Essays on Empirical Macroeconomics and Expectations

by

Mardoqueo Arteaga
M.A., Fordham University, 2020
B.A., Montclair State University, 2018

Dissertation submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy

to the

Faculty of the Department of Economics
Graduate School of Arts and Sciences
Fordham University

New York
May 2023

I dedicate this dissertation to Nathaniel and Brandon Arteaga, as well as the original $\sim$800,000 DACA recipients in the United States. *Si se puede.*

Acknowledgments

Nervous — very, very dreadfully nervous I had been for most of my life that I would not "make it". Defying expectations, this dissertation stands testament less so to my perseverance and insatiable curiosity than to the steadfast support I have received on the journey. No one can go from naive enthusiasm to the pinnacle of success in any professional sphere without those who take their time to guide, mentor, and befriend.

Foremost thanks to Arunima Sinha who let my imagination run wild. Thank you for encouraging me to see this endeavor through from its infancy to fruition. I also extend my heartfelt gratitude to Ralf Hepp and Jerome Lahaye for their support. To Johanna Francis, I offer special thanks for her unwavering friendship; you are a cornerstone of my triumph in this program.

My experience was enriched by the diverse communities that embraced me. Among these are my first year cohort, the Graduate Student Support Group, the Graduate Student Council (GSC), and GSAS. It was an honor to be elected to the GSC and share many memories with Fordham's remarkable students. I am indebted to the distinguished Deans with whom I collaborated with for instilling in me the principles of service and leadership, and I extend my warmest thanks to those in GSAS who ensured I never felt alone: Ann Gaylin, Sara Lehman, Chelsea Smith, and Mary Louise Santoli.

Outside of Fordham, I thank a multitude of individuals who provided the often unsung personal support we need when undertaking the impossible. I am forever grateful to Vidya Atal and Jim Anderson for making my education a reality; I stand upon their shoulders today. Others include Fabian Wander, John Peppas, Chris D'Amato, Roger Ferguson, David Coble, Alberto Brea, Sherif Assef, Maria Lynch, and the Oliasami Family. I also thank the cavalcade of friends that made the five years of my early to mid 20s an adventure: Kevin Jean-Pierre, Jeffrey Yozwiak, Kent Bhupathi, Jay Shah, Carlos Panjon, Cam Hill, Will Storie, Brandon Harris, Meeraj Patel, Sarah Dimichino, Francesco Montagnani, Erica Schaumberg, Joy Mishra, and Alexey Skorokhodov.

Lastly, a thank you to my mother and her partner who, in their own way, always wanted the best for me. To my spiritual guide and confidant, Joseph, who found and saved me from myself many times; to my partner Nicole who allows me to see my reflection in a more forgiving mirror; to my brother Nathaniel, whose strength is unparalleled, and to my youngest brother, Brandon, who has one of the biggest hearts - thank you for believing in me and making this stage of my life worthwhile.

As a formerly undocumented immigrant from El Salvador, I cannot celebrate this intellectual milestone in isolation from the adversities I have surmounted. This achievement may appear to defy statistical probability, but it is a feat that was not accomplished in solitude. We move, and live, and have our being all thanks to those we surround ourselves with. This is especially true if you are like me: no preconception of what to do, no blueprint to follow, no net to fall back on — new. The new needs friends and I am so grateful to those who have been mine. I strive to pay it back in this life.

Epigraph

Tengo que decir de dónde vine,
porque todos los que conmigo llegaron
han olvidado aquel país sin cuerpos.

Claudia Lars (1899-1974)
Salvadoran poet born as Margarita del Carmen Brannon Vega

…In the self-forgetfulness of concentrated attention
the door opens for you into a pure living intimacy.
A shared, timeless happiness, conveyed by a smile, a wave of the hand.
Thanks to those who have taught me this.
Thanks to the days which have taught me this.

Dag Hammarskjöld (1905-1961)
Swedish economist, diplomat, and the 2nd United Nations Secretary-General

Contents

List of Tables

List of Figures

1
Introduction

The global economy is a complex and dynamic system, in which the interactions among various agents shape the patterns of economic growth. One of the key aspects of these interactions is the formation of expectations and beliefs, which play a crucial role in driving economic decisions and the business cycle. This dissertation has three distinct but interconnected chapters, each investigating different facets of expectations and beliefs in the context of macroeconomic outcomes, credit markets, and household sentiment. By examining these aspects, it aims to provide a comprehensive understanding of how expectations and beliefs are formed, evolve, and influence economic activity.

The first chapter, Monetary Policy Announcements and Household Expectations of the Future, explores the impact of FOMC announcements on household expectations measured by the Survey of Consumer Expectations from the New York Federal Reserve Bank. By using an event study approach around a three week window of the announcement, the research estimates the causal relationship behind movements in expectations related to interest rates, inflation, and commodity prices. This chapter contributes to the literature by finding novel estimates of central bank communication through its announced decisions onto household expectations, as well as providing medium-run estimates of how long those immediate impacts last.

The second chapter, Credit Market Expectations and the Business Cycle: Evidence from a Textual Analysis Approach, explores the relationship between credit spread expectations and macroeconomic outcomes, using a novel textual analysis approach to derive a proxy for historical credit spread expectations. By analyzing the front pages of the Wall Street Journal from 1919Q1 to 2022Q3 using natural language processing and topic models, the research constructs a measure of credit spread expectation errors and investigates their ability to predict future GDP, unemployment, and private domestic investment. This chapter contributes to the literature by demonstrating the value of textual analysis in uncovering the drivers of credit market sentiment, as well as providing empirical support for behavioral models that posit elevated sentiment as a precursor to declines in economic activity.

The third chapter, Household Sentiment Analysis through a Hierarchical Bayesian Latent Class Model, shifts the focus to the heterogeneity in beliefs and expectations among different economic agents, specifically households, using the Survey of Consumer Expectations (SCE). This chapter employs a Latent Dirichlet Analysis for Survey Data (LDA-S) approach to model heterogeneity in beliefs as differences in individual information choice. It also shows how these belief types have economic significance in predicting inflation expectations of households. By estimating multiple respondent belief types and exploring their association with demographic and personal variables, this study uncovers valuable insights into how households form expectations and how these expectations can influence economic decision-making and policy efficacy.

By examining these issues from different perspectives and employing novel methodologies, this dissertation highlights the importance of expectations and beliefs in shaping economic activity and underscores the need for further research to better inform policy-making and economic forecasting.

[T]he Federal Reserve's ability to influence economic conditions today depends critically on its ability to shape expectations of the future, specifically by helping the public understand how it intends to conduct policy over time, and what the likely implications of those actions will be for economic conditions.

Vice Chair Janet L. Yellen, April 2013 (Yellen, 2013)

2

Monetary Policy Announcements and Household Expectations of the Future

Household expectations about economic conditions, such as inflation, drive a wide range of decisions that include saving, borrowing, and consumption. To proactively manage economic overheating or slack, then, the central bank must strive to relay information about the future policy path while concurrently understanding how these expectations change. This facet of central banking is especially crucial in a low interest rate environment where the central bank is limited in its ability lowering rates, also known as an effective lower bound. In theory, transparent communication makes policy decisions more predictable and creates more policy options at the effective lower bound (Blinder, 2009), but in practice, this effectiveness hinges on the communication actually changing beliefs of economic agents in the right direction and in the adequate amount of time. Recent studies on how effective the central bank is in steering expectations are contradictory: one strand finds that central bank communication is effective in changing not only the expectations of households but also their future behavior (Binder, 2017a; Kryvtsov and Petersen, 2021; Ehrmann and Wabitsch, 2022) while the competing narrative is that the same communication has little effect on both (Lamla and Vinogradov, 2019; Coibion et al., 2020c, 2022). This paper presents novel evidence on monetary policy announcement effectiveness on a variety of household expectations; it distinguishes between various types of monetary policy communications that are unified or multidimensional as well as exploiting microeconometric and time series methods to estimate these effects at different time horizons.

In this paper, I explore the effects of different types of monetary policy announcements by the Federal Open Market Committee (hereafter FOMC) on household expectations in the United States over the period 2013 to 2021. I use the micro data on household expectations from the Survey of Consumer Expectations by the Federal Reserve Bank of New York. Monetary policy announcement identification takes into account the competing considerations that a unified measure that is simple to interpret would be more easily digestible to the public while maintaining that policy has become more multidimensional ever since the global financial crisis of 2008-09. I contribute to the understanding of household reactions to policy decisions by testing whether these monetary policy announcements have meaningful impacts on household beliefs, specifically future macroeconomic conditions, such as inflation, and their own financial conditions, such as income.

The key driver of the relevance of household expectations in economic analysis is that expectations will translate into behavior for saving, borrowing, and consumption; those behaviors cannot be directly observed given the context of this analysis but expectations about these variables can forewarn whether or not households will plan to react to monetary policy changes. The primacy of understanding how announcements affect household expectations is highlighted by the current economic climate. Recent years of low interest rates and lower bound constraints on the nominal interest rates by central banks lend themselves to an environment where managing public expectations is crucial (Coibion et al., 2020a,b,c);

scholars and policymakers alike, such as in the introductory quote, have expressed the importance of understanding the efficacy of communication on guiding expectations.

The current literature in this space has discovered a number of relevant patterns that guide my approach. Lab experiments, such as the one's that are currently a mandate for the European Central Bank (Kryvtsov and Petersen, 2021), have given insights on how different types of communications can influence expectations for economic variables such as inflation. Haldane and McMahon (2018) show that populations without specialized background in economics can be effectively communicated with via simple messaging, and the academic literature has also made usage of randomized controlled trials (RCTs) to test communication strategies. Coibion et al. (2022) show that information provided in an RCT to a household relating to the Federal Reserve Bank's inflation target can be just as effective as the monetary policy announcements that are made by the FOMC. Combined, the literature gives rise to two stylized facts: (i) central banks need to have clear messaging if they are to be understood (indeed, announcements made by the FOMC are seldom simple in their structure; other monetary authorities, like the European Central Bank, publish statements that have been found to require up to 15 years of formal education to read (Coenen et al., 2017)); and (ii) in isolation, like in an experiment, the central bank is guaranteed to pass along their signal to the recipient and thus establish a causal interpretation. That isolated incidents result in causal interpretations can be a natural upside to this approach but it is also the largest downside of these studies for seldom are households isolated in the information they acquire in daily life.

This paper mutes the isolated experiment concern of the causality studies by analyzing the survey responses from the SCE relative to FOMC meeting announcements as an exogenous and unrelated occurrence; causality here is derived from the tight time window around the FOMC meeting announcement and from the random timing in this window when the survey responses are elicited. I use a wide range of expectations afforded by the SCE to understand which household expectations are affected, if any. Further, I aggregate the panel data from the survey at the monthly level and use local projections to estimate the varying time effects of monetary policy announcements over a 12 month horizon; these medium run effects may help disentangle the immediate reactions (of which over- and under-reactions have been documented in this space) with longer term effects due to information rigidities.

My main results are that expectations for the probability in increasing interest rates one year ahead, one year ahead inflation, and one year ahead home price growth are robustly affected by a number of monetary policy measures. A one standard deviation tightening surprise in the Federal Funds Rate, for instance, leads to a downwards revision of one year ahead inflation expectations by 0.21% of its overall mean. This finding complements the results from Lewis et al. (2019) who discover a similar downward revision after a surprise Federal Funds Rate tightening using the Gallup consumer surveys. A one standard deviation surprise in the unified monetary policy measure leads to an upwards revision of one year ahead inflation expectations by 3.6% of its overall mean. Similar to Lewis et al. (2019) and Lamla and Vinogradov (2019), I find no effects of monetary policy announcements on longer term (24- to 36-months ahead) expectations for inflation, home price growth, or a variety of commodity prices that households would find relevant to their overall financial health. In the corroborating analysis, the local projections affirm the event study results and I find that there are delayed responses to the different measures of monetary policy; this implies that households take time to digest monetary policy announcements. Lastly, I also explore how these policy measures affect a proxy for media channels, namely using Google Trends. I find that the measures from the event study and local projections once more appear salient in affecting public interest, with changes in the Federal Funds Rate and in the unified monetary policy factor significantly increasing search intensity in public searches.

Outline. In the next section, I present the context that this research has in the literature in more detail. Section 2.2 presents the survey data and the measures of FOMC meeting announcements, as well as the process by which I narrow down the sample of my analysis. Section 2.3 lays out both of empirical specifications I follow for FOMC meeting announcement effects, with Section 2.4 presenting the baseline results. Section 2.5 provides a brief discussion into transmission channels, and Section 2.6 concludes.

2.1 Context in the Literature

This paper contributes to two active research areas. The first concerns itself with central bank communication with economic agents in the general sense. Ever since the global financial crisis, scholarship began to focus on the efficiency of monetary policy through its means of being disseminated amongst the public; Chairman Alan Greenspan, who once prided himself on *"mumbling with great incoherence"* when speaking about policy in a public setting (Resche, 2004), was by 2003 explicitly encouraging that the Fed manage expectations in a more transparent way (Blinder et al., 2008). Central banks have since been actively involved in setting expectations through their communications to the public.

Research coming out of central banks has focused on how different types of economic agents are able to process policy announcements in a variety of ways. Campbell et al. (2012) study how these types of announcements affect inflation and unemployment using the Blue Chip Indicators forecast survey and find downward revisions of these forecasts are preceded by tightening policy announcements; Blue Chip solicits projections for key economic variables, including quarterly changes in the CPI and the civilian unemployment rate, from about 50 private forecasters. Nakamura and Steinsson (2018) study the same but for Blue Chip forecasts of real GDP, finding that tightening policy announcements are associated with significant upward revisions in forecasts. Yet, for households, this story is not entirely explored. In fact, the prevailing sentiment is that communication is largely ineffective due to low media coverage and poor economic/financial literacy (Binder, 2017a,b). Corroborating this narrative, Coibion et al. (2020c) use the Michigan Survey of Consumers to find almost no change in the percentage of respondents who say they heard about the news even after selected major policy announcements such as the launch of QE1, QE2, or even the announcement of the 2% inflation target by the FOMC in 2012. On the other hand, experimental evidence documents that forecasts, or beliefs of the future, change once they are provided with direct information about monetary policy. Coibion et al. (2022) use the Nielsen Online survey to provide participants with a direct summary of the FOMC meeting outcome and find that this treatment has a meaningful effect on inflation expectations. Coibion et al. (2020a) use a randomized control trial in the Nielsen Homescan panel to see if Forward Guidance is effective and find that while information treatments about current and one year ahead interest rates are significant, treatments beyond one-year have no effect on households' expectations of inflation, mortgage rates, and unemployment. This complements work about the existence of limited knowledge of interest rates and inflation and that communication from a central bank must take into account the heterogeneity of knowledge in its target audience (Rumler and Valderrama, 2019), but does little to clarify which narrative is true for announcements given in a vacuum.

Lewis et al. (2019) opt to use a high frequency approach through the daily Gallup Consumer Survey to assess FOMC policy news on household sentiment by using local projections. Of all the announcement shocks they subject their time series to between 2008 and 2017, they find that news covering the federal funds rate has significant negative effects on consumer confidence directly following FOMC meeting announcements; by contrast, they find no significant effect for surprises to announcements regarding forward guidance and asset purchases. Claus and Nguyen (2020) use a latent factor model relying on co-movements of elicited expectations on days the meeting announcements come out to identify unobserved news shocks driven by policy changes for Australian consumers; they find statistically significant reactions in economic conditions, unemployment, family finances and readiness to spend immediately following monetary policy shocks. I extend the local projections methodology from the former to include the US-based Survey of Consumer Expectations which specifically solicits expectations through an implied density forecast method, and use the shocks to different measures of monetary policy to analyze the medium-run dynamics of these expectations. To account for the unobserved news shocks driven by policy changes, I include a novel, unified measure of monetary policy news shocks proposed by Bu et al. (2021).

Lamla and Vinogradov (2019) use a random sample of the US population through a survey platform and find that monetary policy announcements have no effect on respondents' inflation and interest rate expectations, and have higher confidence in their grasp of what the true values of these variables are (read: less dispersion in their subjective uncertainty); the identification strategy rests on a tight, two-day window around FOMC meeting announcements for three years between 2015 and 2018. Fiore et al. (2021), who extend the observation window for household surveys to 42 days around an FOMC announcement, study expectation responses to different measures of monetary policy announcements but also isolate policy episodes such as the "Taper Tantrum" which occur too early on in my sample period (mid to late 2013) to be relevant; the study finds that only interest rate expectations are affected by FOMC meeting announcements. I extend this methodology of tight observation windows through the natural survey design and pair it with a larger sample of FOMC meeting announcements (a total of 68). I am also exploiting a different set of outcome variables solicited by the SCE, specifically interest rate expectations 12 months from survey date, inflation expectations 12 months and 24 - 36 months from survey date, expected change in national home prices 12 months and 24 - 36 months from survey date, household income 12 months from survey data, household spending 12 months from survey date, and expected change in various commodity prices 12 months from survey date. These time horizons are set by the survey and gauge how respondents feel about future economic conditions from a macroeconomic and personal standpoint. These expectations represent a number of key variables whose relationship to each other is predicted by topics such as the standard consumption Euler equation and the Phillips curve.

The second area of research relates to exploring the efficacy of unconventional monetary policies on the short-term nominal interest rate when faced with a constraint of an effective lower bound. The main focus in this area has largely been related to financial market participants and professional forecasters, with some research branching out to explore the effects on the macroeconomy (Campbell et al., 2012; Giannoni et al., 2015). Campbell et al. (2019) specifically highlights the role of imperfect information as a source that limits the efficacy of policies such as forward guidance, citing that the central bank's power to shape expectations at longer horizons falls dramatically with the introduction of a noisy environment. On the other hand, studies such as Swanson (2021); Bu et al. (2021) argue that unconventional policies have been effective as substitutes for conventional policy, citing effects to both the macroeconomy and more nuanced measures such as home-ownership. The former explores the effect of forward guidance and quantitative easing on macro variables, finding that their effect has been commensurate to conventional policies during the effective lower bound environment, while the latter uses these unconventional policies to predict home-ownership sentiment. I focus on the interaction of these unconventional policies with household data to see how the general public is reacting instead.

2.2 Data

My analysis for FOMC meeting announcement effects combines expectations data from the NY Fed's Survey of Consumer Expectations (SCE) with monthly monetary policy shocks.

2.2.1 Survey of Consumer Expectations (SCE)

The Survey of Consumer Expectations (SCE) began in June 2013 and is a nationally representative online survey from the United States comprising of a rotating group of approximately between 1,200 and 1,400 household heads.[1] The objective from the New York Fed is to solicit households for a variety of quantitative measures over a wide range of economic outcomes (e.g. inflation, income, household finance, and interest rates).[2] Armantier et al. (2017b) provide the specific components of expectations solicitation and the survey design; a main feature is that households are phased out of the survey once they reach 12 months of participation and new respondents are drawn each month from a stratified sampling procedure designed to maintain a demographically and socioeconomically representative sample of the population. Table 2.1 compares the socioeconomic and demographic distribution of the SCE respondents in my analysis with the U.S. Population numbers as per the U.S. Census Bureau from the latest 2022 numbers. Each respondent answers the various sections of the survey which includes an expectations module on macroeconomic and household level variables. Adoption of the SCE has been used in applications to tackle economic questions from political economy work in polarization (Armantier et al., 2017a), analyzing trends in consumer credit access (Armantier et al., 2018), estimating the elasticity of inter-temporal substitution (Crump et al., 2020), and in forming dynamic models of expectations formations (Fuster et al., 2018; Bellemare et al., 2020). The sample used at the time of this paper is up to December 2021.

Besides demographic information, the SCE elicits various measures of beliefs that are at the one-year (12 months ahead) and three-year (24 - 36 months ahead) horizons. The variables I focus on in these horizons include their: (i) one year ahead probability of a higher interest rate, (ii - iii) one and three year ahead expected inflation rate, (iv) one year ahead expected home price change, (v) three year ahead expected home price point prediction, (vi) one year ahead expected change in household income, (vii) one year ahead expected change in household spending, and (viii) one year ahead expected commodity price change point prediction that includes how much more/less households expect to pay for gallon of

[1] This phrasing of 'household head' is defined as being the main owner or renter of the household home.

[2] Reports and publications from the Federal Reserve Bank of New York using this data almost exclusively rely on using the survey-weighted interpolated mean of the series in question.

gas, food, and medical care. I group responses (i - v) as Macroeconomic Expectations and responses (vi - viii) as Personal Financial Expectations.

Macroeconomic Expectations variables are elicited in particular ways which I expand on below. The probability of interest rates increasing over the next 12 months from survey date is elicited as a point forecast, and this method is also used for the three year (24 - 36 months ahead) point forecast for changes to average home prices and commodity prices. A key innovation in this survey, which makes this a unique data set to analyze, is the robust solicitation of consumer implied density forecasts for inflation and home price expectations. As of this writing, no other major household survey solicits expectations in this manner. For the 12 months and 24 - 36 months ahead inflation expectation, as well as the 12 months ahead home price change expectation, the survey respondent is presented with a set of probability bin ranges that they can answer which must sum up to 100%. The respondent is asked to assign probabilities that the variable of interest will *increase 12% or more, increase between 8% - 12%, increase between 4% - 8%, increase between 2% - 4%, increase between 0% - 2%, decrease between 0% - 2%, decrease between 2% - 4%, decrease between 4% - 8%, decrease between 8% - 12%, decrease 12% or more*. This implied density forecast from respondents is a much more robust method of measuring beliefs than point forecasts and requires more thinking to answer; to decrease the chance of error, respondents cannot go on with the survey if the probabilities do not sum to 100%. Solicitation of the expectations in this way not only allows to capture a respondents' individual inflation density mean, but also their individual inflation uncertainty. To obtain the mean and subjective uncertainty, Armantier et al. (2017b) approach the task by fitting a generalized beta distribution to responses. The mean implied by these distributions is used as the measure of household inflation expectations and I follow this practice in my analysis for consistency with other existing studies using SCE data; further following convention when using household survey data, I also trim the data at the 1% and 99% percentiles. Table 2.2 contains a more thorough write up of the selected expectation variables in the SCE that my analysis will make use of.

A Note about Panel Conditioning in the SCE

Throughout the survey, the SCE references a framework of probability that is used to give respondents a basis for answering when assigning a **percent chance**. Specifically, respondents are given the following guideline near the beginning of the survey.[3]:

> **Q3Intro:** In some of the following questions, we will ask you to think about the **percent chance** of something happening the future. Your answers can range from 0 to 100, where 0 means there is absolutely no chance, and 100 means that it is absolutely certain.
>
> For example, numbers like:
>
> - 2 and 5 percent may indicate *almost no chance*
> - 18 percent or so may mean *not much chance*
> - 47 or 52 percent change may be a *pretty even chance*
> - 83 percent or so may mean a *very good chance*
> - 95 or 98 percent chance may be *almost certain*

Repeated application of this kind of survey guidance can affect how a respondent thinks about their answers, and thus the responses given over time; this is known as panel conditioning. Binder (2019) and Zhao (2022) find this occurring with inflation forecasts and uncertainty in the survey used in this analysis, with respondents in the early phase of participation in survey rounds predictably revising their inflation expectations downwards the longer their tenure in the survey, until they flatten out. This occurs regardless of what the true inflation dynamics are. These kinds of patterns raise questions to the validity of using survey responses and their properties without some sort of guidance as to when conditioning effects taper off, if ever.

Altig et al. (2020) tackle this issue by a non-parametric specification that allows an unrestricted relationship between survey results and the number of previous completions in their Survey of Business Uncertainty, while Fiore et al. (2021) and Zhao (2022) look for equality in probability distributions using

[3] Q3Intro copied verbatim as what a respondent would see from Armantier et al. (2017c, p. 5). The remaining questions of interest are found in the Appendix.

a Kolmogorov-Smirnov test. Without the uncertainty measures that the Survey of Business Uncertainty provides, I opt to follow the latter in order to test whether or not the empirical distributions of the survey tenure are statistically similar between cohorts who were surveyed before and those surveyed after an FOMC announcement meeting. This test helps me determine by how much to limit my sample so that panel conditioning due to longer survey tenure is not the reason trends in expectations unrelated to economic circumstances arise. Using the entire sample, I find that there is significant difference in the distribution of survey tenure between those surveyed ex-ante and those ex-post of the FOMC announcement meeting; the null hypothesis that the samples come from equal distributions is rejected with a p-value of 0.0024. Repeating this test on the range of survey tenures until I cannot reject the null (up through 12 months) leads me to limit the sample to respondents that have completed at least seven survey rounds. Then, the maximum number of observations that any household will have in the sample is if they do months seven through twelve, or six survey responses; the average tenure of my respondents throughout this time period is 9.16 months. Further, I only use surveys which have been completely filled out so that each respondent in my sample has given their expectation responses for all of the variables I include in my analysis; this strict condition brings the amount of total observations to 49,985.

Table 2.3 provides the descriptive statistics about each of the aforementioned expectations variables for the respondents who meet the criteria discussed. Generally, respondents expect higher interest rates in the next twelve months about one in three times, and they seem to overshoot the general price stability mandate by the Federal Reserve (2%). Additionally, they expect house prices to continue rising, in the short and medium term. They generally expect that their percent increase in household income will not exceed their increase in household spending (implying that their dollar will not outpace inflation), and they expect the price of a gallon of gas and food to increase; for a commodity like medical care, there is more uncertainty about whether it will increase or not. Commodities like these last three can be thought of as a proxy for the headline inflation rate (Binder, 2018), which the Federal Reserve has been increasingly relying on to define its official price stability target (FOMC, 2022). Traditionally, monetary policy has focused on targeting the core inflation rate but recent considerations note that a core inflation target may have adverse or no effects on policy communication with the general public (Powell, 2022).

2.2.2 Monetary Policy Announcements

I select measures for monetary policy announcement shocks that are available during the sample period (June 2013 - December 2021) guided by recent discussions in this space that try to combine two important considerations. First, limited knowledge and willingness from the general public to pay attention to the way the central bank makes announcements calls to question whether there should be a unified measure that is simple to interpret for anyone without the assumed training. Second, a unified measure may not be informative enough to understand by which channel monetary policy is being effective (or not). Policy has become much more multi-dimensional ever since the Global Financial Crisis and one single FOMC announcement or decision may have different effects to the yield curve. To mediate between the two considerations, I adopt a number of monetary policy shock measures that range from naive to multidimensional.

I start with a simple approach by using the change in the Effective Federal Funds Rate, a measure that is calculated and simplified to reflect how it stands at the end of each month. Acknowledging that the effective lower bound places a constraint on this measure, I also take into account the change of the Shadow Rate. Both of these measures are easily accessible and can be extended to fit my sample period from Wu and Xia (2016). Being able to measure the change in these rates allow the analysis to capture varying treatment intensities.

To incorporate the multidimensional aspect of monetary policy during the sample period, I extend the analysis by considering the three monetary policy factors calculated in Swanson (2021) who applies a factor model to 2-, 5-, and 10-year Treasury yields, along with assets with maturities below one year. For the time period between 1991 and 2019, Swanson estimates the top three factors which end up being able to explain 94% of the changes in interest rate responses in a 30 minute window around FOMC announcements. These factors, which deliver both a multi-dimensional view of monetary policy as well as a quantification of changes that financial market participants were not anticipating, correspond to the changes to the Federal Funds Rate (FFR), Forward Guidance (FG), and Large Scale Asset Purchases (LSAP), respectively. I follow the methodology and data sources in Bauer and Swanson (2022) to extend the three factors to the sample period for my analysis. By using these factors, then, I have a rough measure of the degree to which different policies implemented by the Federal Reserve affect respondent expectations.

The unified version of the multidimensional consideration would be to find a series that provides a single-factor, summary measure of the decisions taken, or not, during the FOMC policy announcements. For this, I also include the single-factor measures from Bu et al. (2021) which represents the combined effect of all the news that are outcomes on FOMC meeting days; this approach contains no significant information effect which is usually a confounding measure since FOMC meeting announcements contain both this and monetary policy shocks. This measure is also shown to represent an average effect of the changes to the Federal Funds Rate, Forward Guidance, and LSAPs following the FOMC meeting announcement.

I list the six monetary policy measures per FOMC Meeting days ranging from June 19, 2013 through December 15, 2021 in Table 2.4, and graph them through time in Figure 2.2. For readability, all the measures have been rounded to two decimal points in the table. For the analysis following in Section 2.3, all of the monetary policy measures have been standardized $\sim N(0,1)$ so that the magnitudes of the coefficients in Section 2.4 can be compared since the objective is to see to what degree monetary policy is being processed by the general public (if at all). Also included are the number of days between two adjacent FOMC announcement meetings, as well as the number of responses that the survey gathered in the three week window before and after the meeting.

2.3 Empirical Specification of FOMC Announcement Effects

I follow a standard identification method for announcement effects on expectations of survey respondents that has been used in the literature. Event studies, such as in Lewis et al. (2019), Lamla and Vinogradov (2019), Fiore et al. (2021) or Swanson (2021), make usage of a set time window around survey responses given at various frequencies before and after the policy announcement is given. Specifically, I split the window into two three-week cohorts occurring before and after an FOMC announcement meeting; this three-week time window maximizes the amount of observations the sample affords the analysis and I show the pattern of average responses for a subset of expectations in Figure 2.1. This window also is also guided by the days between FOMC announcements; the median amount of time between the 68 FOMC meeting announcements in my sample period is 42 days. If the meeting coincides with a day that survey responses are elicited, I leave them out of the analysis since I cannot determine the time that they were gathered relative to the meeting. Households are randomly assigned to three batches of when the survey module is sent out to them; each subsequent exposure to the survey is sent out to the households with the aim being that there are equal spacing between each round. This implicitly works so that no discernible pattern will exist when looking at which households are filling out the surveys in this three-week window.

Splitting up the observations into the two cohorts, those who give a response in the three-week window *before* the FOMC announcement meeting and those who give a response in the three-week window *after* the FOMC announcement meeting, I perform a baseline analysis on the following:

$$Y^e_{i,t} = \beta_t + \beta' \times C_{i,t} \times \mathbf{M_t} + \delta X_{i,t} + \epsilon_{i,t} \tag{2.1}$$

where $Y^e_{i,t}$ is the response given by individual i over their expectation of the future variable at month t (grouped by Macroeconomic Expectations and Personal Financial Expectations) that is acquired from the both cohorts c, β' is a $1 \times k$ row vector of the coefficients of interest, $\mathbf{M_t}$ is a $k \times 1$ column vector of the monetary policy measures, and $C_{i,t}$ is a dummy variable that takes on the value of 1 if a response is elicited from an individual in the cohort after the FOMC meeting announcement, and 0 otherwise. Included are also the β_t cohort specific constants, a $\delta X_{i,t}$ term that includes month and individual respondent fixed effects as well as various household controls such as age, levels of household income, levels of education, state, household size, numeracy level, and region, and the $\epsilon_{i,t}$ related error term; standard errors are clustered at the individual respondent level. I perform these regressions on the treatment of being in the different cohorts ("After FOMC") and interact this cohort with the various measures of naive to multidimensional monetary policy announcements: (i) a simple dummy variable indicating if the federal funds rate was increased (taking on a value of 1 if yes, 0 otherwise), (ii) the actual quantitative change in the federal funds rate, (iii) the change in the Shadow Rate (following the track by Wu and Xia (2016)),

(iv) the three policy factors calculated in Swanson (2021), and (v) the unified measure of monetary policy as per Bu et al. (2021).

A note about controlling for individual and cohort specific effects: given the tenure component of individuals (up to five months of responses), individual respondent fixed effects will control for time invariant factors that could impact the level of expectations elicited. On a larger scale, the cohort specific fixed effects will control for all information that is common amongst the cohort groups; this control supports the assumption that my analysis is centered around, which is that the only differentiator is the information treatment provided by the FOMC meeting announcement.

2.3.1 Local projections approach

While the baseline results focus on the survey responses by household participants in the immediate timing after FOMC meeting announcements, I would be remiss to disregard the effects of lagged reactions that are drawn from the literature surrounding information rigidities (Coibion and Gorodnichenko, 2015; Coibion et al., 2017). In that strand of literature, households have been found to often require time to process new information coming their way due to behavioral frictions which often take the form of delayed understanding or rational inattention. In this corroborative section, I extend my analysis to estimate the medium-term dynamic effects of policy announcements on my subset of household expectations following Lewis et al. (2019). Following the aggregation methods deployed by the New York Fed, I aggregate household expectations at a monthly frequency and estimate these policy announcement effects through local projections after Jordà (2005). For $0 \leq h \leq 12$ months,

$$Y^e_{t+h} = \beta'_h \times \mathbf{M}_t + \delta_h X_t + \varepsilon_{t+h} \tag{2.2}$$

where $Y^e_{h,t}$ is the aggregated response given by households in month t over their expectations (still grouped by Macroeconomic Expectations and Personal Financial Expectations), β'_h is a $1 \times k$ row vector of the coefficients of interest, and $\mathbf{M_t}$ is a $k \times 1$ column vector of the monetary policy measures surprises at month t. Included is also the $\delta_h X_t$ term that includes three months of policy announcement surprises and two lags of the expectations variable (chosen by the Akaike Information Criteria), the short-term and long-term interest rate, and a credit spread. Within month values for the control variables are not included in order to allow effects of the announcements on all control variables, and I construct the 90% confidence bands using Newey-West standard errors to control for heteroskedasticity and serial correlation.

I scale the impulses in this section to correspond to a 25 basis-point change in the reference rate for each of the six measures of monetary policy; this selection follows the literature but is entirely arbitrary. Monetary policy surprises of this type are rather rare and were most recently seen during the rate hikes of 2022 which is outside of the time range of this analysis. The advantage in scaling the responses to this fixed change in rates is that it will facilitate comparisons between the six measures. Lastly, I want to note that the stimulus provided by a 25 basis-point change in the different measures is not comparable across policy instruments as they would affect different parts of the economy.

2.4 Baseline Announcement Results

I break down the results from the baseline regression in Equation 2.1 in order of the expectations variables chosen. Each regression takes into account the treatment effect of it being solicited after the FOMC meeting announcement interacted with a dummy variable which tracks if the federal funds rate has increased or not month over month by denoting any increases with a 1 and 0 otherwise; then, ”After FOMC” can be taken to mean the treatment effect for whenever the month over month values of the federal funds rate decreases or stays the same. Any increase in the federal funds rate is taken as a sort of tightening of monetary policy as per the FOMC meeting announcement.

Macroeconomic Variables. Table 2.5 shows the results of the model for the one year ahead and three year ahead (from survey date) inflation expectations. In columns (2) through (4), a lack of a tightening announcement leads a negative revision of the one year ahead inflation expectation variable, on average, when taken together with a detected change in the shadow rate, the monetary policy surprise factors, and the unified measure, respectively. In column (2), this suggests that variation in the shadow rate is not directly affecting inflation expectations but a lack of tightening is leading households to revise their expectations downward; in this case, the average downward revision is 0.066, which represents a

0.21% decrease from the overall mean. This same story is told in column (3) with an average downward revision of 0.097, representing a 3% decrease from the overall mean. In column (4), however, the unified measure of monetary policy has a significant effect on inflation expectations with an upward revision of 0.113 when a one standard deviation above the mean unified shock of this nature occurs. This represents a 3.6% upward revision from the overall mean in one year ahead inflation expectations. This unified measure is described as being an average of the monetary policy surprises from column (3), but also isolates the monetary policy shocks and includes 'news'. Then, the significance of this effect can be thought of as being driven by an external measure not captured by the traditional monetary policy factors in the literature but rather a different kind of news on monetary policy. As for the three year ahead inflation expectation variables, all columns (5) through (8) have coefficients that are close to zero and not significant, which supports other previous findings that policy announcements generally have short lived effects on inflation expectations Fiore et al. (2021).

Table 2.6 shows the results of the model for the one year ahead and three year ahead (from survey date) home price expectations. While the method for eliciting both responses differs (the one year ahead one uses the subjective probability distribution approach detailed in Section 2.2), the answers for both are the same: where do respondents think average home prices nationwide will be at in the two time horizons? Columns (1) through (4) show the results for the one year ahead home price change and finds that positive changes in the federal funds rate, the federal funds rate factor, the large scale asset purchase factor, and the unified measure of monetary policy all decrease the expectation of home price changes. Specifically in column (1), for every one standard deviation above its mean, a positive month over month change in the federal funds rate decreases the one year ahead home price change expectation by -0.06 percentage points; the mean home price growth expectation is 4.22%, so this effect corresponds to an expected 1.42% decline. Similarly, in column (3), the effect of the one year ahead home price growth expectation corresponds to an expected 1.54% and 2.06% decline to the home price growth rate in response to a one standard deviation tightening of the federal funds rate factor and large scale asset purchase factor, respectively. Lastly, in column (4), this effect corresponds to a 1.47% decline to the home price growth rate in response to a one standard deviation tightening of the unified monetary policy measure. With a number of the monetary policy measures confirming a similar story, one year ahead expectations for home price growth arguably decline with tightening policy; this corroborates recent advances that suggest home ownership forces households to pay more attention to interest rates and monetary policy due to their investment (Ahn et al., 2022). As shown in columns (5) through (8), none of the monetary policy measures affect the three year ahead home price point prediction in a significant way.

Table 2.7 shows the results of the model for the one year ahead (from survey date) probability of an increase in interest rates. Since the response is elicited as a probability of an increase, it is not possible to take the results as a quantitative conjecture on the marginal effects of the monetary policy announcement on the household expectations as a level of this variable; instead, I can draw conclusions about the relative impact of different measures to conclude if any one monetary policy method is significantly affecting this expectation variable. Column (1) shows that the estimated probability of increases to the interest rates on savings accounts over the next 12 months of the survey date falls by 0.98 percentage points, on average, when no tightening announcement is given. The average probability of interest rates for savings accounts rising in the next twelve months for the entire range is 31.44, so this effect represents a 3.12% decrease in the probability of rising interest rates for this time horizon. When an increase of the federal funds rate occurs, the significance of the effect vanishes (albeit positive at about 2.29). Changes in the shadow rate as proposed by Wu and Xia (2016) also yield a similar story with no effect detected both in the no tightening announcement dummy and the interaction with the quantitative change in the shadow rate in column (2). Moving forward with the high-frequency financial market monetary policy announcement factors proposed by Swanson (2021), I break up the dimensions in column (3) and find that only the federal funds rate factor robustly affects the probability of a higher interest rate in twelve months. Specifically, a federal funds rate factor one above its mean by one standard deviation significantly increases expectations of a higher interest rate on savings accounts by 0.62 percentage points, on average. Looking again at the average probability of interest rates for savings accounts rising in the next twelve months for the entire sample, this effect represents a 1.98% increase in the probability of rising interest rates for this time horizon. Variation in the other two factors concerning with forward guidance and large scale asset purchases do not affect interest rate expectations significantly. Lastly, I present column (4) with the unified factor from Bu et al. (2021) which is constructed as a measure that averages the aforementioned monetary policy factors along with all the relevant news given on that day (and, as the authors state, ”has no significant information shock effect”). Conditional on being treated by the FOMC meeting announcement, the probability of interest rates for savings accounts rising in the next twelve months rises by 0.427 percentage points for a one standard deviation unified monetary policy shock above its mean. Compared to the mean of the entire range, this represents a 1.36% increase in the probability of rising interest rates for this time horizon. In all, the lack of a tightening announcement has a larger relative impact on respondents expecting an increase in the probability of higher interest rates on savings accounts than any of the other policy measures.

Together, the results from the aforementioned paragraphs suggest that expectations about the interest rate on savings accounts, one year ahead inflation expectations, and one year ahead home price growth expectations all are affected by the various measures of monetary policy changes. Expectations about interest rates are affected by policy tightening/easing on the federal funds rate and its related high frequency monetary policy federal funds rate factor as per Swanson (2021), as well as by the unified measure of monetary policy that takes into account other relevant news passed along the day of the FOMC meeting announcement as per Bu et al. (2021). One year ahead inflation expectations are only affected by this same unified measure which captures news on monetary policy but does not disentangle its source. One year ahead home price growth expectations are affected by policy tightening/easing on the federal funds rate and its related federal funds rate factor, the large scale asset purchase factor, and the unified measure.

Personal Financial Variables. Table 2.8 shows the results of the model for the one year ahead (from survey date) percent change in household income and spending. In other words, how much are these measures of forecasted income and spending being affected by the monetary policy announcements? As evidenced from columns (1) through (8), the answer is hardly. All coefficients are generally close to zero and not significant. This corroborates a narrative that has been found in the literature, namely, that the central bank, by focusing on core inflation, leaves out spending variables that have larger weights for households as in D’Acunto et al. (2019). These results support the claim that households do not connect changes in measures of monetary policy to their daily spending and income. Additionally, this empirical finding stands in contrast to the laboratory experiments that found significant effects on announcements for household employment and consumption expectations in a noisy environment (Kryvtsov and Petersen, 2021; Coibion et al., 2022); even with the tight window around the FOMC meeting announcement, news is not as easily interpreted outside of these lab settings.

To further break down the claim in the preceding paragraph, Table 2.9 shows the results of the model for the one year ahead (from survey date) percent change in three household relevant commodity prices: a gallon of gas, food, and medical care. Columns (1) through (4), for a gallon of gas, are all close to zero and insignificant coefficients to either tightening or easing monetary policy. The same goes for columns (5) through (8) for food, and columns (9) through (12) for medical care. Regardless of the commodity chosen in this subset, the monetary policy measures are not affecting household expectations. Households, on average, do not connect how different kinds of policy announcements can filter through to their daily purchases.

2.4.1 Baseline Local Projections Approach

Figures 2.3 through 2.8 show the local projections for each of the aggregated monthly expectations variables to the measures of monetary policy through the method introduced in section 2.3.1. The confidence bands are following Lewis et al. (2019) and correspond to 90%, using Newey-West standard errors to control for heteroskedasticity and serial correlation. The responses are scaled such that impacts correspond to a 25 basis-point change.

Figures 2.3 and 2.4 show the response of the various expectations to the changes in the federal funds rate and shadow rate, respectively, as calculated by Wu and Xia (2016). A positive surprise to both leads to a significant effect on income; the change in the federal funds rate affects after income expectations positively after several months and peaking at 7, while the change in shadow rate decreases income expectations slightly. Neither effect is large relative to the average expectation of increasing income one year ahead which sits at 3.89% for the whole sample. Shadow rate surprises also decrease inflation expectations for the 24- and 36-months ahead horizon but only become significant after a couple of months, peaking at around the 8 month horizon (again, the relative effect is small). Combined with the fact that changes corresponding to 25 basis-points are equal to about 4x and 1.5x the standard deviations for the delta federal funds rate and delta shadow rate shocks, these effects are marginally significant at best. Curiously, the change in federal funds on expectations on a gallon of gas are significant, peaking around 4 months at 0.5 which represent a 8.1% increase in the one year ahead expectations for the price of a gallon of gas. This, however, decreases by almost the same amount 9 months after.

The Swanson (2021) monetary policy shock factors are shown in figures 2.5 through 2.7. The first finding, in Figure 2.5, is that there is a counter-intuitive permanent (to the month range) response in the probability of higher interest rates one year ahead to a positive surprise in the federal funds rate factor. Unlike the immediate and significant response that increased the probability in rising interest rates along with an increase in this factor as shown in Table 2.7, the local projection shows a negative response that appears 4 months after the announcement and stays significant through the 12 month range, peaking at 7 months valued at -0.75. Given that the probability of increasing interest rates is 31.44, this represents a 2.4% decrease. On the flip side, the surprise also leads to a delayed increase in other expectations measures including one year ahead inflation, 24- to 36-months ahead inflation, one year ahead income and spending, and one year ahead price growth of food. The other factors, shown in figures 2.6 and 2.7, are generally small and have no significant effect.

Lastly, the Bu et al. (2021) are shown in Figure 2.8. Here, too, the response in the probability of higher interest rates one year ahead to a positive surprise in the unified measure is negative, delayed but significant from months 6 through 10. It peaks at around month 8 with a value of -0.10, representing a 2.6% decrease from the average throughout the whole sample. The other expectations variables responses are generally small and have only marginally significant effects.

Overall, the evidence presented in this section is broadly in line with the results from the previous section except that I decompose the timing that these effects take place for these variables. The change in federal funds rate, federal funds rate factor, and the unified measure again come out as significant, the others have no or only smaller (and mostly delayed) effects. Expectations for interest rates, inflation in the short run (one year ahead), and income all have varying degrees of a significant response to these various policies.

2.5 Discussion on the Role of Media

As mentioned in section 2.1, the literature in this space has focused around doing laboratory experiments or event study approaches that take on the form similar to this analysis for various expectation solicitations. the former has researchers providing information to survey respondents that allow the analysis to estimate the effect of that isolated information treatment. In event study frameworks, like this, there is no way to account for the noise that enters a households' information set when developing responses. I can control for things like panel conditioning, as in section 2.2.1, but I do not know nor can control what a household sees. It would be difficult to argue that the FOMC meeting announcement is the only source of information a household is subjected to, whether it by their press conferences or website information. Instead, in this section, I want to briefly explore the possibility of news reaching households using traditional media to see if the types of coverage announcements receive affect the expectations solicited.[4] An in depth analysis exploring these kinds of media channels would deviate the central question of this analysis and so instead I opt to use a simple measure of news coverage by using Google Trends data. Google Trends measures and analyzes the popularity of different searches done on the internet across time and, by searching relevant keywords, can provide a simple measure of how popular different policy announcements are and when. While simple, this can be seen as a proxy for a more robust media transmission channel.

I show the search interest for different keyboards related to the FOMC and the Federal Reserve's policies during the time range between June 2013 and December 2021 in Figure 2.9. Google Trends normalizes searches to a relative intensity of 100, to which other searches are scaled to. For example, in sub-figure (b) when it shows "Quantitative Easing", we see the largest search interest for this term following the remarks made by then Federal Reserve Chair Ben Bernanke in the later months of 2013

[4] For some introductory work in this space, we can look at Bianchi et al. (2019) who look at the role of Twitter in affecting central bank independence and Lüdering and Tillmann (2020) who do a textual analysis of monetary policy news coverage for asset prices

about scaling back its bond purchasing program. Similarly, for inflation in sub-figure (a), we see the largest interest for inflation near the tail end of 2021 when the onset of COVID-19 began to circulate fears.

To measure the effects of different types of policy announcements on search interest, I opt to regress these interest time series on the absolute values of the monetary policy announcements such that it allows me to take into account their magnitude but not their direction. Generated interest is the measure from Google Trends, not if the coverage is good or bad. I assume that the household is subjected to all of these measures at the same time as opposed to the tiered approach I used in the event study and local projection analysis.

Table 2.10 shows the simple regression results from the aforementioned analysis. Column (1) indicates that a change in the Federal Funds Rate announcement is significantly related with higher search intensity for the keyword of "FOMC"; this higher intensity can translate to higher public interest in understanding what the FOMC announcement means for the future. Column (2) shows that a similar story happens here for the search "Federal Funds Rate", with changes in the Federal Funds Rate, the Forward Guidance factor as per Swanson (2021), and the unified monetary policy measure by Bu et al. (2021) all significantly affecting the search intensity. One curious direction comes from the Forward Guidance Factor which has a negative impact on search intensity for the "Federal Funds Rate" term. Given that Forward Guidance as a policy tool is not easily understood by the general public, this suggests that a confusing policy detracts from public interest in the relevant monetary policy target.

Overall, the results confirm that announcements regarding changes in the federal funds rate as well as the unified measure are more likely to reach the general public and therefore households that are surveyed, leading to more general interest than any of the other policy measures.

2.6 Conclusion

In this paper, I analyze the effect of various measures of monetary policy via FOMC meeting announcements on household expectations. Current research on exploring household expectations often contradicts each other, with one school of thought completely refuting the efficacy of monetary policy on affecting expectations and the other suggesting there are different channels by which it does. Regardless of this contradiction, considering household expectations as a means to gauge the efficacy of central bank communication is timely not only from the policy perspective (as per the mandates by the most recent Federal Reserve Chairs, Janet Yellen and Jerome Powell) but also because they matter for economic activity. For instance, households often take part of wage bargaining processes that imply they take income, spending, and savings decisions that are influenced heavily by their expectations about future conditions.

This analysis makes usage of the Survey of Consumer Expectations and exploits a timing window surrounding an FOMC meeting announcement. By comparing the responses in this window before and after the announcement, I find that monetary policy announcements robustly affect household expectations of the future for interest rates, one year ahead inflation, and one year ahead home price growth. These effects are found by using a range of monetary policy measures going from simple to multidimensional. This result stands in contrast to the aforementioned literature in experimental settings; noise matters and information is lost when households are being subjected to a variety of different news. A quick analysis of Google Trends search intensity also yields a similar story.

Additionally, I explore the timing of the announcement effects by using a local projections approach and again find that there are significant responses to interest rate and one year ahead inflation expectation. In this analysis, one year ahead income expectations also show a degree of significant responses to monetary policy announcements. The key takeaway from both the event study and local projections approach is that not all these measures are understood equally. Changes in the federal funds rate, the federal funds rate factor, and the unified measure of policy that contains no additional information effect all continuously come out as significant for a number of household expectations. However, no wider range of expectations for commodity prices, spending levels, or house growth in the long run are affected; households are not connecting how policies may affect their economic circumstances in the future as strongly as the literature using other event studies may suggest.

These findings contribute to the discussion about the efficacy of central bank communication with the general public, particularly highlighting that there are communication challenges that exist for the central bank moving forward. While measures taken as unconventional have been heavily relied on for the better part of the last 15 years, there is scant evidence that these policies are understood by the general public and do not generate the kind of interest in the policy targets the central bank aims for. As such, are there channels by which expectations are generated more routinely for households? And, if so, are expectations changing in a systematic way that reflects how the central bank aims to conduct policy over time? I leave these questions for further research.

Table 2.1. SCE Respondent Characteristics vs US Population

	SCE	Population		SCE	Population
Age			**Race and Ethnicity**		
Under 40	27.3%	37.3%	White (Non-Hispanic)	70.6%	60.1%
40 - 60	37.8%	29.3%	Black or African American (Non-Hispanic)	10.9%	12.5%
Over 60	34.9%	33.4%	Hispanic/Latino/Spanish Origin	12.7%	18.5%
			Asian or Other	5.8%	8.9%
Gender			**Household Income**		
Male	52.5%	50.8%	Under $50K	35.8%	37.8%
Female	47.5%	49.2%	$50K to $100K	36.5%	28.6%
			Over $100K	27.7%	33.6%
Region			**Education**		
Midwest	23.9%	20.7%	High School or Less	11.7%	37.9%
Northeast	21.8%	17.3%	Some College	31.7%	27.1%
South	32.7%	38.3%	College or More	56.6%	35.0%
West	21.6%	23.7%			

Notes: SCE column is representing the subset of the survey respondents as per those with a tenure of 7 months or more; Population column represents the values for the U.S. population as obtained from the U.S. Census Bureau. Total number of individual respondents is 10,741 over the time range June 2013 through December 2021. Besides a disparity in the education category, the sampling done by the New York Fed, once filtered to account for the correct type of tenure and provided that the respondent answers all parts of the survey, provides a distribution that is relatively similar to the U.S. population.

Table 2.2. Selected Expectation Variables, Abbreviations and Ranges in the SCE

Catalog	(Short) Question Text	Variable	Range
[Macroeconomic Expectations]			
Q5new	What do you think is the percent chance that **12 months from now** the interest rate in the U.S. will be higher than it is now?	One Year Ahead Probability of Higher Interest Rate	0-100%
Q9	Now we would like you to think about the different things that may happen to inflation **over the next 12 months**... In your view, what would you say is the percent chance that, **over the next 12 months...**	One Year Ahead Expected Inflation Rate	$\mathbb{R}$, sum to 100%
Q9c	And in your view, what would you say is the percent change that, **over the 24 and 36 months from survey date], ...**	Three Year Ahead Expected Inflation Rate	$\mathbb{R}$, sum to 100%
C1	And in your view, what would you say is the percent chance that, **over the next 12 months**, the average home price nationwide will...	One Year Ahead Expected Home Price Change	$\mathbb{R}$, sum to 100%
C2part2	**Over the 12-month period between 24 and 36 months from survey date**, I expect the average home price to [increase/decrease] by _%	Three Year Ahead Expected Home Price Point Prediction	$\mathbb{R}$
[Personal and Financial Expectations]			
Q25v2part2	**Over the next 12 months**, I expect my total household income to [increase/decrease] by _%	One Year Ahead Expected Change in Household Income	$\mathbb{R}$
Q26v2part2	**Over the next 12 months**, I expect my total household spending to [increase/decrease] by _%	One Year Ahead Expected Change in Household Spending	$\mathbb{R}$
C4Info	**Twelve months from now**, what do you think will have happened to the price of the following items: (i) gallon of gas, (ii) food, (iii) medical care	One Year Ahead Expected Commodity Price Change Point Prediction	$\mathbb{R}$

Notes: Expectations solicited from the SCE, time range from June 2013 through December 2021. Questions Q5new, Q25v2part2, Q26v2part2, C2part2, and C4Info are all point predictions as respondents are asked by how much the average variable will change over the specified time period and they give a single-value forecast. Questions Q9, Q9c, and C1 are presented alongside probability bins that ask about the percent chance that the variable will increase/decrease by either 12% or more; by 8% to 12%; by 4% to 8%; by 2% to 4%; by 0% to 2%. A generalized beta distribution is fitted to the responses of each participant, and then the mean of this distribution is calculated to obtain the expectation.

Table 2.3. SCE Descriptive Statistics

	Panel	Mean	S.D.	Min	Max
Macroeconomic Expectations					
Interest Rate 12mo Ahead	Overall	31.44	(25.72)	0.00	100.00
	Between		(22.03)	0.00	100.00
	Within		(13.72)	-75.13	83.33
Inflation Rate 12mo Ahead	Overall	3.84	(4.90)	-25.20	36.30
	Between		(4.25)	-25.20	27.60
	Within		(2.71)	-38.50	37.00
Inflation Rate 24-36mo Ahead	Overall	3.72	(4.97)	-27.00	36.30
	Between		(4.30)	-25.20	26.70
	Within		(2.80)	-39.20	41.70
Home Price Change 12mo Ahead	Overall	4.22	(5.62)	-25.20	36.30
	Between		(4.78)	-25.20	27.80
	Within		(3.31)	-41.70	40.00
Home Price Point Change 24-36mo Ahead	Overall	5.09	(5.87)	-10.00	20.00
	Between		(4.68)	-10.00	20.00
	Within		(3.45)	-15.65	24.97
Personal Financial Expectations					
Household Income 12mo Ahead	Overall	3.89	(5.74)	-20.00	35.00
	Between		(4.98)	-20.00	30.00
	Within		(4.17)	-24.79	27.98
Household Spending 12mo Ahead	Overall	4.01	(6.33)	-20.00	25.00
	Between		(5.67)	-20.00	25.00
	Within		(4.31)	-22.45	29.65
Commodity Price Change 12mo Ahead					
Gallon of Gas	Overall	6.18	(6.51)	-5.00	25.00
	Between		(5.97)	-5.00	25
	Within		(5.46)	-2.12	5.09
Food	Overall	5.65	(4.24)	-1.00	20.00
	Between		(4.09)	1.00	20.00
	Within		(3.95)	1.00	15.00
Medical Care	Overall	11.03	(9.85)	-30.00	40.00
	Between		(10.73)	-10.00	40.00
	Within		(7.89)	-10.00	23.00
Total Observations					49,985
Number of Unique Respondents					10,741
Average Tenure (Month) of Respondents					9.16

Notes: Descriptive statistics of the panel data obtained from the SCE. Respondents are included if they have completed 7 months or more of the survey, and are phased out by the survey design to complete no more than 12 months.

Table 2.4. Monetary Policy Measures per FOMC Meeting

		Wu and Xia (2016)		Swanson (2021) Factors			Bu et al. (2021)		
FOMC Meeting	Days Since Previous	Δ FFR	Δ Shadow Rate	FFR	FG	LSAP	UMPS	Before	After
Jun 19, 2013	-	-0.02	0.30	0.16	1.28	1.96	0.05	154	142
Jul 31, 2013	42	0.02	-0.55	0.09	0.08	-0.23	0.02	178	302
Sep 18, 2013	49	-0.01	-0.14	0.08	-1.34	-2.55	-0.05	314	244
Oct 30, 2013	42	0.01	-0.05	0.10	0.08	0.33	-0.01	259	437
Dec 18, 2013	49	0.00	-0.13	0.21	0.02	0.63	-0.01	311	189
Jan 29, 2014	42	0.00	-0.24	0.22	-0.04	-0.24	0.01	183	393
Mar 19, 2014	49	0.00	-0.08	0.06	1.04	0.57	0.10	315	219
Apr 30, 2014	42	0.03	-0.27	0.15	0.12	0.04	-0.01	341	356
Jun 18, 2014	49	0.01	0.10	0.09	0.41	-0.16	-0.02	323	319
Jul 30, 2014	42	-0.01	0.05	0.15	-0.09	-0.23	-0.03	457	483
Sep 17, 2014	49	0.00	0.09	0.07	0.75	0.16	0.01	423	316
Oct 29, 2014	42	0.00	0.00	0.09	0.88	-0.01	0.07	418	417
Dec 17, 2014	49	-0.02	0.35	0.29	-1.54	0.50	0.02	317	334
Jan 28, 2015	42	0.00	0.15	0.16	-0.14	-0.14	0.02	396	420
Mar 18, 2015	49	0.00	0.17	0.19	-2.42	-0.77	-0.04	310	339
Apr 29, 2015	42	0.02	0.21	0.20	0.31	0.87	-0.04	404	387
Jun 17, 2015	49	0.00	0.03	0.09	-0.65	0.14	-0.07	342	353
Jul 29, 2015	42	0.00	0.11	0.06	0.48	0.20	0.00	370	335
Sep 17, 2015	50	-0.01	0.18	-0.53	-1.53	-0.64	-0.04	323	315
Oct 28, 2015	41	0.00	0.21	0.11	1.80	-0.05	0.06	390	351
Dec 16, 2015	49	0.12	0.26	0.31	-0.02	-0.54	0.02	353	273
Jan 27, 2016	42	0.09	0.14	0.01	-0.46	-0.06	-0.02	392	361
Mar 16, 2016	49	-0.04	-0.02	-0.11	-1.81	0.04	-0.07	333	275
Apr 27, 2016	42	0.05	-0.10	0.10	0.33	-0.25	0.00	370	364
Jun 15, 2016	49	0.01	-0.07	0.04	-0.78	0.19	-0.03	336	319
Jul 27, 2016	42	0.00	0.05	0.09	0.16	-0.32	0.01	368	326
Sep 21, 2016	56	-0.01	0.06	-0.39	-0.18	-0.47	0.02	333	343
Nov 2, 2016	42	0.00	-0.09	0.12	0.18	-0.05	-0.01	340	366
Dec 14, 2016	42	0.24	-0.01	0.03	1.39	0.24	0.08	327	376
Feb 1, 2017	49	0.01	0.01	0.13	-0.38	0.13	0.00	401	444
Mar 15, 2017	42	0.25	0.24	0.25	-1.31	0.03	-0.02	401	376
May 3, 2017	49	0.00	0.18	0.19	0.40	0.00	0.03	322	364
Jun 14, 2017	42	0.23	0.03	0.32	0.35	0.01	0.03	330	355
Jul 26, 2017	42	0.01	0.02	0.10	-0.21	-0.21	-0.03	404	339
Sep 20, 2017	56	-0.01	0.00	0.05	1.17	-0.12	0.03	370	332
Nov 1, 2017	42	0.00	0.13	0.14	0.14	0.02	0.03	373	354
Dec 13, 2017	42	0.26	0.12	0.20	-0.21	-0.17	-0.01	378	296
Jan 31, 2018	49	0.01	-0.01	0.18	0.25	0.16	0.04	376	445
Mar 21, 2018	49	0.32	0.10	0.12	0.11	0.37	-0.01	380	367
May 2, 2018	42	0.01	0.06	0.16	-0.19	-0.10	-0.02	302	383
Jun 13, 2018	42	0.21	0.12	0.02	0.84	0.10	0.01	370	349
Aug 1, 2018	49	0.00	0.09	0.19	-0.05	-0.06	-0.01	332	402
Sep 26, 2018	56	0.27	0.15	0.31	-0.19	0.04	0.01	343	386
Nov 8, 2018	43	0.00	0.10	0.13	0.27	-0.06	0.02	330	379
Dec 19, 2018	41	0.20	0.13	0.50	-0.04	-0.48	0.04	323	357
Jan 30, 2019	42	0.00	-0.04	0.13	-0.67	0.08	-0.05	383	424
Mar 20, 2019	49	0.03	-0.05	0.36	-1.22	-0.18	-0.02	408	310
May 1, 2019	42	-0.05	0.01	-0.02	-0.69	0.06	0.04	391	420
Jun 19, 2019	49	0.00	-0.23	0.48	-2.02	0.71	-0.05	422	353
Jul 31, 2019	42	0.00	-0.01	0.15	0.14	-0.07	0.06	358	345
Sep 18, 2019	49	-0.23	-0.06	0.12	-0.13	-0.03	0.04	393	334
Oct 30, 2019	42	-0.32	-0.30	0.12	-0.19	-0.12	0.02	418	355
Dec 11, 2019	42	-0.01	-0.03	-0.12	0.20	0.10	0.00	356	369
Jan 29, 2020	49	0.04	0.02	-0.10	-0.13	0.08	0.00	391	422
Apr 29, 2020	91	-0.03	-0.19	0.12	-0.10	-0.03	-0.04	403	426
Jun 10, 2020	42	0.03	-0.08	0.12	-0.22	-0.06	0.00	279	392
Jul 29, 2020	49	0.02	-0.15	0.13	-0.04	-0.10	-0.03	387	407
Sep 16, 2020	49	0.00	-0.18	-0.15	0.09	0.06	-0.01	248	358
Nov 5, 2020	50	0.00	-0.40	0.12	0.17	-0.08	0.02	319	350
Dec 16, 2020	41	0.00	-0.06	-0.12	0.00	0.11	-0.01	307	312
Jan 27, 2021	42	-0.02	-0.13	0.12	-0.16	-0.11	0.01	381	374
Mar 17, 2021	49	-0.01	-1.08	0.13	0.17	-0.05	-0.06	291	338
Apr 28, 2021	42	-0.01	-0.24	0.12	-0.13	0.04	-0.02	353	360
Jun 16, 2021	49	0.03	0.17	0.10	-0.53	-0.07	0.06	328	309
Jul 28, 2021	42	-0.01	-0.06	-0.14	0.07	0.02	-0.01	315	334
Sep 22, 2021	56	0.00	-0.01	0.14	0.31	-0.09	0.03	283	334
Nov 3, 2021	42	0.00	-0.15	-0.14	0.00	0.11	0.00	299	347
Dec 15, 2021	42	0.00	0.69	0.21	-0.10	-0.13	-0.02	321	318
Shock Mean		0.02	-0.01	0.07	-0.05	-0.01	0.00		
Shock S.D.		0.06	0.15	0.13	0.56	0.32	0.03		

Notes: FOMC Meeting indicates the day of an FOMC conference as categorized by the Federal Reserve Board website. All estimates have been rounded to two decimal points in this table for ease of readability. The change in the Effective Federal Funds Rate and the change in the Shadow Rate are direct extensions of the data from Wu and Xia (2016) which has been discontinued after February 2022 but does not affect the analysis. The three factors from Swanson (2021) include FFR (Federal Funds Rate), FG (Forward Guidance), and LSAP (Large Scale Asset Purchases); after June 19, 2019, the factors have been extended from their original estimation by using the methodology and data sources from Bauer and Swanson (2022). Bu et al. (2021) provide the Unified Monetary Policy Shock (UMPS) measure that contains no significant information effect.

Table 2.5. [Results] Macroeconomic Expectations: Inflation (12m and 24-36m)

	12m Ahead, Expected Inflation Rate				24-36m Ahead, Expected Inflation Rate			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
After FOMC	-0.048	-0.066*	-0.097*	-0.084**	-0.077	-0.083	-0.109	-0.034
	(0.038)	(0.039)	(0.055)	(0.040)	(0.082)	(0.041)	(0.054)	(0.039)
After FOMC $\times \Delta$ FFR	0.080				0.011			
	(0.105)				(0.099)			
After FOMC $\times \Delta$ Shadow Rate		-0.013				-0.020		
		(0.063)				(0.061)		
After FOMC $\times$ Federal Funds Rate Factor			0.002				-0.018	
			(0.062)				(0.055)	
After FOMC $\times$ Forward Guidance Factor			-0.042				-0.095	
			(0.040)				(0.069)	
After FOMC $\times$ Large Scale Asset Purchases Factor			0.023				-0.007	
			(0.087)				(0.012)	
After FOMC $\times$ Unified Factor				0.113*				-0.090
				(0.061)				(0.055)
Individual and FOMC Fixed Effects	✓	✓	✓	✓	✓	✓	✓	✓
Total Observations	49,985	49,985	49,985	49,985	49,985	49,985	49,985	49,985
Number of Respondents	10,741	10,741	10,741	10,741	10,741	10,741	10,741	10,741
R^2	0.611	0.611	0.611	0.611	0.593	0.593	0.593	0.593

Notes: Fixed effects regressions. Columns (1), (2), (5), and (6) use the monetary policy measures for the Federal Funds Rate and Shadow Rate as measured in Wu and Xia (2016). Columns (3) and (7) use the monetary policy factors as measured in Swanson (2021). Columns (4) and (8) use the unified policy shock as measured in Bu et al. (2021). Standard errors are in parenthesis and are clustered at the respondent level.

* p < 0.10, ** p < 0.05, *** p < 0.01

Table 2.6. [Results] Macroeconomic Expectations: Home Prices (12m Change and 24-36m Point Prediction)

	12mo Ahead, Price Change				24-36m Ahead,Point Prediction			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
After FOMC	-0.137	-0.064**	-0.195**	-0.051**	-0.012	-0.067	-0.045	-0.055
	(0.095)	(0.024)	(0.089)	(0.040)	(0.098)	(0.043)	(0.031)	(0.037)
After FOMC $\times \Delta$ FFR	-0.060*				-0.109			
	(0.036)				(0.087)			
After FOMC $\times \Delta$ Shadow Rate		-0.044				-0.028		
		(0.058)				(0.049)		
After FOMC $\times$ Federal Funds Rate Factor			-0.065*				-0.043	
			(0.035)				(0.027)	
After FOMC $\times$ Forward Guidance Factor			-0.043				-0.035	
			(0.036)				(0.047)	
After FOMC $\times$ Large Scale Asset Purchases Factor			-0.087*				0.055	
			(0.049)				(0.041)	
After FOMC $\times$ Unified Factor				-0.062**				-0.052
				(0.027)				(0.046)
Individual and FOMC Fixed Effects	✓	✓	✓	✓	✓	✓	✓	✓
Total Observations	49,985	49,985	49,985	49,985	49,985	49,985	49,985	49,985
Number of Respondents	10,741	10,741	10,741	10,741	10,741	10,741	10,741	10,741
R^2	0.581	0.581	0.581	0.581	0.509	0.509	0.509	0.509

Notes: Fixed effects regressions. Columns (1), (2), (5), and (6) use the monetary policy measures for the Federal Funds Rate and Shadow Rate as measured in Wu and Xia (2016). Columns (3) and (7) use the monetary policy factors as measured in Swanson (2021). Columns (4) and (8) use the unified policy shock as measured in Bu et al. (2021). Standard errors are in parenthesis and are clustered at the respondent level.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.7. [Results] Macroeconomic Expectations: Interest Rate, 12m

	12mo Ahead, Higher Interest Rate			
	(1)	(2)	(3)	(4)
After FOMC	-0.975**	0.051	0.073	0.053
	(0.379)	(0.120)	(0.120)	(0.120)
After FOMC $\times \Delta$ FFR	2.286			
	(0.741)			
After FOMC $\times \Delta$ Shadow Rate		0.406		
		(0.156)		
After FOMC $\times$ Federal Funds Rate Factor			0.623***	
			(0.211)	
After FOMC $\times$ Forward Guidance Factor			0.199	
			(0.194)	
After FOMC $\times$ Large Scale Asset Purchases Factor			0.287	
			(0.175)	
After FOMC $\times$ Unified Factor				0.427*
				(0.194)
Individual and FOMC Fixed Effects	✓	✓	✓	✓
Total Observations	49,985	49,985	49,985	49,985
Number of Respondents	10, 741	10, 741	10, 741	10, 741
R^2	0.643	0.643	0.643	0.643

Notes: Fixed effects regressions. Columns (1) and (2) use the monetary policy measures for the Federal Funds Rate and Shadow Rate as measured in Wu and Xia (2016). Column (3) uses the monetary policy factors as measured in Swanson (2021). Column (4) uses the unified policy shock as measured in Bu et al. (2021). Standard errors are in parenthesis and are clustered at the respondent level.

* p < 0.10, ** p < 0.05, *** p < 0.01

Table 2.8. [Results] Personal Financial Expectations: 12mo Ahead Percent Change in Household Income and Spending

	12mo Ahead, Income				12mo Ahead, Spending			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
After FOMC	0.024	-0.049	-0.053	-0.065	-0.108	0.087	0.071	0.046
	(0.097)	(0.056)	(0.043)	(0.059)	(0.149)	(0.056)	(0.049)	(0.031)
After FOMC $\times \Delta$ FFR	-0.304				0.066			
	(0.056)				(0.115)			
After FOMC $\times \Delta$ Shadow Rate		0.025				-0.005		
		(0.041)				(0.047)		
After FOMC $\times$ Federal Funds Rate Factor			-0.025				0.066	
			(0.044)				(0.041)	
After FOMC $\times$ Forward Guidance Factor			0.023				0.024	
			(0.067)				(0.039)	
After FOMC $\times$ Large Scale Asset Purchases Factor			-0.057				0.165	
			(0.039)				(0.099)	
After FOMC $\times$ Unified Factor				-0.005				0.069
				(0.051)				(0.077)
Individual and FOMC Fixed Effects	✓	✓	✓	✓	✓	✓	✓	✓
Total Observations	49,985	49,985	49,985	49,985	49,985	49,985	49,985	49,985
Number of Respondents	10,741	10,741	10,741	10,741	10,741	10,741	10,741	10,741
R^2	0.698	0.698	0.698	0.698	0.603	0.603	0.603	0.603

Notes: Fixed effects regressions. Columns (1), (2), (5), and (6) use the monetary policy measures for the Federal Funds Rate and Shadow Rate as measured in Wu and Xia (2016). Columns (3) and (7) use the monetary policy factors as measured in Swanson (2021). Columns (4) and (8) use the unified policy shock as measured in Bu et al. (2021). Standard errors are in parenthesis and are clustered at the respondent level.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 2.9. [Results] Personal Financial Expectations: 12mo Ahead Percent Change Commodity Prices

	Gallon of Gas				Food				Medical Care			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)
After FOMC	-0.007	0.002	0.008	-0.009	-0.014	0.034	0.061	-0.008	-0.023	0.042	-0.020	-0.113
	(0.061)	(0.001)	(0.006)	(0.006)	(0.019)	(0.022)	(0.042)	(0.005)	(0.035)	(0.025)	(0.015)	(0.087)
After FOMC $\times \Delta$ FFR	0.043				0.019				0.007			
	(0.055)				(0.033)				(0.016)			
After FOMC $\times \Delta$ Shadow Rate		0.010				-0.031				-0.001		
		(0.018)				(0.313)				(0.019)		
After FOMC $\times$ Federal Funds Rate Factor			-0.026				0.004				0.085	
			(0.016)				(0.003)				(0.074)	
After FOMC $\times$ Forward Guidance Factor			-0.077				-0.009				0.127	
			(0.055)				(0.014)				(0.159)	
After FOMC $\times$ Large Scale Asset Purchases Factor			-0.410				-0.071				0.0122	
			(0.246)				(0.043)				(0.097)	
After FOMC $\times$ Unified Factor				-0.008				-0.054				-0.136
				(0.121)				(0.060)				(0.149)
Individual and FOMC Fixed Effects	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓
Total Observations	49,985	49,985	49,985	49,985	49,985	49,985	49,985	49,985	49,985	49,985	49,985	49,985
Number of Respondents	10,741	10,741	10,741	10,741	10,741	10,741	10,741	10,741	10,741	10,741	10,741	10,741
R^2	0.476	0.476	0.476	0.476	0.574	0.574	0.574	0.574	0.398	0.398	0.398	0.398

Notes: Fixed effects regressions. Columns (1) (2) use the monetary policy measures for the Federal Funds Rate and Shadow Rate as measured in Wu and Xia (2016). Column (3) uses the monetary policy factors as measured in Swanson (2021). Column (4) uses the unified policy shock as measured in Bu et al. (2021). Standard errors are in parenthesis and are clustered at the respondent level.

* p < 0.10, ** p < 0.05, *** p < 0.01

Table 2.10. [Results] Monetary Policy Announcements and Google Trends Search Intensities

	FOMC	Federal Funds Rate	Inflation	Quantitative Easing	Monetary Policy
	(1)	(2)	(3)	(4)	(5)
Δ FFR	54.43**	77.31***	-10.70	-22.84	-0.60
	(24.56)	(27.51)	(23.79)	(31.74)	(22.10)
Δ Shadow Rate	-5.39	2.91	12.49	6.75	-5.77
	(10.71)	(11.99)	(10.37)	(13.84)	(9.63)
Federal Funds Rate Factor	-14.73	9.51	-20.69	17.63	22.47
	(14.05)	(15.74)	(13.61)	(18.16)	(12.64)
Forward Guidance Factor	-3.65	-10.98**	0.04	0.42	-6.13
	(4.05)	(4.54)	(3.92)	(5.24)	(3.64)
Large Scale Asset Purchases Factor	-5.26	-1.54	0.59	-1.97	5.08
	(5.04)	(5.65)	(4.89)	(6.52)	(4.54)
Unified Factor	43.63	233.53**	-5.71	109.84	56.86
	(82.71)	(92.62)	(80.10)	(106.88)	(74.41)
R^2	0.074	0.175	0.044	0.036	0.082
F Ratio	1.269	3.395	0.741	0.596	1.421

Notes: Results based on regressing the absolute value of various types of monetary policy announcement measures on Google Trends Search Interest (by relative intensity) over June 2013 through December 2021. The keywords are "FOMC", "Federal Funds Rate", "Inflation", "Quantitative Easing" (which is a stand in for "Large Scale Asset Purchases due to the relative popularity and interchangeability of the two words), and "Monetary Policy". Robust standard errors are in parentheses.

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.001$

Figure 2.1: Average Responses surrounding Event Study FOMC Time Window

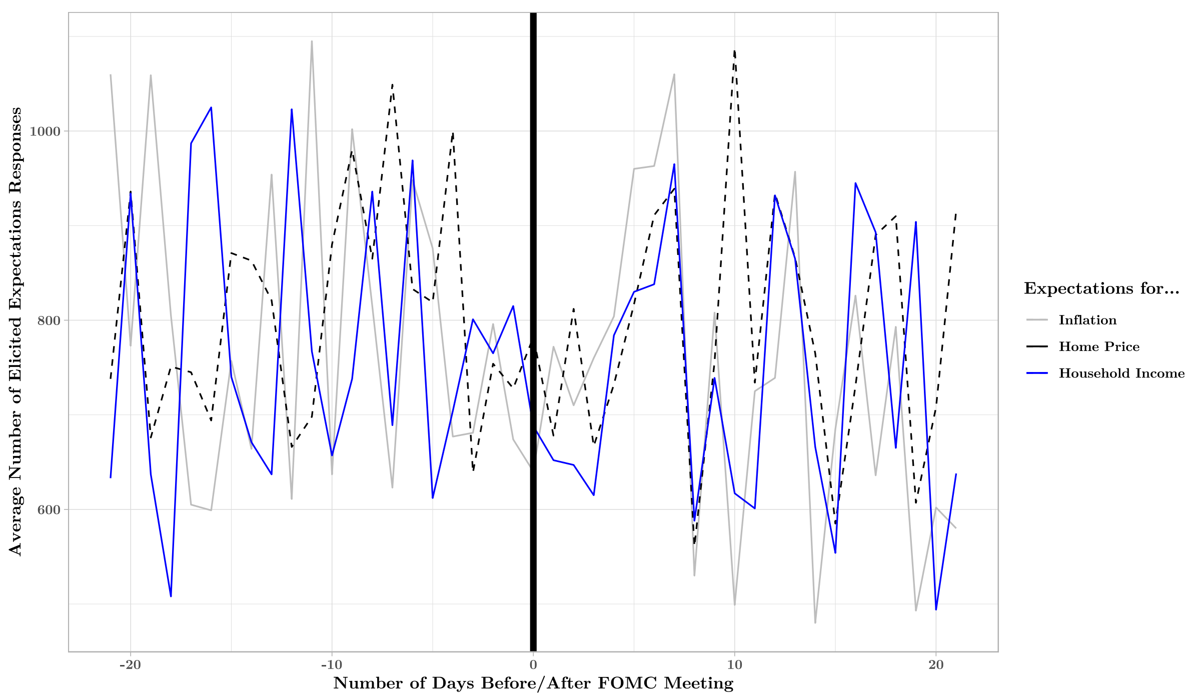

Notes: Average number of responses across a subset of the elicited expectations from the SCE. The thick black line represents the day of the FOMC meeting announcement; the symmetrical number of responses around this set time window of three-weeks before and after captures about the same number of responses for nearly all the expectations variables, including the ones not included in this figure. The subset here represents solicited expectations for One Year Ahead Expected Inflation Rate from survey date, Home One Year Ahead Expected Home Price Change from survey date, and One Year Ahead Expected Change in Household Income from survey date.

Figure 2.2: Measures of Monetary Policy Announcements (Shocks)

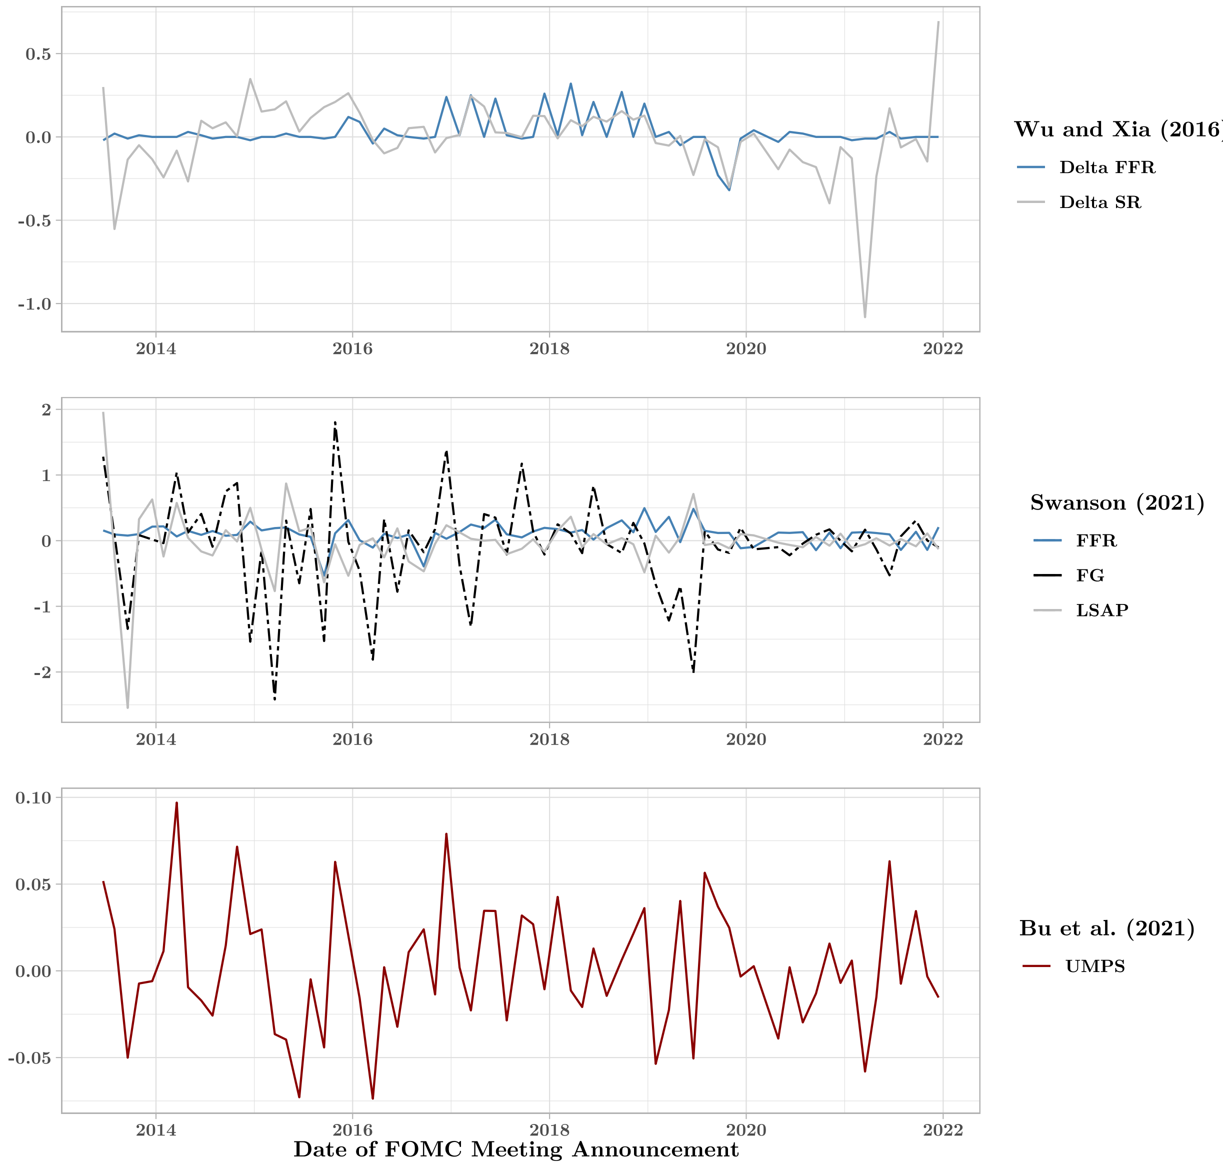

Notes: Six measures of monetary policy that are used as the exogenous shocks to expectations for the analysis. The top panel shows the month over month change in the Federal Funds Rate and the Shadow Rate as calculated by Wu and Xia (2016). The middle panel shows the Swanson (2021) monetary policy surprise factors with an extension to December 2021 by using the methodology outlined in Bauer and Swanson (2022). Lastly, the unified measure of monetary policy announcements are taken from the shock series of Bu et al. (2021).

Figure 2.3: Expectation Impact following Announcement (Changes in Federal Funds Rate)

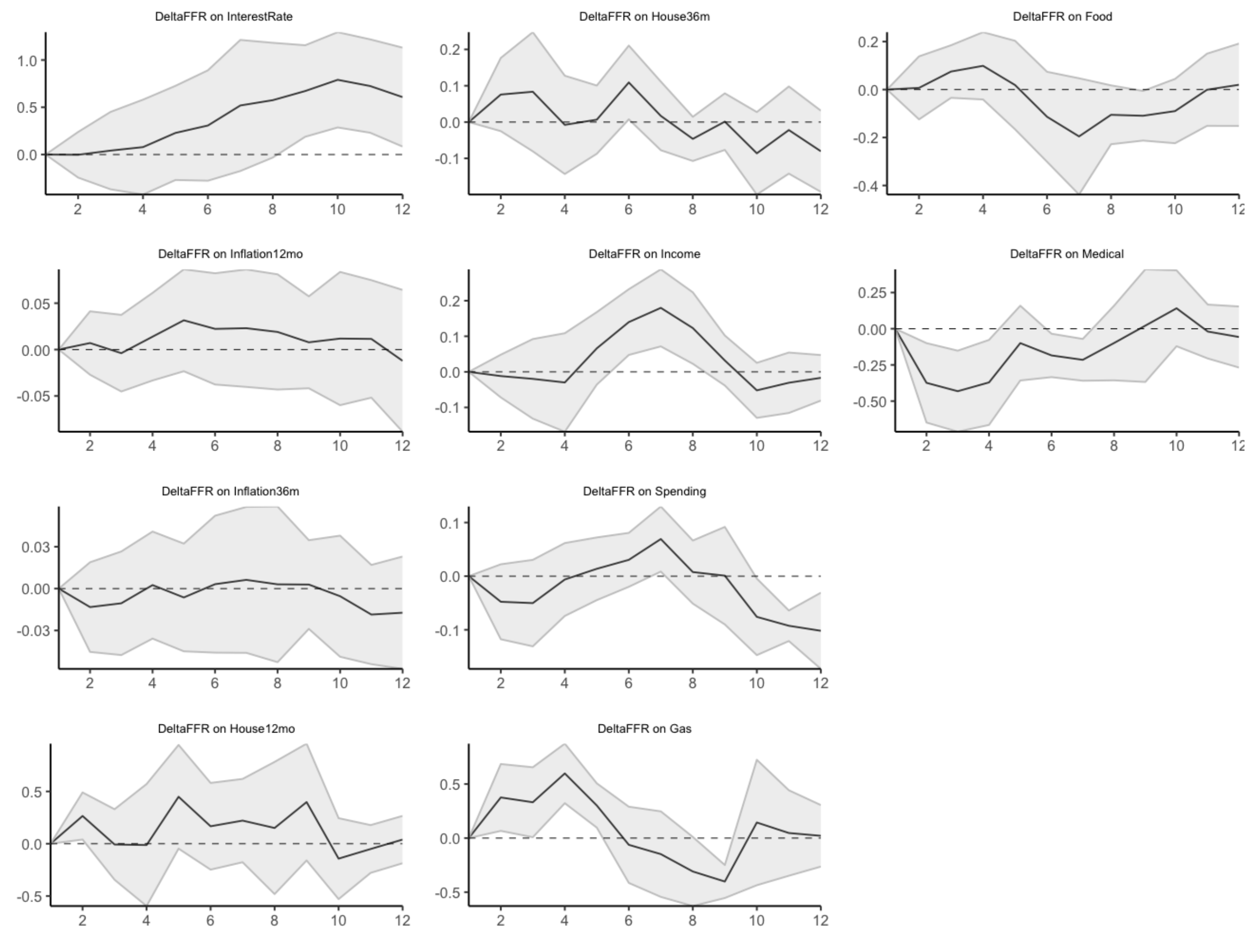

Notes: Estimates based on local projections for up to twelve months in Equation 2.2 of the various expectations variables in the SCE analysis on the monetary policy surprise coming from the change in the Federal Funds Rate calculated by Wu and Xia (2016). Responses are scaled to a shock corresponding to a 25 basis-point increase in the respective rate. Changes in the response variable correspond to the distinct level each expectation is elicited in, such as probability of interest rate increasing for the top left panel. Shaded areas denote 90% Newey-West confidence intervals.

Figure 2.4: Expectation Impact following Announcement (Changes in Shadow Rate)

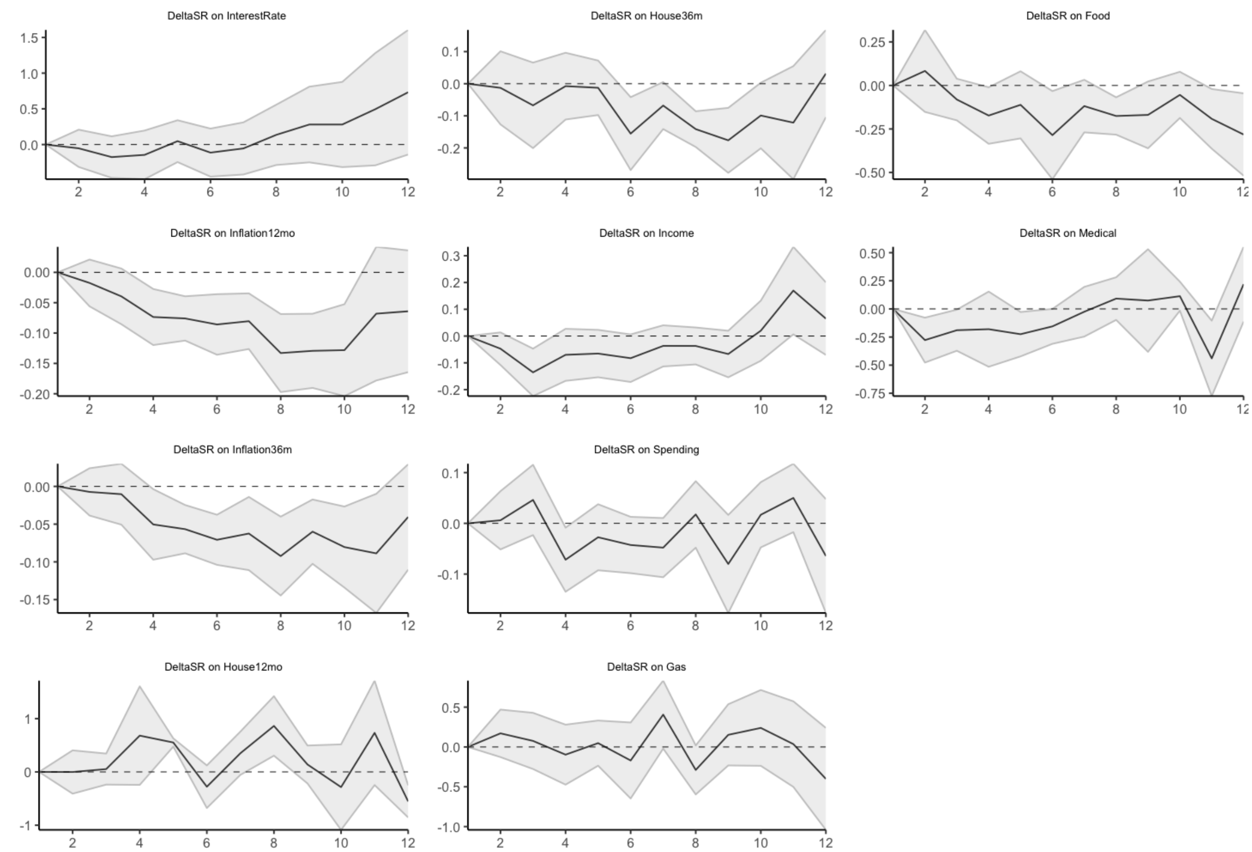

Notes: Estimates based on local projections for up to twelve months in Equation 2.2 of the various expectations variables in the SCE analysis on the monetary policy surprise coming from the change in the Shadow Rate calculated by Wu and Xia (2016). Responses are scaled to a shock corresponding to a 25 basis-point increase in the respective rate. Changes in the response variable correspond to the distinct level each expectation is elicited in, such as probability of interest rate increasing for the top left panel. Shaded areas denote 90% Newey-West confidence intervals.

Figure 2.5: Expectation Impact following Announcement (Federal Funds Rate)

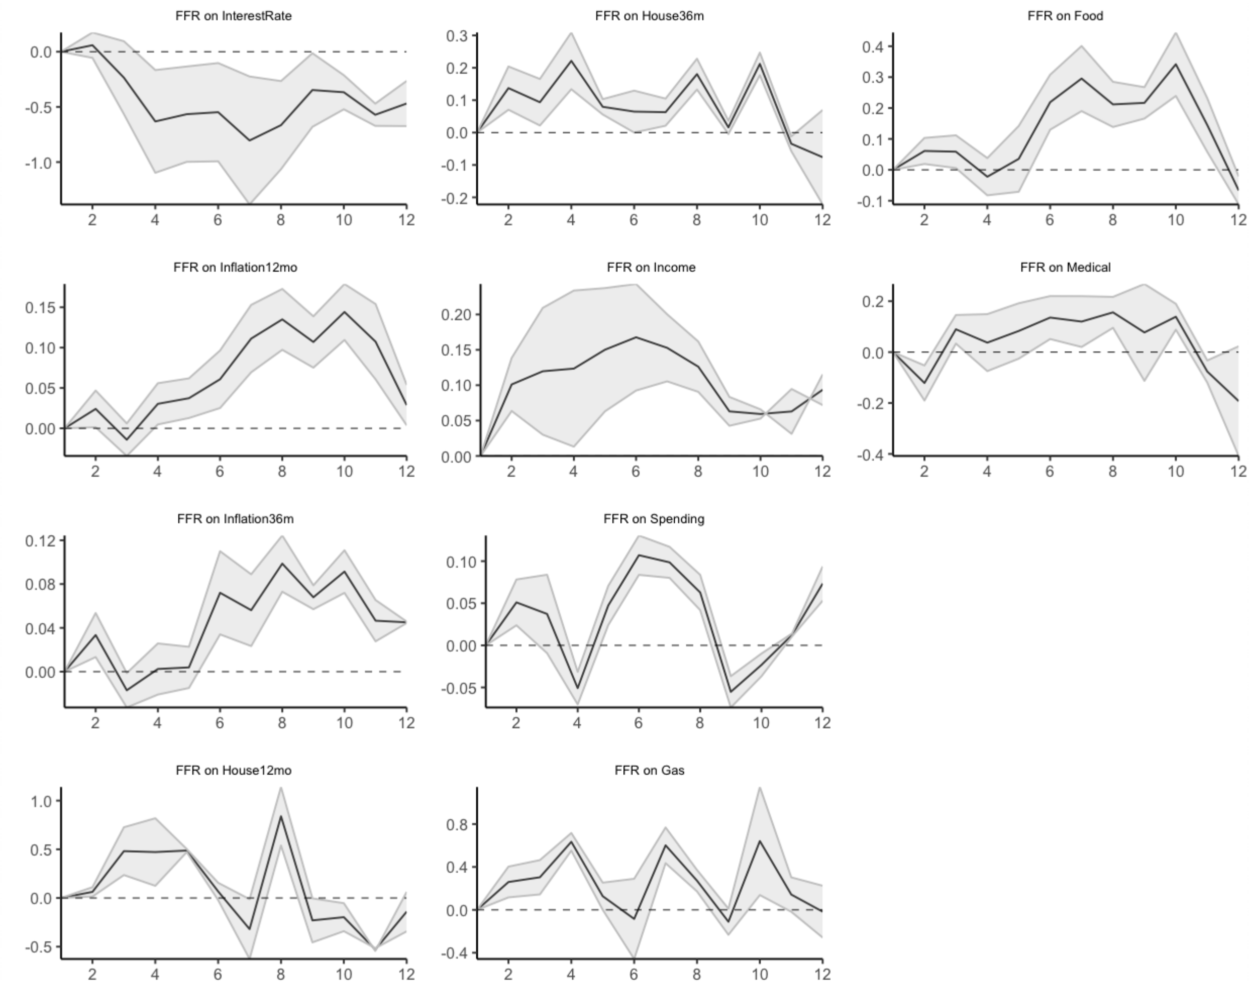

Notes: Estimates based on local projections for up to twelve months in Equation 2.2 of the various expectations variables in the SCE analysis on the monetary policy surprise coming from the the Federal Funds Rate factor by Swanson (2021). Responses are scaled to a shock corresponding to a 25 basis-point increase in the respective rate. Changes in the response variable correspond to the distinct level each expectation is elicited in, such as probability of interest rate increasing for the top left panel. Shaded areas denote 90% Newey-West confidence intervals.

Figure 2.6: Expectation Impact following Announcement (Forward Guidance)

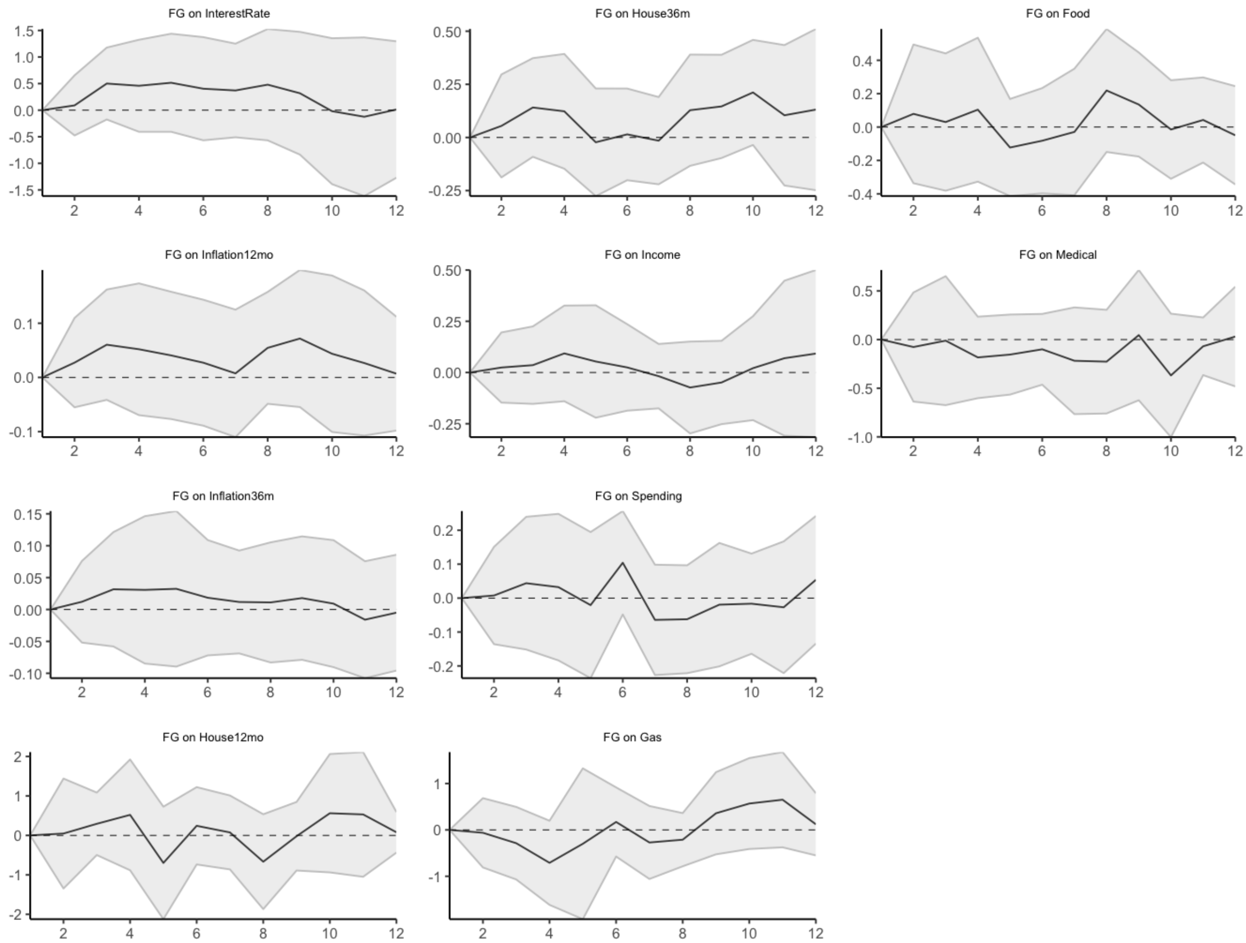

Notes: Estimates based on local projections for up to twelve months in Equation 2.2 of the various expectations variables in the SCE analysis on the monetary policy surprise coming from the the Forward Guidance factor by Swanson (2021). Responses are scaled to a shock corresponding to a 25 basis-point increase in the respective rate. Changes in the response variable correspond to the distinct level each expectation is elicited in, such as probability of interest rate increasing for the top left panel. Shaded areas denote 90% Newey-West confidence intervals.

Figure 2.7: Expectation Impact following Announcement (Large Scale Asset Purchases)

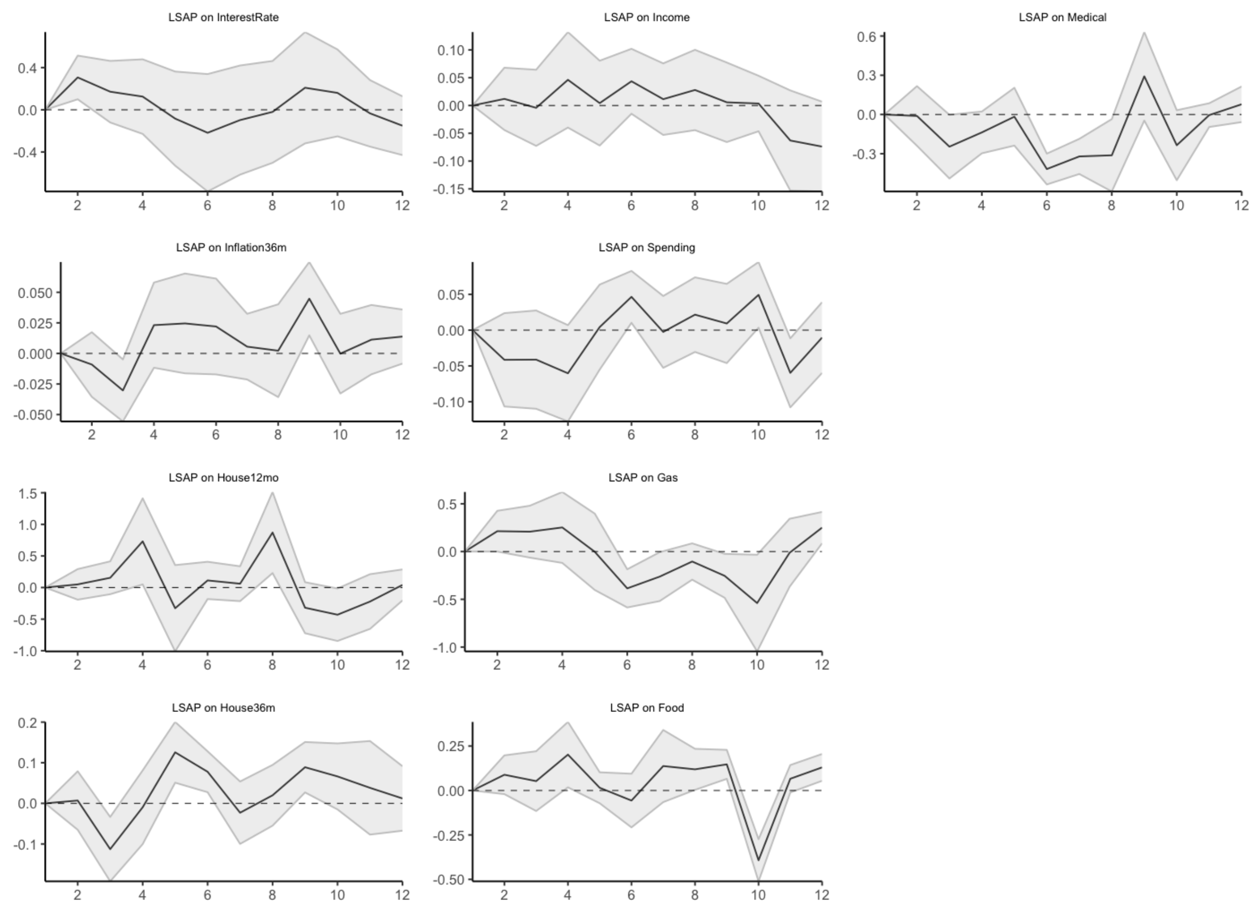

Notes: Estimates based on local projections for up to twelve months in Equation 2.2 of the various expectations variables in the SCE analysis on the monetary policy surprise coming from the the Large Scale Asset Purchases factor by Swanson (2021). Responses are scaled to a shock corresponding to a 25 basis-point increase in the respective rate. Changes in the response variable correspond to the distinct level each expectation is elicited in, such as probability of interest rate increasing for the top left panel. Shaded areas denote 90% Newey-West confidence intervals.

Figure 2.8: Expectation Impact following Announcement (Unified Policy Measure)

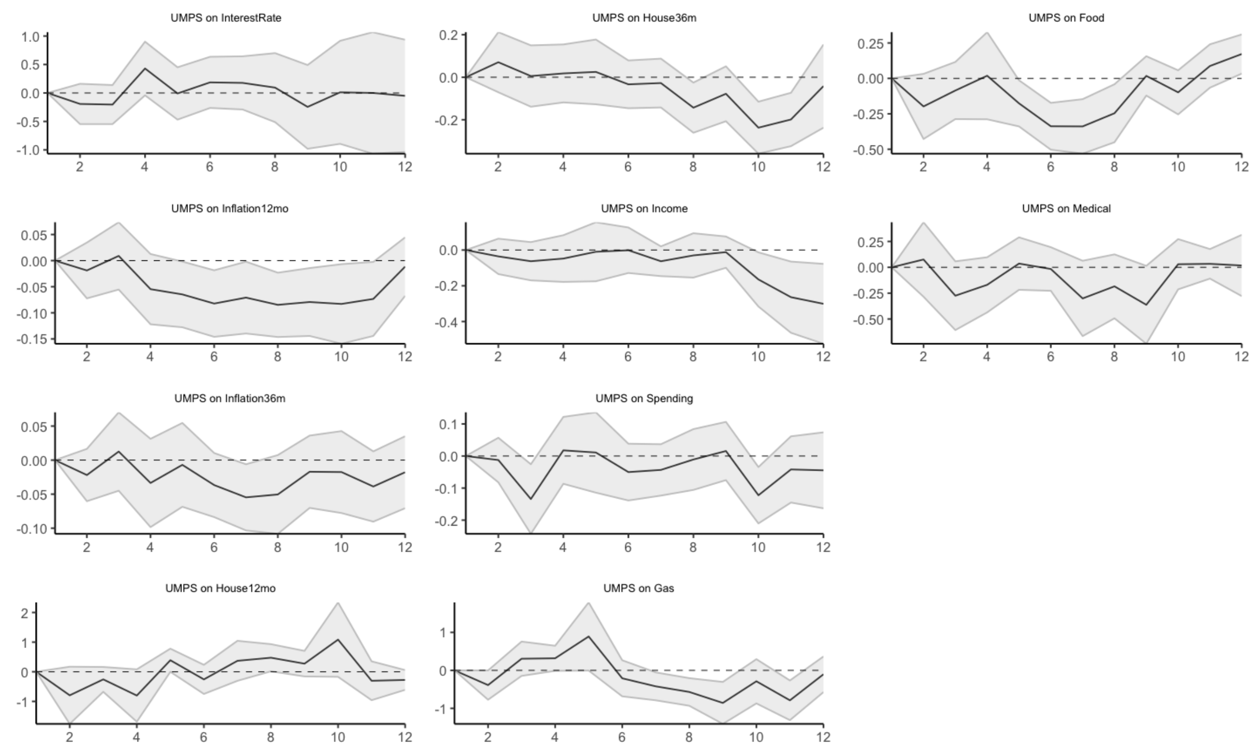

Notes: Estimates based on local projections for up to twelve months in Equation 2.2 of the various expectations variables in the SCE analysis on the monetary policy surprise coming from the the unified policy measure by Bu et al. (2021). Responses are scaled to a shock corresponding to a 25 basis-point increase in the respective rate. Changes in the response variable correspond to the distinct level each expectation is elicited in, such as probability of interest rate increasing for the top left panel. Shaded areas denote 90% Newey-West confidence intervals.

Figure 2.9: Search Intensity for five keywords surrounding the FOMC via Google Trends

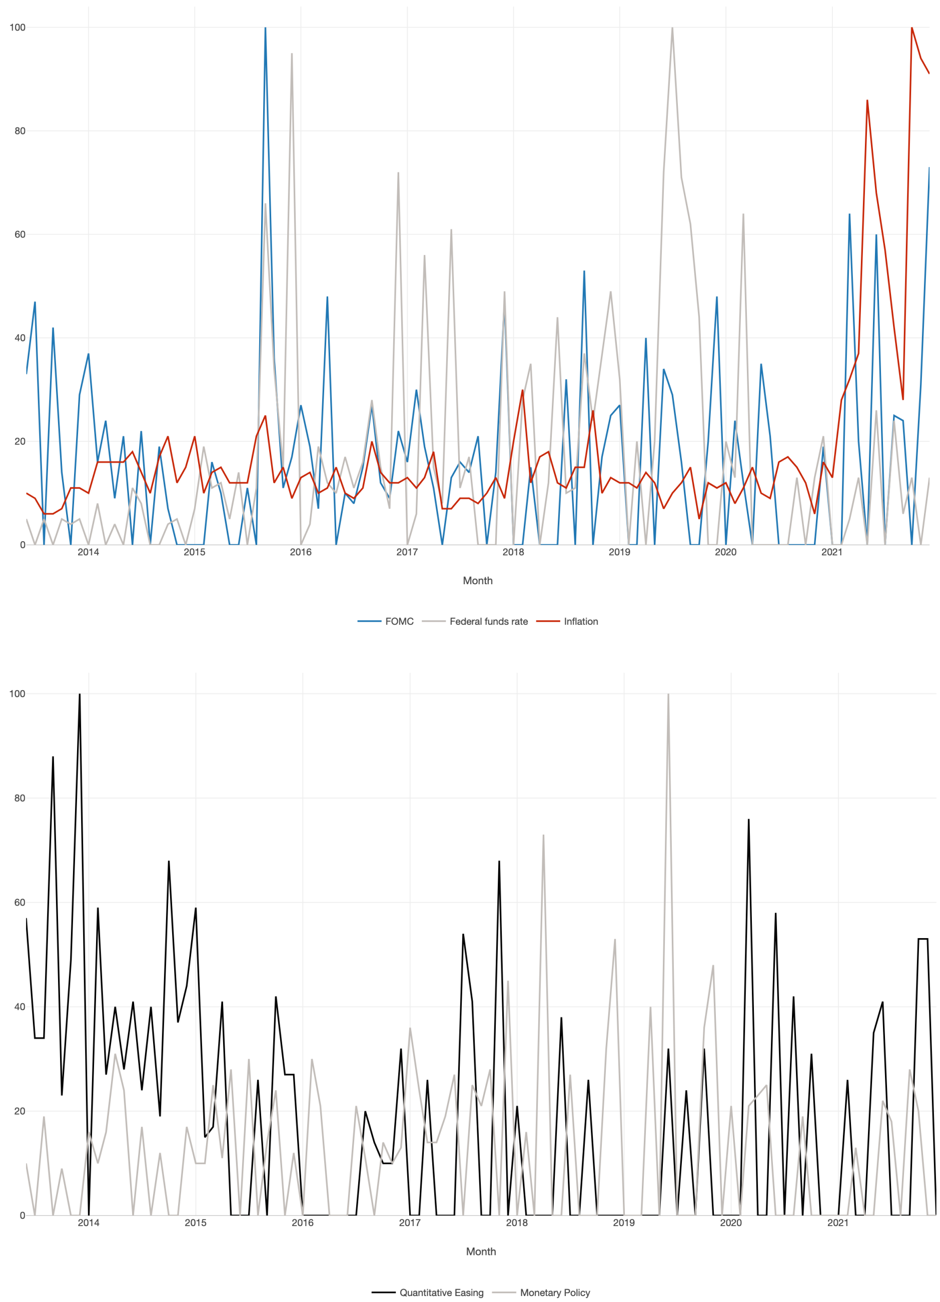

Notes: (a) Keywords here are "FOMC", "Federal Funds Rate", and "Inflation". (b) Keywords here are "Quantitative Easing" and "Monetary Policy". The time period covers the event study analysis from June 2013 to December 2021. Google Trends creates a point of relative intensity (=100) to show the highest intensity and scales the rest of the intensities relative to this.

…[My] results, which come from a sample spanning the period from January 1973 to December 2012, are striking. Upward moves in excess bond premium – again, those corresponding to a widening of credit spreads–are very informative about the future evolution of the real economy …I have to believe that our macro models will ultimately be more useful as a guide to policy if they build on a more empirically realistic foundation with respect to the behavior of interest rates and credit spreads.

Governor Jeremy C. Stein, March 2014 (Stein, 2014)

3

Credit Market Expectations and the Business Cycle: Evidence from a Textual Analysis Approach

The Great Recession led to renewed interest in the relationship between credit expansion and the macroeconomy, with various parties attempting to confirm or refute competing narratives about the causes and propagation of financial crises. Economists noted that developed countries often experience alternating periods of growth and decline in real and financial activity. These so called "boom-bust cycles" are characterized by high levels of investment, output, and leverage, as well as low credit spreads; conversely, what often follows is a reversal in which credit spreads rise and investment, output growth, and leverage decline (Schularick and Taylor, 2012; López-Salido et al., 2017). In addition, credit spreads are widely used by economists, policymakers, and market practitioners as a measure of financial strain, and changes in credit spreads are often seen as a leading indicator of future economic activity, which is conveyed as much in the introductory quote.

Traditionally, theories on the causes of financial instability have often centered on the amplification of shocks that can sometimes be traced back to underlying fundamental factors, but may also arise from financial shocks like a spike in required returns or increased uncertainty (Bernanke and Gertler, 1989; Bianchi, 2011; Eggertsson and Krugman, 2012; Arellano et al., 2019). Specifically, in models with financial market frictions, changes in credit spreads can reflect shifts in the effective supply of funds, which can then affect future economic outcomes. A disruption in the financial market, for instance, could lead to a reduction in the supply of credit, causing credit spreads to widen and leading to a subsequent reduction in spending and production.

Dissent abounds. From both theorists and empiricists, a competing explanation in understanding financial instability emphasizes the role of non-rational beliefs, such as excessive optimism during good times, leading to over-expansion of credit and investment (Minsky, 1977; Gennaioli et al., 2016; Bordalo et al., 2018, 2019, 2020a,b). When beliefs subsequently cool off, credit markets tighten and real activity declines, leading to increased default rates. Credit spreads play a critical role in shaping investor expectations about future credit defaults, and changes in credit spreads can reflect shifts in investor sentiment (Jordá et al., 2013; Greenwood and Hanson, 2013; Baron and Xiong, 2017). In these models, excessively narrow credit spreads can lead to expansions of credit and increased real economic activity, but these patterns will typically reverse when future economic outcomes disappoint investors. This recent research has found that the measured expectations of a broad range of economic agents systematically deviate from rationality, tending to be overly optimistic during good times and then reverting. The key difference between these two types of theories is whether agents know the objective probability distribution in equilibrium. Rational agents should understand the cyclical nature of the credit market and take this into account when forming their expectations. This is why it is crucial to study expectations data directly.

In this paper, I provide empirical evidence on how the errors between realized credit spreads and their forecasts can predict future macroeconomic outcomes. To do so, I first predict a series of past credit spread expectations from 1919Q1 to 2022Q3 by applying textual analysis through natural language processing and topic models in statistical machine learning on front pages of the Wall Street Journal. Then, following the methodology for textual factors in the natural language processing space, I derive the credit spread expectation error which I use as the independent variable in a series of quarterly predictive regressions for GDP, unemployment, and private domestic investment from 1948Q1 until 2022Q3.

I focus on the BAA credit spread which is the difference in yield between the risk-free 10-Year Treasury Bond Yield and the BAA Corporate Bond Yield. Credit spread dynamics refer to the changes in the size of the credit spread over time and it can provide useful information about investors' perceptions of risk and the overall health of the economy because the size of the credit spread is directly related to the level of risk that investors are willing to take on. Wide spreads indicate that investors are demanding a higher return to compensate for the increased risk (less willing to take on risk), while narrow spreads indicate that investors are willing to accept a lower return for the potential reward (more willing to take on risk). The former can be a sign of economic uncertainty or instability as investors are more cautious about the potential risks and rewards of different investments while the latter can be a sign of economic growth and stability, as investors are more confident in the potential returns of different investments.

Credit spread expectations refer to the anticipated changes in the size of the credit spread over a time. These expectations can be based on a variety of factors, such as changes in the level of risk in the economy, changes in interest rates, or changes in investor sentiment. For example, if investors expect the level of risk in the economy to increase, they may anticipate that the credit spread for risky bonds will widen, and they may adjust their investment strategies accordingly. Similarly, if investors expect interest rates to rise, they may anticipate that the credit spread for risky bonds will narrow, and they may adjust their investment strategies accordingly. In theory, analyzing credit spread expectations allow investors and policymakers to gain insight into the market's expectations for changes in the level of risk and the overall health of the economy. Unfortunately, this data is largely collected from surveys and has had limited collection in the past. For instance, the widely used Blue Chip Financial Forecasts only started soliciting forecasts of credit spreads in 1999, covering just two recessions.

By deriving my own proxy for forecasts of credit spreads through a textual analysis approach, I can analyze their empirical relationship with a variety of macroeconomic indicators and the business cycle. I find that overly optimistic sentiment in credit spreads, associated with higher expectation errors, predict a decline in economic activity across three macroeconomic indicators up to four quarters ahead. Specifically, a one-standard deviation jump in the error for credit spread expectations is associated with, at most, a predicted 3% decline in GDP growth, a 1.68% increase in the unemployment rate, and a 2.9% decline in private domestic investment growth. These results are robust to a variety of controls that are common in the literature and include lagged values of the respective indicators, the most recent values for the BAA credit spread and CPI inflation rate, and the changes associated with the 3-month and 10-Year Treasury Yields. The findings empirically corroborate the story in the behavioral models which posit that elevated sentiment predates declines in economic activity while also suggesting that the textual analysis approach in creating historical credit spread expectations and their errors through time has value and provides a proxy for sentiment.

Outline. In the next section, I present the context from the literature most closely aligns to this investigation. In Section 3.2, I present the data acquired from various sources as well as the method by which I construct my own expectations data using statistical machine learning. Section 3.3 presents the empirical specification that set out to predict macroeconomic outcomes using expectations errors. I present the baseline results in Section 3.4, and finally Section 3.5 concludes.

3.1 Context In Literature

This paper contributes to two active research areas. The first is regarding behavioral approaches to predicting business cycles using credit spreads. Whether it be from a theoretical perspective as in Kubin et al. (2019) or by creating a factor that summarizes credit sentiment as in Leiva-León et al. (2022), this literature is most poignantly summarized in Bordalo et al. (2018). There, credits spreads play a critical role as expectations about future credit defaults are over influenced by current news. Bordalo et al. use the Blue

Chip Financial Forecast data to document predictability in credit spread expectation errors and revisions between 1999Q1 and 2014Q4. As a follow up to their analysis, and as a motivational exercise for my textual analysis approach, I perform the same test on data from the Federal Reserve Bank of Philadelphia's Survey of Professional Forecasters (SPF) from 2010Q1 to 2022Q3. Under the assumption of rational expectations (and knowledge of the data generating process), the forecast errors (actual credit spread minus expectation of credit spread) from the professional forecasters should be unpredictable from past data but Figure 3.1 suggests otherwise. Similar to Bordalo et al.'s findings using Blue Chip Financial Forecasts, the SPF corroborates the narrative that when the current spread is low, the expected future spread is too low (actual > forecast = positive error), and when the current spread is high, the expected future spread is too high (actual < forecast = negative error). This visual motivates a simple econometric test of predictability with results in Table 3.1 wherein I regress the actual credit spread (averaged over the next four quarters), the current forecaster expectation (averaged over the next four quarters), and the forecast error (actual spread minus spread expectation) all on the current spread. The evidence does not fit well with the idea that people have rational expectations and instead suggests that analysts' forecasts tend to go through cycles of growth and decline. Column (3) shows that when the current spread is high, the higher the forecast relative to the realization is (thus producing a negative error); the same is true in reverse. When bond markets are doing well (low spreads), expectations are too optimistic and tend to go back to more realistic levels in the future, which can lead to a downturn in the bond markets.

In this space, my paper is most similar to Gilchrist and Zakrajsek (2012) and López-Salido et al. (2017). The first finds that credit spreads contain valuable information about the economy, and that much of this information comes from changes in excess bond risk premiums. While some of the variation these premiums can be explained by frictions in the traditional rational expectations approach, it could also be seen as sentiment. Gilchrist and Zakrajsek use a "bottom-up" approach to construct a credit spread index that allows them to accurately measure investors' expectations of future economic outcomes, despite the presence of time-varying risk premiums, spanning back to the mid-1970s. They use a sample of US non-financial firms covered by the S&P's Compustat database and the Center for Research in Security Prices (CRSP). López-Salido et al. (2017) measure credit-market sentiment based on the expected return to bearing credit risk from 1929 to 2015; more specifically, they use the ex-ante predictable component of corporate bond returns as a proxy for this sentiment and find that elevated sentiment from $t-2$ years ago is associated with a decline in economic activity in years t and $t+1$.[1] López-Salido et al. address concerns that changes in economic activity are driven by external non-financial factors, and that credit spreads simply reflect these changes in advance. This predictable component of changes in credit spreads therefore reflects past shifts in investor sentiment. The current analysis combines the construction of a proxy for credit spread expectations approach from Gilchrist and Zakrajsek with the time span and the predictive regression framework from López-Salido et al.. By contrast, I employ textual analysis from a news source to construct my historical credit spread expectations, which I assume can be used as a direct measure of credit-market sentiment. This sentiment in the corporate bond market is likely to be a key channel of economic transmission that is closely tied to perceptions of credit risk in the financial system. In order to train my machine learning model, I also assume that the dynamics of the observable credit spread expectations remain consistent throughout the entire sample.

The second strand concerns itself with applications of text analysis and machine learning to the fields of economics and finance. For a contemporary introduction, see Gentzkow et al. (2019) who detail the overarching steps by which most of the research in this space, including my own analysis, follows. In essence, Step 1 is to represent the text $\mathcal{D}$ as a numerical array $\mathcal{W}$. Step 2 is to map $\mathcal{W}$ to the predicted values $\tilde{\mathcal{V}}$ of unknown outcomes $\mathcal{V}$. Step 3 is to use the predicted $\tilde{\mathcal{V}}$ in either descriptive or causal analysis. There are various methods by which to use each step, and their applications have spanned widely.

From using text-mining techniques to extract sentiment from financial statements and examining asset returns (Ke et al., 2019; Yue and Jing, 2022), to quantifying political stance and risk using newspaper prints and digital media (Giavazzi et al., 2020; Caldara and Iacoviello, 2022), to even analyzing text from FOMC announcements to study the effects of transparency and measuring monetary policy surprises (Hansen et al., 2018; Shapiro and Wilson, 2019; Doh et al., 2020; Gorodnichenko et al., 2021), the text analysis approach is gaining traction in economics due to computational efficiency. I contribute to this area by following a number of key papers to predict expectation error of credit spreads through time.

[1] In their analysis, elevated sentiment means that the expected return to bearing credit risk is low.

First, following Doh et al. (2020), I use vector embedding and employ the cosine similarity approach to represent similar texts within the front pages of the Wall Street Journal (WSJ). This choice of newspaper is guided by similar studies filling in historical variables using text from WSJ such as Manela and Moreira (2017) who fill in the CBOE S&P 500 Volatility Index (aka VIX), and Kelly et al. (2021) who do the same for nonfarm payroll employment and housing starts. Next, following the Textual Factors model from Cong et al. (2020), I employ Locality Sensitive Hashing (LSH) to cluster similar embeddings that most closely predict expectation errors; LSH is a method borrowed from the neural information processing literature and has been used in a variety of social science applications (Andoni et al., 2015). Lastly, I use a Latent Dirichlet Allocation technique to uncover topics in the unstructured WSJ text data without linking themes to particular word lists prior to my estimation. This approach is used in Hansen et al. (2018) to uncover latent themes in the text database of FOMC transcripts and is meant to produce a number of textual factors by which I can quantitatively calculate loadings, or how much a document represents a certain topic, that can serve as a proxy for the sentiment regarding credit spreads over time.

3.2 Data and Construction of Explanatory Variables

After describing the acquired data from public databases, I proceed with the construction of my explanatory variables which are textual factors that can be thought of as proxies for sentiment through time. These proxies are generated using front page articles from the Wall Street Journal which have been used in a variety of economic studies I describe in Section 2.1. The acquired data on BAA credit spread dates back to January 1919 and is available until the end of the third quarter of 2022 (1919Q1 to 2022Q3), totaling 1,245 months, or 415 quarters. The Survey of Professional Forecasters began acquiring BAA credit spread expectations in January 2010 until present day; matching the end of the third quarter limit from the realized credit spread means it spans 2010Q1 to 2022Q3 and covers a total of 153 months, or

51 quarters.[2] The macroeconomic data obtained from other public data sources such as FRED span, at most, 1948Q1 to 2022Q3, which is 897 months or 299 quarters.

3.2.1 Acquired Data

My entire analysis relies on publicly available data for which historical realized values are plentiful. Sources such as the Federal Reserve Economic Database (FRED) and Capital Markets Data provide the realized values of my key indicator of credit market conditions which is the difference between the Moody's seasoned BAA Corporate Bond Yield and the 10-Year Treasury Bond Yield (interchangeably referred to as the BAA credit spread). This measure can be calculated from data ranging back to January 1919 until the most recent complete quarter of 2022 as of writing. Much less plentiful are data on the forecast values of the BAA Corporate Bond Yield, which, when looked at its difference with the 10-Year Treasury Bond Yield, can be referred to as the forecast of the BAA credit spread (or BAA credit spread expectation). Proprietary sources, such as Moody's historical BAA Corporate Bond yield forecasts and the forecasts from the Blue Chip Financial Forecasts services, have been often used in the aforementioned literature. One advantage of using the Blue Chip Financial Forecasts, for instance, is that information is solicited monthly from 40 panelists working in major financial institutions such as S&P Global, Goldman Sachs, and Citibank; the Blue Chip data, and expectation data on this variable in general, is limited and for this service is available from January 1999 onward. Forecasts for the current quarter up to five future quarters are obtained and averaged to derive a consensus forecast that can be used to derive a BAA credit spread expectation; the difference between this expectation and the realized BAA credit spread is called the expectation error.

I opt to use a similar forecast solicitation survey through the Federal Reserve Bank of Philadelphia which, since January 2010, has been collecting forecasts of the BAA Corporate Bond Yield in their Survey of Professional Forecasters (SPF). The SPF consists of a limited panel of professional forecasters (anywhere from 30 to 45 respondents per survey round) and they give their quarterly projections on ma-

[2] In comparison, the longest survey that includes credit spread expectations is the proprietary Blue Chip Financial Forecasts which have span from 1999Q1. Studies that have used it for credit spread analysis, such as Bordalo et al. (2018), analyze up until 2014Q4 which spans 192 months, or 64 quarters, or about 25% more quarterly data on BAA credit spread expectations than the SPF as of this analysis.

jor macroeconomic indicators up to four quarters ahead. Most forecasters use a mathematical model, adjusted with their subjective judgments, to ascertain their projection while many update forecasts at a monthly frequency (Stark, 2013). I obtain the average point forecasts of the BAA Corporate Bond Yield that are forecasters' projections of the current quarter t to up to four quarters $(t+4)$ ahead forecast, given that they have up until $t-1$ quarter's information. Mathematically, the SPF provides credit spread $CS_{t+\tau|t-1}$ for $\tau = 0, 1, 2, 3, 4$. Coupled with the realized 10-Year Treasury Bond Yield, I can calculate a measure of BAA credit spread expectation in much the same way as I would if I were using proprietary data.

For the analysis of forecasting GDP and other macroeconomic variables, I follow standard sources such as FRED and the ArchivaL Federal Reserve Economic Data (ALFRED) from the St. Louis Federal Reserve Bank. The most historical vintage on GDP is from 1929 and is provided on a yearly basis. Given that the SPF uses quarterly data, I decide to limit my analysis on the most common vintage for GDP data which begins in 1948. On a quarterly basis, I take GDP, unemployment, the CPI Inflation Rate, the 3-Month Treasury Yield, and the 10-Month Treasury Yield from 1948Q1 to 2022Q3, which is 897 months or 299 quarters.

3.2.2 Constructed Data — *The Wall Street Journal*

Most newspapers, magazines, and other sources of information where historical expectations data could be reported are found in scanned images that archives and other historical data repositories allow. To create a more thorough range of data on credit spread expectations, I would need to access these images and extract the data from them. With recent advances in technology, we now have optical character recognition (OCR) software that allows me to perform this function efficiently. I use Textract, a machine learning process hosted by Amazon Web Services (AWS), to create my own textual factor loadings from inputted text data from the Wall Street Journal front pages on expectations for the implied BAA credit spread from January 1919 (1919Q1) to September 2022 (2022Q3). This will complement the Survey of Professional Forecasters data by the Federal Reserve Bank of Philadelphia which covers from 2010Q1 until 2022Q3 by creating a proxy for sentiment on credit spread expectations; I will use the textual factor loadings from the realized credit spread expectations from 2010Q1 to 2022Q3 as the training data and then predict those past credit spread expectations from 1919Q1 to 2009Q4. Since the Wall Street Journal is printed six days a week (excluding 8 federal holidays), I simplify the analysis by averaging the textual factors within a given month (a median of 26 observations per month) or within a given quarter (a median of 78 observations per quarter).

Sections 3.2.3 to 3.2.5 follow the Textual Factors model from Cong et al. (2020) which quantitatively represents text data while preserving interpretability through its informational structure; I leave the main derivations to their paper and discuss the main points below. The goal is to create a series of textual factors that can be loaded into a regression model to estimate the expectation error through time. These factors are created as vectors that quantitatively explain main variations in texts using locality of words and similarities across texts as the main drivers. In some sense, they can also be thought of as a proxy for sentiment across time. This process is computed in three steps using text data as an input. I represent the process through a simple diagram in Figure 3.2 and go into the high level thought process of each step below before describing the algorithmic details in the proceeding subsections.

First, as I explain in Section 3.2.3, I will create vector representations of words that account for their semantic and syntactic meanings. These ”word embedded” vectors will result in a rather large data set where words in various documents are represented by a vector according to distance and similarity with each other whose fixed dimension is set by a neural network; this process is guided by the assumption that words with similar meanings are often used together. For a simple example, consider a scenario where the rule is to create a vector representation using average frequency of neighboring words, and the fixed dimension is the top 5 most frequent words. In a list of words such as *credit*, *stimulus*, and *Senate*, we might find that the words most frequently neighboring *credit* and *stimulus* are *money*, *card*, and *check*. For *Senate*, it might be *decision* and *bill*. Following our rule, the word *credit* might be represented by $[4, 8, 5, 2, 1]$ which means *money* neighbors the word an average of 4 times, *card* neighbors the word an average of 8 times, *check* neighbors the word an average of 5 times, *decision* neighbors the word an average of 2 times, and *bill* neighbors the word an average of 1 time. In the same way, *stimulus* might be represented by $[5, 2, 4, 2, 2]$ and *Senate* might be represented by $[0, 0, 0, 9, 10]$. Note that the actual rule will be more complex than just frequency and rely on the context that the word is being used in. A large collection (often billions of words) gets fed to the model and a sliding window is used to capture the words that lie on either side of the word to determine its context; this context is then represented by its embedding vector that is controlled by a negative sampling process to update weights given to surrounding words. Because words with a similar context usually have closely-linked meanings, such words will end up having similar embedding vectors too. Even in the above example, dimensionality is already a concern; with just three words and with a fixed dimension of 5 words, we have created a 3 x 5 matrix that represents a document; in this case, a document is the front page of a Wall Street Journal volume which would have W number of words to represent.

Second, as I explain in Section 3.2.4, I first tackle the high dimensionality problem from Step 1 in Step 2 by clustering similar words together using a method called Locality Sensitive Hashing. Cascitti et al. (2022) provide a timely and layman-friendly article in the computational science space but the main idea behind this method is to use a hash function that takes an arbitrary amount of data and approximately categorizes it according to bins; these bins are chosen according to the context of the word space it is analyzing based on locality, or distance, to other words. Then, the clustering is more likely to happen for input text values that are close together rather than for inputs that are far apart. For example, consider the word *amazon*. This word, and its attached vector representation, could mean the worlds largest natural rainforest if words such as *jungle* and *biome* were nearby versus referring to the multinational tech company if words such as *alexa* and *.com* were nearby. The word would then be clustered differently based on its locality. Following the simple example from earlier, the three words *credit*, *stimulus*, and *Senate* would be clusters if their locality of neighboring words in that particular Wall Street Journal page were similar enough, and they would become the representative clusters for that document. In essence, I take the W number of words and divides them into K clusters.

Lastly, with these word vector embeddings from Step 1 clustered in different bins according to their locality in Step 2, I employ a topic model to reduce dimensionality by representing different words according to similar topics per document. This process is done via a Latent Dirichlet Allocation and will allow me to find documents according to whatever query I desire. For example, if I want to query the phrase *credit spread*, LDA will search for that phrase amongst the topics it has attributed different documents to. This is where the clustering from Step 2 serves it purpose: the clusters guide the topic search so that the algorithm does not need to comb through all the words. In essence, the K word clusters get fed into the algorithm as an educated guess of what the topics representing that document means, reducing computational time to finding the most optimal topic representation of that document. From the earlier example, a query with the term *credit spread* will pick up the document represented by the three clusters of *credit*, *stimulus*, and *Senate* since *credit* is a match, but it will rank that document less than one whose three clusters are, say, *credit*, *spread*, and *recession*. Thus, the algorithm will identify topics, or textual factors, that are represented by a set of words and the relative frequency distribution of their frequency in a given body of documents. Given that I am looking for specific topics in a body of words related to the BAA credit spread expectation, I can require that the query searches through clusters which have relevant words such as *BAA, bond, Treasury, yield, expectations, spread,* etc.

3.2.3 Word Embedding

The initial stage in any textual analysis is to summarize or represent the words that are present in the texts; this is called embedding. Less complex and count-based statistical models for textual analysis in the social sciences frequently use the one-hot vector encoding representation where words (dubbed N-grams) are treated as extremely high dimensional vectors/indices over a vocabulary with only one 1 and lots of 0s, omitting any consideration of the semantic relations among words. The flaw here is that we end up with words treated as independent units; *economic boom, upturn, strong markets* would all be treated as unrelated, which is inaccurate. To overcome this, the literature turns to semantic vector-space models with real-valued vector representations to create one-hidden-layer neural-network models; starting with Mikolov et al. (2013), this approach has gained traction in economics research to create a proxy for sentiment through these vector representations (Cheng et al., 2022; Dubovik et al., 2022; Fano and Toschi, 2022). Specifically, I follow the Word2Vec version of word embedding, pre-trained by Google on its Google News dataset, to let me filter out common typos, words and phrases that are too frequent (words such as *the, it, then*), and other common words not associated with the sentiment I want to capture (for example, *table of contents, what's news, DOW JONES, wsj.com* and other recurring words on the front page). This processing method for newspaper data is similar to the one done in Manela and Moreira (2017). Once fed documents, the Word2Vec model loops on the words therein to map them to a real-valued p-dimensional vector less than the size of the document vocabulary $|V|$ to a create a learned embedding vector $w \in \mathbb{R}^{p \times V}$. It will then calculate the distances between vectors w_i and w_j with a metric called cosine similarity, defined as

$$\text{similarity}(w_i, w_j) = \arccos \frac{\langle w_i, w_j \rangle}{|| w_i |||| w_j ||} = \frac{\sum_{i=1}^{n} w_i \sum_{j=1}^{n} w_j}{\sqrt{\sum_{i=1}^{n} w_i^2} \sqrt{\sum_{j=1}^{n} w_j^2}} \tag{3.1}$$

where $w_{i,j}$ are the two vectors being compared with their different indices and where a higher cosine similarity implies a more similar vector representation (and thus a more similar semantic meaning in the context of the sentences it appears in). This produces distances in a range $[-1, 1]$ denoting total opposites to exactly the same. For example, if the algorithm is trained on the words *Apple, Bill Gates*, it might spit out $\text{similarity}(\textit{Microsoft}, \textit{SteveJobs}) = 0.781$. I pre-process the words as per Manela and Moreira (2017) to reduce the dictionary of words I consider and then, following Mikolov et al. (2013), use the Word2Vec model to generate a 300×1 vectors for each remaining word that appears in the data from the front pages of the Wall Street Journal through time.

3.2.4 Clustering

To cluster the $w \in \mathbb{R}^{p \times V}$ word vectors, I turn to Locality Senstitive Hashing which returns the nearest-neighbor information through constructing a series of hash functions H that assert the similarity of items in order to put them into bins. Conceptually, on a 2D space, you can think of these bins as being created by random lines the algorithm is generating to cut up the observations based on similarity of the observations. In essence, the hash functions are generally claiming that vectors are similar when they are close together. More specifically, for any random element $h(\cdot) \in H$, which is our case are the word embedded vectors,

$$P[h(w_i) = h(w_j)] = p_1, \quad \text{for any } w_i, w_j \text{ such that } d(w_i, w_j) \leq d_1$$
$$P[h(w_i) = h(w_j)] = p_2, \quad \text{for any } w_i, w_j \text{ such that } d(w_i, w_j) \geq d_2$$

For any well defined LSH family $H(d_1, d_2, p_1, p_2)$, p_1 is the probability of retrieving points that are close to a query point and thus I would like to maximize this value; a query point, in this analysis, is a word such as *credit*. We can think of $1 - p_1$ as the occurrence of false negatives meaning that some points that are closer than distance d_1 to the query point won't be retrieved in the sample. Meanwhile, p_2 is the probability of retrieving points that are further than desired to our query point; p_2 is the probability of false positives which I want to minimize. With this in mind, then we can extrapolate that if similarity(w_i, w_j) is high, then $P[h(w_i) = h(w_j)]$ is high as well. Following the tech leader example, if similarity$(Microsoft, SteveJobs) = 0.781$, then we can infer that the probability that the two words are in the same bin is relatively high, too. To define how the hash function family is generated to make the bins, I turn to the random hyperplane projection method which produces a spherically symmetric random vector r of unit length from the $p-$dimensional space using a signum function; in my analysis, since I follow Mikolov et al. (2013), $p = 300$. More concretely, the hash function family for vectors w can be described as

$$h_r(w) = sgn(\langle r, w \rangle), \quad \text{for } r \text{ randomly sampled from the unit sphere } S^{p-1}$$

Using this method, I reduce the highly-dimensional vectors of word embeddings by generating hash functions with good performance in finding the nearest neighbors of a query word, thus creating my bins, or clusters, K.

3.2.5 Topic Modeling

Topic modeling is a statistical technique that uses the distribution of words in a provided texts to identify their underlying semantic structure and generate a set of topics. For example, in my analysis, a Wall

Street Journal front page following Black Monday might contain words like "trade deficit", "financing mergers," and "SEC Chairman" more frequently than other front pages, while a front page about Apple's valuation might contain words like "cash generating," "visionary," and "Tim Cook" more often. The Wall Street Journal may have common topics through time such as crisis, wars and conflicts, natural disasters and healthcare, the Federal Reserve or monetary policy, taxes or Congress, technology, and employment. In practical research, the amount of documents in a given area makes it near impossible to manually categorize them by topic. However, topic modeling can help us automatically identify the topics discussed in each document by observing their word distributions; for this, I follow the Latent Dirichlet Allocation (LDA).

Following the economics literature using computational linguistics such as Bholat et al. (2015), Hansen et al. (2018), and Tobback et al. (2017), I assume a simple, two-distribution data-generating process where each of my Wall Street Journal front pages are generated from a distribution over a collection of topics. In turn, each topic is generated from a distribution of words in a dictionary; many authors opt to feed this dictionary the words most relevant to their analysis by creating something based on frequency while others use pre-trained dictionaries such as the aforementioned one used by Google. Given the rich vocabulary often used in the Wall Street Journal, the computation of the distributions would be onerous. However, by feeding the clusters from Section 3.2.4 as the topic dictionary, I significantly reduce the search complexity of the topic word distributions. Mathematically, and following Cong et al. (2020), let the notation $\beta_k \sim \text{Dirichlet}(\eta)$ be a multinomial distribution over the dictionary of words for each topic and $\theta_d \sim \text{Dirichlet}(\alpha)$ be a multinomial distribution over K topics for a particular document d. Then, the word-generating process for any document is such that I sample a specific topic $z_{di} \in (1, 2, \ldots, K)$ with $z_{di} \sim \theta_d$ and then sample the observed word $w_{di} \sim \beta_{z_{di}}$ from the entire document vocabulary V. In an expression, this takes the form of

$$[\Theta B]_{dw} := P(w_{di} = w \mid [\theta_d, \beta_1, \beta_2, \ldots, \beta_K]) = \sum_k \theta_{dk}\beta_{kw}$$

where I am calculating the probability of word w_{di} to be equal to word w in the defined dictionary space, and where $\Theta = [\theta_1, \theta_2, \ldots, \theta_D]' \in \mathbb{R}^{D\times K}$ and $B = [\beta_1, \beta_2, \ldots, \beta_K]' \in \mathbb{R}^{K\times V}$. Then, denote the product $[\Theta B]_{dw} = N_{dw} \in \mathbb{R}^{D\times V}$ as the number of times a word w appears in a document d such that, in different topics, different words are assigned different weights. This implicitly assumes that topics that allocate similar weights to words are related more closely to each other than those that do not. Then, each document fulfills being described by the two distributions above: the probability that a document covers a certain topic and the probability that topic itself has certain words assigned to it. I follow Mikolov et al. (2013) and limit the size to the top 300 topics the method identifies as important.

3.2.6 Validation of Factors

The output from the above subsections are the K textual factors which are represented by the word support S of the word cluster i (the words which are being used to denote similarity to other topics), the real-valued vector representing the textual factor F_i, and the factor importance d_i (the similarity measure). Together, they are a triplet of information by which Cong et al. (2020) use the following projection to create the loadings of the textual factor i from:

$$x_i^d := \frac{\langle N_{S_i}^d, F_i \rangle}{\langle F_i, F_i \rangle} \tag{3.2}$$

A thought experiment most easily helps understand these textual factor loadings. Unlike structured data documents, which list quantifiable amounts for relevant topics such as stock prices or the 10-Year Treasury Bond Yield, unstructured data has text which may use a certain vocabulary to define the overall sentiment of the topic. For instance, texts from the Wall Street Journal front page surrounding the 2008 Financial Crisis might have centered around discussions covering things like *financial crisis, investor confidence, fiscal policy,* and *bankruptcy*. The x_i^d obtained in Equation 3.2 allow me to assign a quantitative measure to how much the document "loads" on that topic; each document d can then be represented by the loadings $x_1^d, x_2^d, \ldots x_K^d \in \mathbb{R}^K$.

For a visualization of how these loadings translate to analysis, consider the checks on accuracy via the word cluster and loadings chart for the word *recession*. In Figure 3.3, I find that words like *financial*, *depression*, *inflation*, *GDP*, and *unemployment* all have a close proximity to the main word *recession* in the vector space measured by their cosine similarity. Additionally, in sub-figure (b), the textual factors that are generated follow consistent patterns with real-time events that would have those relevant Wall Street Journal front pages related to the topic, such as during the Great Depression, Post WWII spending declines, the oil crises of the 70s and 80s, and the Great Recession.[3]

3.3 Methodology for Credit Spreads and Macroeconomic Indicators

My methodology for this section can be broken into two parts. First, I predict past credit spreads through the entire span of available realized BAA credit spreads by first following the predictive regression methodology for credit spreads in Bordalo et al. (2018) and then enhancing that with my textual factor loadings. I predict the BAA credit spread expectations from 1919Q1 to 2009Q4 by first training my model with the SPF data from 2010Q1 to 2022Q3. Then, with the predicted credit spread, I can generate my series of the credit spread expectation error by subtracting out the realized credit spread. With the predicted credit spread expectation error series, I am then able to run predictive regressions on changes in macroeconomic indicators across different time horizons. In short, I use the textual factors from Equation 3.2 to predict BAA credit spread expectations in Equation 3.3, which in turn allow me to calculate the expectation error that I treat as the independent variable in Equation 3.4.

3.3.1 Historical Credit Spread Expectations

Similar to Manela and Moreira (2017) who use regressions to predict a Volatility Index (VIX) using n-gram frequencies from their body of text, I will do the same to predict BAA Credit Spread Expectations using the textual factor loadings from my body of text. Given the high number of topics that will be involved in the analysis, I opt to use penalized Lasso regressions so that only the topics that most contribute

[3]To further decrease the noise from the textual loading generation process, I follow Manela and Moreira (2017) and omit other words and phrases found regularly in the Wall Street Journal such as *"business and finance", "world wide", "what's news", "table of contents", "masthead", "other", "no title", and "financial diary".*

to the credit spread expectations stand out.[4] Then, using the BAA Credit Spread Expectations data from the SPF covering 2010Q1 through 2022Q3, I estimate the following:

$$E[CS]_t = \alpha + \mathbf{x}_t^D + \eta_t, \quad t = 1, 2, \ldots, T \tag{3.3}$$

where the BAA credit spread expectations $E[CS]_t$ for quarter t are being calculated by the textual factor loadings $\mathbf{x}_t^D$ estimated during the training time period; is a K vector of regression coefficients which, from the LDA using the inputted text, uses 300 topics, and T is the training time period for $T = 51$ observations in the SPF sample. I compare this with the Bordalo et al. (2018) model from Section 2.1 where $E[CS]_t = \alpha + \gamma[CS]_t + \varepsilon_t$. Next, I perform the same analysis as in Equation 3.3 but this time predicting the variable over the entire span of the textual factor loadings from 1919Q1 to 2009Q4. From this, I can calculate the predicted credit spread expectation error $\widehat{E[CSE]}_t$ for the entire time span by comparing it to the realized BAA credit spread.

3.3.2 Predicting Macroeconomic Indicators

Similar to López-Salido et al. (2017) who use the expected return to bearing credit risk to predict macroeconomic indicators, I follow their methodology to predict changes in GDP, unemployment, and domestic investment using the changes in predicted credit spread expectation error as follows:

$$\Delta y_{t+h} = \beta_0 + \beta_1 \Delta\widehat{E[CSE]}_t + \gamma' \mathbf{x_{t-1}} + v_t \tag{3.4}$$

where Δy_{t+h} will be, in different variants, the log-difference of real GDP per capita, the change in unemployment rates, and the log-difference in domestic investment, all over the course of quarter t to horizon h.[5] $\Delta\widehat{E[CSE]}_t$ is the change of the predicted credit spread expectation error from quarter $t-1$ to t, and the controls $\mathbf{x_{t-1}}$ include the credit spread in quarter $t-1$, the log-difference of each predicted indicator y_t from quarter $t-2$ to $t-1$, the CPI inflation rate in quarter $t-1$, and the changes in both the

[4]Lasso stands for Least Absolute Shrinkage and Selection Operator. This approach works well with my assumption that there are a select number of significant variables that influence credit spread expectations while the rest are close to zero.

[5]As per Subsection 3.2.1, the horizons are defined as $h = 0, 1, 2, 3, 4$

3-month and 10-year Treasury yields from quarter $t-2$ to $t-1$. Newey-West standard errors are estimated with the Newey and West (1994) automatic lag-selection that allows correction for heteroskedasticity and autocorrelation.

This regression approach is merely predictive and cannot pinpoint the casual relationship between the expectation error and the macroeconomic indicators. Instead, we can interpret the coefficients as those that will predict a certain path for the indicator in the future. For example, a negative and statistically significant coefficient on the change in predicted credit spread expectation error is predicting that a positive expectation error can predict negative GDP growth in the corresponding time horizon.

3.4 Results

I break down my results into first a discussion about the textual factors and the validity of how I fill in the historical values for credit spread expectation error, and then the results from my predictive regressions for macroeconomic indicators.

The time series of the actual and predicted errors are in Figure 3.4. As a comparison of the strength of my textual factors for BAA credit spread expectation errors, I take the Bordalo et al. (2018) test as a basis to see how the model with all historical data fits into the finding that current spreads overly predict expectation error (Column (3) in Table 3.1, panel A, where $E[CSE]_t = \beta_0 + \beta_1[CS]_t + \epsilon_t$). When filling in the historical expectation error for credit spreads calculated with the textual factors, I find that the current credit spread is a significant contributor to the error with a coefficient of 0.456 (standard error of 0.079). In comparison, without the textual factors, the current credit spread is a significant contributor to the error with a coefficient of 0.564 (standard error of 0.093). Since this was a machine learning exercise, I look to the in-sample and out of sample R^2 as a means to express variation through time. To do this, I follow the standard machine learning literature and conduct a k-fold cross-validation to estimate the skill my model had on the unobserved credit spread expectation error. By convention, I leave $k = 10$ and split my sample into 10 groups and find that the in-sample R^2 without textual factors is 0.57 but with the factors it jumps to 0.69, a 21.5% increase. The more important test is the out-of-sample R^2 as it will show

how the model can predict data it has not yet seen. In this case, my out-of-sample R^2 without textual factors is 0.41 but with the factors it jumps to 0.52, a 26.8% increase!

I present the results from the predictive regression for growth of Real GDP in Table 3.2. The columns each signify which quarter horizon is being calculated, with $h = 0$ meaning the current forecast, or nowcast, and each subsequent column denoting that many quarters ahead up to four quarters, or one full year, ahead. I find that the change in my predicted credit spread expectation error is statistically significant at the one year horizon, with a negative predictor that implies a one standard deviation jump in expectation error is associated with a predicted real GDP growth decline of 0.03 percentage points. This corroborates the finding in Bordalo et al. (2018) that an increase in expectation error (when the forecasts of credit spreads get higher than the actual realized values, or when sentiment is overly optimistic) predates a decline in economic activity. Similar to López-Salido et al. (2017), I also find that the most recent credit spread and changes in the 10-Year Treasury Yield positively impact real GDP growth at the one year horizon.

I present the results from the predictive regression for changes in the unemployment rate in Table 3.3. The columns each signify which quarter horizon is being calculated, with $h = 0$ meaning the current forecast, or nowcast, and each subsequent column denoting that many quarters ahead up to four quarters, or one full year, ahead. I find that the change in my predicted credit spread expectation error is statistically significant in the nowcast as well as throughout the quarter forecasts up until the one year ahead, all positively significant. In other words, when the expectation error continues to increase (when sentiment is elevated), there is an associated increase in the unemployment rate. With magnitudes, a one-standard deviation jump in the credit spread expectation error leads to an average of 1.68% increase in the unemployment rate over the next year (an average calculated using the coefficients associated with the variable for the nowcast and the subsequent h quarters ahead). The story complements the one told in Table 3.2. I find that the most recent spread has a negative association with a one standard deviation in the change from the most recent spread leading to a -0.51% drop in the unemployment rate. No other controls have statistically significant predictive power across the horizons.

Lastly, I present the results from the predictive regression for growth in private domestic investment in Table 3.4. The columns each signify which quarter horizon is being calculated, with $h = 0$ meaning the current forecast, or nowcast, and each subsequent column denoting that many quarters ahead up to four quarters, or one full year, ahead. I find that the change in my predicted credit spread expectation error is statistically significant for the one year horizon, suggesting that a one-standard deviation jump in credit spread expectation error predicts a 2.9% decline in one year. Similarly, the coefficients for the most recent credit spread show a negative association with the domestic investment growth. This result suggests that changes in private domestic investment respond more strongly following periods associated with overly optimistic sentiment in the market. No other controls have statistically significant predictive power across the horizons. In all, the three aforementioned results are consistent with overly optimistic sentiment in the credit market predating declines in economic activity.

The results here corroborate the story of the behavioral models described in the introduction where there are systematic biases in expectations and whereby errors to the credit spread expectations of the future are seemingly correlated with macroeconomic indicators. However, there is little that my textual factors approach can say regarding the different channels of transmission that these errors are playing into. That said, the historical filling in of errors in credit spread expectations through time remain consistent with the shorter expectations data we do have in the recent decades.

3.5 Conclusion

In this paper, I analyze how credit spread expectation errors can predict future macroeconomic indicators through different time horizons. To do so, I focus on the BAA credit spread which provides useful information about investors' perceptions of risk and the overall health of the economy. Current research looking at the interplay between credit markets and the macroeconomy provides two approaches where one focuses on shock amplification that affects underlying fundamental factors and the other emphasizes the role of beliefs, such as excessive optimism, as a reason for business cycles. Unifying the two approaches may be timely as the economy presently moves into a decade that has been marked so far by a systemic shock (COVID-19) and largely global events such as war and investor uncertainty. To do so, it is a timely endeavour to find empirical evidence of how this interplay works since even policy makers agree on how credit spreads affect the market and should be interwoven in policy. The difficulty arises in acquiring historical data that can be used in economic analysis for periods that cover more than just the recent two decades.

This analysis makes usage of the Survey of Professional Forecasters to first motivate the systematic biases narrative from 2010 through 2022, and then uses the data to train a machine learning model that uses textual analysis on Wall Street Journal title pages, such as in Figure 3.6, to derive textual factors that I use to calculate errors in credit spread expectations. The resulting series, which I think of as a stand in for sentiment, is then used in predictive regressions to forecast macroeconomic indicators on a quarterly basis. I find that increases in the errors for credit spread expectations, associated with over optimism in the credit market, generally predict downturns in economic activity.

These findings, further confirming predictable patterns in credit spread expectation errors, suggest real implications for policy makers who may try to abate economic activity declines by incorporating credit market analysis in their considerations. Additionally, the proxy for sentiment here is derived from a popular news source for market participants but speaks little to the drivers of sentiment and how it manifests into actual market behavior. To do so would require more work than the current analysis but proves useful guidance if we can hope to provide quantitative advice to policy makers.

Table 3.1. [Motivation] Predictability Tests on Credit Spreads (Actual, Forecast, and Error)

	Actual Spread (1)	Forecast Spread (2)	Error (Actual - Forecast) (3)
Current Spread	0.32*	0.51***	-0.31***
	(0.19)	(0.11)	(0.08)
Constant	1.16**	1.23**	1.82***
	(0.49)	(0.51)	(0.37)
Total Observations	51	51	51
R^2	0.64	0.54	0.26

Notes: Quarterly time series regressions following Table I in Bordalo et al. (2018). In columns (1) - (3), the independent variable is the actual credit spread averaged over quarters $t-4$ to $t-1$ prior to the forecast given in quarter t. Then, (1) is the actual credit spread averaged between the Q1 and Q4 forecasts ahead of survey date (Actual average between $t+1$ to $t+4$), (2) is the forecasts of credit spreads averaged between Q1 and Q4 forecasts ahead of survey date (Forecast average between $t+1$ to $t+4$), and (3) is the forecast error (actual minus forecast) of credit spreads. Credit spread forecasts are the consensus forecasts computed from the Survey of Professional Forecasters spanning 2010Q1 to 2022Q3. Newest-West standard errors, with the automatic bandwidth selection from Newey and West (1994), are in parentheses.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.001$

Table 3.2. [Results] Predictive Regressions for Log-Difference of Real GDP (Growth)

	Real GDP Growth (h quarters ahead)				
	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$
	(1)	(2)	(3)	(4)	(5)
Δ Predicted Expectation Error	-0.010	-0.018	-0.019	-0.016	-0.029***
	(0.006)	(0.013)	(0.012)	(0.011)	(0.010)
Most Recent Credit Spread	0.006	0.012	0.008	0.011	0.014**
	(0.004)	(0.009)	(0.012)	(0.007)	(0.006)
Most Recent log Δ GDP	0.459***	0.516***	0.613***	0.514***	0.556***
	(0.084)	(0.080)	(0.091)	(0.063)	(0.079)
Most Recent CPI	0.071	0.089	0.084	0.101	0.076
	(0.051)	(0.068)	(0.083)	(0.85)	(0.127)
Δ 3-Month Treasury Yield	-0.169	-0.097	0.107	-0.018	0.004
	(0.141)	(0.157)	(0.154)	(0.149)	(0.163)
Δ 10-Year Treasury Yield	-0.478	-0.461	-0.459	-0.513*	-0.610*
	(0.398)	(0.349)	(0.363)	(0.267)	(0.340)
R^2	0.419	0.261	0.213	0.224	0.233

Notes: Quarterly time series regressions with sample period 1948Q1 through 2022Q3. All specifications include a constant (not reported). Δ Predicted Expectation Error is the change of the predicted credit spread expectation error from quarter $t-1$ to t, Most Recent Credit Spread is the credit spread in quarter $t-1$, Most Recent log Δ GDP is the log-difference of GDT from quarter $t-2$ to $t-1$, Most Recent CPI is the CPI inflation rate in quarter $t-1$, Δ 3-Month Treasury Yield is the change from quarter $t-2$ to $t-1$, and Δ 10-Year Treasury Yield is the change from quarter $t-2$ to $t-1$. Heteroskedasticity- and autocorrelation-consistent asymptotic Newey-West standard errors are reported in parentheses and use the automatic lag selection method of Newey and West (1994).

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.3. [Results] Predictive Regressions for Changes in Unemployment Rate

	Change in Unemployment(h quarters ahead)				
	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$
	(1)	(2)	(3)	(4)	(5)
Δ Predicted Expectation Error	1.871**	1.654**	1.665**	1.598**	1.601**
	(0.903)	(0.815)	(0.820)	(0.781)	(0.766)
Most Recent Credit Spread	-0.471	-0.416	-0.443	-0.439	-0.509**
	(0.291)	(0.281)	(0.274)	(0.279)	(0.219)
Most Recent Δ Unemployment	0.761***	0.707***	0.699***	0.774***	0.797***
	(0.292)	(0.264)	(0.248)	(0.282)	(0.287)
Most Recent CPI	0.026	0.039	0.040	0.033	0.031
	(0.018)	(0.026)	(0.027)	(0.023)	(0.019)
Δ 3-Month Treasury Yield	-0.004	0.001	0.000	-0.001	-0.003
	(0.004)	(0.003)	(0.001)	(0.002)	(0.004)
Δ 10-Year Treasury Yield	-0.011	-0.009	-0.010	-0.011	-0.007
	(0.009)	(0.010)	(0.012)	(0.016)	(0.008)
R^2	0.319	0.147	0.189	0.109	0.192

Notes: Quarterly time series regressions with sample period 1948Q1 through 2022Q3. All specifications include a constant (not reported). Δ Predicted Expectation Error is the change of the predicted credit spread expectation error from quarter $t-1$ to t, Most Recent Credit Spread is the credit spread in quarter $t-1$, Most Recent Δ Unemployment is the change in unemployment rate from quarter $t-2$ to $t-1$, Most Recent CPI is the CPI inflation rate in quarter $t-1$, Δ 3-Month Treasury Yield is the change from quarter $t-2$ to $t-1$, and Δ 10-Year Treasury Yield is the change from quarter $t-2$ to $t-1$. Heteroskedasticity- and autocorrelation-consistent asymptotic Newey-West standard errors are reported in parentheses and use the automatic lag selection method of Newey and West (1994).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 3.4. [Results] Predictive Regressions for Log-Difference of Domestic Investment (Growth)

	Domestic Investment Growth (h quarters ahead)				
	$h = 0$	$h = 1$	$h = 2$	$h = 3$	$h = 4$
	(1)	(2)	(3)	(4)	(5)
Δ Predicted Expectation Error	-0.019	-0.026	-0.021	-0.024	-0.029***
	(0.014)	(0.019)	(0.016)	(0.015)	(0.011)
Most Recent Credit Spread	-0.014**	-0.016**	-0.011**	-0.018**	-0.020**
	(0.007)	(0.007)	(0.005)	(0.008)	(0.010)
Most Recent log Δ Investment	0.185**	0.197*	0.251*	0.213*	0.194*
	(0.082)	(0.116)	(0.137)	(0.112)	(0.104)
Most Recent CPI	0.017	0.009	0.007	0.016	0.004
	(0.010)	(0.006)	(0.009)	(0.012)	(0.008)
Δ 3-Month Treasury Yield	0.008	0.006	0.003	0.001	0.005
	(0.012)	(0.005)	(0.009)	(0.011)	(0.012)
Δ 10-Year Treasury Yield	0.0012	0.002	0.005	0.013	0.007
	(0.013)	(0.011)	(0.015)	(0.012)	(0.017)
R^2	0.213	0.189	0.181	0.193	0.204

Notes: Quarterly time series regressions with sample period 1948Q1 through 2022Q3. All specifications include a constant (not reported). Δ Predicted Expectation Error is the change of the predicted credit spread expectation error from quarter $t-1$ to t, Most Recent Credit Spread is the credit spread in quarter $t-1$, Most Recent log Δ Investment is the log-difference of private domestic investment from quarter $t-2$ to $t-1$, Most Recent CPI is the CPI inflation rate in quarter $t-1$, Δ 3-Month Treasury Yield is the change from quarter $t-2$ to $t-1$, and Δ 10-Year Treasury Yield is the change from quarter $t-2$ to $t-1$. Heteroskedasticity- and autocorrelation-consistent asymptotic Newey-West standard errors are reported in parentheses and use the automatic lag selection method of Newey and West (1994).
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Figure 3.1: Predictable errors in forecasts of credit spreads.

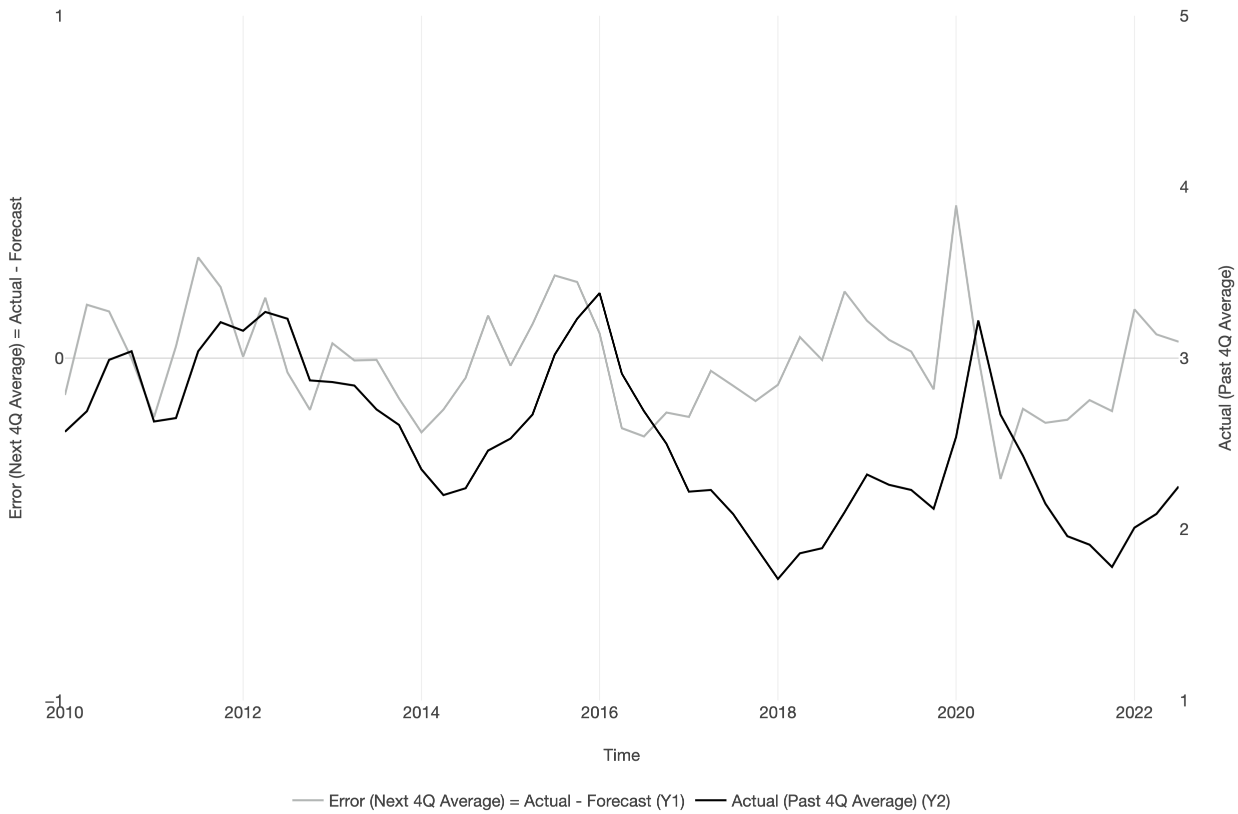

Notes: Quarterly time series plot following Figure 1 in Bordalo et al. (2018). In each quarter t, the gray line shows credit spread expectation errors (actual minus forecast) averaged over quarters $t+1$ to $t+4$ (left scale), and the black line shows the actual credit spread average over quarters $t-4$ to $t-1$, where $t-1$ is the latest quarterly credit spread prior to the forecast (right scale). Credit spread forecasts are the consensus forecasts computed from the Survey of Professional Forecasters spanning 2010Q1 to 2022Q3.

Figure 3.2: Method Diagram to obtain Textual Factor Loadings

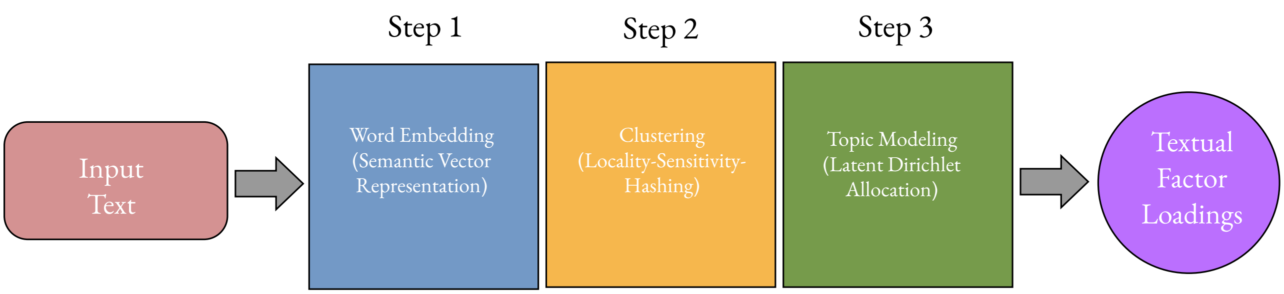

Notes: Diagram depicting the process of obtaining the Textual Factor Loadings that are used in Equation 3.2. Input Text is coming from the digital versions of newspapers obtained through the OCR process. Step 1 transforms the text into vector representations (text embeddings) using the *word2vec* function. Step 2 clusters these vector forms of words in a way that reduces dimensionality. Step 3 creates topics from the clustered word vectors. The result, the textual factors, allow me to create loadings on various topics that are fed into predictive regressions in Section 3.3.

Figure 3.3: Accuracy check on Constructed Variable Data following the Textual Factors Model

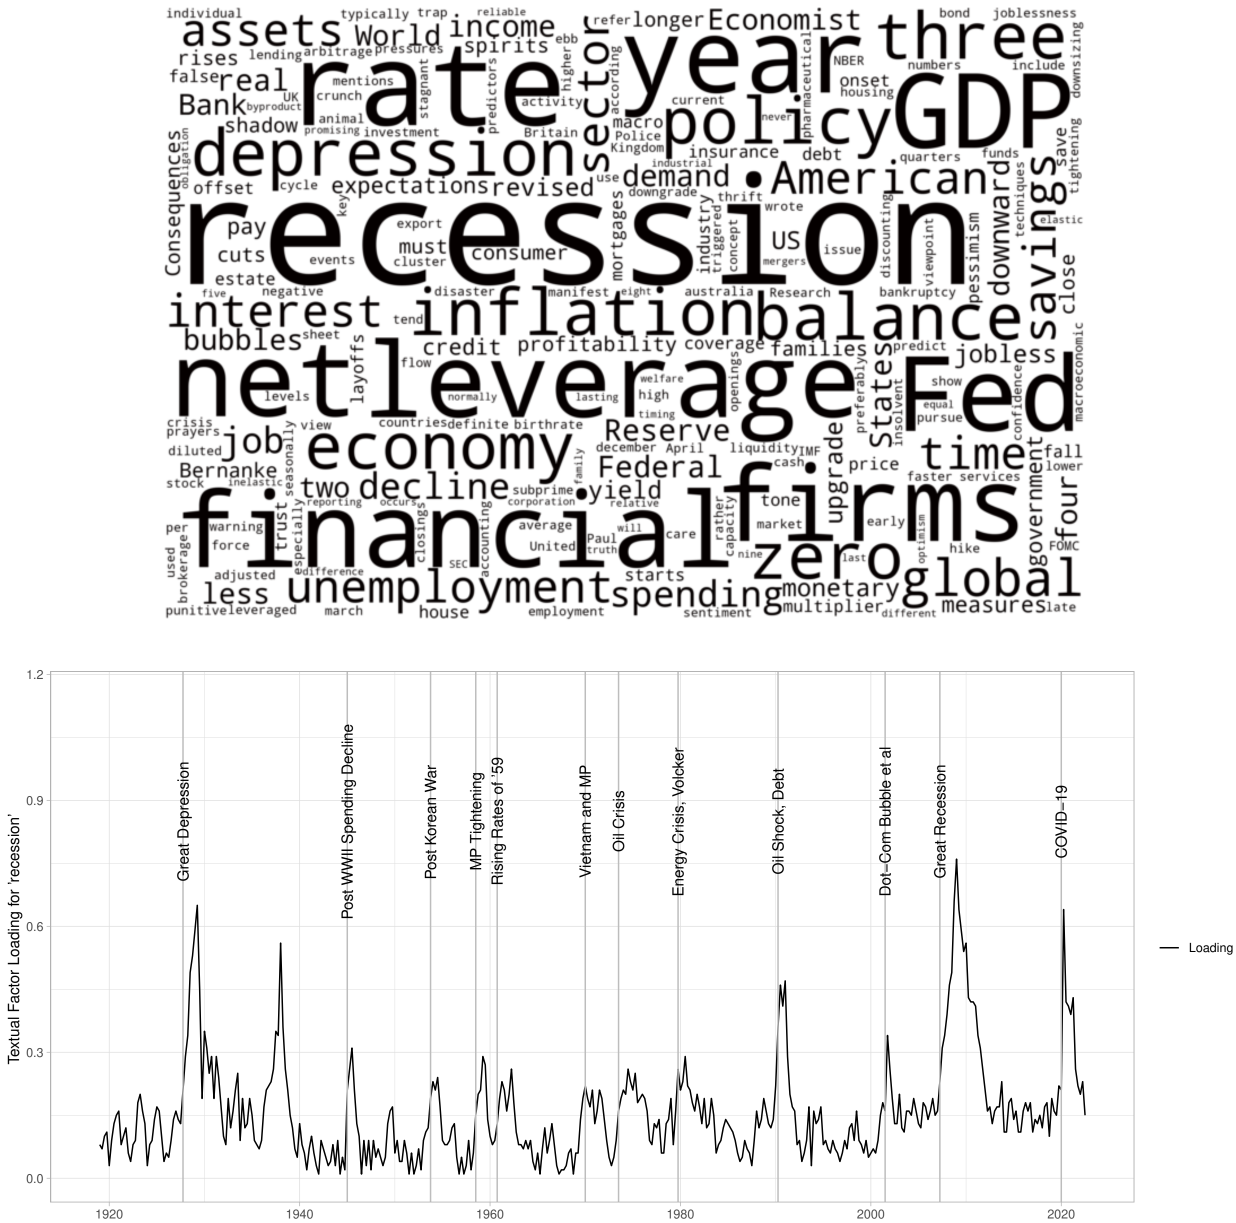

Notes: (a) Sample cluster based on the computational linguistics approach in Section 3.2 for the word ”recession”. (b) Loadings on the textual factors for the word ”recession” through time using the Wall Street Journal front page data as the input source data for the algorithm. Data processing follows the filtering methodology of Manela and Moreira (2017) to leave out common phrases and stop-words including, but not limited to, ”an”, ”the”, ”it”.

Figure 3.4: Comparison between Actual Credit Spread Expectation Error and Predicted Errors using Textual Factors, 2010Q1 - 2022Q3.

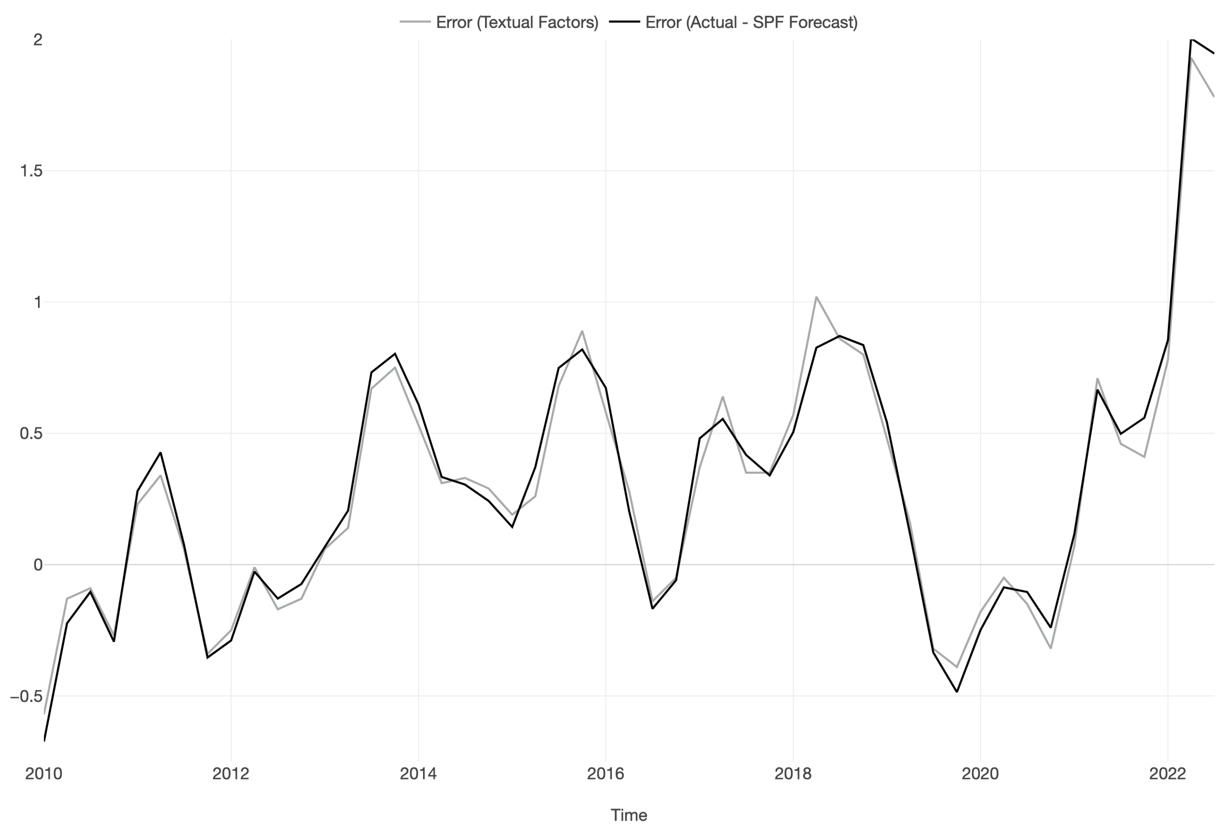

Notes: Quarterly time series plot. The gray line shows credit spread expectation errors (actual minus forecast) from the consensus forecasts computed from the Survey of Professional Forecasters. The black line shows the predicted credit spread expectation errors calculated using the textual factors.

Figure 3.5: Predicted Credit Spread Expectation Error, 1919Q1 - 2022Q3.

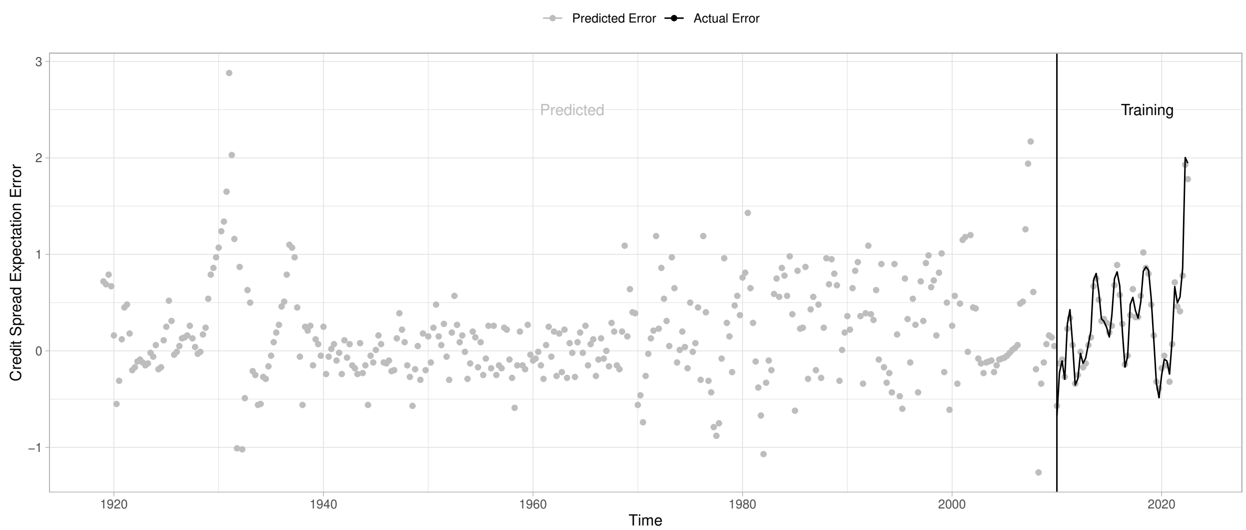

Notes: Solid line is the realized expectation error between SPF consensus forecasts . Dots are credit spread expectations error as derived from the penalized Lasso regressions in Equation 3.3 that derive the expected BAA credit spread expectations from textual factor loadings and are used to calculate the error relative to historical 10-Year Treasury Yields. The Training sub-sample, 2010Q1 to 2022Q3, is used to estimate the dependency between the textual factors and implied expectation error. The Predicted sub-sample includes all earlier observations for which BAA Credit Spread Expectations and, hence, Credit Spread Expectations Error, are not available.

Figure 3.6: Sample document of the Wall Street Journal Front Page; November 9, 1988.

Notes: Throughout its publication, the WSJ has had a number of fixed entities on the front page. For example, the "What's News—" section has either been in the middle-left or left-most column of the front page and dedicates brief snippets to news topics. Over time, graphics became more prominent on the front page, cutting the amount of text as well.

4

Household Sentiment Analysis through a Hierarchical Bayesian Latent Class Model

In recent years, the use of survey-based expectations data has become increasingly common in economic research. These surveys are designed to gather information about the expectations of different economic agents, such as private households, firms, and professional forecasters. The data collected from these surveys can be used to analyze the properties of expectations and their impact on economic decision-making and efficacy of central bank communication (for recent examples, see Coibion et al. (2020a, 2022); Arteaga (2022b); Armantier et al. (2022b); Weber (2022)). Heterogeneity in individual beliefs, or expectations, is believed to contribute significantly to the observed differences. However, beliefs are not easily observed directly which has to an over reliance on the point and density forecasts of respondents, provided the survey is designed to allow such solicitation. Outside of those questions, a significant portion of surveys are comprised of categorical questions, such as the Michigan Survey of Consumers wherein over 75% of the questions are categorical. Despite this, methods designed for analyzing categorical data on beliefs are quite limited and often provide an aggregate summary without modeling any heterogeneity, a critique exemplified in Manski (2004) and Pesaran and Weale (2006). The former goes as far to suggest that the observed heterogeneity between individuals can be related to differences in information processing. This processing difference may be key to understanding why beliefs, and subsequent economic outcomes, vary across individuals. Various studies aim to understand the interplay in observed differences by tying expectations to differences in demographic and personal variables (see Manski (2018) for a review), but there are whole sections of gathered data that receive a cursory glance due to the limitations of categorical analysis. I suggest that there are key unobserved differences hidden in this data not currently considered by researchers or policy makers which may be useful in guiding conversations about expectations formation and in the policy making sphere.

In this paper, I estimate multiple respondent belief types from the Survey of Consumer Expectations (SCE) by modeling heterogeneity in beliefs as differences in individual information choice. I extend the research on household sentiment by using Latent Dirichlet Analysis for Survey Data (LDA-S), a hierarchical Bayesian statistical model, to operationalize belief types as latent classes. The LDA-S approach takes the categorical survey questions in the SCE and allows for an economic interpretation to the unobserved heterogeneity, providing a variety of useful results including the probability that a household observed at time t belongs to a certain belief type and the most likely response a household would to a given question conditional on when they took the survey. These latent belief types can be considered auxiliary information about households through belief type probabilities that reduce the dimensionality of including categorical responses in economic analysis.

In a thorough survey of recent advances in central bank communication with the public, Blinder et al. (2022) argue that socioeconomic backgrounds are relevant factors to understanding the general public as groups, in so much that banks should tailor their focus on the groups with the lowest levels of knowledge (a *'most common denominator'*). They further cite evidence that lower level groups are most heavily influenced by the media, such as television and newspapers, and note that headline news garners more public attention than central bankers ever do. This mismatch between a public that has mixed understanding about where to acquire news and a communicator that does not know the best course in which to target its messaging leads to a natural desire to understand what defines this 'most common denominator'. I pose that uncovering latent groups within survey responses is an effective alternative to simple segmenting the public into their demographic backgrounds, as is commonly done in this research space. Recent studies that have tried to capture what constitutes 'trust' in central banking efforts have failed to find significant impacts from age, household incomes, or occupation, but rather find that news consumption and political ideology do (Brouwer and de Haan, 2022). I follow this finding to explore if this is the case with the SCE using a novel approach which I motivate below.

The hierarchical Bayesian latent class model I use in this analysis is also called a mixed membership model, which are often used to cluster discrete data with high dimensions in applications such as marketing and textual analysis. The basic ideas are that the data are grouped such that each group is modeled with a *mixture*. The mixture components are shared across all the members of the group, but the mixture proportions vary across groups. This explicitly assumes both homogeneity and heterogeneity; for the present analysis, I focus on the expectations heterogeneity found in the SCE. To get a clearer grasp of the intuition behind the LDA-S methodology, consider a thought experiment about information from different news sources. Every month, there are multiple news articles that convey different sentiments about the economy but only a number of these are relevant to the survey questions in the SCE (which focus on financial well-being, inflation, credit access, etc). Some articles might have an optimistic tone while others may have a pessimistic one. This approach proposes that an individual's response to the SCE depends on the prevalence of a particular type of article at that specific time (the time-specific effect) and their own idiosyncratic preference for that type of article (the individual-specific effect). The proportion of optimistic and pessimistic news about the economy varies over time; during times of economic hardship, for instance, it is easier to access negative news and incorporate that information into beliefs.

An individual's choice of news source determines their belief type such that the model can estimate the expected responses of individuals who have absorbed different types of news sources.

I uncover three different belief types that are can be broadly defined as 'inconsistent/uncertain', 'pessimistic', and 'optimistic' during the June 2013 through April 2022 time period. These belief types are characterized by distinct response behavior patterns to the categorical questions posed in the SCE over a variety of macro and personal expectations. The 'inconsistent/uncertain' belief type is characterized by relative positive outlooks in personal expectations for household income and financial state concurrent with pessimism about higher inflation, spending, and worsening credit conditions over the same time horizon. The 'pessimistic' belief type is characterized by a supply-side (or stagflationary) interpretation to changes in macroeconomic variables with responses expecting higher inflation, lower income, and lower growth (through deteriorating credit conditions). This belief type follows the characterization found in Candia et al. (2020) for households in advanced economies. The 'optimistic' belief type closely follows many of the traits for the 'inconsistent/uncertain' belief type but is markedly different by the response behavior looking at improving credit conditions. This third belief type is also the most prevalent in the sample and is positively correlated with other popular indices of sentiment, such as the OCED Consumer Confidence Index. My results show that these belief types are strongly associated with the timing of the survey, following the reasoning that information acquisition of news sources plays a significant part in shaping expectations for households.

I then take the latent belief types and proceed with variable and model selection methods to see if they add any information to models without them. I find statistically significant relationships between the latent belief types and the 12-month ahead inflation expectations variable solicited in the SCE. Particularly, I find that the probability of a household belonging to Belief Type 2, the pessimistic one, increases their inflation expectations forecast by almost 1 percentage point; this represents almost 25% of the average inflation expectations forecast in my sample data. Conversely, I find the opposite to be true for Belief Type 3, continuing to show an overreaction to information sources, such as news, which has been heavily documented for professional forecasters such as Bordalo et al. (2022).

Outline. In the next section, I present the context that this research has in the literature in more detail. Section 4.2 lays out the survey data and econometric model for survey expectations based on information acquisition and about it can be applied to the SCE, with Section 4.3 presenting identification of belief types. Section 4.4 presents and discusses how the belief types are associated to 12-month inflation expectations. Section 4.5 concludes.

4.1 Context In Literature

This paper connects to three kinds of research areas, the first of which deals in operationalizing variable responses in the SCE to obtain more information about household expectations. Many approaches using the SCE data focus on the inflation, home price, and credit access expectation forecasts at various horizons through econometric or machine learning techniques. The SCE has been a timely innovation from the New York Federal Reserve Bank collecting expectations over macro variables such as inflation and home prices, as well as calculating their uncertainty through subjective density forecasts as detailed in Armantier et al. (2013, 2017c). A recurrent research goal has been in using these density forecasts in event studies to ask questions about the efficacy of central bank communication which, by and large, show muted effects on households (Fiore et al., 2021, 2022; Armantier et al., 2022a; Arteaga, 2022b). These studies often take the mean density forecasts for inflation at the 12 month and 24-36 month ahead horizons as a dependent variable and look for patterns within solicited responses based around windows where the Federal Open Market Committee (FOMC) announcement occurs, controlling for the demographic differences across participants. Another focus in this space is using the density forecasts for perceived risks and uncertainty. This kind of approach is marked by using the solicited expectations to quantify uncertainty in households to relate them to consumption (such as in Binder (2017c); Ryngaert (2022)) or with perceived economic risks contributing to unanchoring expectations which is yet again another consideration for central banks (such as in Ryngaert (2023)). In all of these studies, the categorical responses from the SCE are, at best, used as control variables. About one fourth of the questions in the Survey of Consumer Expectations (demographic and special add-in modules excluded) are categorical in nature yet they have received scant analysis in the existing literature that has leveraged the data. Even the New York Federal Reserve Bank's

SCE website reports the responses as simple percentages aggregated cross-sectionally through time. I expand the analysis possible to the SCE by applying a novel novel method which can summarize the set of categorical measures in an economically interpretable way, allowing me to describe households in more detail than previous studies. This allows me to separate belief types that differ only in terms of a few but important dimensions and look at the public through a different lens of characterization than just their demographic backgrounds.

The second area is that on associated belief types in central bank and policy making considerations. The use of surveys in macroeconomics has generally led to new ways of characterizing households in various spaces. van der Cruijsen and Samarina (2023) use survey data from the Eurozone to establish household pattern classifications of trust in order to gauge how effective ECB policies are in the face of competing news stories in public discourse. The information acquisition model I modify to apply to the SCE explicitly accounts for these competing news sources and classifies the public into belief types that come from this information. Breitenlechner et al. (2023) use the Michigan Survey of Consumers to distinguish households between those that expect higher inflation and express less willingness to spend on durables during low interest rates and those that do not. In contrast, I find that households across all the belief types the data uncovers act on their inflation beliefs by responding to an increase in their consumption, indicated by their household spending.

The third and last area is that on LDA applications for economic analysis. While the adaptations have been limited, advances in this space include extracting sentiment from financial statements and then analyzing asset returns as in Yue and Jing (2022), which classify statements by the level of relevance to key drivers in returns. LDA through textual analysis is also prevalent in analyzing text from FOMC announcements to study the effects of central bank transparency and measuring the degree of monetary policy surprises (Hansen et al., 2018; Shapiro and Wilson, 2019; Doh et al., 2020). Recently, LDA approaches have dealt with newspaper articles in trying to extract new measures of expectations data such as for the BOE VIX volatility index (Manela and Moreira, 2017), nonfarm payroll employment and housing starts Kelly et al. (2021), and credit spread forecast errors Arteaga (2022a). These approaches have maintained using the LDA approach for textual analysis, uncovering latent topics in documents. To uncover latent classes from survey responses, Munro and Ng (2022) extend the LDA space by using multinomial distributions to explicitly take into account categorical data, specifying prior distributions that give structure to how group membership affects information source choice and how, in turn, the information choice affects the response given in a survey. They apply this model to the Michigan Survey of Consumers to add more nuance to the published aggregate Consumer Sentiment Index and to show that including latent classes augments the heterogeneous returns to education in an extension of the Card (1993) study using the National Longitudinal Survey of Young Men. As of writing, the only other study using such an approach is Kugler et al. (2022) who use the LDA on survey data from the German National Educational Panel Study to uncover latent parenting styles and their effects on parent-style interactions and measures on cognitive skills. I take the LDA approach for survey data and modify it for usage with the Survey of Consumer Expectations (SCE) to uncover latent belief types that do not rely on the demographic information given by respondents. In doing so, I extend LDA applications for usage in survey data and in macroeconomics, particularly to obtain auxiliary information on households that can be used in guiding policy-making.

4.2 Latent Dirichlet Analysis for Survey Data

The Survey of Consumer Expectations (SCE) is a monthly gauge of household expectations that has run since June 2013. Each month, a rotational panel of individuals is drawn with each individual able to be on the panel for a total of up to 12 months. As households are phased out, new respondents for the SCE are chosen on a monthly basis from the Consumer Confidence Survey hosted by The Conference Board; these individuals are chosen so that they meet representative demographic targets similar to the ones in the American Community Survey. For this analysis, I use the latest microdata release of the SCE (from June 2013 through April 2022) and limit my sample to the cross-sectional subset of respondents when they first start off answering questions. This creates a pooled cross-section of individuals who answer the SCE questionnaire per month.

To identify belief types, I rely on $J = 10$ categorical questions about household expectations that can be broadly divided into four categories. The first deals with financial conditions of the household. A household is asked specifically how they view their own financial state at the time of survey completion versus a year prior, and how they view that state will evolve a year from survey date. The second deals with credit access and asks the same type of questions: how does the household feel about the nature by which people obtain credit (loans, credit cards, mortgages, etc) today versus a year ago from survey date, and how they think that will evolve a year from survey date. The third deals with inflation expectations and asks their belief about the probability they will experience inflation or deflation in the year ahead as well as the time period 24 - 36 months ahead. The last category deals with sub-specific beliefs of the first group: the household is asked if they believe their household income, spending, taxes paid, and home prices nationwide will increase or decrease in the next year ahead from survey date. Table 4.1 summarizes the questions (in order of appearance) and shows the response behavior of 19,025 individual households that responded to the SCE in the time period of this analysis.

For all of the questions, there is a degree of response heterogeneity where the distribution is concentrated around one answer. Given that this is an aggregated number across the entirety of the respondents, I want to analyze what best grouping to think about unobserved heterogeneity and therefore compute the p-values of Chi-Square Tests in the right panel of the table. These tests are run on the whole sample and use the household demographic data collected by the SCE to see what characteristics are contributing to the differences in response behavior. Following the categories of the SCE, and those commonly used in the literature, I look at the age, numeracy level, region, education, and income of the households. Using a significance level of $p = 0.05$, I see that different characteristics are significant for different response groupings. For example, the χ^{REG} column shows that there are substantial differences in response behaviors for the year ahead financial state beliefs (second row) between households in different regions. However, no one characteristic is the main generator of all the differences between the household responses. This conclusion has led to recommendations of focusing on multi-layered approaches to communication by policy-makers, such as in Muñoz-Murillo et al. (2020). Indeed, further analysis on the nature of dependency between these common household demographic variables considered in many studies show that only a few truly contribute to the response differences (namely age, numeracy, and income — see Appendix A.2 for more a more detailed discussion). Given this, and the nature of the pooled cross-section, I test the date factor and find that there are substantial differences in responses given at different dates (the right most column in the table). As such, I use this as my grouping variable.

The four categories of questions are important to differentiate between as basic economic intuition would help guide logical conclusions of how one response influences another. For instance, a household responding that their financial state would be better in the next year ahead could attribute this to a number of factors that create a bettering of economic conditions such as lower inflation, higher income, lower spending, easier access to credit for investments, and lower taxes. Similarly, beliefs in higher inflation ahead could be coupled with lower spending. However, without a more thorough analysis of the drivers of unobserved heterogeneity, any one story of 'logical conclusions' may be proven wrong. For example, Duca-Radu et al. (2021) use a survey covering a 17 country panel of over 2 million observations to document a 'logical contradiction' wherein consumers believing in higher inflation report their willingness for higher spending at the same time. Without seemingly clear economic intuition guiding how households think about these variables in relation with each other, I pose that the data itself can reveal belief types and their dynamics.

My main goal is to explain heterogeneity found in categorical responses in the Survey of Consumer Expectations given the time of survey completion (motivated by the differences in Table 4.1). I apply an adapted version of Latent Dirichlet Analysis (LDA) for Survey Data, a mixture of the Latent Dirichlet Allocation approach introduced by Blei et al. (2003) and the LDA-E for expectations data by Munro and Ng (2022) to connect unobserved heterogeneity with observed characteristics and survey responses. This approach explicitly acknowledges the categorical nature of the survey responses and can provide an economic interpretation of the unobserved heterogeneity therein.

Assume that a survey consists of N individual households indexed by h, and that each household belongs to one of $d_h \in \mathbb{G} = \{1, \ldots, G\}$ observable groups. In a case of a dynamic model where surveys are conducted repeatedly, even with different samples of households each time, these groups can be thought of as the time T when the surveys are collected such that this can also be written as $d_h \in \mathbb{T} = \{1, \ldots, T\}$. There are a total of J discrete survey responses in the Survey of Consumer Expectations where each question j is comprised of L_j possible responses. Households will choose their most appropriate response v

from $x_{hj} \in \mathbb{L}_j = \{1, \ldots, L_j\}$ for each question j which is dependent on the information set processed by the household. In traditional topics modeling when estimating topics in a body (corpus) of documents, the singular value decomposition of a word-document frequency matrix is notated by $\boldsymbol{Y}_D$. A probabilistic variation of this, introduced by Hofmann (1999), treats the document-specific mixtures over topics as a fixed parameter and documents as a fixed collection. Instead of using the frequency of word occurrences in a document, I analyze the frequency of responses to questions in grouped households. As such, the frequency matrix, mapped from the discrete response data $\boldsymbol{X}$, is given by $\boldsymbol{Y}_T = (Y_{G1}, \ldots, Y_{GJ})$ of dimension $G \times L$ for possible response $L = \sum_{j=1}^{J} L_j$. The model of information acquisition that follows is based on the idea of sequential choice. This means that in order to make a decision, a household goes through a series of steps where they acquire more information before ultimately making their choice. This model follows Ruiz et al. (2017) who point out that hierarchical models—which are often used to represent complex decision-making processes—can be explained using economic models of sequential choice. In other words, the way households gather and process information can be viewed as a rational economic process, even in situations where the decision-making seems more complex or nuanced.

4.2.1 An Information Acquisition Model

An individual household h chooses which of the K sources of information determines their belief type $z_h \in \mathbb{K} = \{1, \ldots, K\}$ to consume by maximizing their utility U which is based off of $\boldsymbol{u}_{g,:} \in \mathbb{R}^k$, a group affinity for the information source, and $e_{hk} \in \mathbb{R}$, an individual specific effect that allows the household to deviate from their group:

$$z_h = \arg \max_{k \in 1,\ldots,K} U_h(k) = \arg \max_{k \in 1,\ldots,K} \left(\sum_{j=1}^{K} 1(k = j)(u_{d_h,j} + e_{hj}) \right) \tag{4.1}$$

where $u_{d_h,j}$ denotes group affinity of $d_h = g$ for response $j = k$. This chosen source of information in turn determines an individual household's belief type z_h. The observed heterogeneity of a household's group affinity d_h and unobserved heterogeneity of an household's belief type are linked by a random variable π_{gk} that calculates the probability to choose information source $z_h = k$ given group affinity

$d_h = g$:

$$\pi_{gk} = \mathbb{P}(z_h = k \mid d_h = t) = \mathbb{P}\left(u_{gk} + e_{hk} = \max_{j \in \mathbb{K}}(u_{tj} + e_{hj})\right) \tag{4.2}$$

where the probability that an individual household h selects an information source k is calculated as $u_{gk} + e_{hk} - u_{gj} - e_{hj} \geq 0$ for all $j \in K$. Then, the information source k influences the response to survey question j made by the household so that it maximizes their score function for each response:

$$x_{hj} = \arg\max_{v \in 1,\ldots,L^j}\left(\sum_{u=1}^{L_j} 1(v = u)\left(q^j_{z_h,u} + s^j_{hu}\right)\right) \tag{4.3}$$

where the information source effect $\boldsymbol{q}^j_{k,:} \in \mathbb{R}^{L^j}$ is drawn independently for each k from some distribution $\mathcal{Q}$, while the individual-specific effect $s^j_{vu} \in \mathbb{R}$ is drawn independently for each h, j, v from distribution $\mathcal{S}$. Then, the probability that an individual household h with information source $z_h = k$ believes that option v is the most appropriate response to survey question j is given by the random variable β^j_{kv}, defined as:

$$\beta^j_{kv} = \mathbb{P}(x_{hj} = v \mid z_i = k) = \mathbb{P}\left(q^j_{z_h,v} + s^j_{hv} = \max_{u \in \mathbb{L}_j}(q^j_{z_h,u} + s^j_{hu})\right). \tag{4.4}$$

where $\boldsymbol{u}_{g,:}$ is independent over g with a distribution $\mathcal{F}^g_u$, and e_{hk} is independent over i and k with distribution $\mathcal{F}_e$.

I refer to Munro and Ng (2022) for the conditional independence properties that follow exactly the same here, but mention that since we neither directly observe the components in Equations 4.1 and 4.3 (namely, $\boldsymbol{u}_{g,:}$, $\boldsymbol{e}_{h,:}$, $\boldsymbol{q}^j_{k,:}$ and $\boldsymbol{s}^{j_h}$) nor their distributions, the probabilities in Equations 4.2 and 4.4, the $\pi_{g,:}$ and $\beta^j_{k,:}$, are treated as random. Furthermore, to complete the model, I assume that $\pi_{g,:}$ and $\beta^j_{k,:}$ are defined by a multinomial distribution and, following not knowing the distributions $\mathcal{S}$ and $\mathcal{Q}$ are, specify what sort of belief structures are most likely through Dirichlet priors with hyperparameters $\alpha_{g,:} \in \mathbb{R}^K$

and $\eta^j_{k,:} \in \mathbb{R}^{L_j}$. In summary, the model is defined by the following hierarchical statistical model

$$\begin{aligned} z_h \mid \pi_{d_h,:} &\sim \text{Multinomial}\left(\pi_{d_h,:}\right) \\ x_{hj} \mid \beta, z_i &\sim \text{Multinomial}\left(\beta^j_{z_h,:}\right) \\ \pi_{d_h,:} &\sim \text{Dirichlet}\left(\alpha_{d_h,:}\right) \\ \beta^j_{z_h,:} &\sim \text{Dirichlet}\left(\eta^j_{z_h,:}\right) \end{aligned}$$

where individual households $h = 1, \ldots, N$ and categorical survey responses in the SCE are indexed by $j = 1, \ldots, \mathcal{J}$ to create an $N \times \mathcal{J}$ matrix of survey response data. For each N, we further observe a set of outcomes x_{hj} for $j = 1, \ldots, \mathcal{J}$ where there exists an optimal response v. As such, the joint distribution of the model is defined as

$$p(\beta, \Pi, \boldsymbol{z}, \boldsymbol{d}, \boldsymbol{X}) = \prod_{j=1}^{\mathcal{J}} \prod_{k=1}^{K} p(\beta^j_{k,:}) \prod_{g=1}^{G} p(\pi_{g,:}) \prod_{h=1}^{N} \pi_{d_h, z_h} \prod_{j=1}^{\mathcal{J}} \beta^j_{z_h, x_{hj}}$$

The variables and some of their representation can be summarized in the following table:

Variable in the Model	**Representation**
Households in survey	N, total
Outcome Dimension	$x_{hj} \in 1, \ldots, L_j$
Frequency Matrix	$\boldsymbol{Y}_G$ (group response)
Mixture Size	G, number of groups
Outcomes per Household	$\mathcal{J} \geq 1$ responses in $x_{h,:}$
Outcome distribution	$\beta^j_{k,:}$ for x_{ij} with $z_h = k$
Latent Size	K, information sources
Optimal response	v response of h to question j
Class assignment	z_h, information/belief type via K
Membership	d_h membership of household h in group

4.2.2 Considerations for Estimation

A few considerations to think about for this model in the context of the survey data at hand. Having assignment parameter $\boldsymbol{z}$ being an $N \times 1$ vector assumes that, for each time observed, there is only one classification of belief type per observation. In other words, each household is not a mixture of belief types but rather one set belief type at that time. As such, the model allows for information source selection probabilities to vary across households, while still assuming that each household is a member of only a single belief type. This simplifies interpretation and identification, making it easier to understand the underlying patterns in the data.

The information acquisition model in the prior subsection can be estimated using Monte Carlo Markov Chain (MCMC) methods, particularly using the Gibbs Sampler which iteratively samples each variable from its conditional distribution, itself conditional on all other variables. In this specification, the survey responses are modeled as group-specific mixtures over K belief types, each characterized by the multinomial distributions over survey responses. Gibbs sampling is a method that works really well for sampling information from conditional distributions and as such is often used in Bayesian inference approaches. Each iteration comprises of three steps:

1. $[z_h \mid \boldsymbol{x}_{h,:}, \beta, \pi_{d_h,:}]$ is sampled from a multinomial distribution
2. $[\beta \mid \eta, \boldsymbol{x}, \boldsymbol{z}]$ is sampled from a Dirichlet distribution
3. $[\pi_{g,:} \mid \alpha, \boldsymbol{x}, \boldsymbol{z}]$ is sampled from a Dirichlet distribution

In each iteration, the new variables created are used immediately such that draws of z_h depend on the values of β and $\pi_{d_h,:}$ from the prior iteration, whereas β and $\pi_{g,:}$ depend on z_h from the current iteration. Given this process, any number of iterations run must take into account the initial transient period that most certainly biases the system and thus I opt to burn the 10,000 thousand iterations. In total, I conduct 50,000 iterations and base my results on the sample averages over the whole process.

To estimate the model, I need to make assumptions about the hyperparameters of the Dirichlet distributions as well as the number of belief types; in short, $\alpha_{g,:}$, $\eta^j_{k,:}$, and K must be specified. the first two hyperparameters specify prior beliefs about the importance of the group-specific terms ($\boldsymbol{u}_{g,:}$, $\boldsymbol{q}^j_{k,:}$) relative to the individual-specific ones $\boldsymbol{e}_{h,:}$, $\boldsymbol{s}^j_{h,:}$. For example, I could specify that $\alpha_{gk} < 1$ to betray a belief that households of the same observable group g are likely to choose the same information and therefore the same belief type k. Then, this implies that observed heterogeneity tightly links to unobserved heterogeneity.[1] Or I could choose to specify $\eta^j_{kv} < 1$ to betray a belief that households who choose the same information and therefore the same belief type k are likely to all respond the same way to each question; the same logic holds in reverse. Following Kugler et al. (2022), I settle on $\alpha_{gk} = 1$ for all groups g and information sources k. This is referred to as an uninformative prior, which means that it doesn't impose any strong assumptions about the relationship between group membership and information acquisition. The prior explicitly assumes that all groups have an equal chance of acquiring information from any source, as is possible when dealing with households in the United States. Following Munro and Ng (2022), I settle on $\eta^j_{kv} = 1$ for $k \neq v$, and $\eta^j_{kv} = 10$ otherwise. The choice of this prior explicitly captures the idea that each information source is associated with a correctness score on one response that is higher than any other information source for at least one question in the survey. In other words, each latent class has a strong association with one of the levels of the categorical variable. For example, in the case of the SCE, the first question is about whether the respondent thinks they (and any family living with them) are financially better or worse off than they were a year ago, with the first categorical response being 'much worse off'. Following this, I can interpret one of the belief types estimated in the MCMC procedures as a 'pessimistic' one.

Lastly, I follow Kugler et al. (2022) and choose the optimal K belief types according to the minimum of an approximated Bayesian Information Criterion (BIC). Specifically, I define $L(\hat{\theta}_k)$ be the maximum likelihood value of the data, where θ represents the set of parameters in the model. I use the posterior mean $\tilde{\theta}_k$ from the MCMC draws (the maximum likelihood value of the parameters), observations N, and the number of model parameters p_k when there are k classes in the model to consider the following BIC:

$$\widetilde{BIC}_k = -L(\tilde{\theta})_k + \frac{p_k}{2}\log(N).$$

4.2.3 Application to the SCE

In the 107 months of observations, which I use as my group variable (i.e., $G = \{1, \ldots, 107\}$), there are a total of $N = 19,025$ unique household respondents who answer the survey once in this time period.

[1] The inverse would be implied if $\alpha_{gk} < 1$; a similar logic follows for η^j_{kv}.

Of these, I focus on $\mathcal{J} = 10$ categorical survey questions, four of which have $L = 5$ possible responses and six of which have $L = 2$ possible responses. Together, the data suggests that $K = 3$.[2] The $\mathcal{J} = 10$ questions will each have a β^j associated with them such that they correspond with the probability that a household in each belief type K will select a response v for that question. I detail the questions and associated probability representation in the following table:

SCE Question	β^j_{kv}
1. Do you think you (and any family living with you) are financially better or worse off these days than 12 months ago?	β^1_{kv}
2. Do you think you (and any family living with you) will be financially better or worse off 12 months from now than you are these days?2.	β^2_{kv}
3. Compared to 12 months ago, do you think it is generally harder or easier these days for people to obtain credit or loans?	β^3_{kv}
4. And looking ahead, do you think that 12 months from now it will generally be harder or easier for people to obtain credit or loans than it is these days?	β^4_{kv}
5. Over the next 12 months, do you think there will be inflation or deflation?	β^5_{kv}
6. Over the 12-month period between 24-36 months (from survey date), do you think there will be inflation or deflation?	β^6_{kv}
7. Over the next 12 months, I expect my total household income to...	β^7_{kv}
8. Over the next 12 months, I expect my total household spending to...	β^8_{kv}
9. Twelve months from now, I expect my total taxes to...	β^9_{kv}
10. Over the next 12 months, I expect the average home price to...	β^{10}_{kv}

Questions 1 and 2 about financial conditions for the household have a scale such that response $v \in [1, 5]$ where, in order, the choices read *Much Worse Off, Somewhat Worse Off, About the Same, Somewhat Better, Better Off*.

Questions 3 and 4 about beliefs over credit accessibility have a scale such that $v \in [1, 5]$ where, in order, the choices read *Harder, Somewhat harder, Equally easy or hard, Somewhat easier, Easier*.

Questions 5 and 6 about beliefs over inflation or deflation in the next 12 and 24-36 months have a scale such that $v \in [1, 2]$ where, in order, the choices read *Inflation, Deflation (the opposite of inflation)*.

Questions 7, 8, 9, and 10 about beliefs over household income, spending, taxes paid, and home prices nationwide over the next 12 months have a scale such that $v \in [1, 2]$ where, in order, the choices read *Increase by 0% or more, Decrease by 0% or more*.

[2]The back of the envelope calculations for that are a useful barometer of the maximum K comes from ruling out under-identification, which follows $G(L - \mathcal{J}) \geq K(L - \mathcal{J}) + G(K - 1) \approx K \leq 3.89$, as proposed for LDAs by Anandkumar et al. (2015); the K chosen by BIC is 3.

4.3 Identification of Belief Types

To recap the approach in the preceeding section, the LDA for Survey Data approach imposes a structure on observable group indicators and individual household responses in the SCE by assuming that households optimally choose belief types (via their sources of information K) given their group membership to when they respond G, and optimally select responses in the SCE given their belief type. The optimal choice v from all possible responses $x_{hj} \in \mathbb{L}^j = \{1, \ldots, L^j\}$ for each question j is affected by individual effects and group commonalities first or belief type commonalities in the second case. The individual effects allow respondents to deviate from the choices usually made by other households answering in the same month or belief type.

The results from the LDA show that Belief Type 3 is chosen the most (51.5%), followed by Belief Type 1 (33%) and Belief Type 2 (15.5%). Through time, I plot the average probabilities of an observation being recorded at a given month assigned to a certain belief type (π_{gk}) in Figure 4.1. This pattern is also seen from the probability density of the observations in each belief type (or the density of z_h), which I plot in Figure 4.2. To better interpret the belief types uncovered in the data, I show the probability for an individual household with belief type $z_h = k$ to choose v as their response to survey question j, in other words $\beta^j_{k,:}$, in Figures 4.4 to 4.7. I discuss each more closely below.

In Figure 4.1, the plot shows π_{gk} through time, with an obvious preference in Belief Type 3 throughout most of the period. Spikes in Belief Type 2 appear to follow major economic disruptions such as the US Government Shutdown in late 2013 and the COVID-19 recession in the first and second quarter of 2020. There seems to be persistence in the degree of the distribution for Belief Type 2 after this latter disruption, with an average probability of Belief Type 2 occurring of 23.3% for the last 25 months of the analysis compared to the average of 13% during the preceding 81 months. Belief Type 1 also follows a similar pattern of increasing after disruptions, albeit not the persistence following a disruption exhibited by Belief Type 2. The probability density plot in Figure 4.2 contributes further insight on the relative likelihood of observations belonging to one of the belief types; Belief Type 3 clearly dominates throughout the entire period.

To give each Belief Type more meaning through an economic interpretation, I depict the probabilities $\beta^j_{k,:}$ in Figures 4.3 to 4.8. Figures 4.3 and 4.4 show the typical response behavior for Belief Type 1 across the categorical questions in the SCE. For the first question, households in this Belief Type are characterized by most likely responding with the third option, that they are *About the Same* financially as they were in the previous year ($\beta^1_{1,3} = 44.86\%$). For question 2, they are also more likely to continue to think they will be *About the Same* in the following year ($\beta^2_{1,3} = 44.88\%$). They are also more likely to respond that it is *Somewhat harder* to obtain credit than it was a year ago with a probability of $\beta^3_{1,2} = 65.2\%$, and that it will be *Somewhat harder* in the next year than it is now to do the same with a probability of $\beta^4_{1,2} = 60.1\%$. Belief Type 1 is also more likely to respond that there will be *Inflation* in the next year and in three years ahead, as well as that their income, spending, taxes paid, and home prices nationwide will all *Increase by 0% or more* in the next year, all with probabilities between 84.72% and 94.63% ($84.72\% < \beta^5_{1,1}, \beta^6_{1,1}, \beta^7_{1,1}, \beta^8_{1,1}, \beta^9_{1,1}, \beta^{10}_{1,1} < 94.63\%$).

Belief Type 2 respondents are marked by higher probabilities to respond with worse economic outcome beliefs, as shown in Figures 4.5 and 4.6. They are most likely to respond that they are *Somewhat worse off* financially than a year ago and will be *Somewhat worse off* financially in the next year versus where they are now with probabilities of $\beta^1_{2,2} = 47.85\%$ and $\beta^2_{2,2} = 46.34\%$, respectively. They are also more likely to respond about deteriorating credit conditions, with credit being *Somewhat harder* to obtain now versus a year ago (with probability $\beta^3_{2,2} = 32.86\%$) and credit being *Somewhat harder* to obtain a year from now (with probability $\beta^4_{2,2} = 39.97\%$). Belief Type 2 respondents are also likely to think there will be *Inflation* 12 and 24 - 36 months ahead (with probabilities $\beta^5_{2,1} = 87.96\%$ and $\beta^6_{2,1} = 84.79\%$), and that their household income will *Decrease by 0% or more* in the next 12 months ($\beta^7_{2,2} = 55.11\%$). Despite this, they are more likely to respond that they will see an *Increase by 0% or more* to their household spending ($\beta^8_{2,1} = 63.45\%$), the taxes they pay ($\beta^9_{2,1} = 81.53\%$), and home prices nationwide ($\beta^{10}_{2,1} = 73.39\%$).

Belief Type 3 respondent patterns are shown in Figures 4.7 and 4.8. They are most likely to respond that they are *About the same* financially than a year ago and will be *Somewhat better* off financially in the next year versus where they are now with probabilities of $\beta^1_{3,3} = 41.14\%$ and $\beta^2_{3,4} = 41.01\%$, respectively.

They are also more likely to respond that credit conditions are stable, with credit being *Equally easy or hard* to obtain now versus a year ago (with probability $\beta^3_{3,3} = 49.98\%$) and the same to obtain a year from now (with probability $\beta^4_{3,3} = 53.89\%$). Belief Type 3 is also more likely to respond that there will be *Inflation* in the next year and in three years ahead, as well as that their income, spending, taxes paid, and home prices nationwide will all *Increase by 0% or more* in the next year, all with probabilities of between 84.05% and 93.52% ($84.05\% < \beta^5_{3,1}, \beta^6_{3,1}, \beta^7_{3,1}, \beta^8_{3,1}, \beta^9_{3,1}, \beta^{10}_{3,1} < 93.52\%$).

To summarize the key differences between Belief Types, I compute the Rao distance between the probabilities, i.e. between $\beta^j_{k,:}$ and $\beta^j_{m,:}$ for $k \neq m$, to find where the these types differ most from each other, depicted in Table 4.2.[3] The table, which includes only the five biggest differences between the types, shows that beliefs over how credit conditions have evolved from the past year until now and how they will evolve into the next year are the biggest differences between Belief Type 1 and Belief Type 3. This wedge is the most prominent difference amongst all Belief Types, and the only substantial difference between the 1 and 3, showing that Belief Type 1 and 3 are similar in many aspects. Belief Type 2 and 3's differences are driven by their beliefs over income a year from now and how, it appears, it will affect them financially a year from now. Credit condition beliefs are also marked differences. The differences between Belief Type 1 and 2 mirror the ones between 2 and 3, except that household spending is now more prominent. The differences in household income a year from now, financial conditions (a year ago vs now and now vs a year from now), and beliefs over credit obtainment from a year ago til now are also less substantial between Belief Type 2 and 1 than between Belief Type 2 and 3.

Taken together, the results and differences between the Belief Types lead me to the following conclusions. Belief Type 2 is markedly the most dissimilar of the three and exhibits beliefs that trend pessimistic about economic conditions across the board. The dual response behavior narrative of expecting higher inflation and worse income in the future takes a stagflationary view; this is corroborated by the beliefs over harder credit access (i.e. deteriorating credit conditions) which would suggest slow growth. The household beliefs here mirror the behavior pattern findings of Candia et al. (2020), who find that households in advanced economies take a supply-side interpretation to changes in macroeconomic variables.

[3]The Rao distance is a measure of dissimilarity, computed as the square root of the Kullback-Leibler divergence between two probability distributions; the closer to 0, the more similar to each other. See Rao (1992) for a detailed introduction.

This type of interpretation often concludes with negative income effects, which can depress economic activity. As such, I define Belief Type 2 as 'pessimistic'; households obtain macroeconomic information that feeds their negative sentiment, reporting such beliefs over time. This explains the rising proportion of households of this Belief Type during economic disruptions.

In contrast, Belief Types 1 and 3 display a dual higher inflation and higher income belief pattern. They also exhibit a high probability of increased spending over the following year, mirroring results from van der Cruijsen and Samarina (2023) who find that European consumers with higher inflation expectations are more likely to increase their household spending. The biggest difference between the two is their belief over credit conditions, with Belief Type 1 exhibiting more pessimism. Without other marked differences, I take this inconsistent response behavior as the defining trait from Belief Type 1: households of this type believe that higher inflation and harder credit access, arguably worsening economic conditions, will not affect their financial state and respond their belief in higher income and spending. This type appears to view external macro conditions as separate from their idiosyncratic ones, and as such conclude their response behavior is 'uncertain': households obtain macroeconomic information that feeds a dual narrative of tougher conditions externally while improved financial conditions internally.

Lastly, for Belief Type 3, the consistency of their economic intuition leads me to conclude they are more 'optimistic' in their economic outlook. Despite their tendency to respond there will be higher inflation in the future, they have the strongest belief in higher income and believe credit conditions will either improve or stay the same in the near future. Belief Type 3 is also the most common in the sample and, as such, I take this type to be indicative of other sentiment indices.

In summary, I conclude that Belief Type 1 is characterized by an inconsistent, uncertain sentiment, Belief Type 2 is characterized by a broadly pessimistic sentiment, and Belief Type 3 is characterized by a broadly optimistic sentiment. To see how these beliefs are correlated with aggregate economic conditions, I plot them alongside other commonly used metrics in Figure 4.9. Subplot (a) pairs Belief Type 1 with the index of Monetary Policy Uncertainty (MPU) conceived by Husted et al. (2020). This re-scaled index is a news-based index of monetary policy uncertainty that captures the degree of uncertainty that the public perceives about Federal Reserve policy actions and their consequences. It explicitly bridges the periods of conventional and unconventional monetary policy making, apt for my sample period. While there are some similarities, the Kendall's rank correlation is only moderately negative ($\tau = -0.254$), implying that this uncertainty sentiment is not wholly being explained by uncertainty in monetary policy. For the sections that are correlated, lowered monetary policy uncertainty is associated with a higher probability that a household will respond with behaviors marked in Belief Type 1.

Subplot (b) pairs Belief Type 2 with a re-scaled index of the Unemployment Rate as per the Federal Reserve's Economic Database. The Kendall's rank correlation is decently positive ($\tau = 0.649$) and the similarities in their co-movements are visually intuitive. This supports the 'pessimistic' characterization of Belief Type 2 and further purports that households of this type follow popular macroeconomic indicators and use them in their expectations setting.

Subplot (c) pairs Belief Type 3 with the re-scaled OCED Consumer Confidence Index for the United States, one of the widely used measures that tracks consumer confidence. This indicator provides an indication of future developments of households' consumption and saving, based upon answers regarding their expected financial situation, their sentiment about the general economic situation, unemployment and capability of savings. Belief Type 3 has a strong positive correlation with this index ($\tau = 0.724$) and also exhibits persistence even after declines (as seen from the period following the COVID-19 recession).

4.4 Heterogenous Beliefs and Inflation Expectations

As a means to analyze an application of the belief types from Section 4.3, I estimate a simple static model using the pooled cross-section observations from the SCE. I define the dependent variable as the 12-months ahead inflation expectations solicited through the density forecast method. To motivate my analysis, I will use a model that includes all the prominent SCE demographic variables typically treated as independent features (Fiore et al., 2021; Arteaga, 2022b; Ryngaert, 2022), and add in the three belief types so that I can perform variable selection. I use the best subset, forward stepwise, backward stepwise selection, and 10-Fold Cross Validation methods for selection and estimation of the parameters in my model. In total, I am using 19,025 observations.

My simple static model is based on the demographic variables included in the SCE that are the respondent's age, gender, marriage status, whether they identity as Latino/Hispanic, level of household income, level of education, numeracy (defined in the SCE as 'high' or 'low' as a degree of their maths ability), and the region they live in, all in a control vector $\boldsymbol{X}$. The three belief types added in make this be a total of an eleven variable model, of which I perform the aforementioned variable selection methods. In short, I use the following model: $Y_t = \beta_d \boldsymbol{X}_t + \beta_1 B1 + \beta_2 B2 + \beta_3 B3 + \epsilon$ to study how different subsets of variables in the model perform relative to each other, thereby exploring whether the belief types are informative. I compare the models based off their in sample test error given that the model with the highest number of features will always have the highest R^2; this makes that measure be a poor estimate for comparing the best model among a collection of models with different numbers of features. To compare the models, instead, I use the Adjusted R^2, Cp, and BIC measures, as well as the test mean square error (MSE) obtained from the 10-Fold Cross Validation, all found in the Appendix A.2.1.

Figures A.3 through A.5 show the results of the variable selection methods. In Figure A.3, the adjusted R^2 and Cp approaches continually selects a 10 variable model throughout the best subset, forward stepwise, and backward stepwise methods, though the highest adjusted R^2 and lowest Cp are relatively similar across the 8 through 10 variable models. The BIC selects between a 4 and 6 variable model throughout the best subset, forward stepwise, and backward stepwise methods, with relatively similar lowest BICs between the 4 through 8 variable models. In Figure A.4, I plot the variables that are selected each time. I find that the control for identifying as Latino or Hispanic (Q34), Married (Q38), and most of the Regions are not selected in a majority of the specifications. To further whittle down the variables, I perform 10-Fold Cross Validation and plot the mean squared errors from the cross validation exercise relative to the number of variables selected in Figure A.5; the plot shows that the lowest error is obtained with an eight or ten variable model, which is consistent with the previous findings. Guided by parsimony, I opt to stick with the eight variable model in the rest of the analysis. Those variables are the demographic variables of respondent's age, gender, education, household income, numeracy, and the probabilities of belonging to Belief 1, Belief 2, and Belief 3. With an optimal eight variable model, I also run the model on the demographic variables only and then compare the MSE with the model with the added belief types. I compare the MSE using the validation set and 10-Fold Cross Validation approach. For the validation set approach, the demographics only model has an MSE of 31.157 while the model with all eight variables has an MSE of 30.796. Similarly, the 10-Fold Cross Validation MSE for the demographics only model is 31.144 while the eight variable one has an MSE of 30.811. Both tests support the eight variable model and I proceed as such. As another test for the eight variable model, I estimate the MSE for out of sample validation and do so by splitting up my observations into training (80%) and testing (20%) sets. A lower MSE value will indicate better model performance but if the out-of-sample MSE is significantly higher than the in-sample MSE's found above, the model may be over-fitting. I find that the out-of-sample MSE is 28.761, which makes me feel confident in my decision with this modeling.

Lastly, I present the results from the optimal linear model in Table 4.3 with only the significant coefficients showing.[4] The first column shows the results for the linear model without the belief types added in, showing the significance of respondent age, gender, household income, and education levels, all significant at the 1%. The level of numeracy in this specification is not significant at any of the specified levels. The second column shows the results with the belief types added in. Here, the significance of the demographics relatively stay the same but the low numeracy level jumps to be significant at the 5% level. Additionally, Belief Types 1 and 2 are also significant at the 1% levels. In the third column, I show the results for a varying coefficient model wherein I allow the relationship between the belief types and 12-month inflation expectations to vary smoothly over the dates in the analysis; this model can be represented by $Y = \beta_d \boldsymbol{X} + \sum_{j=1}^{B} B_j \beta_j(T) + \epsilon$, where the coefficients of the $B = 1, 2, 3$ Belief Types, β_j, are allowed to change smoothly with the date T.

My main results are that the inclusion of the Belief Types in the regression are not only producing significant relationships with the 12-month inflation expectations variable, but also increase the adjusted R^2 by nearly twice than the model without (12.9% vs 23.5%). More specifically, I find that Belief Type 2, the pessimistic one, generates higher estimates for inflation. In other words, households whose responses correspond to the Belief Type 2 profile are going to respond with higher estimates of inflation than their other belief type counterparts. This holds true for the baseline model with beliefs plus the varying coef-

[4] I also ran a LASSO specification but found a larger MSE for that model, 31.78, than the current linear specification. Without enough evidence supporting this approach, I opt for the linear model I show in the analysis.

ficient model. I find that the increase in probability for a household to belong in Belief Type 2 increases the average 12-month inflation expectations by 0.934 percentage points, corresponding to 22.7% of the average 12-month inflation expectations forecast given by respondents (about 4.2%). This coefficient skyrockets under the varying coefficient model to 3.64, corresponding to over 85% of the average 12-month inflation expectations forecast in the period. Conversely, I find that the Belief Type 3 probability, the optimistic one, leads to lower inflation expectation forecasts also corresponding to about 25% of the average given in the sample.

Together, my takeaways are that households belonging to different belief types have a statistically significant relationship to influence the respondent 12-month inflation expectation. This highlights the need for central bank communication to take into consideration the type of information households are consuming. In other words, if there is a small yet varying probability that households throughout a time period are consuming negative news and thereby fitting the Belief Type 2 profile, then those households should be targeted more heavily so that their inflation expectations are not so heavily skewed upwards.

4.5 Conclusion

In conclusion, this paper has demonstrated the importance of considering the heterogeneity in beliefs among households when analyzing the Survey of Consumer Expectations (SCE). Through the use of Latent Dirichlet Analysis for Survey Data (LDA-S), I identified three distinct belief types and characterized them as 'inconsistent/uncertain,' 'pessimistic,' and 'optimistic.' These belief types exhibit unique response patterns in their expectations about macroeconomic and personal financial conditions through the categorical questions in the SCE, indicating that households' economic expectations are shaped by the information they consume.

I further show that these belief types are economically significant when predicting inflation expectations. The results show that incorporating belief types in the analysis significantly improves the explanatory power of models predicting households' 12-month inflation expectations. The inclusion of belief types almost doubled the adjusted R-squared in the linear model and, moreover, households with pessimistic beliefs (Belief Type 2) are more likely to have higher inflation expectations, while those with optimistic beliefs (Belief Type 3) tend to have lower inflation expectations. As different belief types are shown to have a statistically significant impact on respondents' 12-month inflation expectations, it becomes crucial for central banks to consider the type of information households are consuming and tailor their communication accordingly. For instance, households fitting the Belief Type 2 profile, characterized by pessimistic views, may require more targeted communication to prevent their inflation expectations from being heavily skewed upwards.

To further advance the understanding of household expectations and their impact on economic conditions, future research could explore how the relationship between belief types and inflation expectations evolves over time, as well as the effect of different macroeconomic shocks on households' beliefs. Additionally, examining the role of central bank communication in shaping these beliefs and finding ways to target households with specific belief types could lead to more effective policy interventions. Ultimately, understanding the nature of these belief types and their impact on economic behavior can lead to more effective policy measures and better-targeted communication strategies.

Table 4.1. [Summary] Household Response Behaviors (Distributions, June 2013 - April 2022)

	Much worse	Somewhat worse	Same	Somewhat better	Much better	χ^{AGE}	χ^{NUM}	χ^{REG}	χ^{EDU}	χ^{INC}	χ^{DATE}
Financially better or worse vs 12mo ago?	0.05	0.21	0.40	0.27	0.06	0.01	0.05	0.07	0.06	0.00	0.00
Financially better or worse 12mo ahead?	0.03	0.15	0.40	0.35	0.08	0.00	0.08	0.03	0.05	0.14	0.00
	Much harder	**Somewhat harder**	**Equal**	**Somewhat easier**	**Much easier**						
Obtaining credit now vs 12mo ago?	0.11	0.29	0.35	0.21	0.04	0.06	0.02	0.11	0.07	0.03	0.00
Obtaining credit 12mo from now?	0.09	0.31	0.39	0.19	0.03	0.08	0.03	0.06	0.05	0.01	0.00
	Inflation		**Deflation**								
Inflation or deflation 12mo ahead?	0.89		0.11			0.05	0.00	0.07	0.06	0.09	0.00
Inflation or deflation 24-36mo ahead?	0.87		0.13			0.09	0.02	0.08	0.08	0.08	0.00
	Increase by 0% or more		**Decrease by 0% or more**								
Household income 12mo ahead?	0.86		0.14			0.01	0.04	0.17	0.13	0.08	0.00
Household spending 12mo ahead?	0.81		0.19			0.05	0.02	0.28	0.17	0.07	0.00
Taxes 12mo ahead?	0.86		0.14			0.16	0.07	0.11	0.11	0.06	0.00
Home prices 12mo ahead?	0.85		0.15			0.03	0.14	0.09	0.19	0.07	0.00

Notes: This table summarizes the measures on responses to the main categorical questions in the SCE. The panel in the middle depics the survey responses. The right panel shows the p-values of a Chi-Square Test with the null hypothesis assuming the independence between each solicitation and a number of household characteristics. The first, χ^{AGE}, is the indicator of age groups as detailed in the SCE (below 40, between 40 and 60, and over 60); the second, χ^{NUM}, is the indicator of high or low numeracy that a household exhibits (understanding of basic economics and mathematics skills as tested by the SCE module); the third, χ^{REG}, is an indicator for region location of the household (Midwest, NorthEast, South, West); the fourth, χ^{EDU}, is an indicator for household head education level (high school, some college, college); the fifth, χ^{INC}, is an indicator for household income (below \$50K, between \$50K and \$100k, and over \$100k); the sixth, χ^{DATE}, is an indicator for date of the household responses which is what I ultimately use as the group indicator in the analysis.

Table 4.2. Largest Differences between Belief Types (Rao Distance between $\beta^j_{k,:}$ and $\beta^j_{m,:}$ for $k \neq m$)

	Type 2	**Type 3**
Type 1	Income higher or lower a year from now (0.70) Financially better or worse a year from now (0.52) Financially better or worse than a year ago (0.42) Credit easier or harder to obtain than a year ago (0.37) Spending higher or lower a year from now (0.30)	Credit easier or harder to obtain than a year ago (0.83) Credit easier or harder to obtain a year from now (0.70) Financially better or worse than a year ago (0.15) Financially better or worse a year from now (0.07) Home prices higher or lower a year from now (0.05)
Type 2		Income higher or lower a year from now (0.69) Financially better or worse than a year ago (0.55) Financially better or worse a year from now (0.55) Credit easier or harder to obtain a year from now (0.54) Credit easier or harder to obtain a year ago (0.52)

Notes: This table summarizes the five biggest differences between each of the Belief Types uncovered from the survey responses. These differences are computed by using the Rao Distance (Rao, 1992) and can be thought of measures of dissimilarity. The score next to each of the questions where the types differ are in order of most dissimilar to least dissimilar. The Rao Distance is read the same way, with a value of 1 meaning strongly dissimilar and a value of 0 being not dissimilar at all. This table shows that Belief Types 1 and 3 are mostly similar with their marked difference coming from their beliefs over credit conditions. Belief Type 2 has moderate to strong dissimilarities with the other two.

Table 4.3. [Results] Belief Types and Inflation Expectations

	Baseline (1)	Belief (2)	Varying (3)
Age	0.011***	0.018***	0.012***
	(0.002)	(0.003)	(0.002)
Gender	-0.894***	-0.891***	-0.899***
	(0.087)	(0.081)	(0.086)
Education	-0.232***	-0.231***	-0.237***
	(0.030)	(0.031)	(0.031)
Income	-0.049**	-0.063***	-0.060***
	(0.017)	(0.019)	(0.017)
Low Numeracy	-0.184	-0.234**	-0.2136**
	(0.123)	(0.099)	(0.098)
Belief 1		0.148	0.090***
		(0.107)	(0.001)
Belief 2		0.934**	3.654***
		(0.389)	(0.895)
Belief 3		-1.027**	-1.133**
		(0.470)	(0.486)
Adjusted R2	0.1293	0.2354	0.3491
Observations			19,025

Notes: This table summarizes the three specifications that I run with the linear model selected with the optimal variables through various validation methods. The dependent variable is the 12-month inflation expectations. In the first column, the specification only includes the demographics data collected from the SCE. In the second column, the Belief Types are added into the specification. The third column is a varying coefficient model where the belief types vary according to the date.
Significance: * for p<0.1, ** for p<0.05, *** for p<0.01.

Figure 4.1: Probability of Observations Across Time Belonging to Belief Types (π_{gk})

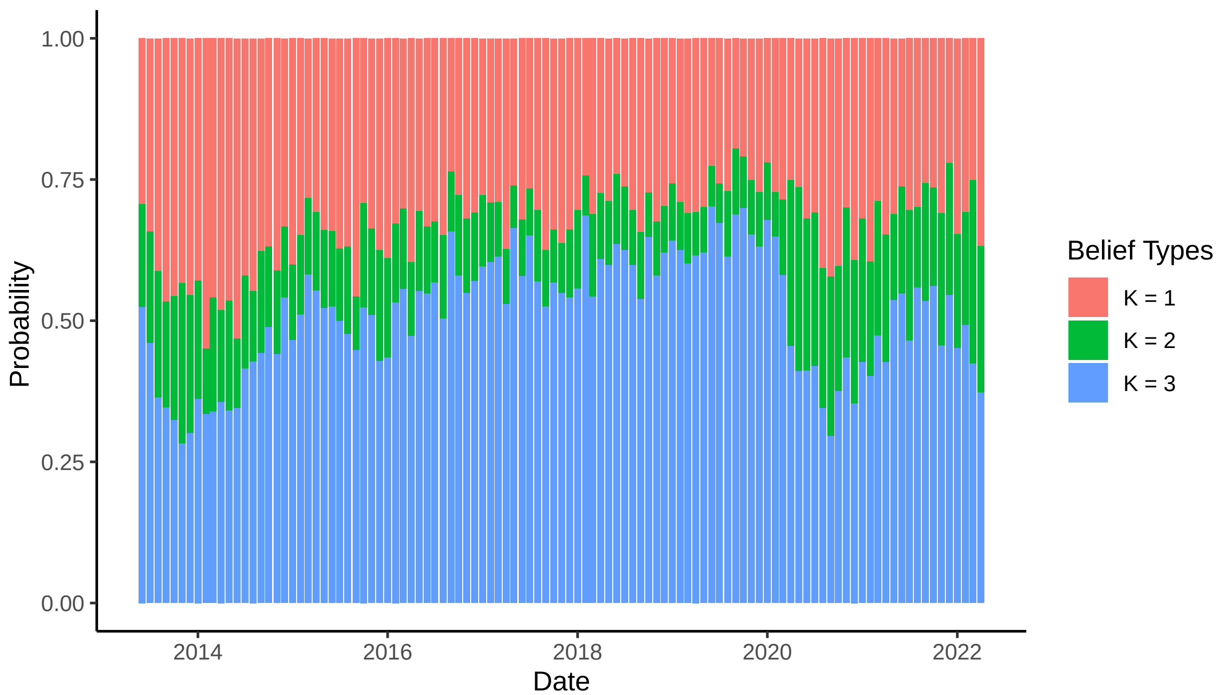

Notes: Estimated representation of the probability for household h of group g belonging to class k. On average, Belief Type 1 is chosen 33% of the time, Belief Type 2 is chosen 15.5% of the time, and Belief Type 3 is chosen 51.5% of the time.

Figure 4.2: Probability Density of the Observations in each Belief Type (Density of z_h)

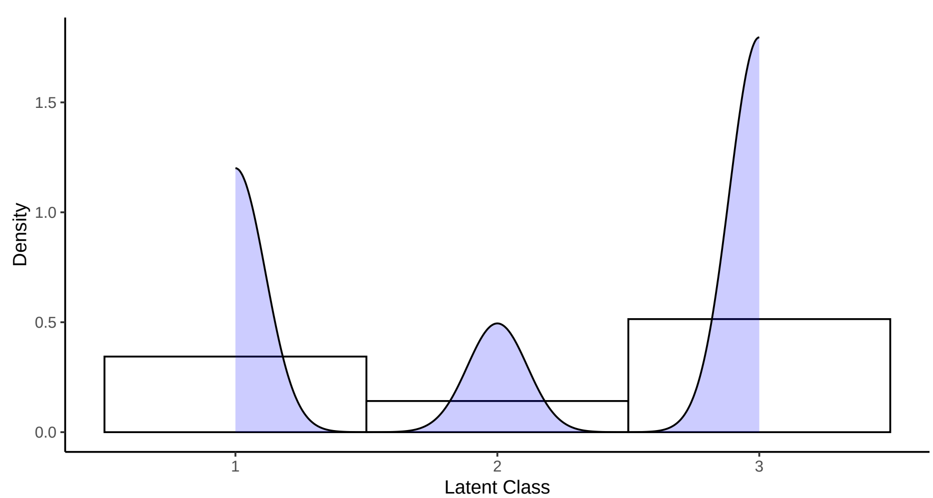

Notes: Probability Density showing that Belief Type 3 has a much higher density and is therefore more likely to be chosen by any observed household.

Figure 4.3: Probability of Response Given Belief Type 1 ($K = 1$): Categorical (β^j_{1v})

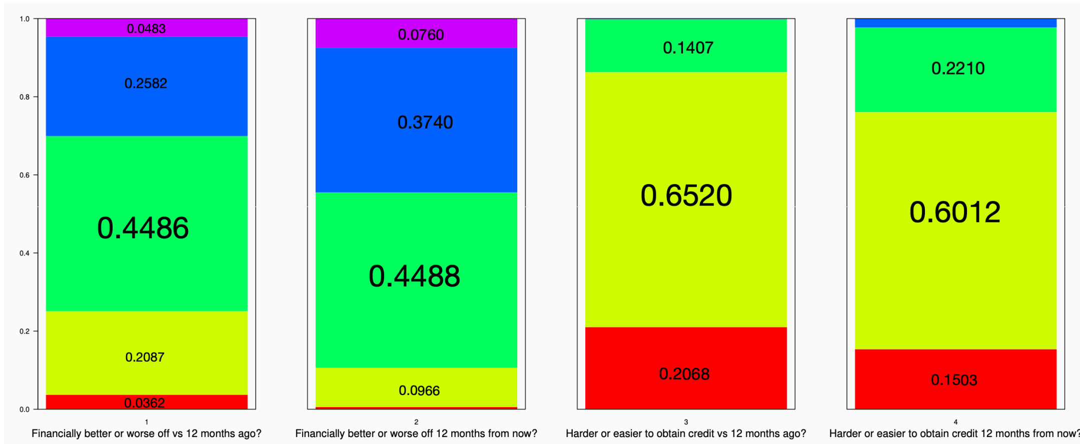

Notes: **[Scale]** Red: Much Worse Off / Harder; Yellow: Somewhat worse off / Somewhat harder; Green: About the same / Equally easy or hard; Blue: Somewhat better / easier; Purple: Much Better Off / Easier.

Figure 4.4: Probability of Response Given Belief Type 1 ($K = 1$): Binary (β^j_{1v})

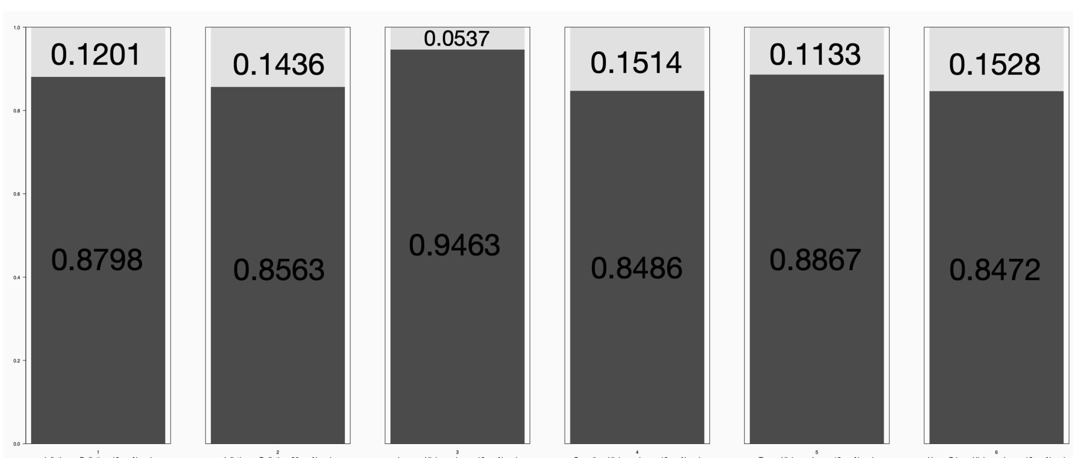

Notes: Scale is based off darker color being the first option e.g. in 'Inflation or Deflation, 12mo Ahead', the darker color represents the respondent selected there will be inflation over the next 12 months.

Figure 4.5: Probability of Response Given Belief Type 2 ($K = 2$): Categorical (β^j_{2v})

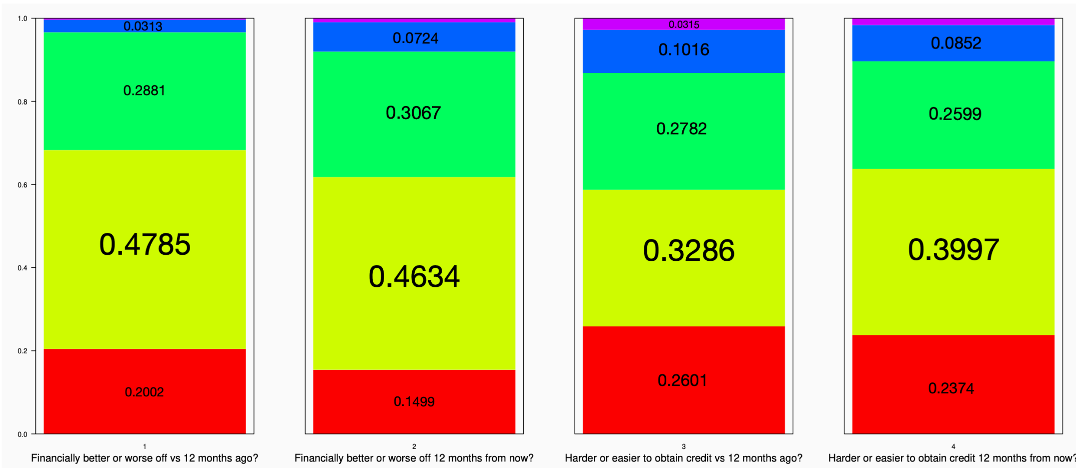

Notes: **[Scale]** Red: Much Worse Off / Harder; Yellow: Somewhat worse off / Somewhat harder; Green: About the same / Equally easy or hard; Blue: Somewhat better / easier; Purple: Much Better Off / Easier.

Figure 4.6: Probability of Response Given Belief Type 2 ($K = 2$): Binary (β^j_{2v})

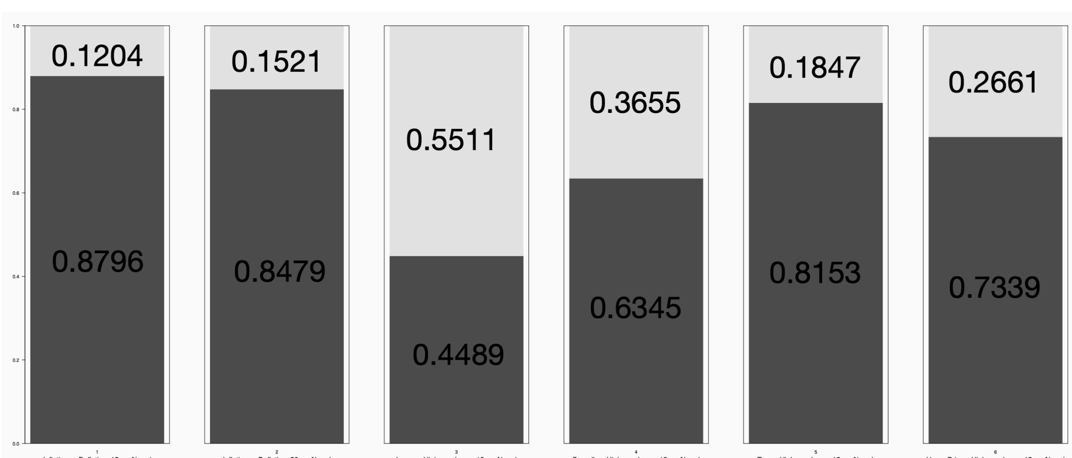

Notes: Scale is based off darker color being the first option e.g. in 'Inflation or Deflation, 12mo Ahead', the darker color represents the respondent selected there will be inflation over the next 12 months.

Figure 4.7: Probability of Response Given Belief Type 3 ($K = 3$): Categorical (β^{j}_{3v})

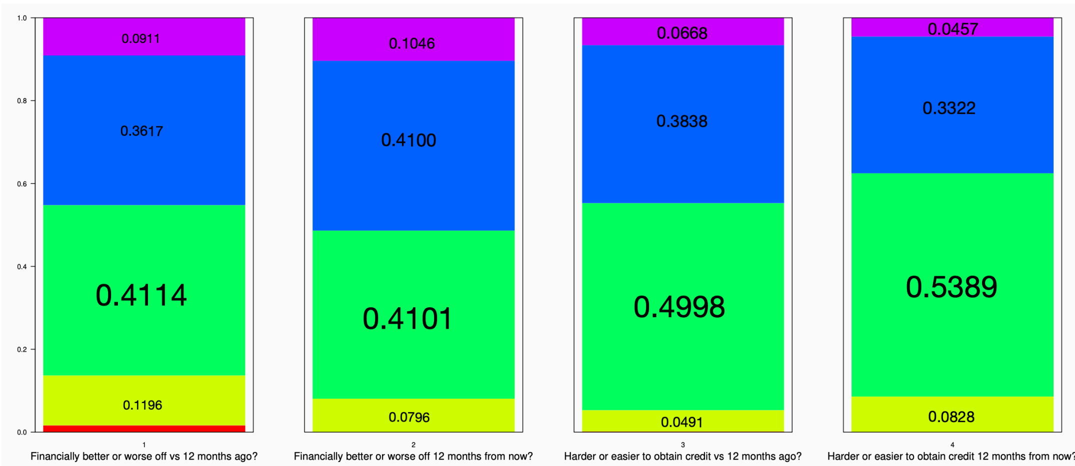

Notes: **[Scale]** Red: Much Worse Off / Harder; Yellow: Somewhat worse off / Somewhat harder; Green: About the same / Equally easy or hard; Blue: Somewhat better / easier; Purple: Much Better Off / Easier.

Figure 4.8: Probability of Response Given Belief Type 3 ($K = 3$): Binary (β^{j}_{3v})

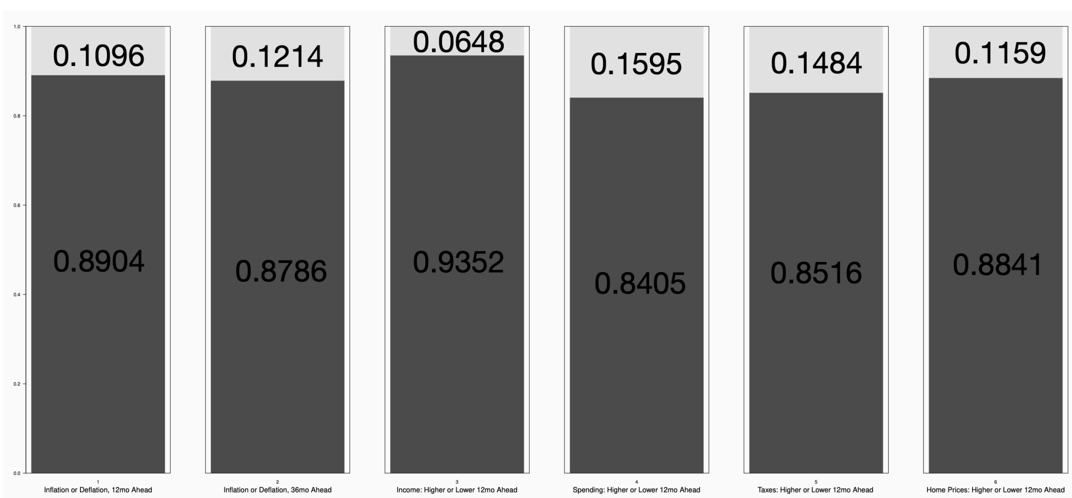

Notes: Scale is based off darker color being the first option e.g. in 'Inflation or Deflation, 12mo Ahead', the darker color represents the respondent selected there will be inflation over the next 12 months.

Figure 4.9: Statistical Model Indices for the Survey of Consumer Expectations; $K = 3$

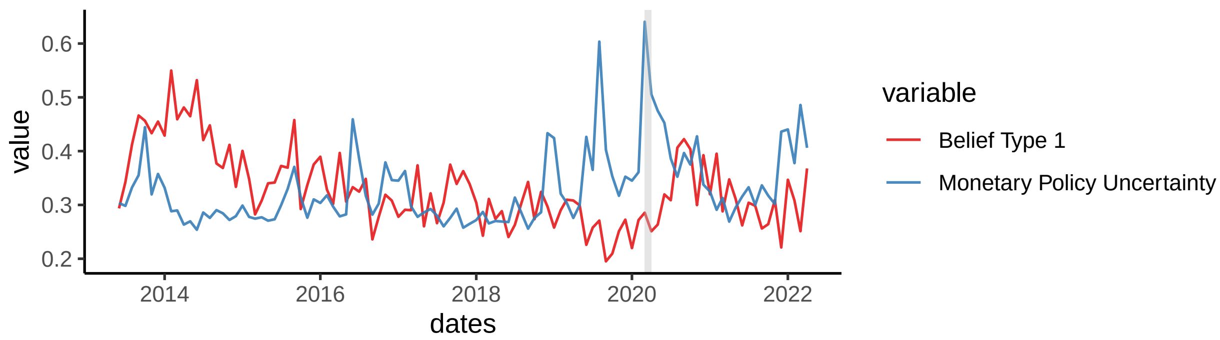

(a) Uncertain Sentiment

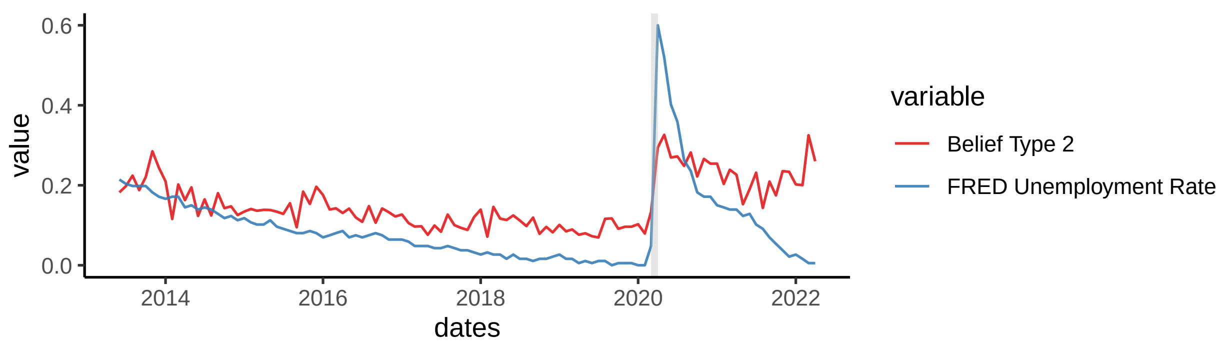

(b) Negative Sentiment

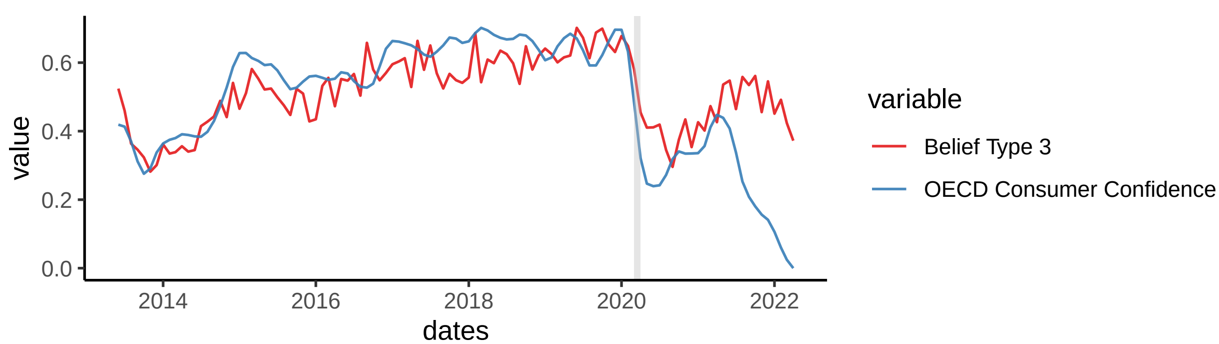

(c) Positive Sentiment

Notes: The belief type indices correspond to the posterior means of the proportions of the $K = 3$ components in the model. These probabilities represent the probability that a randomly selected observation belongs to each cluster at each date. Belief type 1 ($K = 1$) is plotted against a scaled version of the Monetary Policy Uncertainty (MPU) Index from Husted et al. (2020). Belief type 2 ($K = 2$) is plotted against a scaled version of the FRED Unemployment Rate. Belief type 3 ($K = 3$) is plotted against a scaled version of the OCED Consumer Confidence Index for the United States. All date ranges are from June 2013 through April 2022.

5

Conclusion

In this dissertation, I have explored the role of monetary policy announcements, sentiment in credit markets, and household belief types in shaping expectations and their effects on macroeconomic indicators. The three studies presented provide valuable insights into the complex relationships between central bank communication, sentiment, and households' understanding of economic conditions.

In the first study, I examined the impact of monetary policy announcements on household expectations of interest rates, inflation, and home price growth. The results show that while households do respond to monetary policy announcements, they struggle to fully comprehend and react to unconventional policy measures. Expectations for the probability in increasing interest rates one year ahead, one year ahead inflation, and one year ahead home price growth are robustly affected by a number of monetary policy measures. A one standard deviation surprise in the Federal Funds Rate leads to a downwards revision of one year ahead inflation expectations by 0.21% of its overall mean. A one standard deviation surprise in the unified monetary policy measure leads to an upwards revision of one year ahead inflation expectations by 3.6% of its overall mean. I find no effects of monetary policy announcements on 24- to 36-months ahead expectations for inflation, home price growth, or a variety of commodity prices that households would find relevant to their overall financial health.

In the second study, I used machine learning techniques to analyze textual data from the Wall Street Journal and derive factors related to credit spread expectation errors, which served as a proxy for sentiment. The findings demonstrated that over-optimism in credit markets, as indicated by increasing expectation errors, is generally associated with downturns in economic activity. Specifically, a one standard deviation change in the credit spread expectation error is associated with a predicted real GDP growth rate decline of 3%, a predicted 1.68% increase in the unemployment rate, and a predicted 2.9% decrease in domestic investment. These results suggest that sentiment plays a significant role in shaping macroeconomic outcomes.

Finally, in the third study, I identified three distinct household belief types: inconsistent/uncertain, pessimistic, and optimistic. Each of these belief types displayed unique response patterns to macroeconomic and personal expectations. The analysis revealed that the timing of the survey and the information acquisition from news sources played a significant role in shaping these expectations. The inclusion of belief types almost doubled the adjusted R-squared in a model predicting 12-month ahead inflation expectations and showed that households with pessimistic beliefs (Belief Type 2) are more likely to have higher inflation expectations, while those with optimistic beliefs (Belief Type 3) tend to have lower inflation expectations.

Taken together, these studies emphasize the importance of effective central bank communication, sentiment in credit markets, and the diversity of household beliefs in understanding and predicting macroeconomic outcomes. By deepening our understanding of these factors, policymakers and researchers can work towards developing more effective strategies for managing economic stability and growth.

A

A.1 Appendix for Chapter 1

Q5new

What do you think is the percent change **12 months from now** the average interest rate on savings accounts will be *higher* than it is now?

Instruction H2

Ruler & Box

If no response: error E1

Q9

Now we would like you to think about the different things that may happen to inflation **over the next 12 months**. We realize that this question may take a little more effort.

In your view, what would you say is the percent chance that, over the next 12 months...

Instruction H4.

The rate of inflation will be 12% or higher (bin 1)	___ percent chance
The rate of inflation will be between 8% and 12% (bin 2)	___ percent chance
The rate of inflation will be between 4% and 8% (bin 3)	___ percent chance
The rate of inflation will be between 2% and 4% (bin 4)	___ percent chance
The rate of inflation will be between 0% and 2% (bin 5)	___ percent chance
The rate of deflation (opposite of inflation) 0% and 2% (bin 6)	___ percent chance
The rate of deflation (opposite of inflation) 2% and 4% (bin 7)	___ percent chance
The rate of deflation (opposite of inflation) 4% and 8% (bin 8)	___ percent chance
The rate of deflation (opposite of inflation) 8% and 12% (bin 9)	___ percent chance
The rate of deflation (opposite of inflation) will be 12% or higher (bin 10)	___ percent chance
TOTAL	**100**

If no response: error E1

If sum not equal to 100: "Your total adds up to XX" followed by error msg E3.

Q9c

And in your view, what you say is the percent change that, **over the [Month, Year - 24 months from survey date] and [Month, Year - 36 months from survey date], ...**

Instruction H4.

The rate of inflation will be 12% or higher (bin 1)	___ percent chance
The rate of inflation will be between 8% and 12% (bin 2)	___ percent chance
The rate of inflation will be between 4% and 8% (bin 3)	___ percent chance
The rate of inflation will be between 2% and 4% (bin 4)	___ percent chance
The rate of inflation will be between 0% and 2% (bin 5)	___ percent chance

The rate of deflation (opposite of inflation) 0% and 2% (bin 6)	___ percent chance
The rate of deflation (opposite of inflation) 2% and 4% (bin 7)	___ percent chance
The rate of deflation (opposite of inflation) 4% and 8% (bin 8)	___ percent chance
The rate of deflation (opposite of inflation) 8% and 12% (bin 9)	___ percent chance
The rate of deflation (opposite of inflation) will be 12% or higher (bin 10)	___ percent chance
TOTAL	**100**

If no response: error E1

If sum not equal to 100: "Your total adds up to XX" followed by error msg E3.

C1

And in your view, what would you say is the percent chance that, **over the next 12 months**, the average home price nationwide will...

Instruction H4.

Increase by 12% or more (bin 1)	___ percent chance
Increase by 8% to 12% (bin 2)	___ percent chance
Increase by 4% to 8% (bin 3)	___ percent chance
Increase by 2% to 4% (bin 4)	___ percent chance
Increase by 0% to 2% (bin 5)	___ percent chance
Decrease by 0% to 2% (bin 6)	___ percent chance
Decrease by 2% to 4% (bin 7)	___ percent chance
Decrease by 4% to 8% (bin 8)	___ percent chance
Decrease by 8% to 12% (bin 9)	___ percent chance
Decrease by 12% or more (bin 10)	___ percent chance
TOTAL	**100**

If no response: error E1

C2part2

By about what percent do you expect the average home price to [increase/decrease as in C2] over that period?

Instruction H9

Over the 12-month period between [Month, Year - 24 months from survey date] and [Month, Year - 36 months from survey date],

I expect the average home price to [increase/decrease as in C2] by _%

If no response: error E1

Q25v2part2

By about what percent do you expect your total household income [increase/decrease as in Q25v2]? Please give your best guess.

Instructions H9.

Over the next 12 months, I expect my total household income to [increase/decrease] by _%.

If no response: error E1

Q26v2part2

By about what percent do you expect your total household spending [increase/decrease as in Q26v2]? Please give your best guess.

Instructions H9.

Over the next 12 months, I expect my total household spending to [increase/decrease] by _%.

If no response: error E1

C4Info

Twelve months from now, what do you think will have happened to the price of the following items?*Instructions H11.*

I expect...

The price of a gallon of gas to have increased by (1) _ OR decreased by _%

The price of food to have increased by (2)	_ OR decreased by _%
The price of medical care to have increased by (3)	_ OR decreased by _%
The cost of a college education to have increased by (4)	_ OR decreased by _%
The cost of renting a typical house/apartment to have increased by (5)	_ OR decreased by _%
The price of gold to have increased by (6)	_ OR decreased by _%

If no response: error E9

A.2 Appendix for Chapter 3

Figure A.1: Nature of the Dependency between household demographics and responses, categorical

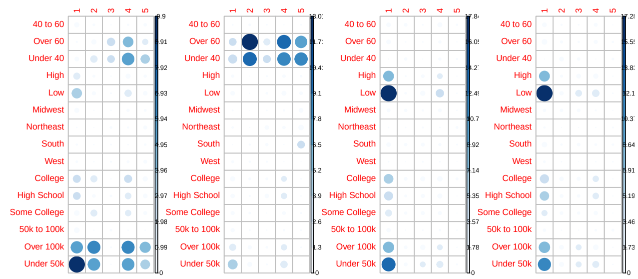

Notes: The figure plots the largest contributors to the Chi-Square dependency tests between the responses for Questions 1 - 4. These are, in order, *"Financially better or worse off vs 12 months ago", "Financially better of worse off 12 months from now", "Easier or harder to obtain credit vs 12 months ago", "Easier or harder to obtain credit 12 months from now"*. The results show that not all of the demographic information collected in the SCE is informative about the nature of dependency for the answers, and that not one characteristic is informative about all of them. Age, levels of financial numeracy (scored via a special module in the SCE), and income are the predominant contributors.

Figure A.2: Nature of the Dependency between household demographics and responses, binary

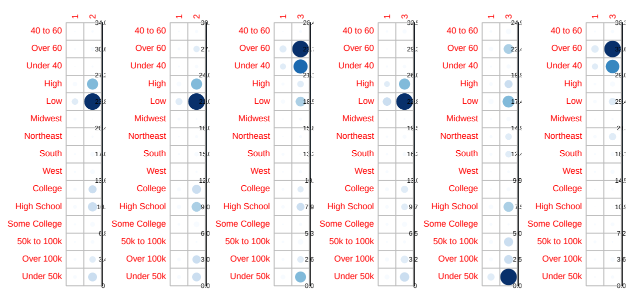

Notes: The figure plots the largest contributors to the Chi-Square dependency tests between the responses for Questions 5 - 10. These are, in order, *"Inflation or Deflation 12 months from now", "Inflation or Deflation 24 - 36 months from now", "Increase or decrease in household income 12 months from now", "Increase or decrease in household spending 12 months from now", "Increase or decrease in taxes paid 12 months from now", "Increase or decrease in home prices nationwide 12 months from now"*. The results show that not all of the demographic information collected in the SCE is informative about the nature of dependency for the answers, and that not one characteristic is informative about all of them. Age, levels of financial numeracy (scored via a special module in the SCE), and income are the predominant contributors.

A.2.1 SCE Variable Selection including Belief Types

Figure A.3: Variable Selection including Belief Types

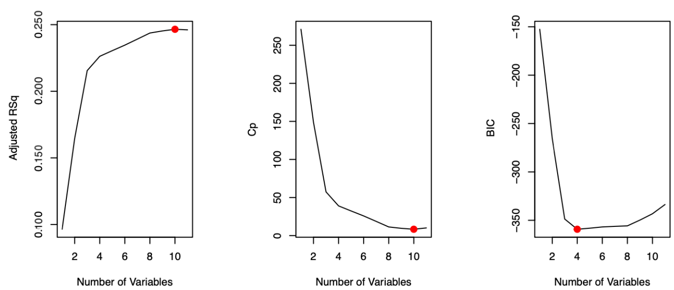

(a) Best Subset Method

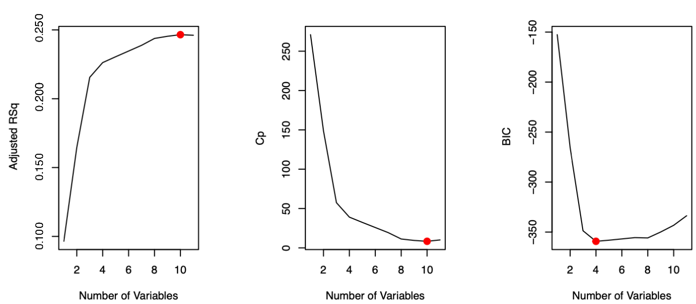

(b) Forward Stepwise

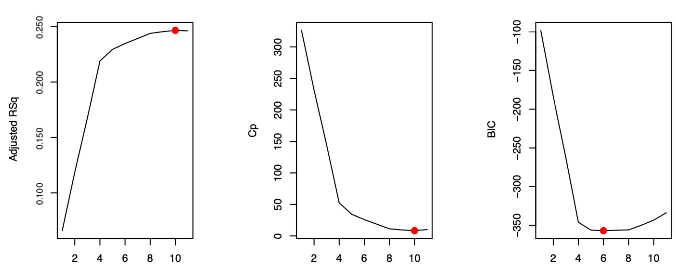

(c) Backward Stepwise

Notes: Different subsets of models and associated measures.

Figure A.4: Variable Selection including Belief Types, Full Test

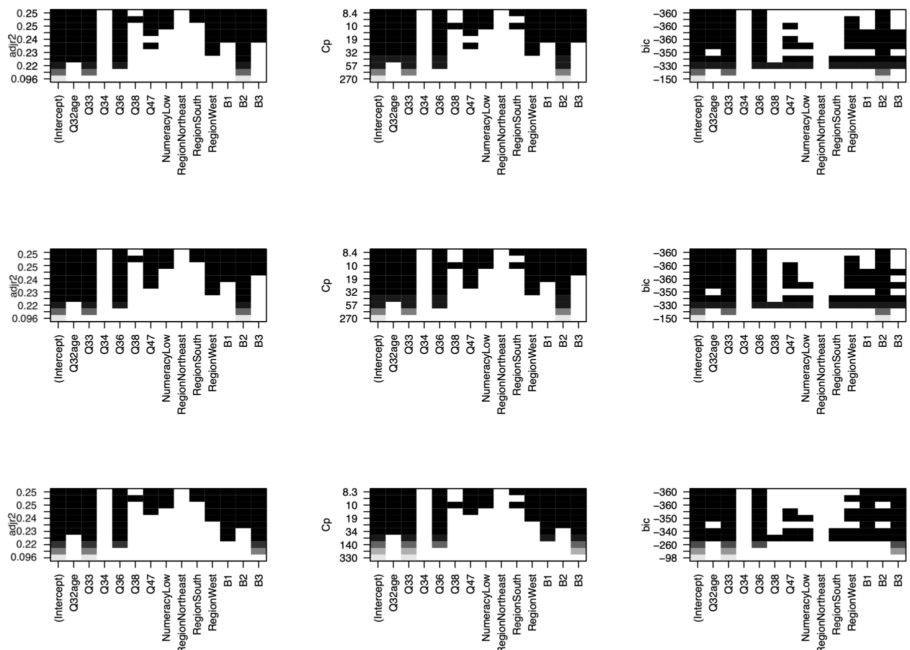

Notes: Plots depicting the adjusted R squared, Cp, and BIC measures per variable included in the model for the Best Subset (Row 1), Forward Stepwise (Row 2), Backward Stepwise (Row 3) methods. The black boxes along the top row of each plot show the variables that were selected by the method.

Figure A.5: Variable Selection, 10-Fold Cross Validation

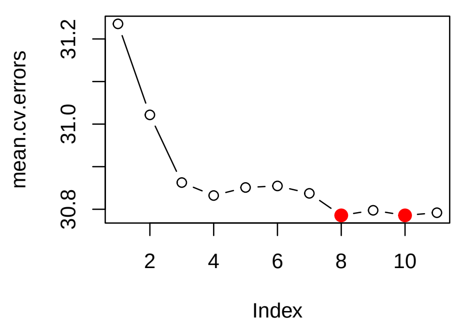

Notes: Plot showing the mean test error for the 10-fold cross validation approach; the 8 and 10 variable model minimize the MSE.

Bibliography

Ahn, H. J., Xie, S., and Yang, C. (2022). Effects of Monetary Policy on Household Expectations: The Role of Homeownership. Technical report, Finance and Economics Discussion Series 2022-065, Washington: Board of Governors of the Federal Reserve System.

Altig, D. E., Barrero, J. M., Bloom, N., Davis, S. J., Meyer, B., and Parker, N. B. (2020). Surveying Business Uncertainty. FRB Atlanta Working Paper 2019-13, Federal Reserve Bank of Atlanta.

Anandkumar, A., Foster, D. P., Hsu, D., Kakade, S. M., and Liu, Y.-K. (2015). A spectral algorithm for latent dirichlet allocation. *Algorithmica*, 72(1):193–214.

Andoni, A., Indyk, P., Laarhoven, T., Razenshteyn, I., and Schmidt, L. (2015). Practical and optimal lsh for angular distance. In Cortes, C., Lawrence, N., Lee, D., Sugiyama, M., and Garnett, R., editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.

Arellano, C., Bai, Y., and Kehoe, P. J. (2019). Financial Frictions and Fluctuations in Volatility. *Journal of Political Economy*, 127(5):2049–2103.

Armantier, O., Boumahdi, F., Koşar, G., Somerville, J., Topa, G., Van der Klaauw, W., Williams, J. C., et al. (2022a). What do consumers think will happen to inflation? Technical report, Federal Reserve Bank of New York.

Armantier, O., Conlon, J. J., and der Klaauw, W. V. (2017a). Political Polarization in Consumer Expectations. Liberty Street Economics 20171215, Federal Reserve Bank of New York.

Armantier, O., Koşar, G., Somerville, J., Topa, G., der Klaauw, W. V., and Williams, J. C. (2022b). The Curious Case of the Rise in Deflation Expectations. Staff Reports 1037, Federal Reserve Bank of New York.

Armantier, O., Scally, J., Smith, K., and der Klaauw, W. V. (2018). Just Released: A Closer Look at Recent Tightening in Consumer Credit. Liberty Street Economics 20181203, Federal Reserve Bank of New York.

Armantier, O., Topa, G., der Klaauw, W. V., and Zafar, B. (2013). Introducing the FRBNY Survey of Consumer Expectations: Survey Goals, Design and Content. Liberty Street Economics 20131204b, Federal Reserve Bank of New York.

Armantier, O., Topa, G., der Klaauw, W. V., and Zafar, B. (2017b). An overview of the Survey of Consumer Expectations. *Economic Policy Review*, (23-2):51–72.

Armantier, O., Topa, G., der Klaauw, W. V., and Zafar, B. (2017c). An overview of the Survey of Consumer Expectations. *Economic Policy Review*, (23-2):51–72.

Arteaga, M. (2022a). Credit market expectations and the business cycle: Evidence from a textual analysis approach. Working paper, Department of Economics, Fordham University.

Arteaga, M. (2022b). Monetary policy announcements and household expectations of the future. Working paper, Department of Economics, Fordham University.

Baron, M. and Xiong, W. (2017). Credit Expansion and Neglected Crash Risk. *The Quarterly Journal of Economics*, 132(2):713–764.

Bauer, M. D. and Swanson, E. T. (2022). A Reassessment of Monetary Policy Surprises and High-Frequency Identification. In *NBER Macroeconomics Annual 2022, volume 37*, NBER Chapters. National Bureau of Economic Research, Inc.

Bellemare, C., Tossou, R. K., and Moran, K. (2020). The Determinants of Consumers' Inflation Expectations: Evidence from the US and Canada. Staff Working Papers 20-52, Bank of Canada.

Bernanke, B. and Gertler, M. (1989). Agency Costs, Net Worth, and Business Fluctuations. *American Economic Review*, 79(1):14–31.

Bholat, D., Hans, S., Santos, P., and Schonhardt-Bailey, C. (2015). *Text mining for central banks*. Number 33 in Handbooks. Centre for Central Banking Studies, Bank of England.

Bianchi, F., Kind, T., and Kung, H. (2019). Threats to Central Bank Independence: High-Frequency Identification with Twitter. NBER Working Papers 26308, National Bureau of Economic Research, Inc.

Bianchi, J. (2011). Overborrowing and Systemic Externalities in the Business Cycle. *American Economic Review*, 101(7):3400–3426.

Binder, C. (2017a). Fed speak on main street: Central bank communication and household expectations. *Journal of Macroeconomics*, 52(C):238–251.

Binder, C. (2017b). Federal Reserve Communication and the Media. *Journal of Media Economics*, 30(4):191–214.

Binder, C. (2019). Panel Conditioning in the Survey of Consumer Expectations. Technical report.

Binder, C. C. (2017c). Measuring uncertainty based on rounding: New method and application to inflation expectations. *Journal of Monetary Economics*, 90(C):1–12.

Binder, C. C. (2018). Inflation expectations and the price at the pump. *Journal of Macroeconomics*, 58(C):1–18.

Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022.

Blinder, A. (2009). Talking about monetary policy: the virtues (and vice?) of central bank communication. BIS Working Papers 274, Bank for International Settlements.

Blinder, A. S., Ehrmann, M., de Haan, J., and Jansen, D.-J. (2022). Central Bank Communication with the General Public: Promise or False Hope? NBER Working Papers 30277, National Bureau of Economic Research, Inc.

Blinder, A. S., Ehrmann, M., Fratzscher, M., De Haan, J., and Jansen, D.-J. (2008). Central bank communication and monetary policy: A survey of theory and evidence. *Journal of Economic Literature*, 46(4):910–45.

Bordalo, P., Gennaioli, N., Ma, Y., and Shleifer, A. (2020a). Overreaction in Macroeconomic Expectations. *American Economic Review*, 110(9):2748–2782.

Bordalo, P., Gennaioli, N., Ma, Y., and Shleifer, A. (2020b). Overreaction in Macroeconomic Expectations. *American Economic Review*, 110(9):2748–2782.

Bordalo, P., Gennaioli, N., Porta, R. L., and Shleifer, A. (2019). Diagnostic Expectations and Stock Returns. *Journal of Finance*, 74(6):2839–2874.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2018). Diagnostic Expectations and Credit Cycles. *Journal of Finance*, 73(1):199–227.

Bordalo, P., Gennaioli, N., and Shleifer, A. (2022). Overreaction and Diagnostic Expectations in Macroeconomics. *Journal of Economic Perspectives*, 36(3):223–244.

Breitenlechner, M., Geiger, M., and Scharler, J. (2023). Monetary policy announcements, consumers' inflation expectations, and readiness to spend. *Macroeconomic Dynamics*, page 1–22.

Brouwer, N. and de Haan, J. (2022). Trust in the ECB: Drivers and consequences. *European Journal of Political Economy*, 74(C).

Bu, C., Rogers, J., and Wu, W. (2021). A unified measure of Fed monetary policy shocks. *Journal of Monetary Economics*, 118(C):331–349.

Caldara, D. and Iacoviello, M. (2022). Measuring Geopolitical Risk. *American Economic Review*, 112(4):1194–1225.

Campbell, J. R., Evans, C. L., Fisher, J. D., and Justiniano, A. (2012). Macroeconomic Effects of Federal Reserve Forward Guidance. *Brookings Papers on Economic Activity*, 43(1 (Spring):1–80.

Campbell, J. R., Ferroni, F., Fisher, J. D., and Melosi, L. (2019). The limits of forward guidance. *Journal of Monetary Economics*, 108(C):118–134.

Candia, B., Coibion, O., and Gorodnichenko, Y. (2020). Communication and the Beliefs of Economic Agents. NBER Working Papers 27800, National Bureau of Economic Research, Inc.

Card, D. (1993). Using Geographic Variation in College Proximity to Estimate the Return to Schooling. NBER Working Papers 4483, National Bureau of Economic Research, Inc.

Cascitti, J., Niebler, S., Müller, A., and Schmidt, B. (2022). Rnacache: A scalable approach to rapid transcriptomic read mapping using locality sensitive hashing. *Journal of Computational Science*, 60:101572.

Cheng, B., Ling, X., and Qiang, Y. (2022). Deep representation learning of text-based economics and other knowledges. *China Journal of Econometrics*, 2(1):58.

Claus, E. and Nguyen, V. H. (2020). Monetary policy shocks from the consumer perspective. *Journal of Monetary Economics*, 114(C):159–173.

Coenen, G., Ehrmann, M., Gaballo, G., Hoffmann, P., Nakov, A., Nardelli, S., Persson, E., and Strasser, G. (2017). Communication of monetary policy in unconventional times. Working Paper Series 2080, European Central Bank.

Coibion, O., Georgarakos, D., Gorodnichenko, Y., Weber, M., and Weber, M. (2020a). Forward Guidance and Household Expectations. CESifo Working Paper Series 8118, CESifo.

Coibion, O. and Gorodnichenko, Y. (2015). Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts. *American Economic Review*, 105(8):2644–2678.

Coibion, O., Gorodnichenko, Y., Knotek, E. S., and Schoenle, R. (2020b). Average Inflation Targeting and Household Expectations. *SSRN Electronic Journal*.

Coibion, O., Gorodnichenko, Y., Kueng, L., and Silvia, J. (2017). Innocent Bystanders? Monetary policy and inequality. *Journal of Monetary Economics*, 88:70–89.

Coibion, O., Gorodnichenko, Y., Kumar, S., and Pedemonte, M. (2020c). Inflation expectations as a policy tool? *Journal of International Economics*, 124(C).

Coibion, O., Gorodnichenko, Y., and Weber, M. (2022). Monetary Policy Communications and Their Effects on Household Inflation Expectations. *Journal of Political Economy*, 130(6):1537–1584.

Cong, L. W., Liang, T., Yang, B., and Zhang, X. (2020). Analyzing Textual Information at Scale. In *Information for Efficient Decision Making Big Data, Blockchain and Relevance*, World Scientific Book Chapters, chapter 10, pages 239–271. World Scientific Publishing Co. Pte. Ltd.

Crump, R. K., Eusepi, S., Tambalotti, A., and Topa, G. (2020). Subjective Intertemporal Substitution. Staff Reports 734, Federal Reserve Bank of New York.

Doh, T., Song, D., and Yang, S.-K. X. (2020). Deciphering Federal Reserve Communication via Text Analysis of Alternative FOMC Statements. Research Working Paper RWP 20-14, Federal Reserve Bank of Kansas City.

Dubovik, A., Fiedler, C., and Parakhonyak, A. (2022). Temporal Patterns in Economics Research. CPB Discussion Paper 440, CPB Netherlands Bureau for Economic Policy Analysis.

Duca-Radu, I., Kenny, G., and Reuter, A. (2021). Inflation expectations, consumption and the lower bound: Micro evidence from a large multi-country survey. *Journal of Monetary Economics*, 118(C):120–134.

D'Acunto, F., Malmendier, U., Ospina, J., and Weber, M. (2019). Exposure to Daily Price Changes and Inflation Expectations. NBER Working Papers 26237, National Bureau of Economic Research, Inc.

Eggertsson, G. B. and Krugman, P. (2012). Debt, Deleveraging, and the Liquidity Trap: A Fisher-Minsky-Koo Approach. *The Quarterly Journal of Economics*, 127(3):1469–1513.

Ehrmann, M. and Wabitsch, A. (2022). Central bank communication with non-experts – A road to nowhere? *Journal of Monetary Economics*, 127(C):69–85.

Fano, S. and Toschi, G. (2022). COVID-19 pandemic and the economy: sentiment analysis on Twitter data. *International Journal of Computational Economics and Econometrics*, 12(4):429–444.

Fiore, F. D., Goel, T., Igan, D., and Moessner, R. (2022). Rising household inflation expectations: what are the communication challenges for central banks? BIS Bulletins 55, Bank for International Settlements.

Fiore, F. D., Lombardi, M. J., and Schuffels, J. (2021). Are households indifferent to monetary policy announcements? BIS Working Papers 956, Bank for International Settlements.

FOMC (2022). 2022 FOMC Longer-Run Goals and Monetary Policy Strategy. Statement, Board of Governors of the Federal Reserve System (U.S.).

Fuster, A., Perez-Truglia, R., Wiederholt, M., and Zafar, B. (2018). Expectations with endogenous information acquisition: An experimental investigation. Working Paper 24767, National Bureau of Economic Research.

Gennaioli, N., Ma, Y., and Shleifer, A. (2016). Expectations and investment. BIS Working Papers 562, Bank for International Settlements.

Gentzkow, M., Kelly, B., and Taddy, M. (2019). Text as Data. *Journal of Economic Literature*, 57(3):535–574.

Giannoni, M., Patterson, C., and Negro, M. D. (2015). The Forward Guidance Puzzle. 2015 Meeting Papers 1529, Society for Economic Dynamics.

Giavazzi, F., Iglhaut, F., Lemoli, G., and Rubera, G. (2020). Terrorist Attacks, Cultural Incidents and the Vote for Radical Parties: Analyzing Text from Twitter. Working Papers 659, IGIER (Innocenzo Gasparini Institute for Economic Research), Bocconi University.

Gilchrist, S. and Zakrajsek, E. (2012). Credit Spreads and Business Cycle Fluctuations. *American Economic Review*, 102(4):1692–1720.

Gorodnichenko, Y., Pham, T., and Talavera, O. (2021). The Voice of Monetary Policy. NBER Working Papers 28592, National Bureau of Economic Research, Inc.

Greenwood, R. and Hanson, S. G. (2013). Issuer Quality and Corporate Bond Returns. *Review of Financial Studies*, 26(6):1483–1525.

Haldane, A. and McMahon, M. (2018). Central Bank Communications and the General Public. *AEA Papers and Proceedings*, 108:578–583.

Hansen, S., McMahon, M., and Prat, A. (2018). Transparency and Deliberation Within the FOMC: A Computational Linguistics Approach. *The Quarterly Journal of Economics*, 133(2):801–870.

Hofmann, T. (1999). Probabilistic latent semantic indexing. In *Proceedings of the 22nd annual international ACM SIGIR conference on Research and development in information retrieval*, pages 50–57.

Husted, L., Rogers, J., and Sun, B. (2020). Monetary policy uncertainty. *Journal of Monetary Economics*, 115(C):20–36.

Jordà, O. (2005). Estimation and Inference of Impulse Responses by Local Projections. *American Economic Review*, 95(1):161–182.

Jordá, O., Schularick, M., and Taylor, A. M. (2013). When Credit Bites Back. *Journal of Money, Credit and Banking*, 45(s2):3–28.

Ke, Z. T., Kelly, B. T., and Xiu, D. (2019). Predicting Returns With Text Data. NBER Working Papers 26186, National Bureau of Economic Research, Inc.

Kelly, B., Manela, A., and Moreira, A. (2021). Text Selection. *Journal of Business & Economic Statistics*, 39(4):859–879.

Kryvtsov, O. and Petersen, L. (2021). Central bank communication that works: Lessons from lab experiments. *Journal of Monetary Economics*, 117(C):760–780.

Kubin, I., Zörner, T. O., Gardini, L., and Commendatore, P. (2019). A credit cycle model with market sentiments. *Structural Change and Economic Dynamics*, 50(C):159–174.

Kugler, P., Kroczek, M., and Zühlke, A. (2022). Parenting Styles, Socioeconomic Status and (Non-) Cognitive Skills. IAW Discussion Papers 138, Institut für Angewandte Wirtschaftsforschung (IAW).

Lamla, M. J. and Vinogradov, D. V. (2019). Central bank announcements: Big news for little people? *Journal of Monetary Economics*, 108(C):21–38.

Leiva-León, D., Perez-Quiros, G., Sapriza, H., and Zakrajšek, E. (2022). Introducing the Credit Market Sentiment Index. *Richmond Fed Economic Brief*, 22(33).

Lewis, D. J., Makridis, C., and Mertens, K. (2019). Do Monetary Policy Announcements Shift Household Expectations? Staff Reports 897, Federal Reserve Bank of New York.

Lüdering, J. and Tillmann, P. (2020). Monetary policy on twitter and asset prices: Evidence from computational text analysis. *The North American Journal of Economics and Finance*, 51(C).

López-Salido, D., Stein, J. C., and Zakrajšek, E. (2017). Credit-Market Sentiment and the Business Cycle. *The Quarterly Journal of Economics*, 132(3):1373–1426.

Manela, A. and Moreira, A. (2017). News implied volatility and disaster concerns. *Journal of Financial Economics*, 123(1):137–162.

Manski, C. F. (2004). Measuring Expectations. *Econometrica*, 72(5):1329–1376.

Manski, C. F. (2018). Survey Measurement of Probabilistic Macroeconomic Expectations: Progress and Promise. *NBER Macroeconomics Annual*, 32(1):411–471.

Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013). Efficient estimation of word representations in vector space. In *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*.

Minsky, H. P. (1977). The Financial Instability Hypothesis: An Interpretation of Keynes and an Alternative to"Standard" Theory. *Challenge*, 20(1):20–27.

Munro, E. and Ng, S. (2022). Latent Dirichlet Analysis of Categorical Survey Responses. *Journal of Business & Economic Statistics*, 40(1):256–271.

Muñoz-Murillo, M., Álvarez Franco, P. B., and Restrepo-Tobón, D. A. (2020). The role of cognitive abilities on financial literacy: New experimental evidence. *Journal of Behavioral and Experimental Economics (formerly The Journal of Socio-Economics)*, 84(C).

Nakamura, E. and Steinsson, J. (2018). High-Frequency Identification of Monetary Non-Neutrality: The Information Effect. *The Quarterly Journal of Economics*, 133(3):1283–1330.

Newey, W. K. and West, K. D. (1994). Automatic Lag Selection in Covariance Matrix Estimation. *Review of Economic Studies*, 61(4):631–653.

Pesaran, M. H. and Weale, M. (2006). Survey Expectations. In Elliott, G., Granger, C., and Timmermann, A., editors, *Handbook of Economic Forecasting*, volume 1 of *Handbook of Economic Forecasting*, chapter 14, pages 715–776. Elsevier.

Powell, J. (2022). FOMC Meeting (June 14 - 15, 2022.). Press conference transcript, Board of Governors of the Federal Reserve System (U.S.).

Rao, C. R. (1992). *Information and the Accuracy Attainable in the Estimation of Statistical Parameters*, pages 235–247. Springer New York, New York, NY.

Resche, C. (2004). Investigating 'greenspanese': from hedging to 'fuzzy transparency'. *Discourse & Society*, 15(6):723–744.

Ruiz, F., Athey, S., and Blei, D. (2017). Shopper: A probabilistic model of consumer choice with substitutes and complements. *The Annals of Applied Statistics*, 14.

Rumler, F. and Valderrama, M. T. (2019). Inflation literacy and inflation expectations: Evidence from Austrian household survey data. *Economic Modelling*, 87(November 2018):8–23.

Ryngaert, J. M. (2022). Inflation disasters and consumption. *Journal of Monetary Economics*, 129(S):67–81.

Ryngaert, J. M. (2023). Balance of Risks and the Anchoring of Consumer Expectations. *Journal of Risk and Financial Management*, 16(2):1–18.

Schularick, M. and Taylor, A. M. (2012). Credit Booms Gone Bust: Monetary Policy, Leverage Cycles, and Financial Crises, 1870-2008. *American Economic Review*, 102(2):1029–1061.

Shapiro, A. H. and Wilson, D. J. (2019). The Evolution of the FOMC's Explicit Inflation Target. *FRBSF Economic Letter*.

Stark, T. (2013). SPF panelists' forecasting methods: A note on the aggregate results of a November 2009 special survey. Research working paper, Federal Reserve Bank of Philadelphia.

Stein, J. C. (2014). Incorporating Financial Stability Considerations into a Monetary Policy Framework: a speech at the International Research Forum on Monetary Policy, Washington, D.C., March 21, 2014. Speech 796, Board of Governors of the Federal Reserve System (U.S.).

Swanson, E. T. (2021). Measuring the effects of federal reserve forward guidance and asset purchases on financial markets. *Journal of Monetary Economics*, 118(C):32–53.

Tobback, E., Nardelli, S., and Martens, D. (2017). Between hawks and doves: measuring central bank communication. Working Paper Series 2085, European Central Bank.

van der Cruijsen, C. and Samarina, A. (2023). Drivers of trust in the ECB during the pandemic. *Applied Economics*, 55(13):1454–1476.

Weber, M. (2022). Subjective inflation expectations of households. *Business Economics*, 57(4):217–221.

Wu, J. C. and Xia, F. D. (2016). Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound. *Journal of Money, Credit and Banking*, 48(2-3):253–291.

Yellen, J. L. (2013). Communication in Monetary Policy: a speech at the Society of American Business Editors and Writers 50th Anniversary Conference, Washington, D.C., April 4, 2013. Speech 625, Board of Governors of the Federal Reserve System (U.S.).

Yue, T. and Jing, M. (2022). Sentiment Analysis Based on Bert and Transformer. In Yang, H., Qiu, R., and Chen, W., editors, *AI and Analytics for Public Health*, Springer Proceedings in Business and Economics, pages 317–328. Springer.

Zhao, Y. (2022). Internal consistency of household inflation expectations: Point forecasts vs. density forecasts. *International Journal of Forecasting*.

Abstract

MARDOQUEO ARTEAGA

M.A., Fordham University

B.A., Montclair State University

Essays on Empirical Macroeconomics and Expectations

Dissertation directed by Arunima Sinha, Ph.D.

This dissertation examines expectations and beliefs in macroeconomic topics through three interconnected studies. First, I analyze the effects of monetary policy announcements on household expectations between 2013 and 2021 using an event study and local projections. I find that the absence of tightening announcements decreases expectations of one-year ahead interest rates by 3.1%, while tightening announcements increase one-year inflation expectations by up to 3.6% and decrease one-year ahead home price growth expectations by 2.0%. Second, I explore the interplay between errors in credit spread expectations and macroeconomic indicators from 1948 to 2022. Using textual analysis on Wall Street Journal front pages, I fill in credit spread expectations and find that one-standard deviation jumps in their forecast errors are associated with declines in economic activity by up to a 3% decline in GDP growth during the sample period. Last, I analyze heterogeneity in categorical expectations data from the Survey of Consumer Expectations through a hierarchical Bayesian latent class model. I identify three belief types and demonstrate their economic significance through co-movements with widely used indices of sentiment and explanatory power for inflation expectations. Taken together, these studies emphasize the importance of effective central bank communication, sentiment in credit markets, and the diversity of household beliefs in understanding and predicting macroeconomic phenomena.

Vita

Mardoqueo (Marc) Eduardo Arteaga Lainez, son of *Mardoqueo Arteaga Delgado* and *Mirna Iveth Lainez Martinez,* was born in Santa Tecla, El Salvador in January 1996. He immigrated to the United States in 2001 to escape gang violence and finished his secondary education in New Jersey. In 2012, he was granted Deferred Action for Childhood Arrivals (DACA) by the Obama Administration.

After finishing high school near the top of his class and working as an overnight janitor, hotel engineer, car salesman, and tutor, he networked to raise private donor scholarships to attend college despite his immigration status. In May 2018, he graduated from Montclair State University with a B.A. in Economics and a minor in Mathematics, delivering the senior address to the Feliciano School of Business.

That same year, he enrolled in the Ph.D. program in Economics at Fordham University in New York, where he worked as the Graduate Assistant to the Economics Department Chair. A year later he became a permanent resident in the United States. In 2021, he was elected Vice President of the Graduate Student Council and served as Council President the following year.

During his graduate school training, he successfully published poetry in the London based *Poetically Magazine* and served on the business advisory board for Warren County Community College. He also worked for a variety of organizations including TIAA, DiMassimo Goldstein, REEF Technology, and the Central Bank of Chile. After graduating, he will start his applied research career as an Economist in the Economic & Statistical Consulting practice at KPMG.

This manuscript was typeset using LaTeX, originally developed by Leslie Lamport and based on Donald Knuth's TeX. The body text is set in 11 point Egenolff-Berner Garamond, a revival of Claude Garamont's humanist typeface. The title page symbol is an homage to my Mesoamerican ancestors and is taken from the `pgfornament` package by Alain Matthes. A template that can be used to format a PhD thesis with this look and feel has been released under the permissive MIT (x11) license, and can be found online at github.com/suchow/Dissertate or from its author, Jordan Suchow, at suchow@post.harvard.edu.