Demand-Based Subjective Expected Returns

Alessandro Crescini   Fabio Trojani   Andrea Vedolin

March 2024

Abstract

This paper proposes a theoretical framework for recovering investors' subjective beliefs/expected returns using holdings data and option prices under the weak assumption of no-arbitrage. We empirically document that the statistical properties of subjective expected returns on the market differ wildly across investor type and depend crucially on their portfolio composition. While expected returns estimated from price data alone suggest that expected returns are highly volatile and countercyclical, including holdings data can imply returns that are less volatile and procyclical. More specifically, we show that the expected returns inferred from public investor beliefs increase in bad times when they become the net suppliers of crash insurance in option markets, mirroring price-based estimates. Financial intermediaries' expected returns decrease during bad times when they become the net buyers of crash protection when their constraints bind. Our findings are in line with the survey literature that documents large heterogeneity in measures of expected returns.

Keywords: expected returns, options, portfolio holdings, recovery
JEL Classification: G12, G40

Comments are very welcome. We are grateful to Alireza Tahbaz-Salehi, Ruslan Tuneshev (discussant), and participants at the 2024 Derivatives and Asset Pricing Conference for helpful comments and suggestions.
University of Geneva, email: alessandro.crescini@unige.ch.
University of Geneva and Swiss Finance Institute, email: fabio.trojani@alphacruncher.com.
Boston University, NBER, and CEPR, email: avedolin@bu.edu.

1 Introduction

Canonical estimates of the expected return on the market inferred from asset prices suggest that the expected return on the market rises significantly during crises periods and is highly volatile. Estimates of expected returns from survey data, however, are less volatile and can be pro-, a-, or counter-cyclical depending on investor type, see, e.g., Greenwood and Shleifer [2014], Nagel and Xu [2023], and Dahlquist and Ibert [2023], respectively. Price-based measures of expected returns are estimated assuming that the representative agent is unconstrained and ignore information on investors' portfolio holdings. Recent empirical evidence, however, posits that holdings data should be informative about investors' risk perceptions, see, e.g., Giglio et al. [2021].

In this paper, we propose a theoretical framework for recovering beliefs/expected returns from prices and holdings data for heterogeneous investors under the weak assumption of no-arbitrage. More specifically, we theoretically show that investors' subjective expected return can be directly inferred from their holdings and the corresponding option prices in real-time. We empirically document substantial heterogeneity in expected return estimates across investor types. Most importantly, we find that expected returns recovered from holdings data can deviate in interesting ways from price-based measures. Using transaction-level data on buy and sell order on S&P500 index options, we show that the expected return of financial intermediaries mirrors measures of expected returns recovered from asset returns alone: it is very volatile and rises in times of distress when intermediaries' constraints bind. Public investors (retail and institutional) who trade against financial intermediaries are net demanders of deep-out-of-the-money puts (insurance) in normal times, see, e.g., Gârleanu, Pedersen, and Poteshman [2008] and Chen, Joslin, and Ni [2019]. Contrary to intermediaries, we find that their expected returns are less volatile and drop during financial crises when they become net suppliers of puts. In line with a large literature that studies measures of expected returns inferred from survey data, we conclude that the dynamics of subjective measures of expected returns vary greatly across investors.

In arbitrage-free markets, prices are the expected value of future payoffs discounted by some stochastic discount factor (SDF) M. The expectation is computed under the probability measure $\mathbb{P}$ supported by M. While $\mathbb{P}$ encodes the investor's subjective belief, the SDF encodes her risk preferences. Standard methods extract agents' beliefs from asset prices under some assumptions for M. These methods, however, ignore information about quantities (such as portfolio holdings, trading flows, open interest, etc.) which, different from prices, are available on a granular level, that is, for each investor.

Our "demand-based" belief recovery extracts $\mathbb{P}$ without the need to specify a specific asset pricing model by leveraging investor-level data on holdings together with option prices.

More specifically, we assume that investors with potentially heterogeneous beliefs can hold wealth shares in the market index and a family of options written on the market index. Our main theoretical result posits that subjective expected returns under $\mathbb{P}$ are only a function of investors' holdings and option prices. As is well-known, a sufficiently large arbitrage-free cross-section of options implicitly founds a pricing measure $\mathbb{Q}$, which can be fully determined by observed option prices (Breeden and Litzenberger [1978]). Since holdings are observable on an investor basis and any payoff under $\mathbb{Q}$ can be recovered from option prices directly, we obtain a measure of subjective expected returns in real-time.

With this methodology, we obtain SDF projections that are functions of the index and the options returns. Therefore, the ensuing expected market returns may be very volatile, with sign and cyclicality properties that depends on the contingent state of the economy. The shapes of the SDF projections also span a large variety of functional forms. For instance, we can recover loss averse investors with time-varying risk aversion, who expect a relatively stable market in the future and contribute to a low premium. These agents are equipped with a strongly asymmetric U-shaped SDF projection. These agents take short positions on out-of-the-money options. Similar intuition also allows us to recover SDF projections that are monotonically decreasing, monotonically increasing, or tilde-shaped. Earlier literature ignores options because they are in zero-net supply. Since in reality, options are non-redundant, we show that holdings in option portfolios are informative about investors' beliefs. For example, a larger investment in the options corresponds to conservative investors that are progressively more sensitive to higher-order risk factors and trade options to reallocate them profitably. Not only the sign, but also the cyclicality of the market risk premium is endogenous to investor's belief. In order to illustrate our theoretical framework, we make use of buy and sell orders of large investors in option markets.

More specifically, we use our results to gain insights about the beliefs and risk perceptions of two groups of option market participants: public investors (retail and institutional) and intermediaries. To this end, we leverage the CBOE Open-Close Database which records daily buy and sell orders per investor category for every option. Real-time holdings data, allows us to recover each investor's beliefs such that the solution to our recovery problem is aligned with the observed portfolio positions.

We summarize our empirical findings as follows. First, we find that public investors and market makers have complementary patterns with regards to the shape of their SDFs. For example, during normal times, market makers have U-shaped SDFs as a function of expected returns, while public investors have inverted U shapes. The U-shape comes from the fact that during normal times, market makers hold large short positions in calls and puts. They are hence exposed to both changes in the underlying on the up- and the downside. Public investors who act as net demanders of these options, have the opposite patterns. The patterns are strikingly different during "crisis" periods such as the Great Financial Crisis or Covid. We find that in November 2008, public investors' SDF projection is monotonically decreasing, while market makers' SDF is flat for negative returns and increasing for positive returns. The reason for that is that during that period, market makers effectively became net demanders of downside protection as their financial constraints started to bind.

Second, the shifting portfolio holdings and exposures to downside risk during crisis periods across the two investors has large effects on the properties of expected returns. We find that the ensuing expected returns are pro-cyclical and very volatile for market makers. In fact, we observe that expected returns become negative during crises. Intuitively again, this happens because of the large long put positions that they hold on their portfolios. Public investors, on the other hand, have countercyclical expected returns because they make the market for crash insurance during bad times.

To summarize, our empirical findings provide intuition for why some of the literature has documented different cyclicality patterns in surveys and price-based measures. Our focus on portfolio holdings rather than just prices alone explains why the time-series properties of beliefs can drastically change depending on investor type.

Related Literature. This paper is related to several strands of the literature. Starting from the seminal work of Ross [2015] a growing literature has proposed ways to recover investors' subjective beliefs, see, e.g., Borovička, Hansen, and Scheinkman [2016], Jensen, Lando, and Pedersen [2019], among others for recent refinements of the Ross [2015] recovery theorem. Our framework differs from these papers in at least three ways: First, we do not constrain the recovered belief with model-specific assumptions on the SDF nor agent preferences. Second, we include demand-based data instead of just asset pricing data to extract investor-specific beliefs. Third, our framework allows us to recover conditional beliefs in real-time.

Chen, Hansen, and Hansen [2020], Ghosh and Roussellet [2023], and Korsaye [2024] use survey data in addition to price data to recover the representative agent's belief and study their properties relative to a rational expectations framework. As we show, holdings data allows us to recover beliefs on a much more granular level, that is at the investor level. More generally, our theoretical framework also allows for the inclusion of survey data. However, long time-series of granular survey data is hard to obtain.

Our paper is also related to the literature that studies the option demand of heterogeneous investors. For example, Chen, Joslin, and Ni [2019] document how variation in the net demand of deep OTM put options between intermediaries and public investors is driven by intermediaries' constraints. Almeida and Freire [2022] show how net option demand helps explain the pricing kernel puzzle. And Farago, Khapko, and Ornthanalai [2021] study a heterogeneous agent economy to explain index put trading volumes. We complement this literature by estimating intermediaries' and public investors' beliefs from observed option demand.

Our paper is also closely related to the literature that makes use of higher-order moments of the market return and options to recover measures of the expected return. Even tough these papers do not explicitly recover investors' beliefs, their results are nested in our framework. For example, Martin [2017], Martin and Wagner [2019], and Gao and Martin [2021] derive lower bounds on expected returns for stocks by assuming that the risk premium on an asset and its derivatives can be inferred from the allocation of a growth-optimal portfolio that maximizes an investor's long-run growth. Similarly, Gormsen and Jensen [2022] study physical (as opposed to risk-neutral) moments as perceived by a power utility investor. Gandhi, Gormsen, and Lazarus [2023] study the term structure of expected returns inferred from option prices and find that long-term expected returns are (excessively) countercyclical and volatile. Tetlock [2023] assumes that the SDF is given by $1/\hat{R}$, where $\hat{R}$ is the return of a portfolio of higher-order moments of R whose weights come from regressing the variance premium on some risk-neutral moments to obtain point estimates of the expected return. Our findings show that the results in these papers are sensitive to the assumption of the optimal weights in the investor's portfolio. While Martin [2017] assumes that investors hold 100% of their wealth in the market (and none in the derivatives themselves), Tetlock [2023] allows for holdings in both the market and the derivatives where the growth-optimal weights are calculated from the data. In our setting, we do not need to make any assumptions about utilities or wealth. In particular, our estimation framework incorporates information from actual weights as provided by transaction level data.

Our paper contributes to an empirical literature studying beliefs of heterogeneous investors. Dahlquist and Ibert [2023] document large heterogeneity in asset managers' beliefs, while Giglio et al. [2021] study the relationship between retail investors' beliefs and portfolio holdings. Meeuwis et al. [2022] document that political orientation determines households' beliefs and portfolio allocation into risky assets. Ghosh, Korteweg, and Xu [2022] recover heterogeneous beliefs from the cross-section of stock returns. Our paper is different from these papers since we recover beliefs from price and holdings data jointly, allowing us to measure beliefs for a long-time series at the daily frequency for large investors.

Finally, our paper contributes to the demand-based asset pricing literature starting with the seminal work of Koijen and Yogo [2019]. Similar to our approach, asset-demand systems impose constraints such that holdings data is matched and market clearing holds in equilibrium. While that literature is mainly interested in how heterogeneous investors affect movements in asset prices, our focus is on expected returns.

Outline. The rest of the paper is organized as follows. The key idea of our paper is that holdings data is informative about investors' risk perceptions. We illustrate this idea in an intuitive example in Section 2. Section 3 presents a general theoretical framework where we show how to infer investor's beliefs from holdings data. Section 4 contains our main empirical results. All proofs and some additional mathematical details are provided in the Appendix.

Additional results are gathered in an Online Appendix.

2 Illustrative Example

The key idea of our paper is that portfolio holdings are informative about risk perceptions of market participants. To provide some intuition, we start with an example to illustrate how portfolio holdings affect SDFs. Assume heterogeneous investors who hold a growth-optimal portfolio. That is, investors maximize expected long-run wealth. As is well-known, in this case, the growth-optimal return is the reciprocal of the stochastic discount factor, i.e., $M^\star = 1/R^\star$, where $R^\star$ is the return of the optimal portfolio according to the investor's subjective view, see, e.g., Long [1990]. Even though investors have the same preferences and are subject to the same constraints, the optimal portfolio can vary across investors because they may have different beliefs.

To set a benchmark, assume there exists a specific constrained utility-maximization problem whose solution is a portfolio fully invested in the market. In that case, the SDF takes the following form: $M^0 := 1/R$, where R is the return on the market index. This is the case studied in Martin [2017]. A priori, there is no reason to exclude other traded assets (say, options) from the optimal portfolio. In fact, ample empirical evidence in the literature shows that options are non-redundant securities and demand for options can be in the order of trillions of dollars, especially following market crashes.[1] In that case, the corresponding optimal portfolio will have non-zero positions in the index options, and the return $R^\star$ will be different from R (and in turn $M^\star \neq M^0$).

Let M^* be the SDF consisting of a portfolio being long some calls or puts. Figure 1 compares the value of $M^\star$ relative to M^0, as a function of the only state variable R for calls (left panel) and puts (right panel) across different moneyness. For illustrative purposes, we assume that the investor holds a portfolio consisting of 90% in the index and 10% of calls or puts.[2] As is immediately evident, even for a small investment in options (with respect to the investment in the underlying) and small fluctuations of the market return, the ratio $M^0/M^\star$ is different from 1. While for in-the-money (ITM) calls, the ratio is not very different from 1, $M^0/M^\star$ rises significantly for at-the-money (ATM) and out-of-the-money (OTM) calls and puts. For example, assuming the market excess return is +20% (-20%), the ratio increases to 17 (25) for OTM calls (puts).

[1] Theoretically, positive net demand for options can arise in settings with heterogeneous beliefs with frictions and market incompleteness. Johnson, Liang, and Liu [2016] study empirical drivers of option demand and argue that the primary reason for the high demand in index option is transfer of unspanned crash risk.

[2] In later sections, we will use transaction-level data from CBOE to track portfolio holdings in real-time.

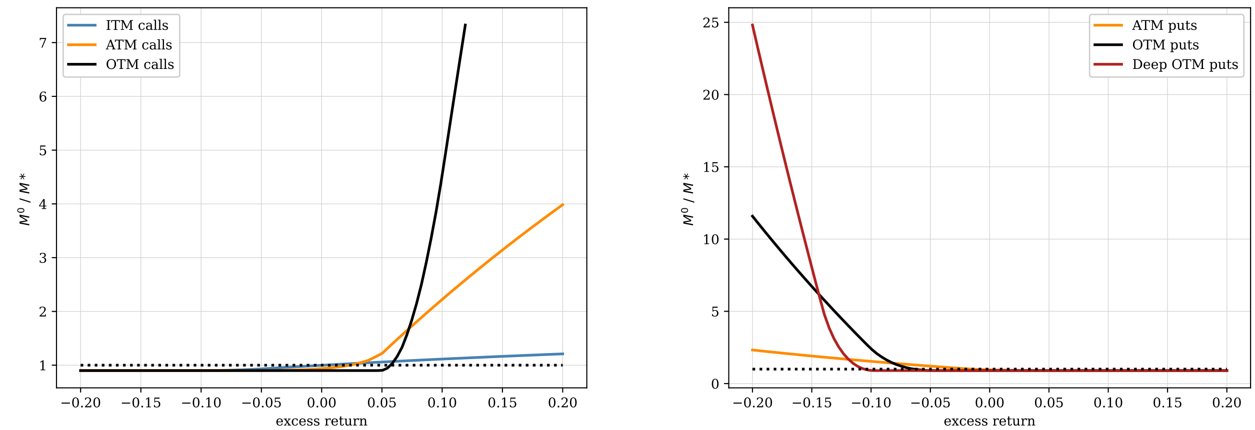

Figure 1. Ratio between Benchmark SDF M^0 and $M^\star$
Notes: This figure plots $M^0/M^\star$ as a function of the excess return on the market. $M^\star = 1/R^\star$, where $R^\star$ is the return of a portfolio investing 90% of the wealth in the underlying and 10% in an equally-weighted portfolio of calls (left plot) and puts (right plot) with different moneyness. ATM options have moneyness $m = K/S \in [.95, 1.05]$ (being S the current index value), ITM calls (OTM puts) are with $.9 \leq m < .95$, OTM calls with $1.05 < m \leq 1.1$.

Intuitively, we can interpret the ratio, $M^0/M^\star$, as the probability distortion one has to introduce to recover probability measure $\mathbb{P}^\star$ (supported by SDF $M^\star$) from the benchmark $\mathbb{P}^0$ (as supported by SDF M^0). For instance, the belief of an investor who optimally chooses to be long in puts, is more left-skewed than the benchmark. Investing in puts, shifts the probability mass uniformly from the region of positive market return towards the region of negative returns proportional to the strikes and the moneyness. In general, investors who assign higher weights to extreme events have higher demand for deep OTM puts. The reverse holds when the investors is long calls.

In a next step, we study the effect of investors' demand on perceived risk premia. In Figure 2, we plot the time-series of the expected market return for different $M^\star$ implied by options with different moneyness together with the benchmark log investor case who is 100% invested in the index (M^0).

The upper panel plots the perceived risk premia for investors who in addition to the index, hold call options with different moneyness. As we can see, the patterns mirror the benchmark case almost one-for-one. Expected returns increase in bad time and are highly volatile. Intuitively, the size of the risk premium itself increases relative to the benchmark since the options represent a levered trade on the underlying itself. The portfolio with ATM calls has the highest premium, since ATM calls move one-for-one with the underlying market index.

We can juxtapose this pattern with inferred risk premia from investors' who are long in puts. As can be seen from the lower panel, being long in puts decreases the exposure to market risk, and as a consequence, the corresponding risk premium is lower. In fact, the risk premium even becomes negative. As discussed before, holding puts reflects the view of investors holding more left-skewed beliefs, who expect higher (negative) fluctuations of the market. From their perspective, the risk-return ratio given by holding the index alone is not profitable. Accordingly, the volatility of the premium is higher with respect to what we recover under M^0. The contribution of the puts may be enough to offset the premium towards a pro-cyclical pattern, whose size and cyclicality properties depend on the moneyness, and on the amount of wealth invested in the puts.

To be more specific, in Figure 3, we vary the amount of wealth invested in OTM puts from 5% to 20%. We notice that the recovered market premium is pro-cyclical when the wealth fraction invested in OTM puts is 20%. Higher investments in insurance means that market downturns are ex-ante perceived as less risky.

We conclude that not only the size, but also the sign, cyclicality, and volatility of the market risk premium depend on asset demand. While these results are based on hypothetical portfolios, in our empirical section, we will use transaction level data on buy and sell orders to track investors' beliefs over time.

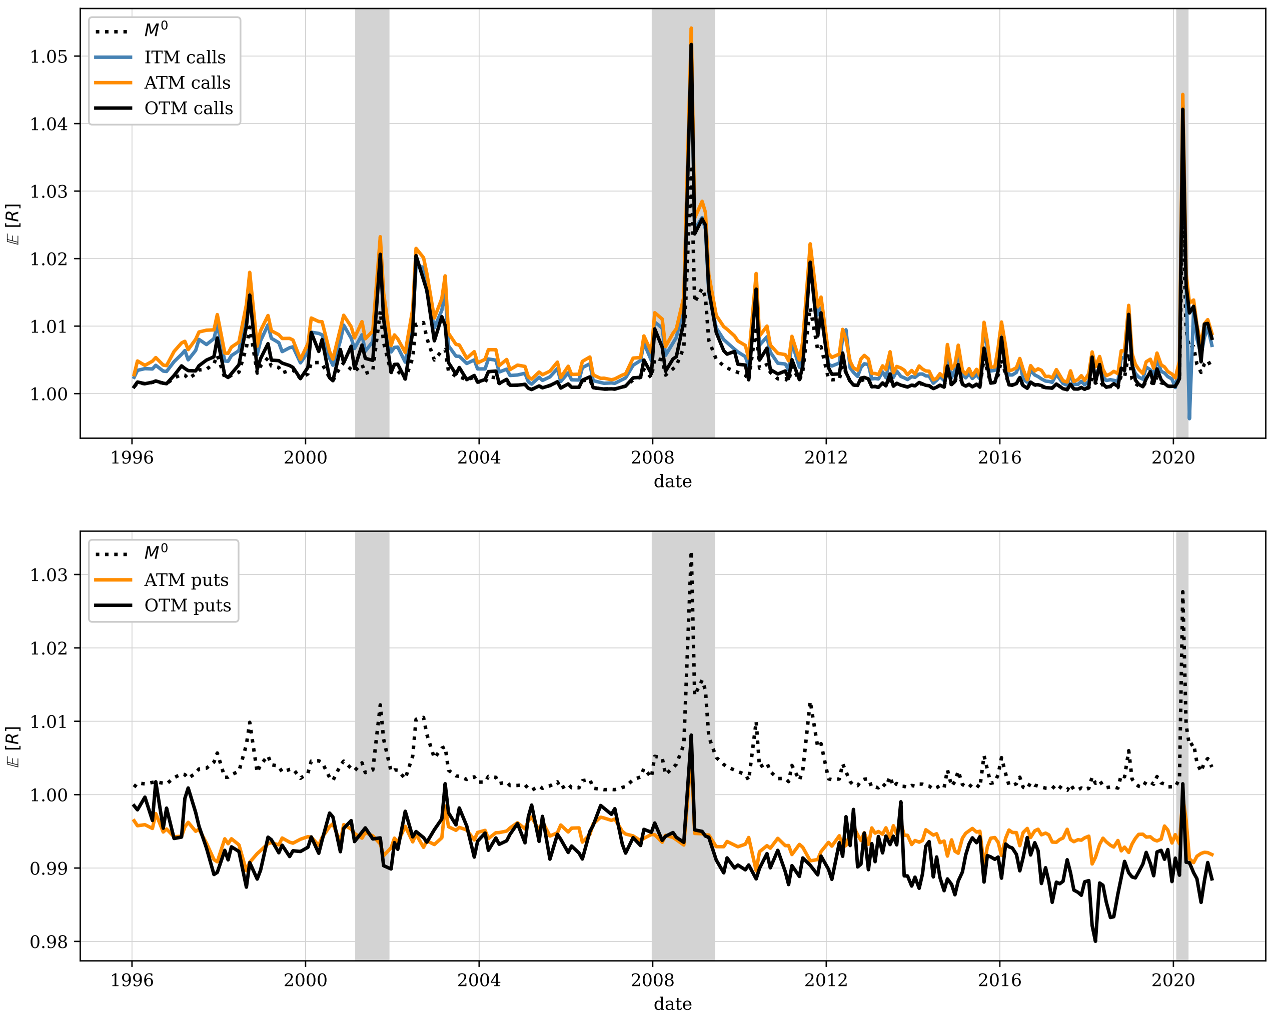

Figure 2. Monthly Time-Series of Expected Market Return

Notes: This figure plots the expected market return recovered from different stochastic discount factors. The dotted line represents the benchmark $M^0 = 1/R$, where M^0 is 100% invested in the index. The other lines arise from $M^\star = 1/R^\star$, where $R^\star$ is the return of a portfolio investing 90% in the market and 10% in an equally-weighted portfolio of calls (upper panel) and puts (lower panel) with different moneyness (see caption of Figure 1). Grey areas indicate NBER recession periods.

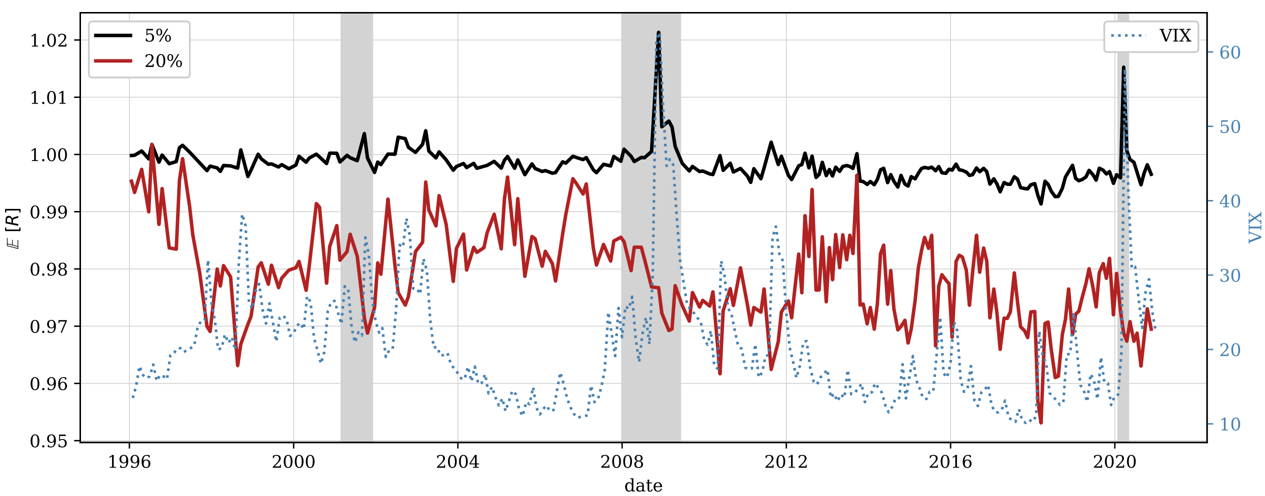

Figure 3. Monthly Time-Series of Expected Market Return

Notes: This figure plots the expected market return recovered from different stochastic discount factors. Each supporting portfolio invests a fraction $1 - x$ of the wealth into the market, and x into an equally-weighted portfolio of OTM puts $(.9 \leq K/S < .95)$; x is either 5% or 20%. The blue dotted line displays the implied volatility index (VIX). Grey areas indicate NBER recession periods.

3 Theoretical Framework

Consider an economy consisting of m different investors indexed by i, each with logarithmic utilities, but potentially different beliefs. Let $\mathbb{E}_i[\cdot]$ denote the subjective expectation operator with respect to the belief of investor i. Motivated by our illustrative example presented before, all investors can trade three types of assets, a risk-free asset with return R_f, a risky asset with forward return R, and an entire family of options on the risky asset with a continuum of strike prices. In this paper, we assume that the risky asset is the S&P 500 index and the options are European calls and puts. Let $\mathbf{R}^e = \mathbf{R} - \mathbf{1}$ be the excess forward return of the index and options.

Our goal is to recover the physical belief $\mathbb{P}_i$ for investor i under the minimal assumptions stated above to infer the subjective expected return on the market, $\mathbb{E}_i[R]$. Since investors have logarithmic utility, it immediately implies that one can define an agent-specific SDF M_i that prices all assets from the perspective of agent i as follows:

$$M_i = \left(1 + \boldsymbol{\theta}_i' \mathbf{R}^e\right)^{-1}, \tag{1}$$

where $\boldsymbol{\theta}_i$ is the share of wealth invested in the market index and the options by investor i. No arbitrage implies that

$$\mathbb{E}_i[M_i \mathbf{R}^e] = \mathbf{0}.$$

Several remarks are in order. Our set-up is similar to Martin [2017], who also assumes a log utility investor. The crucial difference to Martin [2017] is that he assumes that the optimal portfolio held by the log investor consists of 100% in the market. A consequence of this assumption is that all options are redundant as otherwise the growth-optimal portfolio would contain positions in options. In contrast, our framework does not assume such redundancy. However, our setting nests Martin [2017]'s case if θ is 1 for the market and zero otherwise. Another special case is Tetlock [2023] who assumes that the investor can hold (integer) power-contracts written on the market index. In general, however, we do not observe holdings of power-contracts. What we do observe is the holdings of plain vanilla options, that is, calls and puts.

In the following, we will consider two different cases. First, the fact that portfolio holdings are observable in the data, allows us to directly recover agent i's expected return of the market as a function of holdings and option prices. Second, it is reasonable to assume that holdings data is measured with some error. The reason for this is twofold. First, we only observe a subset of the "true" portfolio of investors. For example, while we observe the holdings on calls and puts for the S&P500, we do not observe the holdings on other derivatives, such as

SPY options (that is on the ETF tracking the S&P500).[3] Second, our data is aggregated to the monthly frequency and across different types of investors (retail and institutional). We will therefore consider a second case where we assume that portfolio holdings are observed with some measurement error.

3.1 Subjective Expected Returns

As is well-known, an arbitrage-free cross-section of options suffices for the existence of a probability measure $\mathbb{Q}$, which determines the price of any payoff X that is replicable by a delta-hedged option portfolio (see, e.g., Acciaio et al. [2016]). In our application, the pricing measure $\mathbb{Q}$ is a forward probability between times t and $t+1$. This trivially implies that $\mathbb{E}^{\mathbb{Q}}[\mathbf{R}^e] = \mathbf{0}$.

Since we can obtain the risk-neutral expectations, we can now relate this risk-neutral expectation to agent i's subjective expectation via the following relationship:

$$\mathbb{E}^{\mathbb{Q}}[X] = \frac{\mathbb{E}_i[M_i X]}{\mathbb{E}_i[M_i]}$$

for any payoff X. Therefore, we can use the observed portfolio weights $\boldsymbol{\theta}'$ to construct the SDF M_i in (1) and obtain our main result of the paper.

Proposition 1 (Subjective Expected Returns)**.** *Assume that portfolio weights $\boldsymbol{\theta}_i$ are observed without error, then the subjective expected return on the market of investor i is given by:*

$$\mathbb{E}_i[R] = \mathbb{E}^{\mathbb{Q}}[(1 + \boldsymbol{\theta}'\mathbf{R}^e)R] \tag{2}$$

We can directly compute equation (2) from the data, since we observe the portfolio holdings $\boldsymbol{\theta}_i$ and we can calculate the risk-neutral expectation of $(1 + \boldsymbol{\theta}'\mathbf{R}^e)R$ using the Carr and Madan [2001] formula. Notice that we can recover the implied equity premium in Martin [2017], again by assuming that the weights θ_i are zero except for the market, where the weight is equal to 1. In that case, the expected return of investor i is equal to $\mathbb{E}^{\mathbb{Q}}[R^2]$ which is the risk-neutral variance measured under the forward-measure $\mathbb{Q}$. We can now study the case with measurement error.

3.2 Subjective Expected Returns With Measurement Error

In this section, we provide bounds on the subjective return on the market assuming that portfolio weights $\boldsymbol{\theta}_i$ are observed with error. Since we want to constrain the amount by which the

[3] Moussavi, Xu, and Zhou [2024] show that market makers tend to hold positions in both markets.

optimal portfolio weights can deviate from the observed weights, we impose some constraint such that

$$d(\boldsymbol{\theta}, \boldsymbol{\theta}^\star) \leq \delta \,, \tag{3}$$

where $\boldsymbol{\theta}^\star$ are the growth-optimal portfolio weights and for some convex discrepancy function $d(\cdot,\cdot) \geq 0$ such that $d(x, y) = 0$ if and only if $x = y$. Intuitively, δ measures the amount that optimal weights can deviate from the observed weights. To be concrete, let's assume a L^2-norm. In that case, $d(\boldsymbol{\theta}, \boldsymbol{\theta}^\star) = \frac{1}{2}||\boldsymbol{\theta}^\star - \boldsymbol{\theta}||_2^2$. To solve for the subjective expected return, we now have to solve the following optimization problem:

$$\inf_{\boldsymbol{\theta}^\star \in \mathbb{R}^N} \left\{ \mathbb{E}^{\mathbb{Q}}[(1 + \boldsymbol{\theta}^{\star\prime}\mathbf{R}^e)R] + \lambda \left(\frac{1}{2}||\boldsymbol{\theta}^\star - \boldsymbol{\theta}||_2^2 - \delta \right) \right\} \,. \tag{4}$$

This leads us to our second main result which are closed-form solutions on the bounds of subjective expected returns.

Proposition 2 (Bounds on Subjective Expected Returns)**.** *Define $f(R) := (1 + \boldsymbol{\theta}'\mathbf{R}^e)R$. Assume that portfolio weights $\boldsymbol{\theta}$ are observed with error, in that case, the lower bound for the subjective return on the market for investor i is:*

$$\mathbb{E}_i^L[R] = \mathbb{E}^{\mathbb{Q}}[f(R)] + \mathbb{E}^{\mathbb{Q}}[\boldsymbol{\theta}'\boldsymbol{R}^e f(R)] - \sqrt{2\delta}||\mathbb{E}^{\mathbb{Q}}[\boldsymbol{R}^e f(R)]||_2 \,. \tag{5}$$

The upper bound for the subjective return on the market for investor i is given by:

$$\mathbb{E}_i^U[R] = \mathbb{E}^{\mathbb{Q}}[f(R)] + \mathbb{E}^{\mathbb{Q}}[\boldsymbol{\theta}'\boldsymbol{R}^e f(R)] + \sqrt{2\delta}||\mathbb{E}^{\mathbb{Q}}[\boldsymbol{R}^e f(R)]||_2 \,. \tag{6}$$

It is obvious that if the measurement error is assumed to be zero, i.e., $\delta = 0$, that equations 5 and 6 coincide with equation 2. Intuitively, we can interpret the lower bound as the most conservative assessment of the subjective expected return for any investor who's portfolios align in a neighborhood around the observed portfolios, θ.

4 Empirical Analysis

This section describes the data used and how we empirically implement our main theoretical results.

4.1 Data

To empirically implement our theory, we make use of the CBOE Open-Close dataset that provides daily buy and sell volumes of SPX options since 1996 aggregated by type of position (opening/closing) and origin: (i) customer; (ii) broker-dealer; (iii) firms; and (iv) market makers. We aggregate these daily volumes to cumulative positions for the last three categories and label them market makers. The label 'broker-dealer' is available only from 2011.

The volume data comes without pricing information. We obtain end-of-day option prices from the OptionMetrics database and use bid- and ask-prices to compute mid-point prices.

Throughout our empirical analysis, we use monthly frequency and horizons. Portfolios investing in the risk-free asset, S&P 500, and monthly OTM options are assembled at date t and held to maturity without intermediate rebalancing before the third Friday of the expiration month. We assume that investors build delta-hedged strategies in which the investment in the index is the negative of the delta of the portfolio at time t, see, e.g., Gayda, Grünthaler, and Harren [2023] and Baltussen, Jerstegge, and Whelan [2024].

4.2 Implementation

In order to implement the expression in equation 2, we apply the Carr and Madan [2001] formula to approximate $\mathbb{E}^{\mathbb{Q}}[(1+\boldsymbol{\theta}'\mathbf{R}^e)R]$. More specifically, let $X(K)$ be the payoff of an option with strike K, and recall that $f(R) := (1+\boldsymbol{\theta}'\mathbf{R}^e)R$, then the subjective expected return is given by:

$$\begin{aligned}\mathbb{E}_i[R] &= \mathbb{E}^{\mathbb{Q}}[f(R)], \\ &\approx \mathbb{E}^{\mathbb{Q}}[f(1) + f'(1)(R-1) + \sum f''(k)x(k)\Delta k],\end{aligned}$$

where $x(k) := \frac{X(K)}{F}$ with F being the forward price of the market index, and $k = \frac{K}{F}$.

4.3 Subjective Expected Returns with No Measurement Error

We start from our main result provided in Proposition 1: the variation of expected returns over time as perceived by different agents assuming that the portfolio weights are observed without error. Figures 4 and 5 plot expected returns for customers and market makers over time together with the benchmark case of a full investment in the index.

We notice strikingly different dynamics of the expected return measures across customers and market makers but also distinct from the benchmark case not just in size but also comove-

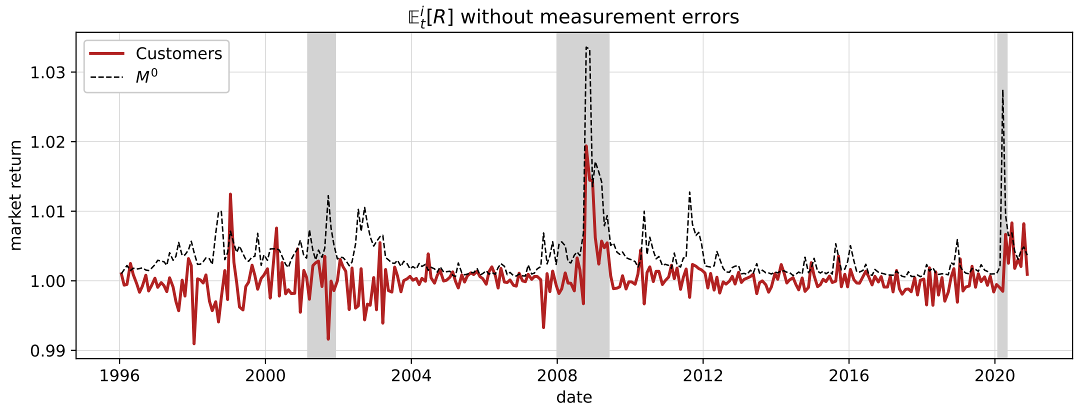

Figure 4. Subjective Expected Market Return Customers
Notes: This figure plots the time-series of the expected market return as recovered by the benchmark SDF $M^0 = 1/R$ and the $M^\star$ supported by customers' positions in delta-hedged options. Frequency is monthly and gray bars indicate NBER recessions.

ment. For example, customers' expected returns are highly volatile and increase in bad times, especially the 2008 Great Financial Crisis and Covid where the expected returns increased by 2% monthly, but not during the 2001 dot-com bubble. Subjective expected returns for the market makers drop during the 2008 crisis and the 2021 Covid period but are elevated during the 2001 crisis. The expected return recovered by M^0 is by construction counter-cyclical only. The correlation between the time-series of the benchmark and the customers is 36% and -36% for market makers.

Intuitively, our results align with the observation that public investors are typically net buyers of OTM options, but they become net sellers of OTM options during periods of higher volatility when market makers become constrained. We will discuss these results in more details later when turning to investors' SDFs. Notice also, that the expected return is below the case of Martin [2017] almost everywhere, and it is slightly less volatile (around 0.4% vs. 0.3% for both investors). This is intuitive, since our earlier results showed that options portfolio can mitigate exposure to the market risk.

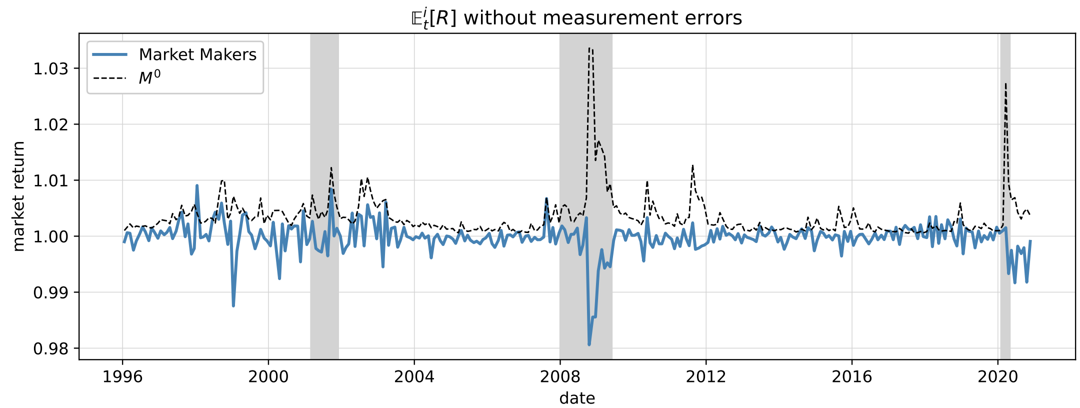

Figure 5. Subjective Expected Market Return Market Makers
Notes: This figure plots the time-series of the expected market return as recovered by the benchmark SDF $M^0 = 1/R$ and the $M^\star$ supported by market makers' positions in delta-hedged options. Frequency is monthly and gray bars indicate NBER recessions.

4.4 Expected Returns with Measurement Errors

We now turn to the case when portfolio holdings contain some measurement error. We solve the linear optimization problem in equation 4 with $\delta = 0.01$ to determine the most conservative and the maximum expected market return compatible with investors' observed positions.

Figures 6 and 7 plot the resulting time-series of subjective expected returns. More specifically, the upper plots show the most conservative subjective expected return for each investor; the lower plots show all the admissible values that lie between the minimum and the maximum. These act as lower and upper bounds for the subjective expected return perceived by all the possible investors whose portfolios are aligned (to some degree) with the observed positions of customers and market makers. In a sense, the lower bound represents the expectation of the "most pessimistic" investor in the group - or, equivalently, it represents the "worst-case" expectation that customers and market makers may formulate.

From our illustrative example discussed before, we expect that the lower bound is attained with a portfolio that hedges volatility risk with long positions in calls and puts. The upper bound, however, is most likely supported by portfolios which tend to have short positions in calls and puts. Therefore they will be more sensitive to periods of high volatility. Indeed the range of admissible values is larger during times of increasing volatility. This explains why the lower bound is largely pro-cyclical (correlation with the M^0 time-series is -78% for market makers and -48% for customers) and the upper bound is mostly counter-cyclical (correlations 48% for market makers and 78% for customers).

Interestingly, as the observed weights are opposite between customers and market makers,

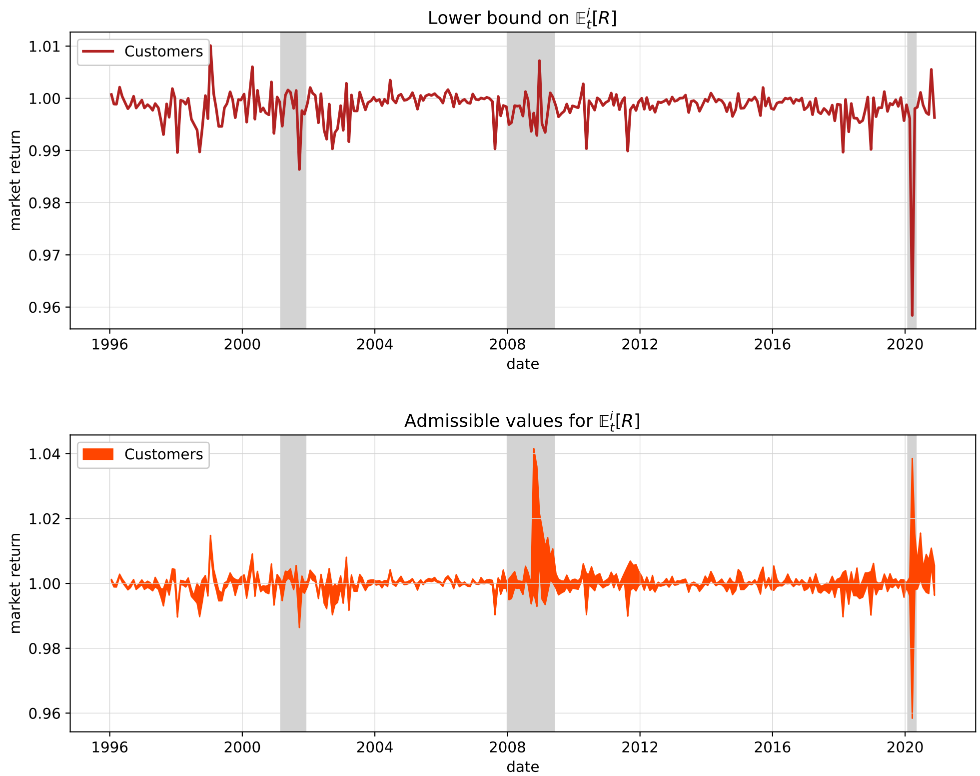

Figure 6. Bounds Subjective Expected Market Return Customers
Notes: This figure plots the time-series of the lower bound on the subjective expected market return (upper plot) and all the possible values for the expected market return (lower plot) as recovered by SDFs compatible with customers' positions in delta-hedged options. Frequency is monthly and gray bars indicate NBER recessions.

the lower bound for each investor class in perfectly negatively correlated with the upper bound of the other investor.

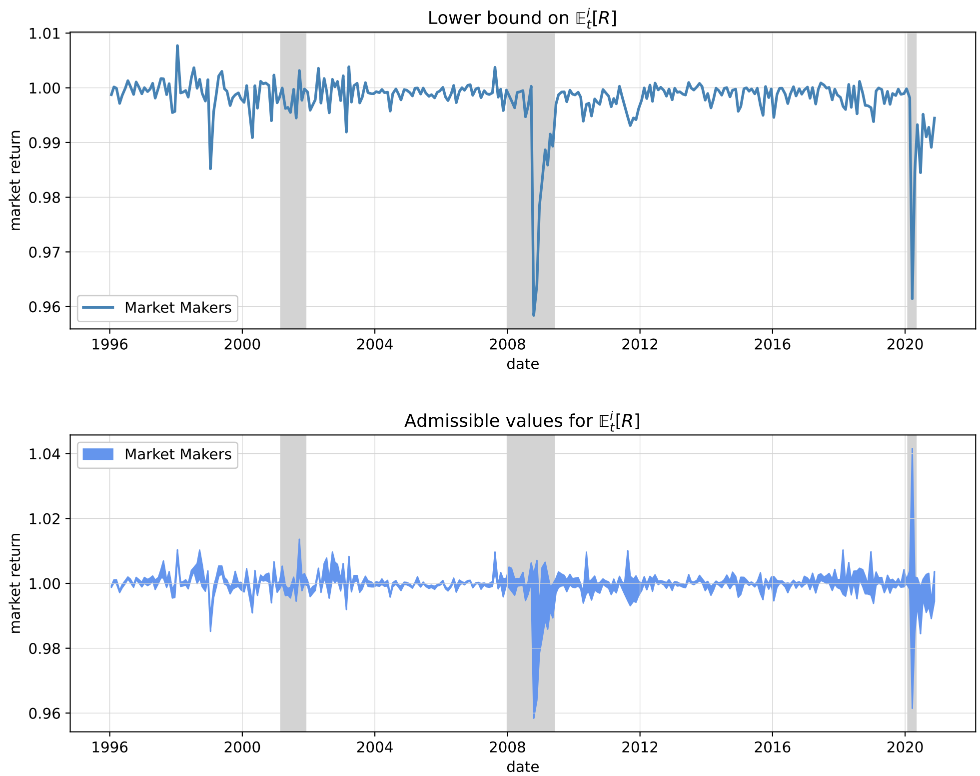

Figure 7. Bounds Subjective Expected Market Return Market Makers
Notes: This figure plots the time-series of the lower bound on the subjective expected market return (upper plot) and all the possible values for the expected market return (lower plot) as recovered by SDFs compatible with market makers' positions in delta-hedged options. Frequency is monthly and gray bars indicate NBER recessions.

4.5 Demand-Based SDFs

To build intuition, we now recover the SDFs of our two investor groups on "crisis" days. To this end, we compare the SDF shapes with and without measurement errors in the same setting as above, on October 2008 and April 2020. Figures 10 and 11 depict the results, while the observed portfolio weights are summarized in Figures 8 and 9.

In both instances, we find that customers are overall short in OTM calls and puts, while market makers are long. Moreover, the amount invested in puts is significantly larger than in calls. This echoes earlier findings in Chen, Joslin, and Ni [2019] who argue that while market makers are net sellers of insurance in normal times, they become net demanders in bad times when their financial constraints bind. The different compositions in the portfolios have large effects on the ensuing expected returns. Customers perceive larger risk premia as they have increased their exposure to the index volatility. It is 2.2% in 2008 and 0.7% in Covid. Correspondingly, market makers perceive -2.2% and -0.7%. Negative risk premium is intuitive since intermediaries hold a portfolio of crash insurance, they are protected against market

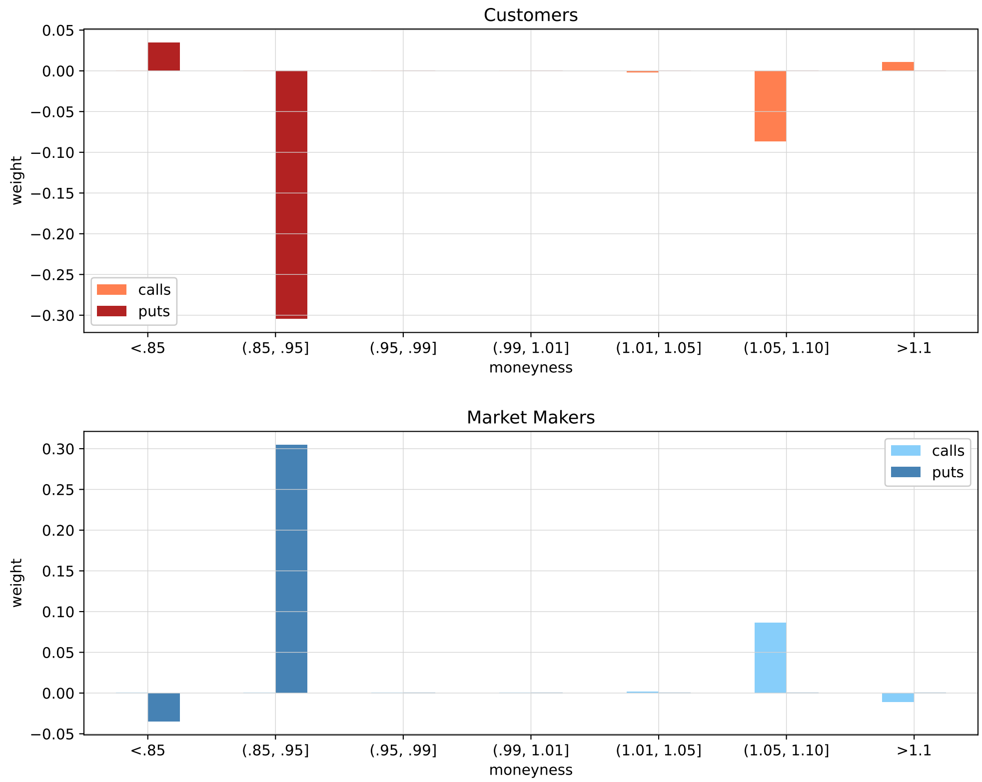

Figure 8. Portfolio Weights October 2008
Notes: This figure shows the distribution of the portfolio weights for OTM calls and puts in the observed option portfolio of customers (upper panel) and market makers (lower panel), across different levels of moneyness, on October 2008.

downturns.

Looking at the SDF shapes confirm these observations. The SDF of the customers is increasing in the upside risk and decreasing in the downside region, because of the positive weights assigned to the options. At the end it exhibits a $U-$shape. The opposite holds for market makers. This is evident during the 2008 financial crisis, when even if we consider deviations due to measurement errors, the supporting SDFs do preserve the shape. Interestingly, this does not happen in April 2020: ceteris paribus, the optimized weights may differ enough to invert the convexity of the SDF. For example, the market makers' SDF can become $U-$shaped, leading to high maximum expected return of 4% monthly. That's why the range of admissible values for the investors on April 2020 is the highest in the dataset.

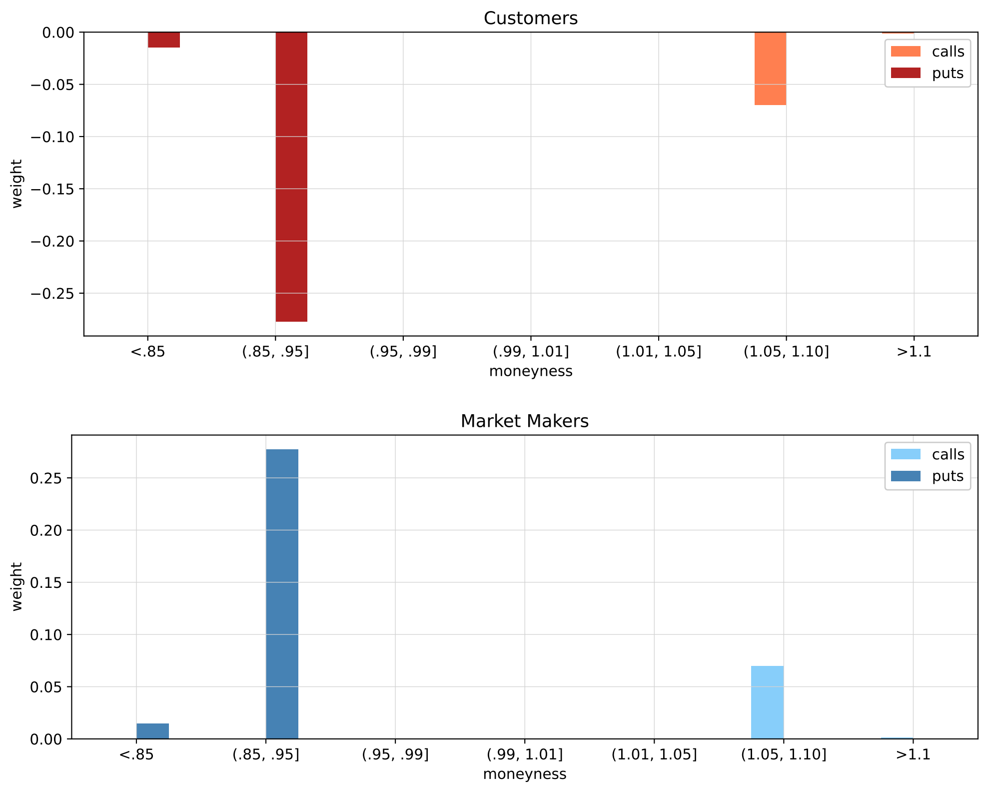

Figure 9. Portfolio Weights Covid
Notes: This figure shows the distribution of the portfolio weight for OTM calls and puts in the observed option portfolio of customers (upper panel) and market makers (lower panel), across different levels of moneyness, on April 2020.

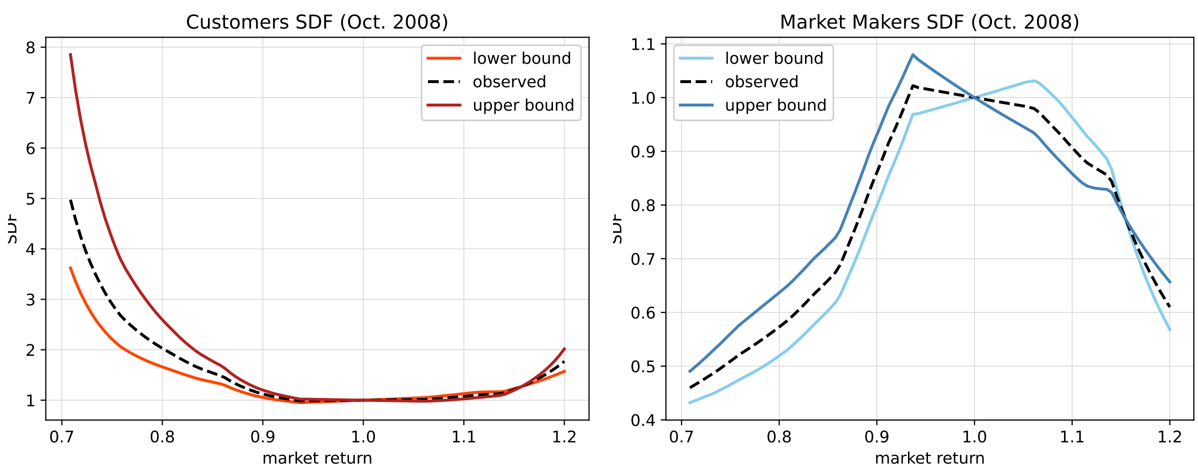

Figure 10. Recovered SDFs
Notes: This figure plots the stochastic discount factors $M^\star$ recovered for market makers and customers on October 2008, as function of the market return. The black line is the SDF recovered from the observed investors' portfolios. The coloured lines display the SDFs supporting the minimum and the maximum expected market return attainable under the constraints.

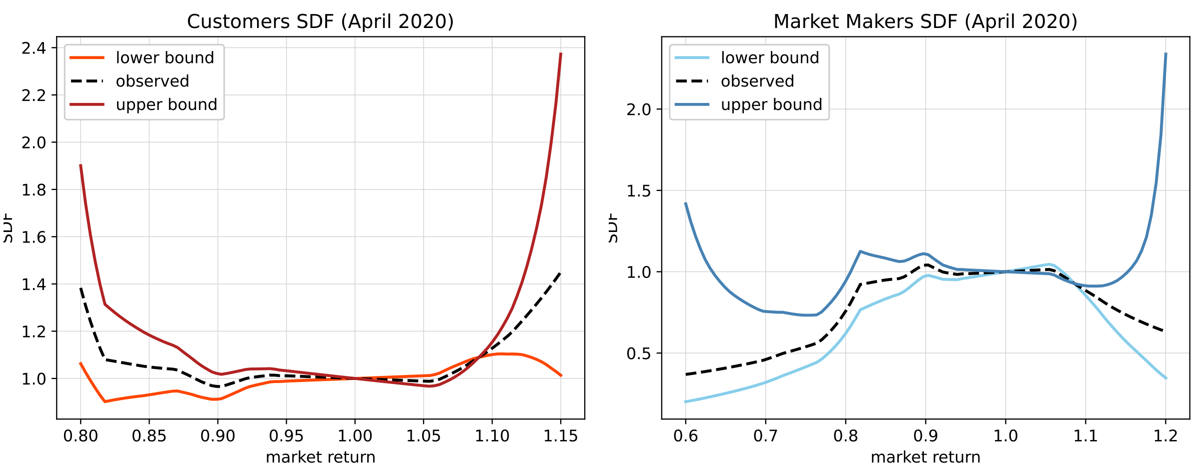

Figure 11. Recovered SDFs

Notes: This figure plots the stochastic discount factors $M^{\star}$ recovered for market makers and customers on April 2020, as function of the market return. The black line is the SDF recovered from the observed investors' portfolios. The coloured lines display the SDFs supporting the minimum and the maximum expected market return attainable under the constraints.

5 Conclusion

In this paper, we propose a theoretical framework for recovering investors' beliefs using demand-based data. Information about investors' holdings allows us to pin down the beliefs of individual investors when observing a cross-section of option prices. Our main empirical result is that the size, dynamics, and cyclicality properties of belief-implied expected returns vary significantly across investor types. Using granular transaction data on buy and sell orders of financial intermediaries and public investors, we show that beliefs are heterogeneous and the implied expected returns vary considerably across the two investors.

For example, customers' expected returns mirror those estimated from price data alone: they increase in bad times and display higher volatility. Public investors' expected returns are less volatile and drop in bad times, hence, are procyclical. Intuitively, this follows from the fact that in normal times, financial intermediaries are the net suppliers of insurance (puts) on the index. In bad times, when their constraints bind, they become net demanders of crash insurance.

Earlier literature rationalizes the procyclicality of expected returns elicited from survey data with models of return extrapolation, see, e.g., Nagel and Xu [2023]. Our paper is agnostic about the origins of the particular time-series patterns of expected returns recovered from the data. Future work could explore the link between our beliefs measures and those elicited from survey data and study deviations from the rational expectations benchmark.

A Proofs and Derivations

Proposition 3 (Upper and lower bounds on expected payoffs of a log investor)**.** *Suppose $\boldsymbol{\theta} \in \Theta$, where Θ is some closed convex set, indexes a log investor holding an optimal portfolio $\boldsymbol{\theta}$, with return R, and having belief $\mathbb{P}$. Further let $f(R)$ be some payoff depending on R. Then, the following upper and lower bounds hold:*

$$\mathcal{U}(f) := \sup_{\boldsymbol{\theta} \in \Theta} \mathbb{E}^{\mathbb{Q}}[Rf(R)] \geq \mathbb{E}_i[f(R)] , \tag{7}$$

and

$$\mathcal{L}(f) := \inf_{\boldsymbol{\theta} \in \Theta} \mathbb{E}^{\mathbb{Q}}[Rf(R)] \leq \mathbb{E}_i[f(R)] . \tag{8}$$

Proof. Since portfolio $\boldsymbol{\theta}$ is optimal for the log utility investor and the associated belief $\mathbb{P}$, we get:

$$\mathbb{E}_i[f(R)] = \mathbb{E}^{\mathbb{Q}}[Rf(R)] , \tag{9}$$

where by construction $\mathbb{E}^{\mathbb{Q}}[R] = 1$. Given convex set Θ of admissible portfolio weights, the worst case expected payoff across admissible maximum growth portfolios is:

$$\mathcal{L}(f) := \inf_{\boldsymbol{\theta} \in \Theta} \mathbb{E}_i[f(R)] = \inf_{\boldsymbol{\theta} \in \Theta} \mathbb{E}^{\mathbb{Q}}[Rf(R)] . \tag{10}$$

Analogously, the best case expected payoff across admissible maximum growth portfolios is:

$$\mathcal{U}(f) := \sup_{\boldsymbol{\theta} \in \Theta} \mathbb{E}_i[f(R)] = \sup_{\boldsymbol{\theta} \in \Theta} \mathbb{E}^{\mathbb{Q}}[Rf(R)] . \tag{11}$$

This concludes the proof. ■

Corollary 1 (Upper and lower bounds on expected payoffs from observed investor's holding)**.** *In the context of Proposition 3, suppose that $\boldsymbol{\theta}_0^\star$ is the optimal portfolio of a log investor with belief $\mathbb{P}$, which is however not observable. Assume further that there exists an observable portfolio $\boldsymbol{\theta}_0$ such that*

$$d(\boldsymbol{\theta}_0, \boldsymbol{\theta}_0^\star) \leq \delta , \tag{12}$$

for some convex discrepancy function $d(\cdot,\cdot) \geq 0$ such that $d(x,y) = 0$ if and only if $x = y$. Then, the upper and lower bounds in Proposition 3 are such that:

$$\mathcal{L}(f) = \inf_{d(\boldsymbol{\theta},\boldsymbol{\theta}_0)\leq\delta} \mathbb{E}^{\mathbb{Q}}[Rf(R) \leq E_i[Rf(R)] \leq \mathcal{U}(f) = \sup_{d(\boldsymbol{\theta},\boldsymbol{\theta}_0)\leq\delta} \mathbb{E}^{\mathbb{Q}}[R(\boldsymbol{\theta})f(R)] \,. \tag{13}$$

In the case where $\delta = 0$, i.e., there is no portfolio measurement error, then

$$\mathcal{L}(f) = \mathbb{E}_i[f(R)] = \mathcal{U}(f) \,.$$

Example 1. If $d(\boldsymbol{\theta},\boldsymbol{\theta}_0) = \frac{1}{2}||\boldsymbol{\theta}-\boldsymbol{\theta}_0||_2^2$, then:

$$g_{\mathcal{L}(f)}(\lambda) = \inf_{\boldsymbol{\theta}\in\mathbb{R}^N} \left\{ \mathbb{E}^{\mathbb{Q}}[R(\boldsymbol{\theta})f(R)] + \lambda\left(\frac{1}{2}||\boldsymbol{\theta}-\boldsymbol{\theta}_0||_2^2 - \delta\right)\right\} \,. \tag{14}$$

This gives the optimality condition:

$$\mathbf{0} = \nabla g_{\mathcal{L}(f)}(\lambda) = \mathbb{E}^{\mathbb{Q}}[\boldsymbol{R}^e f(R)] + \lambda(\boldsymbol{\theta}-\boldsymbol{\theta}_0) \,, \tag{15}$$

and, whenever the constraint is binding:

$$\frac{1}{2}||\mathbb{E}^{\mathbb{Q}}[\boldsymbol{R}^e f(R)]||_2^2 = \frac{1}{2}\lambda^2||\boldsymbol{\theta}-\boldsymbol{\theta}_0||_2^2 = \lambda^2\delta \,, \tag{16}$$

i.e., an optimal Lagrange multiplier given by:

$$\lambda^\star = \frac{1}{\sqrt{2\delta}}||\mathbb{E}^{\mathbb{Q}}[\boldsymbol{R}^e f(R)]||_2 \,. \tag{17}$$

Therefore, the optimal portfolio supporting the lower bound is such that:

$$\boldsymbol{\theta}^\star = \boldsymbol{\theta}_0 - \frac{1}{\lambda^\star}\mathbb{E}^{\mathbb{Q}}[\boldsymbol{R}^e f(R)] = \boldsymbol{\theta}_0 - \frac{\sqrt{2\delta}}{||\mathbb{E}^{\mathbb{Q}}[\boldsymbol{R}^e f(R)]||_2}\mathbb{E}^{\mathbb{Q}}[\boldsymbol{R}^e f(R)] \,. \tag{18}$$

This gives the closed-form lower bound:

$$\mathcal{L}(f) = g_{\mathcal{L}(f)}(\lambda^\star) = \mathbb{E}^{\mathbb{Q}}[R(\boldsymbol{\theta}^\star)f(R)] = \mathbb{E}^{\mathbb{Q}}[f(R)] + \mathbb{E}^{\mathbb{Q}}[\boldsymbol{\theta}_0'\boldsymbol{R}^e f(R)] - \sqrt{2\delta}||\mathbb{E}^{\mathbb{Q}}[\boldsymbol{R}^e f(R)]||_2 \,. \tag{19}$$

In an analogous vein, we obtain:

$$\mathcal{U}(f) = g_{\mathcal{U}(f)}(\lambda^\star) = \mathbb{E}^{\mathbb{Q}}[f(R)] + \mathbb{E}^{\mathbb{Q}}[\boldsymbol{\theta}_0' \boldsymbol{R}^e f(R)] + \sqrt{2\delta}||\mathbb{E}^{\mathbb{Q}}[\boldsymbol{R}^e f(R)]||_2 \,. \tag{20}$$

Proposition 4 (Lower bound on expected log return of optimally invested wealth)**.** *Suppose $\boldsymbol{\theta} \in \Theta$ indexes a log investor holding an optimal portfolio $\boldsymbol{\theta}$, with return R, and having belief $\mathbb{P}$. Then, the following lower bound holds:*

$$\mathcal{L} := \inf_{\boldsymbol{\theta} \in \Theta} \mathbb{E}^{\mathbb{Q}}[R \log R] \leq \mathbb{E}_i[\log R] \,. \tag{21}$$

Proof. Since portfolio $\boldsymbol{\theta}$ is optimal for the log utility investor and the associated belief $\mathbb{P}$, we get:

$$\mathbb{E}_i[\log R] = \mathbb{E}^{\mathbb{Q}}[R \log R] \,, \tag{22}$$

where by construction $\mathbb{E}^{\mathbb{Q}}[R] = 1$. Given convex set Θ of admissible portfolio weights, the worst case expected log utility over maximum growth portfolios is:

$$\mathcal{L} := \inf_{\boldsymbol{\theta} \in \Theta} \mathbb{E}_i[\log R(\boldsymbol{\theta})] = \inf_{\boldsymbol{\theta} \in \Theta} \mathbb{E}^{\mathbb{Q}}[R(\boldsymbol{\theta}) \log R(\boldsymbol{\theta})] \,. \tag{23}$$

This problem is convex, with solution obtained using standard duality methods. This concludes the proof. ■

Corollary 2 (Lower bound extracted from observed investor's holdings)**.** *In the context of Proposition 4, suppose that $\boldsymbol{\theta}_0^\star$ is the optimal portfolio of a log investor with belief $\mathbb{P}$, which is however not observable. Assume further that there exists and observable portfolio $\boldsymbol{\theta}_0$ such that*

$$d(\boldsymbol{\theta}_0, \boldsymbol{\theta}_0^\star) \leq \delta \,, \tag{24}$$

for some convex discrepancy function $d(\cdot, \cdot) \geq 0$ such that $d(x, y) = 0$ if and only if $x = y$. Then, the lower bound in Proposition 4 becomes:

$$\mathcal{L} = \inf_{\boldsymbol{\theta}} \mathbb{E}^{\mathbb{Q}}[R \log R(\boldsymbol{\theta})] \quad \textit{s.t.} \quad d(\boldsymbol{\theta}, \boldsymbol{\theta}_0) \leq \delta \,. \tag{25}$$

In the case where $\delta = 0$, i.e., there is no portfolio measurement error, then

$$\mathcal{L} = \mathbb{E}_i[\log R(\boldsymbol{\theta}_0^\star)] \,.$$

Proof. The lower bound follows from Proposition 4 once we define $\Theta := \{\boldsymbol{\theta} \mid d(\boldsymbol{\theta}, \boldsymbol{\theta}_0) \leq \delta\}$. In the case where there is no measurement error, $\delta = 0$ and $\Theta = \{\boldsymbol{\theta}_0^\star\}$, i.e.:

$$\mathcal{L} = \mathbb{E}^{\mathbb{Q}}[R \log R] = \mathbb{E}_i[\log R] \, . \tag{26}$$

This concludes the proof. ■

Corollary 3 (Dual formulation)**.** *In the context of Proposition 4, for any $\lambda \geq 0$ it follows:*

$$\mathcal{L} \geq g(\lambda) := \inf_{\boldsymbol{\theta} \in \mathbb{R}^N} \left\{ \mathbb{E}^{\mathbb{Q}}[R(\boldsymbol{\theta}) \log R(\boldsymbol{\theta})] + \lambda(d(\boldsymbol{\theta}, \boldsymbol{\theta}_0) - \delta) \right\} \, . \tag{27}$$

Therefore, $\mathcal{L} \geq \sup_{\lambda \geq 0} g(\lambda)$. Moreover, when suitable Constraints Qualification conditions hold, then $\mathcal{L} = \sup_{\lambda \geq 0} g(\lambda)$. In particular, if there exists $0 < \delta' < \delta$ such that $\mathbb{E}^{\mathbb{Q}}[R(\boldsymbol{\theta}) \log R(\boldsymbol{\theta})] < \infty$ for all $\boldsymbol{\theta}$ such that $d(\boldsymbol{\theta}, \boldsymbol{\theta}_0) < \delta'$ then Slater's Constraint Qualification conditions hold.

Proof. The proof follows with standard Lagrangian duality arguments. ■

Example 2. If $d(\boldsymbol{\theta}, \boldsymbol{\theta}_0) = \frac{1}{2}||\boldsymbol{\theta} - \boldsymbol{\theta}_0||_2^2$, then:

$$g(\lambda) = \inf_{\boldsymbol{\theta} \in \mathbb{R}^N} \left\{ \mathbb{E}^{\mathbb{Q}}[R(\boldsymbol{\theta}) \log R(\boldsymbol{\theta})] + \lambda \left(\frac{1}{2}||\boldsymbol{\theta} - \boldsymbol{\theta}_0||_2^2 \right) \right\} \, . \tag{28}$$

References

Acciaio, B., M. Beiglböck, F. Penkner, and W. Schachermayer. 2016. "A Model-Free Version of the Fundamental Theorem of Asset Pricing and the Super-Replication Theorem." *Mathematical Finance* 26 (2): 233–251.

Almeida, Caio, and Gustavo Freire. 2022. "Demand in the Option Market and the Pricing Kernel." *Working Paper, Princeton University*.

Baltussen, Guido, Julian Jerstegge, and Paul Whelan. 2024. "The Derivative Payoff Bias." *Working Paper, Erasmus University*.

Borovička, Jaroslav, Lars Peter Hansen, and José A Scheinkman. 2016. "Misspecified Recovery." *Journal of Finance* 71 (6): 2493–2544.

Breeden, Douglas T, and Robert H Litzenberger. 1978. "Prices of state-contingent claims implicit in option prices." *Journal of Business*: 621–651.

Carr, Peter, and Dilip Madan. 2001. "Optimal positioning in derivative securities."

Chen, Hui, Scott Joslin, and Sophie Xiaoyan Ni. 2019. "Demand for Crash Insurance, Intermediary Constraints, and Risk Premia in Financial Markets." *Review of Financial Studies* 32 (1): 228–265.

Chen, Xiaohong, Lars Peter Hansen, and Peter G Hansen. 2020. "Robust Identification of Investor Beliefs." *Proceedings of the National Academy of Sciences* 117 (52): 33130–33140.

Dahlquist, Magnus, and Markus Ibert. 2023. "Equity Return Expectations and Portfolios: Evidence from Large Asset Managers." *Working Paper, Stockholm School of Economics*.

Farago, Adam, Mariana Khapko, and Chayawat Ornthanalai. 2021. "Asymmetries and the market for put options." *Swedish House of Finance Research Paper* (21-12).

Gandhi, Mihir, Niels Joachim Gormsen, and Eben Lazarus. 2023. "Forward Return Expectations." *Working Paper, University of Chicago*.

Gao, Can, and Ian Martin. 2021. "Volatility, valuation ratios, and bubbles: An empirical measure of market sentiment." *Journal of Finance* 76 (6): 3211–3254.

Gârleanu, Nicolae, Lasse Heje Pedersen, and Allen M Poteshman. 2008. "Demand-based Option Pricing." *Review of Financial Studies* 22 (10): 4259–4299.

Gayda, Leander, Thomas Grünthaler, and Jan Harren. 2023. "Option Liquidity and Gamma Imbalances." *Working Paper, University of Münster*.

Ghosh, Anisha, Arthur Korteweg, and Qing Xu. 2022. "Recovering Heterogeneous Beliefs and Preferences from Asset Prices." *Working Paper, McGill University*.

Ghosh, Anisha, and Guillaume Roussellet. 2023. "Identifying Beliefs from Asset Prices." *Working Paper, McGill University*.

Giglio, Stefano, Matteo Maggiori, Johannes Stroebel, and Stephen Utkus. 2021. "Five Facts About Beliefs and Portfolios." *American Economic Review* 111 (5): 1481–1521.

Gormsen, Niels Joachim, and Christian Skov Jensen. 2022. "Higher-Moment Risk." *Working Paper, University of Chicago*.

Greenwood, Robin, and Andrei Shleifer. 2014. "Expectations of Returns and Expected Returns." *Review of Financial Studies* 27 (3): 714–746.

Jensen, Christian Skov, David Lando, and Lasse Heje Pedersen. 2019. "Generalized Recovery."

Journal of Financial Economics 133 (1): 154–174.

Johnson, Timothy, Mo Liang, and Yan Liu. 2016. "What Drives Index Options Exposures?" *Review of Finance*: 1–33.

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

Korsaye, Sofonias Alemu. 2024. "Investor Beliefs and Market Frictions." *Working Paper, Johns Hopkins University*.

Long, John B. 1990. "The Numeraire Portfolio." *Journal of Financial Economics* 26 (1): 29–69.

Martin, Ian. 2017. "What is the Expected Return on the Market?" *Quarterly Journal of Economics* 132 (1): 367–433.

Martin, Ian, and Christian Wagner. 2019. "What is the Expected Return on a Stock?" *Journal of Finance* 74 (4): 1887–1929.

Meeuwis, Maarten, Jonathan A. Parker, Antoinette Schoar, and Simester Duncan. 2022. "Belief Disagreement and Portfolio Choice." *Journal of Finance* 77 (6): 3191–3247.

Nagel, Stefan, and Zhengyang Xu. 2023. "Dynamics of Subjective Risk Premia." *Journal of Financial Economics* 150 (2): 1–24.

Ross, Steve. 2015. "The Recovery Theorem." *Journal of Finance* 70 (2): 615–648.

Tetlock, Paul C. 2023. "The Implied Equity Premium." *Working Paper, Columbia Business School*.