Demand-Based Expected Returns

Alessandro Crescini Fabio Trojani Andrea Vedolin

December 2024

Abstract

This paper proposes a theoretical framework for recovering investors' subjective beliefs using holdings data and option prices under the assumption of no-arbitrage. We recover investor-specific expected returns and risk and study how to measure consensus belief when beliefs are heterogeneous. Using buy and sell orders on S&P500 options, we empirically document that the statistical properties of subjective expected returns and Sharpe ratios differ wildly between investor types 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 expected returns that are less volatile and acyclical. More specifically, we show that the expected returns inferred from the beliefs of retail and institutional investors can change sign, are acyclical, and are significantly correlate with institutional investor survey measures. Expected returns of market makers correlate highly with price-based measures.

Keywords: subjective expected returns, subjective risk, 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 they increase significantly during crisis periods and are highly volatile. Estimates of expected returns from survey data, however, are less volatile and can be pro-, a-, or countercyclical depending on investor type and level of sophistication, see, e.g., Greenwood and Shleifer [2014], Nagel and Xu [2023], and Dahlquist and Ibert [2024], respectively. Price-based measures of expected returns ignore information on investors' holdings. However, recent research highlights the strong link between holding data and investor beliefs.[1] In this paper, we argue that including information about investors' portfolios is crucial to understanding the dynamics of subjective expected returns and risk.

To this end, we propose a theoretical framework for recovering beliefs of heterogeneous investors from prices and holdings data jointly under the assumption of no-arbitrage. More specifically, we theoretically show that investors' subjective expected return on the market as well as their perceived risk can be directly inferred from option prices and their corresponding option holdings in real-time at granular levels. We empirically document substantial heterogeneity in expected return estimates between investor types. Most importantly, we find that expected returns recovered from holding 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 retail and institutional investors can be significantly smaller and less cyclical than expected returns that are recovered from price data alone.

Intuitively, these investors buy insurance in the options market, which lowers their exposure to market crashes. Financial intermediaries are net suppliers of deep-out-of-the-money puts to public investors, see, e.g., Gârleanu, Pedersen, and Poteshman [2008] and Chen, Joslin, and Ni [2019]. Their expected returns increase during bad times (due to their market exposure via the short puts), while customers' expected returns can be lower (due to their protection). We also find that our measures of expected returns correlate highly with survey measures of expected returns of professional investors, while the correlation with households' expected returns is low. 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 (highly sophisticated) 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

[1] For example, Giglio et al. [2021] document a strong relationship between investor expected returns and portfolio holdings using surveys. Beutel and Weber [2023] provide causal evidence for the link between beliefs and portfolio decisions using experiments. And Egan, MacKay, and Yang [2024] show that belief heterogeneity accounts for the majority of the variation in households' portfolio allocations.

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. However, these methods ignore information about quantities (such as portfolio holdings, trading flows, or open interest), which, unlike prices, are available on a granular level, that is, for each investor.

Our "demand-based" belief recovery extracts $\mathbb{P}_i$ for an investor i 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 $\mathbb{P}$ 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 the 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 not only measures of subjective expected returns in real-time for each investor type, but it also allows us to recover measures of subjective risk.

We also show how to recover the consensus belief in an economy with heterogeneous agents who agree to disagree. In particular, we show that as long as both agents are unconstrained and hold the growth-optimal portfolio, the consensus belief equals the risk-neutral variance.

With this methodology, we obtain SDFs that are joint functions of the index and the options returns. Therefore, the ensuing expected market returns and measures of risk 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 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 it is assumed 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 index 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 allow 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 customers and market makers can have complementary patterns with regard to the shape of their SDFs. For example, during normal times, market makers hold large short positions in calls and puts which exposes them to changes in both the up- and downside. As a consequence, market makers' SDFs are U-shaped as a function of expected returns. Customers, on the other hand, who are net demander of these options, have inverted patterns. These regularities, however, changed dramatically during so-called crisis days when intermediaries' constraints start to bind. For example, we find that in March 2020, customers' SDF projections are U-shaped, while market makers' SDFs are inverted. The reason for this is that customers were net short during this period, selling protection to market makers.

Second, the changing portfolio holdings and exposures to downside risk during crisis periods across the two investors has large effects on the time-series properties of expected returns. We find that the ensuing expected returns are much more highly correlated with price-based measures for market makers than for customers. Customers' expected returns can deviate from price-based measures, often turning negative and significantly smaller in size.

Third, we can also use our framework to recover subjective measures of risk, allowing us to measure subjective risk and return trade-offs. We find that measures of perceived risk are more highly correlated than their expected return measures, while subjective Sharpe ratios are basically uncorrelated (mostly due to the low correlation in expected returns).

Finally, we also study the relation of our expected return measures with survey measures of expected market returns. We find that while the correlation with survey measures of expected returns for individuals or retail investors is negative, the correlation of expected return measures of professional investors with intermediaries' expected return is positive and significant. Finally, we study determinants of our expected return measures. We find that standard predictors of realized returns such as the dividend-price ratio, do not have any statistically significant relation with expected returns. The only variable that loads significantly on expected returns are past returns where the slope coefficient is negative. This implies that when returns are low, both market makers and customers expected expected returns to be high. This is in contrast to survey evidence which has shown that retail investors' expected returns tend to be positively correlated with past returns. We also do not find any statistically significant relation between expected returns and measures of cyclicality, highlighting the acyclical nature of our measures.

Our measures of subjective expected returns and risk can be interpreted as the lower bounds on expected returns and perceived risk by heterogeneous investors, offering intuition for why some of the literature has documented different cyclicality patterns across various surveys. We emphasize that our approach does not recover the "true" beliefs of investors but provides a sensible benchmark for the potential beliefs of large players in the option and index market.

Related Literature. This paper is related to several branches 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 two ways: First, we include demand-based data instead of just asset pricing data to extract investor-specific beliefs. Second, 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 the 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 uses asset prices to recover measures of 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. Chabi-Yo and Loudis [2020] 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 moments (as opposed to risk-neutral) 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 expected returns measures are crucially dependent 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 holdings in both the market and power contracts in 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 that studies the beliefs of heterogeneous investors using surveys. Dahlquist and Ibert [2024] 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 the beliefs of households and the allocation of the portfolio 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. However, we document that intermediaries' subjective measure of returns is highly correlated with survey data of sophisticated investors, whereas households' expectations are not.

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

Outline. The remainder 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 in which we show how to infer subjective expected returns and risk 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 the perception of risk of market participants. To provide some intuition, we start with an example to illustrate how portfolio holdings affect beliefs/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 that 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]. As we argue in our paper, 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 on 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 90% in the index and 10% in an equally-weighted portfolio of calls or puts with the same maturity but different strikes.[3] As is immediately evident, even for a small investment in options (with respect to the investment in the underlying) and small changes in the market return, the ratio $M^0/M^\star$ can wildly differ 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 6 (5) for OTM calls (puts).

Intuitively, if we would like to recover the probability measure $\mathbb{P}^\star$ (supported by SDF $M^\star$) from the probability measure $\mathbb{P}^0$ of an investor fully invested in the market, we can interpret

[2]Theoretically, positive net demand for options can arise in settings with heterogeneous beliefs. For example, Buraschi, Trojani, and Vedolin [2014] show that agents with more pessimistic views about the future growth of the economy demand OTM puts from more optimistic agents. In a setting with frictions and market incompleteness Johnson, Liang, and Liu [2016] show 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.

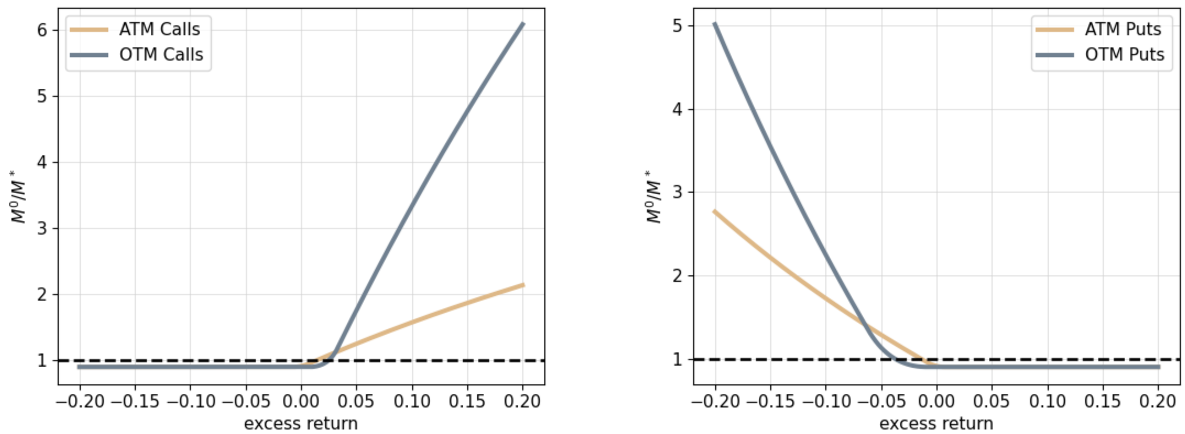

Figure 1. Ratio between Benchmark SDF M^0 and $M^\star$
Notes: This figure plots $M^0/M^\star$ as a function of the excess return on the market. $M^\star = 1/R^\star$, where $R^\star$ is the return of a portfolio investing 90% of the wealth in the underlying and 10% in an equally-weighted portfolio of calls (left plot) and puts (right plot) with different moneyness. ATM options have $|\Delta| \in (0.4, 0.6]$. OTM options have $|\Delta| \in [0.1, 0.4]$.

the ratio $M^0/M^\star$ as the corresponding probability distortion. 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. More generally, investors who assign higher weights to extreme events have higher demand for deep OTM puts. The reverse holds true when the investor is long calls.

In a next step, we study the effect of investors' demand on the time-series properties of subjective expected returns. In Figure C.1, 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).

Panel A plots the perceived expected returns for investors who hold call options in addition to the index. Not very surprisingly, the patterns mirror the benchmark case almost one-for-one. Expected returns increase in bad time, decrease in normal periods, and are highly volatile. Notice that expected returns are considerably higher even with a small investment in options. Intuitively, the size of the expected return increases relative to the benchmark since 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 middle 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) market fluctuations.

From their perspective, the risk-return ratio given by holding the index alone is not profitable - or equivalently, they find OTM puts to be under-priced. Buying puts leads to protection which reflects a pessimistic view and negative expected return. Accordingly, the volatility of the expected return is higher with respect to what we recover under M^0. Notice also that even small investments in puts can lead to a pro-cyclical pattern of the expected return. Given that, we conclude that the size and cyclicality properties of expected returns depend on the options moneyness, returns, 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). In Figure C.1 Panel C, we plot the monthly time-series of expected returns[4] recovered from hypothetical delta-hedged collar and straddle strategies with OTM options, compared to the benchmark case. Results align with our previous example: being long in calls/puts increases the distortion in the tails, the reverse is true for short positions. Thus the option component in a collar strategy significantly reduces the expected return.

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.

[4]When the no-arbitrage condition is violated, we drop the corresponding observation.

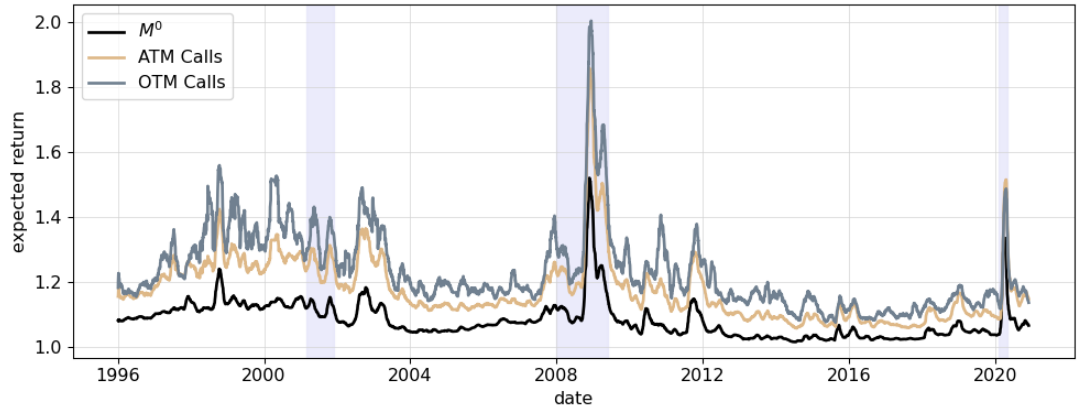

A. Calls

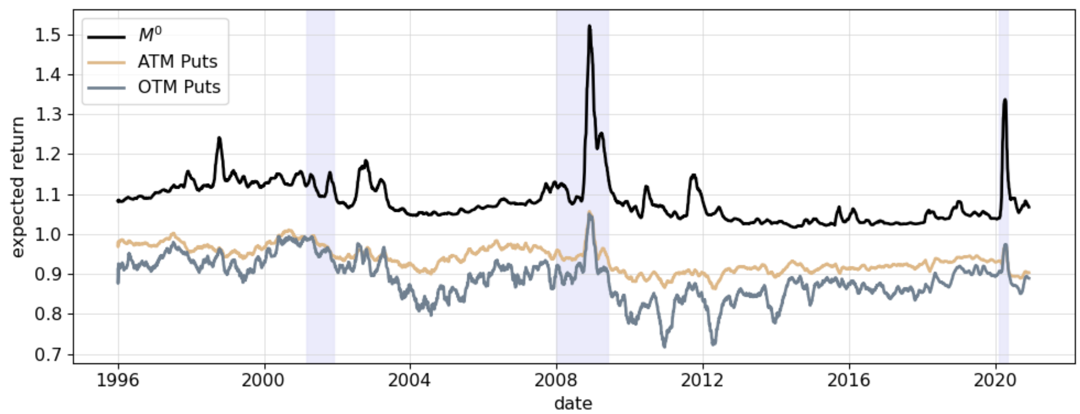

B. Puts

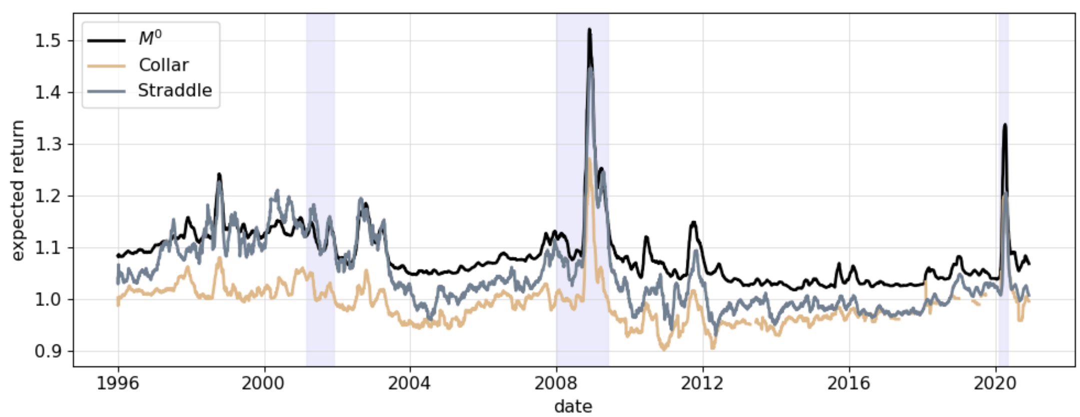

C. Collar and Straddle

Figure 2. Time-Series of Expected Market Return

Notes: This figure plots the expected market return recovered from different stochastic discount factors. In each panel, we plot the expected return recovered by $M^0 = 1/R$, as well as by $M^\star = 1/R^\star$, where $R^\star$ is the return of a portfolio investing 90% in the index and 10% in an equally-weighted portfolio of either ATM or OTM options for calls (Panel A) and puts (Panel B). Panel C shows a «collar» portfolio which is long in the index and in OTM puts, and short in OTM calls; a «straddle» portfolio which is long in the index and in both OTM calls and puts. Frequency is daily, horizon is monthly. Time series are 30-days moving averages. Values are annualized. Grey areas indicate NBER recession periods.

3 Theoretical Framework

We now present a simple theoretical framework to explain how to recover subjective expected returns from options price and holdings 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 written on the risky asset with a continuum of strike prices. Let $\mathbb{E}^i[\cdot]$ denote the subjective conditional 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 $\boldsymbol{R}^e = \boldsymbol{R} - \mathbf{1}$ be the excess forward return of the index and options.

3.1 Subjective Expected Returns

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:

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

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

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

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:

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

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. Intuitively, we interpret equation (2) as the most conservative measure of the expected return of a log investor. Two remarks are in order. First, notice that we do not need to make any assumptions about whether investor i is constrained or not since the constraints are reflected in the portfolio holdings. Second, as we will show later on, even if we do not believe that investors have log utility, identity (2) still provides us with a useful lower bound.

In the following, we will consider two different cases. First, the fact that portfolio holdings are observable in the data at high frequency allows us to directly recover agent i's expected return of the market as a function of holdings and option prices in real time. 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, which does not capture the full risk exposure of the agent. 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 and neither the transactions on other derivatives with the same underlying, such as SPY options (that is on the ETF tracking the S&P500).[5] Since major market makers provide liquidity in both SPX and SPY option markets, we only observe a fraction of their true market exposure.[6] 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 lower and upper bounds on the subjective expected returns representing the most conservative and most aggressive value, respectively.

3.1.1 Log Utility Assumption

Before explaining technical details on how to recover subjective beliefs from the data, one might be worried that the log utility assumption used in equation (1) is unrealistic. Notice that our framework still provides insights into the subjective expected return in the form of a lower bound even if investors do not have log utility. To see this, suppose that the investor has some other utility function, and potentially unknown belief $\mathbb{P}^*$. We say that the negative covariance condition (NCC) holds if:

$$\mathbb{C}ov^*(M^*\boldsymbol{\theta}_i'\boldsymbol{R}, R) \leq 0\ , \tag{3}$$

where $M^* = d\mathbb{Q}/d\mathbb{P}^*$ is the true stochastic discount factor of the non-log investor and $\boldsymbol{\theta}_i$ is her portfolio, as above. In the case where the portfolio is fully invested in the index, we get the

[5]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.

[6]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.

same NCC as in Martin [2017]. In particular, he shows that the NCC is likely to hold empirically and also holds in several leading macro-finance models. If the NCC holds for any portfolio $\boldsymbol{\theta}_i$ of agent i, we get the following lower bound:

$$\mathbb{E}^*[R] \geq \frac{\mathbb{E}^{\mathbb{Q}}[(1+\boldsymbol{\theta}_i'\boldsymbol{R}^e)R]}{\mathbb{E}^{\mathbb{Q}}[1+\boldsymbol{\theta}_i'\boldsymbol{R}^e]} = \mathbb{E}^i[R] \,, \tag{4}$$

where we formally recognize $(1+\boldsymbol{\theta}_i'\boldsymbol{R}^e)/\mathbb{E}^{\mathbb{Q}}[1+\boldsymbol{\theta}_i'\boldsymbol{R}^e]$ as the change of measure from $\mathbb{Q}$ to belief $\mathbb{P}^i$.[7] This change of measure has the functional form of a log-type SDF as in (1). The right-hand side of inequality (4) is the recovered expected return in (2). Therefore, the $\mathbb{E}^i[R]$ we extract under the log-utility assumption additionally provides a lower bound for the subjective moment of the agent holding the same portfolio but with non-log utility.

3.2 Recovering Subjective Expected Returns

We can now discuss how to recover subjective beliefs from options. 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 $\mathbb{E}^{\mathbb{Q}}[\boldsymbol{R}^e] = \boldsymbol{0}$.

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 $(1+\boldsymbol{\theta}_i'\boldsymbol{R}^e)R$ using the Carr and Madan [2001] formula.

Before delving into details about implementation, several remarks are in order. Our setup nests 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 a 100% investment in the market return. A consequence of this assumption is that all options are redundant. As we argue in this paper, 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 on monthly options is: 28% for OTM puts, 20% for ATM puts, 30% for OTM calls, and 1% for ATM calls, respectively.

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

[7] But $\mathbb{E}^{\mathbb{Q}}[1+\boldsymbol{\theta}_i'\boldsymbol{R}^e] = 1$.

wise. 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 $\mathbb{Q}$[8]. We will provide an additional characterization of $\mathbb{E}^{\mathbb{Q}}[R^2]$ as a consensus belief in the next paragraphs.

Another related paper is Tetlock [2023] who assumes that the investor can hold (integer) power contracts written on the market index. The paper also 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] estimates "hypothetical" portfolio holdings from expanding window regressions predicting risk-neutral (power) moments with physical counterparts. Our approach is different since we directly observe holdings on plain vanilla calls and puts.

3.3 Bounds on Subjective Expected Returns

Assuming that we observe holdings with no error, subjective expected returns can be recovered via the exact identity in equation (2). For the reasons discussed earlier, when holdings $\boldsymbol{\theta}_i$ are observed with error, we can provide lower and upper bounds on the subjective return on the market. Since in that case, we want to constrain the amount by which the optimal portfolio weights can deviate from the observed weights, we impose some constraint such that

$$d(\boldsymbol{\theta}_i, \boldsymbol{\theta}) \leq \delta\,, \tag{5}$$

where $\boldsymbol{\theta}$ are the growth-optimal portfolio weights, and for some convex discrepancy function $d(\cdot,\cdot) \geq 0$ such that $d(\boldsymbol{x}, \boldsymbol{y}) = 0$ if and only if $\boldsymbol{x} = \boldsymbol{y}$. Intuitively, δ measures the amount that optimal weights can deviate from the observed weights. To be concrete, in the following, we assume a L^2-norm.[9]

In that case, $d(\boldsymbol{\theta}_i, \boldsymbol{\theta}) = \frac{1}{2}||\boldsymbol{\theta} - \boldsymbol{\theta}_i||_2^2$. To determine the subjective expected return, we now have to solve the following optimization problem:

$$\inf_{\boldsymbol{\theta}\in\mathbb{R}^N} \left\{ \mathbb{E}^{\mathbb{Q}}[(1+\boldsymbol{\theta}'\boldsymbol{R}^e)R] + \lambda \left(\frac{1}{2}||\boldsymbol{\theta} - \boldsymbol{\theta}_i||_2^2 - \delta \right) \right\} \ . \tag{6}$$

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

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

[9] The main reason we use an L^2-norm (as opposed to other norms) is for tractability, as in this case, we get closed-form solutions.

Proposition 1 (Bounds on Subjective Expected Returns)**.** *Assume that portfolio weights $\boldsymbol{\theta}_i$ are observed with errors. In that case, the lower bound for the subjective return on the market for investor i is:*

$$\mathbb{E}^i[R] \geq \mathbb{E}^{\mathbb{Q}}[(1+\boldsymbol{\theta}_i'\boldsymbol{R}^e)R] - \sqrt{2\delta}||\mathbb{E}^{\mathbb{Q}}[\boldsymbol{R}^e R]||_2 \,. \tag{7}$$

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

$$\mathbb{E}^i[R] \leq \mathbb{E}^{\mathbb{Q}}[(1+\boldsymbol{\theta}_i'\boldsymbol{R}^e)R] + \sqrt{2\delta}||\mathbb{E}^{\mathbb{Q}}[\boldsymbol{R}^e R]||_2 \tag{8}$$

It is obvious that if the measurement error is assumed to be zero, i.e., $\delta = 0$, then equations (7) and (8) coincide with equation (2). Intuitively, we can interpret the lower bound as the most conservative assessment of the subjective expected return for any investor whose portfolios align in a neighborhood around the observed portfolios $\boldsymbol{\theta}_i$ (equivalently, whose belief is compatible to some extent with the belief reflected by investing in the portfolio $\boldsymbol{\theta}_i$).

3.4 Heterogeneous Beliefs

So far, we studied expected returns for a single log-investor. In the following, we extend our results to a simple economy with n investors, who can allocate their wealth to the aggregate market index (assumed in positive net supply) and to a portfolio of options (assumed in zero net supply).

We denote by R^i the (forward) return on the i-th investor's wealth over horizon $[0, T]$, and by R the (forward) return on the market index. The wealth of investor i' at time t is indicated by W_t^i. Given the market index value I_t at time t, the market clearing condition yields:

$$I_t = W_t := \sum_i W_t^i \,, \tag{9}$$

and

$$R = \frac{I_T}{I_0} = \frac{W_T}{W_0} = \sum_i \frac{W_0^i}{W_0^i} R^i =: \sum_i w_i R^i \,, \tag{10}$$

with w_i being the $i-$th investor's share of the aggregated wealth.

Let further $\mathbb{P}^i$ denote $i-$th investor's subjective probability belief and $\mathbb{E}^i[\cdot]$ the conditional expectations under this belief. The consensus market expected return among all investors is then defined by:

$$\bar{\mathbb{E}}[R] := \sum_i w_i \mathbb{E}^i[R] \,. \tag{11}$$

The next proposition gives a first characterization of the consensus belief in an economy where all market participants are optimally investing in the index and in the option.

Proposition 2. *Let investor j be a log investor optimally invested in the index, as well as in the option market. It then follows, that j's expected market return is given by:*

$$\mathbb{E}^j[R] = 1 + \mathbb{C}ov(R^j, R) = 1 + w_j \mathbb{V}ar^{\mathbb{Q}}(R^j) + \sum_{i \neq j} w_i \mathbb{C}ov^{\mathbb{Q}}(R^j, R^i) \,. \tag{12}$$

If all investors are log investors optimally invested in the index and in the option market, then the consensus belief is given by:

$$\bar{\mathbb{E}}[R] = 1 + \mathbb{V}ar^{\mathbb{Q}}(R) \tag{13}$$

If all index weights w_i are positive, there exists a consensus belief $\bar{\mathbb{P}} := \sum_i w_i \mathbb{P}^i$ such that:

$$\bar{\mathbb{E}}[R] = \mathbb{E}^{\bar{\mathbb{P}}}[R] \,. \tag{14}$$

An obvious special case of Proposition 2 arises when all investors have identical beliefs, in which case $\bar{\mathbb{P}}$ is the common belief of each investor in the economy. Equation (13) shows that the risk-neutral variance (or SVIX) is the consensus expected return in an economy where (i) all investors are unconstrained log-investors and (ii) wealth is optimally invested in the index and options. In that case, the inverse of the optimal growth portfolio (the stochastic discount factor) prices not just the index but also all options.[10]

In general, we do not expect (i) and (ii) to hold in the real world. However, notice that it is still possible to identify a consensus belief aggregated only among all investors who are optimally invested in both the index and the option market. To this end, let $\mathcal{I}$ be the set indexing the subset of such investors having a non-zero optimal wealth allocation to the index and option market. For each investor $i \in \mathcal{I}$ it then follows:

$$\mathbb{E}^i[R] = 1 + \mathbb{C}ov(R^i, R) \,. \tag{15}$$

[10]Juxtaposing this result to Martin [2017], notice that investors in our economy are not restricted to just holding the market index, because of market clearing in the option market (which is in zero net supply).

The associated consensus belief across such investors is analogously defined as:

$$\bar{\mathbb{E}}^{\mathcal{J}}[R] := \sum_{i \in \mathcal{I}} w_i^{\mathcal{I}} \mathbb{E}^i[R] \, , \tag{16}$$

with wealth shares

$$w_i^{\mathcal{I}} = \frac{W_0^i}{\sum_{i \in \mathcal{I}} W_0^i} \tag{17}$$

The following corollary then characterizes the corresponding consensus belief.

Corollary 1. *In the setting of Proposition 2, if there is a subset $\mathcal{I}$ of investors optimally investing in the index and in the option market, then:*

$$\bar{\mathbb{E}}^{\mathcal{I}}[R] = 1 + \mathbb{C}ov^{\mathbb{Q}}(R^{\mathcal{I}}, R) \, , \tag{18}$$

where

$$R^{\mathcal{I}} := \sum_{i \in \mathcal{I}} w_i^{\mathcal{I}} R^i \, . \tag{19}$$

If all weights $w_i^{\mathcal{I}}$ are positive, then:

$$\bar{\mathbb{E}}^{\mathcal{I}}[R] = \mathbb{E}^{\bar{\mathbb{P}}^{\mathcal{J}}}[R] \, , \tag{20}$$

for a consensus belief $\bar{\mathbb{P}}^{\mathcal{J}} := \sum_{i \in \mathcal{J}} \mathbb{P}^i$.

It follows from Corollary 1 that the consensus equity premium is positive if and only if return $R^{\mathcal{J}}$ correlates positively with the market index return under the pricing measure $\mathbb{Q}$. This feature is more likely to emerge, e.g., when the average index holding across these investors is positive. For instance, whenever the net aggregate wealth allocated to options by investors $i \in \mathcal{I}$ is nil, then [11]:

$$\bar{\mathbb{E}}^{\mathcal{I}}[R] = \bar{\mathbb{E}}[R] = 1 + \mathbb{V}ar^{\mathbb{Q}}(R) > 1 \, .$$

Conversely, when these investors additionally have zero exposure to market risk on aggregate, then $\sum_{i \in \mathcal{J}} w_i = 0$, implying $R^{\mathcal{I}} = 0$ and $\bar{E}^{\mathcal{I}}[R] = 1$.

[11] Because of the option market clearing condition, this identity holds, e.g., when set $\mathcal{J}$ contains all option investors in the economy.

3.5 Subjective Risk

Finally, our theoretical framework also allows us to study subjective expectations of higher-order moments beyond returns. To this end, we next study the subjective risk-return trade-off. While the relationship between measures of realized risk and returns is normally weak, recent literature generally reports a strong positive relationship using survey-based measures, see, e.g., Couts, Goncalves, and Loudis [2023] for the aggregate stock and bond markets and Jensen [2024] for individual stocks. Our theoretical framework can be trivially extended to study measures of subjective risk from prices and holdings using equation (2). It follows immediately that subjective risk for investor i is given by:

$$\mathbb{E}^i[R^2] = \mathbb{E}^{\mathbb{Q}}[M_i^{-1}R^2] = \mathbb{E}^{\mathbb{Q}}\left[\left(1+\boldsymbol{\theta}_i'\boldsymbol{R}^e\right)R^2\right] \ , \tag{21}$$

and the subjective volatility reads as $\mathbb{E}^i[R^2] - \mathbb{E}^i[R]^2$. In the following section, we explain how we recover subjective expected returns and risk using equations (2) and (21).

4 Empirical Analysis

This section describes the data used and how we empirically implement our main theoretical results. For simplicity, we always show measures of spot returns.

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 is common practice, we aggregate these daily volumes to cumulative positions for the last three categories[12] 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:

[12] 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. Summary statistics monthly options data.

	Calls				Puts			
	K/S_t		Holdings		K/S_t		Holdings	
	All	OTM	All	OTM	All	OTM	All	OTM
mean	1.02	1.04	-20,607	-9,133	0.89	0.95	65,696	21,890
std	0.09	0.03	57,496	21,140	0.14	0.03	94,310	41,301
min	0.19	1.00	-675,460	-153,762	0.11	0.71	-534,099	-203,685
median	1.02	1.03	-14,879	-5,688	0.92	0.96	46,813	12,862
max	1.82	1.57	494,353	132,687	2.67	1.00	1,282,355	358,137

Notes: This table reports summary statistics for the options data interpolated at maturity of 30 days. The relative moneyness K/S_t is computed over the single option contracts that are traded on every date. «Holdings» are defined as the total customers' opening/closing buy orders minus sell orders cumulated for every option from the issuance to every trading date, and it is aggregated over all the traded options on a single day. For each variable, the first column refers to the full dataset, while the second to OTM options (with $0.1 \leq |\Delta| \leq 0.4$) only. Data runs from January 1996 to December 2020.

with a price less than \$3/8; 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; and 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 time. No-arbitrage filters apply as well.

Throughout our empirical analysis, we use monthly horizons, as the most liquid options over the full sample expire around thirty days (see Figure B.3). We compute every quantity at daily frequency, then we aggregate to monthly moving averages to account for the large turnover in the option holdings[13]. On each day t, separately for calls and puts, we linearly interpolate options volatility, options delta, and investors' holdings, on a grid of strike prices, for the required maturity (30 days if not explicitly stated otherwise).[14] The grid points consist of the n different strikes that are actively traded in t.

Table 1 reports summary statistics on the interpolated option sample for customers' net demand for calls and puts. Holdings are defined as total buy minus total sell orders aggregated over a day and cumulated over time from the option emission (more details follow in Section 4.2.1). On average, customers are net sellers for call options and net buyers for put options. In our empirical analysis, we focus on OTM options having $|\Delta| \in [0.1, 0.4]$ as they account for the largest trading volume (57%, see Figure B.2).Finally, we also exclude very deep-OTM options (with $|\Delta| < 0.1$) because they are more likely to be affected by interpolation errors.

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

[13]Among others things, this depends on the data aggregation, because the label Customers includes many different individual agents.

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

are assembled at date t.

4.2 Implementation

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

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

where $x(k) := X(K)/F$ with F being the forward price of the market index, and $k := K/F$ the relative moneyness.

4.2.1 Constructing Option Holdings

Subjective expected returns depend on portfolio holdings, θ_i, where θ_i are the *total* portfolio holdings of investor i as opposed to her portfolio flows. This distinguishes our approach from the more standard demand-based literature that exploits opening positions at time t, that is, e.g., daily buy and sell orders on new contracts. To get a measure of total portfolio holdings, notice that our database reports the total daily opening and closing positions on every option contract. Opening orders represent shocks to demand (flows), while aggregate opening and closing positions represent changes to the holdings. In our theoretical framework, investors' portfolio holdings (not the incremental holdings) reflect their belief about the market. We thus aggregate opening and closing positions from issuance of a contract until date t.[15]

Figure 3 shows (the moving average of) our holdings measure for customers of 30-days expiring options across different moneyness. Since options are in zero-net supply, market makers' holdings are just the mirror image. Panel 3A displays the total aggregate positions on OTM calls and OTM puts per month. Overall, customers are long OTM puts (on average $+2.2\times10^4$ contracts/day) and short OTM calls (-0.9×10^4 contracts/day), with the investment in puts being larger in size than the corresponding in calls almost everywhere. In the run-up to the 2008 crisis, aggregate holdings for OTM put options spike in September 2007, then rapidly drop and then spike up again in August 2008. The two series have correlations -24%. We can argue that the plausible portfolio of the customers is long in the index, then they buy

[15] 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.

OTM puts for protection and sell (a smaller amount of) OTM calls to finance their hedging positions.

Panels 3B and 3C distinguish between mild-OTM options with $|\Delta| \in (0.2, 0.4]$ and deep-OTM options with $|\Delta| \in [0.1, 0.2]$. Customers always invest more in mild-OTM than in deep-OTM, but in some periods they take opposite net positions across these categories. For instance, customers were short deep-OTM puts during the Great Financial Crisis and between 2015 and 2018, and they were long deep-OTM calls at the end of 2007. The two pairs of time series do not perfectly move together (correlation between mild-OTM and deep-OTM holdings is 40% for calls and 67% for puts). Therefore, in some periods the demands for deep-OTM and mild-OTM puts will generate opposite effects in the customers' subjective expectations. In determining the final measure of $\mathbb{E}^i[R]$, the contribution from deep-OTM options is relatively hidden. First, on aggregate mild-OTM options dominate because they are more traded. Second, deep-OTM puts distort the customers' SDF in the very downside (or upside) region, where the probability mass is typically smaller than in the range of mild-OTM options.

Our holdings proxy does not exhibit any significant change in sign, except in a few periods for deep-OTM options. This contrasts with findings in Chen, Joslin, and Ni [2019] who find that market makers went from being net sellers of DOTM options to net buyers during the Great Financial Crisis in 2008 due to binding constraints. This property is more muted in our aggregate measure because of three reasons. First, we consider portfolio positions instead of demand shocks, as investors also buy and sell options to close previous contracts in their portfolio. Second, we consider a broader range of moneyness while excluding very deep-OTM options (with $|\Delta| < 0.1$). Third, we look at portfolios with only 30-days maturity to build consistent one-period SDFs, instead of aggregating data across different maturities.[16]

[16] We discuss differences in more detail in the Appendix, see Figure B.4.

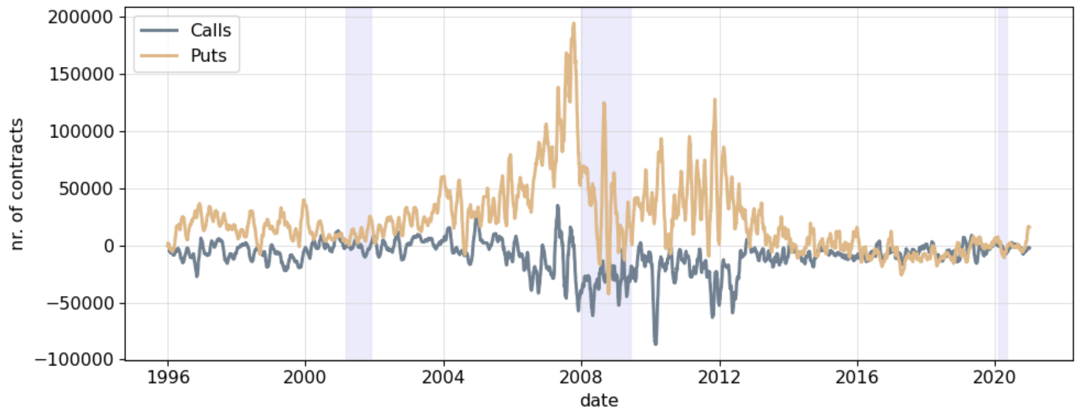

A. OTM Options

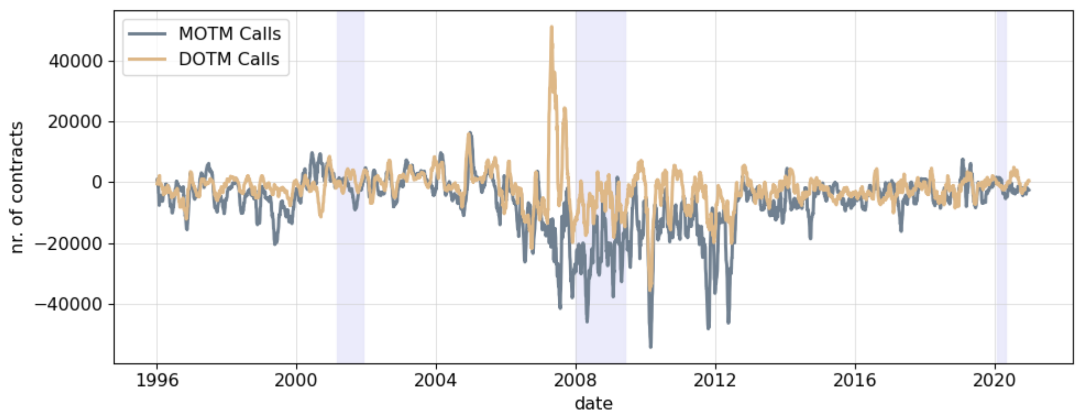

B. OTM Calls

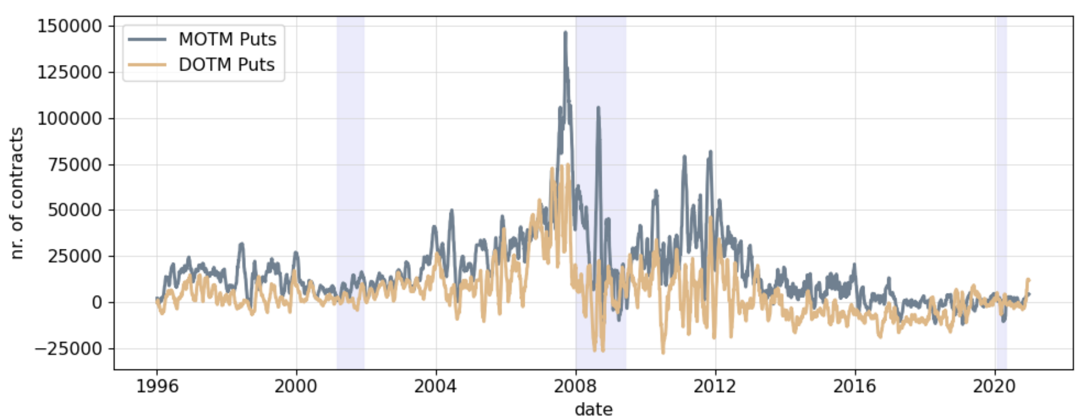

C. OTM Puts

Figure 3. Customers' Holdings of monthly OTM options

Notes: This figure plots the thirty day moving average of customers' portfolio holdings of OTM calls and puts, expiring in 30 days, from 1996 to 2020. Options holdings are the sum of opening and closing positions on the same contract that customers enter from issuance. Mild-OTM options («MOTM») have $|\Delta| \in (0.2, 0.4]$. Deep-OTM options («DOTM») have $|\Delta| \in [0.1, 0.2]$. Holdings in Panel 3A are the sum across moneyness levels shown in Panels 3B and 3C. Grey areas indicate NBER recession periods.

4.3 Subjective Expected Returns

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, say, market makers. Instead, we interpret our measure as a conservative lower bound on subjective expectations of heterogeneous agents. We start from our main identity (2): the variation of expected returns over time as perceived by different agents.

Since we do not know the true real-time investment in the S&P500, we can study the effect of option demand on investors' expected returns as a function of the wealth invested in the index. Empirical evidence suggests that investors hold the underlying and adjust their positions with OTM options. Therefore, we assume that they invest a fraction $\theta_0 \in [0,1]$ in the index and $1-\theta_0$ in the options. Let $\boldsymbol{\omega}$ be the weights of the observed portfolio of options, relative to the wealth allocated in the derivatives, so that $|\boldsymbol{\omega}_i'\mathbf{1}| = 1$. With this portfolio formulation, equation (2) becomes:

$$\mathbb{E}^i[R] = \theta_r + \theta_0 \mathbb{E}^{\mathbb{Q}}[R^2] + (1-\theta_0)\mathbb{E}^{\mathbb{Q}}[\boldsymbol{\omega}'\boldsymbol{R}R] \, , \tag{22}$$

where θ_r is the investment in the risk-free asset.[17] If investors are net long in options, then $\theta_r = 0$; if they are net short, the profits from selling derivatives become savings in the risk-free account. If $\theta_0 = 1$, we obtain a portfolio that is fully invested in the index; therefore, we recover the stochastic discount factor M^0, the benchmark SDF. Under this belief, expected returns are always counter-cyclical and positive, by construction.[18]

In light of this decomposition, investors' subjective expected returns can be interpreted as a leading term (the variance of the market return) plus a belief distortion whose size depends on the option portfolio adjustment. More formally, we can re-write this term as the covariance between the option portfolio and the index:

$$\mathbb{E}^{\mathbb{Q}}[\boldsymbol{\omega}'\boldsymbol{R}R] = \mathbb{C}ov^{\mathbb{Q}}(\boldsymbol{\omega}'\boldsymbol{R}, R) + \boldsymbol{\omega}'\mathbf{1} \, . \tag{23}$$

For instance, if the investor is long OTM puts, the correction term will be negative because the returns of OTM puts are negatively correlated with the return on the underlying. Thus, we expect her subjective expected return to be smaller than the benchmark. The investment in the index, θ_0, measures the relative impact between the variance term and the covariance. Intuitively, this is how in our framework portfolio adjustments are able to reflect belief distortions from the benchmark. It is straightforward to see that $\theta_0 < 1$ decreases the expected

[17] More specifically, $\theta_r = (1-\theta_0)(1-\boldsymbol{\omega}'\mathbf{1})$.
[18] See Figure B.5 in the Appendix.

return, $\mathbb{E}^i[R]$.

Figure 4 depicts the belief distortion $\mathbb{E}^{\mathbb{Q}}[\boldsymbol{\omega}'\boldsymbol{R}R]$ recovered from the data.[19]. It represents the maximum possible correction generated by the customers' positions in the options. Notice that it is below the risk-free rate almost everywhere, often negative with large peaks. The average is -0.7 annualized, therefore negatively contributing to customers' expected returns. It is also quite erratic, as it reflects both changes in returns and (especially) in option holdings $\boldsymbol{\omega}$, which are not persistent. For example, its standard deviation is 2.92 compared to 0.06 of the benchmark. The correlation between the two series is -12%. As expected, the time-varying position in the options gives rise to a globally almost acyclical pattern, with phases (like during and immediately after a crisis) that seem procyclical.

This extreme result is perhaps not surprising because we are considering the implausible case of investors who take significant positions in the options and none in the underlying. It is worth highlighting that these portfolios do not support valid SDFs, so they cannot be used to extract expected returns. In other words, $\mathbb{E}^{\mathbb{Q}}[\boldsymbol{\omega}'\boldsymbol{R}R]$ does not represent the subjective expected return of the investor who optimally chooses to hold only options ($\theta_0 = 0$).

To recover realistic $\mathbb{E}^i[R]$ we now build portfolios with a large and positive index investment. Recall from our earlier findings that customers are mostly long in OTM puts and short in OTM calls, as they hold the index and buy protection, potentially selling calls to finance their costly positions in the puts.

To ensure that the recovery does not lead to violations of the no-arbitrage condition, for every date, we select the minimum value of θ_0 such that the corresponding M_i is strictly positive in its domain.[20] We label such a threshold, $\bar{\theta}_0$. Then, we characterize investors' subjective expected returns when varying the fraction invested in the index, $\alpha \in [0, 1]$:

$$\theta_0 = \bar{\theta}_0 + \alpha \cdot (1 - \bar{\theta}_0) \, . \tag{24}$$

To explore different configurations, we keep α fixed so that the effective index investment θ_0 changes over time. Results are shown in Figure 5. Summary statistics are reported in Table 2.

From Figure 5B, we see that the threshold $\bar{\theta}_0$ (corresponding to case $\alpha = 0$) is always greater than 0 and sometimes close to 1, being 72% on average. Then, assuming $\bar{\theta}_0 = 0$ leads to unfeasible estimates of the stochastic discount factor. We can consistently extract subjective expected returns only on taking a significant long position in the underlying. For instance, between 2016 and 2019, the optimal customers' portfolio compatible with M_i is almost fully invested in the index ($\bar{\theta}_0 \approx 1$). The averages are reported in the last column of Table 2.

[19] By market clearing, market makers' positions will be $-\boldsymbol{\omega}$ and the belief distortion is the negative of Figure 4.

[20] In principle, since M_i may be not monotonic, there is also a maximum level of index investment, but it turns out in the data that this limit is always greater than 1.

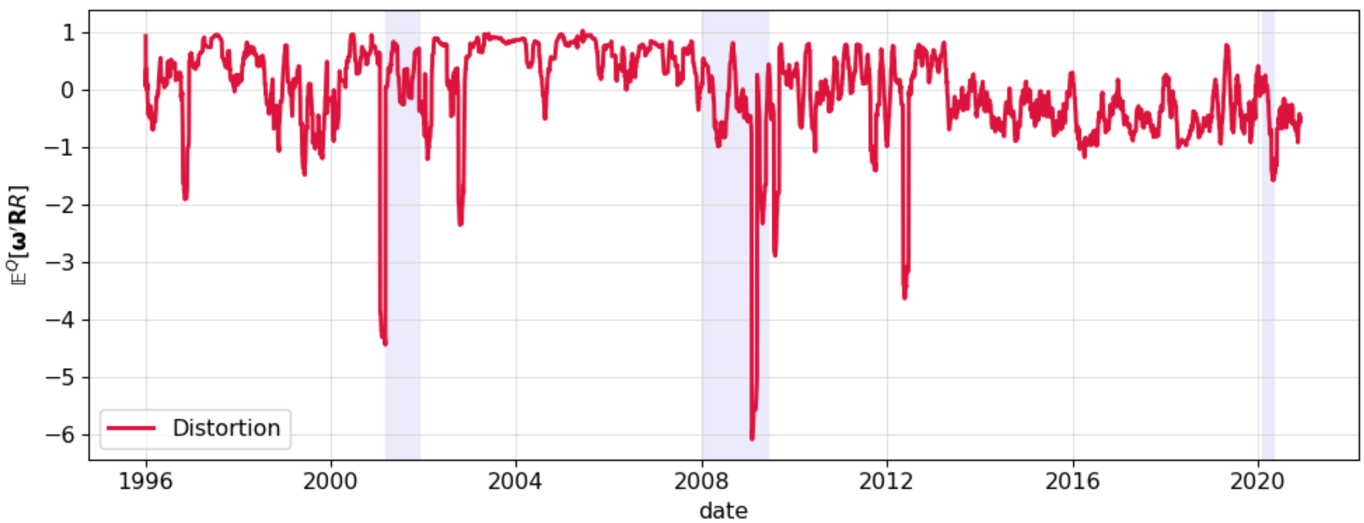

Figure 4. Belief distortion from Customers' Options Holdings
Notes: This figure plots the time-series of the belief distortion $\mathbb{E}^{\mathbb{Q}}[\boldsymbol{\omega}'\boldsymbol{R}R]$, where $\boldsymbol{\omega}$ are the observed Customers' holdings in options. Data are the 30-day moving average of daily holdings. Data runs from January 1996 to December 2020. Gray bars indicate NBER recessions.

From Figure 5A, we see that for customers, larger investments in options lead to smaller expected returns. The largest discrepancies emerge in periods of larger belief distortion $\mathbb{C}ov(\boldsymbol{\omega}'\boldsymbol{R}, R)$ (e.g. during crises) and in periods of smaller index investment $\bar{\theta}_0$ (e.g. between 2003 and 2006).

Volatility increases because subjective expectations are exposed to frequent portfolio rebalancing. This is confirmed in Table 2, where we report persistence coefficients (AR(1)). Our measure of $\mathbb{E}^i[R]$ is more noisy because it absorbs shocks both to option prices and to portfolio rebalancing (i.e. option demand). This second effect is obviously absent in $\mathbb{E}^0[R]$. Notice also that the correlation with the benchmark remains positive but the minimum value only achieves 9%. This confirms that the cyclicality of subjective expected returns may not be perfectly aligned with M^0.

Intuitively, our results can be explained by customers' portfolio compositions. Customers are long in the index and they buy protection through OTM puts, potentially selling OTM calls to finance their long (and costly) hedging positions. As a consequence, their exposure to market risk is more mitigated, coherently with having more pessimistic beliefs (relative to investors who are fully invested in the market) and therefore resulting in lower expected market returns.

We follow an analogous approach to recover market makers' expected returns. Empirical evidence suggests that intermediaries trade delta-hedged strategies; see, e.g., Amayaa et al. [2024] and Baltussen, Jerstegge, and Whelan [2024]. To account for that, we add an extra

Table 2. Summary statistics of Customers' expected returns $\mathbb{E}^i[R]$ different α.

α	mean	std	min	median	max	corr (%)	AR(1)	index (avg.)
1	1.082	0.058	1.017	1.069	1.521	100	0.82	1
95%	1.061	0.056	0.977	1.049	1.472	96	0.77	0.99
90%	1.041	0.058	0.911	1.032	1.433	85	0.69	0.97
80%	1.004	0.071	0.792	1.006	1.366	59	0.58	0.94
50%	0.913	0.124	0.520	0.933	1.337	20	0.51	0.86
0	0.810	0.194	0.258	0.838	1.345	5	0.49	0.72

Notes: This table reports summary statistics for the 30-day moving average of annualized expected returns recovered by different portfolios built on the observed options positions of the Customers. «corr» is the correlation with the benchmark $\mathbb{E}^0[R]$. «index (avg.)» is the average investment in the index for each time-series. Portfolios differ across the amount of wealth invested in the index, expressed as a function of α. The case $\alpha = 1$ is the benchmark recovered by M^0 ($\theta_0 = 1$). The case $\alpha = 0$ corresponds to the minimum investment in the index compatible with the no-arbitrage condition ($\theta_0 = \bar{\theta}_0$). Data is daily, values are annualized. Data runs from January 1996 to December 2020.

component in the underlying coming from delta-hedging, which varies over time but depends only on the observed option positions to market makers' portfolio. Finally, we determine the minimum admissible investment in the index $\bar{\theta}_0$ for the delta-hedged option portfolio at every point in time.[21] Since market makers take opposite positions of customers, they are typically net short OTM puts and long OTM calls, implying a negative delta hedge (- 54% on average).

Results are shown in Figure 6 and Table 3. The time-series of the recovered expected returns are very similar to each other and to the benchmark. Compared to customers, they are more counter-cyclical, less volatile, and more persistent. The belief distortion due to the positions in the options is positive for market makers, thus we would expect higher expected returns. However, the large negative delta-hedging compensates for this effect. Indeed, the covariance correction $\mathbb{C}ov^{\mathbb{Q}}(\boldsymbol{\omega}'\boldsymbol{R}, R)$ is a first-order perturbation in $\boldsymbol{\theta}_i$ that delta-hedging should neutralize. Also, market makers have non-zero positions in the risk-free asset where they invest the profits made by selling the index and/or the options (while for customers we find mostly $\theta_r = 0$). Figure 6B shows the fraction of market makers' wealth invested in the index. Compared to customers (Figure 5B), it is always smaller and may become negative, with $\bar{\theta}_0$ close to 0. Lower α, imply smaller and more volatile index investments. This is to balance the volatile option portfolio $\boldsymbol{\omega}$. As a result, the final expected return time series is less volatile, more persistent, and similar to the benchmark. To some extent, for sufficiently large values of α, the SVIX itself can be interpreted as a measure of market makers' view.

We conclude by emphasizing that equation (2) holds only for the unconstrained agent

[21]That is, for every date t, let Δ_0 be the wealth fraction invested in the index due to the delta-hedging, relative to the total wealth invested in the options. Therefore $|\boldsymbol{\omega}'\mathbf{1} + \Delta_0| = 1$ and typically $\Delta_0 < 0$. The total portfolio of market makers will invest $\theta_0 + (1 - \theta_0)\Delta_0$ in the index, $(1 - \theta_0)\boldsymbol{\omega}$ in the options, the rest in the risk-free.

Table 3. Summary statistics of Market Makers' expected returns $\mathbb{E}^i[R]$, different α.

α	mean	std	min	median	max	corr (%)	AR(1)	index (avg.)
1	1.082	0.058	1.017	1.069	1.521	100	0.82	1
95%	1.078	0.056	1.011	1.067	1.493	100	0.83	0.93
90%	1.074	0.053	1.006	1.065	1.465	99	0.85	0.86
80%	1.045	0.049	0.989	1.059	1.410	97	0.87	0.72
50%	1.045	0.043	0.936	1.043	1.262	78	0.86	0.30
0	1.012	0.048	0.779	1.021	1.104	29	0.69	-0.40

Notes: This table reports summary statistics for monthly averaged annualized expected return recovered by different portfolios built on the observed delta-hedged options positions of Market Makers. «corr» is the correlation with the benchmark. «index (avg.)» is the average investment in the index for each time series, including the hedging. Portfolios differ from the amount of wealth invested in the index, expressed as function of α. Case $\alpha = 1$ is the benchmark recovered by M^0 ($\theta_0 = 1$). Case $\alpha = 0$ corresponds to the minimum investment in the index compatible with the no-arbitrage condition ($\theta_0 = \bar{\theta}_0$). Data is daily, values are annualized. Data runs from January 1996 to December 2020.

who optimally invests in the market. Intuitively, if the agent does not face constraints, then her optimal portfolio $\boldsymbol{\theta}_i$ is a reflection of her own belief. This assumption is likely to fail for market makers, because they merely act as a counterparty to customers' trades (and they are likely subject to many constraints). Their profits do not come from trading the optimal option portfolio according to their belief. Therefore, the expected returns we recover may reflect a spurious belief of market makers.

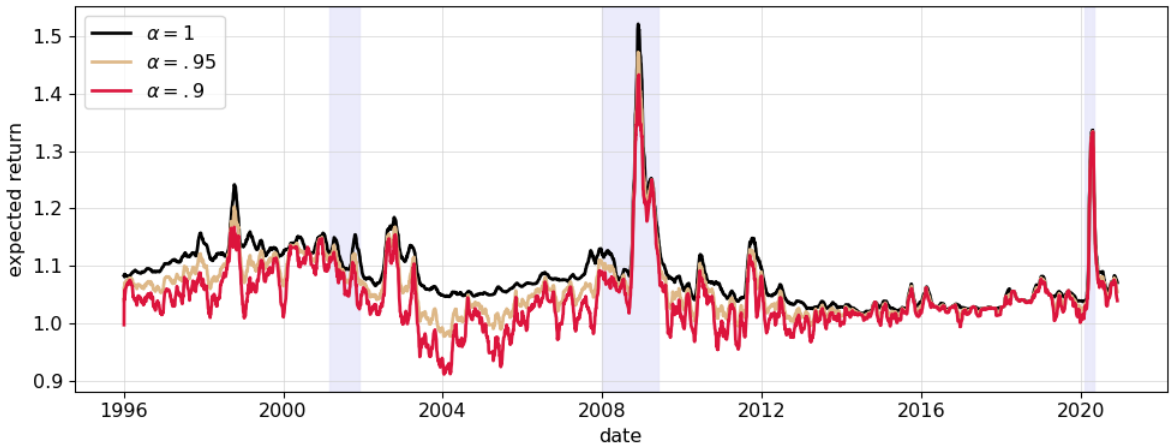

A. Lower Bound Customers Subjective Expected Returns

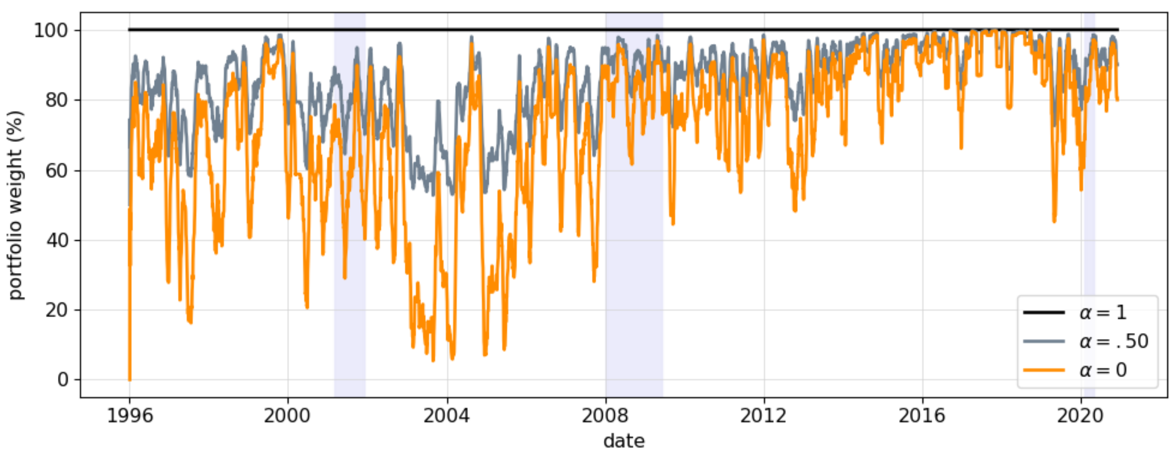

B. Customers' Investment in the Index

Figure 5. Customers' subjective expected returns different α

Notes: Panel 5A plots the time-series of the expected market return implied by customers' options holdings for different levels of index investment. Case $\alpha = 1$ displays the monthly expected return recovered by the benchmark $M^0 = 1/R$. Frequency is daily, horizon is monthly, values are annualized. Panel 5B plots customers' investment in the index as fraction of their total wealth, for various choices of α. Case $\alpha = 0$ corresponds to the minimum investment $\bar{\theta}_0$ such that $M_i > 0$ in its support. Data are monthly moving averages. Gray bars indicate NBER recessions.

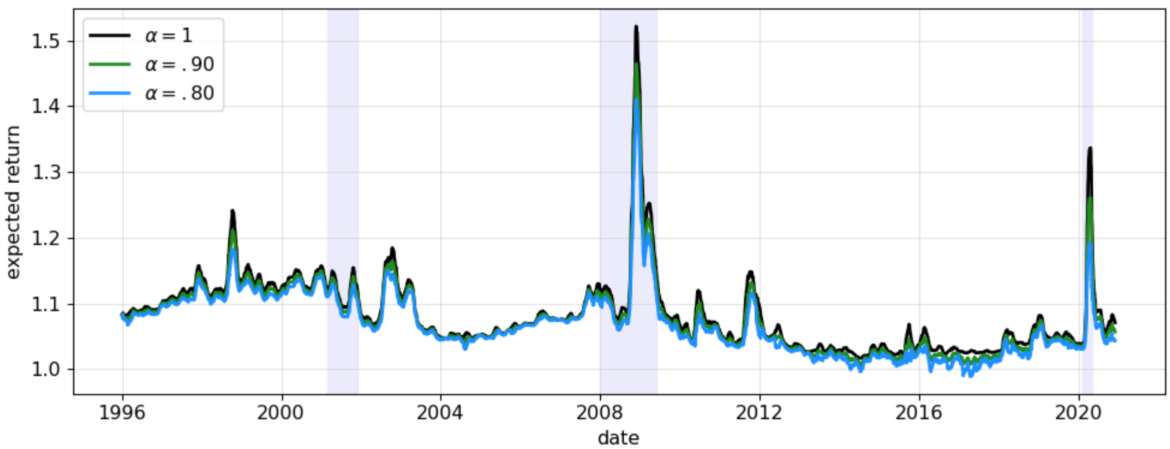

A. Lower Bound Market Makers Subjective Expected Returns

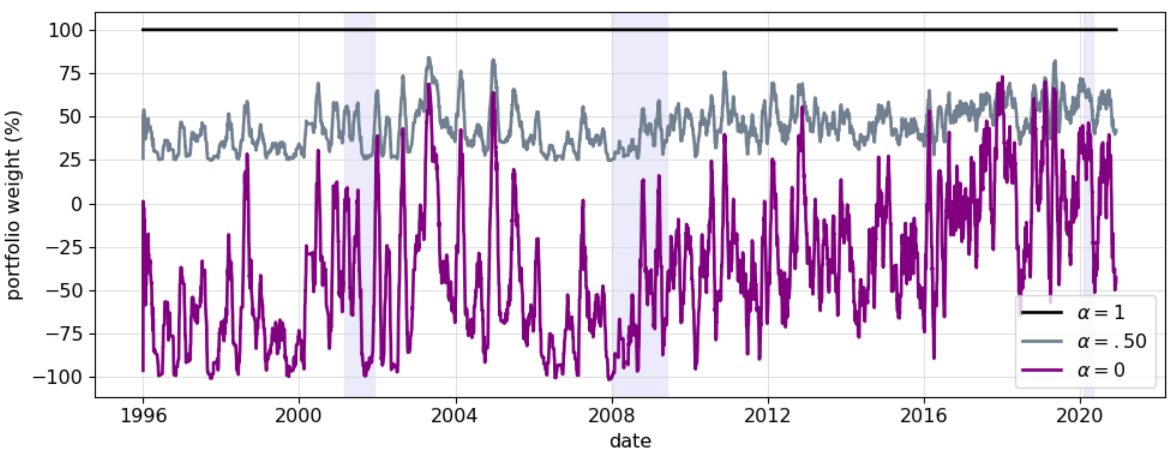

B. Market Makers' Investment in the Index

Figure 6. Market Makers' subjective expected returns on varying of α and $\bar{\theta}_0$

Notes: Panel 6A plots the time-series of the expected market return implied by market makers' options holdings fully delta-hedged, for different levels of index investment. The case $\alpha = 1$ displays the monthly expected return recovered by the benchmark $M^0 = 1/R$. Frequency is daily, horizon is monthly, values are annualized. Panel 6B plots the market makers' total investment in the index (including delta-hedging) as fraction of their total wealth, for various choices of α. The case $\alpha = 0$ corresponds to the minimum investment such that $M_i > 0$ in its support. Data are 30-days moving averages. Gray bars indicate NBER recessions.

4.3.1 Consensus Belief

We now turn to extract the consensus belief among investors following Corollary 1. Recall that if we include both customers and market makers into the computation, equation (13) shows that the consensus belief will be equivalent to the expected return recovered under the benchmark M^0, because of market clearing in the options market. This is confirmed in Figure 7A. Thus, it is more reasonable to interpret the benchmark as the aggregate market view when all the market participants are unconstrained, rather than the subjective view of an investor who optimally decides to not invest in the options.

Alternatively, we can study how much customers' and market makers' deviate from the consensus. Intuitively, this measures the degree of belief heterogeneity in the market. To this end, we use the following (scaled) dispersion measure:

$$\mathbb{D}_t(R) := \frac{1}{\sqrt{\mathbb{V}ar^{\mathbb{Q}}(R)}} \sum_{i \in \mathcal{I}} w_i \left| \mathbb{E}^i[R] - \bar{\mathbb{E}}[R] \right| . \tag{25}$$

We divide by the forward variance of the index to isolate the belief heterogeneity from a mechanical increase due to market volatility. The ensuing dispersions are shown in Figure 7B. The dispersion is mainly driven by customers' expected returns. It is larger in the first part of the sample up to 2012, then it starts decreasing when investors have smaller leverage in the options. Even accounting for the market variance, the belief dispersion has a correlation 49% with the benchmark and it peaks during crisis. The degree of heterogeneity is naturally larger when the option investment is larger, and it goes up with the covariance between R and the excess return $R^i - R^{\mathcal{I}}$.

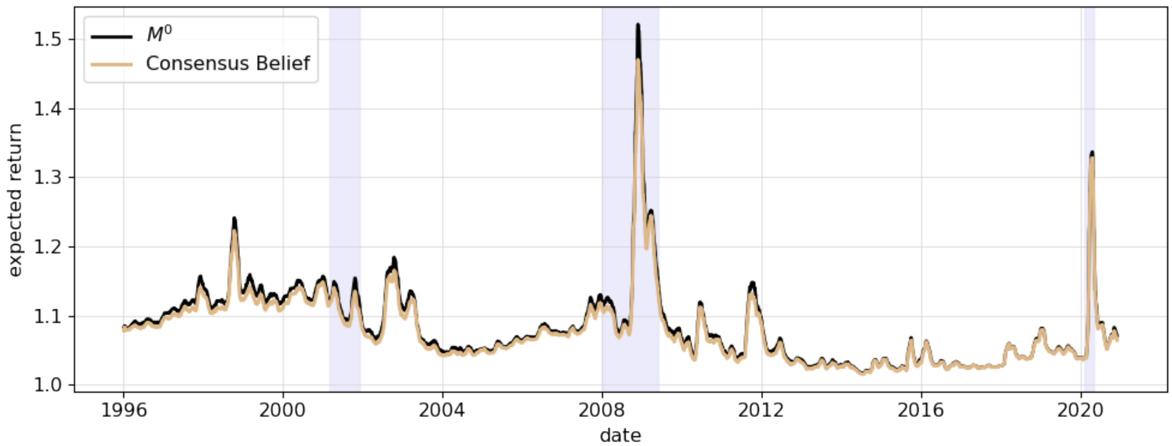

A. Consensus Belief between Customers and Market Makers

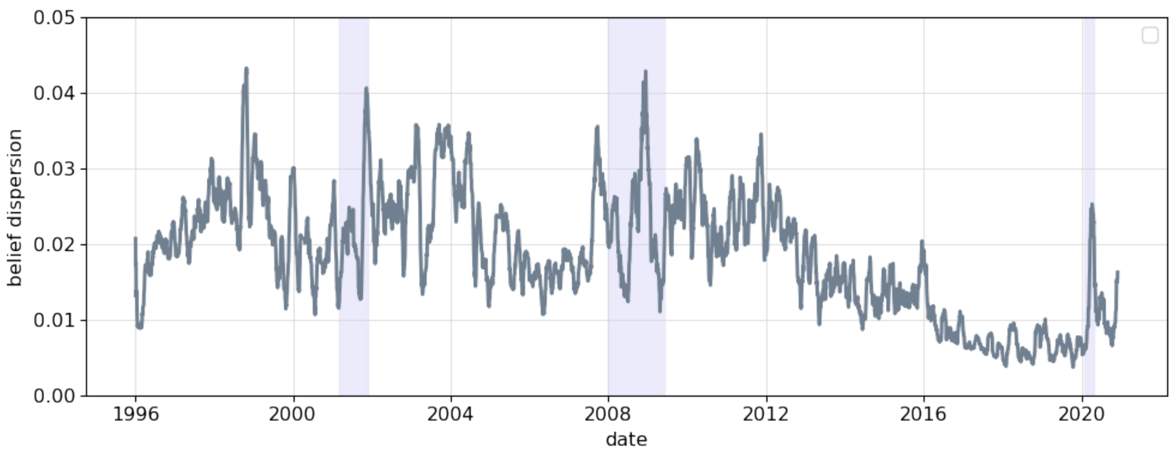

B. Belief Dispersion in the market

Figure 7. Belief aggregation between Customers and Market Makers
Notes: Panel 7A plots the time-series of the consensus belief between customers and market makers, namely the value-weighted average of their subjective expected returns. Panel 7B plots the time series of the belief dispersion, namely the value-weighted distance between the consensus belief and the subjective returns of customers and market makers, rescaled by the market forward volatility. We consider $\alpha = 0.9$. Frequency is daily, horizon is monthly. Values are annualized 30-day moving averages. Gray bars indicate NBER recessions.

4.4 Bounds on Subjective Expected Returns

The main role of market makers consists of liquidity provision. Market makers are required to quote both buy and sell prices, enabling continuous liquidity in the market. To fulfill this role, they must absorb buying and selling pressures, which can lead to imbalances in their inventories. These deviations from optimal inventory levels expose market makers to risks that necessitate hedging. However, hedging is both costly and challenging due to market imperfections, such as discrete trading and sudden price jumps in the underlying asset. Consequently, these inventory risks and hedging costs influence an option's liquidity and are reflected in the bid-ask spread, which serves as compensation for market makers providing liquidity. In the following, we use options' bid-ask spreads to bound expected returns.

To this end, we solve the linear optimization problem in equation (6) to determine the most conservative and the maximum expected market return compatible with investors' observed positions. We set δ equal to half of the average bid-ask spread in the options cross-section at every date t. As before, we average daily data to obtain monthly averages.

To make the results comparable, we compute the upper and lower bound on the belief distortion $\mathbb{E}^{\mathbb{Q}}[\boldsymbol{\omega}'\boldsymbol{R}R]$, namely we optimize on the option weights $\boldsymbol{\omega}$ by keeping the index investment θ_0 equal to 0. Figure 8 shows the results for customers[22]. The red area is the time-varying range of admissible belief distortion that are compatible with customers' observed positions. Since the index investment is fixed, there is no possibility to increase option leverage, but only to rearrange portfolio weights across options up to some limits given by the value of δ. Therefore, the lower bound is close to the case without measurement error (Figure 4) because this belief correction has already a negative impact on the subjective expected return. Intuitively, the lower bound represents the expectation of the "most pessimistic" investor —- or equivalently, it represents the "worst-case" expectation that investors can formulate. In contrast, the upper bound is often greater than 1 and contributes to increasing the perceived expected return.

From our illustrative example discussed earlier, we expect that the lower bound is attained with a portfolio that hedges volatility risk with long positions in calls and puts. However, the upper bound is most likely supported by portfolios that tend to have short positions in calls and puts. This explains why the lower bound is more procyclical (the correlation with $\mathbb{E}^0[R]$ is -34%) than the case without measurement errors.

[22] By market clearing, the bounds for market makers will be the negative of the plot.

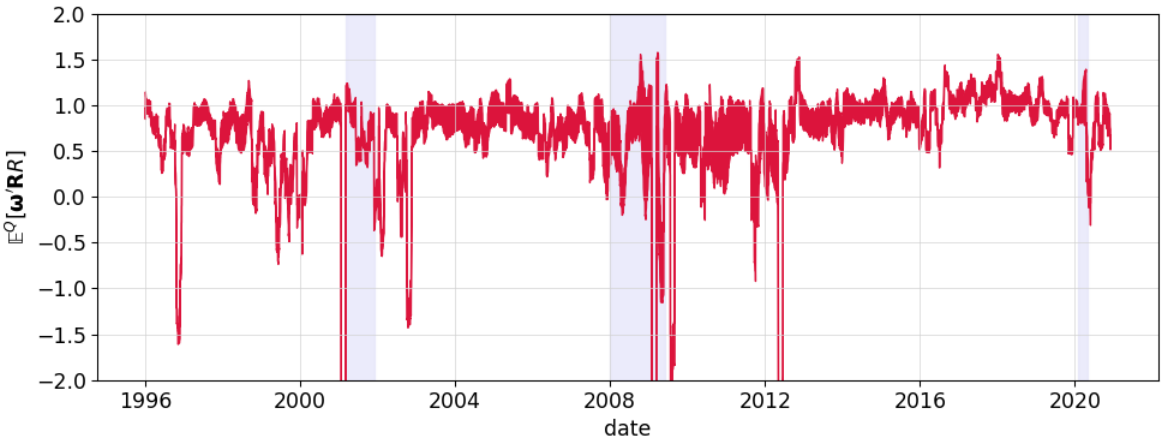

Figure 8. Admissible values for Customers' belief distortion $\mathbb{E}^{\mathbb{Q}}[\boldsymbol{\omega}'\boldsymbol{R}R]$
Notes: This figure plots the range of admissible maximum belief distortions $\mathbb{E}^{\mathbb{Q}}[\boldsymbol{\omega}'\boldsymbol{R}R]$ as recovered by SDFs compatible with the Customers' observed positions in options. Lower bound and upper bound are computed by solving (6) and imposing no investment in the index nor in the risk-free asset. Frequency is daily and horizon is monthly. Data runs from January 1996 to December 2020. Gray bars indicate NBER recessions.

4.5 Survey Data and Time-Series Properties

We now want to investigate the relation of our expected return measures with survey measures and study their cyclicality properties in more detail. To make our measures comparable with survey measures, we construct excess returns and subtract the one-month Treasury Bill rate.[23] We use the following three surveys about expected market returns: The individual investor series by Nagel and Xu [2023], the Graham and Harvey CFO survey [Ben-David, Graham, and Harvey 2013], and the Livingston Survey available from the Federal Reserve Bank of Philadelphia. Nagel and Xu [2023] extend the UBS/Gallup survey backward and forward using other surveys such as the Conference Board survey, and the Michigan Survey of Consumers and is available at the quarterly frequency.[24] The Graham and Harvey survey polls financial officers about the one-year expected return on the S&P500 and is also available at the quarterly frequency. Finally, the Livingston Survey is released in June and December every year and polls economists at financial, non-financial, and academic institutions, as well as labor organizations, government, and insurance companies.

Figure 9 plots expected return measures for market makers and customers together with three survey measures. As with regards to the overall co-movement, we find that, for $\alpha = 0.9$, the correlation between market makers' expected (excess) return and the Livingston Survey

[23]Notice that most surveys measure expected stock returns over a one-year horizon. We do not observe many options with maturity of one year. We therefore use a one-month horizon (as before).
[24]We download data from Zhengyang Xu's webpage.

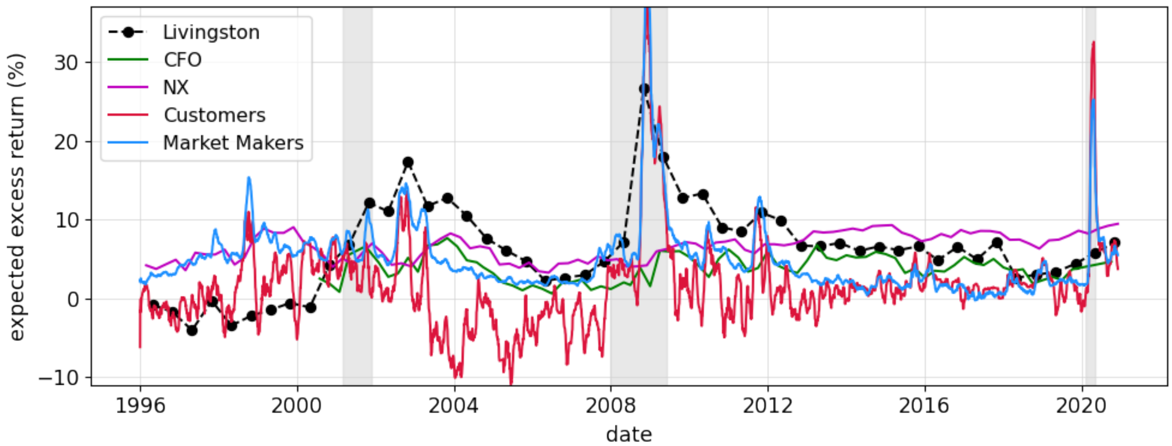

Figure 9. Subjective Expected Market Return from Survey Data, Customers and Market Makers
Notes: This figure plots the time-series of the expected market premium (p.a. in %) for market makers and customers (with $\alpha = 0.9$) together with survey data from Nagel and Xu [2023] (NX), the Graham and Harvey survey (CFO), and the Livingston Survey. Data runs from January 1996 to December 2020. Gray bars indicate NBER recessions.

is 54% (42% for customers' expected return), 1% with the CFO survey (-11% for customers' expected returns), and -22% with the Nagel and Xu [2023] measure of expected returns (2% for customers' expected return). Our findings echo Dahlquist and Ibert [2024] who study surveys from asset managers and find that their expectations behave quite different from those of retail investors. Similarly, market makers and their customers in the SPX market are highly sophisticated hedge funds and banks and their expectations about the market's return are highly negatively correlated while the correlation with financial institutions (Livingston Survey) is positive.

We now turn to studying market makers' and customers' expected return determinants. A large literature has studied drivers of investors' expectations by estimating time-series regressions from expected returns on standard predictors of realized returns. Interestingly, most survey measures do not load significantly on standard predictors, see, e.g., Nagel and Xu [2023]. We follow this literature and regress the expected return measures on the S&P500 P/D ratio, the consumption wealth ratio (CAY) of Lettau and Ludvigson [2001], as well as net equity expansion (NTIS) from Welch and Goyal [2008], calculated as the ratio of twelve-month moving sums of net issues by NYSE-listed stocks divided by the total market capitalization of NYSE stocks at the end of the twelve-month window. In all of our regressions, we also include past realized returns as in Greenwood and Shleifer [2014]. Panel A of Table 4 reports the results.

None of the standard-returned predictor coefficients is statistically significant. This result is in line with Nagel and Xu [2023], who also find that subjective risk premia do not load on standard predictors. However, past realized returns are highly statistically significant on subjective returns of market makers and customers with coefficients of similar size.

Table 4. Determinants and Cyclicality Expected Returns

	Panel A: Determinants						Panel B: Cyclicality							
	CAY		DP		NTIS		IP		TERM		DEFAULT		F1	
Expected Returns: Market Makers														
Const	0.08	0.09	0.07	0.08	0.08	0.11	0.07	0.10	0.08	0.09	0.09	0.09	0.07	0.08
(p-value)	(0.00)	(0.00)	(0.01)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Coeff	1.88	1.81	-1.29	-2.32	1.36	3.06	1.21	-3.05	-0.70	-0.96	4.88	2.84	2.95	1.28
(p-value)	(0.21)	(0.24)	(0.56)	(0.13)	(0.60)	(0.08)	(0.71)	(0.20)	(0.57)	(0.31)	(0.03)	(0.23)	(0.05)	(0.23)
R^e_{past}		-3.60		-4.35		-4.57		-4.58		-3.75		-2.31		-2.87
(p-value)		(0.01)		(0.01)		(0.01)		(0.00)		(0.02)		(0.01)		(0.01)
Adj. R^2	0.11	0.29	0.03	0.28	0.02	0.28	0.00	0.22	0.01	0.20	0.18	0.22	0.13	0.20
Expected Returns: Customers														
Const	0.04	0.06	0.04	0.05	0.03	0.06	0.04	0.06	0.05	0.07	0.05	0.06	0.04	0.05
(p-value)	(0.00)	(0.00)	(0.01)	(0.00)	(0.01)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Coeff	0.17	0.11	0.57	-0.33	-1.24	0.15	3.03	-0.56	-1.28	-1.54	4.54	2.20	3.00	1.33
(p-value)	(0.72)	(0.82)	(0.75)	(0.77)	(0.44)	(0.89)	(0.16)	(0.69)	(0.26)	(0.05)	(0.01)	(0.28)	(0.01)	(0.16)
R^e_{past}		-3.70		-3.80		-3.75		-3.87		-3.84		-2.65		-2.87
(p-value)		(0.00)		(0.00)		(0.00)		(0.00)		(0.00)		(0.00)		(0.00)
Adj. R^2	-0.00	0.24	0.00	0.24	0.02	0.24	0.05	0.24	0.03	0.29	0.19	0.26	0.17	0.26

Notes: This table reports estimated coefficients from regressing market makers' and customers' expected returns on determinants and measures of cyclicality. Data runs from January 1996 to December 2020.

The negative slope coefficient is interesting, because most survey measures load positively on past realized returns. For example, Greenwood and Shleifer [2014] show that when the returns realized in the past are high, investors expect higher returns going forward. The authors interpret this finding as evidence for extrapolation. Using subjective expected returns of market makers and customers, we find that they actually expect lower expected returns.

As discussed earlier, the cyclicality properties of expected returns of market makers' and customers may depend on the composition of their portfolios. To study this relation more formally, we run time-series regressions of expected returns on measures of cyclicality. To proxy for cyclicality, we take industrial production growth (IP), the 10-year minus 3-month Treasury term spread (TERM), the default spread defined as the difference between Moody's BAA and AAA corporate bond yields (DEFAULT), and the real factor of Ludvigson and Ng [2009] (F1). Moving to Panel B in Table 4, we find that, with the exception of the term spread, none of the regression coefficients is statistically significant when controlling for past realized returns. We therefore conclude that the cyclicality properties of subjective expected returns of market makers and customers is muted.

5 Subjective Measures of Risk

We can now study subjective risk as described in equation (21). The belief correction term for the second moment of the market return is proportional to $\mathbb{C}ov^{\mathbb{Q}}(\boldsymbol{\omega}'\boldsymbol{R}, R^2)$. Long positions in

Table 5. Summary statistics of Customers' subjective volatility and Sharpe ratio, for various choices of α.

α	mean	std	min	median	max	corr (%)
Volatility						
1	0.221	0.104	0.094	0.190	0.934	91
90%	0.253	0.103	0.093	0.244	0.931	86
80%	0.269	0.105	0.091	0.265	0.919	81
50%	0.286	0.110	0.087	0.279	0.872	72
0	0.281	0.133	0.210	0.260	1.221	57
Sharpe ratio						
1	0.226	0.068	0.124	0.218	0.523	76
90%	0.060	0.134	-0.351	0.790	0.483	36
80%	-0.054	0.200	-0.821	-0.016	0.481	17
50%	-0.341	0.390	-1.517	-0.269	0.477	-4
0	-0.912	0.850	-3.623	-0.733	0.476	-17

Notes: This table reports summary statistics for the time-series of 30-days moving averages of subjective volatility and Sharpe Ratio recovered from the observed positions of customers, for different choices of α. «corr» is the correlation with the benchmark $\mathbb{E}^0[R]$. Values are annualized. Data runs from January 1996 to December 2020.

both OTM calls and puts will increase investor's exposure to the underlying volatility. Therefore, we expect the subjective volatility perceived by the customers to increase for $\alpha \to 0$, while the risk premium decreases. Consequently, the subjective Sharpe ratio will decrease when customers invest proportionally more in their option portfolio. The results shown in Figure 10 and Table 5 confirm this intuition.

The larger discrepancies between the volatility recovered by M^0 and those with M_i occur during normal times, because there is proportionally more leverage in the options and the benchmark volatility already spikes during crisis. That is why the Sharpe ratio is significantly lower during during normal times and closer to the benchmark in periods of financial distress. This is in line with the procyclical trends we have already observed for the subjective expected returns when α is small. For example, the subjective Sharpe ratio with $\alpha = 0.9$ is often negative, particularly before 2000 and between 2003 and 2008, while it approaches the benchmark case during the Great Financial Crisis and in Covid. The correlation of the Sharpe ratio time series with $\mathbb{E}^0[R]$ goes from 76% (with $\alpha = 1$) to -17% (with $\alpha = 0$). In general, both low Sharpe ratios reflect investors' conservative view of the market.

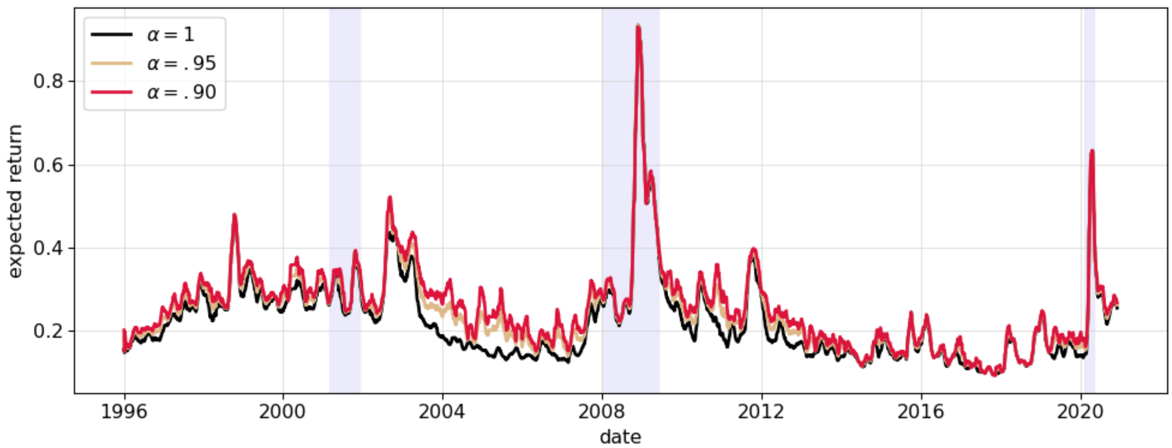

A. Subjective Volatility

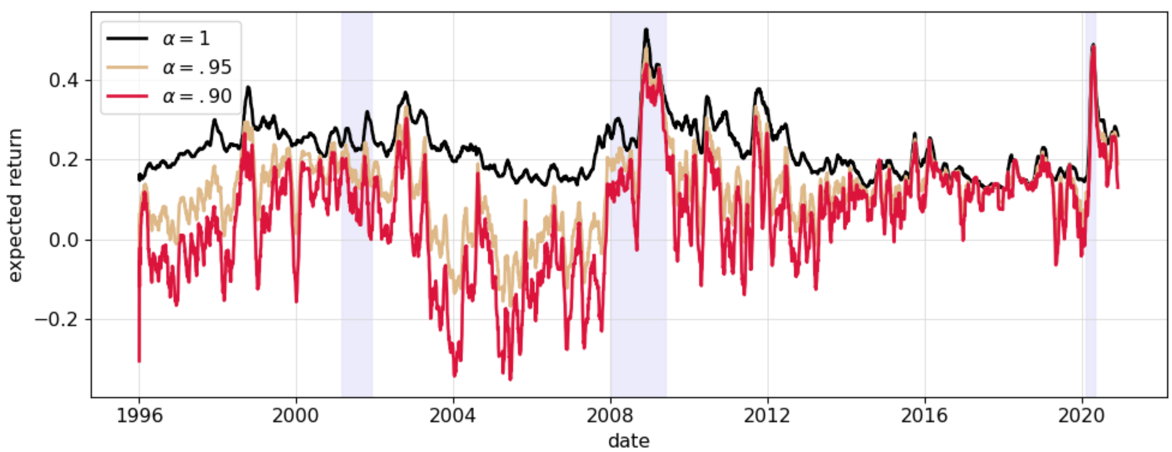

B. Subjective Sharpe Ratio

Figure 10. Subjective Volatility and Sharpe Ratio of Customers
Notes: This figure plots the time-series of the subjective volatility (upper panel) and the subjective Sharpe Ratio (bottom panel) as recovered through the SDF supported by the observed positions of customers, for different levels of underlying investment indexed by α. Frequency and horizon are monthly, values are annualized. Gray bars indicate NBER recessions.

6 Demand-Based SDFs

As our final exercise, we now recover the SDFs of our two investor groups and show how they depend on their portfolio compositions. To this end, we show the SDF shapes and the customers' portfolio weights[25] on options in two complementary situations, summarized in Figures 11 and 12. Their optimal portfolio is selected with an index investment equal to the minimum admissible level ($\alpha = 0$). Obviously, the larger the index investment, the more M_i becomes similar to $M^0 = 1/R$.

On both dates, we find that the investment in puts is larger than in calls, but the type of investment is opposite. Figure 11 shows a day in October 2012 when customers are long in puts and short in calls. We already know that on average this reflects a typical customers' portfolio of OTM options. In this case, your usual SDF does not decrease monotonically (as M^0). More specifically, it is increasing in the downside region because the risk has been hedged by buying protection through the OTM puts; it is increasing in the upside region because they are now exposed to losses in states of high market returns by selling OTM calls. The SDF of the market makers should have a complementary form, but appears flatter and more similar to M^0 due to the delta hedging.

While this reflects the typical pattern, we can observe many other interesting SDF shapes as a function of investors' portfolio compositions. Figure 12 is a prime example of the case where customers have taken overall negative positions in calls and more diversified positions in puts. The overall volume in calls is also greater than in puts. The picture refers to March 30, 2020 during the Covid crisis. This echoes previous findings in Chen, Joslin, and Ni [2019], who argue that while market makers are net providers of insurance in normal times, they become net demanders in bad times when their financial constraints bind. Although this effect is not evident at the monthly frequency, it can be recovered at more granular frequency. Figure 11: shows an opposite pattern where customers' SDF is now increasing in the upside and decreasing in the downside risk, exhibiting a clear U-shaped form. Market makers' SDF is conversely peaked in the middle, and decreasing towards extreme values in both directions. Unlike Figure 11, the hump shape is now more evident, although the index investment is still big and negative (as well as the delta hedging component).

[25] Since options are in zero net supply, market makers' aggregated holdings are just the mirror image.

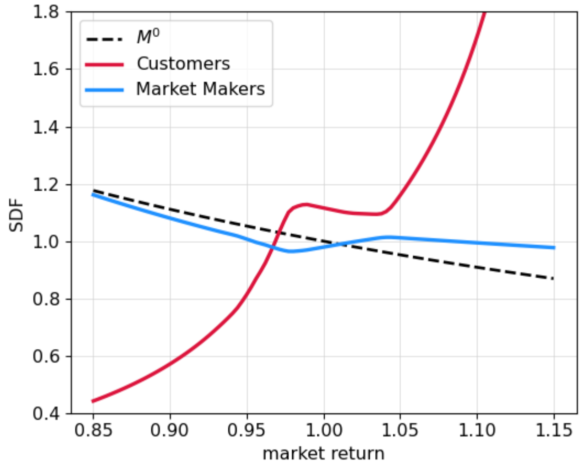

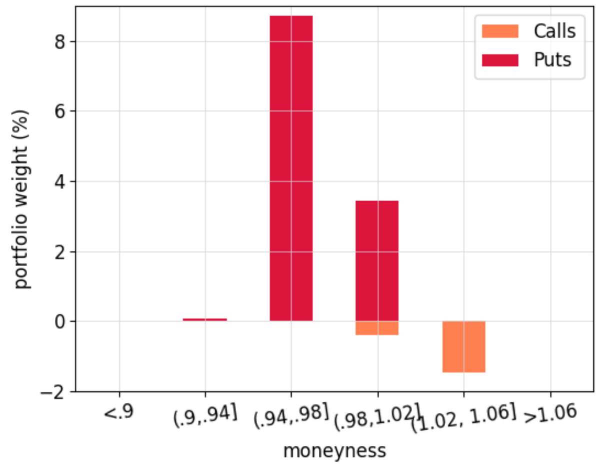

A. SDF **B.** Customers' Option Investment

Figure 11. Recovered SDF and Option Portfolio Weights
Notes: This Figure plots the stochastic discount factor recovered from customers and market makers' observed positions in 30-days expiring options on October 22, 2012 (left panel), and the corresponding distribution of the customers' portfolio weights for OTM calls and puts across different levels of moneyness (right panel). The index investment is equal to $\bar{\theta}_0$ + delta-hedging (only for market makers), resulting in 90% for customers and -98% for market makers. Portfolio weights are summed in every bin.

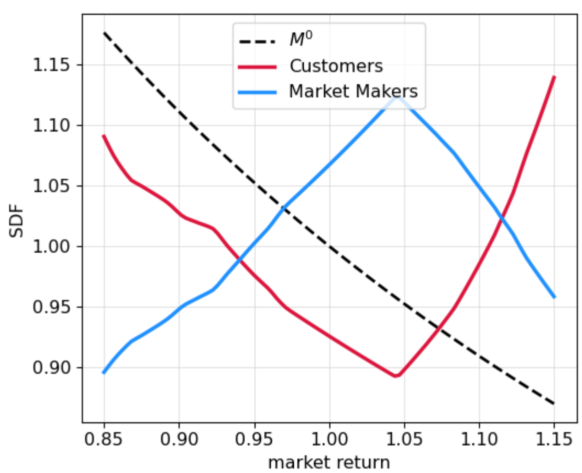

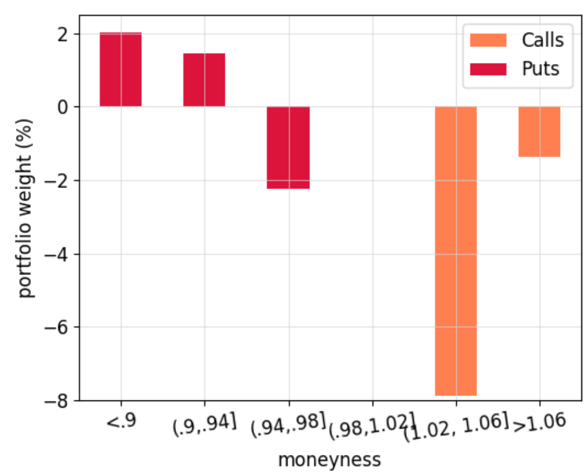

A. SDF **B.** Customers' Option Investment

Figure 12. Recovered SDF and Option Portfolio Weights
Notes: This Figure plots the stochastic discount factor recovered from customers and market makers' observed positions in 30-days expiring options on March 30, 2020 (left panel), and the corresponding distribution of the customers' portfolio weights for OTM calls and puts across different levels of moneyness (right panel).The index investment is equal to $\bar{\theta}_0$ + delta-hedging (only for market makers), resulting in 92% for customers and -106% for market makers. Portfolio weights are summed in every bin.

7 Conclusion

A large literature has documented different time series properties of survey measures of expected returns depending on the level of sophistication. Our findings show that even among investors with very high levels of sophistication such as hedge funds (who can act both as liquidity providers and demander), expected returns can differ wildly. Measures of expected returns that uniquely rely on pricing information are ill suited to explain such a large heterogeneity.

In this paper, we propose a theoretical framework for recovering investors' beliefs using demand-based data. Information about investors' holdings allows us to recover possible 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 and subjective Sharpe ratios vary significantly between 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 as they depend on the structure of the underlying portfolio and on the state of the economy.

While market makers' expected returns are highly correlated with survey measures of sophisticated agents, they are not correlated with households' expectations. The opposite holds true for customers. Finally, we find that the cyclicality properties of subjective expected returns are muted.

Our setting abstracts from frictions such as trading costs. However, the sensitivity of portfolio holdings to expectations depends on such costs. For example, Giglio et al. [2021] show that such sensitivity increases as investors face lower costs. In future work, we plan to jointly model survey data, holdings, and prices in the presence of frictions.

A Proofs and Derivations

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

$$\mathcal{U}(f) := \sup_{\boldsymbol{\theta}\in\Theta} \mathbb{E}^{\mathbb{Q}}[R_\theta f(R)] \geq \mathbb{E}^i[f(R)] \,, \tag{26}$$

and

$$\mathcal{L}(f) := \inf_{\boldsymbol{\theta}\in\Theta} \mathbb{E}^{\mathbb{Q}}[R_\theta f(R)] \leq \mathbb{E}^i[f(R)] \,. \tag{27}$$

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

$$\mathbb{E}^i[f(R)] = \mathbb{E}^{\mathbb{Q}}[R_\theta f(R)] \,, \tag{28}$$

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

$$\mathcal{L}(f) := \inf_{\boldsymbol{\theta}\in\Theta} \mathbb{E}^i[f(R)] = \inf_{\boldsymbol{\theta}\in\Theta} \mathbb{E}^{\mathbb{Q}}[R_\theta f(R)] \,. \tag{29}$$

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

$$\mathcal{U}(f) := \sup_{\boldsymbol{\theta}\in\Theta} \mathbb{E}^i[f(R)] = \sup_{\boldsymbol{\theta}\in\Theta} \mathbb{E}^{\mathbb{Q}}[R_\theta f(R)] \,. \tag{30}$$

This concludes the proof. ■

Corollary 2 (Upper and lower bounds on expected payoffs from observed investor's holdings)**.** *In the context of Proposition 3, suppose that $\boldsymbol{\theta}_i^\star$ 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 $\boldsymbol{\theta}_i$ such that*

$$d(\boldsymbol{\theta}_i, \boldsymbol{\theta}_i^\star) \leq \delta \,, \tag{31}$$

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

$$\mathcal{L}(f) = \inf_{d(\boldsymbol{\theta},\boldsymbol{\theta}_i)\leq\delta} \mathbb{E}^{\mathbb{Q}}[R_{\theta} f(R)] \leq E^{i}[f(R)] \leq \mathcal{U}(f) = \sup_{d(\boldsymbol{\theta},\boldsymbol{\theta}_i)\leq\delta} \mathbb{E}^{\mathbb{Q}}[R_{\theta} f(R)] \,. \tag{32}$$

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

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

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

$$g_{\mathcal{L}(f)}(\lambda) = \inf_{\boldsymbol{\theta}\in\mathbb{R}^N} \left\{ \mathbb{E}^{\mathbb{Q}}[R_{\theta} f(R)] + \lambda\left(\frac{1}{2}\|\boldsymbol{\theta} - \boldsymbol{\theta}_i\|_2^2 - \delta\right)\right\} \,. \tag{33}$$

We can express the return on the investor's portfolio as $R_{\theta} = 1 + \boldsymbol{\theta}'\boldsymbol{R}^e$ where $\boldsymbol{R}^e$ is a vector of excess returns. This gives the optimality condition:

$$\mathbf{0} = \nabla g_{\mathcal{L}(f)}(\lambda) = \mathbb{E}^{\mathbb{Q}}[\boldsymbol{R}^e f(R)] + \lambda(\boldsymbol{\theta} - \boldsymbol{\theta}_i) \,, \tag{34}$$

and, whenever the constraint is binding:

$$\frac{1}{2}\|\mathbb{E}^{\mathbb{Q}}[\boldsymbol{R}^e f(R)]\|_2^2 = \frac{1}{2}\lambda^2\|\boldsymbol{\theta} - \boldsymbol{\theta}_i\|_2^2 = \lambda^2\delta \,, \tag{35}$$

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

$$\lambda^{\star} = \frac{1}{\sqrt{2\delta}}\,\|\mathbb{E}^{\mathbb{Q}}[\boldsymbol{R}^e f(R)]\|_2 \,. \tag{36}$$

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

$$\boldsymbol{\theta}^{\star} = \boldsymbol{\theta}_i - \frac{1}{\lambda^{\star}}\mathbb{E}^{\mathbb{Q}}[\boldsymbol{R}^e f(R)] = \boldsymbol{\theta}_i - \frac{\sqrt{2\delta}}{\|\mathbb{E}^{\mathbb{Q}}[\boldsymbol{R}^e f(R)]\|_2}\,\mathbb{E}^{\mathbb{Q}}[\boldsymbol{R}^e f(R)] \,. \tag{37}$$

This gives the closed-form lower bound:

$$\mathcal{L}(f) = g_{\mathcal{L}(f)}(\lambda^{\star}) = \mathbb{E}^{\mathbb{Q}}[R_{\theta^{\star}} f(R)] = \mathbb{E}^{\mathbb{Q}}[f(R)] + \mathbb{E}^{\mathbb{Q}}[\boldsymbol{\theta}_i'\boldsymbol{R}^e f(R)] - \sqrt{2\delta}\|\mathbb{E}^{\mathbb{Q}}[\boldsymbol{R}^e f(R)]\|_2 \,. \tag{38}$$

In an analogous vein, we obtain the upper bound:

$$\mathcal{U}(f) = g_{\mathcal{U}(f)}(\lambda^\star) = \mathbb{E}^{\mathbb{Q}}[R_{\theta^\star} f(R)] = \mathbb{E}^{\mathbb{Q}}[f(R)] + \mathbb{E}^{\mathbb{Q}}[\boldsymbol{\theta}_i' \boldsymbol{R}^e f(R)] + \sqrt{2\delta}\|\mathbb{E}^{\mathbb{Q}}[\boldsymbol{R}^e f(R)]\|_2 \,. \tag{39}$$

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_θ, and having belief $\mathbb{P}$. Then, the following lower bound holds:*

$$\mathcal{L} := \inf_{\boldsymbol{\theta}\in\Theta} \mathbb{E}^{\mathbb{Q}}[R_\theta \log R_\theta] \leq \mathbb{E}^i[\log R_\theta] \,. \tag{40}$$

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

$$\mathbb{E}^i[\log R_\theta] = \mathbb{E}^{\mathbb{Q}}[R_\theta \log R_\theta] \,, \tag{41}$$

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

$$\mathcal{L} := \inf_{\boldsymbol{\theta}\in\Theta} \mathbb{E}^i[\log R_\theta] = \inf_{\boldsymbol{\theta}\in\Theta} \mathbb{E}^{\mathbb{Q}}[R_\theta \log R_\theta] \,. \tag{42}$$

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

Corollary 3 (Lower bound extracted from observed investor's holdings)**.** *In the context of Proposition 4, suppose that $\boldsymbol{\theta}_i^\star$ 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 $\boldsymbol{\theta}_i$ such that*

$$d(\boldsymbol{\theta}_i, \boldsymbol{\theta}_i^\star) \leq \delta \,, \tag{43}$$

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

$$\mathcal{L} = \inf_{\boldsymbol{\theta}} \mathbb{E}^{\mathbb{Q}}[R_\theta \log R_\theta] \quad \textit{s.t.} \quad d(\boldsymbol{\theta}, \boldsymbol{\theta}_i) \leq \delta \,. \tag{44}$$

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

$$\mathcal{L} = \mathbb{E}^i[\log R_{\theta^\star}] \,.$$

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

$$\mathcal{L} = \mathbb{E}^{\mathbb{Q}}[R_\theta \log R_\theta] = \mathbb{E}^i[\log R_\theta] \,. \tag{45}$$

This concludes the proof. ■

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

$$\mathcal{L} \geq g(\lambda) := \inf_{\boldsymbol{\theta} \in \mathbb{R}^N} \left\{ \mathbb{E}^{\mathbb{Q}}[R_\theta \log R_\theta] + \lambda(d(\boldsymbol{\theta}, \boldsymbol{\theta}_i) - \delta) \right\} \,. \tag{46}$$

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* $\mathbb{E}^{\mathbb{Q}}[R_\theta \log R_\theta] < +\infty$ *for all* $\boldsymbol{\theta}$ *such that* $d(\boldsymbol{\theta}, \boldsymbol{\theta}_i) < \delta'$ *then Slater's Constraint Qualification conditions hold.*

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

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

$$g(\lambda) = \inf_{\boldsymbol{\theta} \in \mathbb{R}^N} \left\{ \mathbb{E}^{\mathbb{Q}}[R_\theta \log R_\theta] + \lambda \left(\frac{1}{2}\|\boldsymbol{\theta} - \boldsymbol{\theta}_0\|_2^2 \right) \right\} \,. \tag{47}$$

Proof of 2. We first obtain, using investor j optimality conditions for investment in the index and in the (complete) option markets:

$$\mathbb{E}^j[R] = \mathbb{E}^{\mathbb{Q}}[R^j R] = \mathbb{E}^{\mathbb{Q}}\left[R^j \sum_i w_i R^i\right] = \sum_i w_i \mathbb{E}^{\mathbb{Q}}[R^j R^i] \, .$$

The market clearing condition in the option market further implies $\sum_i w_i = 1$, which gives:

$$\mathbb{E}^j[R-1] = \sum_i w_i \mathbb{E}^{\mathbb{Q}}[R^j R^i - 1] = \sum_i w_i \mathbb{C}ov^{\mathbb{Q}}(R^j, R^i) = 1 + \mathbb{C}ov(R^j, R) \, .$$

Furthermore,

$$\bar{\mathbb{E}}[R] - 1 = \sum_j w_j \mathbb{E}^j[R-1] = \sum_j w_j \mathbb{C}ov^{\mathbb{Q}}(R^j, R) = \mathbb{V}ar^{\mathbb{Q}}\left(\sum_i w_i R^i\right) = \mathbb{V}ar^{\mathbb{Q}}(R) \, .$$

Finally, if all index weights are positive, then:

$$\bar{\mathbb{E}}[R] = \sum_i w_i \mathbb{E}^i[R] = \mathbb{E}^{\bar{\mathbb{P}}}[R] \, .$$

This concludes the proof. ■

B Additional Figures

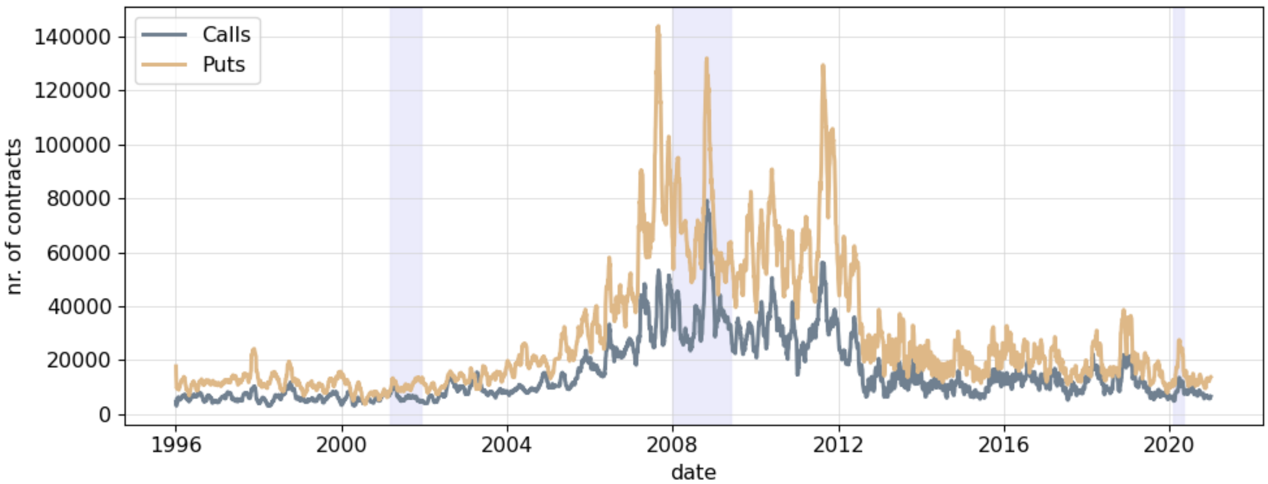

Figure B.1. Trading Volume in 30-days OTM options across time
Notes: This figure plots the time-series of the trading volume in OTM calls and puts, expiring in 30 days, from 1996 to 2020. Data are thirty-days moving average of daily observations. Gray bars indicate NBER recessions.

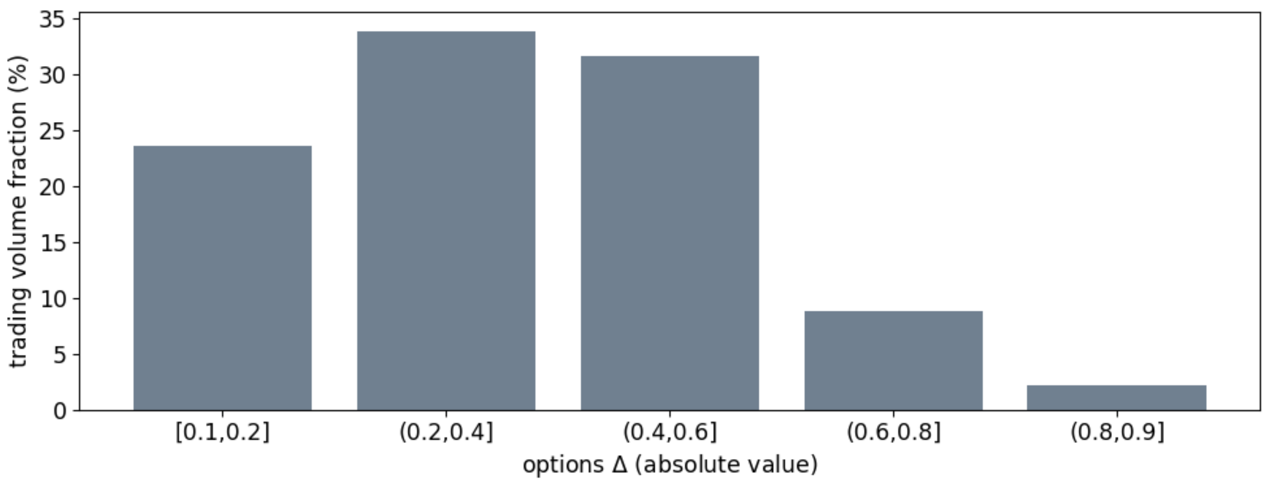

Figure B.2. Trading Volume in 30-days OTM options across moneyness
Notes: This figure plots the distribution of the trading volume in calls and puts expiring in 30 days, across different levels of moneyness identified by their $|\Delta|$. Each value is reported as fraction of the total trading volume. Data runs from 1996 to 2020.

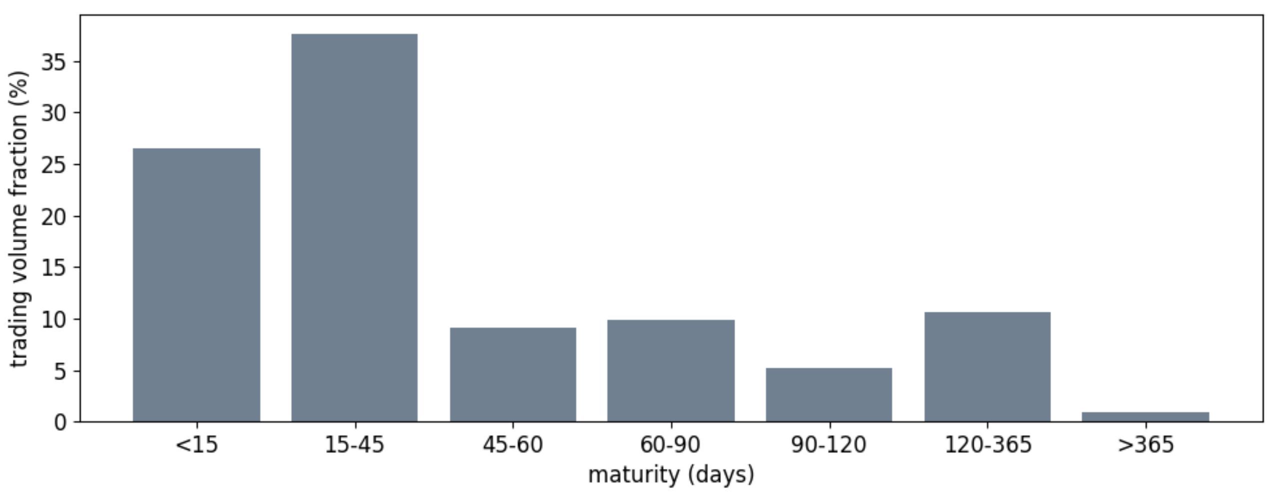

Figure B.3. Trading Volume in OTM options across maturity
Notes: This figure plots the distribution of the trading volume in OTM calls and OTM puts, across different maturities. Each value is reported as fraction of the total trading volume. Data runs from 1996 to 2020.

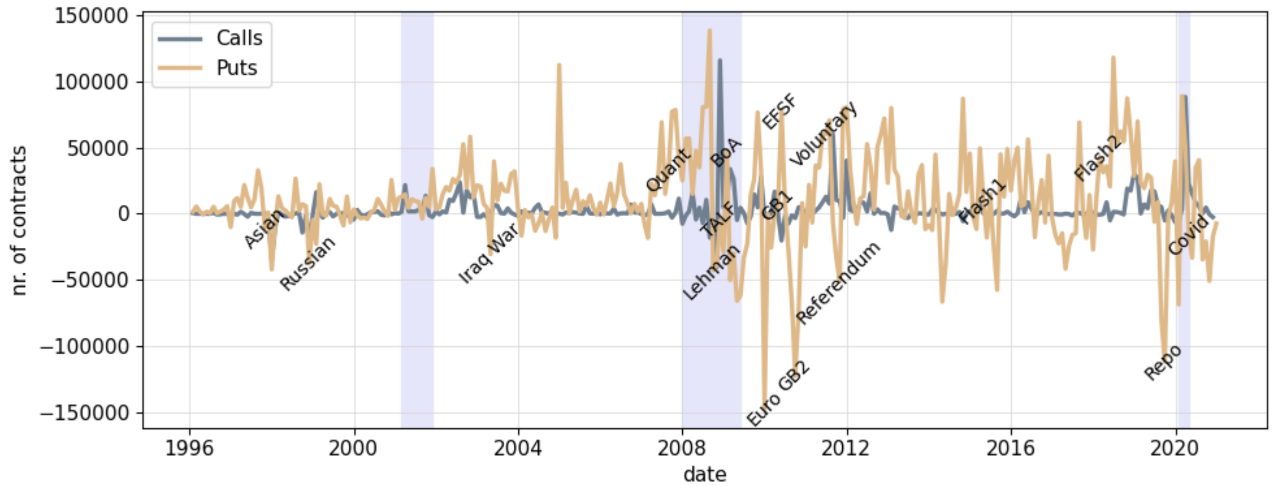

Figure B.4. Customers' demand flows for deep-OTM options at any maturity
Notes: This figure plots the time-series of our proxy for customers' demand flows on deep-OTM calls and puts (with $K/S_t \leq 0.85$ for puts and $K/S_t \geq 1.15$ for calls), at any maturity between 7 and 500 days. Option demand flows are the sum of opening positions on the same contract recorded every day. Daily data are summed over monthly basis. Gray bars indicate NBER recessions. Major financial events are highlighted in the graph.

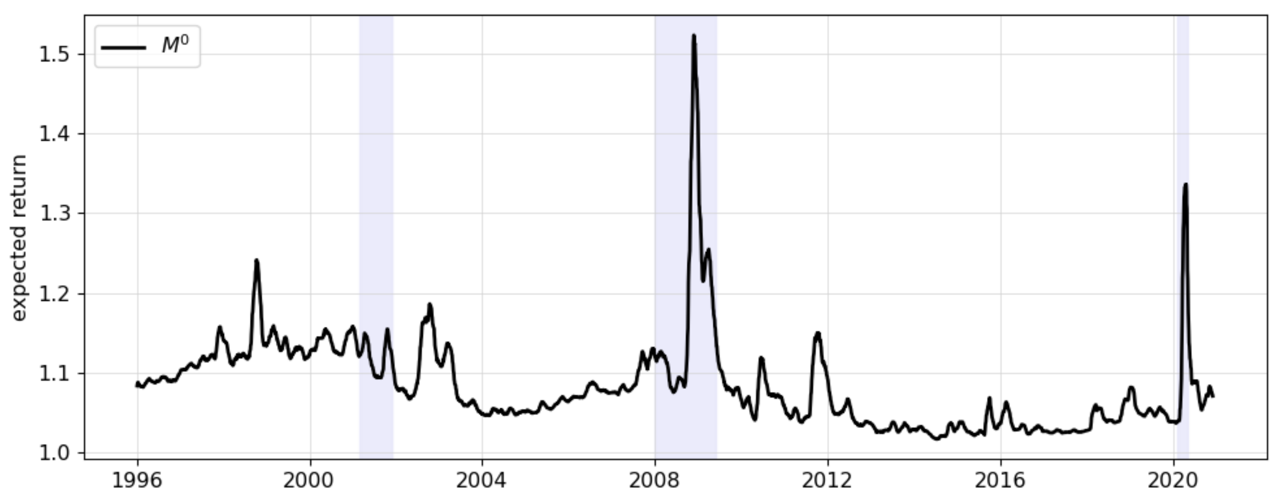

Figure B.5. $\mathrm{E}^0[R]$
Notes: This figure plots the time-series of $\mathrm{E}^0[R]$, i.e. the expected return recovered by $M^0 = 1/R$. The graph shows the 30-days moving average of daily recovered expected returns, over the horizon of one month. Gray bars indicate NBER recessions.

C Option Trading During the Financial Crisis

In the following, we study customers' holdings across different maturities during the Great Financial Crisis. We interpolate options expiring in 15, 30, 60, and 90 days, and we collect data separately for mild- and deep-OTM puts. Figure C.1 plots the corresponding holdings. We note that customers always hold mild-OTM options at any maturity, except for a small period immediately after the Lehman default when they actually sold puts and increased inventory on three-month options. Before approaching the crisis and immediately after, the series display strong co-movement, and seasonality patterns are particularly evident for long-term options. Conversely, holdings for deep-OTM puts depict a different pattern: on some dates (like immediately before the Lehman bankruptcy), customers bought short-term puts while selling long-term ones, as if anticipating the upcoming market turbulence; during the core of the crisis, they took negative positions in options at any maturity (up to the end of 2008, around April 2009) or they started buying long-term puts again. MOTM puts represent the dominant component in customers' portfolios; therefore, we expect that during the crisis the term structure of their subjective expected returns is decreasing, especially at longer horizons. Instead, the expectations at two weeks are more aligned to the monthly, having more pronounced peaks among the deep-OTM puts but being flatter among the mild-OTM.

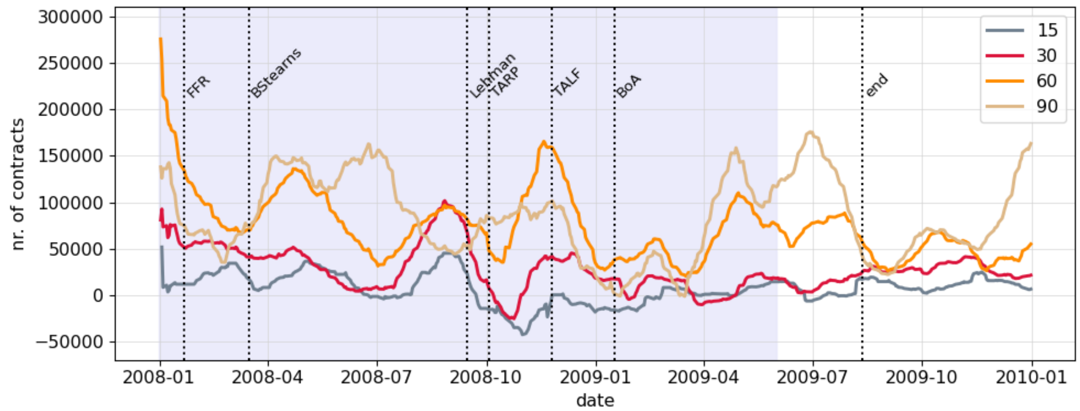

A. Mild-OTM Puts

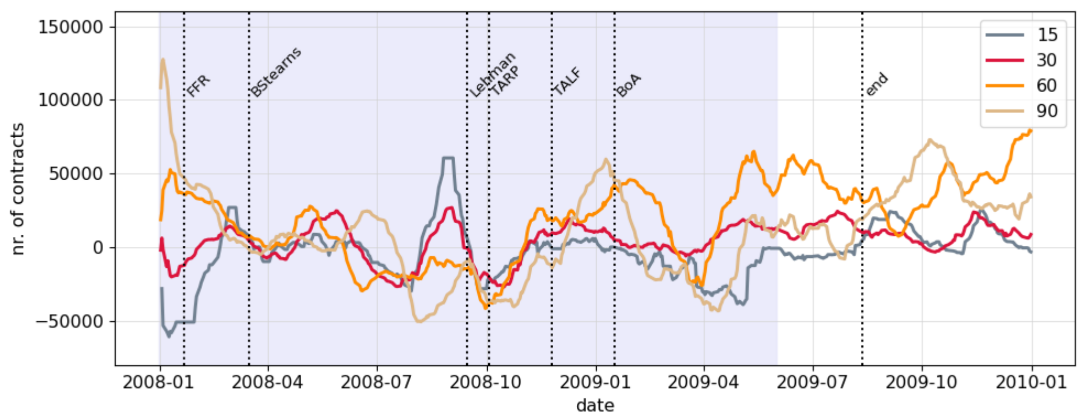

B. Deep-OTM Puts

Figure C.1. Customers' Holdings of monthly OTM puts during the Great Financial Crisis
Notes: This figure plots the time-series of our proxy for customers' portfolio holdings of OTM puts during the period around the Financial Crisis in 2008. Options expire in 15, 30, 60 or 90 days. Options holdings are the sum of opening and closing positions on the same contract that customers enter from issuance. Mild-OTM options have $|\Delta| \in (0.2, 0.4]$. Deep-OTM options have $|\Delta| \in [0.1, 0.2]$. The plots display the thirty-days moving average of the holdings. Grey areas indicate NBER recession periods. Some important events are highlighted.

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*.

Amayaa, Diego, Pedro A. Garcia-Aresb, Neil D. Pearson, and Aurelio Vasquez. 2024. “0DTE Index Options and Market Volatility: How Large is Their Impact?” *Working Paper, UIUC*.

Bakshi, Gurdip, Charles Cao, and Zhiwu Chen. 1997. “Empirical Performance of Alternative Option Pricing Models.” *Journal of Finance* 52 (5): 2003–2049.

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

Ben-David, Itzhak, John R. Graham, and Campbell R. Harvey. 2013. “Managerial Miscalibration.” *Quarterly Journal of Economics* 128: 1547–1584.

Beutel, Johannes, and Michael Weber. 2023. “Beliefs and Portfolios: Causal Evidence.” *Working Paper, Deutsche Bundesbank*.

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.

Buraschi, Andrea, Fabio Trojani, and Andrea Vedolin. 2014. “When Uncertainty Blows in the Orchard: Comovement and Equilibrium Volatility Risk Premia.” *Journal of Finance* 101 (1): 101–137.

Carr, Peter, and Dilip Madan. 2001. “Optimal Positioning in Derivative Securities.” *Quantitative Finance* 1: 19–37.

Chabi-Yo, Fousseni, and Johnathon Loudis. 2020. “The Conditional Expected Market Return.” *Journal of Financial Economics* 137 (3): 752–786.

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.

Couts, Spencer J., Andrei S. Goncalves, and Johnathan Loudis. 2023. “The Subjective Risk and Return Expectations of Institutional Investors.” *Working Paper, Ohio State University*.

Dahlquist, Magnus, and Markus Ibert. 2024. “Equity Return Expectations and Portfolios: Evidence from Large Asset Managers.” *Review of Financial Studies* 37: 1887–1929.

Egan, Mark, Alexander MacKay, and Hanbin Yang. 2024. “What Drives Variation in Investor Portfolios? Estimating the Roles of Beliefs and Risk Preferences.” *Working Paper, Harvard Business School*.

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.

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.

Jensen, Theis Ingerslev. 2024. "Subjective Risk and Return." *Working Paper, Yale University*.

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*.

Lettau, Martin, and Sydney Ludvigson. 2001. "Consumption, Aggregate Wealth, and Expected Stock Returns." *Journal of Finance* 56: 815–849.

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

Ludvigson, Sydney, and Serena Ng. 2009. "Macro Factors in Bond Risk Premia." *Review of Financial Studies* 22: 5027–5067.

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.

Welch, Ivo, and Amit Goyal. 2008. "A Comprehensive Look at The Empirical Performance of Equity Premium Prediction." *Review of Financial Studies* 21: 1455–1508.