Using Demand Estimation to Evaluate the Drivers of Household Investment Performance*

Daniel Graves†

May 22, 2026

Abstract

A large behavioral finance literature documents the mistakes made by individual investors, with only a handful of papers indicating that retail traders might be skilled. This paper investigates retail investment skills and beliefs by building a random coefficients-based structural model of household direct stock ownership and estimating it using only aggregate data. I account for heterogeneity in household return expectations, varying attention to specific stocks, and shared sentiment among households. Using Generalized Empirical Likelihood via a Mathematical Program with Equilibrium Constraints (GEL-MPEC) as described theoretically by Conlon (2013), I estimate household demand parameters in each quarter from 1984 to 2021 and in the process compute average household sentiment about each stock: the extra idiosyncratic return the average household expects to earn from investing in a particular stock. I find that contrary to popular belief, retail traders display skill in evaluating large stocks, with sentiment-based portfolios modestly but positively predicting future returns. However, households' limited attention, macroeconomic beliefs, and overtrading undo the positive gains from investment analysis and result in portfolios that underperform the stock market on a net basis. Investors may understand product quality, but they are poor macroeconomic traders and their skill is ultimately undone by overtrading. This paper's results suggest a nuanced view of household trading skills and biases.

JEL Codes: G12, G11, G40, G23, G14, C34, C36, C60

*I thank my dissertation advisors Nicholas Barberis, Eduardo Dávila, John Geanakoplos, and Stefano Giglio for their continued support and guidance as well as early feedback on this project during graduate school. I also thank Steven Berry, John Campbell, Xiaohong Chen, Philip Haile, Sam Hanson, and Adi Sunderam as well as conference and seminar participants at the Booth Household Finance conference, Yale Economics, and HBS for helpful conversations and comments. I would also like to thank the Yale Center for Research Computing and Harvard's FAS research computing group for the use of computing infrastructure. I gratefully acknowledge support of the NSF GRFP; this material is partly based upon work supported by the National Science Foundation Graduate Research Fellowship under Grant No. DGE-1752134. I also am grateful to the Harvard Data Science Initiative for providing grant support for this project.

†Harvard University, 1805 Cambridge St, Cambridge, MA 02138. Email: danielgraves@fas.harvard.edu

1 Introduction

Household investments in common stocks are one of the main sources of savings for retirements, home purchases, college tuition, and future consumption. Distortions in household beliefs about common stocks therefore have a significant impact on household welfare, as irrational beliefs and limited attention can lead to suboptimal portfolios and therefore diminished risk sharing and price discovery versus standard rational benchmarks such as the CAPM. Moreover, household investment mistakes can have significant welfare consequences and have a strong impact on individual households' financial health. An extensive behavioral finance literature has explored investors' departures from rational benchmarks, with overconfidence (e.g. Barber and Odean (2000)), the disposition effect (e.g. Barberis and Xiong (2009) and Shefrin and Statman (1985)), realization utility (e.g. Barberis and Xiong (2012) and Frydman et al. (2014)), extrapolative expectations (e.g. Barberis, Greenwood, Jin and Shleifer (2015) and Barberis, Greenwood, Jin and Shleifer (2018)), limited attention (e.g. Barber and Odean (2008), Engelberg and Parsons (2011) and Da, Engelberg and Gao (2011)), and non-standard preferences (Barberis and Huang (2008)), among other phenomena, hypothesized as playing a role in the chronic underperformance of households in their direct stock investments. Some papers have found that household trades positively predict future returns, while others argue that the evidence points to underperformance on a trade-weighted basis.

Yet one major unstudied question in the literature is whether households' mixed record in the stock market is not merely a morass of nonstandard preferences, biased beliefs, and faulty inferences, but rather an amalgamation of opposing directional effects that each have a distinct behavioral or informational foundation. In other words, we must disentangle which beliefs and tendencies of households positively versus negatively contribute to their performance. Indeed, various patterns in household trading can exist in combination: individuals might overtrade due to overconfidence in their abilities and be driven to purchase attention-grabbing stocks to their own detriment, but also might use private information effectively to gain an advantage over uninformed traders. Additionally, traders might place trades on the basis of market-timing or factor-timing decisions, such as a bet on a stock market recovery. An alternative hypothesis, however, is that investors might merely resemble the canonical noise traders of De Long, Shleifer, Summers and Waldmann (1990), with no private information.

This paper builds a model of household demand for common stocks, develops an estimation methodology, and then applies the new methodology to study which behavioral and informational factors contribute positively or negatively to household investment performance. The first question addressed by this paper is how to model household demand for common stocks given that we only possess aggregate data and households are generally short sale constrained with heterogeneous beliefs about returns. I show how under a random coefficients specification of beliefs about returns and a plausible assumption about the covariance matrix used by households, we can create an invertible demand system in the spirit of Berry, Levinsohn and Pakes (1995) (henceforth,

"BLP"). With this model in hand, I show how to circumvent the significant computational issues associated with estimating this model by applying a version of the Generalized Empirical Likelihood via a Mathematical Program with Equilibrium Constraints (GEL-MPEC) approach of Conlon (2013). This GEL-MPEC approach, which utilizes a novel instrument as well as standard BLP instruments, enables me to estimate (under the model's assumptions) household beliefs about returns on a quarterly basis. Specifically, I estimate household return expectations as a function of different asset characteristics such as log price, log book value, momentum, and market beta, as well as idiosyncratic beliefs about assets that are not connected to the specified observables ("household sentiment"). Furthermore, we can obtain a time series of dispersion of each component of these beliefs, which is also identified by the paper's moment conditions. Finally, I use the time series of estimates for asset-specific household sentiment as well as the estimates of beliefs about characteristics to conduct a series of return predictability tests to determine which beliefs of retail investors are helpful and which are harmful to investment performance. These tests show that in contrast to many empirical studies of retail investor trading, households benefit from their sentiment, with idiosyncratic beliefs about stock returns positively predicting returns once we account for heterogeneous attention and beliefs about observable characteristics. On the other hand, beliefs about factor returns (e.g. the return to the overall market or momentum) are a modest drag on household stock market performance. The overall picture of household investments in common stocks is therefore decidedly mixed. As was already clear from a long literature, households suffer from chronic overtrading, with short-term round-trip trading dragging down investment performance while benefiting the institutions intermediating this trade (e.g. wholesalers). This paper adds an additional negative channel: the evidence is inconsistent with households having market-timing beliefs that are beneficial to their investment performance, with such beliefs serving as a modest drag on performance. The positive channel added in this paper is the more unique one: households' stock-specific sentiment positively predicts returns, with portfolios formed solely on the basis of household sentiment significantly outperforming those that only incorporate beliefs about factor returns (e.g. returns to momentum or beta exposure).

To estimate unobservable household beliefs about individual stocks and separate them from beliefs about observable stock characteristics, we need to formulate a model of household demand and then estimate it, with this model and estimation being the primary methodological contribution of this paper. I start from a random coefficients specification in the spirit of Berry, Levinsohn and Pakes (1995): household return beliefs are a heterogeneous continuum. We have an average household belief about returns to a given stock's idiosyncratic component or to exposure to a given risk factor such as market beta, momentum, and value, but we additionally have household-specific beliefs about returns to factors and specific stocks. In Berry, Levinsohn and Pakes (1995) in the context of consumer auto purchase decisions, a household might have an idiosyncratic taste for larger, family-friendly vehicles as well as a particularly high idiosyncratic utility gain from a Toyota Previa, beyond the average household's opinion of large vehicles, family-friendly vehicles, and Toyota Previas. In this paper's model, a household might have an idiosyncratically positive

slope of its return expectations with respect to momentum and beta characteristics, beyond the average household's expected returns to momentum and beta, as well as an idiosyncratically high expected return to investments in Coca Cola stock, beyond what the average household believes it will earn from Coca Cola specifically.

The main contrast between this paper's model and that of Berry, Levinsohn and Pakes (1995) is how these random coefficients translate into demand. In Berry, Levinsohn and Pakes (1995), households engage in a discrete choice problem where they either do not purchase a car or do purchase a single car, that car being the car that provides the greatest level of utility. Given a continuum of consumers, one can integrate across all of the consumers' choices to find aggregate demand. In this paper, by contrast, each household will typically demand multiple different stocks, following the basic principle of diversification. Creating a demand system from such behavior would typically be challenging, but this paper's focus on households allows us to make a simplifying assumption that is consistent with the behavioral finance literature. Barberis, Huang and Thaler (2006) and others have argued that households engage in "narrow framing" with their investments and gambles. In other words, stock investments are viewed largely in isolation as a "good" or "bad" investment as opposed to determined as part of a comprehensive optimization problem that rationally accounts for covariance. Here, I assume that households view the covariance matrix as diagonal, with households considering only an average of 10 stocks in this paper's main specification. When a household engages in short sale constrained and leverage constrained mean variance optimization with respect to such a small consideration set and a diagonal covariance matrix, they are fundamentally picking stocks with the best perceived risk-reward ratios and shrinking expected returns by a sufficient amount that the selected portfolio does not exceed the maximum leverage level permitted. Given this demand for a given household, we can then integrate over all households to create a vector representing aggregate demand for each stock.

With a demand model established, the paper turns to estimation and properties of the household demand function. After showing that the demand function of households is consistent with the existence of a unique equilibrium under certain conditions on institutional and household investors, I show that the household demand system satisfies the requisite conditions for invertibility. In this paper's context, invertibility means that we cannot have two different household sentiment vectors (average household idiosyncratic return expectations) that are both consistent with a given level of demand and a given set of demand parameters (e.g. beliefs about factor returns and the dispersion of beliefs). Instead, we are left with a one-to-one mapping so that if we fix the parameters, there is a unique sentiment vector that maps to the observed demand.

Estimating demand, however, is still a significant computational challenge because of the large number of markets (stocks), which leads to traditional approaches such as the nested fixed point (NFXP) approach with GMM failing, despite its popularity in the IO literature. The Lipschitz constant is simply too large to achieve convergence in a reasonable time period. Instead, this paper exploits the theoretical framework of Conlon (2013) and recasts the estimation problem as the solution of a Generalized Empirical Likelihood (GEL) estimation via a Mathematical

Program with Equilibrium Constraints (MPEC). With a set of IVs in hand, the GEL-MPEC approach enables us to estimate demand and recover the corresponding household sentiment vector. A single estimation still takes around 100 CPU hours, and we must conduct the estimation independently for approximately 150 different time periods, but GEL-MPEC is nonetheless a significant computational improvement versus the NFXP approach within the context of this paper's model. The key intuition, detailed in Conlon (2013), is that GEL with an objective function consisting of squared probability weights is the same estimator as continuously updated GMM (CUE). By expressing the optimization problem in its primal form, we are able to obtain a convex objective function subject to highly non-convex equilibrium constraints. Using numerical optimizers, this is a soluble problem.

With the household portfolio demand model and estimation procedure in hand, I move on to data construction and estimation. First, we must compute aggregate household direct ownership in stocks, which can be imputed at a quarterly level by taking the market equity of a given stock, subtracting institutional ownership from 13F filings, subtracting insider ownership, and adding back aggregate short interest. Next, the paper specifies stock characteristics for the random coefficient model of beliefs, which in the main specification are taken to be log market equity, log book value, momentum (twelve month returns excluding the most recent month), market beta, and a constant. As detailed in Section 4, I exploit a novel instrumental variable for momentum (which can be endogenous because changes in household sentiment can create price changes) as well as traditional BLP instruments to identify household beliefs via a classic exclusion restriction. In particular, I use conditional moment restrictions to generate unconditional moment restrictions that are then used in estimation.

After estimating the model in each time period, the paper undertakes two main forms of analysis: (1) studying the estimated beliefs of households regarding traditional factors such as value, momentum, and market beta and (2) analyzing the performance and relative performance of counterfactual portfolios that we would obtain by isolating specific components of household demand while excluding others. The results of (1) provide insight into the preferred portfolio attributes of households and the results of (2) offer strong evidence against the hypothesis that household sentiment is a negative driver of household investment performance.

The estimated coefficients reveal that households on average are value investors: all else equal, they believe that stocks with higher book-to-market ratios will deliver greater returns. By contrast, households swing from being pro-momentum to anti-momentum in their holdings, with no clearly discernible pattern except as relates to volatility: during times of high market volatility, households shift more in the direction of expecting positive returns to momentum. On average, however, household portfolios reflect beliefs that tilt away from stocks with strong recent returns. Finally, households appear to on average place a positive premium on beta but barely do so: contrary to the CAPM, the typical household does not expect that increased exposure to beta will result in higher returns. This need not imply that households have a poor understanding of returns to beta, as the empirical evidence itself is consistent with an extraordinarily weak relationship between beta

and returns in the cross section (see e.g. Frazzini and Pedersen (2014)).

The evidence that households are on average weakly anti-momentum in their holdings appears to contradict past research indicating that households engage in momentum trading. Yet the evidence is not contradictory. First, this paper shows the *average* belief of households about a characteristic as opposed to the belief of those who hold a given stock. One can recast this object as a latent belief of households since it is not directly observable and is not necessarily the prime impetus for trading. A given stock may indeed have concentrated ownership among pro-momentum traders. Second, and most importantly, *trading* and *ownership* are two distinct phenomena. Ownership is of course the byproduct of decades of trading, but it is fundamentally a longer-term object and it ignores short-term round trip trades. On the other hand, households turn over their portfolios with striking regularity, almost fully turning over their portfolios each year in the large discount brokerage data of Barber and Odean (2000). We can simultaneously observe households trading in the direction of momentum while displaying a neutral tendency with respect to momentum in their ownership. Moreover, to the extent that both stock characteristics and beliefs about the returns associated with particular characteristics fluctuate over time, shifting household ownership and retail trader behavior can be partly explained by attention and its interaction with these effects as opposed to wild swings in underlying sentiment about stocks (the unobservable characteristics).

While household beliefs about observable characteristics are a major object of interest in the literature and this paper, the primary focus of this paper's empirical analysis is the unobservable characteristic of stocks, which I call "household sentiment." Conceptually, this object is the average household's expectation about the additional returns they will receive for investing in a particular stock, independent of its observable characteristics. In other words, if we took a cheap-appearing momentum stock with a high beta, and took away households' beliefs about value, momentum, and beta, what is the remaining stock-specific (idiosyncratic) return expectation for this stock? Much of the literature is dedicated to debating whether households have skill in *trading*, but here we are studying whether households have skill in stock *evaluation*.

The demand estimations reveal that household portfolios benefit from stock-specific sentiment but are hurt by beliefs about macroeconomic factors. I construct portfolios that (a) resemble observed household holdings by quarter, (b) represent what household demand would look like if the effects of beliefs about observable characteristics were removed, and (c) represent what household demand would look like if sentiment were removed. Additionally, I test the effects of removing disagreement among households so that each household holds the same portfolio (disagreement interacts with the short sale constraints, per Miller (1977), to create demands and, in turn, prices). I find that while disagreement in and of itself appears to have little effect on aggregate household performance, sentiment and characteristic-based demand have meaningful effects. I begin by testing quarterly buy-and-hold strategies, generating quarterly portfolios at each quarter's end and then determining daily returns as if one invested at the start of the quarter and held until the start of the subsequent quarter. I compute daily alphas using the five factors of Fama and French (2015) as well as the momentum from Carhart (1997) ("FF6" model, alternatively "FF5 + UMD").

If we take the aggregate household portfolio at the start of each quarter and hold it until the end of the following quarter, then this strategy generates six-factor alpha indiscernible from zero (0.03% annualized). If we instead strip out the effects of sentiment, leaving behind only observable characteristic-based demand, the portfolio has an annualized alpha of -0.06%, which is also not statistically distinguishable from zero and is weakly negative. On the other hand, a portfolio based entirely on household sentiment with the effects of observable characteristics removed has an annualized six-factor alpha of 1.21%, which has a t-statistic of 2.22 using Newey-West standard errors.

We can test the difference between sentiment-driven portfolios and the actual portfolios of households (which add in characteristics) by creating a long-short strategy that goes long the sentiment portfolio and short the actual holdings portfolio. When averaged over the most recent four quarterly portfolios, this has an annualized alpha of 1.74%, which is highly statistically significant. The inclusion of beliefs about characteristics detracts from what appears to be modest evidence of stock-picking skill among households. Meanwhile, if we eliminate the effects of heterogeneous attention by simply having each investor form portfolios as if they paid even attention to everything, then the results are qualitatively unchanged. In other words, whereas household trading appears to be strongly and negatively influenced by attention dynamics, the performance of the household portfolio at longer horizons does not appear to be strongly influenced by attention. If we split stocks into "consumer-oriented" (B2C) and "producer-oriented" (B2B) groups on the basis of NAICS codes, we see that the positive boost to performance is indeed driven by selection of B2C stocks, not B2B stocks. This paper's empirical results are therefore consistent with household private information about product and business quality, but inconsistent with household sentiment being driven by biases.

This paper's demand system-based analysis of household ownership, in conjunction with the extensive literature on retail investor trading behavior, advances a nuanced view of household behavior in the stock market. Instead of being a homogeneous group of highly behavioral investors with no skill but prone to overconfidence, extrapolative expectations, overtrading, and the disposition effect, investors might possess diffuse private information and genuine skill. This skill, however, is entirely undone by market-timing beliefs, attention-driven trading, and short-term trading.

One caveat is that this paper's demand system-based approach to household investments is only able to analyze the largest holdings of households, which is to say larger-cap stocks (this paper focuses on the 300 largest holdings of households). While these stocks comprise the vast majority of household ownership in dollar terms, the remainder is still meaningful and might serve as a greater drag on household performance than the larger-cap stocks. Additionally, I am unable to quantify household losses due to short-term trading, though these have already been studied in numerous other papers, particularly Barber and Odean (2000). Nonetheless, this paper's novel approach to household demand estimation could be applied to study a larger cross section of stocks, and the empirical analysis of sentiment versus macro/factor timing beliefs illuminates different channels

versus the overtrading and attention-driven channels documented in the previous literature.

Related Literature I review this paper's contributions in the context of two strands of the financial economics and IO literature: (1) household investment wisdom and mistakes and (2) demand estimation.

Household Investment Wisdom and Mistakes A large literature has addressed whether households ("retail investors") are skilled, informed investors whose trades positively predict future returns or, alternatively, unskilled market participants reacting on impulse to swings in attention, thus earning worse returns than they would have achieved via passive investing. Investors have been shown to overtrade (Barber and Odean (2000), Odean (1999), and Barber, Lee, Liu and Odean (2009)), display overconfidence (Barber and Odean (2001)), and react in a contrarian way to earnings announcements in a way that generates post earnings announcement drift (Luo, Ravina, Sammon and Viceira (2025)). Moreover, retail traders are driven to trade by attention and the media (see e.g. Barber and Odean (2008), Da, Engelberg and Gao (2011), Engelberg and Parsons (2011), Engelberg, Sasseville and Williams (2012), and Graves (2023)). However, the literature has also suggested that retail investors display potential skill or, at the very least, that retail trades appear to predict short-term future returns with a positive sign. Kaniel, Saar and Titman (2008), Kaniel, Liu, Saar and Titman (2012), Kelley and Tetlock (2013), and Boehmer, Jones, Zhang and Zhang (2021), among others, provide evidence that retail flows positively predict returns over short horizons, including leading up to earnings announcements. Meanwhile, Coval, Hirshleifer and Shumway (2021) argue that a subgroup of retail investors display persistent skill and are able to consistently beat the market. Barber, Lin and Odean (2024) attempt to reconcile these two seemingly contradictory strands of the literature, arguing that in the cross section, retail trades produce short-term return predictability, but when these trades are weighted by dollar volume, retail trades negatively predict short-horizon returns.

This paper does not study retail *trades*, but rather *holdings*, and in particular, it does so at a quarterly level. I do not study short-horizon return predictability, so this paper neither confirms nor contradicts any of the past literature that studies retail trades at a short horizon. Instead, this paper aims to study how different components of household investor beliefs drive their investment performance, independently of the effects of overtrading and attention fluctuations that have been previously documented. This paper's methodology implicitly dispenses with within-quarter round-trip trading by households where a stock is purchased and sold within a short period. We are therefore left with only long-horizon beliefs and their corresponding contributions to portfolio returns. By studying different objects from those previously studied, I seek to contribute a new lens through which to view retail trader portfolios.

Demand Estimation Methodologically, this paper's main contribution is a new approach to asset demand estimation that applies to retail investors. Koijen and Yogo (2019) initiated the recent literature on Demand System Asset Pricing (DSAP) by proposing a logit-based model of institutional stock holdings that they estimate using 13F data. Under the assumptions of their

model, they are able to study counterfactual market prices, latent demand for stocks, and which investors are responsible for volatility in the financial markets. They find return predictability in the latent demand of investors, showing a mean-reversion phenomenon where high latent demand leads to lower returns going forward. This paper builds on the insights of Koijen and Yogo (2019) regarding the connection between between latent demand and return predictability.

Like Koijen and Yogo (2019), this paper uses demand estimation to study questions relevant to financial economics. Yet this paper features multiple major differences: (a) I exclusively study household investors as opposed to institutional investors, with this paper lying at the intersection of household finance, behavioral finance, and asset pricing; (b) I model households as a heterogeneous continuum using a random coefficients specification of beliefs as opposed to as a "household sector" like in Koijen and Yogo (2019) (wherein households operate as if they are a single institutional investor); (c) the purpose of this paper is to study the mistakes and skill of household investors as opposed to the demand of institutional investors; and (d) this paper's notion of "household sentiment" is identified as a belief variable under the assumptions made throughout, whereas the "latent demand" object of Koijen and Yogo (2019) can correspond to multiple different forces, one of which is beliefs. This paper's treatment of the household investors' perceived covariance matrix as diagonal would not be appropriate in the context of institutional investors, but is plausible for household investors based on their empirically observed behavior (e.g. extreme portfolio concentration and limited capacity to integrate gambles); likewise, the demand system treatment of large investors in Koijen and Yogo (2019) does not map into a context of a pseudo-continuum of heterogeneous household investors.

The demand estimation methodology used in this paper represents a significant departure from the existing financial economics literature (see e.g. Gabaix et al. (2022), Koijen and Yogo (2023), Koijen and Yogo (2026), and Haddad et al. (2025)), yet is in keeping with a long tradition in the IO literature that dates back to Berry (1994) and, in particular, Berry, Levinsohn and Pakes (1995) ("BLP"). BLP serves as the foundational model for much of new empirical IO and similarly serves as the primary motivation for this paper's model. This paper, however, features several differences: in this paper, households are not solving a discrete choice problem but rather a portfolio allocation problem, and households each face short sale constraints and leverage constraints (they cannot borrow on margin or, alternatively, have restrictions on margin borrowing). I show theoretically how to adapt the general approach of BLP to this paper's setting, while still generating an invertible demand system that satisfies the connected substitutes property of Berry, Gandhi and Haile (2013) that gives us a one-to-one mapping between a particular vector that characterizes household investors' stock-specific sentiments and a given set of demands, conditional on other belief parameters. The methodology used in this paper is based on the GEL-MPEC approach proposed by Conlon (2013), which to the best of my knowledge has only been used empirically in Conlon (2012). Additionally, the notion of "consideration sets" used in this paper is a direct analogue of the approach of Goeree (2008), wherein consumers are only aware of a subset of PCs in making their purchase decision (modulated by advertising). Instead of advertising influencing

consideration set selection, this paper's baseline model uses the attention model of Graves (2023) as an approximation for investor attention: stocks that are predicted to have large dollar volume of both purchases and sales are "high attention" stocks, with a higher probability of being on the radar of a given household investor.

Given that this paper seeks to extract beliefs, it is related as well to Egan, MacKay and Yang (2022), which uses leveraged funds and fee shifters to separately identify risk aversion and return expectations. Whereas Egan, MacKay and Yang (2022) studies the stock market as a whole, this paper focuses on beliefs about the cross-section of returns and therefore uses a different identification strategy via instruments and functional form assumptions in the absence of cost shifters.

This paper can be viewed as the companion paper of Graves (2025) despite its different methodology and research questions. Methodologically, Graves (2025) addresses the asset demands of institutional investors and how to estimate them in the presence of short sale constraints and data censoring (in 13F filings), which pose significant technical challenges. This paper addresses the demands of households and how to estimate them when (a) households are heterogeneous, (b) only approximate aggregate holdings are available, and (c) households face short sale and leverage constraints. Graves (2025) studies the idiosyncratic beliefs of institutional investors, particularly "hidden beliefs" that are embedded within a zero holding (absence from a 13F filing). This paper does not study "hidden beliefs" and "overt beliefs" but rather the "sentiment" of household investors, which can be viewed as an analogous object to the idiosyncratic return expectations decomposed into "hidden" and "overt" beliefs in Graves (2025).

Outline This paper proceeds as follows. Section 2 derives a model of household investor demand and explores its theoretical properties. Section 3 applies a GEL-MPEC approach to show how to estimate this model. Section 4 discusses the required instruments, data construction, and computational approach. Section 5 reviews the estimation results and explores how household beliefs have changed over time. Section 6 studies the skill level of households and the origins of household investment mistakes. Section 7 concludes.

2 A Model of Household Asset Demand

I construct an asset demand model that features atomistic households, then conclude the section by showing how if we add in non-atomistic institutional investors we can construct a demand system that is consistent with market clearing and existence of a unique equilibrium. Households possess heterogeneous beliefs based on observable asset characteristics as well as unobservable characteristics. Using a basic return factor structure following Carhart (1997) and Fama and French (2015), I derive households' optimal portfolio choice under short sale constraints and restricted leverage. Under an assumption based on the "narrow framing" result of Barberis, Huang and Thaler (2006), household optimal demand is shown to be equal to the product of (1) an asset-specific scaling factor equal to the inverse of asset variance and (2) the sum of return beliefs minus

a belief-based constant, censored at zero.

2.1 Environment

The basic environment follows much of the setup of Koijen and Yogo (2019), Graves (2025), the recent DSAP literature, and many other strands of the asset pricing literature. Time is discrete, $t = \{0, 1, ..., T\}$. $N + 1$ assets are traded each period, including one risk-free asset in perfectly elastic supply ($n = 0$) and N risky assets ($n = 1, ..., N$). We denote the set of assets by $\mathcal{N}$. Each risky asset n is in fixed finite supply $S_t(n)$ at time t, pays dividends $D_t(n)$ at time t, and has price $P_t(n)$. Gross returns between $t-1$ and t are expressed as $R_t(n) = \frac{P_t(n)+D_t(n)}{P_{t-1}(n)}$ and market equity is simply $\text{ME}_t(n) = S_t(n) P_t(n)$. Following Graves (2025), I normalize $S_t(n) = 1 \; \forall n, t$ in all theory sections and apply appropriate normalization in all empirical work, as market capitalization is the only relevant variable for the model, not price per share. Market equity for stock n is therefore $P_t(n)$ throughout the paper.[1]

Each asset n has observed characteristics $x_{k,t}(n)$, $k = 0, ...K$, where $x_{K,t}(n) = 1$ is simply a constant. $\omega_{h,t}(n)$ are unobserved characteristics unique to household h. For households, we will express $\omega_{h,t}(n) = \xi_t(n) + \sigma_{\epsilon,t}\epsilon_{h,t}(n)$ where ξ_t represents average household sentiment and $\sigma_{\epsilon,t}\epsilon_{h,t}(n)$ represents household h's disagreement with the average ($\sigma_{\epsilon,t}\epsilon_{h,t}(n) \sim N\left(0, \sigma^2_{\epsilon,t}\right)$). We can therefore express household h's time t information set as a $K+2$-dimensional vector consisting of

$$x^h_t(n) = \begin{bmatrix} x_{0,t}(n) \\ x_{1,t}(n) \\ \vdots \\ x_{K,t}(n) \\ \omega_{h,t}(n) \end{bmatrix}.$$

Letting $w_{h,t} \in \mathbb{R}^N$ denote the vector of portfolio weights for the risky assets, we have that wealth $A_{h,t+1} \in \mathbb{R}^+$ in period $t+1$ for agent h is equal to

$$A_{h,t+1} = A_{h,t}\left(w'_{h,t}R_{t+1} + \left(1 - 1'w_{h,t}\right) R_{t+1}(0)\right). \tag{1}$$

The portfolio problem households solve is given by

$$\max_{w_{h,t} \in \mathcal{S} \subseteq \mathbb{R}^N} \mathbb{E}_{h,t}\left[\log\left(A_{h,T}\right)\right]$$

where $\mathcal{S}$ is a convex set in $N_{\mathcal{S}} \leq N$ dimensions that is determined by leverage constraints, short sale constraints, and a consideration set that restricts assets available for trading. This log utility portfolio problem, as is well established in the literature, results in myopic portfolio choice in each

[1] Retail investors are particularly prone to behavioral biases related to the per-share price of stocks; this paper abstracts from such mechanisms, though they are empirically relevant for price determination. For a more in depth exploration of these tendencies, see e.g. Kumar (2009), Baker, Greenwood and Wurgler (2009), Birru (2015), Birru and Wang (2016), and Cox, Van Ness and Van Ness (2022).

period. We could just as easily microfound the model using classic mean-variance optimization with coefficient $\gamma = 1$ and use returns and covariance of returns instead of the logarithmic versions.

2.2 Household Portfolio Problem

Economists do not see household holdings of the aggregate stock market or of products like ETFs, mutual funds, and index funds. Microdata regarding labor income, money market holdings, bank balances, illiquid assets, bequests, social security, and other factors that weigh into a household's decision making regarding stock investments are similarly elusive. One well established fact, however, is that household single-name equity investments are consistently concentrated beyond what one would expect via standard portfolio optimization theory unless households only invest in single-name equities about which they have extraordinarily high return expectations versus variance. Here I establish a portfolio model of household investment that incorporates these empirical facts as well as the common constraints facing retail investors, namely short sale and leverage constraints.

Household $h \in [0, 1]$ selects an optimal portfolio weight vector $w_{h,t} \in \mathbb{R}^N$ by solving the first order approximation of the following constrained optimization problem:

$$\max_{w_{h,t}(n) \geq 0,\ \sum_n w_{h,t}(n) \leq 1,\ w_{h,t}(n) = 0 \forall n \notin \mathcal{C}_{h,t}} \mathbb{E}_{h,t} \left[\log \left(A_{h,T} \right) \right]. \tag{2}$$

$\mathcal{C}_{h,t} \subseteq \mathcal{N}$ is the consideration set of household h: this is the set of stocks that h is paying attention to in period t. Let $\mu_{h,t}$ and $\Sigma_{H,t}$ denote the mean and covariance of log excess returns over the risk-free asset conditional on the time t information set of households. Throughout, we assume for parsimony that each household has the same covariance expectations given by $\Sigma_{H,t}$, which we will later assume, for our empirical sections, is based on recent log return covariance over τ months.

In the model retail investors make a simplification that bears a resemblance to the "narrow framing" of Barberis, Huang and Thaler (2006) wherein agents do not fully integrate gambles: instead of optimizing their portfolio by considering the correlations across assets, they assume no correlation and invest by heuristically "diversifying" across opportunities as if the covariance matrix were diagonal. Using the first-order approximation of the optimal solution and the fact that the expected log return covariance matrix $\Sigma_{H,t}$ is diagonal, we can express portfolio weights for household h (see Appendix A) as

$$w_{h,t} = \Sigma_{h,t}^{-1} \max \left\{ 0, \mu_{h,t} - \lambda_{h,t} \mathbf{1} \right\} \tag{3}$$

where the maximum operation is taken elementwise over the vector and

$$\lambda_{h,t} = \max \left\{ \frac{\sum_{n \in C_{h,t} : \Lambda_{h,t}(n) = 0} \sigma_t^{-2}(n) \, \mu_{h,t}(n) - 1}{\sum_{n_1 \in C_{h,t}, \Lambda_{h,t}(n_1) = 0} \sigma_t^{-2}(n_1)}, 0 \right\} \tag{4}$$

where $\{n \in C_{h,t} : \Lambda_{h,t}(n) = 0\}$ is the set of all stocks n such that the short sale constraint is slack

(the corresponding Lagrange multiplier is zero).

2.3 Consideration Sets

Consideration set $\mathcal{C}_{h,t}$ is drawn from a distribution with independent asset inclusion probabilities given by $g_t \in [0,1]^N$ (multinomial resampling). The full measure of households then contains every possible subset of assets excluding those with $g_t(n) = 0$. Consideration set $\mathcal{C}_{h,t}$ has probability weight

$$\Pr(\mathcal{C}_{h,t}) = \prod_n g_t(n)^{1_{n\in\mathcal{C}_{h,t}}} (1 - g_t(n))^{1_{n\notin\mathcal{C}_{h,t}}}. \tag{5}$$

This probability weighting scheme interacts with the diagonal covariance matrix assumption in a subtle way; with full information and a diagonal covariance matrix, households would just "diversify away" all of their risk by investing in virtually every stock about which they have a modestly positive return expectation. With narrow framing but a limited consideration set, households do engage in a modicum of diversification but nonetheless take concentrated positions in their highest-conviction stocks; this type of portfolio is consistent with the empirical evidence.[2]

2.4 Household Beliefs: Random Coefficients

Households do not just possess heterogeneous consideration sets; they also have different characteristic-based return expectations and asset-specific sentiment. I focus on the case where all dimensions of heterogeneity are independent for ease of exposition, though the model can be suitably modified to allow for correlations at the cost of additional notation.[3] I use a random coefficients model of expected log excess return beliefs in the spirit of the industrial organization literature's discrete choice models. Specifically, I will assume that

$$\mu_{h,t} = \sum_{k=0}^{K} \bar{\beta}_{k,t} x_{k,t} + \xi_t + \sum_{k=0}^{K} \sigma_{k,t} x_{k,t} \nu_{k,h,t} + \sigma_{\epsilon,t} \epsilon_{h,t} \tag{6}$$

where $\epsilon_{h,t} \in \mathbb{R}^N$ and $\nu_{h,t} = (\nu_{0,h,t}, \nu_{1,h,t}, \ldots, \nu_{K,h,t})' \in \mathbb{R}^{K+1}$ are vectors capturing household-specific sentiment and characteristic-based beliefs respectively. $\xi_t \in \mathbb{R}^N$ contains a vector of unobservable (to the econometrician) characteristics. Throughout the rest of the paper we will

[2] If we consider the case of a household investor following standard mean-variance optimization while paying attention to all stocks (the so-called FIRE assumption or "Full Information Rational Expectations"), then we struggle to match the empirical data wherein households frequently only own several stocks (see e.g. Barber and Odean (2000)). Such concentrated portfolios despite diversification benefits are only consistent with rational investors who possess extreme beliefs, and such portfolios with two to three stocks are generally not observed among institutional investors who are presumably more "rational." When attention is limited and retail investors only focus on a small group of ten to twenty stocks, then idiosyncratic risk dominates market beta as a source of portfolio volatility when households hold positive views of only a subset of these stocks. If stocks possess betas that fall in a generally narrow range (as is empirically observed), then the narrow framing assumption is not far from a more standard "rational" benchmark that incorporates attention effects.

[3] The model's tractability is not altered by alternative assumptions about correlations across dimensions of beliefs, though we would require either (a) assumptions about these correlations or (b) additional moments, preferably from micro data, that would pin down these correlations.

call ξ_t "household sentiment" since it represents the average return beliefs of households that are not driven by observable factors. Note that the objects $\mu_{h,t}$ and $\epsilon_{h,t}$ are N-dimensional and therefore include beliefs about stocks that are not in agent h's consideration set $\mathcal{C}_{h,t}$. This inclusion is by design: the belief vectors are latent objects that provide the return expectations that would exist for stock n if agent h's attention to stock n were switched on. It is beyond the scope of this paper to evaluate mechanisms by which the process of attention activation will in and of itself influence opinions of a given stock; this paper's belief specification nonetheless accommodates rich heterogeneity in beliefs about stocks and characteristics across households. We now proceed to the setup for aggregating household demand.

2.5 Total Household Asset Demand

We do not know the amount held in the outside asset, which in this case is the risk-free asset that is in perfectly elastic supply. Likewise, I do not model the choices of households to invest in single-name equities or index funds. I therefore take as a given the amount invested in single-name equities in a given period, and do not seek to model a household's mutual fund selection process or broader asset allocation decisions.[4]

Let $A_{H,t}$ be the total assets of households devoted to stocks within single-stock portfolios (which includes the outside asset) and define $\mathcal{C}_{h,t}$ to be a diagonal matrix such that the nth diagonal element is 1 if $n \in C_{h,t}$ and 0 otherwise. We then have that the total dollar amount invested in each equity is given by

$$A_{H,t} \int C_{h,t} \Sigma_{H,t}^{-1} \max\left\{0, \mu_{h,t} - \lambda_{h,t}\mathbf{1}\right\} dh \tag{7}$$
$$= A_{H,t} \int C_{h,t} \Sigma_{H,t}^{-1} \left(\max\left\{ 0, \sum_{k=0}^{K} \bar{\beta}_{k,t} x_{k,t} + \xi_t + \sum_{k=0}^{K} \sigma_{k,t} x_{k,t} \nu_{k,h,t} + \sigma_{\epsilon,t} \epsilon_{h,t} - \lambda_{h,t}\mathbf{1} \right\} \right) dh.$$

The integration over households h is with respect to multiple independent dimensions: the consideration set $C_{h,t}$, the idiosyncratic demand preferences $\epsilon_{h,t}$, and the idiosyncratic characteristic preferences $\nu_{k,h,t}$. Because of the diagonal covariance matrix, we can express the aggregate household demand for asset n as

$$A_{H,t} \int_{h:n \in C_{h,t}} \sigma_t^{-2}(n) \left(\max\left\{ 0, \sum_{k=0}^{K} \bar{\beta}_{k,t} x_{k,t}(n) + \xi_t(n) + \sum_{k=0}^{K} \sigma_{k,t} x_{k,t}(n) \nu_{k,h,t} + \sigma_{\epsilon,t} \epsilon_{h,t}(n) - \lambda_{h,t}\mathbf{1} \right\} \right) dh. \tag{8}$$

Note that this equation contains a censored function of many variables, including a Lagrange multiplier, on the right hand side. This equation has no analytic solution and instead must be computed via numerical methods once we specify the parameters.

[4] Understanding how households allocate between actively managed products, index funds, bonds, commodity funds, ETFs, Master Limited Partnerships, private equity, international versus domestic equity, and currencies is an expansive topic for future household finance research that I do not address.

2.6 Toward a Full Asset Demand System (Equilibrium)

We now have a demand function for households that expresses total demand as a function of beliefs and preferences. If we add mean-variance optimizing institutional investors with wealth $A_{i,t}$ and assume that the first characteristic is price ($x_{0,t} = P_t$), then given a specific set of characteristics and preferences, our aggregate demand function in dollars is

$$\begin{aligned} X(P_t) &= \sum_i A_{i,t} w_{i,t} + A_{H,t} \int_h w_{h,t} dh \\ &= \sum_i A_{i,t} D_{i,t}^{-1} \left(P_t \left(\beta_{0,i,t} - \kappa_{i,t} \Psi_{0,i,t} \right) + \sum_{k=1}^{K} x_{k,t} \left(\beta_{k,i,t} - \kappa_{i,t} \Psi_{k,i,t} \right) + \epsilon_{i,t} \right) \\ &+ A_{H,t} \int C_{h,t} \Sigma_{H,t}^{-1} \left(\max \left\{ 0, \left(\bar{\beta}_{0,t} + \sigma_{0,t} \nu_{0,h,t} \right) P_t + \sum_{k=1}^{K} \bar{\beta}_{k,t} x_{k,t} + \xi_t + \sum_{k=1}^{K} \sigma_{k,t} x_{k,t} \nu_{k,h,t} + \sigma_{\epsilon,t} \epsilon_{h,t} - \lambda_{h,t} \mathbf{1} \right\} \right) dh \\ &= J(P_t) + H(P_t) \end{aligned} \tag{9}$$

where $J(P_t) \in \mathbb{R}^N$ and $H(P_t) \in \mathbb{R}^N$ denote the aggregate dollar demand vectors for institutions and households respectively (we have for now excluded the leverage constraint for households) and we have followed the notation of Koijen and Yogo (2019). The equilibrium condition is then given by the fixed point $X(P_t) = P_t$.

Definition 1. An equilibrium at time t given by $\mathcal{E} = \left\{ P_t, \{s_{j,t}\}_{j=1}^{J}, \{s_{h,t}\}_h \right\}$ given parameters consists of

1. Market clearing: share demands $\{s_{j,t}\}_{j=1}^{J}$ for institutions and household share demands $s_{h,t}$ for $h \in [0,1]$ such that
$$\sum_j s_{j,t} + \int_0^1 s_{h,t} dh = 1 \tag{10}$$
2. Optimality: given market capitalizations P_t, the share demands for households and institutions solve the respective household and institutional optimization problems previously described conditional on parameters.

We make the following assumption on the factor structure of expected covariance for each institution.

Assumption 1. $\Psi_{0,i,t} = 0 \ \forall i$.

For this assumption to hold, institutions must believe that the common factor in the log return covariance matrix is not driven by market capitalization. The ubiquitous model that uses the market beta of recent returns to determine common risk, for example, is one of many such models that satisfies this assumption.

Our next assumption establishes a condition on the slope of institutional and household demand curves. Let $\|x\|_\infty$ be the usual max norm $\max_n |x_n|$ when $x \in \mathbb{R}^N$ and let $\|A\|_\infty = \max_{x \in \mathbb{R}^N, \|x\|_\infty = 1} \|Ax\|_\infty$ denote the matrix norm induced by the l_∞ norm.

Assumption 2. $\| \sum_i A_{i,t} \beta_{0,i,t} \Sigma_{i,t}^{-1} \|_\infty + \| A_{H,t} \int 1_{n \in \mathcal{C}_{h,t}} \Sigma_{H,t}^{-1} \mathbf{1} | \bar{\beta}_{0,t} + \sigma_{0,t} \nu_{0,h,t} | dh \|_\infty < 1$

The intuition behind the condition is simply that excessively sloped aggregate demand curves result in demand tending toward infinity at a faster rate than prices. If institutions in aggregate, weighted by assets under management and asset return precision expectations, have exceptionally upward sloping demand curves, then as prices tend toward infinity, demand continues to increase to infinity at a faster rate than prices.

Proposition 1. *An equilibrium always exists and is unique under Assumptions (1) and (2) if households are not leverage constrained. If we both (a) add back the leverage constraint for households and (b) assume institutions are also short sale constrained, then equilibrium exists under these assumptions.*

The proof of Proposition 1, given in Appendix B, shows how Assumption 2 is sufficient for the invocation of Banach's fixed point theorem. This constructive proof provides us with a ready method for counterfactual analysis, as we can modify objects such as attention or sentiment and iterate the aggregate demand function X to obtain counterfactual prices. Assumption 2 is a minimally restrictive condition that allows for both downward sloping and upward sloping demand by institutions and households.

It is beyond the scope of this paper to study equilibrium notions, conduct counterfactual analyses, or estimate variance contributions in the spirit of Koijen and Yogo (2019). I analyze equilibrium to show that this paper's model of households as a heterogeneous continuum can be integrated with models of institutional demand to generate a comprehensive demand system that is still consistent with equilibrium notions. The precise model of institutional demand (e.g. logit as in Koijen and Yogo (2019) or short sale constrained with max operators as in Graves (2025)) will determine what form of equilibrium existence and uniqueness results we can establish.

3 Estimation via GEL-MPEC

The objective of household demand estimation is to recover household sentiment dispersion, characteristic preference dispersion, and average beliefs about observable characteristics, represented by the parameter vector $\theta_t = \left(\sigma_{\epsilon,t}, \{\sigma_{k,t}\}_{k=1}^K, \bar{\beta}\right)$. We additionally wish to obtain the vector of unobservable asset characteristics ξ_t, an object that represents average household sentiment about assets. This section describes the estimation process, which draws from the extensive Industrial Organization literature on differentiated product demand estimation in the spirit of Berry, Levinsohn and Pakes (1995). Given aggregate household demand $\hat{H}$, observable asset characteristics, prices P_t, and a vector of instruments, we can theoretically recover the

parameter vector θ_t. First, for alignment with BLP and much of the IO literature, I show how to recover θ_t and ξ_t via a traditional NFXP/GMM approach. Due to computational challenges associated with estimation in this paper's empirical context, however, we must use an alternative estimation framework; I show how to use the GEL-MPEC paradigm of Conlon (2013) to estimate the parameters by recasting the problem. Subsection 3.1 establishes that the fixed point algorithm converges via Banach's Fixed Point Theorem, subsection 3.2 discusses the crucial mean independence identifying assumption for unobservable sentiment, subsection 3.3 reviews identification in light of recent work on connected substitutes in industrial organization, subsection 3.4 details the traditional NFXP/GMM algorithm used to find the estimates $\hat{\theta}_t$ as well as ξ_t, and subsection 3.5 describes the GEL-MPEC setup.

3.1 Obtaining Unobserved Characteristics given Demand Parameters

Recall that aggregate household demand is given by

$$H\left(P_t, \theta_t\right) = A_{H,t} \int C_{h,t} \Sigma_{H,t}^{-1} \left(\max \left\{ 0, \left(\bar{\beta}_{0,t} + \sigma_{0,t} \nu_{0,h,t} \right) P_t \right.\right. \\ \left.\left. + \sum_{k=1}^{K} \bar{\beta}_{k,t} x_{k,t} + \xi_t + \sum_{k=1}^{K} \sigma_{k,t} x_{k,t} \nu_{k,h,t} + \sigma_{\epsilon,t} \epsilon_{h,t} - \lambda_{h,t} \mathbf{1} \right\} \right) dh \tag{11}$$

so that given aggregate household holding data $\hat{H}$ the estimator $\hat{\theta}_t$ must solve $\hat{H} = H\left(P_t, \hat{\theta}_t\right)$. Let $\delta\left(n\right) = \sum_{k=0}^{K} \bar{\beta}_{k,t} x_{k,t}\left(n\right) + \xi_t\left(n\right)$ be the average characteristics of asset n, including unobservables, and let δ be the corresponding vector. Multiplying through both sides of $\hat{H} = H\left(P_t, \hat{\theta}_t\right)$ by $A_{H,t}^{-1} \Sigma_{H,t}$ and adding δ, we obtain via rearrangement

$$\delta = f\left(\delta\right) \\ = \delta + A_{H,t}^{-1} \Sigma_{H,t} \hat{H} - \int C_{h,t} \left(\max \left\{ 0, \delta + \sum_{k=0}^{K} \sigma_{k,t} x_{k,t} \nu_{k,h,t} + \sigma_{\epsilon,t} \epsilon_{h,t} - \lambda_{h,t} \mathbf{1} \right\} \right) dh. \tag{12}$$

Given a particular parameter vector $\tilde{\theta}_t$, we will be able to readily obtain the corresponding δ provided that the function f defined by equation (12) is a contraction mapping. The following proposition, whose proof is in Appendix C, establishes that the Banach Fixed Point theorem can be applied to f.

Proposition 2. *Suppose that households are subject to the leverage restriction $\sum_n w_{h,t}\left(n\right) \leq L_{H,t}$ in addition to short sale constraints and hold portfolios $w_{h,t} = \Sigma_{H,t}^{-1} \max\left\{0, \mu_{h,t} - \lambda_{h,t}\mathbf{1}\right\}$ where $\lambda_{h,t}$ is the Lagrangian multiplier corresponding to the shadow value of relaxing the leverage constraint. Let $f\left(\delta\right) = \delta + A_{H,t}^{-1}\Sigma_{H,t}\hat{H} - \int C_{h,t}\left(\max\left\{0, \delta + \sum_{k=0}^{K}\sigma_{k,t}x_{k,t}\nu_{k,h,t} + \sigma_{\epsilon,t}\epsilon_{h,t} - \lambda_{h,t}\mathbf{1}\right\}\right) dh$. Then f is a contraction mapping.*

This proposition establishes that f is a contraction mapping in the presence of household leverage constraints, where household leverage $L_{H,t}$ can be chosen to be arbitrarily large. With this

function established, we are now able to create a nested fixed point algorithm (NFXP) to estimate the parameter vector.

3.2 Identifying Assumptions

The conditional moment restrictions are

$$\mathbb{E}\left[\xi_{n,t}|z_{n,t}\right] = 0 \tag{13}$$

where $\xi_{n,t} = \xi_t(n)$ for notational convenience and $z_{n,t}$ being a vector of instruments including an instrument for prices $\hat{P}_t$ characterized in Section 4 as well as exogenous characteristics. Because all characteristics-based demand models in this paper include the constant vector $x_{K,t} = \mathbf{1}$, the normalization to 0 in equation (13) is without loss of generality. These restrictions require that the unobservable characteristic for each asset be mean independent of the instruments. Average household sentiment is thus required to be orthogonal to asset characteristics and instruments. Unlike in the Industrial Organization literature, firms cannot choose product characteristics: given the opportunity, every firm would elect to have high profitability, strong price momentum, or a ballooning book value. The assumption that non-price characteristics are exogenously given, while clearly imperfect, makes for a less troubling identifying assumption than when observed product characteristics are continuously modified to suit consumer tastes. Although I repeatedly mention "unobservable characteristics" of assets for consistency with the IO literature, a more precise definition of $\xi_{n,t}$ would be the average private household information (beliefs) about next-period returns of asset n, which should be mean independent of publicly observable data.

Our final assumption is that the coefficient on prices $\bar{\beta}_0$ be negative: this essentially amounts to a minimally restrictive assumption that financial assets are not Giffen goods with respect to the economy in aggregate and households in particular.

Assumption 3. *Households believe returns will be higher when purchasing an asset for a cheaper price:* $\bar{\beta}_{0,t} < 0$ *and* $\sigma_{0,t} = 0$.

3.3 Identification, Invertibility, and Connected Substitutes

Because of the narrow framing of households, where the covariance matrix in the optimization problem is diagonal, we can readily see that the "connected substitutes" property of Berry, Gandhi and Haile (2013) is satisfied: each good substitutes to every other good in that a decrease in the price of asset k causes an increase in the demand for asset k and a decrease in the demand for every other asset $j \neq k$ through the budget constraint, as described in the equilibrium proof in Appendix B. Although our model is not one of discrete choice, the inversion and identification proofs of Berry, Gandhi and Haile (2013) and Berry and Haile (2014) apply here with minimal adaptation. The requisite ingredients for identification are all present: a linear index restriction, scalar unobservables, a sufficient number of instruments, and a mean independence assumption.

Under the conditions of the model, the main theorem of Berry and Haile (2014) provides that the household demand system is identified for all t and the household sentiment vector ξ_t is identified almost surely for all t.

Theorem 1. *Suppose that the household covariance matrix Σ is diagonal ("narrow framing"). Then the demand system satisfies the connected substitutes property of Berry, Gandhi and Haile (2013) and the demand system is invertible.*

Proof. I will show connected substitutes and apply the invertibility result of Berry, Gandhi and Haile (2013). First we can establish a strengthened version of the weak substitutes property: demand H_j is strictly decreasing in ξ_l for all $j \neq l$. We will start from an initial vector of sentiment and perturb the results. Take household h with optimal weights given by w_{h0} and Lagrange multiplier λ_{h0} for the leverage constraint, which might be zero. Given initial ξ_{h0} we have the main component of the integrand (we can ignore the rest) given by $\max\left\{0, \sum_{k=0}^{K} \bar{\beta}_{k,t} x_{k,t} + \xi_t + \sum_{k=0}^{K} \sigma_{k,t} x_{k,t} \nu_{k,h,t} + \sigma_{\epsilon,t} \epsilon_{h,t} - \lambda_{h,0} \mathbf{1}\right\}$. Now suppose we increase the lth component of ξ_{h0} so that $\xi_{h1} = \xi_{h0} + \Delta e_l$ where e_l is the $N \times 1$ basis vector of all zeros with a 1 in the lth row and Δ is a small positive number. Then the mean-variance optimal weights have shifted from $\max\left\{0, \sum_{k=0}^{K} \bar{\beta}_{k,t} x_{k,t} + \xi_t + \sum_{k=0}^{K} \sigma_{k,t} x_{k,t} \nu_{k,h,t} + \sigma_{\epsilon,t} \epsilon_{h,t} - \lambda_{h,0} \mathbf{1}\right\}$ to $\max\left\{0, \sum_{k=0}^{K} \bar{\beta}_{k,t} x_{k,t} + \xi_t + \sum_{k=0}^{K} \sigma_{k,t} x_{k,t} \nu_{k,h,t} + \sigma_{\epsilon,t} \epsilon_{h,t} - \lambda_{h,1} \mathbf{1}\right\}$, with a difference of

$$\lambda_{h,0} - \lambda_{h,1}$$

for stocks $j \neq l$ when $\lambda_{h,0}$ and $\lambda_{h,1}$ are both non-zero and the stock is in the portfolio (otherwise the change is equal to zero for other elements of the vector). We need not consider the measure zero cases where $\lambda_{h,1} > 0$ and $\lambda_{h,0} = 0$. We have that $\lambda_{h,0} - \lambda_{h,1}$ is equal to

$$\sum_{n \in C_{h,t} : \Lambda_{h,t}(n) = 0} \left(\frac{\sigma_t^{-2}(n) \left(\mu_{h,0}(n) - \mu_{h,1}(n) \right)}{\sum_{n_1 \in C_{h,t}, \Lambda_{h,t}(n_1) = 0} \sigma_t^{-2}(n_1)} \right)$$

but $\mu_{h,0}(l) - \mu_{h,1}(l) = -\Delta < 0$ for stock l and otherwise for $n \neq l$, $\mu_{h,1}(n) - \mu_{h,0}(n) = 0$. We therefore note trivially that the difference in the vectors

$$\max\left\{0, \sum_{k=0}^{K} \bar{\beta}_{k,t} x_{k,t} + \xi_{h0} + \sum_{k=0}^{K} \sigma_{k,t} x_{k,t} \nu_{k,h,t} + \sigma_{\epsilon,t} \epsilon_{h,t} - \lambda_{h,0} \mathbf{1}\right\}$$

and

$$\max\left\{0, \sum_{k=0}^{K} \bar{\beta}_{k,t} x_{k,t} + \xi_{h1} + \sum_{k=0}^{K} \sigma_{k,t} x_{k,t} \nu_{k,h,t} + \sigma_{\epsilon,t} \epsilon_{h,t} - \lambda_{h,1} \mathbf{1}\right\}$$

is given by

$$\Delta e_l + \sum_{n \in C_{h,t} : \Lambda_{h,t}(n) = 0} \left(\frac{\sigma_t^{-2}(n) \left(\mu_{h,0}(n) - \mu_{h,1}(n) \right)}{\sum_{n_1 \in C_{h,t}, \Lambda_{h,t}(n_1) = 0} \sigma_t^{-2}(n_1)} \right).$$

This sum is clearly greater than 0 for $n = l$ because

$$\Delta + \sum_{n \in C_{h,t}: \Lambda_{h,t}(n)=0} \left(\frac{\sigma_t^{-2}(n)\left(\mu_{h,0}(n) - \mu_{h,1}(n)\right)}{\sum_{n_1 \in C_{h,t}, \Lambda_{h,t}(n_1)=0} \sigma_t^{-2}(n_1)} \right) = \Delta - \frac{\sigma_t^{-2}(l)\,\Delta}{\sum_{n_1 \in C_{h,t}, \Lambda_{h,t}(n_1)=0} \sigma_t^{-2}(n_1)} > 0$$

but less than zero for $n \neq l$ as already established.

The continuum of households has beliefs that are continuously distributed, so some agents will reach their leverage constraint. Substitution to the outside good is therefore established. We have thus already shown Assumption 3 of Berry, Gandhi and Haile (2013) (connected strict substitution) and the proof is complete. □

3.4 Procedure for Computation of Estimator

We have already seen in subsection 3.1 how to compute via NFXP the unobserved characteristics of the various products ξ_t given a set of parameters $\theta_t = \left(\sigma_{\epsilon,t}, \{\sigma_{k,t}\}_{k=1}^{K}, \bar{\beta}\right)$. We thus define the function

$$G^{N,t}(\theta) = \mathbb{E}\left[H_n(z)\, T(z_n)\, \xi_{n,t}(\theta, P)\right]$$

where our moment conditions ensure that $G^N(\theta_0) = 0$. We can then follow BLP and create the criterion function

$$G_{N,t}(\theta, P) = \frac{1}{N} \sum_{n=1}^{N} H_n(z)\, T(z_n)\, \xi_{n,t}(\theta, P) \tag{14}$$

and define the estimate of θ to be

$$\hat{\theta} = \text{argmin}_\theta \, \| G_{N,t}(\theta, P) \|. \tag{15}$$

The algorithm follows the same outline as BLP:

(1) Select a given θ_0

(2) Compute ξ_t via the nested fixed point algorithm

(a) Integrate by selecting random set of consideration sets and keeping these sets fixed while changing expectations to ensure convergence

(b) Integrate using discretization of distribution of ϵ and ν values, same for each of the different consideration sets

(3) Compute the appropriate instruments and evaluate the GMM criterion function for θ_0

Last, iteration over values of θ is necessary until a global minimum is located. This can be accomplished by using a non-derivative based optimization method such as a simplex algorithm.

The issue with this approach is that for the empirical setting of this paper in which we wish to estimate demand with a large number of stocks (300 in the main specification), the Lipschitz constant is extremely close to 1. This is the underlying motivation for considering the Mathematical Program with Equilibrium Constraints approach of Su and Judd (2012) and Dubé, Fox and Su (2012) wherein we optimize over our criterion function subject to non-convex equilibrium constraint

conditions that set total quantity demanded from the model to the quantity demanded in the data.

3.5 GEL-MPEC Methodology

Following Conlon (2013), which is inspired by Dubé, Fox and Su (2012), I recast the optimization problem in 3.4. In standard GMM, we take the unconditional moment conditions $\mathbb{E}\left[\xi_{n,t} z_{n,t}\right] = 0$ and require that our estimated values for sentiment $\xi_{n,t}$ satisfy this condition. In the Generalized Empirical Likelihood (GEL) framework, we loosen this restriction by considering a probability distribution π that places weights π_n on each data observation n (call this distribution $\mathcal{Q}$) instead of the weight $\frac{1}{N}$ from the empirical distribution (call this flat distribution over N observations $\mathcal{P}$). In GEL, we require that $\sum_n \pi_n \xi_n z_n^k = 0$ instead of $\sum_n \frac{1}{N} \xi_n z_n^k = 0$ for each instrument k, with the π_n needing to be non-negative and sum to one (so that $\mathcal{Q}$ is an actual distribution). We also require that $S\left(\xi, \theta\right) = \hat{S}$ where $S\left(\xi, \theta\right)$ is the share demand of households given unobservable sentiment ξ and parameters θ and $\hat{S}$ is the empirically observed household share demands.

GEL minimizes a distance function subject to these constraints, where the distance function used determines the type of estimator. If we minimize the Kullback-Leibler divergence $K\left(\mathcal{P}, \mathcal{Q}\right)$ then we get the traditional "Empirical Likelihood" estimator. If we minimize the Kullback-Leibler divergence $K\left(\mathcal{Q}, \mathcal{P}\right)$ we obtain the "Exponential Tilting" estimator of Kitamura and Stutzer (1997). On the other hand, if we use a quadratic loss function and minimize $\sum_n \pi_n^2$, we obtain the Continuously Updated Estimator version of GMM (CUE-GMM) of Hansen, Heaton and Yaron (1996), as shown by Newey and Smith (2004). This estimator has many desirable asymptotic properties (see Newey and Smith (2004)), but does not require us to invoke the NFXP algorithm as in direct applications of GMM. Note that if $\pi_n = \frac{1}{N}$ in conjunction with the estimates for the parameters satisfies all of the constraints, then the empirical distribution and the distribution $\mathcal{Q}$ are identical and there is no distance.

The benefits of this recasting of the problem come from not needing to iterate a fixed point algorithm. Moreover, the optimization objective function is now attractively convex. Immediately, however, we note the major issue: the equilibrium constraint condition can be highly non-convex. In the next section I discuss implementation and how modern numerical solvers reduce the burden posed by the problem. Equation (16) provides the GEL-MPEC optimization problem solved in this paper.

$$\begin{aligned} \min_{\bar{\beta}, \{\sigma_k\}_{k=1}^K, \xi, \pi} & \sum_n \pi_n^2 \\ \text{s.t.} \sum_n \pi_n \xi_n z_n^k = 0, \forall k & \\ \pi_n \geq 0 \forall n, \sum_n \pi_n = 1 & \\ S\left(\xi, \theta\right) = \hat{S} & \end{aligned} \tag{16}$$

The solution of this problem provides the distribution $\mathcal{Q}$, represented as the weights π, as well as the sentiment vector ξ, the slope of beliefs $\bar{\beta}$, and the dispersion parameters σ. In the next section, I address the data and instruments used to identify these parameters via the moment condition (13) above.[5]

4 Empirical Approach: Data, Instruments, and Optimization

This section provides details on the data sources, characteristics used in estimation, the construction of instrumental variables, and the computational approach used to solve the GEL-MPEC problem described in Section 3.

4.1 Data

Stock price, return, and shares outstanding data are obtained from the Center for Research in Security Prices (CRSP) Daily Stock Database. I restrict the sample to common equities by filtering to share codes 10 and 11. Quarterly accounting data, including book equity, are drawn from the CRSP/Compustat Merged Fundamentals Quarterly Database accessed through WRDS. To avoid look-ahead bias while limiting the staleness of accounting information, accounting variables are lagged by one quarter relative to the estimation period.

Institutional holdings data are sourced from the Thomson Reuters Institutional Holdings Database (s34), also accessed through WRDS. This database aggregates mandatory 13F filings submitted to the Securities and Exchange Commission (SEC) by investment managers with at least $100 million in qualifying securities. Each filing reports the manager's equity positions at quarter end, including CUSIP identifiers, total shares held, and market values. I use these data to construct measures of aggregate institutional ownership for each stock in each quarter, following the data procedures and methods detailed in Graves (2025).

Aggregate household holdings are not directly observed. Following the residual approach common in the literature, I compute household dollar ownership of stock n at time t as

$$\hat{H}_t(n) = \max\{0,\, \mathrm{ME}_t(n) - \mathrm{InstOwn}_t(n) + \mathrm{Shorts}_t(n) - \mathrm{Insider}_t(n)\}$$

where $\mathrm{ME}_t(n)$ is market equity, $\mathrm{InstOwn}_t(n)$ is aggregate institutional ownership from 13F filings, $\mathrm{Shorts}_t(n)$ is the dollar value of short interest, and $\mathrm{Insider}_t(n)$ denotes insider holdings computed as the difference between total market equity and shares available in the public float. The resulting

[5] When implementing this paper's model empirically, I use a slight modification to the demand described above: households compute the optimal weight vector without leverage constraints (but with short-sale constraints) and then, if the weights sum to more than one, they divide weights by the total sum to leave leverage at a maximum of one. This results in substantially similar portfolios for the vast majority of investors but creates an order of magnitude acceleration in estimation time due to no longer needing to compute Lagrange multipliers for every simulated household at every function call. This highly similar generated demand system also satisfies connected substitutes and is no less plausible a portfolio optimization rule for households to heuristically follow versus the one in the main text.

vector is scaled by the quarterly fraction of household wealth allocated to equities, obtained from the Federal Reserve Economic Data (FRED) database. For each quarter, the 300 stocks with the largest household holdings are selected for estimation.

Consumer versus producer classifications for stocks are assigned using four-digit NAICS codes by utilizing Claude (Anthropic) to approximate the consumer share of different NAICS code groupings. This heuristic approach serves to separate businesses that are predominantly focused on B2B versus B2C. Stocks are classified as consumer-facing if this share exceeds 50% and as producer-facing otherwise.

4.2 Empirical Specification

The estimation targets the parameter vector $\theta_t = \left(\bar{\beta}, \sigma_{\text{book}}, \sigma_{\text{mom}}, \sigma_\epsilon\right)$ and the unobservable sentiment vector ξ_t for each quarter. The belief specification is linear in five observable asset characteristics: log market equity $x_0(n)$, log book equity $x_1(n)$, momentum $x_2(n)$, CAPM beta $x_3(n)$ computed from six months of daily returns, and a constant $x_4(n) = 1$. The mean belief of a household about asset n is therefore $\bar{\beta}' x(n) + \xi_t(n)$ where $\bar{\beta} = \left(\bar{\beta}_0, ..., \bar{\beta}_4\right)$ is the vector of mean preference coefficients and $\xi_t(n)$ is the unobservable component capturing average household sentiment about stock n. All characteristics except the constant are standardized to have unit variance and zero mean prior to estimation.[6]

Household heterogeneity in beliefs is characterized by random coefficients on log book equity and momentum, with dispersion parameters σ_{book} and σ_{mom}, and through an idiosyncratic preference shock with dispersion σ_ϵ. For tractability, I do not allow for random coefficients on the slope of beliefs with respect to log market cap (no heterogeneity in how beliefs vary as a function of market cap). Idiosyncratic return variance for each stock is estimated from the past 126 trading days of log excess returns net of the market component, with the variance then annualized. The estimation proceeds via the GEL-MPEC formulation of equation (16), which jointly solves for the parameter vector θ_t, the unobservable sentiment vector ξ_t, and the GEL probability weights π by minimizing $\sum_n \pi_n^2$ subject to the market clearing constraint $S(\xi, \theta) = \hat{S}$ and the moment conditions $\sum_n \pi_n \xi_n z_n^k = 0$ for each instrument k. The demand function S is computed by simulating 300,000 households whose consideration sets are drawn from the attention model described in Subsection 4.3 and whose preference shocks are drawn from the heterogeneity distribution. Each household optimizes over its consideration set subject to short-sale constraints and a leverage restriction as described in Section 3.

4.3 Attention Vector Construction

A key input to both the household consideration set model and the instrumental variables is a measure of the attention each stock receives from retail investors. I construct this attention measure using the cross-sectional retail volume prediction model of Graves (2023), which exploits the finding

[6]This normalization is not relevant to the econometric analysis but aids in numerical optimization.

that the cross-sectional distributions of retail purchases and sales on a given day are predicted by the same two factors: recent exchange-wide dollar volume and the salience of recent returns.

The model predicts the fraction of total retail dollar volume on day t that occurs in stock n as a multiplicative function of four lagged volume share variables, five return salience rankings, and the stock's fractional market capitalization. Specifically, the predicted fractional volume for order type $Z \in \{\text{Buys}, \text{Sells}\}$ is

$$RetFrac_{t,n}^{Z} = c \times \prod_{k=1}^{4} \left(x_{k,n}^{V}\right)^{\alpha_{k,Z}^{V}} \times \prod_{k=0}^{4} \exp\left(-\eta_{k,Z}^{R}\, j_{n,k}\right) \times (fracMktCap_{n,t})^{\alpha_{5,Z}} \tag{17}$$

where $x_1^V, ..., x_4^V$ are lagged volume shares at the daily, two-day, weekly, and monthly horizons; $j_{n,k} \in [0,1]$ is the normalized salience ranking of stock n for return window k (with 1 the most salient), computed using the salience function of Bordalo, Gennaioli and Shleifer (2012) applied to contemporaneous, one-day, two-day, weekly, and monthly return windows; $fracMktCap_{n,t}$ is the stock's share of total market capitalization; and the parameters α^V, η^R, and α_5 are estimated by OLS after taking logs. This specification can be estimated by OLS in logs, as shown in Graves (2023). The parameter vectors $\hat{\beta}^{Buys}$ and $\hat{\beta}^{Sells}$ are estimated on the subpenny-identified retail trade data of Graves (2023) over 2013–2021 using the BJZZ methodology and held fixed in this paper.

Given these estimated parameters, I compute predicted daily attention for each stock in the full CRSP common equity universe by evaluating equation (17) separately under the buy and sell coefficient vectors and averaging:

$$\widetilde{a}_t(n) = \frac{1}{2}\left(\widehat{RetFrac}_{t,n}^{Buys} + \widehat{RetFrac}_{t,n}^{Sells}\right). \tag{18}$$

The daily predictions are then normalized to sum to one across stocks on each day and averaged within each calendar quarter to obtain the quarterly attention vector:

$$a_t(n) = \frac{1}{|Q_t|} \sum_{d \in Q_t} \frac{\widetilde{a}_d(n)}{\sum_m \widetilde{a}_d(m)} \tag{19}$$

where Q_t denotes the set of trading days in quarter t. The resulting vector a_t assigns to each stock a non-negative attention weight that sums to one across the universe and captures how concentrated retail trading activity is in each equity during the quarter.

The assumption underlying the use of a_t as a basis for household consideration sets and for constructing instruments is that the cross-sectional distribution of retail attention is well-proxied by the distribution of retail volume. This is a natural assumption: a household that devotes more attention to a stock is more likely to trade it, so cross-sectional variation in retail volume fractions reflects cross-sectional variation in attention. The attention vector a_t is computed entirely from exchange-wide volumes, publicly available return data, market capitalizations, and parameters estimated from identified retail trades. Crucially, the construction does not use any information

about household portfolio holdings, the dependent variable in the demand estimation.

For the attention vector to serve as a valid component of instrumental variables, we require that the variation it induces in the instruments be orthogonal to unobservable household sentiment ξ_t. The main potential threat is that a shock to sentiment about stock n simultaneously increases both household demand and exchange-wide volume, which would contaminate the attention measure. Two features of the construction mitigate this concern. First, the attention model's volume inputs are aggregate exchange-wide volumes—not retail volumes—so that institutional trading, which constitutes the vast majority of daily dollar volume, dominates the variation. Second, the return salience variables capture the *ranking* of a stock's return relative to the cross section rather than its level, so that uniform shifts in beliefs that move all prices proportionally leave the salience rankings unchanged. The remaining identifying assumption is that the cross-sectional pattern of aggregate volume and return salience is not systematically driven by household private sentiment conditional on the observable characteristics already included in the model.

4.4 Instrument Construction

The estimation employs eight instrumental variables for the $N = 300$ stocks in each quarter, yielding a just-identified system with eight moment conditions corresponding to the eight parameters (five mean coefficients and three dispersion parameters). The instruments, outside of exogenous characteristics, fall into three categories: a price instrument, a momentum instrument based on attention, and BLP instruments based on within-sector sums of characteristics.

Price instrument. As in Graves (2025), I construct an instrument for log market equity using counterfactual institutional demand derived from the consideration sets of institutional investors. For each stock n, the instrument aggregates the hypothetical dollar demand that would result if each institutional investor allocated its assets equally or in proportion to book equity across its consideration set, using the consideration sets detailed in Graves (2025). The consideration sets are determined using the taxonomy of rigid and dynamic managers from Graves (2025). The instrument is

$$z_1(n) = \log\left(1 + \hat{P}_t^{EQ}(n) + \hat{P}_t^{R,EQ}(n) + \hat{P}_t^{BE}(n) + \hat{P}_t^{R,BE}(n)\right)$$

where $\hat{P}_t^{EQ}$ and $\hat{P}_t^{BE}$ are the counterfactual demand from dynamic institutions under equal- and book-value weighting, respectively, and $\hat{P}_t^{R,EQ}$ and $\hat{P}_t^{R,BE}$ are the corresponding quantities from rigid institutions. Variation in this instrument is driven by the exogenous structure of institutional consideration sets, which shifts equilibrium prices independently of household private beliefs.

Momentum instrument. The second instrument exploits the year-over-year change in the attention that stock n receives from household investors. Specifically,

$$z_3(n) = a_t(n) - a_{t-4}(n)$$

where $a_t(n)$ is the quarterly attention measure constructed in Subsection 4.3 using the retail volume prediction model of Graves (2023). This variable captures shifts in estimated household awareness over the past year and is plausibly exogenous to contemporaneous household sentiment because it is constructed from aggregate exchange-wide volumes and return salience rankings rather than from household portfolio holdings, as discussed in Subsection 4.3.

Exogenous characteristics and BLP instruments. Log book equity $z_2(n)$, CAPM beta $z_4(n)$, and a constant $z_5(n) = 1$ serve as their own instruments. The remaining instruments are constructed in the BLP tradition by computing within-sector sums of characteristics, using four-digit NAICS codes to define sectors. These are the same-sector sum of attention $z_6(n)$, the level of attention $z_7(n)$, and the same-sector sum of log book equity $z_8(n)$. In the spirit of Berry, Levinsohn and Pakes (1995), these instruments use the characteristics of competing stocks within a stock's sector to provide exogenous variation.

4.5 Optimization

The GEL-MPEC optimization problem (16) is solved separately for each of the 149 quarters from 1984 to 2021 using the Knitro nonlinear programming solver. Each quarterly problem involves 608 parameters — five mean coefficients, three dispersion parameters, 300 unobservable sentiment values, and 300 GEL probability weights—along with 308 nonlinear constraints comprising 300 market clearing conditions and eight moment conditions. The objective function is quadratic and convex, but the equilibrium constraints linking demand to the parameter vector and sentiment are highly nonlinear. To provide good initial conditions, I first solve the nested fixed point (NFXP) mapping of Proposition 2 at the initial parameter guess to obtain a starting value for ξ_t. When the initial Knitro solve does not converge, the NFXP is re-solved at the returned parameter values to generate an improved starting point, and Knitro is restarted.

Each quarterly optimization is executed on the Harvard Faculty of Arts and Sciences Research Computing (FASRC) cluster. Within each run Knitro parallelizes constraint evaluations across all available cores via MATLAB's Parallel Computing Toolbox. The demand function evaluation requires simulating portfolio choices for 300,000 households, making parallelization essential. Across the full sample of 149 quarters the procedure is run independently for three specifications — the attention model specification, the market-capitalization-based specification, and the alternative-characteristics specification — with results compared for robustness (see Appendix H). Additional details on the Knitro configuration are provided in Appendix E.

4.6 Alternative Demand Specifications

In addition to the main specification, I estimate two alternative demand specifications that serve as robustness checks. The first replaces the attention model with a simpler market-capitalization-based measure of household awareness; the second uses an alternative set of observable characteristics. Both specifications are estimated via the same GEL-MPEC procedure described above.

No-Attention specification. The no-attention specification ("NoAttn") retains the same five observable characteristics as the main specification—log market equity $x_0(n)$, log book equity $x_1(n)$, momentum $x_2(n)$, CAPM beta $x_3(n)$, and a constant $x_4(n) = 1$—and the same eight instrumental variables. The parameter vector is identical: five mean coefficients $\bar{\beta}$, two dispersion parameters σ_{book} and σ_{mom}, and an idiosyncratic dispersion parameter σ_ϵ. The sole difference is in the construction of household consideration sets: rather than using the retail volume prediction model of Graves (2023) described in Subsection 4.3, the probability that stock n enters a household's consideration set is proportional to its market capitalization,

$$a_t^{\text{NoAttn}}(n) = \frac{\text{ME}_t(n)}{\sum_m \text{ME}_t(m)} \times \bar{C}$$

where $\bar{C}$ is the target mean consideration set size. This specification tests whether the estimated sentiment values and resulting portfolio strategies are driven by the specific attention model or whether they are robust to a simpler, market-capitalization-based proxy for household awareness. All other aspects of the estimation—the objective function, the market clearing constraints, and the optimization procedure—remain unchanged.

Alternative characteristic specification. The alternative specification ("AltSpec") modifies both the observable characteristics and the instrument set. In place of the five characteristics of the main specification, the belief function is linear in four observable asset characteristics: the market-equity-to-book-equity ratio $\tilde{x}_0(n) = \text{ME}_t(n)/\text{BE}_t(n)$, momentum $\tilde{x}_1(n)$, CAPM beta $\tilde{x}_2(n)$, and a constant $\tilde{x}_3(n) = 1$. By replacing log market equity and log book equity with their ratio, this specification captures the value dimension of beliefs through a single variable while reducing the number of mean coefficients from five to four.

Because the characteristic space differs, the dispersion structure is also modified. The alternative specification allows for heterogeneity in beliefs over three characteristics—ME/BE, momentum, and beta—with dispersion parameters $\sigma_{\text{ME/BE}}$, σ_{mom}, and σ_β, in addition to the idiosyncratic dispersion σ_ϵ. This yields a total of eight parameters (four mean coefficients plus four dispersion parameters), matching the main specification's count and preserving the just-identified structure of the estimation.

The instrument set is adapted to reflect the new characteristic space. The price instrument is rescaled by book equity to match the ME/BE characteristic:

$$\tilde{z}_1(n) = \frac{\hat{P}_t^{EQ}(n) + \hat{P}_t^{R,EQ}(n) + \hat{P}_t^{BE}(n) + \hat{P}_t^{R,BE}(n)}{\text{BE}_t(n)}$$

where the numerator is the same counterfactual institutional demand used in the main specification's price instrument but is now divided by book equity rather than appearing inside a log transformation, so that the instrument is relevant for the ME/BE ratio. The momentum instrument $\tilde{z}_2$ remains the year-over-year change in attention. CAPM beta and the constant serve as their own instruments. The BLP-style instruments are the level of attention, the same-sector

sum of the attention change, the same-sector sum of beta, and the count of stocks in other sectors. After normalization, eight instruments identify the eight parameters. The attention model is the same as in the main specification.

5 Household Demand Estimates

In this section, I discuss the results of the demand estimations, including the estimated household demand coefficients and their trends over time. These coefficients capture households' return expectations regarding factor exposures. Specifically, households display a consistently positive belief in a value premium, a fluctuating and on average neutral belief about returns to momentum, and a weakly positive average belief about returns to beta. I also discuss the persistence of stock-specific sentiment in the data.

5.1 Summary Statistics

Table 1 provides summary statistics of total household assets, assets held in the top 300 holdings, assets held outside the top 300 holdings, share of household assets invested in direct ownership of common stocks, largest household holdings in dollars, and smallest household holdings in dollars. As is readily apparent from the table, while household ownership in common stocks has increased significantly over time in dollar terms, household ownership in common stocks has declined relative to institutional ownership. Indeed, the time period covered by this paper saw a significant increase in relative institutional ownership as retail traders migrated to fund, index, and ETF vehicles from direct ownership. This trend is documented in greater detail in Graves (2025).

The most striking change over the data sample, however, is not the steady increase in household assets directly invested in common stocks, but rather the increased concentration of its ownership in a smaller group of stocks. In the earlier time windows of the data sample, from 1984-1988, 1989-1992, and 1993-1996, we see that smaller holdings represent an average of 28%, 23%, and 26%, respectively, of total investment in common stocks. By the last three time windows of 2013-2016, 2017-2020, and 2021, smaller holdings only represent an average of 13%, 10%, and 11%, respectively. This paper's focus on conducting demand estimation for the top 300 common stock holdings therefore includes the vast majority of household holdings in dollar terms throughout the sample, and the overwhelming majority of household ownership in the later years of the sample.

5.2 Beliefs over Time

Figure 1 plots the estimated coefficients for log market equity $x_0(n)$, log book equity $x_1(n)$, momentum $x_2(n)$, CAPM beta $x_3(n)$, and a constant term x_4 in each quarter from 1984Q4 through 2021Q4. These characteristics are specifically chosen to be close to the specification of Koijen and Yogo (2019), which uses stock characteristics inspired by Fama and French (2015), only here I have substituted the model of Carhart (1997). This change is due to the literature's past emphasis

Table 1: Household Sector Summary Statistics

	Total Assets	Top 300 Assets	Non-Top 300 Assets	Equity Fraction	Largest Holding	Smallest Holding
1984–1988	$721B	$520B	$201B	10.0%	$40.3B	$0.46B
1989–1992	$919B	$712B	$207B	12.3%	$30.9B	$0.55B
1993–1996	$1545B	$1148B	$397B	15.8%	$41.4B	$0.93B
1997–2000	$3609B	$2954B	$655B	21.4%	$133B	$1.7B
2001–2004	$3237B	$2733B	$504B	16.7%	$142B	$1.5B
2005–2008	$3424B	$2883B	$541B	17.4%	$217B	$1.6B
2009–2012	$3368B	$2852B	$516B	15.3%	$196B	$1.5B
2013–2016	$4713B	$4086B	$627B	19.2%	$244B	$2.0B
2017–2020	$5861B	$5257B	$604B	21.9%	$439B	$2.1B
2021	$9900B	$8847B	$1053B	26.8%	$923B	$3.7B

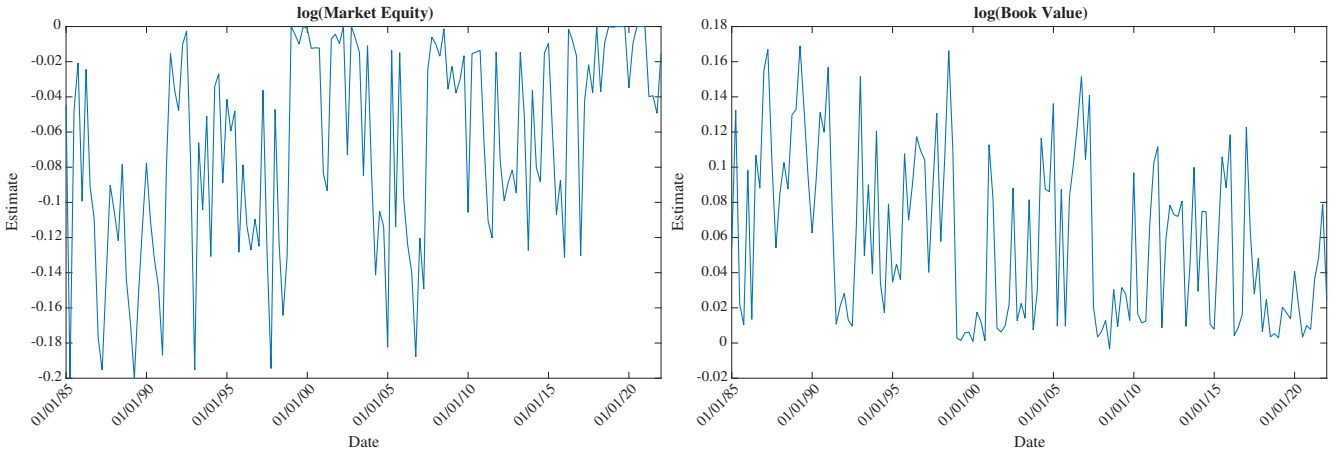
Note: This table presents summary statistics for the household sector's direct stock holdings from 1984 to 2021. Total Assets reports total household common stock holdings in billions of dollars, computed as the residual of market capitalization minus institutional holdings (from 13F filings) minus insider holdings plus short interest. Top 300 Assets and Non-Top 300 Assets decompose total holdings into the 300 largest household stock positions and the remainder. Equity Fraction is the share of household wealth allocated to directly held common stocks, obtained from the Federal Reserve's Financial Accounts of the United States (FRED). Largest Holding and Smallest Holding report the dollar value of the single largest and 300th largest household stock position, respectively. Household holdings are defined for common stocks (CRSP share codes 10 and 11) with positive market capitalization. Each statistic is averaged over all quarters within the designated four-year window; the 2021 row reports a single-year average. Data sources include CRSP, Thomson Reuters 13F institutional holdings, FTSE Russell Index data (for insider ownership imputation), and the Federal Reserve Financial Accounts.

on momentum trading and extrapolation among household investors, resulting in the necessity of including momentum and creating an instrumental variable to control for its potential endogeneity. I use the baseline estimation procedure and specification detailed in Section 4, which includes the attention model in subsection 4.3. Results using the market cap-based attention model and an alternative specification involving the market-to-book ratio and an additional dispersion parameter are contained in Appendix H and Appendix I, respectively.

Several patterns are visible in these time series.[7] First, household beliefs about returns to risk factors do not appear to be stable over time. Instead, average beliefs appear to oscillate significantly and include protracted higher and lower periods as well. Second, average beliefs about returns to momentum appear to fluctuate between positive and negative without a significant directional tendency, in contrast to some of the previous literature's emphasis on households as pro-momentum extrapolators (or, in other cases, counter-momentum liquidity providers). Beliefs about beta are similarly close to neutral on average, with an extended positive period in late 2020 through 2021 after the start of the pandemic (the so-called "meme stock" craze era). Parameter estimates in any one time period should be interpreted cautiously, as they are dependent on model specification.

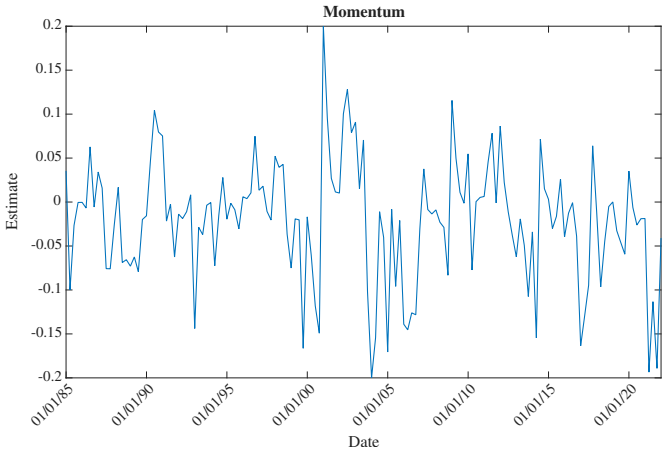
Table 2 reports summary statistics for the estimated demand parameters across all three

[7]For ease of interpretation, recall that all characteristics except for the constant term are normalized to have mean 0 and unit variance.

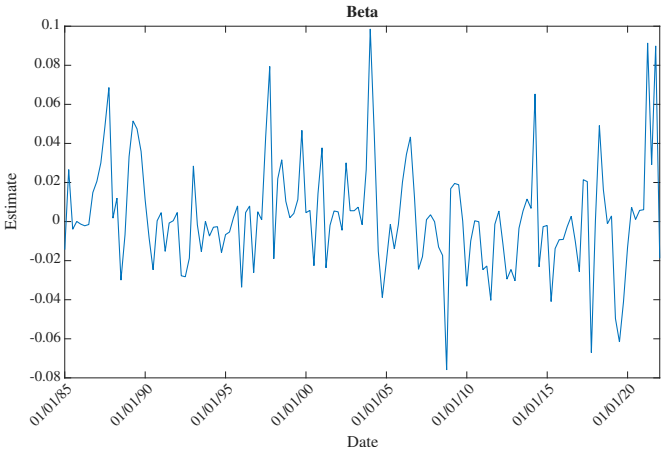


(a) Log Market Equity

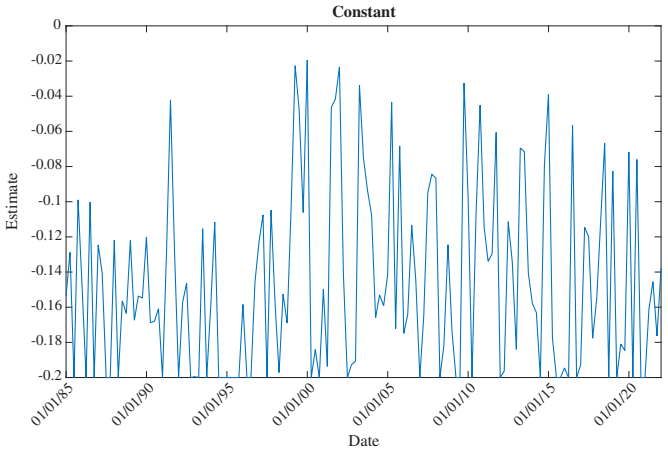
(b) Log Book Value



(c) Momentum



(d) Beta



(e) Constant

Figure 1: Parameter Estimates with Attention: Cross-Sectional Distributions

Note: Each subfigure plots the quarterly time series of an estimated household demand parameter from the baseline household demand model (including attention effects), from 1984 to 2021. The five parameters shown are the coefficients on log market equity, log book value, momentum, beta, and a constant in the household portfolio choice model. Parameters are estimated quarterly via GEL-MPEC using Knitro optimization on the top 300 U.S. stocks by market capitalization, with an average consideration set size of 10 stocks and estimated attention vectors.

Table 2: Demand Estimation Summary Statistics

	Main	NoAttn	AltSpec
Mean coefficients			
$\bar{\beta}_{\log \text{ME}}$	-0.0674 (0.0573)	-0.0580 (0.0570)	—
$\bar{\beta}_{\log \text{BE}}$	0.0593 (0.0476)	0.0487 (0.0484)	—
$\bar{\beta}_{\text{ME/BE}}$	—	—	-0.0944 (0.0735)
$\bar{\beta}_{\text{mom}}$	-0.0207 (0.0686)	-0.0160 (0.0601)	0.0190 (0.0737)
$\bar{\beta}_{\text{beta}}$	0.0026 (0.0278)	0.0045 (0.0269)	-0.0076 (0.0351)
$\bar{\beta}_{\text{const}}$	-0.1456 (0.0520)	-0.1613 (0.0524)	-0.1517 (0.0557)
Dispersion parameters			
σ_{book} / $\sigma_{\text{ME/BE}}$	0.0388 (0.0265)	0.0366 (0.0317)	0.0566 (0.0457)
σ_{mom}	0.0668 (0.0434)	0.0883 (0.0448)	0.0460 (0.0477)
σ_{beta}	—	—	0.0438 (0.0456)
σ_{ϵ}	0.0418 (0.0318)	0.0419 (0.0291)	0.0373 (0.0277)

Note: This table reports time-series means and standard deviations (in parentheses) of the estimated demand parameters for each of the three model specifications, computed over the 149 estimated quarters. The Main specification uses the attention-based demand model with five stock characteristics: log market equity, log book equity, momentum, CAPM beta, and a constant. The NoAttn specification uses market capitalization-based attention weights in place of estimated attention vectors but retains the same five characteristics. The AltSpec specification uses an alternative set of four characteristics—the market equity-to-book equity ratio, momentum, CAPM beta, and a constant—while including three dispersion parameters rather than two. All parameters are estimated quarterly via GEL-MPEC on the top 300 U.S. stocks by market capitalization.

specifications. Unlike point-in-time estimates, the averages appear to be qualitatively consistent across the different demand specifications. Furthermore, as shown in Appendix G, even the point estimates align across the paper's two main specifications: the demand estimates using this paper's attention model or an alternative model with attention proportional to market cap are highly correlated in the time series. I now turn to this paper's main object of interest, stock-specific household sentiment.

5.3 Persistence of Sentiment

Before studying the return predictability associated with household sentiment in Section 6, I analyze the autocorrelation patterns in household sentiment. If sentiment were an arbitrary artifact of demand estimation, we would expect minimal persistence. On the other hand, if sentiment captures some combination of genuine stock-specific information held by retail investors or persistent dogmatic beliefs, we would expect it to be significantly autocorrelated. Naturally, the demand estimation methodology, reliance on numerical non-convex optimization, and shifting set of top 300 stocks all introduce noise into the measurement of persistence. The further apart in time we move, the more likely that the 300 stocks used in the estimations have shifted considerably given the threshold approach. Furthermore, we are unable to study persistence in a broader cross-section

of stocks due to the computational limitations.

To measure persistence, I regress $\xi_t(n)$ on $\xi_{t-k}(n)$ for lags $k = 1, \ldots, 12$ using pooled OLS. Because sentiment values for different stocks in the same quarter may be cross-sectionally correlated and sentiment for the same stock across quarters may be serially correlated, I compute two-way clustered standard errors by stock and by quarter following the procedure of Petersen (2009). Specifically, the clustered variance estimator is $V_{\text{twoway}} = V_{\text{stock}} + V_{\text{quarter}} - V_{\text{White}}$, where V_{stock} and V_{quarter} are the one-way cluster-robust variance estimators and V_{White} is the heteroskedasticity-robust (HC1) variance, each with appropriate small-sample corrections. Observations with missing sentiment values in either the dependent or independent variable are excluded, as some stocks fall into and out of the top 300 holdings over time.

Figure 2 reports the results. Household sentiment displays statistically significant persistence. At a lag of one quarter, the autocorrelation coefficient is approximately 0.088, statistically significant with a 95% confidence interval of roughly [0.067, 0.109]. The persistence decays as the lag increases but remains significant: the coefficient is approximately 0.04 at two quarters, 0.03 at four and seven quarters, and approximately 0.02 at eight quarters. The 95% confidence intervals exclude zero for most lags, though at lags of nine and eleven quarters the intervals narrowly include zero. At twelve quarters (three years), the coefficient is approximately 0.02 and is again statistically significant. The persistence of household sentiment and the gradual decay suggest that household sentiment captures stock-specific views that persist over time, even if they are measured with significant noise due to the limitations inherent in the reliance on aggregate data. The results for the two alternative specifications, contained in Appendix F, are both quantitatively and qualitatively similar. All three estimations reveal positive point estimates for every single time lag out to three years (12 quarters). These results are consistent with persistent information-based beliefs of households but inconsistent with mean-reverting noise-trader-like beliefs, particularly when combined with the results that follow in Section 6 on return predictability.

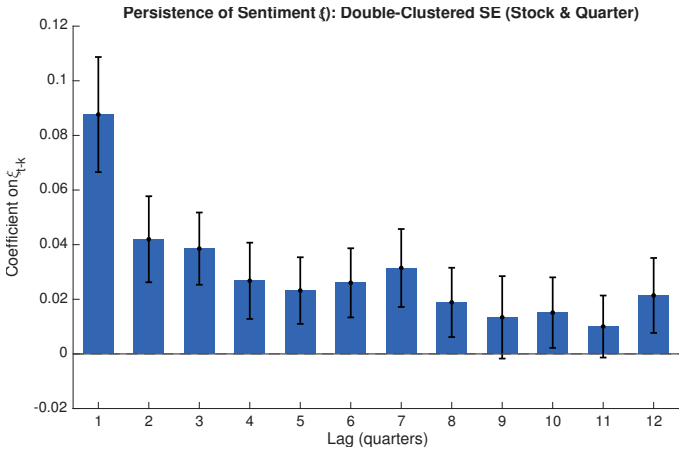


Figure 2: Persistence of Household Sentiment: Autocorrelation Profile

Note: This figure plots the estimated autocorrelation coefficients of household sentiment $\xi_t(n)$ at lags of 1 through 12 quarters. Each bar represents the slope coefficient from a pooled OLS regression of $\xi_t(n)$ on $\xi_{t-k}(n)$ across all stocks and quarters with nonzero sentiment values. Error bars depict 95% confidence intervals constructed from two-way clustered standard errors (by stock and by quarter) following Petersen (2009). Sentiment is estimated quarterly via GEL-MPEC using Knitro optimization on the top 300 U.S. stocks by market capitalization from 1984Q4 through 2021Q4.

6 Household Beliefs and Return Predictability

In this section, I study one of the core objects recovered from this paper's demand estimation: the "sentiment" of households about different stocks. Sentiment corresponds to the average household's expected return belief that is idiosyncratic to a specific stock. In other words, sentiment is the unobservable belief component that is unrelated to beliefs about factor returns, which is to say the observable characteristics of a stock such as momentum or beta. Here, I start by describing how I generate counterfactual portfolios that establish what demand would have looked like if a specific component were removed. For example, I study household portfolios under the counterfactuals where all portfolios are driven by beliefs about observable characteristics and where all portfolios are driven exclusively by sentiment. I then proceed to conduct return predictability exercises, showing that household sentiment is a positive contributor to portfolio returns and indeed has predictive power. This predictive power is concentrated in stocks whose firms are largely consumer-oriented.

6.1 Counterfactual Portfolio Generation

To evaluate which components of estimated household beliefs contribute positively or negatively to household returns, I use the quarter-by-quarter parameter estimates from Section 3 to construct counterfactual portfolios that isolate specific channels of the household demand model. Each quarter t, I simulate 300,000 households whose consideration sets $C_{h,t}$ are drawn from the estimated attention model described in Subsection 4.3 and whose preference shocks $(\nu_{k,h,t}, \epsilon_{h,t})$ are drawn from the estimated heterogeneity distribution. For each household h and each stock n in its consideration set, I compute the optimal risky-asset allocation from the mean-variance demand function subject to the short-sale and leverage constraints described in detail in Section 2; I then

aggregate across the simulated households to obtain a vector of counterfactual portfolio weights. By systematically zeroing out different elements of the belief function — the household sentiment vector ξ_t, the observable characteristic-based component of average beliefs $\bar{\beta}'x(n)$, and the preference heterogeneity parameters (σ_{book}, σ_{mom}, σ_ϵ) — I isolate each channel's contribution to aggregate household demand.

The primary object of interest is what I call the *sentiment portfolio*, which removes the observable characteristic-based component $\bar{\beta}'x(n)$ so that each household's demand for stock n depends only on the common sentiment $\xi_t(n)$ and the stock's idiosyncratic return variance $\sigma_t^2(n)$; I still include random coefficients on characteristics but the average directional beliefs are purged from these portfolios. For comparison, I construct three additional portfolios. The *characteristics portfolio* removes sentiment $\xi_t(n)$ entirely, leaving behind only the observable characteristic mean component $\bar{\beta}'x(n)$ and random coefficients, isolating demand driven solely by household responses to stock fundamentals. Finally, the *actual holdings portfolio* weights each stock by its observed share of aggregate household direct equity holdings in the data, providing a benchmark against which the model-based counterfactuals can be evaluated.

I repeat the full counterfactual portfolio exercise under two alternative model specifications described in Subsection 4.6. The no-attention specification replaces the estimated heterogeneous consideration sets with inclusion probabilities proportional to market capitalization, testing whether the results are driven by the specific attention model or are robust to proxying for household awareness via market cap instead of a more complex model. The alternative characteristic specification replaces the five-characteristic belief model (log market equity, log book equity, momentum, beta, and a constant) with a more parsimonious four-characteristic model that uses the market-equity-to-book-equity ratio in place of separate log market equity and log book equity terms and adds disagreement about returns to beta, testing sensitivity to the included characteristics in the linear belief specification as well as the number of dispersion parameters. In this section I first evaluate the performance of sentiment portfolios, characteristics portfolios, and actual portfolios before progressing to an analysis of the potential origin of sentiment-based return predictability.

6.2 Sentiment-Driven Portfolios versus Actual Portfolios and Characteristic-Based Portfolios

To test the performance of these counterfactual portfolios, I record the portfolio composition as of the start of each quarter and collect buy-and-hold daily returns using CRSP delisting-adjusted returns. I compute the returns for the "contemporaneous" or "lag 0" portfolio based on contemporaneous household aggregate holdings as of the last day of the calendar quarter, and I also compute returns based on portfolios formed from 1 through 7 quarters ago. Throughout the paper, a "contemporaneous" or "lag 0" portfolio refers to the former, whereas "lag 0 through 3" or "lag 0 through 7" refer to strategies computed by using the average of four or eight lagged portfolio returns, respectively. By way of illustration, the "sentiment lag 0 through 7" portfolio would take the average returns of the eight portfolios formed contemporaneously through 1.75 years ago. I

analyze the returns through the lens of portfolio Sharpe, the CAPM, and a six-factor model based on adding the momentum factor of Carhart (1997) to the five factors of Fama and French (2015). The six-factor model allows us to determine any "abnormal returns" relative to known "risk factors" in the literature that have been previously shown to have statistically significant risk premia.

The results from the simplest portfolios, the contemporaneous lag 0 portfolios, are contained in Table 3. Perhaps unsurprisingly, the six-factor annualized abnormal returns ("alpha") of the "actual" portfolio is functionally 0 (0.03%). While this may appear surprising in light of past findings of poor retail trader performance, two major points are worth noting. First, this paper's analysis only focuses on the top 300 stocks of households. Second, ownership (holdings) is not the same object as "trading," which captures both ownership and short-term trading. The quarterly structure of the data means that we do not pick up the short-term mistakes of retail traders induced by significant overtrading (see e.g., Heimer and Simsek (2019) and Barber and Odean (2000)), though such activity is a significant negative drag on investor returns and is orthogonal to the beliefs (sentiment) of interest in this paper.

We also see in Table 3 that the sentiment-based portfolio has positive abnormal returns (1.21% annualized with a t-statistic of 2.22) whereas the characteristic-based portfolio has alpha of -0.06% with a t-statistic of -0.11. These results are consistent with sentiment contributing positively to portfolio performance whereas characteristic-based beliefs are a weak negative. Alternatively stated, household sentiment is a positive contributor to household performance whereas the evidence is inconsistent with households benefiting from their characteristic-driven beliefs.

Table 4 shows the effects of eliminating household disagreement (no random coefficients) and additionally the effects of an "even attention" model where every stock regardless of size receives equal average attention from households in portfolio formation. The six-factor results are not statistically significant but nonetheless provide strong evidence against the hypothesis that differential household attention to stocks is a performance albatross: on the contrary, the results suggest that heterogeneous household attention, when combined with other components of household demand, serves as a modest positive for performance. This contrasts with the literature on retail trading, which has shown that attention fluctuations trigger overtrading and are likely associated with welfare losses (see e.g. Barber and Odean (2008) and Graves (2023)).

We now progress to studying portfolio returns for longer lags. Table 5 details the divergence between sentiment-based portfolios, actual portfolios, and characteristic-based portfolios. I include analysis of both "No Disagreement" portfolios (labeled ND) where the random coefficients are eliminated to leave behind a representative agent as well as the benchmark portfolios that still contain disagreement among the households. The purest test of the effects of sentiment is contained in the first column, which tests the sentiment (ND) portfolio against the characteristics (ND) portfolio via a strategy that goes long the no-disagreement sentiment portfolio and short the no-disagreement characteristics portfolio. The annualized alpha of such a strategy is 1.98% with a t-statistic of 2.4. Meanwhile, the long-short strategy that purchases the no-disagreement sentiment portfolio and sells short the actual portfolio of holdings achieves an annualized alpha of 1.37% (t-stat

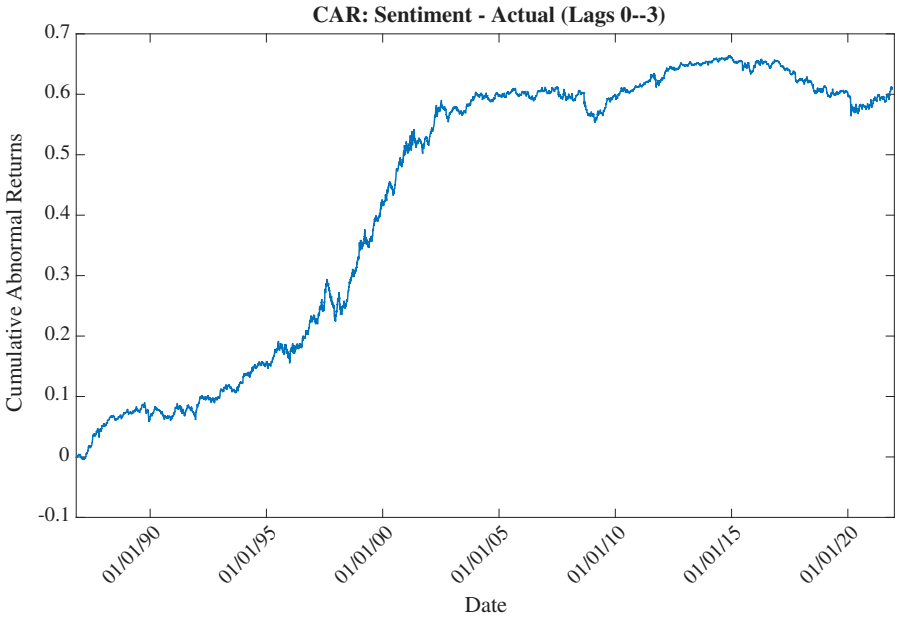


Figure 3: Cumulative Abnormal Returns: Sentiment minus Actual (Lags 0–3)

Note: This figure plots cumulative abnormal returns (CARs) for a long-short strategy that goes long the Sentiment portfolio and short the Actual Holdings portfolio, corresponding to the second and first columns of Table 5. The Sentiment portfolio weights stocks by estimated household sentiment parameters, while the Actual Holdings portfolio weights stocks by observed household equity positions. CARs are computed from daily residuals of a six-factor regression on the Fama-French five factors (MktRF, SMB, HML, RMW, CMA) plus momentum (UMD), with returns averaged over lags 0 through 3 (four quarters). Portfolios are formed from the top 300 U.S. stocks by market capitalization and rebalanced quarterly. The six-factor regression uses Newey-West (HAC) standard errors with a Bartlett kernel and automatic bandwidth selection following Newey and West (1994).

of 2.93). Results are similar when considering portfolios with disagreement, and the gap between sentiment portfolios and characteristics portfolios remains statistically significantly positive. These findings are consistent with household sentiment continuing to have positive predictive power at longer time horizons and are inconsistent with mean-reverting sentiment that would produce the exact opposite patterns. Figure 3 plots the cumulative abnormal returns for the sentiment portfolios against the actual portfolios at lags 0 through 3.

Before exploring the origin of these return prediction results, I conclude this subsection by examining whether we see evidence of sentiment reverting at longer lags. Table 6 provides household counterfactual portfolio return analysis when lagging by up to seven quarters. The sentiment portfolio now has a six-factor alpha of 2.55% (t-stat of 4.74) and remains significantly higher than the alpha of 0.81% for the actual portfolios. The weakly positive and statistically insignificant alpha on the "actual" portfolio is unsurprising given that the lag seven portfolio includes only stocks that were in existence seven quarters ago and therefore are not recent IPOs, which tend to underperform (see e.g. Ritter (1991)). The most striking finding is the strongly significant and ongoing outperformance of the sentiment portfolios. This long-duration positive predictability of sentiment is consistent with households collectively possessing private information about individual stocks but is inconsistent with canonical "noise trader" models of households where sentiment

is either random and i.i.d. or tends to reverse, creating mean reversion in prices. In the next subsection, I test the sources of this edge in performance.

6.3 Household Sentiment and the Consumer vs. Producer Divide

A central question is where this informational content of household sentiment comes from. One potential theory mentioned in the last subsection is that households do not actually possess private information but instead are informed about their own future sentiment. This pattern would be present consistently across different types of stocks when separated by industry groupings, as it would only be plausible if retail traders displayed strong positive autocorrelations in their order imbalances at very long time horizons (a phenomenon that has not been previously documented). On the other hand, households might collectively possess private information about firms. But where would households obtain such private information in aggregate? The answer lies in usage of products and services: if a company produces a product or service that is universally beloved, then households would be more likely to think favorably of the stock ("high sentiment"). By contrast, a company producing a universally loathed product or service would be more likely to find households staying out of its stock entirely. Meanwhile, companies like an oil refiner never face consumers directly and therefore households would be unlikely to possess as strong private information. In this subsection I directly test this hypothesis.

To test whether household sentiment differentially predicts returns based on a company's business orientation (consumer-facing or "B2C" versus business-facing or "B2B"), I classify each 4-digit NAICS industry by the degree to which its products are experienced by typical U.S. households, using a consumer-experience score that ranges from 0 (niche business-to-business products with which few households have direct contact) to 100 (ubiquitous products or services that virtually every household encounters). The classification was performed with the assistance of Claude (Anthropic), a large language model, by prompting it to assign scores to each 4-digit NAICS code based on whether the typical end consumers are households or businesses. Full details of the classification procedure and the prompt used are provided in Appendix J. Industries such as automobile manufacturing (NAICS 3361) and full-service restaurants (NAICS 7221) receive scores near 100, while heavy industrial categories such as iron and steel mills (NAICS 3311) and pipeline transportation (NAICS 4861) receive scores near 0. I then define "consumer stocks" as those whose NAICS-based consumer share is at or above the 50% threshold, and "producer stocks" as those below this cutoff.

Each quarter, I follow the procedures of 6.1 to construct counterfactual portfolios. I create two different portfolios using a 50% threshold based on these scores. First, I create sentiment-based portfolios by looking only within the Consumer grouping and forming a portfolio based only on sentiment. I repeat this within the Producer grouping. I then do the same, but instead of using counterfactual portfolio weights, I use a simple market cap-based weighting scheme within the Consumer and Producer groups. The abbreviations throughout the tables are "Sent" for sentiment-based portfolios and "All" for market-cap-based portfolios.

Table 7 provides results using lags 0 through 7. The first and fourth column present the results for the consumer and producer stock-based sentiment strategies, respectively, while the second and fifth columns present results for the consumer and producer stock-based market-cap-weighted strategies, respectively. The third and sixth columns provide results for long/short strategies that purchase the sentiment-based portfolios and sell short the market-cap-weighted portfolios to separate out the effects of sentiment from the effects of separating stocks into the two categories. The third column gives the most striking results: the consumer-based sentiment portfolio outperforms the market-weighted consumer stock portfolio by 1.94% (t-stat of 3.92) in annualized six-factor alpha, almost 2% per year and highly statistically significant. Meanwhile, the sentiment portfolio for producers only outpaces the market-weighted portfolio by 57 basis points, with the difference not statistically significant. Household sentiment regarding producer stocks does not appear to be a drag on retail investor performance, yet household sentiment regarding consumer stocks appears to be a strict positive. These results are consistent with households possessing information about stocks' quality ("unobservables") from experience. Legendary investor Peter Lynch famously urged households to "invest in what you know" on the basis that consumer encounters with goods and services will produce knowledge of how successful the company will be in the long term.[8] By contrast, the average household has little personal experience with mining companies or steel producers. Figure 4 plots this divergence visually, showing how the long/short strategy based on consumer stock sentiment has outpaced the long/short strategy based on producer stock sentiment in cumulative abnormal returns.

[8]Peter Lynch managed Magellan Fund at Fidelity for thirteen years, yielding an annualized performance of over 29% for clients and significantly outpacing the S&P 500 by nearly a factor of two.

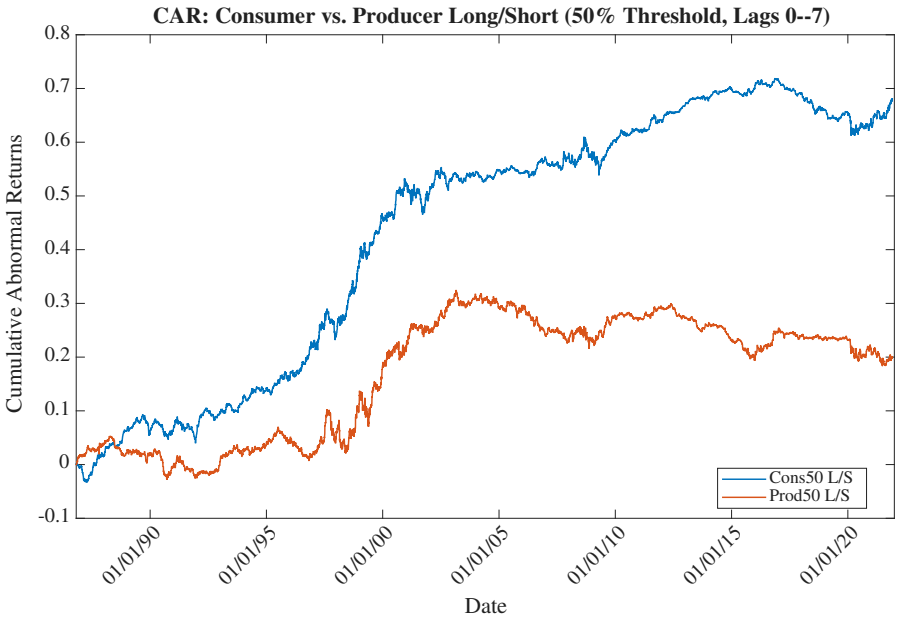


Figure 4: Cumulative Abnormal Returns: Consumer and Producer Long/Short Spreads (50% Threshold, Lags 0–7)

Note: This figure plots cumulative abnormal returns (CARs) for the consumer and producer long/short spread strategies from Table 7. The consumer (producer) long/short strategy goes long the sentiment-weighted consumer (producer) portfolio and short the corresponding market-capitalization-weighted benchmark portfolio, using the 50% NAICS-based consumer spending share threshold. Consumer stocks are those with consumer share of at or above 50%, and producer stocks are those below 50%. CARs are computed from daily residuals of a six-factor regression on the Fama-French five factors (MktRF, SMB, HML, RMW, CMA) plus momentum (UMD), with returns averaged over lags 0 through 7 (eight quarters). Portfolios are formed from the top 300 U.S. stocks by market capitalization and rebalanced quarterly. The six-factor regression uses Newey-West (HAC) standard errors with a Bartlett kernel and automatic bandwidth selection following Newey and West (1994).

Table 3: Core Models: Contemporaneous Lag

	Actual	Sentiment	Char.
A. Raw excess returns			
Average	9.25	9.50	8.70
	(3.52)	(3.27)	(3.22)
T	8867	8867	8867
SR	0.5116	0.4895	0.4882
B. CAPM regressions			
Intercept	-0.04	-0.33	-0.38
	(-0.09)	(-0.50)	(-0.63)
MktRF	0.98	1.04	0.96
	(251.77)	(135.00)	(173.96)
R^2	0.9776	0.9505	0.9610
IR	-0.0158	-0.0774	-0.1085
C. Six-factor regressions			
Intercept	0.03	1.21	-0.06
	(0.10)	(2.22)	(-0.11)
MktRF	0.98	1.00	0.96
	(449.77)	(260.19)	(195.86)
SMB	-0.20	-0.18	-0.14
	(-25.55)	(-22.02)	(-20.92)
HML	-0.01	-0.07	0.00
	(-2.26)	(-7.56)	(0.41)
RMW	0.01	-0.09	-0.01
	(1.61)	(-7.67)	(-0.71)
CMA	0.00	-0.23	-0.02
	(0.32)	(-11.05)	(-0.95)
UMD	0.00	0.00	-0.02
	(0.67)	(0.48)	(-2.71)
R^2	0.9892	0.9689	0.9669
IR	0.0174	0.3540	-0.0192

Note: This table presents portfolio performance for three core model strategies: Actual Holdings (market-cap weighted household direct stock holdings), Sentiment (belief-based model with all parameters), and Characteristics (characteristics-only model with sentiment parameters set to zero). Portfolios are formed from the top 300 U.S. stocks by market capitalization and rebalanced quarterly. Returns are evaluated at the contemporaneous lag (lag 0). Panel A reports annualized mean excess returns, t-statistics, the number of daily observations, and annualized Sharpe ratios. Panel B reports CAPM alpha, market beta, R-squared, and the information ratio. Panel C reports six-factor alpha and factor loadings from a regression on the Fama-French five factors (MktRF, SMB, HML, RMW, CMA) plus momentum (UMD), along with R-squared and the information ratio. All t-statistics are computed using Newey-West (HAC) standard errors with a Bartlett kernel and automatic bandwidth selection following Newey and West (1994).

Table 4: No-Disagreement Models: Contemporaneous Lag

	Sent. (ND)	Sent. (EA, ND)	Char. (ND)	Char. (EA, ND)
A. Raw excess returns				
Average	9.42	8.76	7.69	8.48
	(3.33)	(3.72)	(2.86)	(3.57)
T	8867	8867	8867	8867
SR	0.4986	0.5753	0.4390	0.5575
B. CAPM regressions				
Intercept	-0.08	1.21	-0.75	1.62
	(-0.11)	(1.41)	(-0.65)	(1.14)
MktRF	1.01	0.80	0.89	0.73
	(141.43)	(68.78)	(78.65)	(40.63)
R^2	0.9357	0.9088	0.8589	0.7536
IR	-0.0167	0.2641	-0.1137	0.2140
C. Six-factor regressions				
Intercept	0.74	-1.00	-1.48	-0.99
	(1.12)	(-1.45)	(-1.42)	(-0.86)
MktRF	0.98	0.85	0.91	0.79
	(225.77)	(98.89)	(92.60)	(56.04)
SMB	-0.19	-0.07	-0.08	-0.02
	(-21.72)	(-7.40)	(-6.19)	(-1.49)
HML	-0.08	0.07	0.11	0.20
	(-8.47)	(5.93)	(4.06)	(7.58)
RMW	-0.03	0.19	0.14	0.25
	(-1.97)	(14.58)	(6.49)	(12.10)
CMA	-0.13	0.23	0.13	0.30
	(-6.12)	(13.44)	(4.01)	(8.92)
UMD	0.01	0.03	-0.10	-0.04
	(1.70)	(4.18)	(-5.29)	(-1.93)
R^2	0.9534	0.9376	0.8868	0.8321
IR	0.1809	-0.2624	-0.2504	-0.1595

Note: This table presents portfolio performance for four no-disagreement model variants: Sentiment (No Disagreement), Sentiment (Even Attention, No Disagreement), Characteristics (No Disagreement), and Characteristics (Even Attention, No Disagreement). The no-disagreement variants set belief dispersion parameters to zero, isolating the role of average household sentiment. Even-attention variants use uniform consideration set probabilities rather than estimated attention vectors. Portfolios are formed from the top 300 U.S. stocks by market capitalization and rebalanced quarterly. Returns are evaluated at the contemporaneous lag (lag 0). Panel A reports annualized mean excess returns, t-statistics, the number of daily observations, and annualized Sharpe ratios. Panel B reports CAPM alpha, market beta, R-squared, and the information ratio. Panel C reports six-factor alpha and factor loadings from a regression on the Fama-French five factors (MktRF, SMB, HML, RMW, CMA) plus momentum (UMD), along with R-squared and the information ratio. All t-statistics are computed using Newey-West (HAC) standard errors with a Bartlett kernel and automatic bandwidth selection following Newey and West (1994).

Table 5: Long-Short Spreads: Lags 0 to 3

	Panel A: Long Sent. (ND)			*Panel B: Long Sentiment*		
	− Char. (ND)	− Actual	− Char.	− Char. (ND)	− Actual	− Char.
A. Raw excess returns						
Average	1.38	1.00	0.91	1.25	0.88	0.78
	(1.27)	(1.99)	(1.60)	(1.08)	(1.58)	(1.33)
T	8867	8867	8867	8867	8867	8867
SR	0.2022	0.3213	0.2501	0.1746	0.2626	0.2104
B. CAPM regressions						
Intercept	0.48	0.80	0.57	0.04	0.36	0.14
	(0.45)	(1.62)	(1.00)	(0.04)	(0.71)	(0.24)
MktRF	0.10	0.02	0.04	0.13	0.05	0.07
	(7.76)	(4.34)	(5.75)	(9.96)	(10.78)	(10.44)
R^2	0.0647	0.0159	0.0316	0.1056	0.0880	0.1114
IR	0.0724	0.2578	0.1602	0.0061	0.1135	0.0392
C. Six-factor regressions						
Intercept	1.98	1.37	1.00	2.34	1.74	1.37
	(2.40)	(2.93)	(2.02)	(2.79)	(3.95)	(2.90)
MktRF	0.06	0.01	0.02	0.07	0.02	0.04
	(8.49)	(1.62)	(6.00)	(11.14)	(6.84)	(11.13)
SMB	-0.07	0.01	-0.04	-0.06	0.02	-0.04
	(-5.01)	(2.14)	(-4.40)	(-4.63)	(4.23)	(-3.91)
HML	-0.26	-0.07	-0.10	-0.26	-0.06	-0.10
	(-10.69)	(-6.81)	(-9.03)	(-11.34)	(-8.71)	(-10.42)
RMW	-0.06	-0.01	0.01	-0.14	-0.08	-0.06
	(-3.25)	(-0.80)	(1.14)	(-7.25)	(-10.33)	(-5.66)
CMA	-0.24	-0.12	-0.11	-0.32	-0.20	-0.18
	(-8.25)	(-9.35)	(-6.93)	(-10.67)	(-15.78)	(-11.59)
UMD	0.04	0.00	0.02	0.02	-0.01	-0.00
	(2.50)	(0.81)	(1.99)	(1.28)	(-2.42)	(-0.27)
R^2	0.4323	0.1901	0.2538	0.4895	0.3851	0.4027
IR	0.3855	0.4899	0.3195	0.4575	0.6650	0.4752

Note: This table presents portfolio performance for long-short spread strategies with returns averaged over lags 0 through 3 (four quarters). The strategies take a long position in Sentiment (No Disagreement) or the full Sentiment model and a short position in one of three alternatives: Characteristics (No Disagreement), Actual Holdings, or Characteristics. Portfolios are formed from the top 300 U.S. stocks by market capitalization and rebalanced quarterly. Panel A reports annualized mean excess returns, t-statistics, the number of daily observations, and annualized Sharpe ratios. Panel B reports CAPM alpha, market beta, R-squared, and the information ratio. Panel C reports six-factor alpha and factor loadings from a regression on the Fama-French five factors (MktRF, SMB, HML, RMW, CMA) plus momentum (UMD), along with R-squared and the information ratio. All t-statistics are computed using Newey-West (HAC) standard errors with a Bartlett kernel and automatic bandwidth selection following Newey and West (1994).

Table 6: Core Models: Lags 0 to 7

	Actual	Sentiment	Char.
A. Raw excess returns			
Average	9.50	10.47	9.73
	(3.52)	(3.59)	(3.50)
T	8867	8867	8867
SR	0.5165	0.5402	0.5308
B. CAPM regressions			
Intercept	0.13	0.61	0.38
	(0.19)	(0.90)	(0.59)
MktRF	0.99	1.04	0.99
	(151.28)	(156.95)	(151.04)
R^2	0.9615	0.9575	0.9627
IR	0.0349	0.1531	0.1079
C. Six-factor regressions			
Intercept	0.81	2.55	1.16
	(1.87)	(4.74)	(2.64)
MktRF	0.98	1.00	0.98
	(392.83)	(273.89)	(265.19)
SMB	-0.21	-0.18	-0.15
	(-23.68)	(-15.75)	(-22.95)
HML	0.01	-0.05	0.07
	(1.49)	(-6.35)	(9.97)
RMW	0.02	-0.05	-0.00
	(1.74)	(-3.77)	(-0.11)
CMA	0.05	-0.14	0.02
	(3.48)	(-7.58)	(1.28)
UMD	-0.13	-0.14	-0.13
	(-17.85)	(-15.73)	(-19.67)
R^2	0.9833	0.9760	0.9810
IR	0.3423	0.8479	0.4581

Note: This table presents portfolio performance for three core model strategies: Actual Holdings (market-cap weighted household direct stock holdings), Sentiment (belief-based model with all parameters), and Characteristics (characteristics-only model with sentiment parameters set to zero). Portfolios are formed from the top 300 U.S. stocks by market capitalization and rebalanced quarterly. Returns are averaged over lags 0 through 7 (eight quarters). Panel A reports annualized mean excess returns, t-statistics, the number of daily observations, and annualized Sharpe ratios. Panel B reports CAPM alpha, market beta, R-squared, and the information ratio. Panel C reports six-factor alpha and factor loadings from a regression on the Fama-French five factors (MktRF, SMB, HML, RMW, CMA) plus momentum (UMD), along with R-squared and the information ratio. All t-statistics are computed using Newey-West (HAC) standard errors with a Bartlett kernel and automatic bandwidth selection following Newey and West (1994).

Table 7: Consumer and Producer Portfolios (50% Threshold): Lags 0 to 7

	Panel A: Consumers			*Panel B: Producers*		
	Cons50 (Sent)	Cons50 (All)	Cons50 L/S	Prod50 (Sent)	Prod50 (All)	Prod50 L/S
A. Raw excess returns						
Average	11.38	10.20	1.18	10.29	9.77	0.52
	(4.02)	(3.76)	(2.15)	(2.96)	(2.90)	(1.04)
T	8867	8867	8867	8867	8867	8867
SR	0.5985	0.5460	0.3474	0.4682	0.4549	0.1615
B. CAPM regressions						
Intercept	1.85	0.77	1.07	-0.49	-0.91	0.42
	(2.20)	(1.00)	(1.97)	(-0.38)	(-0.82)	(0.84)
MktRF	1.01	1.00	0.01	1.14	1.13	0.01
	(127.24)	(130.09)	(1.75)	(118.96)	(128.05)	(2.34)
R^2	0.9305	0.9425	0.0035	0.8899	0.9151	0.0036
IR	0.3684	0.1723	0.3171	-0.0666	-0.1451	0.1305
C. Six-factor regressions						
Intercept	3.05	1.11	1.94	1.88	1.31	0.57
	(4.42)	(2.00)	(3.92)	(1.54)	(1.27)	(1.14)
MktRF	0.98	0.99	-0.01	1.10	1.09	0.01
	(185.55)	(225.20)	(-2.35)	(118.63)	(151.73)	(1.51)
SMB	-0.23	-0.25	0.02	-0.05	-0.06	0.01
	(-16.94)	(-20.56)	(3.84)	(-3.33)	(-5.10)	(2.19)
HML	-0.06	0.04	-0.11	-0.15	-0.10	-0.04
	(-5.91)	(2.84)	(-8.87)	(-6.74)	(-6.02)	(-5.15)
RMW	0.05	0.08	-0.03	-0.14	-0.14	0.00
	(3.21)	(5.96)	(-3.27)	(-5.31)	(-6.14)	(0.39)
CMA	-0.13	0.01	-0.14	-0.06	-0.00	-0.06
	(-6.19)	(0.51)	(-8.85)	(-1.80)	(-0.06)	(-4.51)
UMD	-0.11	-0.11	-0.00	-0.16	-0.18	0.02
	(-10.95)	(-13.03)	(-0.19)	(-10.85)	(-12.10)	(3.10)
R^2	0.9541	0.9699	0.2601	0.9029	0.9283	0.0640
IR	0.7485	0.3437	0.6631	0.2742	0.2270	0.1828

Note: This table presents portfolio performance comparing sentiment-weighted and market-capitalization-weighted strategies for consumer-facing and producer-facing stocks, classified using a 50% consumer spending share threshold. Consumer (producer) stocks are those whose NAICS industry consumer spending share is above (below) the threshold. The Sent columns weight stocks by estimated household sentiment parameters; the All columns use market-capitalization weights as a benchmark. The L/S columns report the long-short spread between Sent and All within each group. Portfolios are formed from the top 300 U.S. stocks by market capitalization and rebalanced quarterly. Returns are averaged over lags 0 through 7 (eight quarters). Panel A reports annualized mean excess returns, t-statistics, the number of daily observations, and annualized Sharpe ratios. Panel B reports CAPM alpha, market beta, R-squared, and the information ratio. Panel C reports six-factor alpha and factor loadings from a regression on the Fama-French five factors (MktRF, SMB, HML, RMW, CMA) plus momentum (UMD), along with R-squared and the information ratio. All t-statistics are computed using Newey-West (HAC) standard errors with a Bartlett kernel and automatic bandwidth selection following Newey and West (1994).

7 Conclusion

This paper has devised a new method for estimating the asset demand of households using only aggregate data, observable characteristics, and instrumental variables. I show how to map the canonical Berry, Levinsohn and Pakes (1995) discrete choice demand estimation approach onto a setting with portfolio choice under restrictions on the covariance matrix. Using the GEL-MPEC approach of Conlon (2013), I show how to decompose household demand into (a) stock-specific household sentiment ξ, (b) beliefs about returns to factor exposure $\bar{\beta}$, and (c) disagreement terms that capture the dispersion of beliefs among households. With the methodology established I apply it to estimating household demand from 1984Q4 through 2021Q4. These estimates reveal that household sentiment, in contrast to some of the previous literature's conclusions, positively predicts future returns. Moreover, this potential for private information among households is mostly concentrated in the equity shares of consumer-facing companies such as restaurants or soft drink manufacturers as opposed to business-oriented companies such as refiners or steel mills.

The natural question that arises from this paper's empirical results is whether households should directly invest in common stocks. While one might be tempted to reach such a conclusion given the return predictability results, I argue that the overwhelming evidence continues to point toward households being better suited to index investing. The informational content of household sentiment is in and of itself an interesting object to study in the context of household investor beliefs and mistakes, but is not an object readily available to household investors. Instead, each investor is forced to trade on the basis of his or her limited set of information. Essentially, structural parameters are not contained in any household's information set. Moreover, and most importantly, households generally arrive at their portfolios by way of trading. This paper has focused exclusively on the informational content of *ownership*, which is inherently different from *trading*. Households' tendency to trade attention-grabbing stocks to their own detriment is layered on top of this paper's sentiment results. Moreover, households are unlikely to be able to disentangle their (unhelpful) market timing and factor-based beliefs from their idiosyncratic beliefs. While this paper's empirical evidence does not support encouraging retail traders to trade, it does provide evidence against the narrative that households are collectively devoid of information in their portfolio composition. Their trading behavior as described in the rest of the litearture, however, does not appear to be predominantly informed.

References

Baker, Malcolm, Robin Greenwood, and Jeffrey Wurgler. 2009. "Catering through Nominal Share Prices." *Journal of Finance*, 64(6): 2559–2590.

Barber, Brad M, and Terrance Odean. 2000. "Trading is hazardous to your wealth: The common stock investment performance of individual investors." *The Journal of Finance*, 55(2): 773–806.

Barber, Brad M, and Terrance Odean. 2001. "Boys will be boys: Gender, overconfidence, and common stock investment." *The Quarterly Journal of Economics*, 116(1): 261–292.

Barber, Brad M, and Terrance Odean. 2008. "All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors." *Review of Financial Studies*, 21(2): 785–818.

Barber, Brad M., Shengle Lin, and Terrance Odean. 2024. "Resolving a Paradox: Retail Trades Positively Predict Returns but Are Not Profitable." *Journal of Financial and Quantitative Analysis*, 59: 2547–2581.

Barber, Brad M, Yi-Tsung Lee, Yu-Jane Liu, and Terrance Odean. 2009. "Just how much do individual investors lose by trading?" *Review of Financial Studies*, 22(2): 609–632.

Barberis, Nicholas, and Ming Huang. 2008. "Stocks as Lotteries: The Implications of Probability Weighting for Security Prices." *American Economic Review*, 98(5): 2066–2100.

Barberis, Nicholas, and Wei Xiong. 2009. "What Drives the Disposition Effect? An Analysis of a Long-Standing Preference-Based Explanation." *the Journal of Finance*, 64(2): 751–784.

Barberis, Nicholas, and Wei Xiong. 2012. "Realization utility." *Journal of Financial Economics*, 104(2): 251–271.

Barberis, Nicholas, Ming Huang, and Richard Thaler. 2006. "Individual Preferences, Monetary Gambles, and Stock Market Participation: A Case for Narrow Framing." *American Economic Review*, 96(4): 1069–1090.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer. 2015. "X-CAPM: An extrapolative capital asset pricing model." *Journal of Financial Economics*, 115(1): 1–24.

Barberis, Nicholas, Robin Greenwood, Lawrence Jin, and Andrei Shleifer. 2018. "Extrapolation and bubbles." *Journal of Financial Economics*, 129(2): 203–227.

Berry, Steven. 1994. "Estimating Discrete-Choice Models of Product Differentiation." *RAND Journal of Economics*, 25(2): 242–262.

Berry, Steven, Amit Gandhi, and Philip Haile. 2013. "Connected Substitutes and Invertibility of Demand." *Econometrica*, 81(5): 2087–2111.

Berry, Steven, and Philip Haile. 2014. "Identification in Differentiated Products Markets Using Market Level Data." *Econometrica*, 82: 1749–1797.

Berry, Steven, James Levinsohn, and Ariel Pakes. 1995. "Automobile prices in market equilibrium." *Econometrica: Journal of the Econometric Society*, 841–890.

Birru, Justin. 2015. "Confusion of Confusions: A Test of the Disposition Effect and Momentum." *Review of Financial Studies*, 28(7): 1849–1873.

Birru, Justin, and Baolian Wang. 2016. "Nominal Price Illusion." *Journal of Financial Economics*, 119(3): 578–598.

Boehmer, Ekkehart, Charles M. Jones, Xiaoyan Zhang, and Xinran Zhang. 2021. "Tracking Retail Investor Activity." *Journal of Finance*, 76(5): 2249–2305.

Bordalo, Pedro, Nicola Gennaioli, and Andrei Shleifer. 2012. "Salience Theory of Choice Under Risk." *The Quarterly Journal of Economics*, 127(3): 1243–1285.

Carhart, Mark M. 1997. "On persistence in mutual fund performance." *The Journal of Finance*, 52(1): 57–82.

Conlon, Christopher T. 2012. "A Dynamic Model of Prices and Margins in the LCD TV Industry." *Working Paper*.

Conlon, Christopher T. 2013. "The Empirical Likelihood MPEC Approach to Demand Estimation." Working paper, Columbia University. First version: January 2009.

Coval, Joshua D., David Hirshleifer, and Tyler Shumway. 2021. "Can Individual Investors Beat the Market?" *Review of Asset Pricing Studies*, 11(3): 552–579.

Cox, Justin, Bonnie Van Ness, and Robert Van Ness. 2022. "Stock Splits and Retail Trading." *Financial Review*, 57(4): 731–750.

Da, Zhi, Joseph Engelberg, and Pengjie Gao. 2011. "In search of attention." *The Journal of Finance*, 66(5): 1461–1499.

De Long, J Bradford, Andrei Shleifer, Lawrence H Summers, and Robert J Waldmann. 1990. "Noise Trader Risk in Financial Markets." *Journal of Political Economy*, 98(4): 703–738.

Dubé, Jean-Pierre, Jeremy T. Fox, and Che-Lin Su. 2012. "Improving the Numerical Performance of Static and Dynamic Aggregate Discrete Choice Random Coefficients Demand Estimation." *The Review of Economic Studies*, 79(5): 2231–2267.

Egan, Mark, Alexander MacKay, and Hanbin Yang. 2022. "Recovering Investor Expectations from Demand for Index Funds." *The Review of Economic Studies*, 89(5): 2559–2599.

Engelberg, Joseph, Caroline Sasseville, and Jared Williams. 2012. "Market madness? the case of mad money." *Management Science*, 58(2): 351–364.

Engelberg, Joseph E, and Christopher A Parsons. 2011. "The causal impact of media in financial markets." *The Journal of Finance*, 66(1): 67–97.

Fama, Eugene F., and Kenneth R. French. 2015. "A five-factor asset pricing model." *Journal of Financial Economics*, 116(1): 1–22.

Frazzini, Andrea, and Lasse Heje Pedersen. 2014. "Betting Against Beta." *Journal of Financial Economics*, 111(1): 1–25.

Frydman, Cary, Nicholas Barberis, Colin Camerer, Peter Bossaerts, and Antonio Rangel. 2014. "Using Neural Data to Test a Theory of Investor Behavior: An Application to Realization Utility." *The Journal of Finance*, 69(2): 907–946.

Gabaix, Xavier, Ralph J. Koijen, Federico Mainardi, Sangmin Oh, and Motohiro Yogo. 2022. "Asset Demand of U.S. Households." *SSRN*.

Goeree, Michelle Sovinsky. 2008. "Limited Information and Advertising in the U.S. Personal Computer Industry." *Econometrica*, 76(5): 1017–1074.

Graves, Daniel. 2023. "Attention and the Retail Alignment Puzzle." Working Paper, SSRN No. 4531906.

Graves, Daniel. 2025. "What Lies Beneath Zero: Censoring, Demand Estimation, and Hidden Beliefs." *Working Paper*. Job Market Paper, Yale University.

Haddad, Valentin, Zhiguo He, Paul Huebner, Péter Kondor, and Erik Loualiche. 2025. "Causal Inference for Asset Pricing." *Working Paper*. SSRN No. 5187305. Swiss Finance Institute Outstanding Paper Award 2025.

Hansen, Lars Peter, John Heaton, and Amir Yaron. 1996. "Finite-Sample Properties of Some Alternative GMM Estimators." *Journal of Business & Economic Statistics*, 14(3): 262–280.

Heimer, Rawley, and Alp Simsek. 2019. "Should retail investors' leverage be limited?" *Journal of Financial Economics*, 132(3): 1–21.

Kaniel, Ron, Gideon Saar, and Sheridan Titman. 2008. "Individual investor trading and stock returns." *The Journal of Finance*, 63(1): 273–310.

Kaniel, Ron, Shuming Liu, Gideon Saar, and Sheridan Titman. 2012. "Individual investor trading and return patterns around earnings announcements." *The Journal of Finance*, 67(2): 639–680.

Kelley, Eric K., and Paul C. Tetlock. 2013. "How Wise Are Crowds? Insights from Retail Orders and Stock Returns." *Journal of Finance*, 68(3): 1229–1265.

Kitamura, Yuichi, and Michael Stutzer. 1997. "An Information-Theoretic Alternative to Generalized Method of Moments Estimation." *Econometrica*, 65(4): 861–874.

Koijen, Ralph S. J., and Motohiro Yogo. 2019. "A Demand System Approach to Asset Pricing." *Journal of Political Economy*, 127(4): 1475 – 1515.

Koijen, Ralph S. J., and Motohiro Yogo. 2023. "Understanding the Ownership Structure of Corporate Bonds." *American Economic Review: Insights*, 5(1): 73–92.

Koijen, Ralph S. J., and Motohiro Yogo. 2026. "Exchange Rates and Asset Prices in a Global Demand System." *Journal of Political Economy*. Forthcoming.

Kumar, Alok. 2009. "Who gambles in the stock market?" *The Journal of Finance*, 64(4): 1889–1933.

Luo, Patrick, Enrichetta Ravina, Marco C Sammon, and Luis M Viceira. 2025. "Retail Investors' Contrarian Behavior Around News, Attention, and the Momentum Effect." *NBER working paper*.

Miller, E.M. 1977. "Risk, Uncertainty, and Divergence of Opinion." *Journal of Finance*, 1151–1168.

Newey, Whitney K., and Richard J. Smith. 2004. "Higher Order Properties of Gmm and Generalized Empirical Likelihood Estimators." *Econometrica*, 72(1): 219–255.

Odean, Terrance. 1999. "Do Investors Trade Too Much?" *The American Economic Review*, 89(5): 1279–1298.

Petersen, Mitchell A. 2009. "Estimating Standard Errors in Finance Panel Data Sets: Comparing Approaches." *Review of Financial Studies*, 22(1): 435–480.

Ritter, Jay R. 1991. "The Long-Run Performance of Initial Public Offerings." *The Journal of Finance*, 46(1): 3–27.

Shefrin, Hersh, and Meir Statman. 1985. "The disposition to sell winners too early and ride losers too long: Theory and evidence." *The Journal of finance*, 40(3): 777–790.

Su, Che-Lin, and Kenneth L. Judd. 2012. "Constrained Optimization Approaches to Estimation of Structural Models." *Econometrica*, 80(5): 2213–2230.

Online Appendix

A Optimal Household Portfolio Choice **50**

B Existence of Equilibrium **51**

B.1 First Model with No SSCs for Institutions (Existence and uniqueness, via Banach) . 51

B.2 Institutions have SSCs and Leverage Constraints, Households have 52

C Household Contraction Mapping Function **53**

D Data Construction Details **55**

D.1 Stock-Level Data 55

D.2 Household Holdings 56

D.3 Idiosyncratic Return Variance 56

D.4 Attention Vectors 56

D.5 Sample Selection and Variable Construction 57

E Cluster-Based Estimation using Knitro Package **57**

E.1 Optimization Problem Structure 58

E.2 Knitro Solver Configuration 58

E.3 Cluster Computing Environment 59

E.4 Starting Values and Initialization via NFXP 59

E.5 Non-Convergence and Reoptimization 60

F No Attention and Alternative Specification Sentiment Autocorrelation Analysis **60**

G Robustness: Additional Analysis (Main Specification Comparison) **62**

H Robustness: Using Market Caps instead of Attention Model **64**

I Robustness: Alternative Specifications **68**

J Classification of Companies by Industry into B2B vs B2C **78**

A Optimal Household Portfolio Choice

We solve the more general standard optimization problem

$$\max_{\sum w_{i,t}(n)\leq L_{i,t},\, w_{i,t}(n)\geq -L_{i,t}^{S},\, w_{i,t}(n)=0\forall n\notin\mathcal{H}_{i,t}} \mathbb{E}_{i,t}\left[\log\left(A_{i,T}\right)\right] \tag{20}$$

by forming the Lagrangian

$$\mathcal{L}_{i,t}=\mathbb{E}_{i,t}\left[\log\left(A_{i,T}\right)+\sum_{s=t}^{T-1}\left[\Lambda_{i,s}'\left(L_{i,s}^{S}+w_{i,s}\right)+\lambda_{i,s}\left(L_{i,s}-w_{i,s}'\mathbf{1}\right)\right]\right]. \tag{21}$$

Using the fact that we can decompose the Lagrangian as in Koijen and Yogo (2019) into

$$\mathcal{L}_{i,t}=\log\left(A_{i,t}\right)+\mathbb{E}_{i,t}\left[\sum_{s=t}^{T-1}\log\left(\frac{A_{i,s+1}}{A_{i,s}}\right)+\sum_{s=t}^{T-1}\left[\Lambda_{i,s}'\left(L_{i,s}^{S}+w_{i,s}\right)+\lambda_{i,s}\left(L_{i,s}-w_{i,s}'\mathbf{1}\right)\right]\right] \tag{22}$$

which we rewrite as

$$\mathcal{L}_{i,t}=\log\left(A_{i,t}\right)+\mathbb{E}_{i,t}\left[\sum_{s=t}^{T-1}\log\left(w_{i,s}'R_{s+1}\right)+\sum_{s=t}^{T-1}\left[\Lambda_{i,s}'\left(L_{i,s}^{S}+w_{i,s}\right)+\lambda_{i,s}\left(L_{i,s}-w_{i,s}'\mathbf{1}\right)\right]\right]. \tag{23}$$

We readily see that the optimal portfolio problem is myopic with the solution given by the first order condition

$$\frac{\partial\mathcal{L}_{i,t}}{\partial w_{i,t}}=\mathbb{E}_{i,t}\left[\frac{1}{w_{i,t}'R_{t+1}}R_{t+1}+\Lambda_{i,t}-\lambda_{i,t}\mathbf{1}\right]=0. \tag{24}$$

$$\frac{\partial\mathcal{L}_{i,t}}{\partial w_{i,t}}=\mathbb{E}_{i,t}\left[\frac{1}{w_{i,t}'R_{t+1}}R_{t+1}\right]=\lambda_{i,t}\mathbf{1}-\Lambda_{i,t} \tag{25}$$

as well as the complementary slackness conditions $\Lambda_{i,s}'\left(L_{i,s}^{S}+w_{i,s}\right)=0$ and $\lambda_{i,s}\left(L_{i,s}-w_{i,s}'\mathbf{1}\right)=0$.

For household h's portfolio problem with a shortsale constraint equivalent to setting $L_{h,t}^{s}=0$, we have

$$w_{h,t}\approx\left(\Sigma_{h,t}\right)^{-1}\left[\mu_{h,t}+\Lambda_{h,t}-\lambda_{h,t}\mathbf{1}\right], \tag{26}$$

which can be re-expressed as

$$w_{h,t}=\Sigma_{h,t}^{-1}\max\left\{0,\mu_{h,t}-\lambda_{h,t}\mathbf{1}\right\} \tag{27}$$

since $\Lambda_{h,t}\left(n\right)>0$ if and only if $w_{h,t}=0$ by complementary slackness and $\Sigma_{h,t}$ is assumed diagonal. Using the fact that $\Sigma_{h,t}$ is diagonal,

$$\lambda_{h,t}=\max\left\{\frac{\sum_{n\in C_{h,t}:\Lambda_{h,t}(n)=0}\sigma_{t}^{-2}\left(n\right)\mu_{h,t}\left(n\right)-1}{\sum_{n_{1}\in C_{h,t},\Lambda_{h,t}(n_{1})=0}\sigma_{t}^{-2}\left(n_{1}\right)},0\right\}.$$

B Existence of Equilibrium

B.1 First Model with No SSCs for Institutions (Existence and uniqueness, via Banach)

Proof. (Proposition 1) Let us consider first institutional demand $J(P_t)$ where characteristics $x_{k,t}$, $k > 0$, covariance matrices $\Sigma_{i,t}$, and coefficients $\beta_{k,i,t}$ are all exogenously given ($\Sigma_{i,t}$ does not depend on P_t by assumption). Aggregate dollar demand for assets is given by the vector

$$J(P_t) = \sum_i A_{i,t} w_{i,t} = \sum_i A_{i,t} \Sigma_{i,t}^{-1} \mu_{i,t} = \sum_i A_{i,t} \Sigma_{i,t}^{-1} \left(P_t \beta_{0,i,t} + \sum_{k=1}^{K} x_{k,t} \beta_{k,i,t} + \epsilon_{i,t} \right). \tag{28}$$

Note that $\Sigma_{i,t}^{-1} = H_{i,t} \Sigma_{i,t,all}^{-1} H_{i,t}$ where $\Sigma_{i,t,all}$ is the covariance matrix over all N risky assets and $H_{i,t}$ is a selector matrix with all $0's$ except for $1's$ on the diagonal corresponding to assets $n \in \mathcal{H}_{i,t}$ that are in i's habitat.

Let us now consider two price vectors, $P_{1,t}$ and $P_{2,t}$. Taking the difference, we obtain

$$\begin{aligned} & J(P_{1,t}) - J(P_{2,t}) \\ &= \sum_i A_{i,t} \beta_{0,i,t} \Sigma_{i,t}^{-1} (P_{1,t} - P_{2,t}) \\ &\leq \| \sum_i A_{i,t} \beta_{0,i,t} \Sigma_{i,t}^{-1} \|_\infty \| P_{1,t} - P_{2,t} \|_\infty \end{aligned}$$

where $\|x\|_\infty$ is the usual max norm $\max_n |x_n|$ when $x \in \mathbb{R}^N$ and $\|A\|_\infty = \max_{x \in \mathbb{R}^N, \|x\|_\infty = 1} \|Ax\|_\infty$ is the matrix norm induced by the l_∞ norm.

We next consider aggregate household demand $H(P_t)$, using the convention that $H_n(P_t)$ refers to the nth element of the vector $H(P_t) \in \mathbb{R}^N$. We start by restating the elementwise aggregate household demand function $H_n(P_t)$:

$$\begin{aligned} H_n(P_t) = A_{H,t} \int 1_{n \in \mathcal{C}_{h,t}} \sigma_t^{-2}(n) \max \Bigg\{ & 0, \bar{\beta}_{0,t} P_t(n) + \sum_{k=1}^{K} \bar{\beta}_{k,t} x_{k,t}(n) \\ & + \xi_t + \sigma_{0,t} P_t(n) \nu_{0,h,t} + \sum_{k=1}^{K} \sigma_{k,t} x_{k,t} \nu_{k,h,t} + \sigma_{\epsilon,t} \epsilon_{h,t} \Bigg\} dh \end{aligned}$$

Using the triangle inequality we can bound the difference $H_n(P_{1,t}) - H_n(P_{2,t})$:

$$\begin{aligned} & H_n(P_{1,t}) - H_n(P_{2,t}) \\ &\leq \left(A_{H,t} \int 1_{n \in \mathcal{C}_{h,t}} \sigma_t^{-2}(n) |\bar{\beta}_{0,t} + \sigma_{0,t} \nu_{0,h,t}| dh \right) |P_{1,t}(n) - P_{2,t}(n)| \\ &\leq \| A_{H,t} \int 1_{n \in \mathcal{C}_{h,t}} \Sigma_{H,t}^{-1} \mathbf{1} |\bar{\beta}_{0,t} + \sigma_{0,t} \nu_{0,h,t}| dh \|_\infty \| P_{1,t} - P_{2,t} \|_\infty. \end{aligned}$$

Putting together these two constraints with the condition from Assumption 2, we obtain

$$\begin{aligned}
&\|X\left(P_{1,t}\right)-X\left(P_{2,t}\right)\|_{\infty} \\
&\leq\|J\left(P_{1,t}\right)-J\left(P_{2,t}\right)\|_{\infty}+\|H\left(P_{1,t}\right)-H\left(P_{2,t}\right)\|_{\infty} \\
&\leq\|\sum_{i} A_{i,t}\beta_{0,i,t}\Sigma_{i,t}^{-1}\|_{\infty}\|P_{1,t}-P_{2,t}\|_{\infty} \\
&+\|A_{H,t}\int 1_{n\in\mathcal{C}_{h,t}}\Sigma_{H,t}^{-1}\mathbf{1}|\bar{\beta}_{0,t}+\sigma_{0,t}\nu_{0,h,t}|dh\|_{\infty}\|P_{1,t}-P_{2,t}\|_{\infty} \\
&\left(\|\sum_{i} A_{i,t}\beta_{0,i,t}\Sigma_{i,t}^{-1}\|_{\infty}+\|A_{H,t}\int 1_{n\in\mathcal{C}_{h,t}}\Sigma_{H,t}^{-1}\mathbf{1}|\bar{\beta}_{0,t}+\sigma_{0,t}\nu_{0,h,t}|dh\|_{\infty}\right)\|P_{1,t}-P_{2,t}\|_{\infty} \\
&<\|P_{1,t}-P_{2,t}\|_{\infty}.
\end{aligned}$$

We have thus established all of the conditions of the Banach fixed point theorem. Aggregate demand is a contraction mapping, implying the existence of a unique fixed point P_t such that $X\left(P_t\right)=P_t$. □

B.2 Institutions have SSCs and Leverage Constraints, Households have

Proof. (Proposition 1, continued) In this subsection we consider existence of equilibrium when households are additionally constrained by $\sum_n w_{h,t}\left(n\right)\leq 1$ while institutions are short sale constrained by institution-specific limits $L_{i,t}^{S}\in\mathbb{R}_{+}^{N}$ such that $w_{i,t}\left(n\right)\geq -L_{i,t}^{S}\left(n\right)$ and also constrained by total leverage limits $L_{i,t}\in\mathbb{R}^{+}$ such that $\sum_n w_{i,t}\left(n\right)\leq L_{i,t}$. The aggregate demand function $X\left(P\right)$ is then given by

$$\begin{aligned}
X\left(P_t\right) &= \sum_{i} A_{i,t}w_{i,t}+A_{H,t}\int_{h} w_{h,t}dh \\
&= \sum_{i} A_{i,t}D_{i,t}^{-1}\left(P_t\left(\beta_{0,i,t}-\kappa_{i,t}\Psi_{0,i,t}\right)+\sum_{k=1}^{K} x_{k,t}\left(\beta_{k,i,t}-\kappa_{i,t}\Psi_{k,i,t}\right)+\epsilon_{i,t}-\mathbf{1}\lambda_{i,t}+\Lambda_{i,t}\right) \\
&\quad +A_{H,t}\int C_{h,t}\Sigma_{H,t}^{-1}\max\left\{0,\left(\bar{\beta}_{0,t}+\sigma_{0,t}\nu_{0,h,t}\right)P_t\right. \\
&\quad \left.+\sum_{k=1}^{K}\bar{\beta}_{k,t}x_{k,t}+\xi_t+\sum_{k=1}^{K}\sigma_{k,t}x_{k,t}\nu_{k,h,t}+\sigma_{\epsilon,t}\epsilon_{h,t}-\lambda_{h,t}\mathbf{1}\right\}dh \\
&= J\left(P_t\right)+H\left(P_t\right)
\end{aligned} \tag{29}$$

where we have

$$\kappa_{i,t}=\Gamma_{i,t}'D_{i,t}^{-1}\left(\mu_{i,t}+\Lambda_{i,t}-\lambda_{i,t}\mathbf{1}\right)\left(1+\Gamma_{i,t}'D_{i,t}^{-1}\Gamma_{i,t}\right)^{-1}. \tag{30}$$

We clearly have that the function $X\left(P\right)$ is continuous in its argument since P enters either linearly or as the integral of piecewise linear max operators. We then need to show the second condition, namely that the function maps a convex, compact subset $\mathcal{K}$ of $\mathbb{R}^N$ to itself. First, we note that S^{inv} is strictly positive, so we need only look to bound the optimal market equity

component of the equation. Furthermore, we note that households are constrained to only hold positive positions so that the vector $H(P) \geq 0$. We then must examine $J(P)$ to establish a lower bound. The maximum short allocation that can be made by institution i is given by $L_{i,t}^S$, so we have that regardless of coefficients on market equity given by $\beta_{0,i,t}$, $J(P) \geq -\sum_i A_{i,t} L_{i,t}^S$ and therefore $X(P) \geq -\sum_i A_{i,t} L_{i,t}^S = \underline{P}$. Additionally, we have the obvious constraint $J(P) \leq \sum_i A_{i,t} L_{i,t}$ and $H(P) \leq A_{H,t}$ so that $X(P) \leq \sum_i A_{i,t} L_{i,t} + A_{H,t} = \bar{P}$. Then $X(P) \in \left[\underline{P}, \bar{P}\right] \forall P \in \left[\underline{P}, \bar{P}\right]$ where we have abused notation and this refers to elements taking values in the appropriate ranges so that P lies in the Cartesian product of each element's intervals. This establishes the conditions of the Brouwer Fixed Point theorem and therefore equilibrium always exists. □

C Household Contraction Mapping Function

Proof. In order to do so we show that the function

$$f(\delta) = \delta + A_{H,t}^{-1} \Sigma_{H,t} \hat{H} - \int C_{h,t} \left(\max \left\{ 0, \delta + \sum_{k=0}^{K} \sigma_{k,t} x_{k,t} \nu_{k,h,t} + \sigma_{\epsilon,t} \epsilon_{h,t} - \lambda_{h,t} \mathbf{1} \right\} \right) dh$$

is a contraction mapping we establish that the conditions of the theorem in Appendix I of Berry, Levinsohn and Pakes (1995) hold.

The function of δ given by $h(\delta) = \int C_{h,t} \left(\max \left\{ 0, \delta + \sum_{k=0}^{K} \sigma_{k,t} x_{k,t} \nu_{k,h,t} + \sigma_{\epsilon,t} \epsilon_{h,t} - \lambda_{h,t} \mathbf{1} \right\} \right) dh$, by direct inspection has three obvious properties: (1) it is clearly bounded below by 0, (2) it is continuously differentiable, and (3) for any given ϵ, there exists a $\delta = (-1, -1, ..., -1)\kappa$ for $\kappa > 0$ sufficiently large such that $\| \frac{1}{A_{i,t}} \Sigma_{i,t} ME_{Households}(\delta) \|_1 < \epsilon$. In integrating over the space of all households, because of the multinomial sampling construction of consideration sets, we are summing over a discrete (but large) number of possible consideration sets and then integrating over the distribution of household beliefs for each consideration set. Consideration sets $C_{h,t}$ carry weights given by $\Pr(C_{h,t}) = \prod_n g_t(n)^{1_{n \in C_{h,t}}} (1 - g_t(n))^{1_{n \notin C_{h,t}}}$ for all sets $C_{h,t}$ of order ranging from 0 to N. Within the measure of each of these consideration sets we assume a continuum of households with preferences characterized by $\nu_{h,t} \sim N(0, \Gamma)$ and the n-tuple $\epsilon_{h,t} \sim N(0, \Psi)$ with Γ and Ψ diagonal. With $\delta_t = \sum_{k=0}^{K} \bar{\beta}_{k,t} x_{k,t} + \xi_t$, differentiating f_j with respect to δ_n (and suppressing all time subscripts for notation convenience) we obtain

$$f(\delta) = \delta + A_H^{-1} \Sigma_H \hat{H} - \sum_{C_h} \Pr(C_h) \int \mathbf{1}_{\mu_h > \lambda_h} \odot \mathbf{1}_{C_h} \odot \max \left\{ 0, \mu_h - \frac{1}{\sum_{n_1 \in C_h : \Lambda_h(n_1) = 0} \sigma^{-2}(n_1)} \max \left\{ \sum_{n \in C_h : \Lambda_h(n) = 0} \max \left\{ 0, \sigma^{-2}(n) \mu_h(n) \right\} - L_H, 0 \right\} \right\} dh \tag{31}$$

$$\frac{\partial f_n}{\partial \delta_n} = 1 - \frac{\partial h(\delta)_n}{\partial \delta_n}, j = n$$

and

$$\frac{\partial f_j}{\partial \delta_n} = -\frac{\partial h(\delta)_j}{\partial \delta_n}, j \neq n$$

where

$$h(\delta) = \sum_{C_h} \Pr(C_h) \int \mathbf{1}_{\mu_h > \lambda_h} \odot \mathbf{1}_{C_h} \odot$$

$$\max\left\{0, \mu_h - \frac{1}{\sum_{n_1 \in C_h : \Lambda_h(n_1)=0} \sigma^{-2}(n_1)} \max\left\{\sum_{n \in C_h : \Lambda_h(n)=0} \max\left\{0, \sigma^{-2}(n)\,\mu_h(n)\right\} - L_H, 0\right\}\right\} dh$$

We start with the case $j \neq n$: trivially, increasing the expected return of asset n cannot increase the allocation to asset j because we have assumed a diagonal covariance matrix; it can only serve to tighten the leverage constraint captured by the Lagrange multiplier term λ_h. This means that $\frac{\partial f_j}{\partial \delta_n} > 0$, as we will now explicitly show. We start with the formula for the Lagrange multiplier

$$\lambda_h = \frac{\max\left\{\mathbf{1}\left(\Sigma^{(1,1)}\right)^{-1}\mu_h^{(1)} - L_H, 0\right\}}{\mathbf{1}\left(\Sigma^{(1,1)}\right)^{-1}\mathbf{1}}$$

so we have that $\frac{\partial \lambda_h}{\partial \delta_n} = 0$ if $\mu_h(n) < 0$ and $\frac{\partial \lambda_h}{\partial \delta_n} = \frac{\sigma^{-2}(n)}{\mathbf{1}\left(\Sigma^{(1,1)}\right)^{-1}\mathbf{1}}$ if $\mu_h(n) > 0$ and $\mathbf{1}\left(\Sigma^{(1,1)}\right)^{-1}\mu_h^{(1)} > L_H$. The second constraint determines whether maximum leverage is reached: if asset returns are sufficiently high relative to variance (the left hand term) then the household will desire to purchase more assets than the leverage constraint permits. This will result in all holdings being shrunk by $\Sigma^{-1}\lambda_h$, which captures the degree to which the leverage constraint binds. When $\lambda_h = 0$ the leverage constraint is slack. We then have

$$\frac{\partial h(\delta)_j}{\partial \delta_n} = -\sum_{C_h} \Pr(C_h) \int \left(1_{\mu_h(j) > \lambda_h, \mu_h(n) > \lambda_h, j \in C_h} \frac{\sigma^{-2}(n)}{\sum_{n_1 \in C_h : \Lambda_h(n_1)=0} \sigma^{-2}(n_1)}\right) dh < 0$$

where the inequality follows trivially from the fact that the function being integrated over is strictly positive. Therefore

$$\frac{\partial f_j}{\partial \delta_n} = -\frac{\partial h(\delta)_j}{\partial \delta_n} > 0, j \neq n.$$

Next we turn to the case where $j = n$. We then find

$$\frac{\partial h(\delta)_n}{\partial \delta_n} = \sum_{C_h} \Pr(C_h) \int 1_{\Lambda_h(n)=0, n \in C_h} \left(1 - \frac{\sigma^{-2}(n)}{\sum_{n_1 \in C_h : \Lambda_h(n_1)=0} \sigma^{-2}(n_1)}\right) dh < 1$$

since trivially $0 < \frac{\sigma^{-2}(n)}{\sum_{n_1 \in C_h : \mu_h(n_1) > 0} \sigma^{-2}(n_1)} < 1$ and the household space for each consideration set has measure 1. This establishes that

$$\frac{\partial f_n}{\partial \delta_n} = 1 - \frac{\partial h(\delta)_n}{\partial \delta_n} > 0.$$

We next note that

$$\sum_{k=1}^{N}\frac{\partial f_n}{\partial \delta_k}$$

$$= 1-\sum_{C_h}\Pr\left(C_h\right)\int 1_{\Lambda_h(n)=0,n\in C_h}\left(1-\frac{\sigma^{-2}\left(n\right)}{\sum_{n_1\in C_h:\mu_h(n_1)>0}\sigma^{-2}\left(n_1\right)}\right)dh$$

$$+\sum_{k\neq n}\sum_{C_h}\Pr\left(C_h\right)\int\left(1_{\Lambda_h(k)=0,\Lambda_h(n)=0,k\in C_h}\frac{\sigma^{-2}\left(k\right)}{\sum_{n_1\in C_h:\Lambda_h(n_1)=0}\sigma^{-2}\left(n_1\right)}\right)dh$$

$$= 1-\sum_{C_h}\Pr\left(C_h\right)\int 1_{\Lambda_h(n)=0,n\in C_h}dh+$$

$$\sum_{C_h}\Pr\left(C_h\right)\int\left(\sum_{k=1}^{N}1_{\Lambda_h(k)=0,\Lambda_h(n)=0,k\in C_h}\frac{\sigma^{-2}\left(k\right)}{\sum_{n_1\in C_h:\Lambda_h(n_1)=0}\sigma^{-2}\left(n_1\right)}\right)dh$$

$$< 1-\sum_{C_h}\Pr\left(C_h\right)\int 1_{\Lambda_h(n)=0,n\in C_h}dh+\sum_{C_h}\Pr\left(C_h\right)\int\left(1_{\Lambda_h(n)=0,n\in C_h}\right)dh=1$$

where we have used the fact that $\sum_{k:\Lambda_h(k)=0,k\in C_h}\frac{\sigma_h^{-2}(k)}{\sum_{n_1\in C_h:\Lambda_h(n_1)=0}\sigma_h^{-2}(n_1)}=1$.

We have now established condition (1) of the theorem. We further readily see that there exists $\underline{\delta}=(-1,-1,...,-1)\underline{\kappa}$ such that $f\left(\underline{\delta}\right)>\underline{\delta}$ across all elements since as $\kappa\to\infty$ we already have seen that $H\left(P,\delta\right)\to 0$ as $\delta\to-\infty$. Additionally, if we set $\bar{\delta}=(1,1,...,1)\,\bar{\kappa}$ where $\bar{\kappa}=A_HL_H$ since the maximum market equity investable in any one asset is equal to total household wealth times the leverage constraint, the function $H\left(\delta\right)$ is bounded above by $\bar{\delta}$. We have therefore proven all three conditions of the theorem (we have actually shown the conditions hold with the upper and lower boundary conditions flipped, which is equivalent since the conditions established here preclude the existence of a fixed point on the boundary of $[\underline{\kappa},\bar{\kappa}]^N$). Thus a unique fixed point $f\left(\delta\right)=\delta$ exists. □

D Data Construction Details

This appendix provides additional details on how the dataset used for estimation is constructed. The construction proceeds in several steps: assembling stock-level characteristics from CRSP and Compustat, computing household holdings as a residual, estimating stock-level return variances, constructing attention vectors, and selecting the estimation sample.

D.1 Stock-Level Data

The primary source of stock data is the CRSP Daily Stock Database, from which I obtain daily prices, returns, shares outstanding, and trading volume for all common equities (CRSP share codes 10 and 11). Book equity is drawn from the CRSP/Compustat Merged Fundamentals Quarterly Database and lagged by one quarter to avoid look-ahead bias. Momentum is computed as the cumulative twelve-month return excluding the most recent month. CAPM betas are estimated from 126 trading days (approximately six months) of daily returns regressed on the market excess

return, with the Fama-French daily factors used to construct market excess returns.

D.2 Household Holdings

Aggregate household ownership of each stock is not directly observed and must be inferred. I follow the residual approach: for each stock n at quarter t, household dollar ownership is

$$\hat{H}_t(n) = \max\{0,\, \mathrm{ME}_t(n) - \mathrm{InstOwn}_t(n) + \mathrm{Shorts}_t(n) - \mathrm{Insider}_t(n)\}$$

where $\mathrm{ME}_t(n)$ is market equity in dollars, $\mathrm{InstOwn}_t(n)$ is aggregate institutional holdings from 13F filings, $\mathrm{Shorts}_t(n)$ is dollar short interest, and $\mathrm{Insider}_t(n)$ denotes insider holdings. Insider holdings are computed as the difference between total market equity (in dollars) and the dollar value of shares available in the public float, both of which are obtained at the stock-quarter level. The resulting household holdings vector is then scaled so that its sum across all stocks equals the fraction of total household assets allocated to directly held common stocks, as reported at a quarterly frequency in the Federal Reserve Economic Data (FRED) database. This scaling ensures that the aggregate household equity allocation in the model matches the macroeconomic data.

D.3 Idiosyncratic Return Variance

For each stock and each quarterly estimation date, I estimate the idiosyncratic return variance from the trailing 126 trading days of log excess returns. Specifically, daily log returns in excess of the risk-free rate are computed for each stock, and the sample variance of these excess returns is calculated over the 126-day window ending on the last trading day of the quarter. This variance is then annualized by multiplying by 252. The resulting stock-level variance vector enters the demand model as the idiosyncratic risk parameter governing how concentrated household demand can be in any individual stock: higher-variance stocks require stronger sentiment to attract the same portfolio weight.

D.4 Attention Vectors

The attention vector is a key input that governs which stocks enter each household's consideration set. Its construction is detailed in Section 4.3; here I describe the computational implementation. Using the regression coefficients from the retail volume prediction model of Graves (2023), I first construct five salience ranking matrices for each trading day. Return salience is computed following Bordalo, Gennaioli and Shleifer (2012) at five horizons: the contemporaneous daily return, the one-day lagged return, the two-day lagged return, the prior week's cumulative return (days $t-3$ through $t-7$), and the prior month's cumulative return (days $t-8$ through $t-28$). On each day and for each return horizon, salience scores are calculated using the Bordalo, Gennaioli, and Shleifer salience function with parameters $\theta = 0.01$ and $\delta = 0.7$ applied to each stock's return relative to the leave-one-out cross-sectional mean. Stocks are then ranked by their salience score,

and these rankings are normalized to lie in $[0, 1]$ with 1 denoting the most salient. Four lagged dollar volume share variables are similarly constructed at the daily, two-day, weekly (five-day average), and monthly (21-day average) horizons. The daily predicted attention for each stock is evaluated as a multiplicative function of these lagged volume shares, salience rankings, and the stock's fractional market capitalization, separately under the buy-side and sell-side coefficient vectors. The buy and sell predictions are averaged and normalized to sum to one across stocks within each day. Daily normalized attention scores are then averaged within each calendar quarter to produce the quarterly attention vector a_t.

D.5 Sample Selection and Variable Construction

The estimation is conducted independently for each of 149 quarters, spanning the period from the twentieth available quarter (to allow for the four-quarter lag in the momentum instrument as well as the lags incorporated in the price instrument) through the end of the sample in 2021Q4. In each quarter, the sample is restricted to common equities (CRSP share codes 10 and 11) for which log market equity, log book equity, momentum, and CAPM beta are all non-missing and finite (in other words, book value and market equity are non-zero and well-defined), and for which the relevant attention-based instruments are positive. From this subset, I select the 300 stocks with the largest imputed household dollar ownership. All observable characteristics except the constant are standardized to zero mean and unit variance within each quarter to improve numerical conditioning for the Knitro Optimizer. The instrument matrix consists of eight instruments: the BLP-style price instrument, log book equity, the year-over-year change in quarterly attention, CAPM beta, a constant, the within-sector sum of the quarterly attention vector, the quarterly attention level, and the within-sector sum of log book equity. Sectors are defined at the four-digit NAICS level.

The mean household consideration set size is set to 10, following the evidence in the household finance literature that individual investors hold highly concentrated portfolios. Attention levels for each stock are capped at 0.5 (no stock can appear in more than half of all consideration sets) and floored at 10^{-4}. A further constraint ensures that no stock's attention probability is so low that the observed household holding would require an implausibly large fraction of attentive households to concentrate their portfolios in that stock. Consideration sets are then drawn by assigning each stock to each simulated household independently with the given attention probability. The estimation uses 300,000 simulated households per quarter.

E Cluster-Based Estimation using Knitro Package

This appendix provides implementation details for the numerical optimization underlying the GEL-MPEC demand estimation described in Section 4.2. Each of the 149 quarterly estimation problems is solved using the Artelys Knitro nonlinear programming solver, accessed through its MATLAB MEX interface via the `knitro_nlp` function.

E.1 Optimization Problem Structure

Recall that each quarterly estimation solves the GEL-MPEC problem of equation (16). The parameter vector is of length $2N + 8 = 608$, partitioned as follows: positions 1–5 contain the five mean belief coefficients $\bar{\beta}$ (on log market equity, log book equity, momentum, CAPM beta, and the constant); positions 6–8 contain the three dispersion parameters (σ_{book}, σ_{mom}, σ_{ϵ}); positions 9–308 contain the $N = 300$ unobservable sentiment residuals $\tilde{\xi}_n \equiv \xi_n - \bar{\beta}' x_n$; and positions 309–608 contain the 300 GEL probability weights π_n. The objective function is the sum of squared GEL weights, $\sum_n \pi_n^2$, which is quadratic and convex.

The solver must satisfy 308 nonlinear equality constraints. Of these, 300 are market clearing conditions requiring that the model-implied demand for each stock equals its observed household ownership share. These constraints are evaluated by the nonlinear constraint function, which, for a given candidate parameter vector, computes the demand for each stock by simulating portfolio allocations across 300,000 households, each with a randomly drawn consideration set and preference shocks. Specifically, for each household the constraint function computes return expectations as the sum of the mean belief component $\bar{\beta}' x_n + \tilde{\xi}_n$, the household-specific deviations arising from the random coefficients on book equity and momentum, and the idiosyncratic preference shock, all filtered through the household's consideration set. Each household then allocates wealth to stocks in its consideration set in proportion to the ratio of expected return (floored at zero) to idiosyncratic variance, and the resulting allocations are averaged across all households to yield the model-implied demand vector. The remaining eight constraints are the GEL moment conditions $\sum_n \pi_n \tilde{\xi}_n z_n^k = 0$ for each of the eight instruments k.

In addition to the nonlinear constraints, the problem includes one linear equality constraint requiring that the GEL weights sum to one: $\sum_n \pi_n = 1$. Box constraints are imposed on all parameters. The five mean coefficients are bounded between -0.2 and 0 for the log market equity coefficient (first position) and between -0.2 and 0.2 for the remaining four; the three dispersion parameters are bounded between a small negative number (-0.01 or -0.02) and 0.5; the sentiment residuals are bounded between -1 and 1; and the GEL weights are bounded between 0 and 1.

E.2 Knitro Solver Configuration

Knitro ("Nonlinear Interior point Trust Region Optimization") is a state-of-the-art commercial solver for large-scale nonlinear programming. It implements several algorithmic strategies, including interior-point (barrier) methods and active-set methods. Knitro is invoked via `knitro_nlp` with its default algorithm selection, which allows the solver to choose the most appropriate algorithm variant at runtime.

Because user-supplied analytic gradients are not provided, Knitro approximates all derivatives via finite differences. The `UseParallel` option is set to `true`, which instructs MATLAB's Parallel Computing Toolbox to distribute the finite-difference perturbations across available cores. Each perturbation requires a full evaluation of the nonlinear constraint function, which in turn simulates

the portfolio choices of 300,000 households, making parallelization essential for tractability. The output verbosity level is set to 4 (`outlev=4`) to facilitate monitoring of convergence progress.

E.3 Cluster Computing Environment

Each quarterly estimation is executed on the Harvard Faculty of Arts and Sciences Research Computing (FASRC) cluster. The jobs are configured to utilize 112 CPU cores, which are made available to MATLAB's Parallel Computing Toolbox via `parpool`. At the start of each estimation run, any existing parallel pool is deleted and a fresh pool is opened with the number of workers set to the number of available cores detected by MATLAB (`feature('numcores')`). This ensures that all 112 cores are active during the finite-difference constraint evaluations that dominate the computational cost. A single quarterly estimation requires approximately 100 CPU hours, and across all 149 quarters and all three model specifications (the attention model, the market-capitalization-based specification, and the alternative-characteristics specification), the total computational burden is approximately 45,000 CPU hours.

E.4 Starting Values and Initialization via NFXP

Effective starting values are critical for convergence in problems with highly nonlinear constraints. The initial guess for the parameter vector is constructed as follows. The five mean belief coefficients are initialized to $(-0.1,\ 0.1,\ 0.02,\ -0.03,\ -0.1)$, reflecting economically motivated priors: a negative coefficient on log market equity (smaller stocks have higher expected returns), a positive coefficient on log book equity (value stocks), a weakly positive coefficient on momentum, a small negative loading on beta, and a negative constant. The three dispersion parameters are initialized to $(0.05,\ 0.05,\ 0.05)$, reflecting modest heterogeneity across households. The 300 GEL probability weights are initialized uniformly at $1/N = 1/300$.

The sentiment residuals $\tilde{\xi}_n$ cannot be initialized to an arbitrary value because they must be approximately consistent with the market clearing conditions for the solver to make progress. Accordingly, before the main Knitro optimization begins, I solve the nested fixed point (NFXP) problem of Proposition 2 at the initial parameter guess. The NFXP algorithm iteratively updates the vector $\delta = \bar{\beta}'x + \tilde{\xi}$ using a gradient-based step that adjusts each stock's utility index in proportion to the gap between the variance-adjusted observed share and the model-implied demand, scaled by the marginal effect of utility on demand. The iteration continues until the maximum absolute change in δ falls below 10^{-10} or a maximum of 10,000 iterations is reached. The converged NFXP solution is then decomposed as $\tilde{\xi}_n = \delta_n - \bar{\beta}'x_n$ to obtain the initial sentiment residuals. This procedure ensures that the starting point approximately satisfies the 300 market clearing constraints, placing the initial iterate in the feasible region and dramatically improving convergence.

E.5 Non-Convergence and Reoptimization

Despite the NFXP-based initialization, the initial Knitro solve may fail to converge to a feasible local optimum for some quarters. Knitro signals non-convergence via a nonzero exit flag (`exitflag` $\neq 0$), which may indicate that the algorithm terminated at an infeasible point, exceeded its iteration limit, or encountered numerical difficulties. When this occurs, the procedure invokes a reoptimization step. The returned (possibly infeasible) parameter values from the first Knitro run are taken as the new candidate for $\bar{\beta}$ and σ. The NFXP contraction mapping is then re-solved at these updated parameter values, with the initial delta vector set to $\bar{\beta}'x + 1.4$ (a shifted starting point), yielding a new sentiment residual vector that is consistent with market clearing under the updated coefficients. This provides a substantially improved starting point for the second Knitro run because the mean coefficients and dispersion parameters from the first attempt, even if the full optimization did not converge, are typically much closer to the solution than the original guesses.

The second Knitro run is then launched with the same objective, constraints, bounds, and solver options as the first. Both sets of results (from the initial and reoptimization runs) are saved to disk, enabling post-hoc comparison and selection. If the first run converges (exit flag equal to zero), only its results are retained and the reoptimization step is skipped. This two-stage procedure is applied independently to each of the 149 quarters, and the estimation loop iterates from the 20th available quarter through the 168th (corresponding to the sample period 1984–2021, accounting for the required four-quarter lag for the momentum attention instrument).

F No Attention and Alternative Specification Sentiment Autocorrelation Analysis

Figures 5 and 6 report the autocorrelation profiles of estimated household sentiment under the no-attention and alternative demand specifications, respectively. These figures are the analogs of Figure 2, which reports the autocorrelation profile under the main specification. Under the no-attention specification, in which market-capitalization weights replace the estimated attention model, household sentiment continues to display persistence at short lags, with a pattern of gradual decay qualitatively similar to the main specification. Under the alternative specification, which replaces the separate log market equity and log book equity characteristics with a single market equity-to-book equity ratio and adds a third dispersion parameter, the persistence profile is likewise broadly consistent with the main results. The robustness of the autocorrelation pattern across specifications suggests that the persistence of household sentiment is not an artifact of the particular attention model or characteristic set employed in the main estimation.

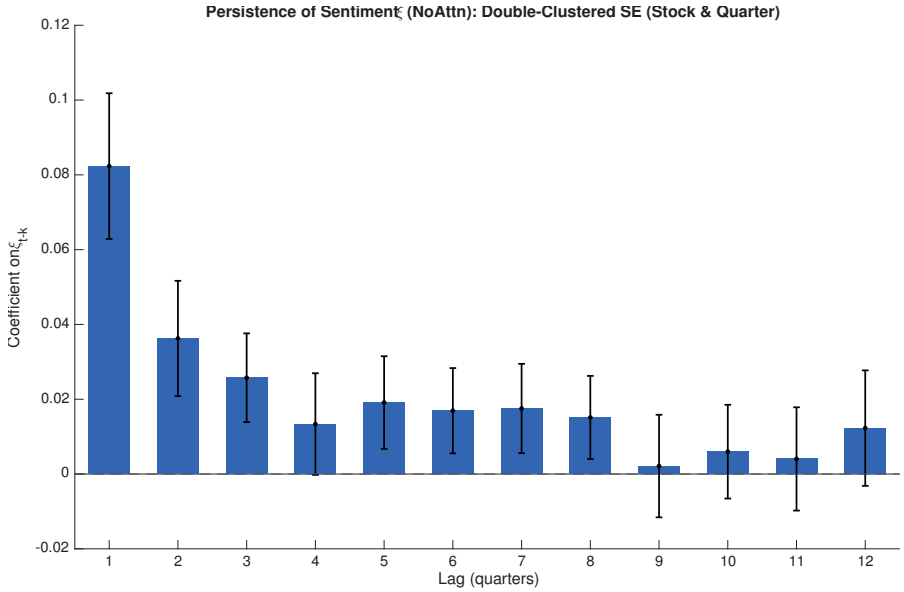


Figure 5: Persistence of Household Sentiment: Autocorrelation Profile (No Attention Specification)
Note: This figure plots the estimated autocorrelation coefficients of household sentiment $\xi_t(n)$ at lags of 1 through 12 quarters under the no-attention specification, in which market-capitalization weights replace the estimated attention model. Each bar represents the slope coefficient from a pooled OLS regression of $\xi_t(n)$ on $\xi_{t-k}(n)$ across all stocks and quarters with nonzero sentiment values. Error bars depict 95% confidence intervals constructed from two-way clustered standard errors (by stock and by quarter) following Petersen (2009). Sentiment is estimated quarterly via GEL-MPEC using Knitro optimization on the top 300 U.S. stocks by market capitalization from 1984Q4 through 2021Q4.

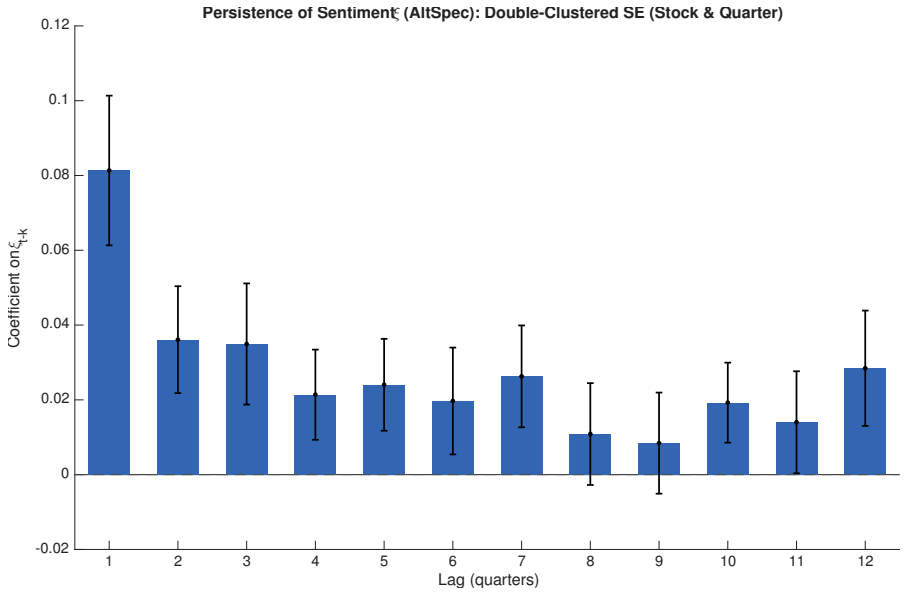
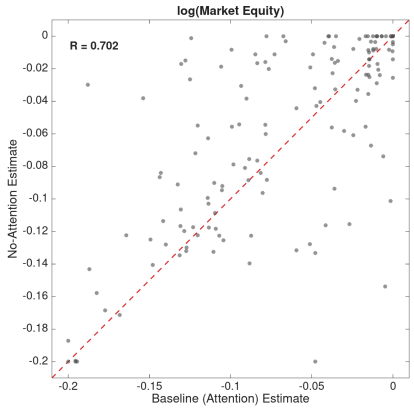


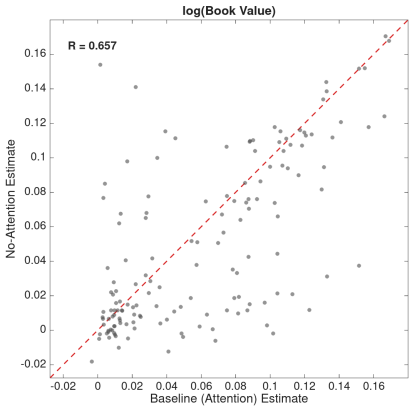
Figure 6: Persistence of Household Sentiment: Autocorrelation Profile (Alternative Specification)
Note: This figure plots the estimated autocorrelation coefficients of household sentiment $\xi_t(n)$ at lags of 1 through 12 quarters under the alternative specification, which uses the market equity-to-book equity ratio, momentum, CAPM beta, and a constant as stock characteristics, with three dispersion parameters. Each bar represents the slope coefficient from a pooled OLS regression of $\xi_t(n)$ on $\xi_{t-k}(n)$ across all stocks and quarters with nonzero sentiment values. Error bars depict 95% confidence intervals constructed from two-way clustered standard errors (by stock and by quarter) following Petersen (2009). Sentiment is estimated quarterly via GEL-MPEC using Knitro optimization on the top 300 U.S. stocks by market capitalization from 1984Q4 through 2021Q4.

G Robustness: Additional Analysis (Main Specification Comparison)

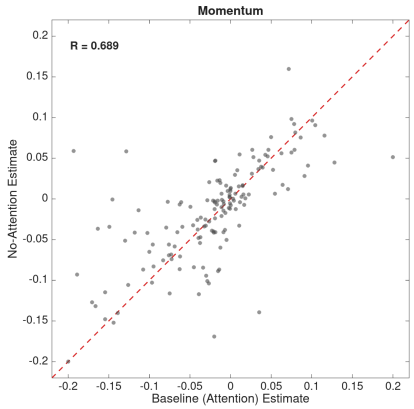
This section features an additional robustness check on the estimation process. The baseline specification in the paper uses the attention model of subsection 4.3. I additionally run analysis using attention that is proportional to market capitalization. Figure 7 details the relationship between the estimates. The bunching at boundaries is due to parameter restrictions used in optimization.



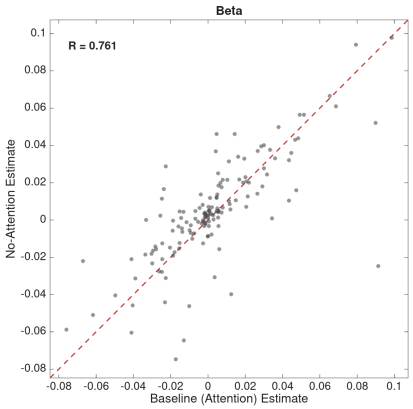
(a) Log Market Equity



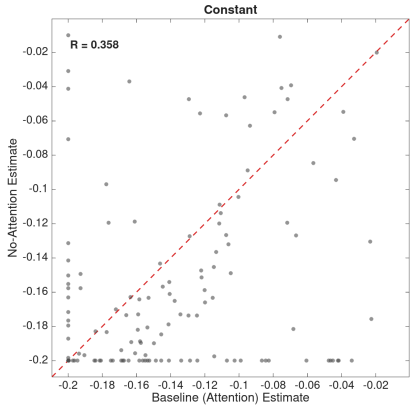
(b) Log Book Value



(c) Momentum



(d) Beta



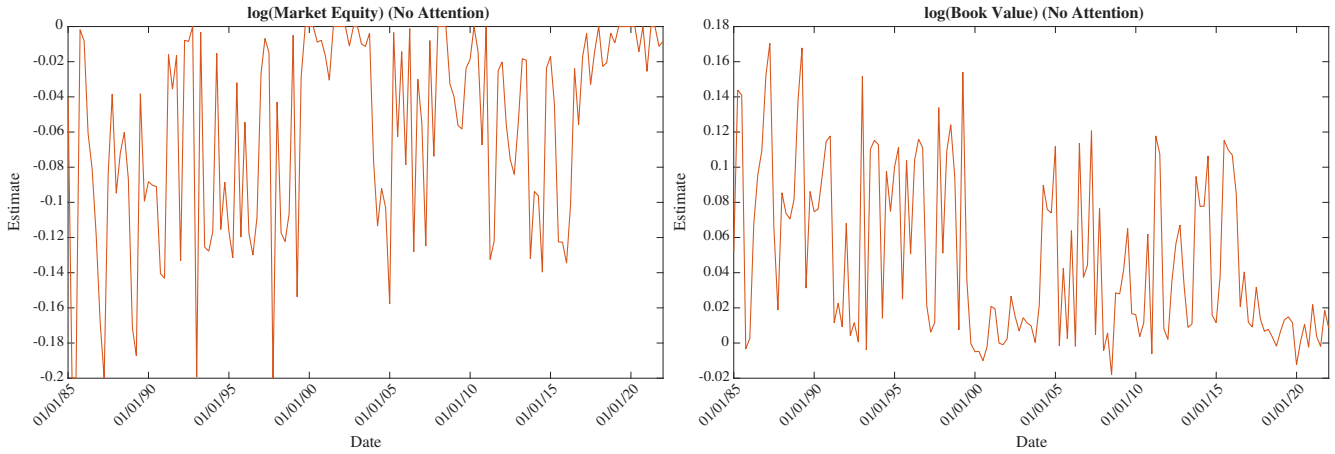
(e) Constant

Figure 7: Scatter Plots of Parameter Estimates, Differing Consideration Set Probabilities

Note: Each subfigure is a scatter plot comparing quarterly household demand parameter estimates from the attention model (horizontal axis) to those from the no-attention model (vertical axis), from 1984 to 2021. The red dashed line is the 45-degree reference. The correlation coefficient R is reported for each parameter. Parameters are estimated via GEL-MPEC using Knitro optimization on the top 300 U.S. stocks by market capitalization.

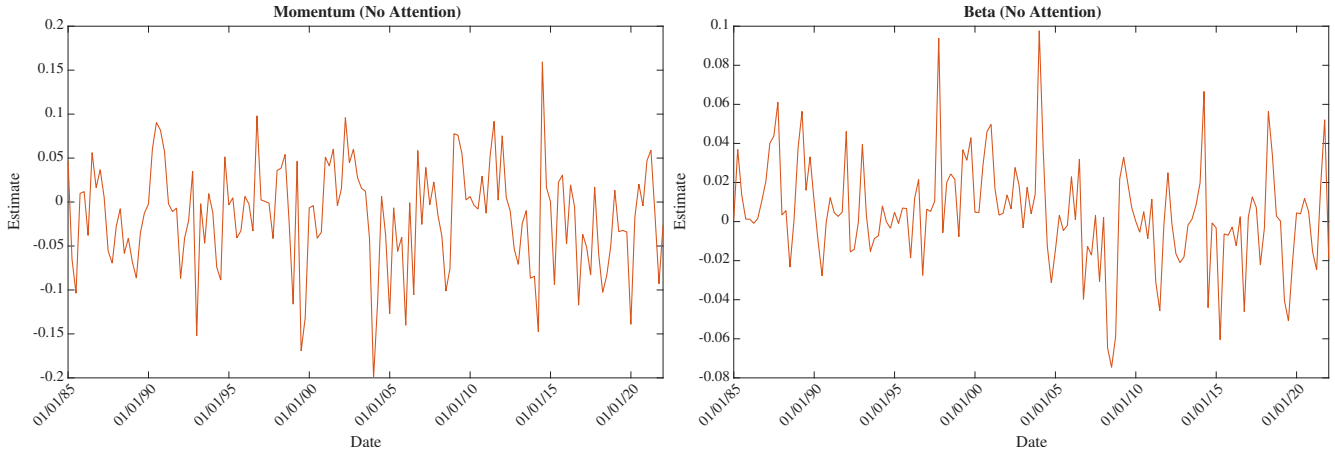
H Robustness: Using Market Caps instead of Attention Model

In this Appendix I provide estimates and results for the model when using market cap-based attention weights instead of the attention model used in the paper's main specification. See subsection 4.6 for more details on the methodology.



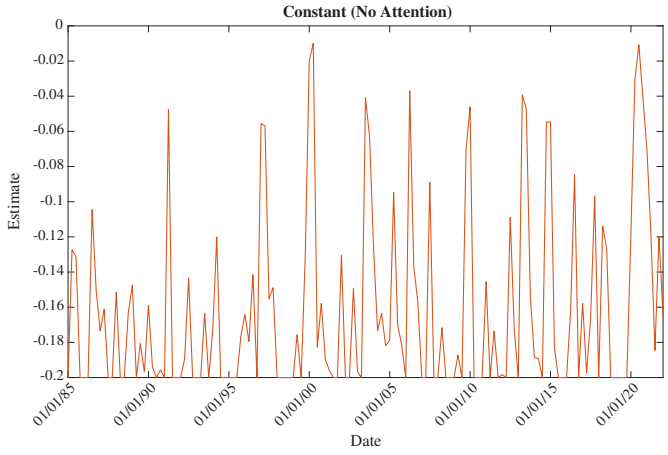
(a) Log Market Equity

(b) Log Book Value



(c) Momentum

(d) Beta



(e) Constant

Figure 8: Parameter Estimates without Attention: Cross-Sectional Distributions

Note: Each subfigure plots the quarterly time series of an estimated household demand parameter from the no-attention model, from 1984 to 2021. The five parameters shown are the coefficients on log market equity, log book value, momentum, beta, and a constant in the household portfolio choice model. Parameters are estimated quarterly via GEL-MPEC using Knitro optimization on the top 300 U.S. stocks by market capitalization, with uniform consideration set probabilities (no estimated attention vectors).

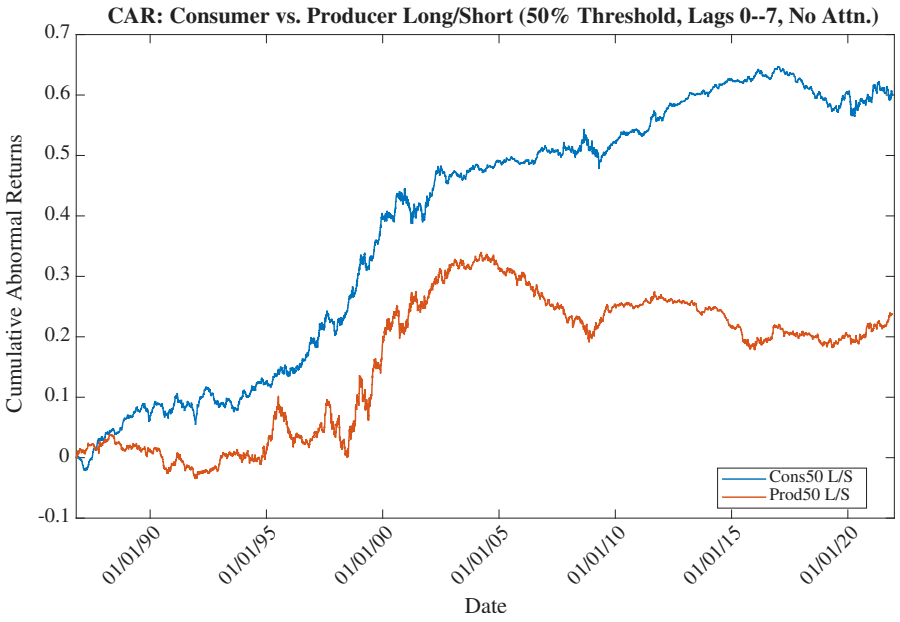


Figure 9: Cumulative Abnormal Returns, No Attention: Consumer and Producer Long/Short Spreads (50% Threshold, Lags 0–7)

Note: This figure plots cumulative abnormal returns (CARs) for the consumer and producer long/short spread strategies using the no-attention specification. The consumer (producer) long/short strategy goes long the sentiment-weighted consumer (producer) portfolio and short the corresponding market-capitalization-weighted benchmark portfolio, using the 50% NAICS-based consumer spending share threshold. Consumer stocks are those with consumer share at or above 50%, and producer stocks are those below 50%. CARs are computed from daily residuals of a six-factor regression on the Fama-French five factors (MktRF, SMB, HML, RMW, CMA) plus momentum (UMD), with returns averaged over lags 0 through 7 (eight quarters). Portfolios are formed from the top 300 U.S. stocks by market capitalization and rebalanced quarterly. The six-factor regression uses Newey-West (HAC) standard errors with a Bartlett kernel and automatic bandwidth selection following Newey and West (1994). This figure uses the no-attention specification, which assigns uniform consideration set probabilities rather than estimating attention vectors. The market-capitalization-weighted benchmark portfolios are shared with the main specification, as they depend only on market equity weights.

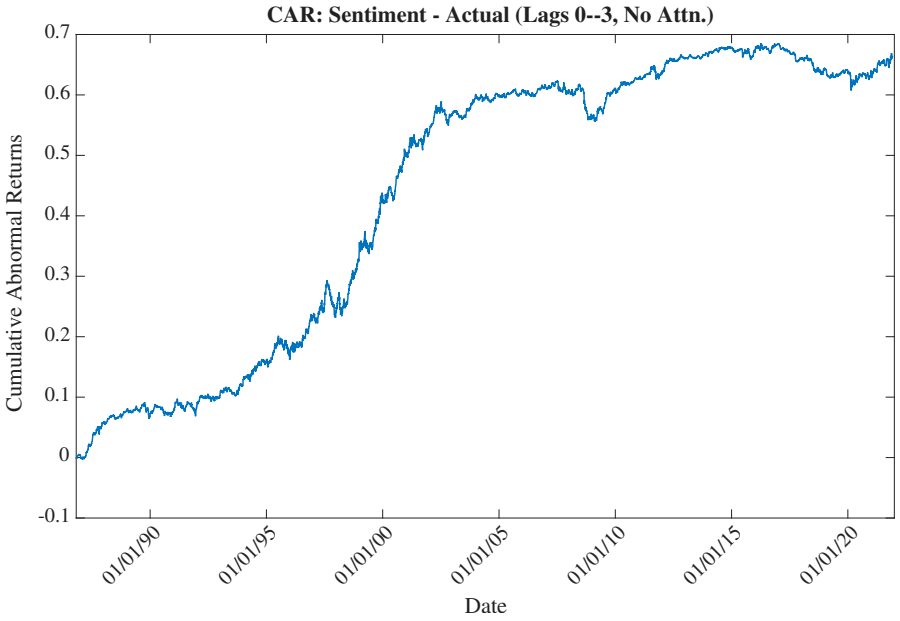


Figure 10: Cumulative Abnormal Returns, No Attention: Sentiment minus Actual (Lags 0–3)

Note: This figure plots cumulative abnormal returns (CARs) for a long-short strategy that goes long the Sentiment portfolio and short the Actual Holdings portfolio, using the no-attention specification. The Sentiment portfolio weights stocks by estimated household sentiment parameters, while the Actual Holdings portfolio weights stocks by observed household equity positions. CARs are computed from daily residuals of a six-factor regression on the Fama-French five factors (MktRF, SMB, HML, RMW, CMA) plus momentum (UMD), with returns averaged over lags 0 through 3 (four quarters). Portfolios are formed from the top 300 U.S. stocks by market capitalization and rebalanced quarterly. The six-factor regression uses Newey-West (HAC) standard errors with a Bartlett kernel and automatic bandwidth selection following Newey and West (1994). This figure uses the no-attention specification, which assigns uniform consideration set probabilities rather than estimating attention vectors.

I Robustness: Alternative Specifications

In this Appendix I provide estimates and results for the model when using an alternative specification with a different IV and ME/BE (the market-to-book ratio) as one of the belief components in the linear specification instead of $\log(ME)$ and $\log(BE)$ separately. See subsection 4.6 for more details on the methodology.

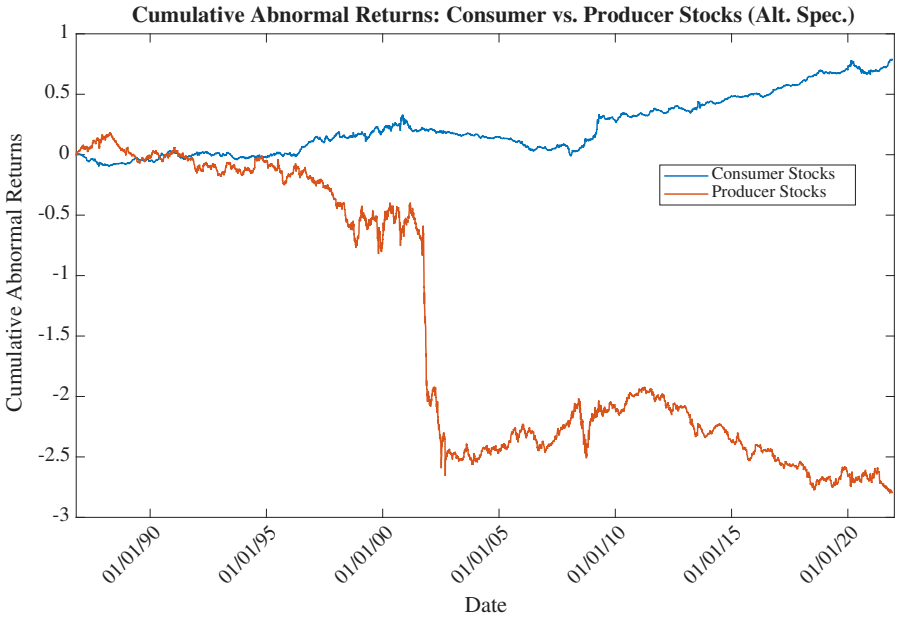


Figure 11: Cumulative Abnormal Returns, Alternative Specification: Consumer vs. Producer

Note: This figure plots cumulative abnormal returns (CARs) for consumer-facing and producer-facing stock portfolios, using the 10% NAICS-based consumer share threshold. Consumer stocks are those with consumer share at or above 50%, and producer stocks are those below 50%. CARs are computed from daily residuals of a six-factor regression on the Fama-French five factors (MktRF, SMB, HML, RMW, CMA) plus momentum (UMD), evaluated at the contemporaneous lag (lag 0). Portfolios are formed from the top 300 U.S. stocks by market capitalization and rebalanced quarterly. The six-factor regression uses Newey-West (HAC) standard errors with a Bartlett kernel and automatic bandwidth selection following Newey and West (1994). This figure uses the alternative specification, which estimates household demand using four characteristics (ME/BE ratio, momentum, beta, and a constant) instead of the baseline five characteristics (log market equity, log book value, momentum, beta, and a constant).

Table 8: Core Models, Alternative Specification: Contemporaneous Lag

	Actual	Sentiment	Char.
A. Raw excess returns			
Average	9.74	10.54	9.66
	(3.68)	(3.71)	(3.26)
T	8867	8867	8867
SR	0.5363	0.5580	0.4873
B. CAPM regressions			
Intercept	0.42	0.98	-0.13
	(0.83)	(1.37)	(-0.13)
MktRF	0.99	1.01	1.04
	(247.09)	(121.01)	(80.50)
R^2	0.9756	0.9480	0.9022
IR	0.1465	0.2278	-0.0207
C. Six-factor regressions			
Intercept	0.01	1.53	0.71
	(0.03)	(2.45)	(0.79)
MktRF	1.00	1.00	1.01
	(476.97)	(144.53)	(112.05)
SMB	-0.19	-0.16	-0.18
	(-55.13)	(-18.44)	(-13.90)
HML	-0.01	-0.06	-0.04
	(-1.33)	(-5.42)	(-2.03)
RMW	0.09	0.05	-0.11
	(14.63)	(2.91)	(-2.27)
CMA	0.03	-0.11	-0.09
	(3.84)	(-5.47)	(-2.00)
UMD	-0.01	-0.03	0.04
	(-1.81)	(-3.17)	(2.73)
R^2	0.9894	0.9600	0.9139
IR	0.0054	0.4060	0.1217

Note: This table presents portfolio performance for three core model strategies: Actual Holdings (market-cap weighted household direct stock holdings), Sentiment (belief-based model with all parameters), and Characteristics (characteristics-only model with sentiment parameters set to zero). Portfolios are formed from the top 300 U.S. stocks by market capitalization and rebalanced quarterly. Returns are evaluated at the contemporaneous lag (lag 0). Panel A reports annualized mean excess returns, t-statistics, the number of daily observations, and annualized Sharpe ratios. Panel B reports CAPM alpha, market beta, R-squared, and the information ratio. Panel C reports six-factor alpha and factor loadings from a regression on the Fama-French five factors (MktRF, SMB, HML, RMW, CMA) plus momentum (UMD), along with R-squared and the information ratio. All t-statistics are computed using Newey-West (HAC) standard errors with a Bartlett kernel and automatic bandwidth selection following Newey and West (1994). This table uses the alternative specification, which estimates household demand using four characteristics (ME/BE ratio, momentum, beta, and a constant) instead of the baseline five characteristics (log market equity, log book value, momentum, beta, and a constant).

Table 9: Core Models, Alternative Specification: Lags 0 to 7

	Actual	Sentiment	Char.
A. Raw excess returns			
Average	9.97	10.71	10.04
	(3.68)	(3.77)	(3.48)
T	8867	8867	8867
SR	0.5395	0.5661	0.5262
B. CAPM regressions			
Intercept	0.55	1.07	0.31
	(0.86)	(1.61)	(0.49)
MktRF	1.00	1.02	1.03
	(150.40)	(164.95)	(154.59)
R^2	0.9608	0.9606	0.9631
IR	0.1517	0.2854	0.0839
C. Six-factor regressions			
Intercept	0.78	2.00	1.81
	(1.89)	(4.05)	(3.66)
MktRF	1.00	1.01	1.00
	(413.33)	(279.70)	(255.94)
SMB	-0.20	-0.18	-0.18
	(-38.85)	(-13.66)	(-25.52)
HML	0.02	-0.07	0.02
	(2.28)	(-9.41)	(2.50)
RMW	0.08	0.07	-0.06
	(6.77)	(5.29)	(-4.76)
CMA	0.05	-0.04	-0.07
	(4.25)	(-2.68)	(-4.02)
UMD	-0.12	-0.13	-0.12
	(-17.83)	(-17.97)	(-14.43)
R^2	0.9838	0.9786	0.9787
IR	0.3299	0.7213	0.6500

Note: This table presents portfolio performance for three core model strategies: Actual Holdings (market-cap weighted household direct stock holdings), Sentiment (belief-based model with all parameters), and Characteristics (characteristics-only model with sentiment parameters set to zero). Portfolios are formed from the top 300 U.S. stocks by market capitalization and rebalanced quarterly. Returns are averaged over lags 0 through 7 (eight quarters). Panel A reports annualized mean excess returns, t-statistics, the number of daily observations, and annualized Sharpe ratios. Panel B reports CAPM alpha, market beta, R-squared, and the information ratio. Panel C reports six-factor alpha and factor loadings from a regression on the Fama-French five factors (MktRF, SMB, HML, RMW, CMA) plus momentum (UMD), along with R-squared and the information ratio. All t-statistics are computed using Newey-West (HAC) standard errors with a Bartlett kernel and automatic bandwidth selection following Newey and West (1994). This table uses the alternative specification, which estimates household demand using four characteristics (ME/BE ratio, momentum, beta, and a constant) instead of the baseline five characteristics (log market equity, log book value, momentum, beta, and a constant).

Table 10: Long-Short Spreads, Alternative Specification: Contemporaneous Lag

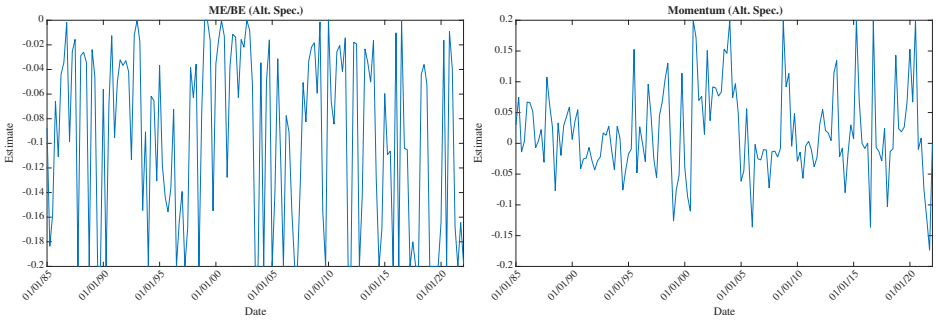
	Panel A: Long Sent. (ND)			*Panel B: Long Sentiment*		
	− Char. (ND)	− Actual	− Char. (EA)	− Char. (ND)	− Actual	− Char. (EA)
A. Raw excess returns						
Average	2.42	1.00	1.87	2.22	0.80	1.67
	(1.38)	(1.27)	(1.23)	(1.35)	(1.25)	(1.17)
T	8867	8867	8867	8867	8867	8867
SR	0.2220	0.2168	0.2043	0.2154	0.2109	0.1935
B. CAPM regressions						
Intercept	2.18	0.84	0.51	1.91	0.57	0.24
	(1.22)	(1.08)	(0.34)	(1.14)	(0.90)	(0.17)
MktRF	0.03	0.02	0.14	0.03	0.03	0.15
	(1.11)	(1.92)	(6.76)	(1.52)	(3.29)	(7.26)
R^2	0.0018	0.0043	0.0817	0.0035	0.0145	0.1024
IR	0.2004	0.1830	0.0581	0.1849	0.1494	0.0287
C. Six-factor regressions						
Intercept	2.33	1.38	2.20	2.48	1.52	2.35
	(1.32)	(1.77)	(1.53)	(1.52)	(2.56)	(1.79)
MktRF	0.02	0.00	0.10	0.02	0.00	0.10
	(1.13)	(0.51)	(5.48)	(1.09)	(0.33)	(5.68)
SMB	0.04	0.03	-0.10	0.04	0.03	-0.10
	(1.76)	(3.09)	(-4.41)	(1.86)	(4.10)	(-4.76)
HML	-0.10	-0.07	-0.18	-0.08	-0.05	-0.16
	(-2.59)	(-4.82)	(-6.12)	(-2.31)	(-5.33)	(-6.09)
RMW	0.12	0.02	0.01	0.07	-0.03	-0.04
	(1.66)	(1.04)	(0.28)	(0.96)	(-1.88)	(-0.74)
CMA	-0.10	-0.10	-0.27	-0.15	-0.14	-0.32
	(-1.39)	(-4.14)	(-5.43)	(-2.05)	(-7.41)	(-6.63)
UMD	-0.05	-0.03	-0.06	-0.05	-0.03	-0.06
	(-1.39)	(-1.86)	(-2.07)	(-1.49)	(-2.60)	(-2.29)
R^2	0.0192	0.0612	0.2013	0.0230	0.1306	0.2450
IR	0.2160	0.3084	0.2689	0.2428	0.4286	0.3120

Note: This table presents portfolio performance for long-short spread strategies at the contemporaneous lag (lag 0). The strategies take a long position in Sentiment (No Disagreement) or the full Sentiment model and a short position in one of three alternatives: Characteristics (No Disagreement), Actual Holdings, or Characteristics (Even Attention). Portfolios are formed from the top 300 U.S. stocks by market capitalization and rebalanced quarterly. Panel A reports annualized mean excess returns, t-statistics, the number of daily observations, and annualized Sharpe ratios. Panel B reports CAPM alpha, market beta, R-squared, and the information ratio. Panel C reports six-factor alpha and factor loadings from a regression on the Fama-French five factors (MktRF, SMB, HML, RMW, CMA) plus momentum (UMD), along with R-squared and the information ratio. All t-statistics are computed using Newey-West (HAC) standard errors with a Bartlett kernel and automatic bandwidth selection following Newey and West (1994). This table uses the alternative specification, which estimates household demand using four characteristics (ME/BE ratio, momentum, beta, and a constant) instead of the baseline five characteristics (log market equity, log book value, momentum, beta, and a constant).

Table 11: Long-Short Spreads, Alternative Specification: Lags 0 to 3

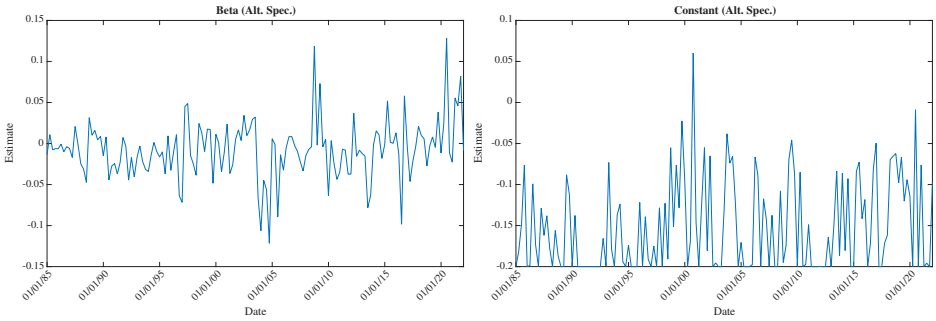
	Panel A: Long Sent. (ND)			*Panel B: Long Sentiment*		
	− Char. (ND)	− Actual	− Char. (EA)	− Char. (ND)	− Actual	− Char. (EA)
A. Raw excess returns						
Average	2.23	1.37	2.26	1.98	1.12	2.01
	(2.19)	(2.66)	(1.92)	(2.14)	(2.47)	(1.77)
T	8867	8867	8867	8867	8867	8867
SR	0.3782	0.4537	0.3219	0.3717	0.4267	0.2983
B. CAPM regressions						
Intercept	1.96	1.28	0.99	1.58	0.90	0.61
	(1.90)	(2.48)	(0.85)	(1.69)	(2.02)	(0.55)
MktRF	0.03	0.01	0.13	0.04	0.02	0.15
	(2.60)	(1.64)	(8.97)	(4.27)	(4.79)	(10.06)
R^2	0.0077	0.0031	0.1209	0.0213	0.0271	0.1607
IR	0.3340	0.4253	0.1506	0.2989	0.3458	0.0982
C. Six-factor regressions						
Intercept	1.82	1.65	2.61	1.88	1.71	2.68
	(1.86)	(3.56)	(2.63)	(2.10)	(4.41)	(2.88)
MktRF	0.03	0.00	0.10	0.03	0.00	0.10
	(3.33)	(0.33)	(8.59)	(3.90)	(1.25)	(9.04)
SMB	0.02	0.02	-0.12	0.02	0.03	-0.12
	(1.55)	(1.79)	(-4.54)	(1.93)	(2.31)	(-4.98)
HML	-0.17	-0.12	-0.24	-0.13	-0.08	-0.21
	(-8.40)	(-11.98)	(-14.38)	(-7.68)	(-11.13)	(-12.69)
RMW	0.12	0.03	0.00	0.07	-0.03	-0.05
	(6.03)	(3.16)	(0.24)	(3.48)	(-3.51)	(-2.97)
CMA	-0.06	-0.07	-0.21	-0.11	-0.12	-0.26
	(-2.02)	(-4.71)	(-7.79)	(-4.04)	(-10.30)	(-10.29)
UMD	-0.01	-0.01	-0.05	-0.01	-0.02	-0.06
	(-0.49)	(-2.02)	(-3.61)	(-0.84)	(-3.40)	(-4.26)
R^2	0.1210	0.2183	0.3774	0.1240	0.2848	0.4230
IR	0.3287	0.6168	0.4708	0.3772	0.7702	0.5219

Note: This table presents portfolio performance for long-short spread strategies with returns averaged over lags 0 through 3 (four quarters). The strategies take a long position in Sentiment (No Disagreement) or the full Sentiment model and a short position in one of three alternatives: Characteristics (No Disagreement), Actual Holdings, or Characteristics (Even Attention). Portfolios are formed from the top 300 U.S. stocks by market capitalization and rebalanced quarterly. Panel A reports annualized mean excess returns, t-statistics, the number of daily observations, and annualized Sharpe ratios. Panel B reports CAPM alpha, market beta, R-squared, and the information ratio. Panel C reports six-factor alpha and factor loadings from a regression on the Fama-French five factors (MktRF, SMB, HML, RMW, CMA) plus momentum (UMD), along with R-squared and the information ratio. All t-statistics are computed using Newey-West (HAC) standard errors with a Bartlett kernel and automatic bandwidth selection following Newey and West (1994). This table uses the alternative specification, which estimates household demand using four characteristics (ME/BE ratio, momentum, beta, and a constant) instead of the baseline five characteristics (log market equity, log book value, momentum, beta, and a constant).



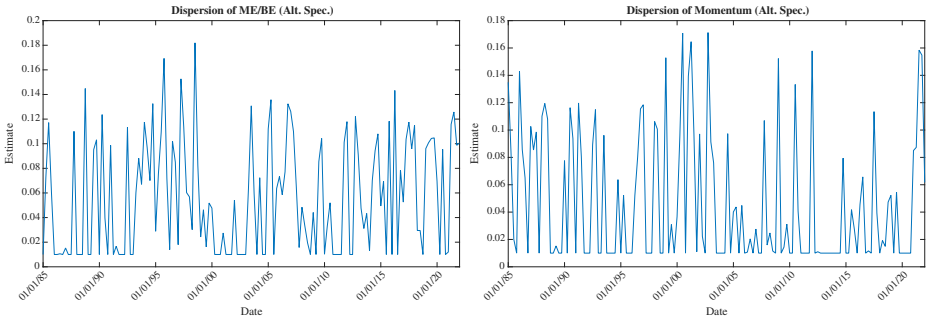
(a) ME/BE

(b) Momentum



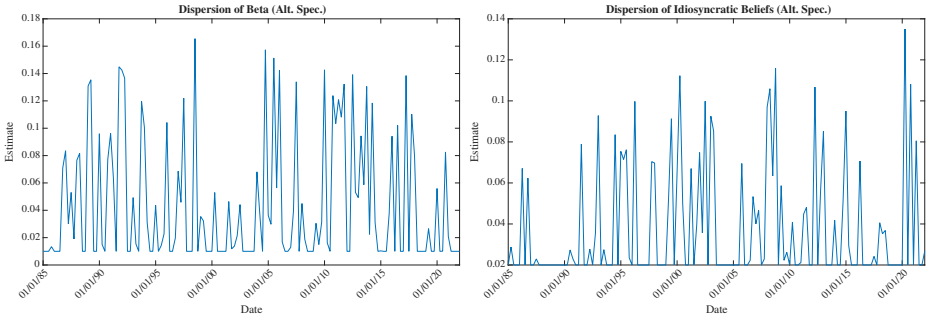
(c) Beta

(d) Constant



(e) Dispersion of ME/BE

(f) Dispersion of Momentum



(g) Dispersion of Beta

(h) Dispersion of Idiosyncratic

Figure 12: Parameter Estimates, Alternative Specification: Cross-Sectional Distributions

Note: Each subfigure plots the quarterly time series of an estimated household demand parameter from the alternative specification model, from 1984 to 2021. The four mean parameters shown are ME/BE ratio, momentum, beta, and a constant. The four dispersion parameters shown are the dispersions of ME/BE ratio, momentum, beta, and an idiosyncratic component. Parameters are estimated quarterly via GEL-MPEC using Knitro optimization on the top 300 U.S. stocks by market capitalization. This figure uses the alternative specification, which estimates household demand using four characteristics (ME/BE ratio, momentum, beta, and a constant) instead of the baseline five characteristics (log market equity, log book value, momentum, beta, and a constant).

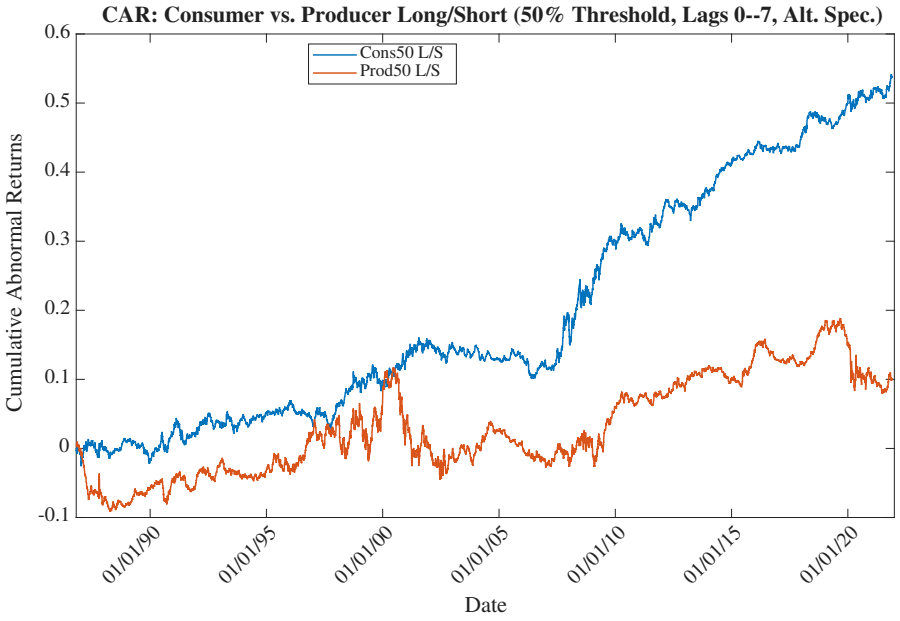


Figure 13: Cumulative Abnormal Returns, Alternative Specification: Consumer and Producer Long/Short Spreads (50% Threshold, Lags 0–7)

Note: This figure plots cumulative abnormal returns (CARs) for the consumer and producer long/short spread strategies using the alternative specification. The consumer (producer) long/short strategy goes long the sentiment-weighted consumer (producer) portfolio and short the corresponding market-capitalization-weighted benchmark portfolio, using the 50% NAICS-based consumer spending share threshold. Consumer stocks are those with consumer share at or above 50%, and producer stocks are those below 50%. CARs are computed from daily residuals of a six-factor regression on the Fama-French five factors (MktRF, SMB, HML, RMW, CMA) plus momentum (UMD), with returns averaged over lags 0 through 7 (eight quarters). Portfolios are formed from the top 300 U.S. stocks by market capitalization and rebalanced quarterly. The six-factor regression uses Newey-West (HAC) standard errors with a Bartlett kernel and automatic bandwidth selection following Newey and West (1994). This figure uses the alternative specification, which estimates household demand using four characteristics (ME/BE ratio, momentum, beta, and a constant) instead of the baseline five characteristics (log market equity, log book value, momentum, beta, and a constant). The market-capitalization-weighted benchmark portfolios are shared with the main specification, as they depend only on market equity weights.

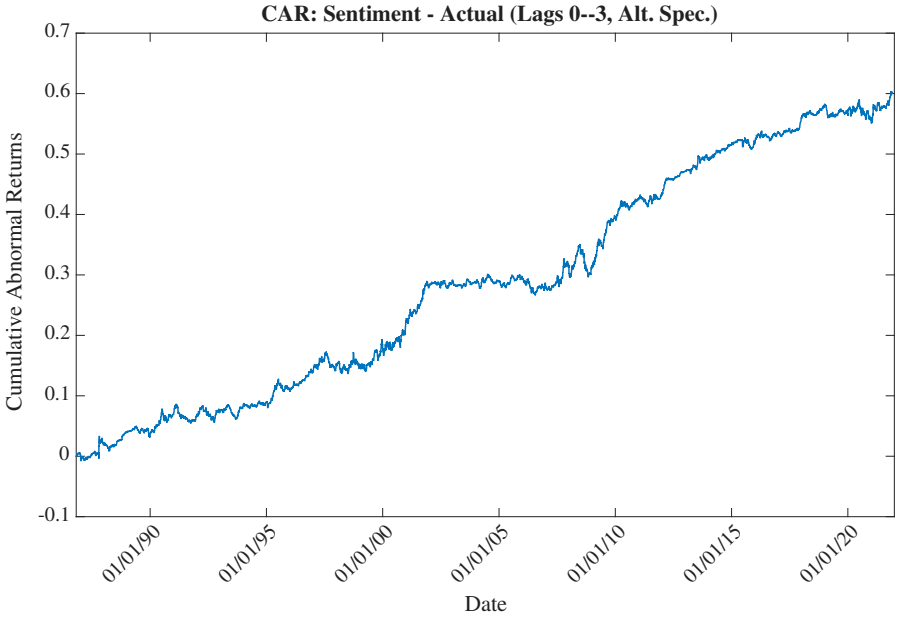


Figure 14: Cumulative Abnormal Returns, Alternative Specification: Sentiment minus Actual (Lags 0–3)

Note: This figure plots cumulative abnormal returns (CARs) for a long-short strategy that goes long the Sentiment portfolio and short the Actual Holdings portfolio, using the alternative specification. The Sentiment portfolio weights stocks by estimated household sentiment parameters, while the Actual Holdings portfolio weights stocks by observed household equity positions. CARs are computed from daily residuals of a six-factor regression on the Fama-French five factors (MktRF, SMB, HML, RMW, CMA) plus momentum (UMD), with returns averaged over lags 0 through 3 (four quarters). Portfolios are formed from the top 300 U.S. stocks by market capitalization and rebalanced quarterly. The six-factor regression uses Newey-West (HAC) standard errors with a Bartlett kernel and automatic bandwidth selection following Newey and West (1994). This figure uses the alternative specification, which estimates household demand using four characteristics (ME/BE ratio, momentum, beta, and a constant) instead of the baseline five characteristics (log market equity, log book value, momentum, beta, and a constant).

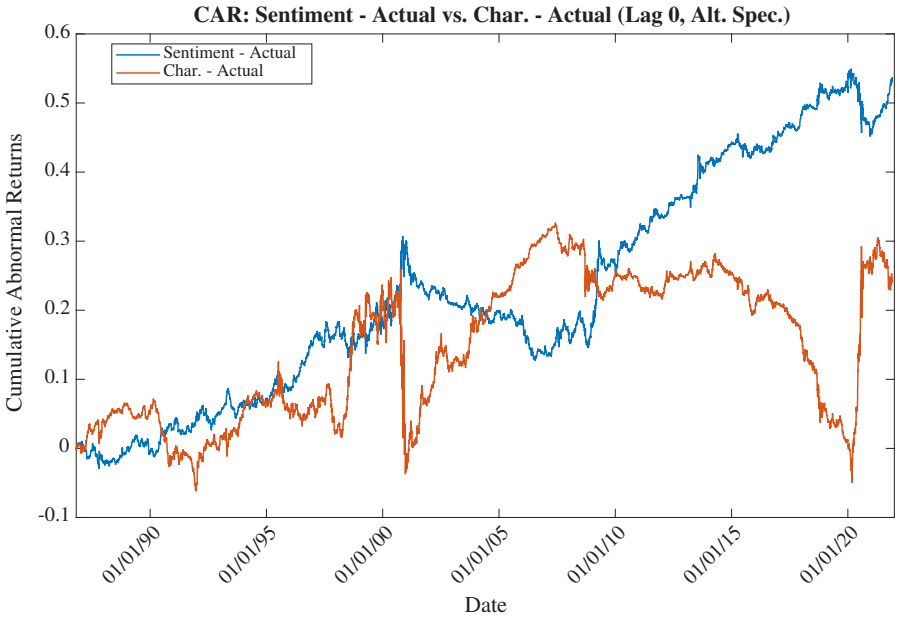


Figure 15: Cumulative Abnormal Returns, Alternative Specification: Sentiment minus Actual and Characteristics minus Actual (Lag 0)

Note: This figure jointly plots cumulative abnormal returns (CARs) for two long-short strategies using the alternative specification. The first strategy goes long the Sentiment portfolio and short the Actual Holdings portfolio; the second strategy goes long the Characteristics portfolio and short the Actual Holdings portfolio. The Sentiment portfolio weights stocks by estimated household sentiment parameters, the Characteristics portfolio weights stocks by observable firm characteristics only, and the Actual Holdings portfolio weights stocks by observed household equity positions. CARs are computed from daily residuals of a six-factor regression on the Fama-French five factors (MktRF, SMB, HML, RMW, CMA) plus momentum (UMD), evaluated at the contemporaneous lag (lag 0). Portfolios are formed from the top 300 U.S. stocks by market capitalization and rebalanced quarterly. The six-factor regression uses Newey-West (HAC) standard errors with a Bartlett kernel and automatic bandwidth selection following Newey and West (1994). This figure uses the alternative specification, which estimates household demand using four characteristics (ME/BE ratio, momentum, beta, and a constant) instead of the baseline five characteristics (log market equity, log book value, momentum, beta, and a constant).

J Classification of Companies by Industry into B2B vs B2C

To classify firms as consumer-facing or producer-facing (in other words, essentially more "B2B" or "B2C" in corporate parlance), I assign each 4-digit NAICS industry a consumer-experience score on a scale from 0 to 100. A score of 0 indicates that functionally no households have direct personal experience with the industry's products or services (e.g., pipelines that transport crude oil, steel mills, or management consulting), while a score of 100 indicates a ubiquitous product or service that nearly every U.S. household has encountered (e.g., grocery stores, car manufacturers, restaurants). The score is designed to embed the breadth of household exposure, which results in industries that are split roughly evenly between household and business customers receiving intermediate scores.

The classification was performed using Claude (Anthropic, Opus 4.6), a large language model, prompted with the following instructions:

> *In the csv file naics4_consumer_shares_dataset_Claude you can find a column that lists 4-digit NAICS codes. I want you to fill in the second column with a variable that takes on a 1 if the typical end consumers of the industry's products are households and 0 if businesses. I am trying to separate which companies have products and services that households would have personal experience with and which companies have products and services that households would not have personal experience with. An alternative version of this would be "niche" versus "broad" usage of the products, with "1" for broad and "0" for niche. Industries whose products are ultimately used by only a smaller number of people, like Oracle making database management software, would be "0" whereas Google or Coca Cola would be "1" because virtually everyone in the US has used Google or seen a Coca Cola product. For each NAICS grouping, give a rating from 0 to 100 where 0 is "very few people will have experience with this industry's services and products" and 100 is "this is a ubiquitous service/product and everyone has some form of experience with this industry." Automobile manufacturers and food service companies would be 100 whereas heavy industrials and oil refiners would be 0. Service industries predominantly providing niche B2B services would be closer to 0 whereas a service industry split 50/50 between consumer services and B2B services would be a 0.5.*

The model assigned scores at the specific 4-digit NAICS level where standard industry definitions exist (e.g., NAICS 3254 Pharmaceutical and Medicine Manufacturing = 85; NAICS 5221 Depository Credit Intermediation = 90; NAICS 4451 Grocery Stores = 95), and applied 2-digit sector-level defaults for codes without a specific mapping. The procedure covered all 2,400 four-digit NAICS codes in the dataset spanning 24 sectors. Representative scores by sector are summarized below:

- **High consumer exposure (scores 80–100):** Retail trade (NAICS 44xx–45xx), accommodation and food services (72xx), healthcare providers (62xx), utilities (22xx), arts and entertainment (71xx), personal care services (8121), postal and courier services (49xx).
- **Moderate consumer exposure (scores 30–70):** Finance and insurance (52xx, with banks scoring 90 and investment funds scoring 20), real estate (53xx), residential construction (2361), passenger transportation (4811, 4851–4859), information and telecommunications (51xx), legal and accounting services (5411–5412).

- **Low consumer exposure (scores 0–20):** Mining and oil extraction (21xx), wholesale trade (42xx), management of companies (55xx), professional and technical services (54xx), primary metals and basic chemicals manufacturing (3311, 3251), pipeline transportation (4861–4869), freight trucking (4841–4842).

Firms are then assigned to the "consumer" portfolio if their 4-digit NAICS consumer-experience score meets or exceeds a threshold of 50% on the 0%-100% scale, and to the "producer" portfolio otherwise.