Demand-Based 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 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. Using buy and sell orders on S&P500 options, we show that the expected returns inferred from retail and institutional investor beliefs increase in bad times when they become the net suppliers of crash insurance in option markets, mirroring price-based estimates. Market makers' 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, Alessandro Melone (discussant), Ruslan Tuneshev (discussant), participants at Indiana University, Purdue University, Boston University, 2024 LUISS Finance Workshop, and 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 ignore information on investors' holdings. In this paper, we argue that including information about investors' portfolios is crucial to understanding the dynamics of subjective expected returns.[1]

To this end, we propose a theoretical framework for recovering beliefs/expected returns for heterogeneous investors from prices and holdings data jointly under the weak assumption of no-arbitrage. More specifically, we theoretically show that investors' subjective expected return on the market can be directly inferred from option prices and their option holdings 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 orders on S&P500 index options, we show that the subjective expected return of financial intermediaries' drops in times of distress, contrary to price-based measures. In normal times, financial intermediaries are net suppliers of deep-out-of-the-money puts (crash insurance) to public investors, see, e.g., Gârleanu, Pedersen, and Poteshman [2008] and Chen, Joslin, and Ni [2019]. During crises times, however, when financial intermediaries' constraints bind, public investors provide crash insurance to intermediaries. As a consequence, estimates of customers' subjective expected returns increase during bad times. 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 can 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, or open interest) 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

[1]For example, recent empirical evidence, documents a strong relationship between investors' expected returns and portfolio holdings using surveys, see Giglio et al. [2021]

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 $\mathbb{P}$ can be directly inferred from investors' holdings and option prices. As is well-known, the risk-neutral pricing measure $\mathbb{Q}$ can be fully determined by observed option prices under no arbitrage (Breeden and Litzenberger [1978]). Since holdings are observable at the investor level and payoffs 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 SDFs that are functions of the index and the options returns. Therefore, the ensuing expected market returns may be less or more volatile depending on portfolio composition, with sign and cyclicality properties that depend 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 or monotonically increasing. 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 deep-out-of-the-money puts corresponds to conservative investors that are progressively more sensitive to higher-order risk factors and trade options to reallocate them profitably. Moreover, we show that 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 merge option price information with buy and sell orders of large investors in option markets.

More specifically, we use our results to gain insights about the beliefs and subjective expected returns 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 most closely related to the literature that makes use of asset prices to recover measures of the expected return. Even tough these papers do not explicitly recover heterogeneous investors' beliefs, some of 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 expected return of an asset can be inferred from the allocation of a growth-optimal portfolio that maximizes an investor's long-run growth. Expected returns are shown to be functions of risk-neutral variance. **?** use a Taylor series expansion of the inverse of the marginal utility to construct lower and upper bounds on the conditional expected excess market return that are functions of higher-order risk-neutral simple return moments. 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 of a log investor is the reciprocal of a combination between the market return and the return of a portfolio of higher-order (risk-neutral) 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 dynamics and statistical properties of measures of expected returns crucially depend on the weights allocated to the basis assets. While Martin [2017] assumes that investors choose to hold 100% of their wealth in the market (and none in the derivatives themselves), Tetlock [2023] allows for holdings in both the market and power contracts on the market. In our setting, we do not need to make any assumptions about the redundancy of option markets and optimal weights since 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 recovering subjective 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 subjective expected returns from holdings and price 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.[2] 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 supported by 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 85% in the index and 15% in calls or puts with the same maturity but different strikes[3]. As is immediately evident, even for a small invest-

[2]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.

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

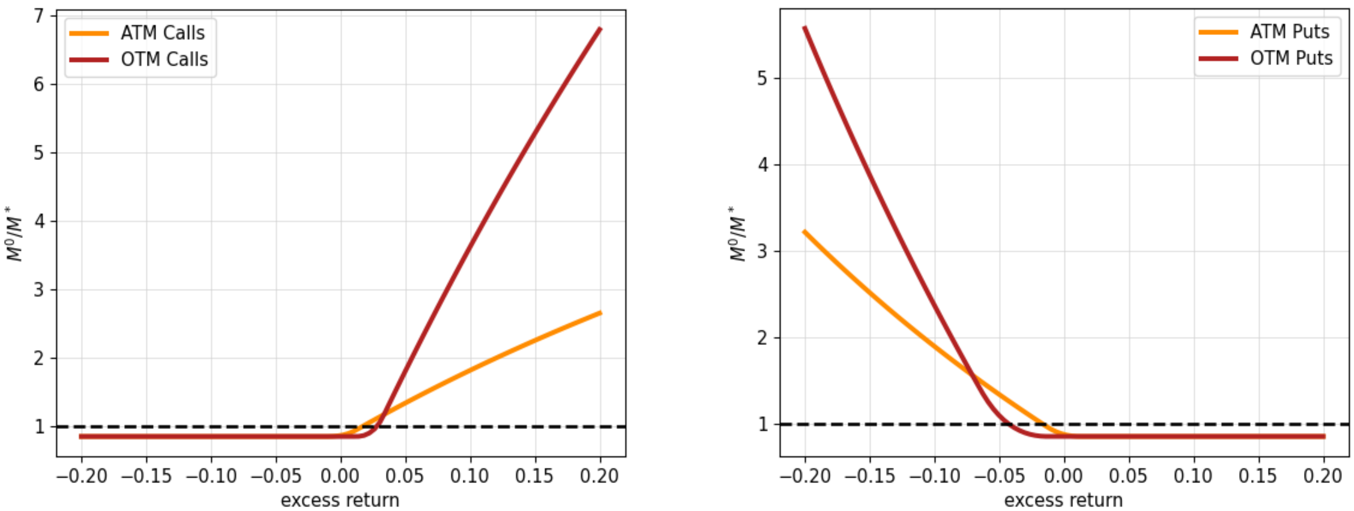

Figure 1. Ratio between Benchmark SDF M^0 and $M^\star$

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 85% of the wealth in the underlying and 15% in an equally-weighted portfolio of calls (left plot) and puts (right plot) with different moneyness. ATM options have $|\Delta| \in (0.375, 0.625)$. OTM options have $|\Delta| \in [0.125, 0.375]$.

ment in options (with respect to the investment in the underlying) and small fluctuations of the market return, the ratio $M^0/M^\star$ is widely different from 1. $M^0/M^\star$ rises significantly for at-the-money (ATM) and especially out-of-the-money (OTM) calls and puts. For example, assuming the market excess return is +20% (-20%), the ratio increases to 7 (5.5) for OTM calls (puts).

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 subjective expected returns. In Figure 2, we plot the time-series of the expected market return for different $M^\star$ implied by options (the same as in Figure 1) 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 hold call options in addition to the index. As we can see, the patterns mirror the benchmark case almost one-for-one. Expected returns increase in bad time, decrease in normal periods and are highly volatile.

The value is considerably higher even with a small investment in the option. Intuitively, the size of the expected return 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 expected returns 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 expected return is lower. In fact, the expected return 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. Buying puts acts as protection mechanism which reflects a pessimistic view and negative risk premium. Accordingly, the volatility of the expected return is higher with respect to what we recover under M^0. The contribution of the puts may be enough to lead to a pro-cyclical pattern. Given that, we conclude that the size and cyclicality properties of expected returns depend on the moneyness, and on the amount of wealth invested in the options.

While instructive, the above examples maybe too stylized. To study a more realistic setting, we now showcase two popular option strategies: collars and straddles. For example, some investors are known to hold the underlying and add protection via longing puts and shorting calls (collar); other investors bet on the underlying volatility by taking long positions in calls and puts (straddle strategy). In Figure 3, we plot the ratio $M^0/M^\star$ for a hypothetical collar and straddle with OTM options, separately in the case with (left panel) and without (right panel) delta-hedging. In the first, the functional form aligns with our previous example: being long in calls/puts increases the distortion in the tails, the reverse is true for short positions. When the linear component is removed by delta-hedging, the effect is significantly mitigated and the relative leverage between the options and the index decreases. In general, the implication of asset demand on the expected quantities is not trivial, as it depends on the composition of the portfolio and on the initial probability distribution.

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

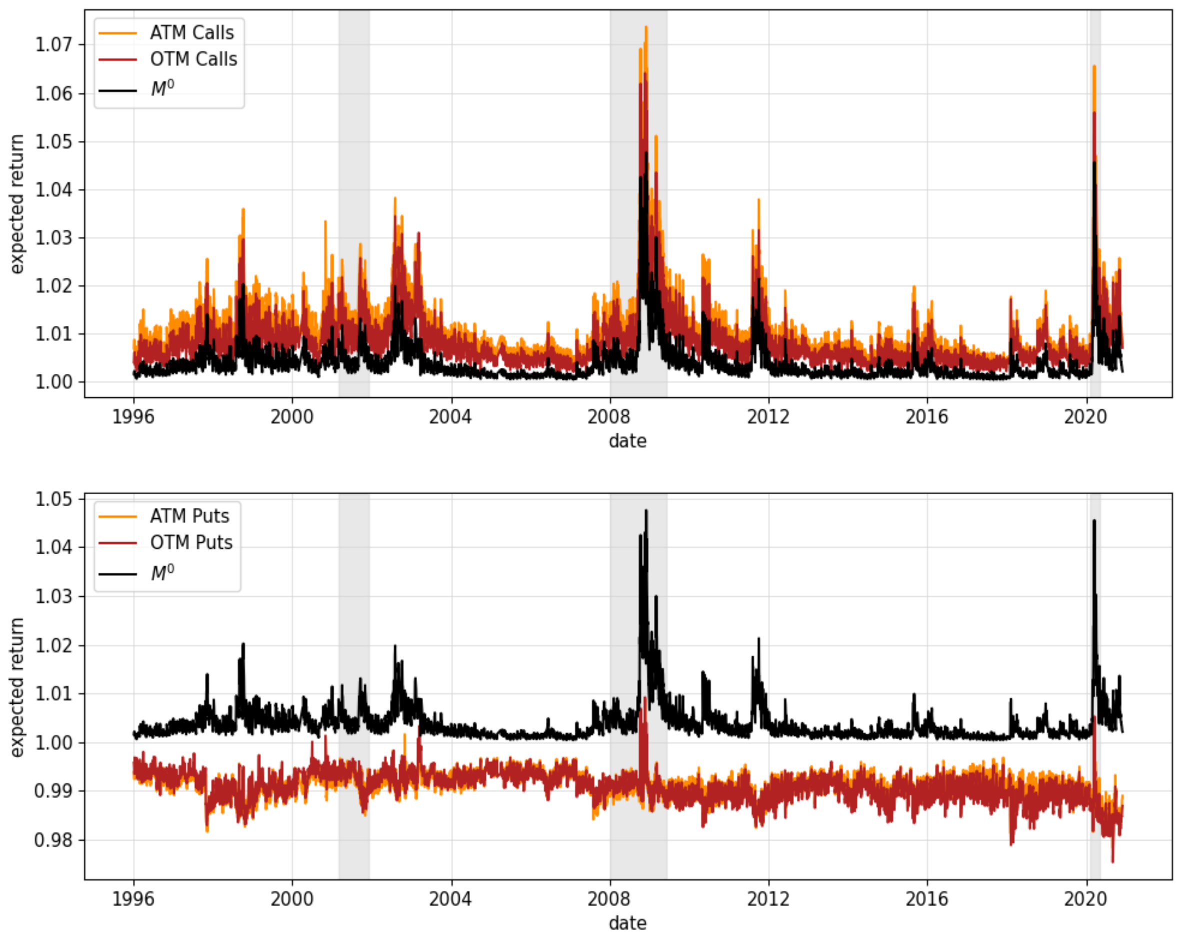

Figure 2. Monthly Time-Series of Expected Market Return
Notes: This figure plots the expected market return recovered from different stochastic discount factors. Data are daily and the horizon is monthly. The black line is the expected return recovered by $M^0 = 1/R$, where R is 100% invested in the index. The other lines arise from $M^\star = 1/R^\star$, where $R^\star$ is the return of a portfolio investing 85% in the index and 15% in an equally-weighted portfolio of ATM (orange line) or OTM (red line) options. The upper panel is with calls, the lower with puts. Grey areas indicate NBER recession periods.

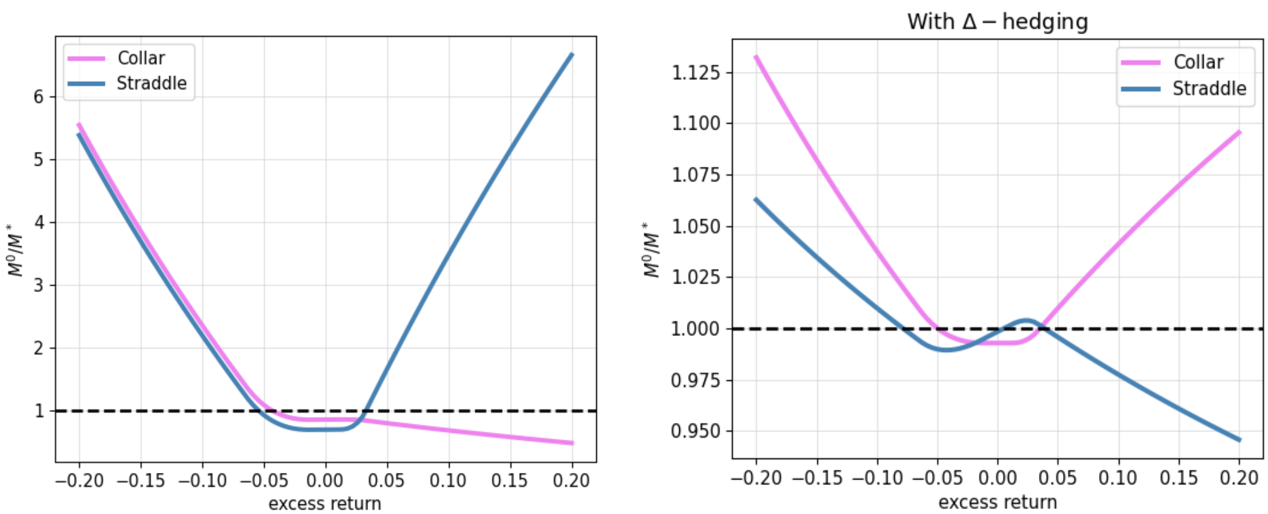

Figure 3. Ratio between benchmark M^0 and $M^\star$

Notes: This figure plots $M^0/M^\star$ as function of the excess return on the market. $M^0 = 1/R$ is fully invested in the index. The «collar» is long in the index and in OTM puts, and short in OTM calls. The «straddle» is long in OTM calls and OTM puts. In the right panel, the portfolios supporting $M^\star$ are added with an extra investment in the index so as to keep the option component exactly delta-hedged.

3 Theoretical Framework

We now present a simple theoretical framework to explain how to recover subjective expected returns from options data. Consider an investor, labeled i, with logarithmic preferences who has access to 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. Let $\mathbb{E}_i[\cdot]$ denote the subjective expectation of this investor over possible states of the world. The investor's subjective beliefs may or may not coincide with the true underlying data-generating process. 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 $\boldsymbol{\theta}_i$ are the portfolio weights in the market index and the options by investor i, see, e.g., Long [1990]. SDF M_i is the reciprocal of the return of the growth-optimal portfolio since it maximizes expected long-run growth of the investor i's wealth.

No arbitrage implies that

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

We can now define a change of measure, $\frac{d\mathbb{Q}}{d\mathbb{P}_i} = M_i$. The subjective expected return of the market for investor i under the physical measure can hence be written as:

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

Equation (2) is the main identity studied in this paper. It relates investor-specific physical beliefs to risk-neutral expectations about prices of assets that can be traded and investor i's holdings. 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 intuition for this is at least twofold. First, we only observe a subset of the "true" portfolio of investors. For example, while we observe the open and close orders on calls and puts for the S&P500 (SPX), we do not observe the holdings neither the transactions on other derivatives with the same underlying, such as SPY options (that is on the ETF tracking the S&P500).[4] Since major market makers provide liquidity in both SPX and SPY option markets, we only observe a fraction of their true market exposure.[5] Second, our data is sampled at high frequency (every 30 minutes), however, we aggregate to the monthly frequency and across different types of customers (retail and institutional) to calculate expected returns. This aggregation will likely lead to further measurement error affecting our estimates. Given this, we consider a second case where we assume that portfolio holdings are observed with measurement error leading to bounds on the subjective expected returns representing the most conservative and maximum value.

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

Given this, 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.

Several remarks are in order. Our set-up is similar to Martin [2017], who derives a lower bound on the expected return assuming an unconstrained rational investor who's risk aversion is at least one. The crucial difference to Martin [2017] is that he imposes that the optimal portfolio held by the investor consists of 100% in the market return. A consequence of this assumption is that all options are redundant. There are at least two reasons why this assumption seems too restrictive. First, several papers show that options are non-redundant assets since they allow to hedge crash risk and demand for OTM puts options is significant. Second, in the data, we find that the SDF defined as the inverse of the market return induces significant pricing errors when pricing options. For example, we find that for the 1996 to 2020 period, the average pricing error is: 29% for OTM puts, 13% for ATM puts, 8% for OTM calls, and 7% for ATM calls, respectively.

In our paper, we do not assume such redundancy. However, notice that our setting nests Martin [2017]'s case if one assumes a log investor and θ equals one for the market and zero

[4]While SPX options trade exclusively on the CBOE, SPY options trade across several exchanges. Institutional public investors mainly trade SPX options (due to larger contract sizes, tax treatment, etc.) while retail investors mostly trade in SPY options.

[5]Moussawi, Xu, and Zhou [2024] show that at least four market makers provide liquidity in both markets simultaneously: Susquehanna Securities, Citadel Securities, Wolverine Trading, and IMC Financial Markets.

otherwise. 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 $\mathbb{Q}$.[6]

Another related paper is Tetlock [2023] who assumes that the investor can hold (integer) power-contracts written on the market index. The paper again makes an assumption about the redundancy of certain option contracts such as non-integer power contracts. In general, however, we do not observe holdings of power-contracts. To circumvent this issue, Tetlock [2023] has to estimate the weights from expanding window regressions predicting risk-neutral (power) moments with physical counterparts. Our approach is straightforward since we directly observe holdings on plain vanilla calls and puts.

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 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 1 (Bounds on Subjective Expected Returns)**.** *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:*

$$\mathbb{E}_i[R] \geq \mathbb{E}^{\mathbb{Q}}[(1 + \boldsymbol{\theta}' \mathbf{R}^e) R] - \sqrt{2\delta} ||\mathbb{E}^{\mathbb{Q}}[\boldsymbol{R}^e R]||_2 \, . \tag{5}$$

[6] If the investor has log utility, Martin [2017]'s lower bound becomes an exact identity since $\text{cov}(M_i R, R) = 0$ in that case.

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

$$\mathbb{E}_i[R] \leq \mathbb{E}^{\mathbb{Q}}[(1+\boldsymbol{\theta}'\mathbf{R}^e)R] + \sqrt{2\delta}||\mathbb{E}^{\mathbb{Q}}[\boldsymbol{R}^e R]||_2\,. \tag{6}$$

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 separately for type of position (opening/closing) and origin: (i) customer; (ii) brokers-dealer; (iii) firm; and (iv) market maker. As common practice, we aggregate these daily volumes to cumulative positions for the last three categories[7] and label them "market makers". Customers include retail and institutional investors. Our data starts in January 1996 and ends in December 2020. The label "broker-dealer" is available only from 2011 and accounts for less than 3% of the trades.

The volume data comes without pricing information. To this end, we obtain end-of-day bid-ask prices from the OptionMetrics database and use best closing bid- and ask-prices to compute mid-point prices. We merge the CBOE Open-Close database with price data and apply standard filters from Bakshi, Cao, and Chen [1997]. That is, we remove option contracts: with price less than \$3/8; expiring in less than 5 days or more than 475 days; with implied volatility smaller than 0.1% or greater than 1; with bid price exceeding ask price; with relative bid-ask spread larger than 1/2; that are traded for less than 5 units. We also filter out events where the sum of the transactions across investors is not zero, which happens for less than 2% of the times. Non-arbitrage filters apply as well.

Throughout our empirical analysis, we use daily frequency and monthly horizons. On each

[7]The CBOE Regulatory Circular defines firms as «OCC clearing member firm proprietary accounts». Thus we aggregate firms and brokers-dealers with market makers because they mainly trade against public customers, although they are not designated as intermediaries.

Table 1. Descriptive statistics options data.

	K/S_t		Customers' net demand		Maturity (days)	
Calls	All	OTM	All	OTM	All	OTM
mean	1.02	1.06	-2,127	-625	59.9	59.6
std. dev.	0.06	0.05	11,127	10,980	76.3	71.7
min.	0.08	1.00	-112,908	-123,146	2	2
median	1.02	1.04	-1,193	-421	33	36
max.	2.03	2.03	61,835	71,017	473	473
Puts						
mean	0.92	0.89	2,093	795	63.9	64.0
std. dev.	0.11	0.10	14,689	15,208	77.8	76.5
min.	0.10	0.10	-140,275	-140,639	2	2
median	0.94	0.92	1708	716	36	37
max.	2.72	1.02	92,694	92,408	473	473

Notes: This table reports summary statistics for the options data. The relative moneyness and the maturity are computed over the single option contracts that are traded on every date. Customers' net demand is instead aggregated over all the options traded on a single day; net demand is defined as the total opening/closing buy orders minus sell orders. For each variable, the first column refers to the full dataset, while the second to OTM options only. Data runs from January 1996 to December 2020.

date t, separately for calls and puts, we linearly interpolate[8] options volatility, options delta and investors' trading orders, on a grid of strike prices, for the required maturity (30 days if not explicitly stated). The grid consists of n uniformly distributed values between the smallest and the largest available strike in t, where n is the number of calls/puts actively traded in t.

Portfolios investing in the risk-free asset, S&P 500, and OTM options with the same maturity 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].

Table 1 reports summary statistics on the option sample for customers' net demand for calls (top panel) and puts (lower panel). Net demand is defined as total buy minus total sell orders aggregated over a day. On average, customers are net sellers for call options and net buyers for put options with an average maturity of 60 days. Traded options on average are out-of-the-money.

[8]When extrapolation is required, we look for the nearest value outside the convex hull in strikes and maturities.

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 let's define $f(R) := (1+\boldsymbol{\theta}'\mathbf{R}^e)R$, then the subjective expected return is given by:

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

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

Constructing Option Holdings

Subjective expected returns depend on portfolio holdings, θ_i, where θ_i are the total portfolio holdings of investor i. This distinguishes our approach from the demand-based literature that exploits opening positions at time t, flows–daily buy and sell orders. Given that, we need a measure of total portfolio holdings. To this end, notice that our database reports the total daily opening and closing positions on every option contract. Opening trading orders represent shocks to demand (flows), while aggregate opening and closing positions represent changes to the holdings. In order to get a measure of investors' holdings, we aggregate opening and closing positions from the issuance of a contract until date t.[9]

Figure B.1 and Figure B.2 in the Appendix compare our holdings proxy with flows for customers, monthly aggregated OTM options. While both measures exhibit a similar pattern before 2008, they diverge after the crisis. While flows remain negative almost everywhere before 2011, then they become progressively larger and positive until a new drop between 2019 and 2020. On the contrary, actual holdings tend to be positive and large after 2010, then they remain highly volatile but with a relatively stable positive average, and they stay positive during Covid as well. This is even more evident for puts that feature larger trading volume. The correlation between the two measures is 43% for calls and 10% for puts.

The time-series properties of option flows line up with the intuition of constrained arbitrageurs during the financial crisis when customers became the net sellers, see, e.g., Chen, Joslin, and Ni [2019]. However, this pattern is much more muted when looking at aggregate portfolios. For instance, while we notice a similar patterns during 2008, this did not repeat

[9]The CBOE Open-Close Database explicitly assigns a unique identification number to every option contract, which can be used to track it day-by-day from issuance to expiration.

during Cvoid. Figure B.3 in Appendix shows the cross-section of monthly demand changes (opening plus closing orders in t) and holdings (as before). Especially for puts, it is quite common to have negative changes in the option demand that are not large enough to offset the sign of the customers' portfolio, whose holdings remain positive.

4.3 Subjective Expected Returns with No Measurement Error

We now have all the ingredients to calculate subjective expected returns for the two investor types. We again stress that our theoretical framework does not allow us to recover the "true" beliefs of a particular investor but we interpret our measure a possible lower bound on subjective expectations. We start from our main identity (2): the variation of expected returns over time as perceived by different agents, assuming that the portfolio weights are observed without error.

We suppose that investors hold a delta-hedged portfolio of OTM options while the remainder of the wealth is invested in the index.[10]

Figure 4 plots expected returns for customers and market makers over time. Figure 5 compares the results with the benchmark case. There are several interesting observations. First we notice strikingly different dynamics of the expected return measures across investors but also distinct from the benchmark case both in size and in comovement. For example, customers' expected returns are less volatile than market makers (0.32% vs. 0.55%), and they can be negative more frequently. Also, customers feature higher expected returns in bad times, especially during Covid (up to 6%) and around the 2001 dot-com bubble, but not during the 2008 Great Financial Crisis, where on some dates, market makers' expected returns are higher. The correlation between the two time-series is 16%.

Second, neither of the two series correlate perfectly with the benchmark: 46% for customers and 59% for market makers, respectively. The expected return recovered through M^0 is by construction counter-cyclical only. Such correlations suggest that the cyclicality of subjective expected returns is contingent on the state of the economy.[11]

Third, subjective expected returns are also more volatile and less persistent than those from M^0. Our measure of $\mathbb{E}_i[R]$ is noisy because it absorbs shocks both to option prices and to portfolio rebalancing (i.e. option demand). This second effect is absent in $\mathbb{E}_0[R]$. The daily autocorrelation of the M^0 time-series stays significantly above 50% for at least 40 days, while it drops below 50% after one day for customers (two days for market makers).

[10]To this end, we re-scale the portfolio such that the unconditional NCC holds with the minimum achievable gap. In cases when the SDF is not strictly positive, we remove this observation.

[11]Note that the effect is partially damped by the additional time-varying investment in the index (due to delta-hedging). For instance, the time-series recovered from the same portfolios but without the delta-hedging component would get 6% correlation for customers and 8% for intermediaries.

From Figure 5, we notice that customers' expected returns are mostly below the case of Martin [2017]. Market makers' are frequently above. 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 studying investors' SDFs.

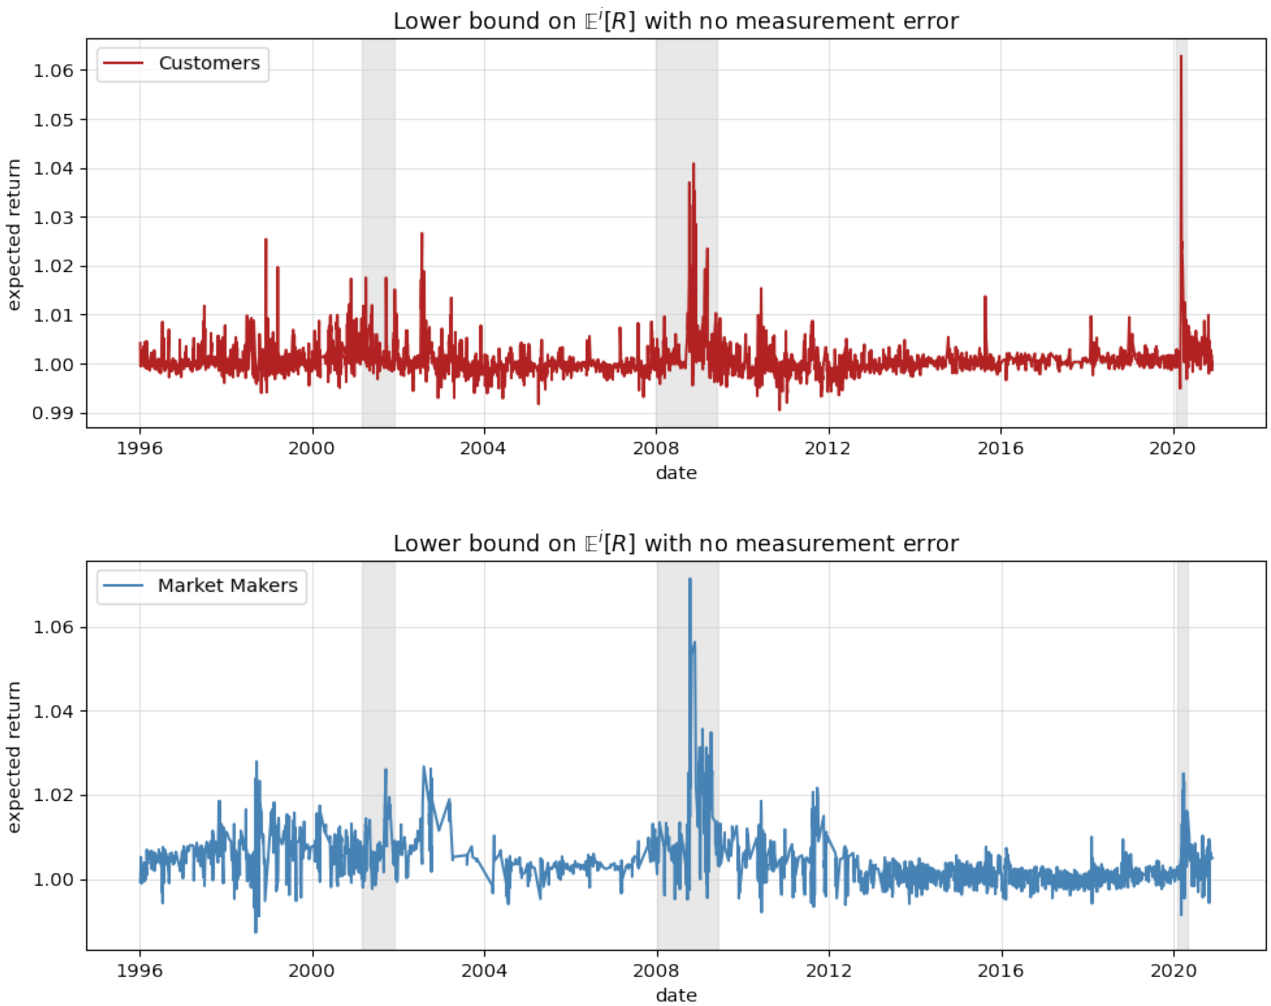

Figure 4. Subjective Expected Market Return with no measurement error
Notes: This figure plots the time-series of the expected market return as recovered through the SDF supported by observed customers' (top) and market makers' (bottom) positions, as described in the main text. Frequency is daily, horizon is monthly. Gray bars indicate NBER recessions. When investors' portfolio return can be negative in its support, the event is excluded.

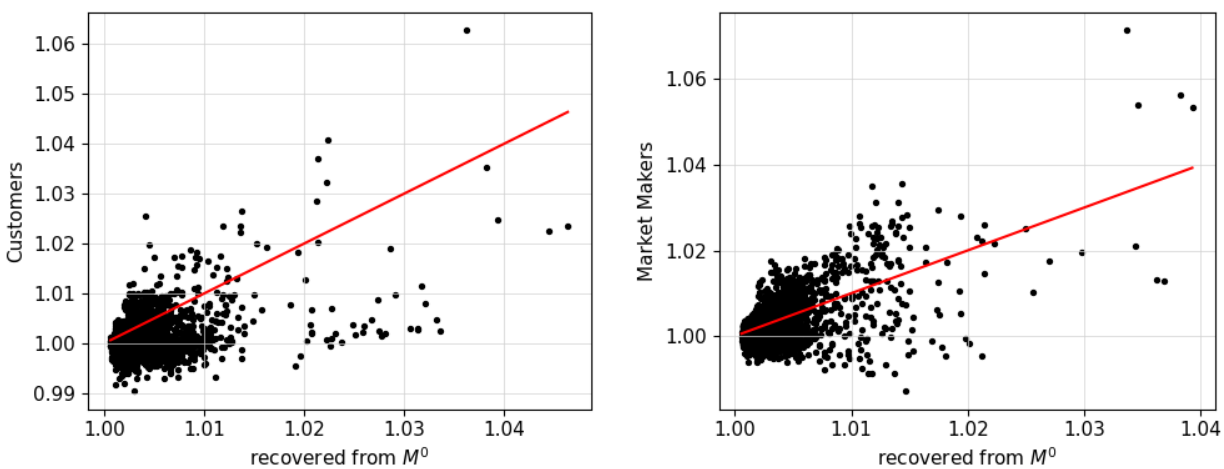

Figure 5. Subjective Expected Market Return with no measurement error vs. Benchmark
Notes: This figure plots the monthly expected return recovered by the benchmark M^0 vs. the subjective expected return for customers (left) and market makers (right). Each point corresponds to a single day measurement. The red line would be for $\mathbb{E}_i[R] = \mathbb{E}^0[R]$.

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) to determine the most conservative and the maximum expected market return compatible with investors' observed positions. δ is equal to half of the average bid-ask spread in the options cross-section at every date t. We also require the total leverage on the options to be not higher than what observed in $\boldsymbol{\theta}$, and the option component to be still delta-hedged[12].

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.

Before 2004, the lower bound is close to those recovered in Figure 4. Between 2008 and 2020, customers' lower bound is widely pro-cyclical (clearly between 2009 and 2012) with two huge negative peaks during the Great Financial Crisis (-3.2%) and Covid (-5.2%). Such behavior is almost everywhere in contrast to the bound extracted in absence of measurement errors. Conversely, market makers' lower bound is more aligned to Figure 4 in sign and cyclicality; it is just smaller in size, being often negative, up to the minimum value of -4.4% during Covid. This very negative risk premium is intuitive since in that case, market makers hold a portfolio of crash insurance, meaning they are protected against market downturns.

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 degree of belief heterogeneity in the market results larger during times of increasing volatility. This explains why the lower bound is largely pro-cyclical (correlation with the M^0 time-series is -23% for customers and 15% for market makers) and the upper bound is mostly counter-cyclical (correlations 73% for customers and 82% for market makers).

Our earlier results show that option portfolios can mitigate exposure to market risk. Due to that, we find that optimized portfolio weights tend to assign larger mass on the option component of the portfolio (compatibly with the delta-hedging). As observed weights are

[12] In the analytic solution in Proposition 1, the lower bound and the upper bound are symmetric with respect to a time-varying middle point $\mathbb{E}_t^{\mathbb{Q}}[(1+\boldsymbol{\theta}'\boldsymbol{R}^e)R]$. Here, the leverage constraint $\|\boldsymbol{\theta}^\star\|_1 \leq \ell$ makes not always possible to find solutions with such features.

the opposite between customers and market makers, the lower bound for each investor is negatively correlated (-29%) with the other, while the correlation is slightly positive (+16%) without measurement errors.

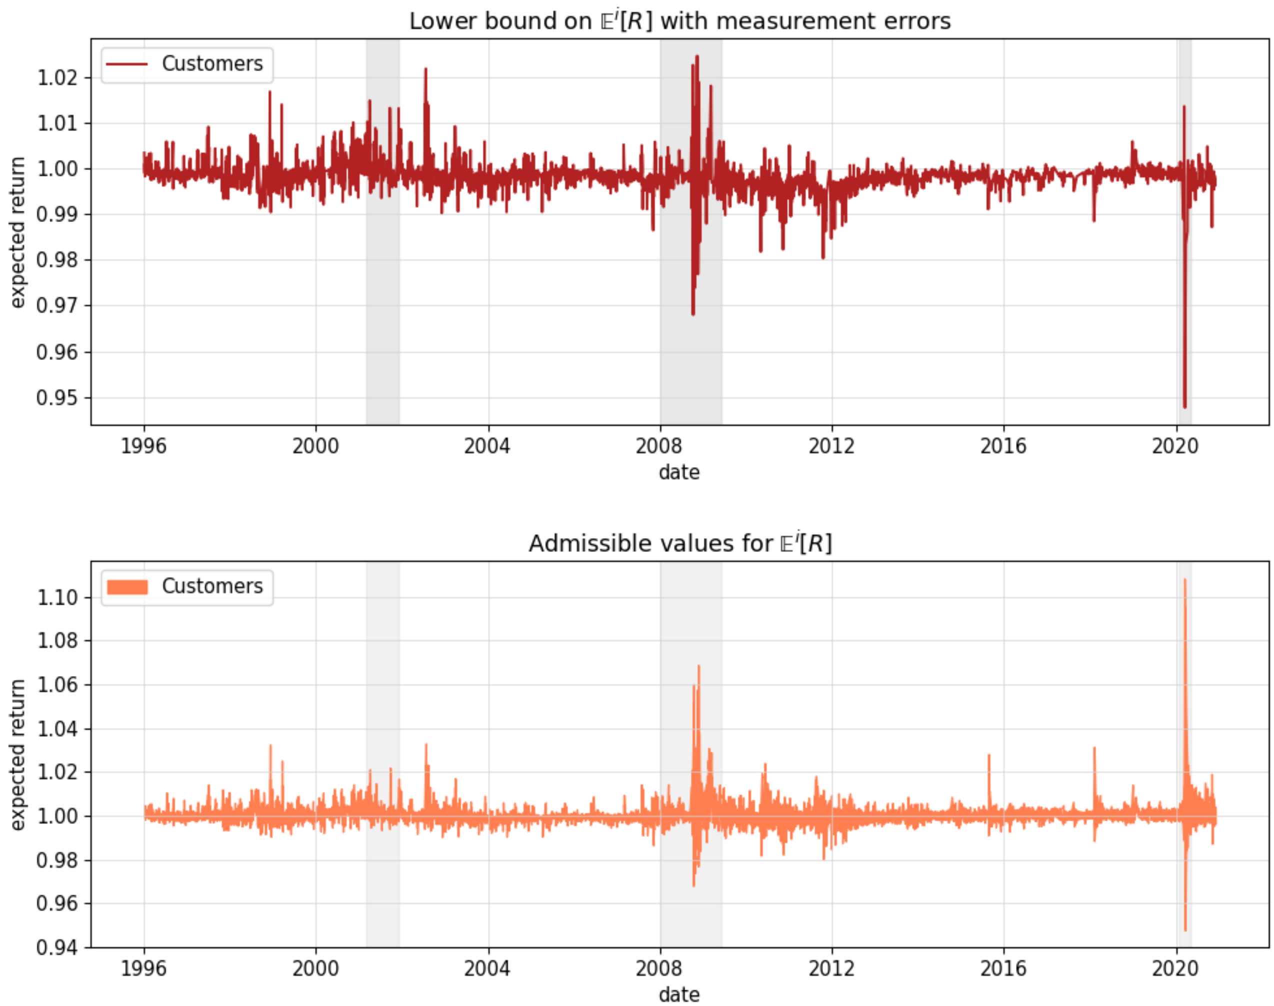

Figure 6. Bounds for Subjective Expected Market Return Customers
Notes: This figure plots the time-series of the lower bound on the subjective expected market return (upper plot) and all the possible values for the expected market return (lower plot) as recovered by SDFs compatible with customers' positions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

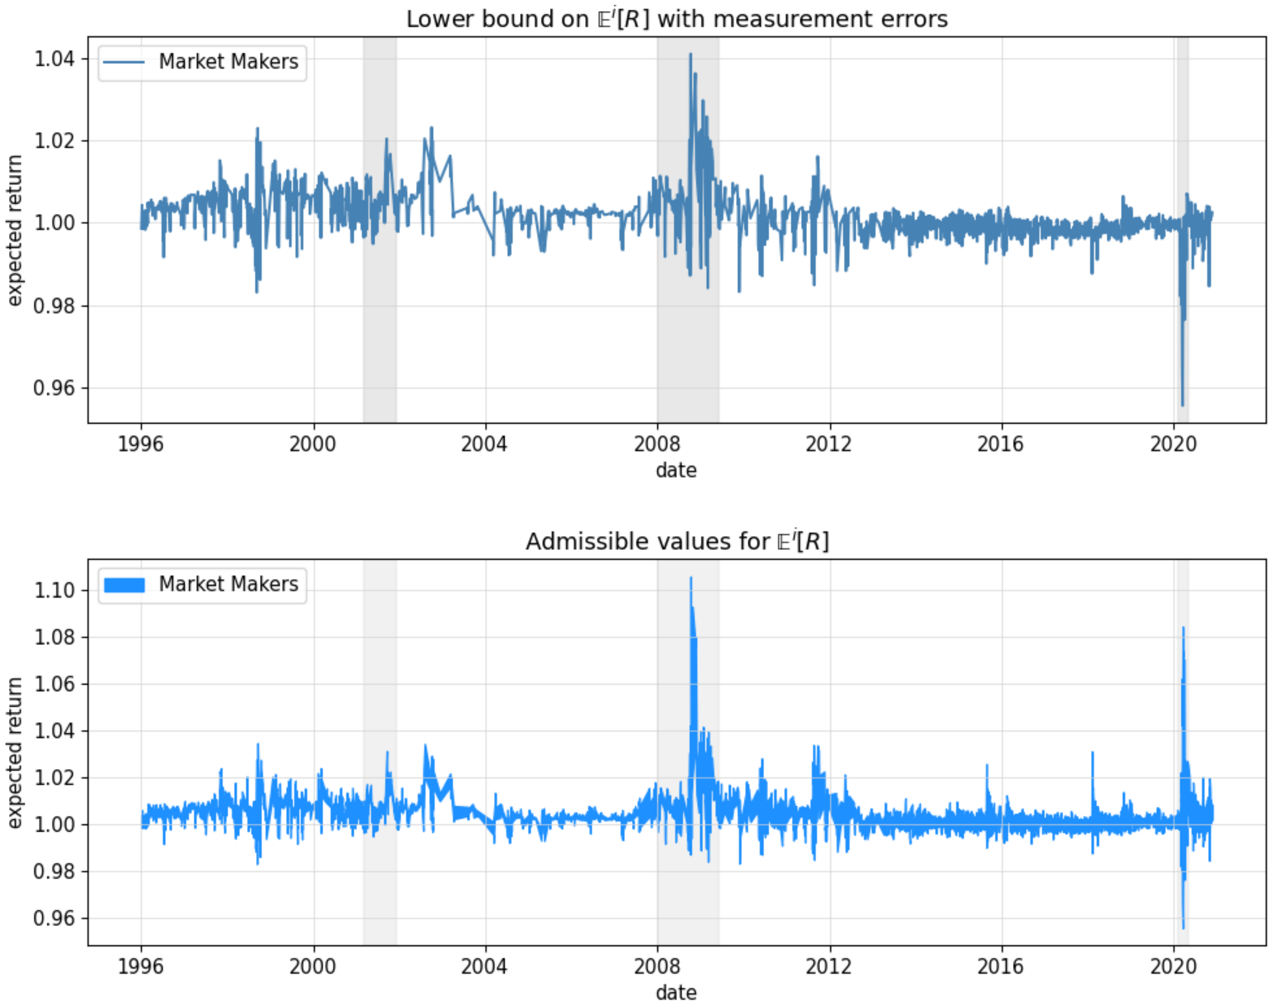

Figure 7. Bounds for Subjective Expected Market Return Market Makers
Notes: This figure plots the time-series of the lower bound on the subjective expected market return (upper plot) and all the possible values for the expected market return (lower plot) as recovered by SDFs compatible with market makers' positions in delta-hedged options. Frequency is daily and horizon is monthly. 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 November 2008 and April 2020, to illustrate two complementary situations appearing in the data. Figures 8 and 9 summarize the investors' portfolio composition in options, while Figures 10 and 9 depict the results.

In both instances, we find that the investment in puts is larger than in calls. November 2008 is a prime example of customers being overall short in OTM calls and puts, while market makers are long. This echoes earlier findings in Chen, Joslin, and Ni [2019] who argue that while market makers are net suppliers of insurance in normal times, they become net demanders in bad times when their financial constraints bind. On the contrary, during the Covid crisis the negative demand shocks absorbed by customers are not large enough to completely offset their total portfolio holdings. As a result, it becomes a mixed combination of long and short positions in calls and puts.

The recovered SDF shapes confirm these observations. In 2008, customers' SDF is increasing in the upside and the downside risk, exhibiting a clear U-shaped form, which becomes even more pronounced when we consider the upper bound with measurement errors. Deviations are not preserved in the lower bound yet, which can change convexity. Market makers' observed SDF is conversely decreasing (but not always monotonically), reflecting a portfolio that is long in options. In 2020, the balanced effect of options and the large investment in the index make the two SDFs quite similar in the shape, even though the customers SDF has larger slope.

The different compositions in the portfolios have significant effects on the ensuing expected returns. If we consider the case with no measurement error, customers perceive larger risk premia (2.9% vs. 2.2%) as they have increased their exposure to the index volatility in 2008; but the two groups formulate very similar expectations (1.3% vs. 1.2%) during Covid.

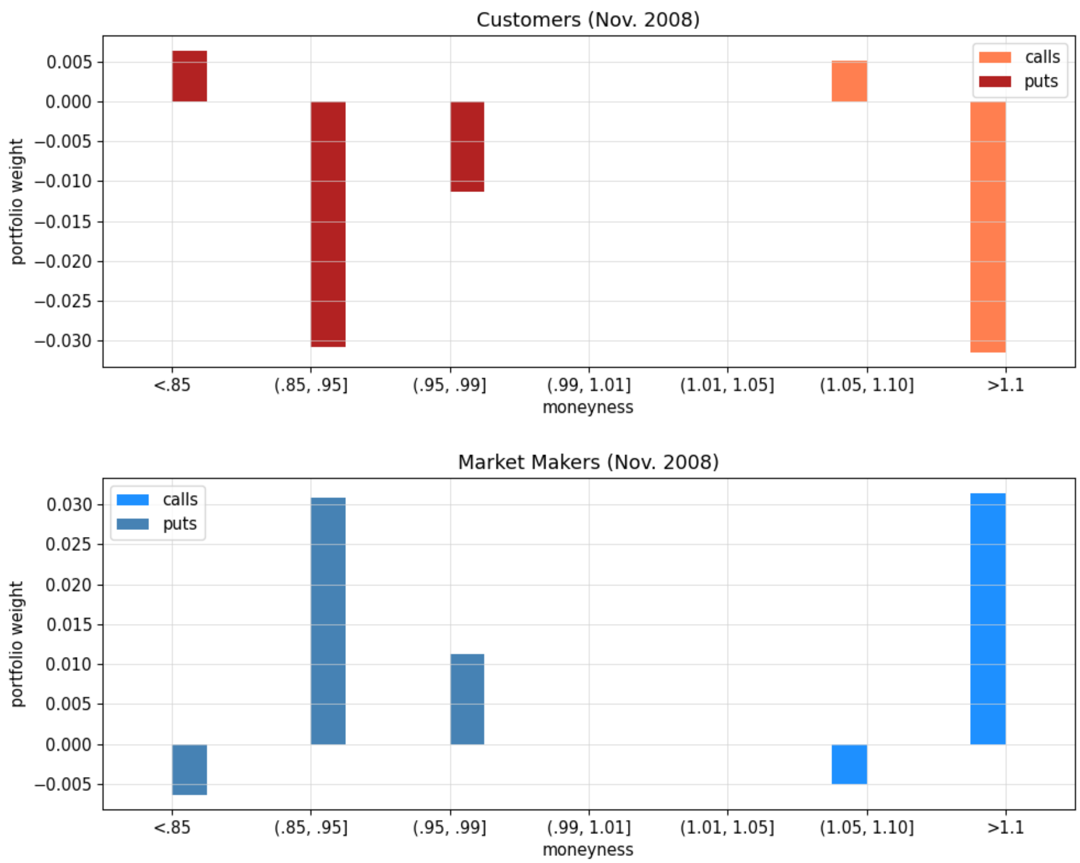

Figure 8. Portfolio Weights 2008 Great Financial Crisis
Notes: This figure shows the distribution of the portfolio weights for OTM calls and puts in the observed delta-hedged option portfolio of customers (upper panel) and market makers (lower panel), across different levels of moneyness, on November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

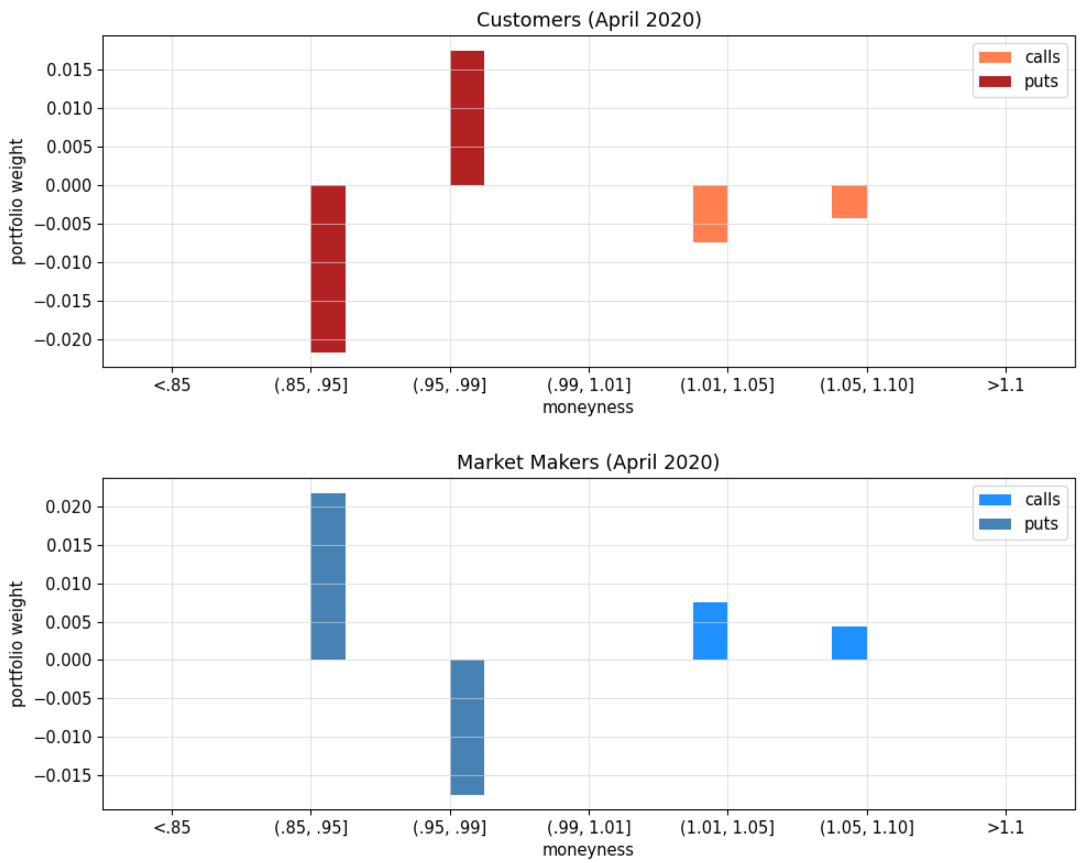

Figure 9. Portfolio Weights Covid

Notes: This figure shows the distribution of the portfolio weights for OTM calls and puts in the observed delta-hedged option portfolio of customers (upper panel) and market makers (lower panel), across different levels of moneyness, on April 2020.

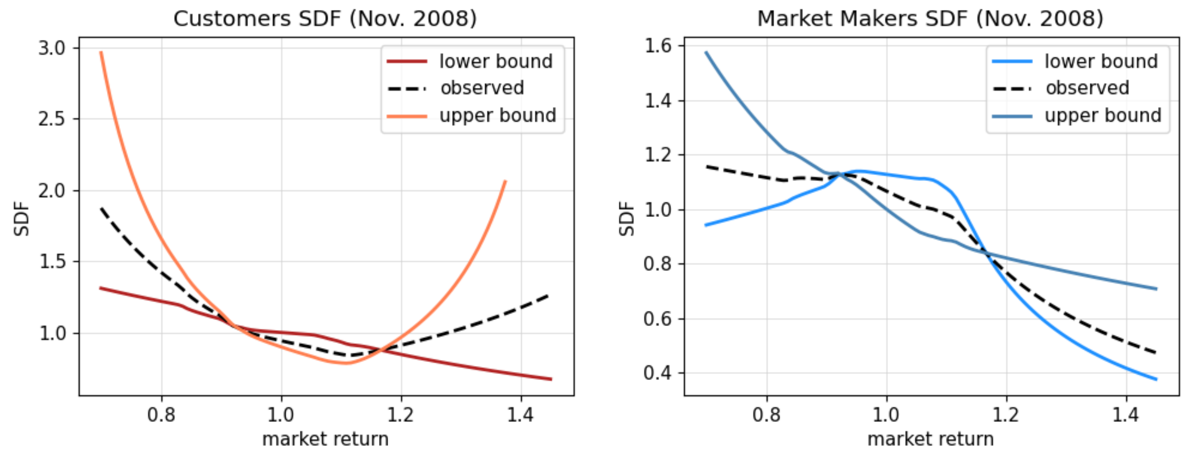

Figure 10. Recovered SDFs 2008 Great Financial Crisis
Notes: This figure plots the stochastic discount factors $M^{\star}$ recovered for customers (left panel) and market makers (right panel) on November 2008, as function of the market return. The black dashed line is the SDF recovered from the observed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

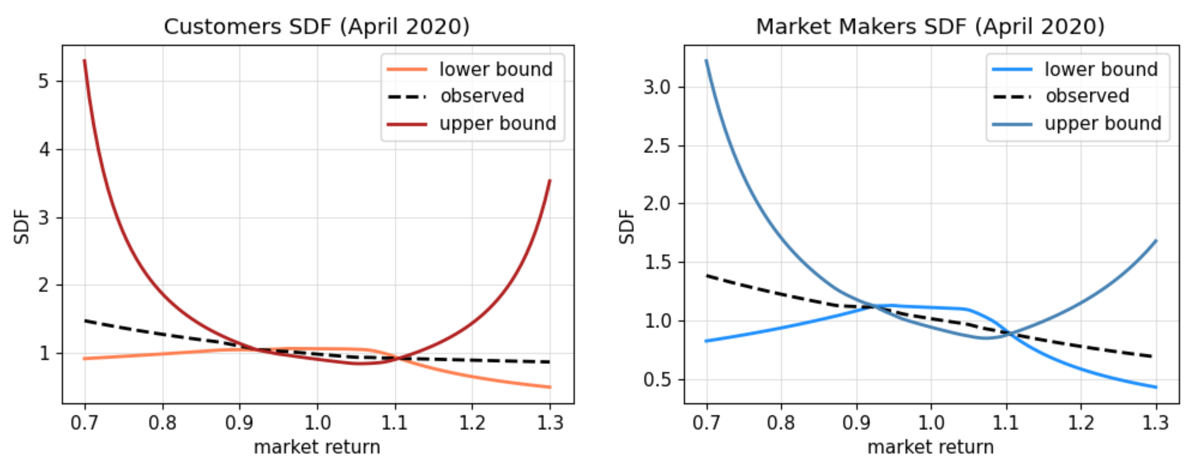

Figure 11. Recovered SDFs Covid
Notes: This figure plots the stochastic discount factors $M^{\star}$ recovered for customers (left panel) and market makers (right panel) on April 2008, as function of the market return. The black dashed line is the SDF recovered from the observed investors' portfolios. 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 may vary considerably across the two investors and significantly depends on the structure of the supporting portfolio and on the state of the economy.

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 2 (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 2, 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 2 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 3 (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 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 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 3 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 3 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 3, 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:

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

B Additional Figures

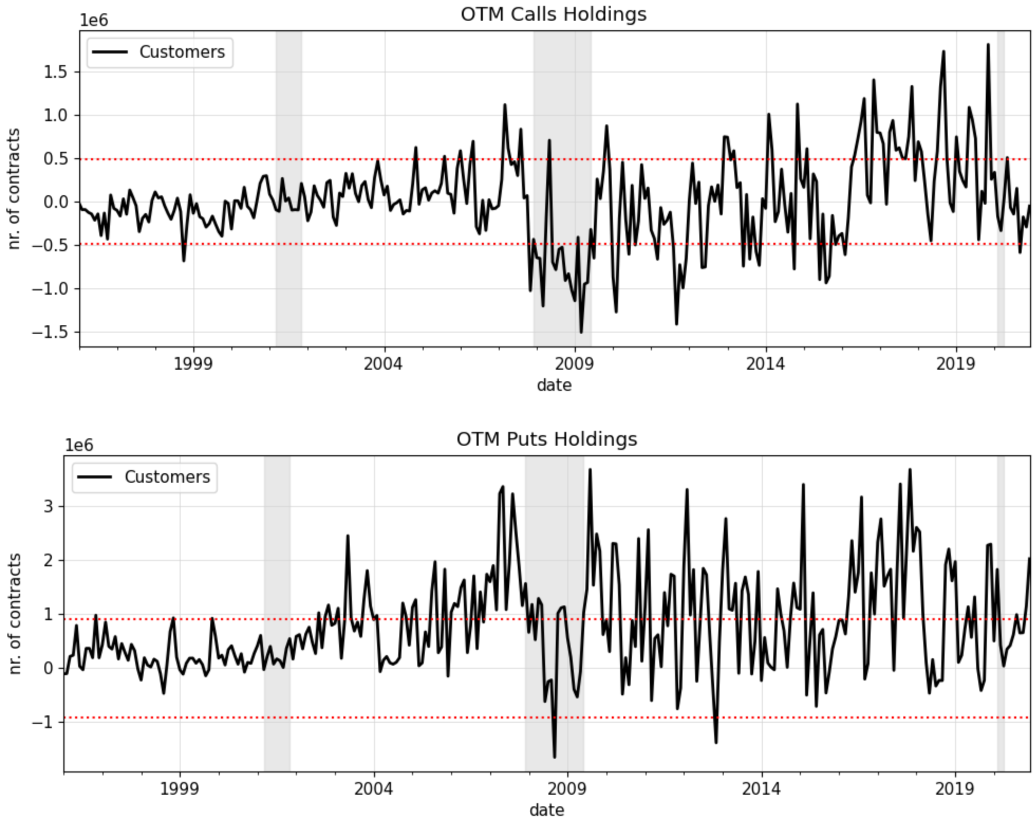

Figure B.1. Customers' holdings of monthly OTM options

Notes: This figure plots the time-series of our proxy for customers' portfolio holdings of OTM calls (upper panel) and puts (lower panel) expiring between 15 and 60 days, on our dataset from 1996 to 2020. Options holdings are the sum of opening and closing positions on the same contract that customers take from the new emission to the expiration date. For readability, the daily data are aggregated on a monthly basis. Gray bars indicate NBER recessions.The red dotted lines depict ±1 standard deviation around 0.

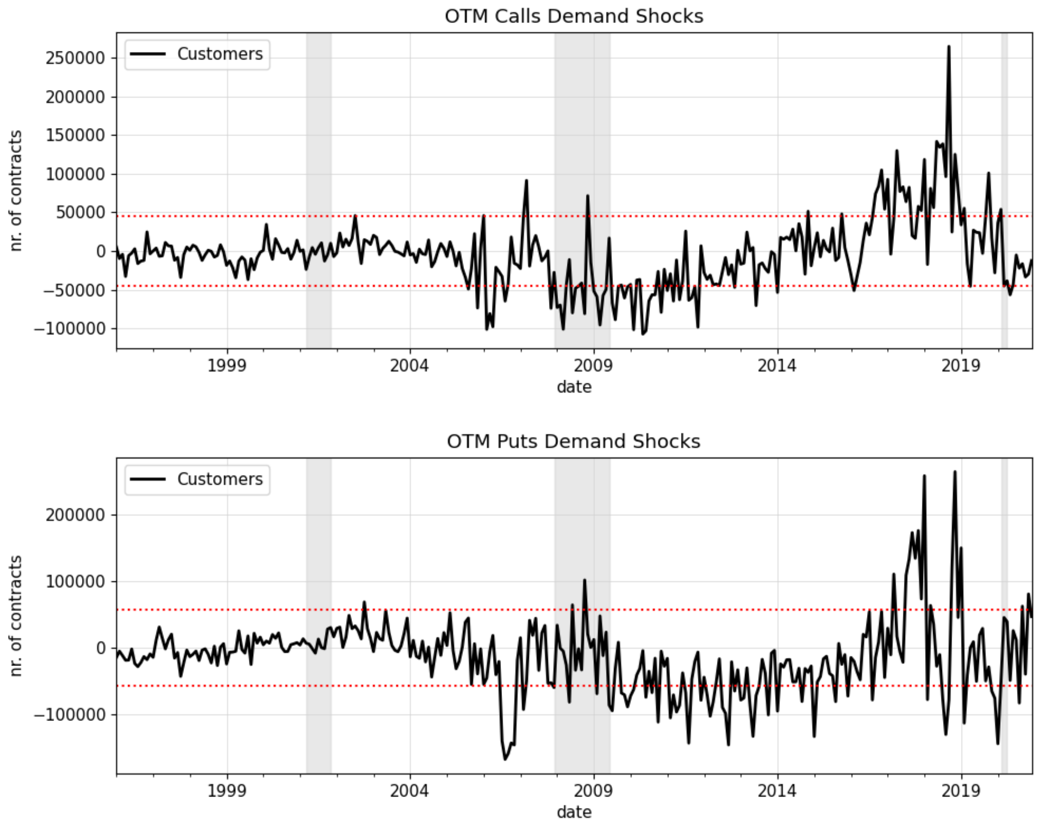

Figure B.2. Customers' demand shocks on monthly OTM options
Notes: This figure plots the time-series of customers' demand shocks for OTM calls (upper panel) and puts (lower panel) expiring between 15 and 60 days, on our dataset from 1996 to 2020. Options demand shocks are the sum of opening positions (buy orders minus sell orders) recorded on each date. For readability, the daily data are aggregated on a monthly basis. Gray bars indicate NBER recessions.The red dotted lines depict ± 1 standard deviation around 0.

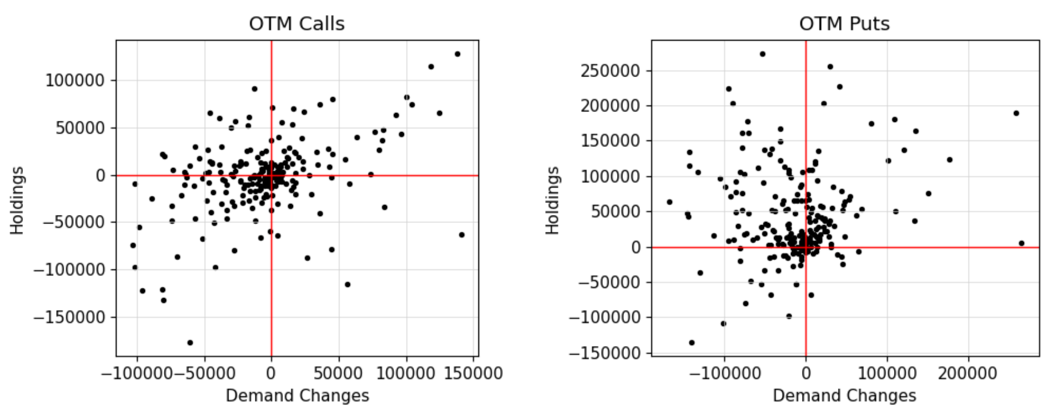

Figure B.3. Customers' demand changes vs. holdings on monthly OTM options
Notes: This figure plots the holdings and the holdings variations in the customers' portfolios for OTM calls (left panel) and puts (right panel) expiring between 15 and 60 days. Data are computed daily from 1996 to 2020, then they are monthly aggregated for readability.

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*.

Bakshi, Gurdip, Charles Cao, and Zhiwu Chen. 1997. "Empirical performance of alternative option pricing models." *The Journal of finance* 52 (5): 2003–2049.

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.

Moussawi, Rahib, Lai Xu, and Zhaoque Zhou. 2024. "A Market Maker of Two Markets: The Role of Options in ETF Arbitrage." *Working Paper, Villanova University*.

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*.