

Household Sentiment Analysis through a Hierarchical Bayesian Latent Class Model *

Mardoqueo Arteaga[†]

March 2023

Abstract

This paper employs Latent Dirichlet Analysis for Survey Data (LDA-S) to identify and classify households into distinct belief types based on their responses in the Survey of Consumer Expectations (SCE). I uncover three belief types – inconsistent/uncertain, pessimistic, and optimistic – characterized by unique patterns of expectations about macroeconomic and personal financial conditions. By incorporating these belief types into a model predicting inflation expectations, I demonstrate a significant improvement in the model's explanatory power. The findings of this study have important implications for central bank communication strategies. 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. Moreover, this research highlights the potential of using latent class analysis techniques to extract valuable information from survey data, which can be applied in various economic contexts.

Keywords: Survey expectations, belief types, latent Dirichlet analysis for survey data, categorical data.

JEL Classification: C83, E71, D83.

*I thank Arunima Sinha and my committee, and the advice from Evan Munro (Stanford) and Martin Kroczeck (Institute for Applied Economic Research, Germany). The usual caveat on error applies.

[†]Economics Department, Fordham University, Bronx, NY 10458. Email: arteaga@fordham.edu

1 Introduction

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. (2020, 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 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 3 presents the survey data and the categorical questions that I focus on in this

analysis. Section 4 lays out the econometric model for survey expectations based on information acquisition and about it can be applied to the SCE, with Section 5 presenting identification of belief types. Section 6 presents and discusses how the belief types are associated to 12-month inflation expectations. Section 7 concludes.

2 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, 2017). 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 (2017); 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 Cruisen 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.

3 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 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 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.

4 Latent Dirichlet Analysis for Survey Data

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 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 $d_h \in \mathbb{G} = \{1, \dots, G\}$ 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 $d_h \in \mathbb{T} = \{1, \dots, T\}$. 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 $x_{hj} \in \mathbb{L}_j = \{1, \dots, L_j\}$ 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 \mathbf{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 \mathbf{X} , is given by $\mathbf{Y}_T = (Y_{G1}, \dots, Y_{GJ})$ 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.

4.1 An Information Acquisition Model

An individual household h chooses which of the K sources of information determines their belief type $z_h \in \mathbb{K} = \{1, \dots, K\}$ to consume by maximizing their utility U which is based off of $\mathbf{u}_{g,:} \in \mathbb{R}^k$, 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:

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

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 $d_h = g$:

$$\pi_{gk} = \mathbb{P}(z_h = k \mid d_h = t) = \mathbb{P} \left(u_{gk} + e_{hk} = \max_{j \in \mathbb{K}} (u_{tj} + e_{hj}) \right) \quad (2)$$

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:

$$x_{hj} = \arg \max_{v \in \{1, \dots, L_j\}} \left(\sum_{u=1}^{L_j} \mathbb{1}(v=u) (q_{z_h,u}^j + s_{hu}^j) \right) \quad (3)$$

where the information source effect $\mathbf{q}_{k,:}^j \in \mathbb{R}^{L_j}$ is drawn independently for each k from some distribution \mathcal{Q} , while the individual-specific effect $s_{vu}^j \in \mathbb{R}$ 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_{kv}^j , defined as:

$$\beta_{kv}^j = \mathbb{P}(x_{hj} = v \mid z_i = k) = \mathbb{P}\left(q_{z_h,v}^j + s_{hv}^j = \max_{u \in \mathbb{L}_j} (q_{z_h,u}^j + s_{hu}^j)\right). \quad (4)$$

where $\mathbf{u}_{g,:}$ is independent over g with a distribution \mathcal{F}_u^g , 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 1 and 3 (namely, $\mathbf{u}_{g,:}$, $e_{h,:}$, $\mathbf{q}_{k,:}^j$ and \mathbf{s}_h^j) nor their distributions, the probabilities in Equations 2 and 4, the $\pi_{g,:}$ and $\beta_{k,:}^j$, are treated as random. Furthermore, to complete the model, I assume that $\pi_{g,:}$ and $\beta_{k,:}^j$ 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 $\eta_{k,:}^j \in \mathbb{R}^{\mathbb{L}_j}$. In summary, the model is defined by the following hierarchical statistical model

$$\begin{aligned} z_h & \mid \boldsymbol{\pi}_{d_h,:} \sim \text{Multinomial}(\boldsymbol{\pi}_{d_h,:}) \\ x_{hj} & \mid \boldsymbol{\beta}, z_i \sim \text{Multinomial}(\boldsymbol{\beta}_{z_h,:}^j) \\ \boldsymbol{\pi}_{d_h,:} & \sim \text{Dirichlet}(\boldsymbol{\alpha}_{d_h,:}) \\ \boldsymbol{\beta}_{z_h,:}^j & \sim \text{Dirichlet}(\boldsymbol{\eta}_{z_h,:}^j) \end{aligned}$$

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

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

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

4.2 Considerations for Estimation

A few considerations to think about for this model in the context of the survey data at hand. Having assignment parameter \mathbf{z} being an $N \times 1$ vector assumes that, for each time observed, there is only

Variable in the Model	Representation
Households in survey	N , total
Outcome Dimension	$x_{hj} \in 1, \dots, L_j$
Frequency Matrix	\mathbf{Y}_G (group response)
Mixture Size	G , number of groups
Outcomes per Household	$J \geq 1$ responses in $x_{h,:}$
Outcome distribution	$\beta_{k,:}^j$ 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

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. $[z_h \mid \mathbf{x}_{h,:}, \boldsymbol{\beta}, \boldsymbol{\pi}_{d_{h,:}}]$ is sampled from a multinomial distribution
2. $[\boldsymbol{\beta} \mid \boldsymbol{\eta}, \mathbf{x}, \mathbf{z}]$ is sampled from a Dirichlet distribution
3. $[\boldsymbol{\pi}_{g,:} \mid \boldsymbol{\alpha}, \mathbf{x}, \mathbf{z}]$ 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 $\boldsymbol{\beta}$ and $\boldsymbol{\pi}_{d_{h,:}}$ from the prior iteration, whereas $\boldsymbol{\beta}$ and $\boldsymbol{\pi}_{g,:}$ 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, $\alpha_{g,:}$, $\eta_{k,:}^j$, and K must be specified. The first two hyperparameters specify prior beliefs about the importance of the group-specific terms ($u_{g,:}$, $q_{k,:}^j$) relative to the individual-specific ones $e_{h,:}$, $s_{h,:}^j$. 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 $\eta_{kv}^j < 1$ 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 $\eta_{kv}^j = 1$ for $k \neq v$, and $\eta_{kv}^j = 10$ 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:

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

¹The inverse would be implied if $\alpha_{gk} < 1$; a similar logic follows for η_{kv}^j .

4.3 Application to the SCE

In the 107 months of observations, which I use as my group variable (i.e., $G = \{1, \dots, 107\}$), there are a total of $N = 19,025$ unique household respondents who answer the survey once in this time period. Of these, I focus on $J = 10$ 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 $J = 10$ 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_{kv}^j
1. Do you think you (and any family living with you) are financially better or worse off these days than 12 months ago?	β_{kv}^1
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?	β_{kv}^2
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_{kv}^3
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_{kv}^4
5. Over the next 12 months, do you think there will be inflation or deflation?	β_{kv}^5
6. Over the 12-month period between 24-36 months (from survey date), do you think there will be inflation or deflation?	β_{kv}^6
7. Over the next 12 months, I expect my total household income to...	β_{kv}^7
8. Over the next 12 months, I expect my total household spending to...	β_{kv}^8
9. Twelve months from now, I expect my total taxes to...	β_{kv}^9
10. Over the next 12 months, I expect the average home price to...	β_{kv}^{10}

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*.

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

5 Identification of Belief Types

To recap the approach in the preceding 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 $x_{hj} \in \mathbb{L}^j = \{1, \dots, L^j\}$ 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 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 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 $\beta_{k,v}^j$, in Figures 4 to 7. I discuss each more closely below.

In Figure 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 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 $\beta_{k,v}^j$ in Figures 3 to 8. Figures 3 and 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 ($\beta_{1,3}^1 = 44.86\%$). For question 2, they are also more likely to continue to think they will be *About the Same* in the following year ($\beta_{1,3}^2 = 44.88\%$). 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 $\beta_{1,2}^3 = 65.2\%$, and that it will be *Somewhat harder* in the next year than it is now to do the same with a probability of $\beta_{1,2}^4 = 60.1\%$. 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% ($84.72\% < \beta_{1,1}^5, \beta_{1,1}^6, \beta_{1,1}^7, \beta_{1,1}^8, \beta_{1,1}^9, \beta_{1,1}^{10} < 94.63\%$).

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

Belief Type 3 respondent patterns are shown in Figures 7 and 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 $\beta_{3,3}^1 = 41.14\%$ and $\beta_{3,4}^2 = 41.01\%$, 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 $\beta_{3,3}^4 = 53.89\%$). 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% ($84.05\% < \beta_{3,1}^5, \beta_{3,1}^6, \beta_{3,1}^7, \beta_{3,1}^8, \beta_{3,1}^9, \beta_{3,1}^{10} < 93.52\%$).

To summarize the key differences between Belief Types, I compute the Rao distance between the probabilities, i.e. between $\beta_{k,:}^j$ and $\beta_{m,:}^j$: for $k \neq m$, to find where these types differ most from

each other, depicted in Table 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. 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 Crujssen 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

³The 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.

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 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).

6 Heterogeneous Beliefs and Inflation Expectations

As a means to analyze an application of the belief types from Section 5, 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 identify 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 \mathbf{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: $Y_t = \beta_d \mathbf{X}_t + \beta_1 B1 + \beta_2 B2 + \beta_3 B3 + \varepsilon$ 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 B.

Figures 12 through 14 show the results of the variable selection methods. In Figure 12, 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 13, 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 14; 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 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

⁴I 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.

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 $Y = \beta_d \mathbf{X} + \sum_{j=1}^B B_j \beta_j(T) + \varepsilon$, 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 coefficient 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.

7 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.

References

- 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.
- 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., 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., 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. (2017). 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.
- Binder, C. C. (2017). Measuring uncertainty based on rounding: New method and application to inflation expectations. *Journal of Monetary Economics*, 90(C):1–12.
- Blei, D. M., Ng, A. Y., and Jordan, M. I. (2003). Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022.
- 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.
- 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).

- 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.
- Coibion, O., Georgarakos, D., Gorodnichenko, Y., Weber, M., and Weber, M. (2020). Forward Guidance and Household Expectations. CESifo Working Paper Series 8118, CESifo.
- 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.
- 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.
- 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.
- 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.
- 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.
- Kelly, B., Manela, A., and Moreira, A. (2021). Text Selection. *Journal of Business & Economic Statistics*, 39(4):859–879.
- Kugler, P., Kroczeck, M., and Zühlke, A. (2022). Parenting Styles, Socioeconomic Status and (Non-) Cognitive Skills. IAW Discussion Papers 138, Institut für Angewandte Wirtschaftsforschung (IAW).
- 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.
- 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).
- 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.
- Rao, C. R. (1992). *Information and the Accuracy Attainable in the Estimation of Statistical Parameters*, pages 235–247. Springer New York, New York, NY.
- 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.
- 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.
- Shapiro, A. H. and Wilson, D. J. (2019). The Evolution of the FOMC’s Explicit Inflation Target. *FRBSF Economic Letter*.
- van der Crujsen, 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.
- 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.

Table 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
Obtaining credit now vs 12mo ago?	Much harder	Somewhat harder	Equal	Somewhat easier	Much easier						
Obtaining credit 12mo from now?	0.11	0.29	0.35	0.21	0.04	0.06	0.02	0.11	0.07	0.03	0.00
	0.09	0.31	0.39	0.19	0.03	0.08	0.03	0.06	0.05	0.01	0.00
Inflation or deflation 12mo ahead?	Inflation		Deflation								
Inflation or deflation 24-36mo ahead?	0.89		0.11			0.05	0.00	0.07	0.06	0.09	0.00
	0.87		0.13			0.09	0.02	0.08	0.08	0.08	0.00
Household income 12mo ahead?	Increase by 0% or more		Decrease by 0% or more								
Household spending 12mo ahead?	0.86		0.14			0.01	0.04	0.17	0.13	0.08	0.00
Taxes 12mo ahead?	0.81		0.19			0.05	0.02	0.28	0.17	0.07	0.00
Home prices 12mo ahead?	0.86		0.14			0.16	0.07	0.11	0.11	0.06	0.00
	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 depicts 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 2. Largest Differences between Belief Types (Rao Distance between $\beta_{k,:}^j$ and $\beta_{m,:}^j$ for $k \neq m$)

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

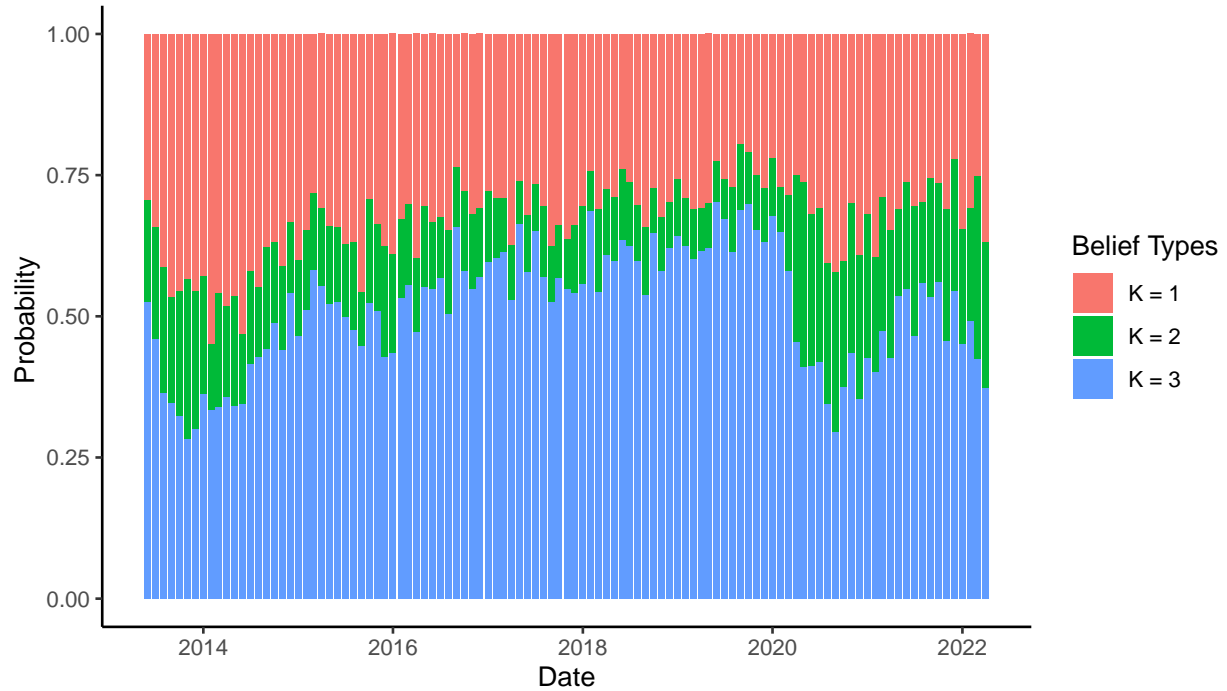
Notes: This table summarizes the five biggest differences between each of the Belief Types uncovered from the survey responses. These differences are computed by using the Rao Distance (Rao, 1992) and can be thought of measures of dissimilarity. The score next to each of the questions where the types 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 3. [Results] Belief Types and Inflation Expectations

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

Figure 1: Probability of Observations Across Time Belonging to Belief Types (π_{gk})



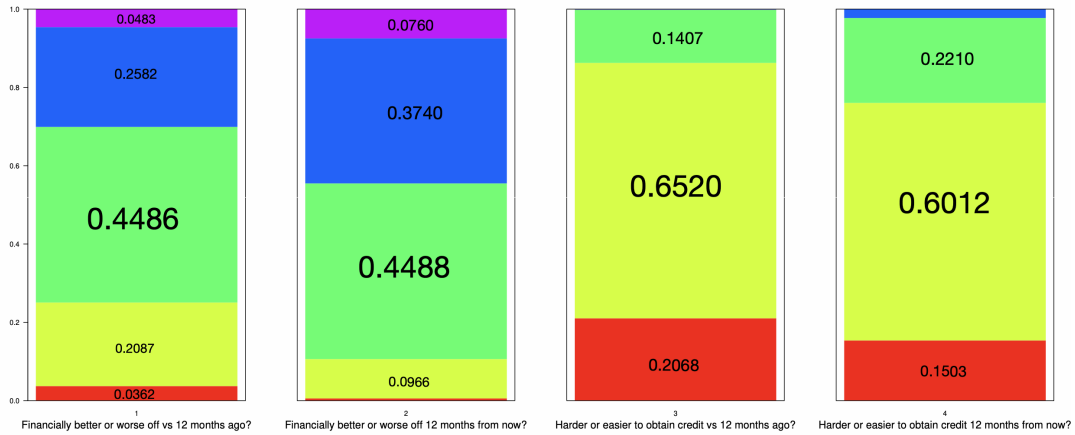
Notes: Estimated representation of the probability for household h of group g belonging to class k . On average, Belief Type 1 is chosen 33% of the time, Belief Type 2 is chosen 15.5% of the time, and Belief Type 3 is chosen 51.5% of the time.

Figure 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 3: Probability of Response Given Belief Type 1 ($K = 1$): Categorical (β_{1v}^j)



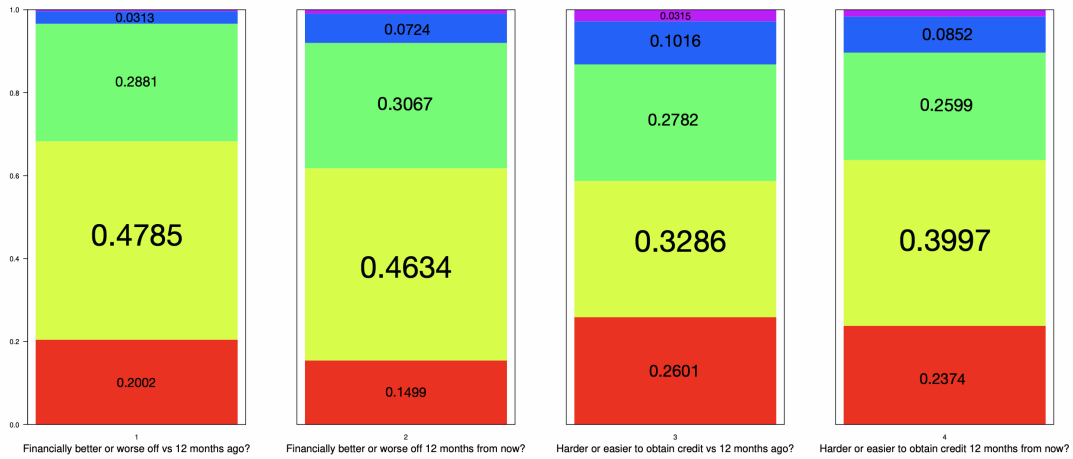
Notes: [Scale] **Red:** Much Worse Off / Harder; **Yellow:** Somewhat worse off / Somewhat harder; **Green:** About the same / Equally easy or hard; **Blue:** Somewhat better / easier; **Purple:** Much Better Off / Easier.

Figure 4: Probability of Response Given Belief Type 1 ($K = 1$): Binary (β_{1v}^j)



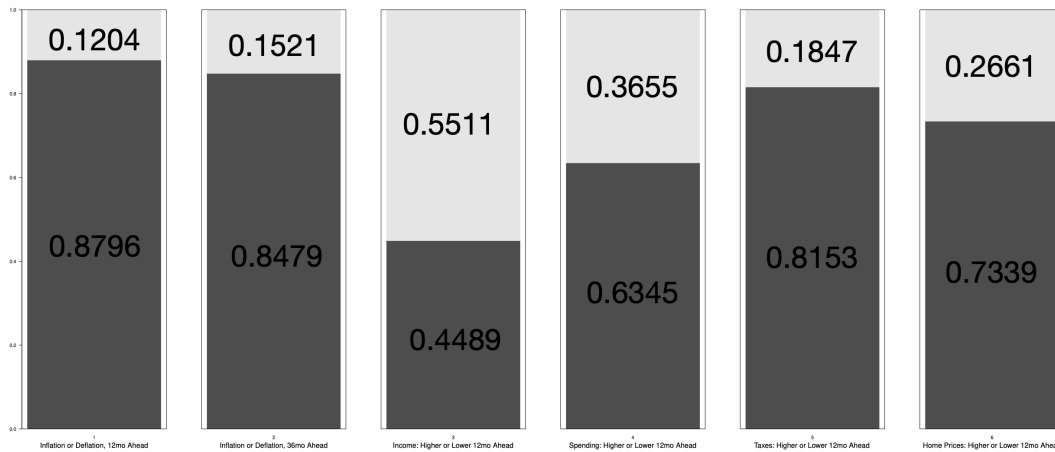
Notes: Scale is based off darker color being the first option e.g. in 'Inflation or Deflation, 12mo Ahead', the darker color represents the respondent selected there will be inflation over the next 12 months.

Figure 5: Probability of Response Given Belief Type 2 ($K = 2$): Categorical (β_{2v}^j)




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 6: Probability of Response Given Belief Type 2 ($K = 2$): Binary (β_{2v}^j)



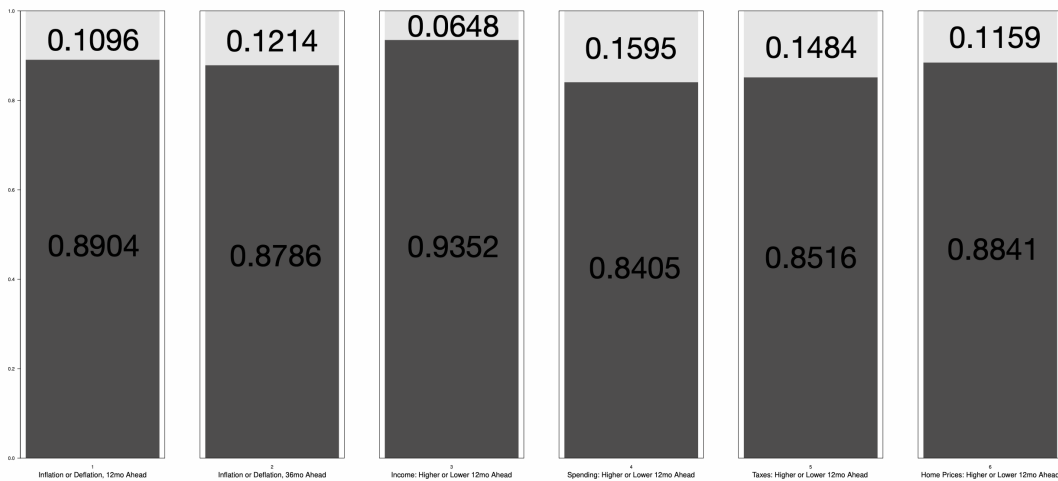
Notes: Scale is based off darker color being the first option e.g. in 'Inflation or Deflation, 12mo Ahead', the darker color represents the respondent selected there will be inflation over the next 12 months.

Figure 7: Probability of Response Given Belief Type 3 ($K = 3$): Categorical (β_{3v}^j)




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 8: Probability of Response Given Belief Type 3 ($K = 3$): Binary (β_{3v}^j)




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



(a) Uncertain Sentiment



(b) Negative Sentiment

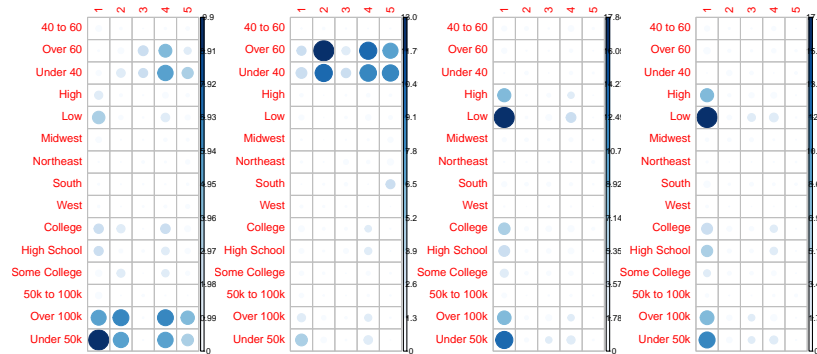


(c) Positive Sentiment

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.

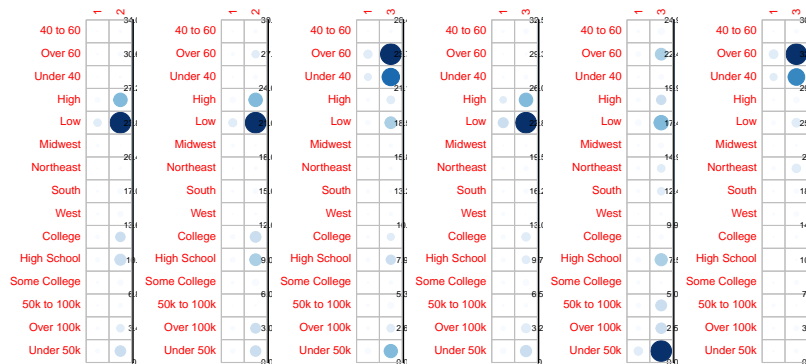
A SCE Response Heterogeneity and Demographic Characteristics

Figure 10: 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 or 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.

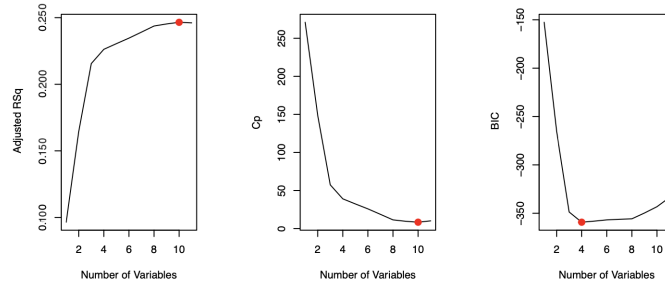
Figure 11: Nature of the Dependency between household demographics and responses, binary



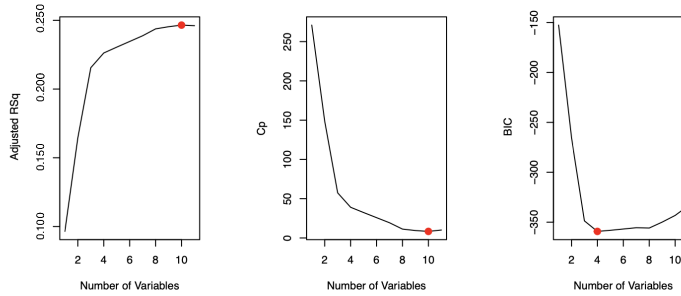
Notes: The figure plots the largest contributors to the Chi-Square dependency tests between the responses for Questions 5 - 10. These are, in order, "Inflation or Deflation 12 months from now", "Inflation or Deflation 24 - 36 months from now", "Increase or decrease in household income 12 months from 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.

B SCE Variable Selection including Belief Types

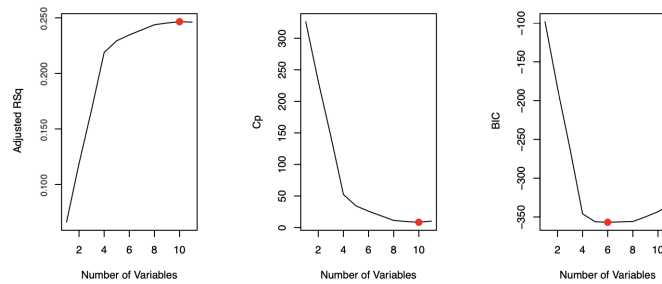
Figure 12: Variable Selection including Belief Types



(a) Best Subset Method



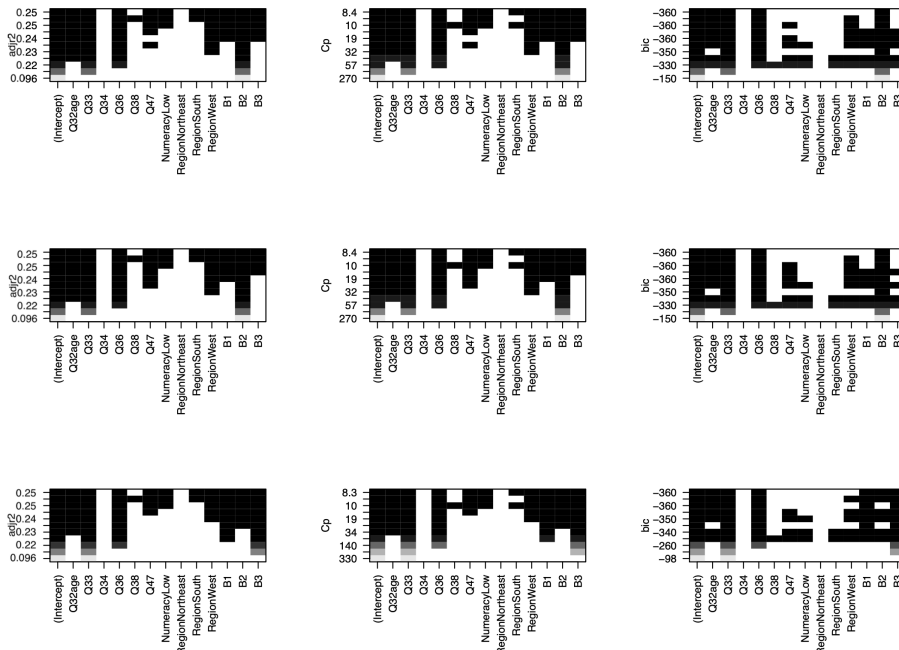
(b) Forward Stepwise



(c) Backward Stepwise


Notes: Different subsets of models and associated measures.

Figure 13: 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 14: 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.