. Williams, 1989, "Why Does Stock Market Volatility Change Over Time?," *Journal of Finance* 44, 1115-1153.

Shleifer, Andrei and Robert W. Vishny, 1997, "The Limits of Arbitrage," *Journal of Finance* 52, 35-55.

Swanson, Eric T. and John C. Williams, 2013, "Measuring the Effect of the Zero Lower Bound on Yields and Exchange Rates in U.K. and Germany," *Journal of International Economics* 73, 251-277.

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, 3154-3185.

Taylor, John B., 1993, "Discretion versus Policy Rules in Practice," *Carnegie-Rochester Series on Public Policy* 39, 195-214.

Tetlock, Paul C., 2011, "All the News That's Fit to Reprint: Do Investors React to Stale Information?," *Review of Financial Studies* 24, 1481-1512.

Tversky, Amos, and Daniel Kahneman, 1974, "Judgement under Uncertainty: Heuristics and Biases," *Science* 185 1124-1131.

Veronesi, Pietro, 1999, "Stock Market Overreaction to Bad News in Good Times: A Rational Expectations Equilibrium Model," *Review of Financial Studies* 12, 975-1007.

Williams, Christopher D., 2014, "Asymmetric Responses to Earnings News: A Case for Ambiguity," forthcoming *Accounting Review*.

Wu, Jing Cynthia and Fan Dora Xia, 2014, "Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound," unpublished paper, The University of Chicago and University of California, San Diego.

Zhou, John C., 2014, "Asset Price Reactions to News at the Zero Lower Bound," unpublished paper, Harvard University.

Zhou, John C., 2015, "The Good, the Bad, and the Ambiguous: The Aggregate Stock Market Dynamics around Macroeconomic News," unpublished paper, Harvard University.

Appendix A

Appendix to Chapter 1


A.1 Intraday Stock and Bond Dollar Volume

Figure A.1 uses volume data to illustrate ES and TY liquidity on days without MNAs. Panels A and B plot the intraday, minute-by-minute and cumulative dollar volumes for stocks and bonds, respectively. The values are averaged across all days in the sample without MNAs. The three dashed vertical lines correspond to 8:30 AM, 9:15 AM, and 10:00 AM ET when MNAs are announced on the days that have such announcements. For both assets, there are negligible or moderate increases in trading activity corresponding to these timestamps.¹ Days without MNAs have substantial liquidity: the average daily cumulative dollar volumes for stock and bond futures are on the order of \$100 billion.

Figure A.2 compares the liquidity between days in the sample with and without MNAs. For both ES and TY, the dollar volumes are greater in the former case than the latter case. The differences in liquidity are particularly large on and immediately after timestamps associated with MNAs; for example, the differences are over \$250 million per minute for ES futures and over \$700 million per minute for TY futures around 8:30 AM ET. Days with MNAs are more liquid even in the early morning hours, though the differences are small in

¹The fact that there are any increases at all is due to the sample selection. What I call days without MNAs are actually days without MNAs from Table 1.1. There are a few types of MNAs not included in Table 1.1 that occur on these days and drive increases in trading activity.

Panel A: Stocks



Panel B: Bonds





Figure A.1: *Intraday Stock and Bond Dollar Volumes on Days without MNAs*

Notes: Panels A and B plot the intraday minute-by-minute and cumulative dollar volumes for stocks (front ES futures contract) and bonds (front TY futures contract), respectively. The values are averaged across all days in the sample without MNAs. The three dashed vertical lines correspond to 8:30 AM, 9:15 AM, and 10:00 AM ET.

Panel A: Stocks



Panel B: Bonds

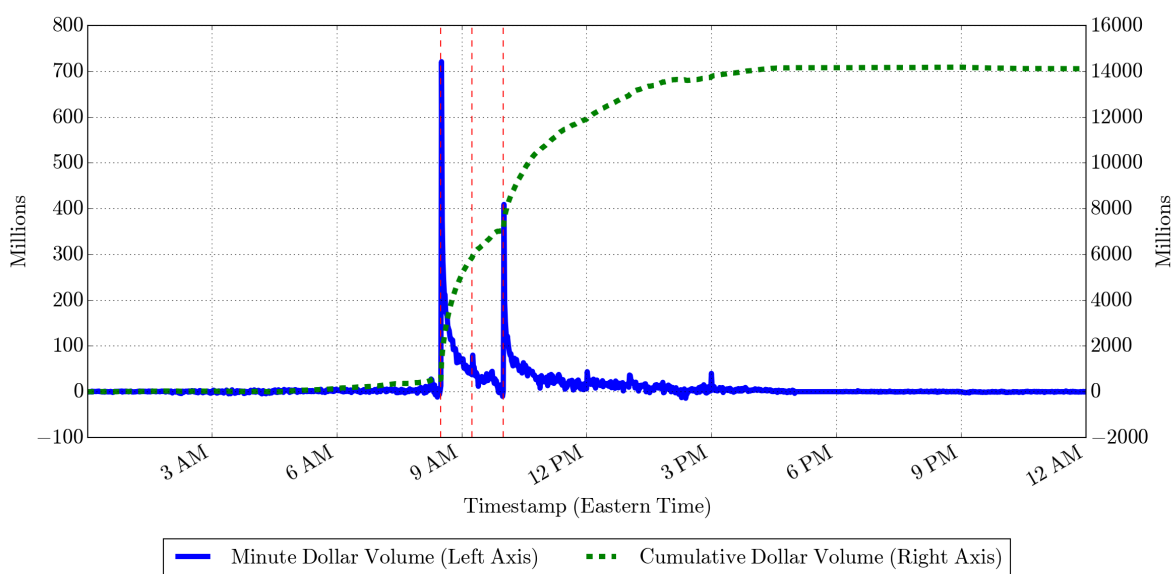


Figure A.2: Differences in Intraday Stock and Bond Dollar Volumes between Days with and without MNAs

Notes: Panels A and B plot the average differences in the intraday minute-by-minute and cumulative dollar volumes for stocks (front ES futures contract) and bonds (front TY futures contract), respectively. The differences are taken between all days in the sample with and without MNAs. The three dashed vertical lines correspond to 8:30 AM, 9:15 AM, and 10:00 AM ET.

magnitude. The average daily cumulative dollar volumes for days with MNAs are over \$8 billion higher for stock futures and \$14 billion higher for bond futures compared to those for days without MNAs.

A.2 Order Imbalances

To classify trades into buys and sells and compute order imbalances, I use the bulk volume classification (BVC) approach developed by Easley, López de Prado, and O'Hara (2013). This trade classification algorithm addresses the challenges presented by modern financial markets dominated by high frequency trading. For example, buying and selling pressure are increasingly less likely to be manifest in aggressive trades (lifting the offer or hitting the bid) and more likely to be manifest in passive trades (such as sitting on the bid or offer). Based on papers such as Easley, López de Prado, and O'Hara (2013) and Panayides, Shohfi, and Smith (2014), the BVC algorithm often outperforms other trade classification algorithms such as that of Lee and Ready (1991). The BVC technique compares favorably in both the correct classification of trades into buys and sells and in identifying order flow linked to underlying information.

The implementation of the BVC algorithm is straightforward: aggregate trades over some interval and split the trades into buys and sells based on the standardized price change between the start and end of the interval. Specifically, for every minute interval τ in which there are trades, I split the volume V_τ into estimated buys \hat{V}_τ^B and sells \hat{V}_τ^S :

$$\hat{V}_\tau^B = V \times t \left(\frac{P_\tau - P_{\tau-1}}{\sigma_{\Delta P}}, df \right), \text{ and}$$

$$\hat{V}_\tau^S = V - \hat{V}_\tau^B.$$

P_τ is the close price of interval τ , $\sigma_{\Delta P}$ is the standard deviation of $P_\tau - P_{\tau-1}$ estimated from July 2003 to March 2014, and $t(\cdot)$ is the CDF of the Student's t distribution with df degrees of freedom. Consistent with Easley, López de Prado, and O'Hara (2013), I set $df = 0.25$. For a longer interval T that covers some number of minute intervals τ_i , the estimated buys

$\hat{V}_T^B = \sum_i \hat{V}_{\tau_i}^B$, the estimated sells $\hat{V}_T^S = \sum_i \hat{V}_{\tau_i}^S$, and the total volume $V_T = \hat{V}_T^B + \hat{V}_T^S$. The order imbalance for interval T is

$$\hat{O}_{I_T} = \frac{\hat{V}_T^B - \hat{V}_T^S}{V_T}.$$

The order imbalance multiplied by the volume gives the estimated excess volume of buys or sells.

Analogous to Table 1.2, I calculate summary statistics of order imbalances in Table A.1. Row $-x$ corresponds to the pre-announcement order imbalance (in bps) from x minutes to 5 minutes before MNAs. Row 0 corresponds to the ± 5 minute announcement order imbalance (in bps) around MNAs. Panel A presents summary statistics for stocks. In the “Full Sample of MNAs,” there are unconditionally positive buy order imbalances both in the 10-minute window around announcements (mean of 50.6 bps) and in the pre-announcement period (mean of up to 22.0 bps).² As shown in the “News Report > 1 ” and “News Report < -1 ” columns, the announcement order imbalance of stocks is positive (mean of 453.3 bps) for greater than 1 standard deviation positive news reports and negative (mean of -808.5 bps) for less than 1 standard deviation negative news reports. The announcement order imbalance is preceded by the pre-announcement order imbalance in the same direction. There is a positive buy order imbalance (up to 60.7 bps on average) in the hours prior to greater than 1 standard deviation positive news reports and a negative sell order imbalance (up to -4.7 bps on average) in the hours prior to less than 1 standard deviation negative news reports.

Panel B presents analogous summary statistics for bonds. In the “Full Sample” of MNAs in Panel B of Table A.1, there are unconditionally negative sell order imbalances on announcements (mean of -21.1 bps) and before announcements (mean of up to -35.2 bps).³ Opposite that of stocks, the order imbalance of bonds is negative (mean of -787.8

²These two results correspond to the announcement premium of stocks around MNAs and pre-announcement drift of stocks before MNAs, respectively. See Fn. (11) and Fn. (12), respectively, from chapter one.

³The former result is statistically insignificant, consistent with the lack of an announcement premium for bonds around MNAs. The latter result is statistically significant, consistent with the presence of the pre-announcement drift of bonds before MNAs. See Fn. (13) from chapter one.

Table A.1: Summary Statistics of Pre-Announcement and Announcement Order Imbalances in Stocks and Bonds

Panel A: Stocks

	Full Sample ($n = 2179$)			News Report > 1 ($n = 291$)			News Report < -1 ($n = 286$)		
	Mean	Std. Dev.	t -stat	Mean	Std. Dev.	t -stat	Mean	Std. Dev.	t -stat
-720	18.2	218.9	3.9	45.0	228.7	3.4	-0.6	191.1	-0.1
-660	18.3	220.9	3.9	43.7	231.2	3.2	-1.6	192.5	-0.1
-600	18.1	222.6	3.8	44.9	234.5	3.3	-1.5	195.1	-0.1
-540	18.5	224.4	3.9	46.6	236.3	3.4	0.0	198.2	0.0
-480	18.9	226.7	3.9	47.6	239.9	3.4	-0.8	200.0	-0.1
-420	18.0	229.9	3.7	47.4	244.1	3.3	-2.1	203.3	-0.2
-360	17.6	236.7	3.5	45.9	251.9	3.1	-3.7	210.8	-0.3
-300	17.5	250.9	3.3	50.8	268.1	3.2	-3.8	228.1	-0.3
-240	17.7	267.3	3.1	50.9	286.4	3.0	-0.1	252.2	0.0
-180	18.2	284.2	3.0	55.7	304.9	3.1	-4.7	272.0	-0.3
-120	21.7	310.6	3.3	55.0	335.1	2.8	5.6	306.0	0.3
-60	22.0	361.1	2.8	60.7	368.9	2.8	17.1	362.7	0.8
0	50.6	1277.2	1.8	453.3	1360.2	5.7	-808.5	1403.0	-9.7

Panel B: Bonds

	Full Sample ($n = 2179$)			News Report > 1 ($n = 291$)			News Report < -1 ($n = 286$)		
	Mean	Std. Dev.	t -stat	Mean	Std. Dev.	t -stat	Mean	Std. Dev.	t -stat
-720	-19.7	190.1	-4.8	-48.3	189.6	-4.3	6.9	201.6	0.6
-660	-20.0	192.4	-4.9	-48.9	190.4	-4.4	7.5	204.6	0.6
-600	-21.1	194.9	-5.0	-50.2	193.4	-4.4	6.2	207.7	0.5
-540	-21.4	196.5	-5.1	-50.6	196.2	-4.4	6.7	207.6	0.5
-480	-23.0	199.1	-5.4	-52.2	198.1	-4.5	5.6	211.1	0.4
-420	-23.6	202.2	-5.5	-53.5	202.3	-4.5	3.5	214.7	0.3
-360	-24.4	209.8	-5.4	-53.2	211.5	-4.3	5.5	222.1	0.4
-300	-26.4	225.3	-5.5	-61.3	227.0	-4.6	3.7	244.4	0.3
-240	-29.5	246.7	-5.6	-66.4	243.5	-4.7	-0.5	265.8	0.0
-180	-31.2	270.8	-5.4	-75.0	270.8	-4.7	-0.1	287.3	0.0
-120	-34.3	304.1	-5.3	-72.4	301.0	-4.1	0.7	325.8	0.0
-60	-35.2	373.4	-4.4	-91.4	368.2	-4.2	-5.9	397.5	-0.2
0	-21.1	1319.8	-0.7	-787.8	1245.3	-10.8	868.2	1284.8	11.4

Notes: Panels A and B present summary statistics for stocks (front ES futures contract) and bonds (front TY futures contract), respectively. Row $-x$ corresponds to the pre-announcement order imbalance (in bps) from x minutes to 5 minutes before MNAs. Row 0 corresponds to the ± 5 minute announcement order imbalance (in bps) around MNAs. The "Full Sample" columns include the full sample of MNAs. The "News Report > 1" columns include only the MNAs with greater than 1 standard deviation positive news reports; that is, $NR_t > 1$. The "News Report < -1" columns include only the MNAs with less than 1 standard deviation negative news reports; that is, $NR_t < -1$. NR_t measures the informational content of MNAs and is defined in Figure 1.1.

bps) on announcements of greater than 1 standard deviation positive news reports and positive (mean of 868.2 bps) on announcements of less than 1 standard deviation negative news reports. Identical to the case of stocks, however, the announcement order imbalance of bonds is preceded by the pre-announcement order imbalance of bonds in the same direction. There is a negative sell order imbalance (up to -91.4 bps on average) in the hours prior to greater than 1 standard deviation positive news reports and a positive buy order imbalance (up to 7.5 bps on average) in the hours prior to less than 1 standard deviation negative news reports.

Table A.2 shows that the pre-announcement conditional pattern in the order imbalance is statistically significant. Panel A runs regressions of pre-announcement and announcement order imbalances on the news report variable NR_t :

$$\hat{O}I_{T,t} = \alpha + \beta NR_t + \epsilon_t.$$

In row $-x$, the left-hand-side variable $\hat{O}I_{T,t}$ is the pre-announcement order imbalance (in bps) for interval T from x minutes to 5 minutes before MNAs. In row 0, $\hat{O}I_{T,t}$ is the announcement order imbalance (in bps) for interval T covering the ± 5 minutes around MNAs. The coefficient β measures the order imbalance sensitivity to a 1 standard deviation news report. For both assets, $\hat{\beta}$ in row 0 is statistically significant: the order imbalance of stocks increases {decreases} 357.7 bps more and that of bonds decreases {increases} 498.4 bps more in the 10 minutes around announcements of 1 standard deviation positive {negative} news reports. For rows $-x$, the majority of the $\hat{\beta}$ for stocks and all the $\hat{\beta}$ for bonds are statistically significant: the order imbalance of stocks increases {decreases} up to 15.9 bps more and that of bonds decreases {increases} up to 24.4 bps more before announcements of 1 standard deviation positive {negative} news reports.

In Panel B, I evaluate how news reports impact order imbalances in a different regression specification:

$$\hat{O}I_{T,t} = \alpha + \beta^+ D_t^+ + \beta^- D_t^- + \epsilon_t.$$

Table A.2: Regressions of Pre-Announcement and Announcement Order Imbalances in Stocks and Bonds

Panel A: Regressions of the Form $\hat{O}I_{T,t} = \alpha + \beta NR_t + \epsilon_t$

	Stocks			Bonds		
	$\hat{\beta}$ (bps)	$\hat{\alpha}$ (bps)	R^2_{adj} (%)	$\hat{\beta}$ (bps)	$\hat{\alpha}$ (bps)	R^2_{adj} (%)
-720	12.0**	18.0***	0.26	-13.3***	-19.5***	0.45
-660	11.8**	18.1***	0.24	-13.5***	-19.7***	0.46
-600	12.0**	17.9***	0.25	-13.6***	-20.8***	0.45
-540	11.7**	18.3***	0.23	-13.7***	-21.1***	0.45
-480	12.0**	18.6***	0.24	-13.9***	-22.7***	0.45
-420	12.1**	17.8***	0.23	-14.2***	-23.4***	0.46
-360	12.4**	17.4***	0.23	-14.7***	-24.1***	0.45
-300	13.4**	17.3***	0.24	-17.1***	-26.1***	0.54
-240	13.8**	17.4***	0.22	-17.2***	-29.2***	0.44
-180	15.9***	17.9***	0.27	-19.5***	-30.8***	0.48
-120	13.4**	21.4***	0.14	-20.6***	-33.9***	0.42
-60	15.2**	21.7***	0.13	-24.4***	-34.7***	0.38
0	357.7***	44.0*	7.89	-498.4***	-11.9	14.40

Panel B: Regressions of the Form $\hat{O}I_{T,t} = \alpha + \beta^+ D_t^+ + \beta^- D_t^- + \epsilon_t$

	Stocks				Bonds			
	$\hat{\beta}^+$ (bps)	$\hat{\beta}^-$ (bps)	$\hat{\alpha}$ (bps)	R^2_{adj} (%)	$\hat{\beta}^+$ (bps)	$\hat{\beta}^-$ (bps)	$\hat{\alpha}$ (bps)	R^2_{adj} (%)
-720	28.3**	-17.3	16.7***	0.21	-29.0**	26.2**	-19.3***	0.47
-660	26.4*	-18.9	17.3***	0.19	-29.2**	27.2**	-19.7***	0.48
-600	28.1**	-18.2	16.8***	0.21	-29.5**	26.8**	-20.7***	0.46
-540	29.9**	-16.7	16.7***	0.21	-29.5**	27.8**	-21.1***	0.47
-480	30.4**	-17.9	17.2***	0.23	-29.5**	28.3**	-22.8***	0.47
-420	31.1**	-18.4	16.3***	0.23	-30.5**	26.5**	-23.0***	0.44
-360	29.6**	-19.9	16.3***	0.21	-28.7**	30.0**	-24.5***	0.43
-300	35.5**	-19.0	15.3**	0.24	-35.7**	29.2**	-25.5***	0.46
-240	36.1**	-14.9	14.8**	0.18	-38.4**	27.5*	-28.0***	0.39
-180	40.2**	-20.2	15.5**	0.23	-46.3***	28.6*	-28.7***	0.44
-120	36.5*	-12.9	18.5**	0.10	-38.7**	34.4*	-33.7***	0.29
-60	44.8*	1.2	15.8*	0.09	-61.2**	24.3	-30.2***	0.31
0	322.4***	-939.3***	130.8***	7.47	-747.2***	908.8***	-40.6	10.41

Notes: See Appendix A.7.3.


D_t^+ and D_t^- are dummy variables equal to 1 if $NR_t > 1$ and $NR_t < -1$, respectively, and $\hat{O}I_{T,t}$ is as defined above. The coefficients β^+ and β^- measure the order imbalance sensitivity to large magnitude news reports that are positive and negative, respectively. For both assets, $\hat{\beta}^+$ and $\hat{\beta}^-$ in row 0 are statistically significant: the order imbalance of stocks increases 322.4 bps {decreases 939.3 bps} more and that of bonds decreases 747.2 bps {increases 908.8 bps} more for announcements of greater than 1 standard deviation positive {negative} news reports.⁴ For rows $-x$, the majority of the $\hat{\beta}^+$ but none of the $\hat{\beta}^-$ for stocks and almost all the $\hat{\beta}^+$ and $\hat{\beta}^-$ for bonds are statistically significant: the order imbalance of stocks increases up to 44.8 bps {decreases up to 20.2 bps} more and that of bonds decreases up to 61.2 bps {increases up to 34.4 bps} more before announcements of greater than 1 standard deviation positive {negative} news reports.⁵

Panels A and B of Figure A.3 graphically illustrate the pre-announcement conditional pattern in the order imbalance for stocks and bonds, respectively. I plot the mean order imbalance (in bps) for the indicated time intervals on the x-axes starting 360 minutes before MNAs. First, consider the results for stocks in Panel A. In the hours before announcements, the order imbalance is lowest before less than 1 standard deviation negative news reports, intermediate before announcements in general, and highest before greater than 1 standard deviation positive news reports. Parallel results hold for bonds in Panel B. In the hours before announcements, the order balance is highest before less than 1 standard deviation negative news reports, intermediate before announcements in general, and lowest before greater than 1 standard deviation positive news reports. The pre-announcement conditional pattern in the order imbalance, in addition to the pre-announcement conditional drift, provides support for sophisticated trading in stocks and bonds around MNAs.

⁴See Fn. (14) from chapter one for discussion of the asymmetry in stocks.

⁵For stocks in Panels A and B, the row 0 $\hat{\alpha}$ is positive and statistically significant. This corresponds to the announcement premium for stocks around MNAs referenced in Fn. (11) from chapter one. The row $-x$ $\hat{\alpha}$ is also positive and generally statistically significant. This is the pre-announcement drift of stocks before MNAs referenced in Fn. (12) from chapter one. For bonds, the row 0 $\hat{\alpha}$ is statistically insignificant, but the row $-x$ $\hat{\alpha}$ is negative and generally statistically significant. As discussed in Fn. (13) from chapter one, there is no announcement premium for bonds around MNAs, but there is a pre-announcement drift of bonds before MNAs.

Panel A: Stocks



Panel B: Bonds




Figure A.3: Pre-Announcement Order Imbalances in Stocks and Bonds

Notes: Panels A and B plot the mean order imbalances (in bps) for the indicated time intervals on the x-axes starting 360 minutes before MNAs for stocks (front ES futures contract) and bonds (front TY futures contract), respectively. “Full Sample” includes the full sample of MNAs. “News Report > 1” includes only the MNAs with greater than 1 standard deviation positive news reports; that is, $NR_t > 1$. “News Report < -1” includes only the MNAs with less than 1 standard deviation negative news reports; that is, $NR_t < -1$. NR_t measures the informational content of MNAs and is defined in Figure 1.1

A.3 Alternative Tests for Anchoring Bias

In Section 1.4, I estimate the anchoring bias regression in Eq. (1.5) for each type of MNA separately. The results are robust to a battery of tests. In particular, the results hold when I estimate the the regression using different anchors and different windows. As indicated in the “Anchor” columns of Table A.3, the alternative tests use the $h = 1$ anchor (last actual data release), $h = 3$ anchor (average of the last 3 actual data releases), or “Best” anchor (anchor that more frequently has a higher adjusted R^2 or anchor that has a higher adjusted R^2 in each iteration). As indicated in the “Window” columns of Table A.3, the alternative tests extend the rolling regression window back the listed number of years, with “All” corresponding to the earliest available date with data. The results in Table 1.4 covered in the body of this paper correspond to the “Best” anchor (anchor that more frequently has a higher adjusted R^2), “All” window specification. Other specifications yield results similar to those in Table 1.4: for most types of MNAs, the mean of $\hat{\gamma}$ is positive and large relative to the standard deviation.

I also estimate the anchoring bias regression grouping all MNAs together. To do so, I standardize the independent and dependent variables in Eq. (1.5) by dividing by $\hat{\sigma}$ such that the variables are in the same (standardized) units across different types of MNAs:

$$NR_t = \kappa + \gamma \left(\frac{F_t - \bar{A}_h}{\hat{\sigma}} \right) + \epsilon_t. \quad (\text{A.1})$$

The left-hand-side variable is just the news report variable NR_t . I estimate the above regression using the different anchors and different windows listed in Table A.3. Table A.4 shows the mean, standard deviation, maximum, and minimum of $\hat{\gamma}$ for the various tests.⁶ As in the case without grouping all MNAs together, the case with grouping yields positive $\hat{\gamma}$ means that are large relative to the standard deviations. Based on both cases, economists’ forecasts exhibit anchoring bias, and differences between economists’ forecasts and past data realizations predict macroeconomic news reports.

⁶The “Best” anchor specifications use the anchor that has a higher adjusted R^2 in each iteration.

Table A.3: *Alternative Tests for Anchoring Bias*

Anchor	Window	Anchor	Window	Anchor	Window
Best	All	1	All	1	3 years
Best	1 year	3	All	3	3 years
Best	2 years	1	1 year	1	4 years
Best	3 years	3	1 year	3	4 years
Best	4 years	1	2 years	1	5 years
Best	5 years	3	2 years	3	5 years

Notes: The table lists different specifications for running the anchoring bias regressions in Eqs. (1.5) and (A.1). As indicated in the “Anchor” columns, the alternative tests use the $h = 1$ anchor (last actual data release), $h = 3$ anchor (average of the last 3 actual data releases), or “Best” anchor (anchor that more frequently has a higher adjusted R^2 or anchor that has a higher adjusted R^2 in each iteration). As indicated in the “Window” columns, the alternative tests extend the rolling regression window back the listed number of years, with “All” corresponding to the earliest available date with data.

A.4 Alternative Constructions of the Predictable and Unpredictable Components of News Reports

Appendix A.3 shows that various specifications of the anchoring bias regression predict macroeconomic news reports. Accordingly, there are many ways to construct the predictable and unpredictable components of the news report variable. I run regressions of pre-announcement (corresponding to Eq. (1.6)) and announcement (corresponding to Eq. (1.7)) stock and bond returns on alternative constructions of these two components.

When MNAs are not grouped together, the regression results are similar to those in Table 1.5, so I do not report the estimated coefficients. In other words, the results in Table 1.5 are robust. In the pre-announcement return regressions, $\hat{\beta}^p$ is positive for stocks, negative for bonds, and statistically significant for both assets: stocks increase {decrease} and bonds decrease {increase} in the hours before news reports predicted to be positive {negative}. This finding suggests that, around MNAs, markets are inefficient over time, as I discuss in Section 1.5.1. In the announcement return regressions, $\hat{\beta}^p$ is statistically significant for stocks but not for bonds: stocks but not bonds react to announcements of the predictable component. This finding suggests that, around MNAs, markets are inefficient across assets,

Table A.4: Predicting Macroeconomic News Reports Grouping All MNAs Together

Anchor	Window	Mean $\hat{\gamma}$	Std. Dev. $\hat{\gamma}$	Max. $\hat{\gamma}$	Min. $\hat{\gamma}$
Best	All	0.11	0.01	0.13	0.10
Best	1 year	0.11	0.06	0.29	-0.03
Best	2 years	0.11	0.05	0.23	0.03
Best	3 years	0.11	0.03	0.16	0.05
Best	4 years	0.11	0.03	0.15	0.06
Best	5 years	0.11	0.02	0.15	0.06
1	All	0.07	0.01	0.09	0.06
3	All	0.11	0.01	0.13	0.10
1	1 year	0.06	0.04	0.17	-0.03
3	1 year	0.11	0.06	0.29	-0.01
1	2 years	0.06	0.03	0.14	0.00
3	2 years	0.11	0.05	0.23	0.03
1	3 years	0.06	0.02	0.10	0.02
3	3 years	0.11	0.03	0.16	0.05
1	4 years	0.06	0.02	0.10	0.03
3	4 years	0.11	0.02	0.15	0.06
1	5 years	0.06	0.02	0.09	0.03
3	5 years	0.11	0.02	0.15	0.06

Notes: For each “Anchor” and “Window” specification described in Table A.3, the corresponding row runs the regression $NR_t = \kappa + \gamma ((F_t - \bar{A}_h) / \hat{\sigma}) + \epsilon_t$. The “Best” anchor specifications use the anchor that has a higher adjusted R^2 in each iteration. NR_t measures the informational content of a given data release, and $\hat{\sigma}$ and NR_t are both defined in Figure 1.1. For the same data release, F_t is the median forecasted value from a survey of economists, and \bar{A}_h is the historical average of the past h realizations. To avoid look-ahead bias, I estimate the regression in a dynamic manner over a changing window.

as I discuss in Section 1.5.2.

When the MNAs are grouped together, the regression results are similar to those in Table A.5. In this table, I construct $NR_{p,t}$ as the one-step-ahead forecast of the “Best” anchor (anchor that has a higher adjusted R^2 in each iteration), “1 year” window specification of the anchoring bias regression. $NR_{u,t}$ is the regression residual. The results of Table A.5 are generally similar to those in Table 1.5. In the pre-announcement return regressions, for example, $\hat{\beta}^p$ is positive for stocks, negative for bonds, and statistically significant for both assets. There is still a “late trading puzzle” over time. In the announcement return regressions, however, $\hat{\beta}^p$ is now statistically insignificant for stocks and bonds, so it appears that neither asset reacts to announcements of the predictable component. This finding (or lack thereof) calls into question whether it is the case that, before announcements, the predictable component of news reports is fully impounded into bonds but only partially impounded into stocks. It may be that the predictable component is fully impounded into both assets. Nonetheless, there is other, independent evidence that stocks are less efficient than bonds. As shown in Section 1.5.2, stocks but not bonds display return momentum between pre-announcement and announcement returns.

A.5 Model with One Risky Asset

Section 1.6.1 sets up the model. Agents in the model at time t choose demand x_t to maximize utility over next period wealth $W_{t+1} = W_t + (p_{t+1} + d_{t+1} - p_t) x_t$, with p_t the ex-dividend price of the risky asset, and d_{t+1} the dividend paid out by the risky asset between t and $t + 1$. The standard framework of the model implies that the optimal demand x_t for agents with a given set of beliefs and information is

$$x_t = \frac{\mathbb{E}[p_{t+1} + d_{t+1}] - p_t}{AV[p_{t+1} + d_{t+1}]} \quad (\text{A.2})$$

Below, I show the derivation for equilibrium prices under the general case for measure $i \in [0, 1]$ of sophisticated traders and measure $1 - i$ of unsophisticated traders.

Table A.5: Regressions of Stock and Bond Returns on the Predictable and Unpredictable Components of Macroeconomic News Reports

	Stocks				Bonds			
	$\hat{\beta}^p$ (bps)	$\hat{\beta}^u$ (bps)	$\hat{\alpha}$ (bps)	R_{adj}^2 (%)	$\hat{\beta}^p$ (bps)	$\hat{\beta}^u$ (bps)	$\hat{\alpha}$ (bps)	R_{adj}^2 (%)
-720	9.3*	1.6	4.3***	0.12	-7.7***	-1.3***	-0.3	0.93
-660	9.2*	1.3	4.1***	0.10	-8.0***	-1.4***	-0.5	1.09
-600	6.2	1.3	4.2***	0.04	-6.9***	-1.4***	-0.7*	0.94
-540	4.9	1.4	4.6***	0.03	-6.2***	-1.4***	-0.9**	0.88
-480	8.4*	1.9*	4.3***	0.17	-6.0***	-1.4***	-1.1***	0.92
-420	8.0	2.2**	3.5***	0.20	-5.9***	-1.4***	-0.9**	0.96
-360	3.9	1.8*	2.8***	0.09	-5.1***	-1.4***	-0.7*	0.94
-300	4.4	2.3**	1.7*	0.23	-4.7***	-1.4***	-0.7*	1.00
-240	2.7	1.9**	1.1	0.14	-4.8***	-1.2***	-0.7**	0.91
-180	9.1***	1.6**	0.3	0.40	-5.6***	-1.1***	-0.7**	1.15
-120	2.0	1.0	0.8	0.02	-3.2**	-1.1***	-0.6**	0.84
-60	6.5***	0.7	0.1	0.31	-1.1	-0.7***	-0.7***	0.56
0	2.2	6.8***	1.1*	6.43	1.6	-5.5***	0.1	8.26

Notes: The table shows results from regressions of pre-announcement and announcement stock (front ES futures contract) and bond (front TY futures contract) returns on the predictable and unpredictable components of news reports $NR_{p,t}$ and $NR_{u,t}$, respectively: $R_t = \alpha + \beta^p NR_{p,t} + \beta^u NR_{u,t} + \epsilon_t$. Regressions are run on the full sample of MNAs, and $n = 2174$. In row $-x$, the left-hand-side variable R_t is the cumulative pre-announcement return (in bps) from x minutes to 5 minutes before MNAs. In row 0, the left-hand-side variable R_t is the ± 5 minute announcement return (in bps) around MNAs. $NR_{p,t}$ is the one-step-ahead forecast of the "Best" anchor, "1 Year" window specification of the anchoring bias regression in Table A.4. $NR_{u,t}$ is the regression residual. t -statistics (not shown) are based on heteroskedasticity-consistent standard errors. *** denotes significance at the 1% level, ** denotes significance at the 5% level, and * denotes significance at the 10% level.

A.5.1 $t = 4$

Working backward in time, it is obvious that $p_4 = 0$, and the efficient-market price of the risky asset right before the payment of dividend d at $t = 4$ is

$$\mathbb{E} [p_4 + d|n] = \alpha n,$$

with

$$\alpha = \frac{\sigma_d^2}{\sigma_d^2 + \sigma_\varepsilon^2}. \quad (\text{A.3})$$

A.5.2 $t = 3$

At $t = 3$, sophisticated traders see the public news report n , which they already know. Their demand is

$$x_{S,3} = \frac{\mathbb{E} [p_4 + d|n] - p_3}{A_S \mathbb{V} [p_4 + d|n]} = \frac{\alpha n - p_3}{A_S \alpha \sigma_\varepsilon^2}.$$

Unsophisticated traders also see n and have demand

$$x_{U,3} = \frac{\mathbb{E} [p_4 + d|n] - p_3}{A_U \mathbb{V} [p_3 + d|n]} = \frac{\alpha n - p_3}{A_U \alpha \sigma_\varepsilon^2}.$$

Imposing the market-clearing condition at $t = 3$ that

$$i x_{S,3} + (1 - i) x_{U,3} = 0,$$

we get that the equilibrium price of the risky asset at $t = 3$ is

$$p_3 = \alpha n.$$

A.5.3 $t = 2$

At $t = 2$, sophisticated traders have advance knowledge of n and can trade on this information in the “late period.” Doing so is risky, however, because there is a shock to holding the risky asset in the form of dividend payout $s_l \sim \mathcal{N} \left(0, \sigma_{s,l}^2 \right)$ between $t = 2$ and $t = 3$. The

sophisticated traders' demand is

$$x_{S,2} = \frac{\mathbb{E}[p_3 + s_l | n] - p_2}{A_S \mathbb{V}[p_3 + s_l | n]} = \frac{\alpha n - p_2}{A_S \sigma_{s,l}^2}.$$

Unsophisticated traders are unaware of n , so their demand is

$$x_{U,2} = \frac{\mathbb{E}[p_3 + s_l] - p_2}{A_U \mathbb{V}[p_3 + s_l]} = \frac{-p_2}{A_U (\alpha^2 (\sigma_d^2 + \sigma_\epsilon^2) + \sigma_{s,l}^2)}$$

Imposing the market-clearing condition at $t = 2$ that

$$i x_{S,2} + (1 - i) x_{U,2} = 0$$

and solving for p_2 yields

$$p_2 = \frac{1}{1 + \left(\frac{1-i}{i}\right) \left(\frac{A_S}{A_U}\right) \frac{\sigma_{s,l}^2}{(\alpha^2 (\sigma_d^2 + \sigma_\epsilon^2) + \sigma_{s,l}^2)}} \alpha n = \frac{1}{1 + \left(\frac{1-i}{i}\right) \left(\frac{A_S}{A_U}\right) \frac{\sigma_{s,l}^2}{(\alpha^2 (\sigma_d^2 + \sigma_\epsilon^2) + \sigma_{s,l}^2)}} p_3,$$

which is Eq. (1.9). I define the coefficient in front of p_3 as k_2 . In the efficient-market specification in which all traders are sophisticated and $i = 1$, $p_2 = p_3$. Otherwise, in the inefficient-market specification in which some traders are unsophisticated and $i \in (0, 1)$, $p_2 < p_3$.

A.5.4 $t = 1$

At $t = 1$, sophisticated traders first obtain advance knowledge of n and have an opportunity to trade in the "early period." Trading is risky, and there is a shock to holding the risky asset in the form of dividend payout $s_e \sim \mathcal{N}(0, \sigma_{s,e}^2)$ between $t = 1$ and $t = 2$. The sophisticated traders' demand is

$$x_{S,1} = \frac{\mathbb{E}[p_2 + s_e | n] - p_1}{A_S \mathbb{V}[p_2 + s_e | n]} = \frac{k_2 \alpha n - p_1}{A_S \sigma_{s,e}^2}.$$

Unsophisticated traders are unaware of n , so their demand is

$$x_{U,1} = \frac{\mathbb{E}[p_2 + s_e] - p_1}{A_U \mathbb{V}[p_2 + s_e]} = \frac{-p_1}{A_U (k_2^2 \alpha^2 (\sigma_d^2 + \sigma_\epsilon^2) + \sigma_{s,e}^2)}.$$

Imposing the market-clearing condition at $t = 1$ that

$$ix_{S,1} + (1 - i)x_{U,1} = 0$$

and solving for p_1 yields

$$p_1 = \frac{1}{1 + \left(\frac{1-i}{i}\right) \left(\frac{A_S}{A_U}\right) \frac{\sigma_{s,e}^2}{(k_2^2 \alpha^2 (\sigma_d^2 + \sigma_\epsilon^2) + \sigma_{s,e}^2)}} k_2 \alpha n = \frac{1}{1 + \left(\frac{1-i}{i}\right) \left(\frac{A_S}{A_U}\right) \frac{\sigma_{s,e}^2}{(k_2^2 \alpha^2 (\sigma_d^2 + \sigma_\epsilon^2) + \sigma_{s,e}^2)}} p_2,$$

which is Eq. (1.10). In the efficient-market specification in which all traders are sophisticated and $i = 1$, $p_1 = p_2 = p_3$. Otherwise, in the inefficient-market specification in which some traders are unsophisticated and $i \in (0, 1)$, $p_1 < p_2 < p_3$.

A.5.5 $t = 0$

Since no information has arrived for either type of traders and all random variables have zero mean,

$$p_0 = 0.$$

A.6 Model with Two Risky Assets

Section 1.6.2 sets up the model. Agents in the model at time t choose demands $x_{t,a}$ and $x_{t,b}$ for risky assets a and b , respectively, to maximize utility over next period wealth $W_{t+1} = W_t + (p_{t+1,a} + d_{t+1,a} - p_{t,a})x_{t,a} + (p_{t+1,b} + d_{t+1,b} - p_{t,b})x_{t,b}$. $p_{t,a}$ and $p_{t,b}$ are the ex-dividend prices of the risky assets, and $d_{t+1,a}$ and $d_{t+1,b}$ are the dividends paid out by the risky assets between t and $t + 1$. Since the random variables in the model corresponding to assets a and b are independent across assets, the optimal demands $x_{t,a}$ and $x_{t,b}$ for agents with a given set of beliefs and information are of the same form as Eq. (A.2), the optimal demand in the model with one risky asset. Similarly, the derivation for equilibrium prices in

the model with two risky assets parallels that for the model with one risky asset. The only difference is the addition of asset-specific subscripts. For example, α_a and α_b are defined by Eq. (A.3) as $\sigma_{d,a}^2 / (\sigma_{d,a}^2 + \sigma_{\epsilon,a}^2)$ and $\sigma_{d,b}^2 / (\sigma_{d,b}^2 + \sigma_{\epsilon,b}^2)$, respectively.

A.7 Notes for Figures and Tables

A.7.1 Figure 1.6

Panel A plots the annualized Sharpe ratios from trading strategies based on perfect foresight of macroeconomic news reports. Panel B decomposes the Sharpe ratios into the annualized standard deviations and the annualized means. For a given $-x$ on the x-axis, the trading strategies take positions in stocks and bonds starting x days or minutes before MNAs. In stocks (front ES futures contract), the strategies go long {short} before positive {negative} news reports of greater than 0.7 standard deviation in magnitude; that is, $|NR_t| > 0.7$. In bonds (front TY futures contract), the strategies go the opposite direction. In both stocks and bonds, positions are unwound 5 minutes after MNAs. NR_t measures the informational content of MNAs and is defined in Figure 1.1.

A.7.2 Figure 1.7

The figure plots the annualized Sharpe ratios from trading strategies based on the predictable component of macroeconomic news reports $NR_{p,t}$ defined in Table 1.5. The strategies take positions in stocks (Panels A and B) and bonds (Panels C and D) starting some number of days (Panels A and C) or minutes (Panels B and D) before MNAs. In stocks (front ES futures contract), the strategies go long {short} before the top {bottom} “Trade Threshold” percentile of $NR_{p,t}$. In bonds (front TY futures contract), the strategies go the opposite direction. Positions are unwound 5 minutes after MNAs for stocks and 5 minutes before MNAs for bonds.

A.7.3 Table A.2

The table shows results from regressions of pre-announcement and announcement stock (front ES futures contract) and bond (front TY futures contract) order imbalances on the news report variable NR_t . Regressions are run on the full sample of MNAs, and $n = 2179$. Panels A and B represent different regression specifications: $\hat{O}I_{T,t} = \alpha + \beta NR_t + \epsilon_t$, and $\hat{O}I_{T,t} = \alpha + \beta^+ D_t^+ + \beta^- D_t^- + \epsilon_t$, respectively. In row $-x$, the left-hand-side variable $\hat{O}I_{T,t}$ is the pre-announcement order imbalance (in bps) for interval T from x minutes to 5 minutes before MNAs. In row 0, the left-hand-side variable $\hat{O}I_{T,t}$ is the announcement order imbalance (in bps) for interval T covering the ± 5 minutes around MNAs. NR_t measures the informational content of MNAs and is defined in Figure 1.1. D_t^+ and D_t^- are dummy variables equal to 1 if $NR_t > 1$ and $NR_t < -1$, respectively, and zero otherwise. t -statistics (not shown) are based on heteroskedasticity-consistent standard errors. *** denotes significance at the 1% level, ** denotes significance at the 5% level, and * denotes significance at the 10% level.

Appendix B

Appendix to Chapter 2

B.1 Conditional Distribution

Let

$$\mathbf{X} = \begin{pmatrix} \mathbf{X}_1 \\ \mathbf{X}_2 \end{pmatrix} \sim \mathcal{N} \left(\begin{pmatrix} \mu_1 \\ \mu_2 \end{pmatrix}, \begin{pmatrix} \Sigma_{11} & \Sigma_{12} \\ \Sigma_{21} & \Sigma_{22} \end{pmatrix} \right).$$

Then the conditional distribution of \mathbf{X}_2 given \mathbf{X}_1 is

$$\mathcal{N} \left(\mu_2 + \Sigma_{21} \Sigma_{11}^{-1} (\mathbf{X}_1 - \mu_1), \Sigma_{22} - \Sigma_{21} \Sigma_{11}^{-1} \Sigma_{12} \right).$$

B.2 Conditional Means of Normal and Squared Normal Random Variables

For $z \sim N(0, \sigma_z^2)$, the conditional mean of the normal random variable is

$$\begin{aligned}
\mathbb{E}[z|z > 0] &= \frac{1}{\mathbb{P}[z > 0]} \frac{1}{\sigma_z \sqrt{2\pi}} \int_0^\infty z \times \exp\left(-\frac{z^2}{2\sigma_z^2}\right) dz \\
&= \frac{\sqrt{2}}{\sigma_z \sqrt{\pi}} \int_0^\infty \sigma_z^2 \times \exp(-y) dy \\
&= -\frac{\sigma_z \sqrt{2}}{\sqrt{\pi}} \times \exp(-y) \Big|_0^\infty \\
&= \frac{\sigma_z \sqrt{2}}{\sqrt{\pi}}.
\end{aligned}$$

The second step uses the change of variables, with $y = z^2/2\sigma_z^2$.

The conditional mean of the squared normal random variable is

$$\begin{aligned}
\mathbb{E}[z^2|z > 0] &= \frac{1}{\mathbb{P}[z > 0]} \frac{1}{\sigma_z \sqrt{2\pi}} \int_0^\infty z^2 \times \exp\left(-\frac{z^2}{2\sigma_z^2}\right) dz \\
&= \frac{\sqrt{2}}{\sigma_z \sqrt{\pi}} \int_0^\infty z^2 \times \exp(-az^2) dz, \quad a = \frac{1}{2\sigma_z^2} \\
&= \frac{\sqrt{2}}{\sigma_z \sqrt{\pi}} \left(\frac{1}{4} \sqrt{\frac{\pi}{a^3}} \times \operatorname{erf}(z\sqrt{a}) - \frac{z}{2a} \times \exp(-az^2) \right) \Big|_0^\infty \\
&= \frac{\sqrt{2}}{\sigma_z \sqrt{\pi}} \frac{1}{4} \sqrt{\frac{\pi}{a^3}} \\
&= \sigma_z^2.
\end{aligned}$$

The third step uses $\operatorname{erf}(\cdot)$ to denote the error function.

B.3 Notes for Figures and Tables

B.3.1 Figure 2.4

Panel A plots a daily time series of the variance risk premium, “VRP,” or $VRP_{t,t+22}$; the squared VIX, “sq. VIX,” or $VIX_t^2/12$; and the physical expectation of realized variance, “cond. var.,” or $\mathbb{E}_t^{\mathbb{P}}[RV_{t+1,t+22}]$. $VIX_t^2/12$ is the square of the day-end CBOE VIX index divided by 12. $\mathbb{E}_t^{\mathbb{P}}[RV_{t+1,t+22}]$ is the one-step-ahead forecasts from regressing realized variance $RV_{t+1,t+22}$ on lagged realized variance and lagged squared VIX in Eq. (2.24):

$$RV_{t+1,t+22} = \underset{(1.145)}{2.130} + \underset{(0.047)}{0.310} \times R_{t-21,t} + \underset{(0.050)}{0.464} \times VIX_{t-22}^2/12 + e_t.$$
 I construct $RV_{t+1,t+22}$ as the sum of 22 daily realized variances between $t + 1$ and $t + 22$ (inclusive), with daily realized variance calculated as the sum of squared five-minute log returns on ES futures from 9:30 AM ET to 4:00 PM ET and the squared close-to-open log return. The variance risk premium is the difference between “sq. VIX” and “cond. var.”: $VRP_{t,t+22} = VIX_t^2/12 - \mathbb{E}_t^P [RV_{t+1,t+22}]$. Panel B plots monthly time series of the same three variables as in Panel A taking a simple average of daily data within each month. Panel C plots monthly time series of the same three variables as in Panels A and B constructed in the following manner. “sq. VIX” is the end-of-month squared VIX $VIX_t^2/12$, with t the last trading day of a given month. $\mathbb{E}_t^P [RV_{t+1,t+22}]$ is the one-step-ahead forecasts from the monthly regression of realized variance on one-month lagged realized variance and one-month lagged squared VIX in Eq. (2.25): $RV_{t+1,t+22} = \underset{(2.762)}{0.195} + \underset{(0.144)}{0.282} \times R_{t-21,t} + \underset{(0.098)}{0.511} \times VIX_{t-22}^2/12 + e_t$. Monthly realized variance $RV_{t+1,t+22}$ is the sum of daily realized variances for the month that includes days $t + 1$ to $t + 22$. The variance risk premium for the month with t as its last trading day is the difference between “sq. VIX” and “cond. var.”: $VRP_{t,t+22} = VIX_t^2/12 - \mathbb{E}_t^P [RV_{t+1,t+22}]$. Note that I keep all notation the same as before, but obviously some months have more trading days than other months. All three panels above cover the whole period from November, 1997 to March, 2014.

B.3.2 Figure 2.7

All panels plot data from Table 2.8 with the y-axis denominated in bps and the x-axis representing variance risk premium VRP_t quintiles, with the 1st quintile corresponding to the smallest values and the 5th quintile corresponding to the largest values. Panel A plots $\hat{\beta}^+$ (“p_beta”) and $\hat{\beta}^-$ (“n_beta”). Panel B plots $\hat{\beta}^+ - \hat{\beta}^-$. Panel C plots $\hat{\alpha}$.

B.3.3 Table 2.9

The tables show results from return predictability regressions using monthly data from November, 1997 to March, 2014. Panel A presents results for the specification $r_{t,t+h}^{MNA} =$

$a_{MNA} + b_{MNA}X_t + \epsilon_{t,t+h}^{MNA}$, and Panel B presents results for the specification $r_{t,t+h}^{nMNA} = a_{nMNA} + b_{nMNA}X_t + \epsilon_{t,t+h}^{nMNA}$. $r_{t,t+h}^{MNA}$ is the annualized sum of the continuously compounded return of the stock market in the ± 5 minutes around MNAs from the end of month t to the end of month $t + h$. $r_{t,t+h}^{nMNA} = r_{t,t+h} - r_{t,t+h}^{MNA}$, in which $r_{t,t+h}$ is defined in Table 2.7. Table displays results for horizons $h = 1, 3, 6, 9, 12, 15, 18, 21, 24$ and three sets of covariates X_t drawn from $\log(P/E)_t$, the log of the cyclically adjusted price-to-earnings ratio from Robert Shiller's website, and the "Monthly Averages" variance risk premium series from Panel B of Figure 2.4. Table reports estimated parameters \hat{b} , associated t -statistics, and adjusted R^2 . t -statistics (shown in parentheses) are based on Newey-West standard errors with $\max\{3, 2h\}$ lags. *** denotes significance at the 1% level, ** denotes significance at the 5% level, and * denotes significance at the 10% level.

Appendix C

Appendix to Chapter 3

C.1 Forward Rate

We know that the forward rate

$$f_{n,t} = (n + 1) y_{n+1,t} - n y_{n,t}$$

with $y_{n,t}$ the log of the yield to maturity of a zero-coupon bond of maturity n at time t . We can write this statement as

$$\begin{aligned} f_{n,t} &= -\log \left(\exp(-i_t) \times \mathbb{E}_t^{\mathbb{Q}} \left[\exp \left(-\sum_{j=1}^n i_{t+j} \right) \right] \right) + \log \left(\exp(-i_t) \times \mathbb{E}_t^{\mathbb{Q}} \left[\exp \left(-\sum_{j=1}^{n-1} i_{t+j} \right) \right] \right) \\ &\approx \mathbb{E}_t^{\mathbb{Q}} \left[\sum_{j=1}^n i_{t+j} \right] - \frac{1}{2} \mathbb{V}_t^{\mathbb{Q}} \left[\sum_{j=1}^n i_{t+j} \right] - \mathbb{E}_t^{\mathbb{Q}} \left[\sum_{j=1}^{n-1} i_{t+j} \right] + \frac{1}{2} \mathbb{V}_t^{\mathbb{Q}} \left[\sum_{j=1}^{n-1} i_{t+j} \right] \\ &= \mathbb{E}_t^{\mathbb{Q}} [i_{t+n}] - \frac{1}{2} \left(\mathbb{V}_t^{\mathbb{Q}} \left[\sum_{j=1}^n i_{t+j} \right] - \mathbb{V}_t^{\mathbb{Q}} \left[\sum_{j=1}^{n-1} i_{t+j} \right] \right). \end{aligned} \tag{C.1}$$

$\mathbb{E}_t^{\mathbb{Q}}[\cdot]$ is the conditional expectation and $\mathbb{V}_t^{\mathbb{Q}}[\cdot]$ is the conditional variance. The first equality comes from the fact that $-n y_{n,t} = p_{n,t}$, with $p_{n,t}$ the log price that equals a discounted expectation of the terminal payoff. The approximation comes from the fact that $\log(\mathbb{E}[\exp(Z)]) \approx \mathbb{E}[Z] + \frac{1}{2} \mathbb{V}[Z]$ for any random variable Z .

Substituting Eq. (3.1) into the first term in Eq. (C.1), we see that

$$\begin{aligned}\mathbb{E}_t^{\mathbb{Q}} [i_{t+n}] &= \mathbb{E}_t^{\mathbb{Q}} [\max \{s_{t+n}, \underline{i}\}] \\ &= \Pr_t^{\mathbb{Q}} [s_{t+n} < \underline{i}] \times \underline{i} + \Pr_t^{\mathbb{Q}} [s_{t+n} \geq \underline{i}] \times \mathbb{E}_t^{\mathbb{Q}} [s_{t+n} | s_{t+n} \geq \underline{i}].\end{aligned}\quad (\text{C.2})$$

Based on Eqs. (3.2) and (3.3), the shadow rate has a conditional normal distribution:

$$s_{t+n} | I_t \sim \mathcal{N} \left(\bar{a}_n + b_n x_t, \sigma_n^{\mathbb{Q}} \right) \quad (\text{C.3})$$

with parameters defined as

$$\begin{aligned}\bar{a}_n &\equiv \delta_0 + \delta_1 \mu^{\mathbb{Q}} \frac{1 - (\rho^{\mathbb{Q}})^n}{1 - \rho^{\mathbb{Q}}}, \\ b_n &\equiv \delta_1 (\rho^{\mathbb{Q}})^n, \text{ and} \\ \sigma_n^{\mathbb{Q}} &\equiv \sqrt{(\delta_1 \sigma_\epsilon)^2 \frac{1 - (\rho^{\mathbb{Q}})^{2n}}{1 - (\rho^{\mathbb{Q}})^2}}.\end{aligned}$$

By assumption that $\delta_1 > 0$ and $\rho^{\mathbb{Q}} > 0$, we know that $b_n > 0$. The standard deviation $\sigma_n^{\mathbb{Q}} > 0$ as well. Using the conditional normal distribution of the shadow rate in Eq. (C.3) and the formula for the first moment of the truncated normal distribution, Eq. (C.2) simplifies:

$$\begin{aligned}\mathbb{E}_t^{\mathbb{Q}} [i_{t+n}] &= \underline{i} + \sigma_n^{\mathbb{Q}} \left(\left(\frac{\bar{a}_n + b_n x_t - \underline{i}}{\sigma_n^{\mathbb{Q}}} \right) \Phi \left(\frac{\bar{a}_n + b_n x_t - \underline{i}}{\sigma_n^{\mathbb{Q}}} \right) + \phi \left(\frac{\bar{a}_n + b_n x_t - \underline{i}}{\sigma_n^{\mathbb{Q}}} \right) \right) \\ &= \underline{i} + \sigma_n^{\mathbb{Q}} g \left(\frac{\bar{a}_n + b_n x_t - \underline{i}}{\sigma_n^{\mathbb{Q}}} \right)\end{aligned}\quad (\text{C.4})$$

for $g(x) = x\Phi(x) + \phi(x)$. In order to derive the second term in Eq. (C.1), Wu and Xia (2014) show and utilize the following approximations:

$$\mathbf{V}_t^{\mathbb{Q}} [i_{t+n}] \approx \Pr_t^{\mathbb{Q}} [s_{t+n} \geq \underline{i}] \mathbf{V}_t^{\mathbb{Q}} [s_{t+n}] \text{ and} \quad (\text{C.5})$$

$$\begin{aligned}\mathbf{C}_t^{\mathbb{Q}} [i_{t+n-j}, i_{t+n}] &\approx \Pr_t^{\mathbb{Q}} [s_{t+n-j} \geq \underline{i}, s_{t+n} \geq \underline{i}] \mathbf{C}_t^{\mathbb{Q}} [s_{t+n-j}, s_{t+n}] \\ &\approx \Pr_t^{\mathbb{Q}} [s_{t+n} \geq \underline{i}] \mathbf{C}_t^{\mathbb{Q}} [s_{t+n-j}, s_{t+n}].\end{aligned}\quad (\text{C.6})$$

$\mathbf{C}_t^{\mathbb{Q}} [\cdot]$ is the conditional covariance. The last approximation uses Bayes' law and assumes

that the shadow rate is persistent such that

$$\Pr_t^{\mathbb{Q}} [s_{t+n-j} \geq \underline{i} | s_{t+n} \geq \underline{i}] \approx 1.$$

Making use of the approximations in Eqs. (C.5) and (C.6), the second term in Eq. (C.1) is

$$\begin{aligned} \frac{1}{2} \left(\mathbb{V}_t^{\mathbb{Q}} \left[\sum_{j=1}^n i_{t+j} \right] - \mathbb{V}_t^{\mathbb{Q}} \left[\sum_{j=1}^{n-1} i_{t+j} \right] \right) &\approx \Pr_t^{\mathbb{Q}} [s_{t+n} \geq \underline{i}] \times \frac{1}{2} \left(\mathbb{V}_t^{\mathbb{Q}} \left[\sum_{j=1}^n s_{t+j} \right] - \mathbb{V}_t^{\mathbb{Q}} \left[\sum_{j=1}^{n-1} s_{t+j} \right] \right) \\ &= \Phi \left(\frac{\bar{a}_n + b_n x_t - \underline{i}}{\sigma_n^{\mathbb{Q}}} \right) \times (\bar{a}_n - a_n) \end{aligned} \quad (\text{C.7})$$

with parameter a_n defined as

$$a_n \equiv \bar{a}_n - \frac{1}{2} (\delta \sigma_\varepsilon)^2 \left(\frac{1 - (\rho^{\mathbb{Q}})^n}{1 - \rho^{\mathbb{Q}}} \right)^2.$$

Simplifying Eq. (C.1) by substituting in expressions for the first term derived in Eq. (C.4) and second term derived in Eq. (C.7),

$$\begin{aligned} f_{n,t} &\approx \underline{i} + \sigma_n^{\mathbb{Q}} g \left(\frac{\bar{a}_n + b_n x_t - \underline{i}}{\sigma_n^{\mathbb{Q}}} \right) + \Phi \left(\frac{\bar{a}_n + b_n x_t - \underline{i}}{\sigma_n^{\mathbb{Q}}} \right) \times (\bar{a}_n - a_n) \\ &= \underline{i} + \sigma_n^{\mathbb{Q}} g \left(\frac{\bar{a}_n + b_n x_t - \underline{i}}{\sigma_n^{\mathbb{Q}}} \right) + \sigma_n^{\mathbb{Q}} \frac{\partial g \left(\frac{\bar{a}_n + b_n x_t - \underline{i}}{\sigma_n^{\mathbb{Q}}} \right)}{\partial \bar{a}_n} \times (\bar{a}_n - a_n) \\ &\approx \underline{i} + \sigma_n^{\mathbb{Q}} g \left(\frac{a_n + b_n x_t - \underline{i}}{\sigma_n^{\mathbb{Q}}} \right). \end{aligned}$$

The last line comes from a first-order Taylor approximation. We thus obtain the forward rate $f_{n,t}$ in Eq. (3.4).

C.2 Comparative Statics

I use the following definition in many of the proofs below:

$$z \equiv \frac{\underline{i} - \bar{a}_n - b_n x_t}{\sigma_n^{\mathbb{Q}}}.$$

C.2.1 Proof of Proposition 1

Eq. (3.5) gives the probability that the ZLB binds at some time $t + n$. As x_t decreases, this probability increases:

$$\frac{\partial}{\partial x_t} \left(\Pr_t^Q [s_{t+n} < \underline{i}] \right) = \phi \left(\frac{\underline{i} - \bar{a}_n - b_n x_t}{\sigma_n^Q} \right) \times \left(-\frac{b_n}{\sigma_n^Q} \right) < 0.$$

As δ_1 decreases, the probability that the ZLB is still binding also increases under certain conditions:

$$\frac{\partial}{\partial \delta_1} \left(\Pr_t^Q [s_{t+n} < \underline{i}] \right) < 0.$$

To see this, expand this probability using the model parameters:

$$\begin{aligned} \Pr_t^Q [s_{t+n} < \underline{i}] &= \Phi \left(\frac{\underline{i} - \bar{a}_n - b_n x_t}{\sigma_n^Q} \right) \\ &= \Phi \left(\frac{\underline{i} - \delta_0 + \delta_1 \mu^Q \frac{1 - (\rho^Q)^n}{1 - \rho^Q} - \delta_1 (\rho^Q)^n x_t}{\delta_1 \sqrt{(\sigma_\epsilon)^2 \frac{1 - (\rho^Q)^{2n}}{1 - (\rho^Q)^2}}} \right) \\ &= \Phi \left(\frac{\underline{i} - \delta_0}{\delta_1 \sqrt{(\sigma_\epsilon)^2 \frac{1 - (\rho^Q)^{2n}}{1 - (\rho^Q)^2}}} + \frac{\mu^Q \frac{1 - (\rho^Q)^n}{1 - \rho^Q} - (\rho^Q)^n x_t}{\sqrt{(\sigma_\epsilon)^2 \frac{1 - (\rho^Q)^{2n}}{1 - (\rho^Q)^2}}} \right). \end{aligned}$$

The second term in the CDF doesn't depend on δ_1 . Assuming that $\underline{i} - \delta_0 > 0$, the first term becomes arbitrarily large as δ_1 nears 0, which implies that the CDF becomes arbitrarily close to 1 in value.

C.2.2 Proof of Proposition 2

To show that $\partial f_{n,t} / \partial x_t > 0$, I take the comparative static

$$\frac{\partial f_{n,t}}{\partial x_t} = \Phi \left(\frac{a_n + b_n x_t - \underline{i}}{\sigma_n^Q} \right) \times b_n > 0.$$

C.2.3 Proof of Proposition 3

To show that $\partial/\partial x_t (\partial f_{n,t}/\partial x_t) > 0$, I take the comparative static

$$\frac{\partial}{\partial x_t} \left(\frac{\partial f_{n,t}}{\partial x_t} \right) = \phi \left(\frac{a_n + b_n x_t - i}{\sigma_n^Q} \right) \times \frac{b_n^2}{\sigma_n^Q} > 0.$$

To show that $\partial/\partial \delta_1 (\partial f_{n,t}/\partial x_t) > 0$, I first make the following claim.

Claim 1 $\partial z/\partial \delta_1 > 0$ for $\delta_1 < \bar{\delta}$ for some $\bar{\delta}$.

The comparative static

$$\frac{\partial}{\partial \delta_1} \left(\frac{\partial f_{n,t}}{\partial x_t} \right) = \phi(z) \times \frac{\partial z}{\partial \delta_1} b_n + \Phi(z) \times (\rho^Q)^n$$

is guaranteed to be positive when $\delta_1 < \bar{\delta}$ because $\partial z/\partial \delta_1 > 0$ and the other terms are all positive as well.

C.2.3.1 Proof of Claim 1

To sign $\partial z/\partial \delta_1$, consider only the numerator since the denominator is positive:

$$\frac{\partial}{\partial \delta_1} (a_n + b_n x_t - i) \times \sigma_n^Q - \frac{\partial}{\partial \delta_1} (\sigma_n^Q) \times (a_n + b_n x_t - i).$$

In the expressions below, the first line simplifies the above expression for the numerator and each successive arrow ignores positive terms that are factored out.

$$\begin{aligned} & \frac{\partial}{\partial \delta_1} (a_n + b_n x_t - i) \times \sigma_n^Q - \frac{\sigma_n^Q}{\delta_1} \times (a_n + b_n x_t - i) \\ \implies & \mu^Q \frac{1 - (\rho^Q)^n}{1 - \rho^Q} - \delta_1 \sigma_\epsilon^2 \left(\frac{1 - (\rho^Q)^n}{1 - \rho^Q} \right)^2 + (\rho^Q)^n x_t \\ & - \left(\frac{\delta_0}{\delta_1} + \mu^Q \frac{1 - (\rho^Q)^n}{1 - \rho^Q} - \frac{1}{2} \delta_1 \sigma_\epsilon^2 \left(\frac{1 - (\rho^Q)^n}{1 - \rho^Q} \right)^2 + (\rho^Q)^n x_t - \frac{i}{\delta_1} \right) \\ \implies & \frac{i - \delta_0}{\delta_1} - \frac{1}{2} \delta_1 \sigma_\epsilon^2 \left(\frac{1 - (\rho^Q)^n}{1 - \rho^Q} \right)^2. \end{aligned}$$

This last line must be positive in order for $\partial z/\partial \delta_1 > 0$. This occurs when

$$\underline{i} - \delta_0 > \frac{1}{2} (\delta_1)^2 \sigma_\epsilon^2 \left(\frac{1 - (\rho^Q)^n}{1 - \rho^Q} \right)^2 \implies \delta_1 < \sqrt{\frac{2}{\sigma_\epsilon} (\underline{i} - \delta_0) \frac{1 - \rho^Q}{1 - (\rho^Q)^n}} \equiv \bar{\delta}.$$

For δ_1 that is not too large, we have $\partial z/\partial \delta_1 > 0$. This proof also relies on the assumption in Proposition 1 that $\underline{i} - \delta_0 > 0$. Assuming that σ_ϵ is not too large, there should be a sufficiently large range of values for δ_1 for which Claim 1 holds.

C.2.4 Proof of Proposition 4

The comparative static $\partial/\partial x_t (\partial/\partial x_t (\partial f_{n,t}/\partial x_t))$ is below:

$$\begin{aligned} \frac{\partial}{\partial n} \left(\frac{\partial}{\partial x_t} \left(\frac{\partial f_{n,t}}{\partial x_t} \right) \right) &= \frac{\partial}{\partial x_t} \left(\frac{\partial}{\partial n} \left(\frac{\partial f_{n,t}}{\partial x_t} \right) \right) \\ &= \frac{\partial}{\partial x_t} \left(\Phi(z) \times \ln(\rho^Q) \times \delta_1 (\rho^Q)^n + \phi(z) \times \frac{\partial z}{\partial n} \delta_1 (\rho^Q)^n \right) \\ &= \delta_1 (\rho^Q)^n \left(\phi(z) \times \ln(\rho^Q) \times \frac{\partial z}{\partial x_t} + \phi'(z) \times \frac{\partial z}{\partial x_t} \frac{\partial z}{\partial n} + \phi(z) \times \frac{\partial}{\partial x_t} \left(\frac{\partial z}{\partial n} \right) \right) \\ &= \delta_1 (\rho^Q)^n \left(\phi(z) \times \ln(\rho^Q) \times \frac{b_n}{\sigma_n^Q} - z \phi(z) \times \frac{b_n}{\sigma_n^Q} \frac{\partial z}{\partial n} + \phi(z) \times \frac{\partial}{\partial x_t} \left(\frac{\partial z}{\partial n} \right) \right) \\ &= \delta_1 (\rho^Q)^n \phi(z) \left(\ln(\rho^Q) \times \frac{b_n}{\sigma_n^Q} - z \times \frac{b_n}{\sigma_n^Q} \frac{\partial z}{\partial n} + \frac{\partial}{\partial x_t} \left(\frac{\partial z}{\partial n} \right) \right). \quad (\text{C.8}) \end{aligned}$$

To sign this expression, I make the following claim.

Claim 2 $\partial z/\partial n > 0$ for $x_t < \bar{x}_1$ for some \bar{x}_1 . Moreover, $\partial/\partial x_t (\partial z/\partial n) < 0$ and is independent of x_t .

Now consider the 3 terms inside the large parentheses in Eq. (C.8). The first term is negative because $\ln(\rho^Q) < 0$ and $b_n/\sigma_n^Q > 0$. The third term is also negative because $\partial/\partial x_t (\partial z/\partial n) < 0$ by Claim 2. Moreover, the first and third terms are independent of x_t by Claim 2. Assume that $x_t < \bar{x}_1$ such that $\partial z/\partial n > 0$ by Claim 2. As x_t decreases, z becomes more negative, which implies that the second term is overall positive. For $x_t < \bar{x}_2$, the

overall expression is positive, as the second term outweighs the first and third terms, which are independent of x_t . Define $\bar{x} \equiv \min \{\bar{x}_1, \bar{x}_2\}$. Then we have the final result that

$$\frac{\partial}{\partial n} \left(\frac{\partial}{\partial x_t} \left(\frac{\partial f_{n,t}}{\partial x_t} \right) \right) > 0 \text{ for } x_t < \bar{x}. \quad (\text{C.9})$$

Evaluating the second part of Proposition 4, the comparative static $\partial/\partial x_t (\partial/\partial x_t (\partial f_{n,t}/\partial x_t))$ is

$$\begin{aligned} \frac{\partial}{\partial n} \left(\frac{\partial}{\partial \delta_1} \left(\frac{\partial f_{n,t}}{\partial x_t} \right) \right) &= \frac{\partial}{\partial n} \left(\phi(z) \times \frac{\partial z}{\partial \delta_1} b_n + \Phi(z) \times (\rho^Q)^n \right) \\ &= \frac{\partial}{\partial n} \left(\phi(z) \times \frac{\partial z}{\partial \delta_1} b_n \right) + \frac{\partial}{\partial n} \left(\Phi(z) \times (\rho^Q)^n \right). \end{aligned} \quad (\text{C.10})$$

I make the following claims.

Claim 3 $\partial/\partial n \left(\Phi(z) \times (\rho^Q)^n \right) > 0$ for $x_t < \bar{x}_3$, for some \bar{x}_3 .

Claim 4 $\partial/\partial n \left(\phi(z) \times (\partial z/\partial \delta_1) b_n \right) > 0$ for $x_t < \bar{x}_4$ for some \bar{x}_4 and for $\delta_1 < \bar{\delta}$ for the same $\bar{\delta}$ in Claim 1.

Assuming that $\delta_1 < \bar{\delta}$ and $x_t < \min \{\bar{x}_3, \bar{x}_4\} \equiv \bar{x}^*$, Claim 3 and Claim 4 imply that both terms in Eq. (C.10) are positive. Then we have the final result that

$$\frac{\partial}{\partial n} \left(\frac{\partial}{\partial \delta_1} \left(\frac{\partial f_{n,t}}{\partial x_t} \right) \right) > 0 \text{ for } x_t < \bar{x}^* \text{ and } \delta_1 < \bar{\delta}. \quad (\text{C.11})$$

Eqs. (C.9) and (C.11) show Proposition 4. When economic conditions are bad enough and the ZLB is sufficiently binding by Proposition 1, the attenuation in the sensitivity of forward rates to news in Proposition 3 is greater for longer maturity forward rates. On the other hand, when economic conditions are good enough and the ZLB does not bind severely by Proposition 1, the attenuation in the sensitivity of forward rates to news in Proposition 3 is smaller for longer maturity forward rates.

C.2.4.1 Proof of Claim 2

The comparative static $\partial z/\partial n$ is

$$\begin{aligned}\frac{\partial z}{\partial n} &= \frac{\partial}{\partial n} \left(\frac{a_n + b_n x_t - \dot{i}}{\sigma_n^Q} \right) \\ &= \frac{\frac{\partial}{\partial n} (a_n + b_n x_t - \dot{i}) \times \sigma_n^Q - \frac{\partial}{\partial n} (\sigma_n^Q) \times (a_n + b_n x_t - \dot{i})}{(\sigma_n^Q)^2}.\end{aligned}$$

We first show that $\partial/\partial x_t (\partial z/\partial n) < 0$ and is independent of x_t by considering only the numerator of $\partial z/\partial n$ and the terms involving x_t since the denominator is positive and independent of x_t :

$$\frac{\partial}{\partial n} (b_n x_t) \times \sigma_n^Q - \frac{\partial}{\partial n} (\sigma_n^Q) \times b_n x_t.$$

In the expressions below, the first line expands the above expression for terms in the numerator related to x_t and each successive arrow ignores positive terms that are factored out:

$$\begin{aligned}& \ln(\rho^Q) \times \delta_1 (\rho^Q)^n x_t \sigma_n^Q - \frac{1}{2} \frac{1}{\sigma_n^Q} (\delta_1 \sigma_\epsilon)^2 \frac{-2 \ln(\rho^Q) \times (\rho^Q)^{2n}}{1 - (\rho^Q)^2} \delta_1 (\rho^Q)^n x_t \\ \implies & \ln(\rho^Q) \times \delta_1 (\rho^Q)^n x_t (\delta_1 \sigma_\epsilon)^2 \frac{1 - (\rho^Q)^{2n}}{1 - (\rho^Q)^2} + (\delta_1 \sigma_\epsilon)^2 \frac{\ln(\rho^Q) \times (\rho^Q)^{2n}}{1 - (\rho^Q)^2} \delta_1 (\rho^Q)^n x_t \\ \implies & \ln(\rho^Q) \times x_t \left(1 - (\rho^Q)^{2n} \right) + \ln(\rho^Q) \times (\rho^Q)^{2n} x_t \\ \implies & \ln(\rho^Q) \times x_t.\end{aligned}$$

Since $\rho^Q < 1$ by assumption of stationarity, $\ln(\rho^Q) < 0$ and $\partial/\partial x_t (\partial z/\partial n) < 0$ and is moreover independent of x_t . Since $\partial z/\partial n$ is decreasing in x_t , for $x_t < \bar{x}_1$ for some threshold \bar{x}_1 , $\partial z/\partial n > 0$.

C.2.4.2 Proof of Claim 3

Expand the comparative static $\partial/\partial n \left(\Phi(z) \times (\rho^Q)^n \right)$:

$$\begin{aligned} \frac{\partial}{\partial n} \left(\Phi(z) \times (\rho^Q)^n \right) &= \phi(z) \times \frac{\partial z}{\partial n} (\rho^Q)^n + \Phi(z) \times \ln(\rho^Q) \times (\rho^Q)^n \\ &= (\rho^Q)^n \left(\phi(z) \times \frac{\partial z}{\partial n} + \Phi(z) \times \ln(\rho^Q) \right). \end{aligned}$$

Thus,

$$\frac{\partial}{\partial n} \left(\Phi(z) \times (\rho^Q)^n \right) > 0 \iff \frac{\Phi(z)}{\phi(z)} < -\frac{1}{\ln(\rho^Q)} \frac{\partial z}{\partial n}. \quad (\text{C.12})$$

We know from Claim 2 that $\partial/\partial x_t (\partial z/\partial n) < 0$ and is independent of x_t , so the right-hand side of the inequality in Eq. (C.12) becomes arbitrarily positive as x_t becomes arbitrarily negative. Moreover, as x_t becomes arbitrarily negative, z becomes arbitrarily negative as well. Using the relationship that

$$\lim_{z \rightarrow -\infty} \frac{\Phi(z)}{\phi(z)} = \lim_{-z \rightarrow \infty} \frac{1 - \Phi(-z)}{\phi(z)} \approx \frac{1}{-z},$$

it's clear that at some point, the left-hand side of the inequality in Eq. (C.12) becomes an arbitrarily small positive number. Thus the inequality in Eq. (C.12) holds for $x_t < \bar{x}_3$ for some \bar{x}_3 , which proves Claim 3.

C.2.4.3 Proof of Claim 4

Expand the comparative static $\partial/\partial n \left(\phi(z) \times (\partial z/\partial \delta_1) b_n \right)$:

$$\begin{aligned} \frac{\partial}{\partial n} \left(\phi(z) \times \frac{\partial z}{\partial \delta_1} b_n \right) &= \frac{\partial}{\partial n} \left(\phi(z) \times \delta_1 (\rho^Q)^n \right) \frac{\partial z}{\partial \delta_1} + \left(\phi(z) \times \delta_1 (\rho^Q)^n \right) \frac{\partial}{\partial n} \frac{\partial z}{\partial \delta_1} \\ &= \left(\phi'(z) \times \frac{\partial z}{\partial n} \delta_1 (\rho^Q)^n + \phi(z) \times \delta_1 \ln(\rho^Q) \times (\rho^Q)^n \right) \frac{\partial z}{\partial \delta_1} \\ &\quad + \left(\phi(z) \times \delta_1 (\rho^Q)^n \right) \frac{\partial}{\partial n} \frac{\partial z}{\partial \delta_1} \\ &= \left(-z\phi(z) \times \frac{\partial z}{\partial n} \delta_1 (\rho^Q)^n + \phi(z) \times \delta_1 \ln(\rho^Q) \times (\rho^Q)^n \right) \frac{\partial z}{\partial \delta_1} \\ &\quad + \left(\phi(z) \times \delta_1 (\rho^Q)^n \right) \frac{\partial}{\partial n} \frac{\partial z}{\partial \delta_1} \\ &= \delta_1 \phi(z) \times (\rho^Q)^n \left(-z \frac{\partial z}{\partial n} \frac{\partial z}{\partial \delta_1} + \ln(\rho^Q) \frac{\partial z}{\partial \delta_1} + \frac{\partial}{\partial n} \frac{\partial z}{\partial \delta_1} \right). \quad (\text{C.13}) \end{aligned}$$

The proof for Claim 1 shows that $\partial z / \partial \delta_1$ is independent of x_t , so the second and third terms of Eq. (C.13) are likewise independent of x_t . Claim 1 also says that $\delta_1 < \bar{\delta}$ implies $\partial z / \partial \delta_1 > 0$. Then for $x_t < \bar{x}_4$, we have z (increasing in x_t) sufficiently negative and $\partial z / \partial n$ (decreasing in x_t based on Claim 2) sufficiently positive such that the first term in Eq. (C.13) is positive and of large enough magnitude that the overall expression $\partial / \partial n (\phi(z) \times (\partial z / \partial \delta_1) b_n) > 0$, which proves Claim 4.

C.3 Notes for Figures

C.3.1 Figure 3.2

Focusing on Panel A, the right-hand plot shows two time series. “fcsts,” plotted on the right y-axis, is the median number of quarters until the first rate hike based on primary dealer surveys by the FRBNY. “beta,” plotted on the left y-axis, comes from a daily 1-year rolling regression that aggregates all 18 MNAs: $\Delta f_{n,t} = \alpha_n + \beta_n S_t + \epsilon_{n,t}$. $\Delta f_{n,t}$ is the basis points change in the forward rate implied from the ED n contract in a ± 5 minute window around a MNA. The right-hand side variable S_t aggregates $S_{i,t}$ for MNA i . $S_{i,t}$ is defined as the actual data less the expected data based on a Bloomberg survey of economists, all normalized by the sample standard deviation of the actual less the expected difference: $S_{i,t} = (A_{i,t} - \mathbb{E}_{t-}[A_{i,t}]) / \hat{\sigma}_i$. “beta” is specifically the $\hat{\beta}_{n,t}$ estimated parameter from running the aforementioned regression on a 1-year window centered at t . The two vertical lines correspond to the dates of the August, 2011 FOMC meeting (when calendar guidance was introduced) and the December, 2012 FOMC meeting (when data thresholds was introduced), respectively. The left-hand plot is simply a scatterplot of “beta” on “fcsts.” Each point in the scatterplot is color-coded based on the timestamp: red for 12/17/08 to 8/9/11 corresponding to the period of general forward guidance (“fwd”); green for 8/10/11 to 12/12/12 corresponding to the period of calendar guidance (“cdr”); and blue for 12/13/12 to 3/6/14 corresponding to the period of data thresholds (“dat”). In a given period, lighter colors indicate earlier dates. For Panels A, B, C, and D, $n = 1, 4, 8, \text{ and } 16$, respectively.

Panel E replaces $\Delta f_{n,t}$ with R_t , the % change in the price level of the TY contract. Thus, $\hat{\beta}_{n,t}$ is measured in basis points for Panels A, B, C, and D but in % for Panel E.

C.3.2 Figure 3.3

Focusing on Panel A, the right-hand plot shows two time series. “FF24,” plotted on the right y-axis, is the implied FF rate (%) from the FF24 futures contract. “beta,” plotted on the left y-axis, comes from a daily 1-year rolling regression that aggregates all 18 MNAs: $\Delta f_{n,t} = \alpha_n + \beta_n S_t + \epsilon_{n,t}$. $\Delta f_{n,t}$ is the basis points change in the forward rate implied from the ED n contract in a ± 5 minute window around a MNA. The right-hand side variable S_t aggregates $S_{i,t}$ for MNA i . $S_{i,t}$ is defined as the actual data less the expected data based on a Bloomberg survey of economists, all normalized by the sample standard deviation of the actual less the expected difference: $S_{i,t} = (A_{i,t} - \mathbb{E}_{t-}[A_{i,t}]) / \hat{\sigma}_i$. “beta” is specifically the $\hat{\beta}_{n,t}$ estimated parameter from running the aforementioned regression on a 1-year window centered at t . The two vertical lines correspond to the dates of the August, 2011 FOMC meeting (when calendar guidance was introduced) and the December, 2012 FOMC meeting (when data thresholds was introduced), respectively. The left-hand plot is simply a scatterplot of “beta” on “FF24.” Each point in the scatterplot is color-coded based on the timestamp: red for 12/17/08 to 8/9/11 corresponding to the period of general forward guidance (“fwd”); green for 8/10/11 to 12/12/12 corresponding to the period of calendar guidance (“cdr”); and blue for 12/13/12 to 3/6/14 corresponding to the period of data thresholds (“dat”). In a given period, lighter colors indicate earlier dates. For Panels A, B, C, and D, $n = 1, 4, 8,$ and $16,$ respectively. Panel E replaces $\Delta f_{n,t}$ with R_t , the % change in the price level of the TY contract. Thus, $\hat{\beta}_{n,t}$ is measured in basis points for Panels A, B, C, and D but in % for Panel E.

C.3.3 Figure 3.4

Focusing on Panel A, the right-hand plot shows two time series. “fcsts,” plotted on the right y-axis, is the median number of quarters until the first rate hike based on primary

dealer surveys by the FRBNY. “beta,” plotted on the left y-axis, comes from a daily 1-year rolling regression that aggregates all 18 MNAs: $R_t = \alpha_s + \beta_s S_t + \epsilon_{s,t}$. R_t is the % change in the price level of the ES contract in a ± 5 minute window around a MNA. The right-hand side variable S_t aggregates $S_{i,t}$ for MNA i . $S_{i,t}$ is defined as the actual data less the expected data based on a Bloomberg survey of economists, all normalized by the sample standard deviation of the actual less the expected difference: $S_{i,t} = (A_{i,t} - \mathbb{E}_{t-} [A_{i,t}]) / \hat{\sigma}_i$. “beta” is specifically the $\hat{\beta}_{s,t}$ estimated parameter from running the aforementioned regression on a 1-year window centered at t . The two vertical lines correspond to the dates of the August, 2011 FOMC meeting (when calendar guidance was introduced) and the December, 2012 FOMC meeting (when data thresholds was introduced), respectively. The left-hand plot is simply a scatterplot of “beta” on “fcsts.” Each point in the scatterplot is color-coded based on the timestamp: red for 12/17/08 to 8/9/11 corresponding to the period of general forward guidance (“fwd”); green for 8/10/11 to 12/12/12 corresponding to the period of calendar guidance (“cdr”); and blue for 12/13/12 to 3/6/14 corresponding to the period of data thresholds (“dat”). In a given period, lighter colors indicate earlier dates. Panel B is identical to Panel A, except the former replaces “fcsts” with “FF24,” the implied FF rate (%) from the FF24 futures contract.