Sentimental Recovery*

Altan Pazarbasi **Paul Schneider** **Grigory Vilkov**

This version: November 29, 2019

Abstract

We extract subjective risk-neutral and physical distributions from option quotes on S&P 500 and VIX futures according to agents' sentiment without assumptions on preferences or underlying processes. The recovered joint distributions imply marginal distributions markedly different from conventional univariate approaches. With these distributions at hand, we devise optimal Sharpe ratio trading strategies in S&P500 and VIX futures markets that are subjective to the agents, and implement them at the observed quotes. The bivariate distributions define important investment opportunities that would not be available considering the two markets separately. Dispersion of beliefs regarding both market and volatility dynamics is related to, and predicts macroeconomic indicators.

Keywords: Recovery, sentiment, market views, volatility trading, market spanning

JEL: G11, G12, G13, G17

*Altan Pazarbasi and Grigory Vilkov are affiliated with Frankfurt School of Finance & Management, `A.Pazarbasi@fs.de` and `vilkov@vilkov.net`; Paul Schneider is affiliated with USI Lugano and Swiss Finance Institute, `paul.schneider@usi.ch`. We received helpful comments and suggestions from Jens Jackwerth. All errors are ours.

Market participants and researchers routinely use observable prices to make inferences about the risk-neutral probability distribution. Before Hansen and Scheinkman (2009) and the seminal recovery theorem of Ross (2015), the formal conversion from the risk-neutral to the physical world had assumed strong conditions on the underlying market in the form of stochastic processes and preferences of participating agents. The growing field of studies on recovery have distilled a set of less strict sufficient conditions to extract physical beliefs and pricing kernels from Arrow-Debreu prices alone. In this pursuit, the majority of existing recovery approaches still assumes a representative agent in a complete Markov economy with a unique risk-neutral measure that is defined solely over the market returns as proxy for economic states. Thus, the typical set of assumptions is quite restrictive and is hard to reconcile with reality.

This paper claims three contributions: First, we significantly relax assumptions commonly made in the recovery literature. We jointly recover risk-neutral and physical probability distributions for agents with heterogeneous beliefs without specifying processes of fundamentals, not using a particular set of preferences, and generally not requiring a complete market. Second, we show quantitatively that agents' sentiment towards the economy is crucial for the recovered quantities; and, hence, standard procedures that do not explicitly model time-varying beliefs of the representative agent forego important information. Last, we explicitly augment volatility dynamics in the definition of an economic state. Since the market volatility index (VIX) is commonly dubbed the "fear index," this second dimension allows us to treat the same market states differently conditional on the expected future market dynamics. For example, a -10% market movement does not necessarily represent a highly priced state if the expected future volatility is relatively low and economic agents anticipate a rapid recovery. It adjusts the agents' recognition of "good" and "bad" states, regularizes the recovery procedure, leading to less noisy estimates and more balanced optimal investment strategies.

We acknowledge that it is not possible to recover a subjective physical distribution from asset prices without making assumptions (Borovička, Hansen, and Scheinkman, 2016; Schneider and Trojani, 2019). Rather then imposing technical assumptions on the underlying processes, we engineer our approach from *economic* assumptions in the following steps: first, we devise an algorithm to extract a risk-neutral (Q) and a corresponding subjective physical (P) joint distribution of the market index and its future volatility consistent with particular sentiment towards the future market performance and expected volatility dynamics from the quotes of traded options. This algorithm is consistent with bid-ask spreads from all observed instruments, and *recovers* these distributions under economic constraints. The title 'sentimental recovery' stems from the particular assumptions we make about the sentiment of the agents, while the procedure is otherwise model-free. We choose the bivariate index return and index expected volatility distribution, because a liquid option market for both underlying instruments exists, and because (as is shown in Chabi-Yo, Garcia, and Renault, 2008), jointly, market index and expected volatility contain economically relevant information that is not contained in the respective marginal distributions. Second, we put our algorithm to work using S&P 500 and VIX option quotes to extract subjective risk-neutral and physical measures for agents with implicit beliefs covering the whole range of realistic market sentiments from extremely bullish to extremely bearish. We suggest a very simple procedure to combine the recovered sentiment-specific distributions into an aggregate one, so that an "aggregate agent" inherits a mixture of these distributions in time-varying proportions and demonstrates sensible beliefs regarding future market and expected volatility dynamics. Adding the second (volatility) dimension to the recovered subjective distributions plays a crucial role in how agents treat "good" and "bad" states: in the presence of the second dimension, the Arrow-Debreu prices change considerably stronger across volatility states than across market index states, and overall the joint state space of market and volatility regularizes the recovery procedure. Depending on subjective beliefs, agents pursue very different optimal trading strategies, and the consideration of the two-dimensional distribution overall leads to more balanced investments. Lastly, we show empirically that dispersion of market-compatible beliefs about future market performance and its volatility are linked to the current economic regime and future investment opportunities. In particular, we highlight that the dispersion of beliefs about expected volatility, which is pivotal in forming heterogeneous pricing kernels, predicts the future state of the real economy.

Literature: Our curiosity is driven by recent advances on *recovery*, a new field in financial economics started by Ross (2015). Recovery is concerned with saying as much as possible about the conditional physical distribution, assuming as little as possible about the underlying economy, and making inference using only a cross section of asset prices, such as options. In the original approach, Ross (2015) assumes a discrete-time, discrete-state Markov economy together with a state-independent pricing kernel to achieve this task. Jensen, Lando, and Pedersen (2019) discard the first assumption, while maintaining the second. Borovička, Hansen, and Scheinkman (2016) criticize the assumption of a state-independent pricing kernel as economically and econometrically unjustified.

The recovery approach closest to ours is explored by Schneider and Trojani (2019). They acknowledge that market incompleteness allows the co-existence of many distributions compatible with observed prices, and suggest to use economically justified constraints on expected profits of tradings strategies for identification of candidate physical distributions. Ignoring any time series information, and with trading strategies in mind, they propose to use the distribution implying the lowest Hansen and Jagannathan (1991) asset pricing bound out of those. Our approach is similar to theirs in that we employ economic constraints and a good-deal bound. However, it is different with respect to many important dimensions. On top of considering a bivariate market, our approach is compatible with *all* observed options, and not just a small number of moments implied by them. Furthermore, we recover also a pricing distribution, jointly with the physical, rather than assuming it to be fixed. These choices turn out to be economically relevant.

There is also literature developing bounds on moments of individual stocks, and market indices, rather than trying to recover the full distribution. These studies comprise, for example, Martin (2017) and Chabi-Yo and Loudis (2019) for the market expected return, Martin and Wagner (forthcoming) and Kadan and Tang (forthcoming) for bounds on individual stock expected returns. Our framework can accommodate these bounds, and others, as input that goes into the determination of the recovered distributions.

We naturally extend the area of research dealing with determining expected higher moments of return and their risk premiums. For example, Martin (2017) and Britten-Jones and Neuberger (2000) for model-free variances, Bakshi, Kapadia, and Madan (2003) for skewness and higher moments, Carr and Wu (2009) for variance risk premium, Kozhan, Neuberger, and Schneider (2012) for skewness and Schneider (2015) for generalized risk premiums, Andersen, Fusari, and Todorov (2015), Bollerslev, Todorov, and Xu (2015), and Bates (2019) for tail (jump) risk premiums, Driessen, Maenhout, and Vilkov (2005, 2009), and Buss and Vilkov (2012) for risk-neutral correlations and their risk premiums. Expected realized moments and their decomposition into building components have been studied extensively, e.g., Andersen, Bollerslev, Diebold, and Labys (2003), Andersen, Bollerslev, and Diebold (2007), among others. Because we identify full bivariate distributions of the market index and its volatility, we can use them to compute any expected moment of index return and its volatility, under both physical and risk-neutral measures, and, consequently, their risk premiums. An important difference between our method and the existing ones is that our estimates are based on agents with particular sentiment towards market dynamics and, thus, allows to study the effect of differences in beliefs on the mentioned quantities.

None of the above mentioned papers considers the joint distribution of index returns and index volatility. There are only few studies that do. Song and Xiu (2016) choose a non-parametric approach using time-series information to investigate the marginal pricing kernel of the S&P 500 index, and employ a fully parametric model to investigate the joint S&P 500 and VIX kernel. Similarly, Bardgett, Gourier, and Leippold (2019) employ a parametric approach using an affine stochastic volatility model for the same purpose. Jackwerth and Vilkov (2018) use information in short-term market index and index volatility options to model a copula-based bivariate distribution, which is calibrated then using longer-term market index options. Our paper combines this literature with the one on recovery to explore the effects of agent sentiment on trading without assuming a time-series model, and to relate dispersion of (reasonable) beliefs to macroeconomic time series.

We also relate to the literature on heterogeneous agents, which are essential for explaining and matching stylized empirical facts about the economy, financial markets, and the behaviour of households. A number of papers study the implications of heterogeneous beliefs for asset allocation and asset pricing. From the point of vocabulary for defining the agents with stable (but heterogeneous beliefs), and some trading implications, the closest study to us is Martin and Papadimitriou (2018), in which the heterogeneity in beliefs leads to sentiment and makes extreme states of nature especially valuable, and also gives rise to variance fluctuations and the variance risk premium. Other studies include Scheinkman and Xiong (2003), Gallmeyer and

Hollifield (2008), Xiong and Yan (2010), Chabakauri (2013), and Bhamra and Uppal (2014), to name just a few. In contrast to these papers we do not solve for a general equilibrium generating asset pricing and asset allocation implications based on fundamentals; instead, we rather go in the opposite direction: we take observed quotes of traded assets as given and solve jointly for the agent-specific physical and pricing probability measures and associated optimal portfolios. The source of heterogeneity in our case are different sentiments towards the most probable region of the future values of market index and its volatility, and we do not restrict the other parameters of an agent, except for imposing a good-deal bound on her likelihood ratio.

The rest of the paper is organized as follows. Section 1 introduces the economic setup and discusses our identification approach, and subsequently the dimensions along which we can investigate empirically. Section 2 deals with data content, sources, processing rules, and preliminary statistics. Section 3 delivers empirical results and analyzes several features of the output, including the variety of market views of various agents, their one and two-dimensional probability distributions, likelihood measures and the resulting optimal trading, and, lastly, the link between dispersion of beliefs and economic conditions. Section 4 gives concluding remarks and directions of further research.

1 Technology

This section introduces the modeling setup and identification approach in Subsection 1.1, and then defines in Subsection 1.2 the notions of an aggregate agent and dispersion of beliefs in our model. Subsection 1.3 discusses the importance of the second (volatility) dimension both for the initial recovery problem and for the field analysis of the recovered probability distributions and likelihood ratios. Subsection 1.4 looks at the economic interpretation of the likelihood ratio of the pricing distribution with respect to the subjective one.

1.1 Setup and Identification Approach

We take the markets with observed quotes as given, and look for both physical and equivalent martingale probability measures defined over a state space $\Omega := \omega_M \times \omega_\sigma$ with $\omega_M \subset \mathbb{R}_+$ being the state space of the S&P 500 market index $F_{M,t}$, and $\omega_\sigma \subset \mathbb{R}_+$ the state space of the VIX futures contract $F_{\sigma,t}$. We work with discrete and bounded state spaces, reflecting the convention of market quotation. For ease of notation, to avoid clutter, we use continuous notation below, despite the discrete formulation. The states are defined relative to the current forward (futures) price of the assets.[1] Thus, the future state $(1, 1)$ corresponds to the state, in which $F_{M,T} = F_{M,t}$, and $F_{\sigma,T} = F_{\sigma,t}$. For the market index, it indicates a realized return over $T - t$; for the VIX, which can not be replicated without significant transaction costs, it is equivalent to the realized basis of the selected futures (which in our sample is around +10% per month). We identify all $i = 1, \dots, n$ observed VIX and S&P 500 options on a given date with their payoff function $x_i : \Omega \mapsto \mathbf{R}$ and their bid (b^i) and ask (a^i) quotes.

The martingale distribution prices the assets. In incomplete markets such as ours, it can be made unique only through selection, despite highly informative option quotes. The physical distribution originates from the sentiment of the agents, under the assumption that it can not be too far away from the martingale distribution. Otherwise, the martingale and the physical distribution taken together would imply unreasonable risk premiums and Sharpe ratios.

We assume different economic agents, each carrying an index $l = 1, \dots, L$. The agents differ through heterogeneous beliefs, and we describe the beliefs by the region $\Omega^l \subseteq \Omega$ of the underlying state space Ω, in which the agents seek to concentrate probability mass corresponding to their beliefs. If the agents were otherwise unconstrained, they would concentrate all probability in their target region Ω^l, but observing the market, this could not be reconciled with observed bid-ask spreads. Concretely, any subjective martingale measure Q^l supported on Ω must generate option prices y_i within option bid-ask spreads $[b_i, a_i]$:

$$y_i = \int_\Omega x_i(s) dQ^l(s), \;\; y_i \in [b_i, a_i], \;\; i = 1, \dots, n, \tag{1}$$

[1] For this we assume that VIX futures prices are approximately equivalent to forward prices.

and any pair P^l, Q^l must exclude unreasonable Sharpe ratios, which we specify in terms of a good-deal bound α:

$$\int_{\Omega} \left(\frac{dQ^l}{dP^l} \right)^2 dP^l \leq \alpha. \tag{2}$$

In addition, following empirical facts about the skewness risk premium, we look for distributions Q^l and P^l, such that the risk-neutral skewness of market returns is smaller than the subjective physical skewness,

$$Skew^{Q^l} \leq Skew^{P^l}, \tag{3}$$

where tradable skewness under a particular measure Z is computed following Schneider and Trojani (2018) as

$$Skew^Z = \int_{\omega_M} \left(4(R - 1 - \sqrt{R} \log R) \right) dZ(R). \tag{4}$$

We denote the set of distributions P and Q that satisfy the constraints (1) to (3) by C. With this set at hand we are able to state the definition of the P and Q distributions subjective to agent l as

$$P^l, Q^l := \arg \max_{P,Q \in C} \int_{\Omega} \mathbf{1}_{s \in \Omega^l} dP(s), \tag{5}$$

where $\mathbf{1}_A$ denotes the indicator associated with the event A. The resulting program is high-dimensional, but convex, and can be solved rapidly.

As a consequence of market incompleteness, each agent has her own pricing (martingale) distribution Q^l. This would be true even in the case where $b_i = a_i$ for $i = 1, \ldots, n$. Contrary to widely accepted practice, we do not attempt to find the risk-neutral distribution pricing all available assets with the smallest error (where the error is typically defined relative to the mid-price); instead, we identify for each agent l the risk-neutral measure Q^l, such that (i) it prices available assets within observed bid-ask spreads, (ii) the good-deal bound given by a limit on a likelihood ratio is satisfied, and (iii) the risk-neutral market return skewness is smaller than its physical counterpart. In addition to the full optimization on the joint state space, we perform the analogous exercise on the marginal state space ω_M, pricing only the options on the market index.

We define a set of agents with views ranging from extreme bearish to extreme bullish, to cover the whole range of realistic sentiments on the market.[2] For example, a bullish agent believes that the market index will go up, and, consequently, that the market volatility will be below the current level (a well-known asymmetric volatility effect, e.g., Bekaert and Wu (2000); Carr and Wu (2017)). A bearish agent, on the contrary, believes that the market index will most probably go down, and that future expected volatility will be at least as high as the current expectation. One can build agents with time-varying views, for example, linked to the current economic situation, a sentiment index, and other variables. In the empirical section we describe several agent definitions, study the decision such agents make in trading, and analyze their revealed characteristics. Note that in general, to build a joint non-parametric risk-neutral distribution between the market index and its future expected volatility, one should observe prices of derivatives written on a function of both index and its volatility, e.g., on a sum or a ratio of two variables. Such instruments are not traded. We choose a distinct path and identify such distributions by explicitly maximizing its mass in view-consistent regions, such that the resulting risk-neutral distribution produces prices within observed bid-ask spreads for both asset classes. See Figure 1 for a stylized example of relatively bullish and bearish views.

Now we quickly discuss, along which dimensions we plan to carry out the analysis.

1.2 Market-Compatible Disagreement, or, Difference in Beliefs

While stipulating particular views of agents, we do not *a-priori* know how different assumptions about beliefs will affect the recovered distributions. When observed bid-ask spreads for available options are very narrow, and the traded strikes cover a wide range of possible market index and expected volatility values, the distributions will be very similar. In the limit, when markets are complete, we arrive at the same (unique) risk-neutral distribution for all agents,

[2] One can certainly imagine agents with even more extreme and completely "crazy" sentiments, but the relative weight of such agents should be negligible for the economy.

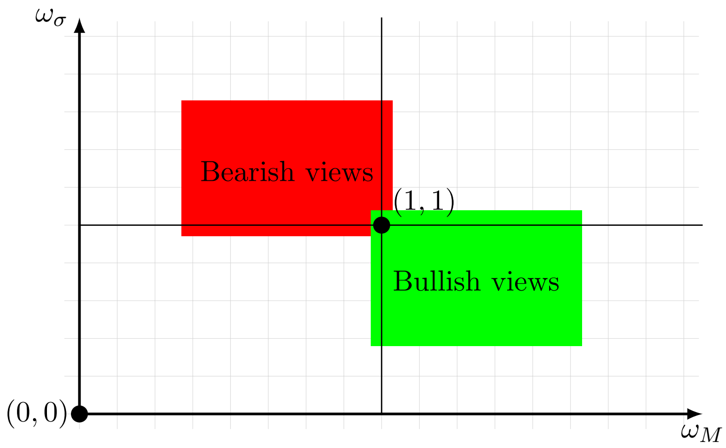

Figure 1: This figure illustrates a two-dimensional state space $\omega_{SP} \times \omega_{VIX}$ with two shaded regions, corresponding to "bullish" (in red) and "bearish" (in green) views. The center of the state space with point $(1,1)$ corresponds to the level of market index and its expected volatility equal to the S&P 500 forward and VIX futures levels.

and then, depending on the good-deal bound and specification of an agent's sentiment, we end up with a number of a agent-specific physical distributions. All three factors will affect the output: admissible risk-neutral distribution, agent's sentiment, and good-deal bound. Thus, the difference between the subjective physical distributions informs us about current deviations of market conditions from the ideal "model" world, namely, the degree of incompleteness and difference in beliefs with respect to market index and its future volatility. Analyzing these differences provides a way to analyze how these market imperfections are related to the state of economy, and even allows us to split the effect along market and volatility dimensions.

Since we have full risk-neutral and physical subjective distributions for all underlying states of nature as an output from our optimization routine, we can analyze the differences between these distributions, or "disagreement," in numerous ways, and some choices need to be made. First of all, we can compare the distributions pairwise and come up with some metric of disagreement, or we can aggregate them into a benchmark, call it "aggregate agent," and analyze then deviations from this benchmark. Second, we can look at deviations in terms of either individual state prices or some aggregate quantity, for example, expected returns. We proceed by combining agents into an aggregate one, and then use as a measure of difference in beliefs the weighted dispersion of agents' expected returns from the benchmark expected return.

To identify the aggregate agent weights we use only the pricing information (that is, risk-neutral distributions) and find weights of the agents such that the weighted average of agent-specific asset prices is as close as possible to the observed market (read, mid) prices. Using only pricing reduces the effect of the potential model misspecification. Thus, we will be looking for agent weights w that minimize the mean-squared pricing error:

$$w_t = \arg\min_{w_t} \frac{1}{N} \sum_{i=1,\ldots,N} \left(P_{i,t}^{observed} - w_t^\top P_{i,t}^{agent}\right)^2 \tag{6}$$

$$s.t. \;\; |w_t|^\top \times \mathbf{1} = w_t^\top \times \mathbf{1} = 1.0,$$

where $P_{i,t}^{observed}$ is the observed price of asset i, and $P_{i,t}^{agent}$ is the vector of agent-specific prices of asset i. The restriction on the sum of absolute weights is equivalent to a non-negativity constraint. While initially the agents are specified in a way to cover the most reasonable range of sentiments, aggregation of agents shall reveal the current prevailing (or, average) sentiment on the market. Using the weighted subjective physical P^l and risk-neutral Q^l measures, we can also analyze the prevailing physical and risk-neutral distributions, change of measure, and optimal trading strategies.

After constructing an aggregate agent, we will analyze the dispersion of beliefs defined as the weighted standard deviation of agent-specific expected returns defined for an underlying asset j at time t as

$$DOB_{j,t} = \sqrt{w_t^\top (\mu_{j,t} - w_t^\top \mu_{j,t})^2}, \tag{7}$$

where $\mu_{j,t}$ is the $L \times 1$ vector of agent-specific expected returns for asset j at a given point in time, and w_t is as before the $L \times 1$ vector of aggregate agent weights.

1.3 Importance of Volatility Dimension

There are two angles from which we can analyze volatility-related instruments. First, the choice of a source of information to extract the view-specific likelihood ratios, and, second, the risk premiums and risk-return trade-off assigned by the respective agents to the dimension in the payoff space. That is, we analyze recovery using more or less data, and then we see how it affects the perception of trading opportunities with and without volatility as instrument.

We already mention above that for forming view-specific distributions, in addition to traditionally used derivatives on the market index, we explicitly account for information from the volatility markets. An interesting question arises, whether the volatility dimension is important for information purposes, that is, if the 2-dimensional state space calibrated to both market index and its future expected volatility better identifies economic states and makes agents better off after portfolios pay off (in terms of any suitable performance metric, in our case the Sharpe ratio). Note that market index and consumption states, which matter for asset pricing, generally do not coincide; though an assumption about their approximate equivalence is often made in the recovery literature. To address this question, we first assume that only the market index states matter, calibrate one-dimensional probability measures and the corresponding likelihood ratios (1-D problem), and study the resulting pricing of market states and optimal trading strategies. We then proceed to experiments with the 2-dimensional state space (2-D problem) to see if adding the future expected volatility markets helps identifying "expensive" states more precisely.

One can always claim that more information is better; however, one should be careful in answering this question unequivocally. While theoretically even a small amount of non-redundant and precise information helps making a more qualified decision, practically (or empirically, in our case), any non-redundant information we extract will be plagued with noise, affecting the investment decisions and potentially over-weighting the positives from adding the extra dimension.

The payoff space that we use to replicate the likelihood ratios helps to determine the redundancy of volatility-based instruments. Theoretically, one can trade volatility and the yield risk premium for variance risk by trading options on the market index [Dumas (1995), Britten-Jones and Neuberger (2000), among others]; practically, however, the expected variance or volatility (for example, VIX) indices cannot be replicated due to market frictions. Trading volatility directly through futures or even volatility-of-volatility (approximately) through VIX futures options can be beneficial. We address this question by comparing the likelihood ratio swap projections using payoff space configurations with and without volatility-based products, as discussed below.

1.4 Change of Measure and its Economic Meaning

The result of solving the above identification problem is a set of risk-neutral and physical joint distributions for the market index and its future volatility, consistent with a particular agent's views, and the corresponding likelihood ratio $\mathcal{L}^l = \frac{dQ^l}{dP^l}$ for each specification of agent l.

While we cannot observe directly the positions each agent l takes in the financial markets, we know from Schneider (2015) what her trading strategy maximizing the Sharpe ratio is, using available assets—a so-called likelihood ratio swap. Such a strategy is given by the projection of the likelihood ratio $\mathcal{L}^l$ on the space of gross returns of available k trading instruments with payoffs $y_j(\omega),\ j = 1, \dots, k$ for $\omega \in \Omega$. The basic idea is that the change of measure is shifting the expected return of traded assets by the risk premium, and its projection on the payoff space yields this risk premium from a particular dimension of risk in an optimal, with respect to the Sharpe ratio, fashion. Concretely, we find an optimal composition $\theta_1^\star, \dots, \theta_k^\star$ of a portfolio of k selected assets by solving the following convex problem for payoffs $y_1, \dots, y_k$ across all states of nature $\omega \in \Omega$:

$$(\theta_1^\star, \dots, \theta_k^\star) = \arg \min_{(\theta_0, \dots, \theta_k)} \int_\Omega \left(\frac{dQ^l}{dP^l} - \theta_0 - \theta_1 y_1(\omega) - \dots - \theta_k y_k(\omega) \right)^2 d\omega \tag{8}$$

$$\text{subject to} \begin{cases} \displaystyle\int_\Omega (\theta_0 - \theta_1 y_1 - \dots - \theta_k y_k)\, dP^l = 1 \\ \displaystyle\int_\Omega (y_i(\theta_0 - \theta_1 y_1 - \dots - \theta_k y_k)\, dP^l = \int_\Omega y_i dQ^l,\ i = 1, \dots, k. \end{cases}$$

We can select any combination of assets that have defined payoffs across our state space. However, due to the fact that a projection is akin to an ordinary least squares multivariate regression (with some constraints though), we should be careful not to include assets with extremely correlated payoffs.

Thus, obtaining the view-specific risk-neutral and physical probability measures and the likelihood ratio giving the change of measure, we can directly study the consequences of the particular views for a trading strategy that is optimal with respect to the Sharpe ratio criterion. In Section 3 below, we consider different portfolios of payoffs such as S&P500 and VIX futures.

2 Data and Procedures

2.1 Data Sources and Filters

The data for the study comprise the S&P500 index, the Chicago Board Options Exchange (CBOE) volatility index VIX, futures on VIX, and options on both S&P500 index and VIX futures. The options data come from the Ivy DB OptionMetrics database for the period from March 2006 to December 2017.[3] We consider call and put options written on the S&P500 index and the VIX futures with the standard monthly maturity (expiring on the third Friday of a month for S&P500, and (usually) on the Wednesday before the third Friday for VIX). We sample these options at the monthly frequency to have approximately one month to maturity. We take all out-the-money (OTM) options (using absolute delta less or equal than 0.5), apply standard filters (remove options with negative bid-ask spreads, zero bids and zero volume), and then keep the maximum number of options compatible with no-arbitrage restrictions.[4] The forward level of S&P500 is determined at the close of each day following the CBOE VIX methodology described in CBOE (2018); and the futures level of VIX is collected as daily closing prices from CBOE Futures Exchange (CFE). As discussed in Section 1, we define the market and future volatility states relative to S&P500 forward and VIX futures levels. For computing payoffs and returns on S&P500-based forwards and options we use the closing index level on the payment date, while for VIX-based futures and options the final settlement value is a Special Opening Quotation (SOQ) of the VIX Index, computed in morning of the expiration. The SOQ levels for our sample are obtained from the CFE web-site.

[3]Note that the data for the S&P500 index and many other underlying assets are available from 1996; however, CBOE introduced VIX options only in 2006, and the first several months, until March 2006, the markets were very thin and did not provide enough data for analysis.

[4]We apply for each set of options on each day the package DIconvex developed by Karagyaur and Schneider and available at https://CRAN.R-project.org/package=DIconvex.

We use monthly data on the following economic data series are from the Federal Reserve Bank of St. Louis (FRED): Economic activity is measured by Chicago Fed National Activity Index (CFNAI) and the Leading Index (USSLIND), real output is measured by the Industrial Production Index (INDPRO), employment is measured by the logarithm of Total Nonfarm Payrolls (PAYEMS), the recession probability by the smoothed US recession probabilities (RECPROUSM156N).

We use the following financial variables: The term spread is defined as the difference between the 10-Year and one-year Treasury Constant Maturity rates, the default spread is the spread between Moody's Seasoned Baa Corporate Bond and 10-Year Treasury Constant Maturity, the dividend yield is the trailing 12-month S&P500 dividend yield from Standard&Poors and Robert Shiller's data.[5]. The risk-free rate and the market factor are from Kenneth French's website.

For comparison, we also include other sentiment measures in our analysis. Uncertainty is measured as the US Economic Policy Uncertainty Index of Baker, Bloom, and Davis (2016) from the FRED database (USEPUINDXM). Sentiment is the investor sentiment of Baker and Wurgler (2006) from Jeffrey Wurgler's website.

2.2 Data Summary Statistics

Table 1 gives an overview of the data used for calibration. Due to the filtering procedures outlined above, we select 141 dates, such that the average time to maturity is around 31 days for both market index and VIX options. The S&P500 options are vastly more liquid, and we generally have many more options on the OTM put side (105 options with up to 38% out-the money) compared to the OTM call side (37 options only up to 14% out-the-money); for VIX we have far fewer options and they are mostly traded on the OTM call side (14 calls vs. 5 puts on average) with the most remote strikes traded at +130% relative to the futures level.

Thus, from the data availability we expect to have more degrees of freedom in the volatility dimension. That is, in the market index dimension the markets seem to be more complete (informally), and agents should select more similar risk-neutral measures, while there are so few reference (pricing) points in the volatility markets that agents' risk-neutral measures in the

[5] See www.multpl.com/s-p-500-dividend-yield

	S&P 500	VIX
Days to Maturity	31.23	30.85
Number of calls	36.98	14.11
Number of puts	104.75	5.03
Moneyness, min	0.62	0.77
Moneyness, max	1.14	2.32
Futures basis, %	-0.082	10.040
Mean return, % p.a.	9.93	-61.36
Volatility, % p.a.	17.36	81.98
Skewness	-1.74	2.77

Table 1: The table reports summary statistics for the options and the underlying forward/futures contracts used in the paper. The reported numbers are the time-series means for 141 trading days from March 2006 to December 2017. *Moneyness* is the ratio of strike to underlying price. *Futures basis* is the difference between the front-month forward (S&P 500) or futures (VIX) price and the underlying price relative to the underlying price. The mean return and volatility for the underlying forwards/futures are annualized.

volatility dimension can be very different. We get back to this observation in the empirical section.

Futures price behavior is without surprises: the basis values are around zero for the S&P500, and 10% for VIX. Mean returns, volatilities, and skewness correspond to values expected for the considered period.

3 Empirical Analysis

Subsection 3.1 analyzes the output from the identification problem in terms of moments of subjective distributions; Subsection 3.2 looks in detail at the changes of measure and discussed the differences between 1-D and 2-D problem results. Subsection 3.3 links the beliefs and optimal trading strategies for agents individually, while Subsection 3.4 constructs an aggregate agent, and analyzes the dispersion of beliefs of individuals.

3.1 Market Views

For the empirical analysis we select several types of agents (from "Extreme Bear" to "Extreme Bull"), with the respective one- and two-dimensional sentiment. The one-dimensional sentiment bounds are specified regarding the future S&P 500 index return, and the two-dimensional sentiment bounds are just a combination of the S&P and VIX futures returns, that is, the changes in the underlying index or its expected volatility (VIX) relative to the respective current futures

Agent label	Market index sentiment	from	to	Future volatility sentiment	from	to
Extreme Bear	$\downdownarrows$	-15.0	0.0	$\upuparrows$	0.0	1.0
Moderate Bear	$\downarrow$	-10.0	5.0	$\uparrow$	0.0	50.0
Flat	$\rightleftarrows$	-3.0	3.0	$\rightleftarrows$	-10.0	10.0
Moderate Bull	$\uparrow$	-5.0	10.0	$\downarrow$	-15.0	0.0
Extreme Bull	$\upuparrows$	0.0	15.0	$\downdownarrows$	-25.0	-10.0

Table 2: This table provides the summary of agents types differentiated by the sentiment towards market and its future volatility. The from/to bounds indicate the ends of a range of the distribution domain, over which an agent maximizes the probability mass under her subjective measure. The numbers are given in terms of percentage deviations from the current forward (for the market index) and from the current futures (for VIX) levels, and, thus, are given in (approximately) monthly terms.

levels. For the S&P500 index it almost coincides with the monthly index return, and for the VIX it depends on the VIX evolution and initial futures basis. One-dimensional bounds always correspond to the two-dimensional ones with integrated out volatility dimension. A summary of the different ranges defining the agents is provided in Table 2. While our definition of the agents is rather *ad hoc*, it is reasonable, and covers all spectra of views on the market from extremely bearish to extremely bullish, with the middle point being the "Flat" agent seeing the market bouncing in a narrow range with relatively stable expected volatility. The joint expected dynamics is roughly compatible with the non-linear asymmetric volatility effect documented for S&P500 and VIX by Jackwerth and Vilkov (2018).

We solve the optimization problem (5) subject to the pricing (1) and good-deal (2) constraints on each date; for the good-deal bound we use $\alpha = 1.01$, which corresponds to the maximal Sharpe ratio of 0.3464 p.a.[6] In addition, we impose a constraint that the skewness premium on the S&P 500 is positive. This constraint reflects the pricing of the leverage effect, negative correlation between index returns and index volatility, which results in implied skewness being more negative than its subjective counterpart. We use for this purpose the definition of skewness from Schneider and Trojani (2018) as defined in equation (4). We solve both the one-dimensional problem considering only market states (on ω_M), and the two-dimensional problem with state space defined by both market and its future expected volatility ($\Omega = \omega_M \times \omega_\sigma$). Our algorithm yields both risk-neutral Q^l and subjective physical P^l conditional distributions for each agent. In Figure 2 we show (for one representative day) the difference in fitted option prices across

[6] The Sharpe ratio for the monthly likelihood ratio swap is computed as $\sqrt{\alpha - 1} \times \sqrt{12}$, see Schneider (2015).

agents, for both market index and VIX futures. The S&P 500 bid-ask spreads are extremely narrow, and one can barely distinguish between the option prices of different agents. The market for options on VIX futures is less liquid, with a substantially wider bid-ask spread. This allows the agents to have more different opinions about the true option prices, and the risk-neutral distributions are then "more different" in the volatility dimension. We can observe, which agent prices VIX options of different moneyness ranges closer to their bid or their ask prices. For example, *Moderate Bear* values especially high slightly OTM calls and OTM puts, while the *Flat* agent values far OTM calls highly, and, thus, will be hedging extreme volatility at prices lower than the current ask price.

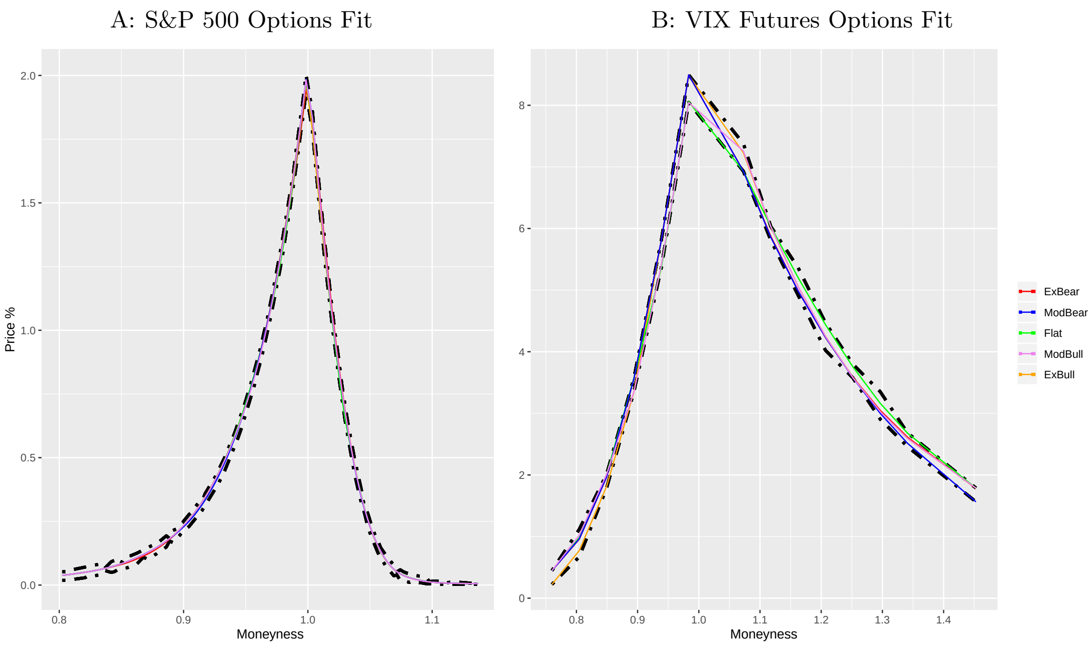

Figure 2: The figure shows an illustration of the 2-D model fit to observed quotes for out-the-money S&P500 and VIX futures options, calibrated on February 19, 2016. Moneyness and prices are defined relative to the forward/ futures level.

The unconditional moments of the resulting distributions (for the 2-D version we report moments in each dimension integrating the second one) are provided in Table 3, and the time-series of market index expected returns under subjective probability measures are provided in Figure 3.

Inspecting unconditional moments, several interesting observations emerge: First, as anticipated, expected market returns go from negative to positive and expected volatility changes go

Agent label	μ^P	σ^P	$Skew^P$	μ^Q	σ^Q	$Skew^Q$
1-D problem: S&P 500 moments						
Extreme Bear	-3.64	18.42	-1.95	0.00	18.67	-2.11
Moderate Bear	0.07	17.24	-1.45	0.00	18.66	-1.87
Flat	0.82	17.51	-1.64	0.00	18.67	-1.87
Moderate Bull	3.26	17.09	-1.44	0.00	18.64	-1.88
Extreme Bull	4.47	17.80	-1.62	0.00	18.60	-1.89
2-D problem: S&P 500 moments						
Extreme Bear	-2.61	18.36	-1.86	0.00	18.66	-1.98
Moderate Bear	-0.10	18.26	-1.80	0.00	18.68	-1.93
Flat	0.72	18.04	-1.82	0.00	18.64	-1.98
Moderate Bull	1.54	18.05	-1.77	0.00	18.68	-1.95
Extreme Bull	2.75	18.06	-1.77	0.00	18.67	-1.94
2-D problem: VIX moments						
Extreme Bear	22.14	109.56	22.66	0.00	108.84	17.81
Moderate Bear	15.99	107.91	19.25	0.00	110.05	17.68
Flat	-2.06	105.59	22.56	0.00	109.92	22.51
Moderate Bull	-10.44	104.33	17.12	0.00	109.88	18.59
Extreme Bull	-19.02	105.43	21.39	0.00	109.50	24.87

Table 3: The table reports the selected moments (mean μ, volatility σ, and skewness *Skew*) of subjective P and risk-neutral Q distributions for agents defined in Table 2, for the market index S&P 500 (for both 1-D and 2-D problems) and its expected volatility VIX (for 2-D problem only). The mean (μ) and volatility (σ) are annualized and given in percentages. *Skew* is computed in non-normalized version using formula (4) and is scaled up by 1e4 for convenience.

from positive to negative at the same time when views change from bearish to bullish. While each agent by definition maximizes probability mass in a very wide region of the subjective physical distribution (see Table 2 for the ranges in *monthly terms*), the expected returns under the subjective physical measures are relatively modest. For example, Extreme Bear believes that the market index loses up to 15% in a month, while her recovered expected return is -3.64% p.a. for 1-D and -2.61% p.a. for 2-D problems, respectively. Second, observed bid-ask quotes of the market and volatility futures options are compatible with an extremely wide range of expected returns and volatility changes. The range for market returns is more than 8% p.a. and 5% for 1-D and 2-D, respectively, and for volatility, agents' expectations differ by more than 40%. For the 1-D case there is also considerable disagreement in subjective physical higher moments (for the market only, surely), while risk-neutral moments look more alike (except for a slightly more negative risk-neutral skewness of the Extreme Bear compared to other agents). When we include the volatility dimension in estimation, the dispersion of their beliefs with respect to

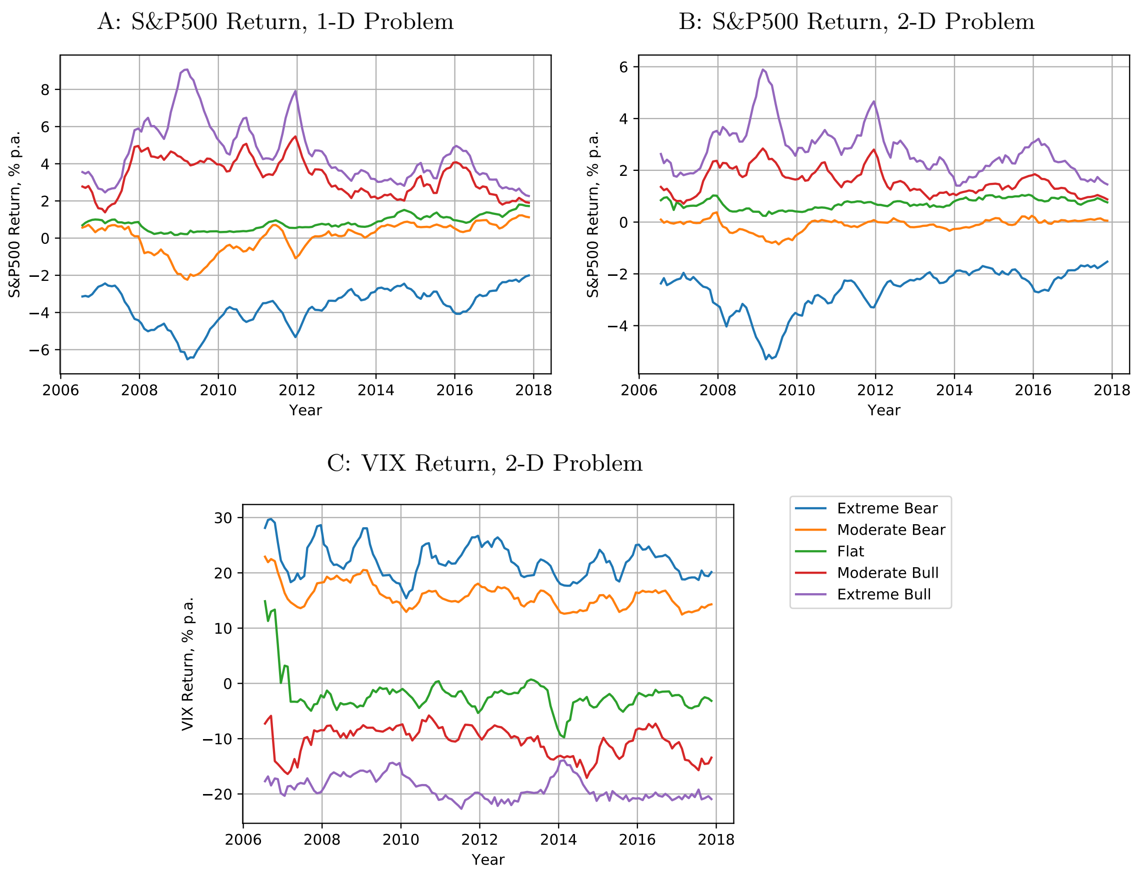

Figure 3: The figure provides the time-series of *ex ante* expected returns for the market index (S&P 500) and VIX futures, implied by the solution of 1-D and 2-D problems. The lines are smoothed using MA(5). The agents' sentiments are defined in Table 2.

both physical and risk-neutral moments of the market index goes down, so agents are more in agreement with respect to higher-order market dynamics; however, there is now a lot of dispersion of beliefs with respect to the future volatility states. Agents disagree on both risk-neutral and physical moments, with the largest effects observed in skewness ranging from 17.12 to 22.66 under P, and from 17.68 to 24.87 under individual subjective Q measures; expected volatilities (i.e., vol-of-vol) under the risk-neutral measures are quite similar across agents, though under subjective physical measures they still ranges from 104% p.a. to almost 110% p.a. Note that all the agents end up finding the distributions with the investment opportunities hitting the same good-deal bound corresponding to a Sharpe ratio of 0.35 p.a., that is, the long-term Sharpe ratio of the aggregate market index.

The series of expected returns in Figure 3 show a lot of dynamics over time coming from changing market conditions. It is the case for all agent types, though for the *Flat* agent with relatively narrow sentiment bounds, the range of expected returns is also narrower compared to the other agents. The agents with opposite sentiments end up with negatively correlated expected return predictions, and, surprisingly, the *Flat* agent is well in agreement with bullish-sentiment agents (correlations of 0.4 to 0.5); the correlations for the 2-D problem are given in Table 4. The expectations of VIX changes are mostly positively correlated amongst agents, even though the levels of these changes are very different, as can be seen from the bottom panel of Figure 3. We also observe that the expected returns extracted from S&P500 options only are typically more extreme compared to the returns from the 2-D problem, when we allow agents to use sentiment towards volatility. For the *Flat* agent, the expected returns for 1- and 2-D problems almost coincide for most of the sample period, and diverge after 2016, with 1-D problem solution being considerably more bullish.

These differences in expected returns for 1- and 2-D problems are intriguing. Theoretically, we understand that a particular change in financial wealth (say, by -5%) may correspond to very different states of the world. Solving the 1-D problem, we cannot distinguish between future investment opportunities in a given state and the 5% loss of financial wealth will always correspond to a high marginal utility state. However, by adding the volatility dimension to recovery, we now differentiate between -5% and low future volatility and -5% and high future volatility, with the former state considered better than the latter. Thus, sentimental recovery takes into account future states in a sense of the ICAPM [Merton (1973)]. Market index options on their own do not provide enough information for it.

Any of the distributions we consider is compatible with the entire option surface. For this reason it is not possible to ex-ante disqualify any of them as unreasonable. From this it follows that any recovery procedure crucially necessitates a choice of subjective beliefs, despite its theoretical motivation to say as much as possible assuming as little as possible. Adding the volatility dimension to the model regularizes the estimates of the market dynamics under both the physical and risk-neutral measures, and allows to differentiate market states by not only

	Extreme Bear	Moderate Bear	Flat	Moderate Bull	Extreme Bull
Correlations of S&P500 Expected Returns					
Extreme Bear	-	0.40	0.42	-0.70	-0.76
Moderate Bear	0.40	-	0.45	-0.06	-0.17
Flat	0.42	0.45	-	-0.08	-0.18
Moderate Bull	-0.70	-0.06	-0.08	-	0.86
Extreme Bull	-0.76	-0.17	-0.18	0.86	-
Correlations of VIX Futures Returns					
Extreme Bear	-	0.69	0.23	0.29	-0.07
Moderate Bear	0.69	-	0.24	0.20	0.01
Flat	0.23	0.24	-	0.41	-0.13
Moderate Bull	0.29	0.20	0.41	-	-0.19
Extreme Bull	-0.07	0.01	-0.13	-0.19	-

Table 4: This table provides the time-series correlations between the model-implied market index expected returns and between the returns on VIX futures for the agents with different sentiment. The expected returns are based on the 2-D problem solution.

realized wealth dynamics, but also by the future investment opportunities (expected volatility in our case).

We will now analyze in more detail the pricing of states by looking at the likelihood ratios, or, in other words, at the change of measure from subjective physical P to a risk-neutral Q.

3.2 Probability Measures and Pricing Kernels

To understand the differences in subjective distributions, we now look at the valuation of states, defined in one dimension in terms of the S&P500 return, or in two dimensions as the joint distribution of the S&P500 return and changes in future expected volatility. While the pricing distributions look quite similar (as expected from the fitted option prices on S&P and VIX futures in Figure 2), major differences can be observed from the likelihood ratios, or, in other words, from pricing kernels used by each agent to value particular states. Recall that the likelihood ratio $\frac{dQ}{dP}$ normalized by the riskfree rate represents the collection of state prices, which in equilibrium models are equal to the marginal utility of an optimizing agent. Bad states are more expensive, because consumption is valued higher when it is scarce.

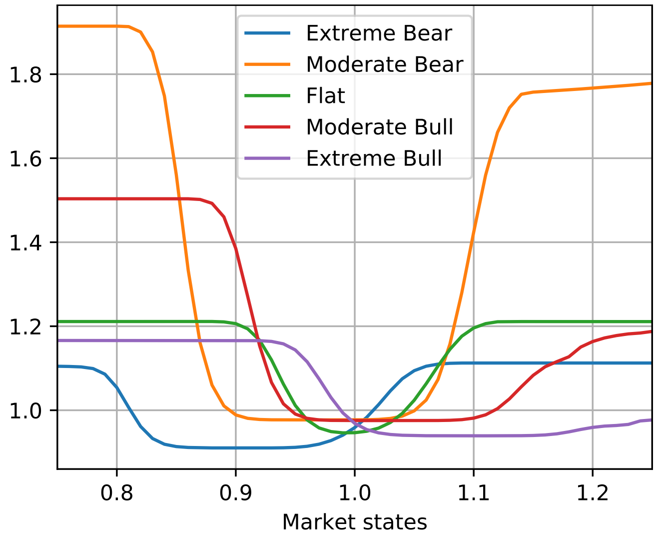

Figure 4: The figure shows the unconditional likelihood ratio $\frac{dQ}{dP}$ for the 1-D problem, evaluated at the state space of the the market index (S&P 500) for agents with different views defined in Table 2. Data preparation procedures are described in footnote 7.

Figure 4 shows the time-series average likelihood ratios for all agents for the states defined by the market monthly realizations from -20% to 20%.[7] All the pricing kernels are U-shaped, confirming existing findings. (see for example, Song and Xiu, 2016; Schneider and Trojani, 2019) Two facts are worth noting: First, the "cheapest" states for each agent l are located in the region, over which she maximizes the probability mass under the subjective physical measure P^l. Second, for the *Flat* agent both tails are equally expensive, that is, the pricing kernel is absolutely symmetric around moneyness of one. For most other agents the left tail represents far more expensive states of the world, compared to the right tail, and the left-to-right tail relative value increases with the more optimistic (bullish) sentiment. Overall, there is great heterogeneity in how states of nature, defined by the market index return, are valued, and how the likelihood ratio of particular states can differ by orders of magnitude. It stresses the importance of stating the precise assumptions behind any recovery procedures.

To investigate likelihood ratios adding the volatility dimension to the definition of a state of nature, we need to visualise three-dimensional data. In Figure 5 we plot the the likelihood ratios

[7] We first linearly interpolate the subjective physical P and risk-neutral Q probabilities (on each date and for each agent) to fill in standard moneyness grid points (from 0 to 1.5 with a 0.01 step), smooth them using Gaussian filter with $\sigma = 2$, then compute the likelihood ratio for each date as the ratio of subjective risk-neutral to physical probability for each state on the grid; we average the likelihood ratios over time for these standard moneyness points, and plot the resulting function after extra smoothing it with Gaussian filter with $\sigma = 0.5$.

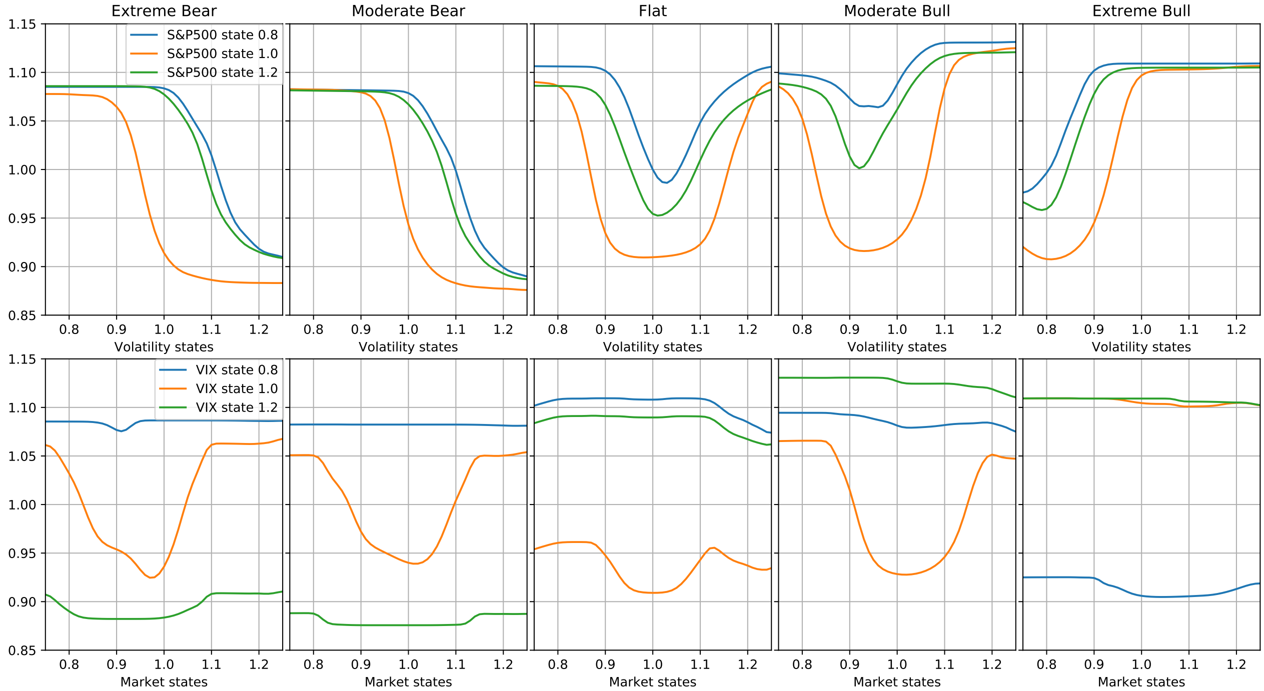

Figure 5: The figure provides the unconditional likelihood ratios $L = \frac{dQ}{dP}$ for the 2-D problem, along the market index (S&P 500) and volatility (VIX futures) states for agents with different views defined in Table 2. Data preparation procedures are described in footnote 8.

for all agents as functions of the volatility (VIX) states for three predetermined market states (S&P500 states 0.8, 1.0, and 1.2), and as a function of market states for three predetermined volatility states (VIX states 0.8, 1.0, and 1.2).[8] The results are striking for two reasons: First, we observe again that initial beliefs, or sentiments, towards the future state of the economy play a crucial role in defining the change of measure. Second, in the 2-D problem, where agents evaluate joint market and future volatility states, the market dimension is important for the state price only when the volatility state is close to the ATM (that is, to the current futures) level. Interestingly, all agents consider market states deviating from the current futures level more expensive (for a given volatility state) compared to the S&P500 state of 1.0. Changes in volatility states command much higher changes in state prices than the equivalent changes in market states. Taking into account that the expected volatility is far more agile compared to the market (with volatility of VIX futures being 3-5 times higher than the volatility of S&P500),

[8]Similar to the procedure used to generate the 1-D unconditional likelihood ratios and described in footnote 7, we perform a two-dimensional linear interpolation of 2-dimensional P and Q in market and volatility dimensions, evaluate at standard moneyness grid points, then select all points along volatility (market) states for a given level of market (volatility), smooth the resulting vectors of probabilities using Gaussian filter with $\sigma = 2$, compute the likelihood ratios, average the result over time and plot it.

such difference in state pricing is remarkable. Conditional on VIX states being -20% and 20% from the current futures level, the market states barely matter anymore. The results seem logical ex post, because the volatility dimension represents the expectation of how calm or turbulent the situation is expected to be in the next period after the realization of the market "state" and thus allows agents to "look into the future." Even a very bad market realization with low expected volatility is not that bad, because the agents anticipate a quick reversal; however, a bad market realization with high expected volatility is really bad. Thus, the expected volatility dimension provides additional information to rank the market states, especially under extreme conditions. The results stress the importance of including the second (expected volatility) dimension into the recovery procedure.

3.3 Beliefs, Trading, and Performance

Each likelihood ratio we extract from data satisfies the same good-deal bound, because the constraint is binding. Likelihood ratios across agents therefore promise the same Sharpe ratio of around 0.35 p.a.

To understand the difference in likelihood ratios in terms of information content (risk premiums and amount of risk), we can also look at the optimal trading strategies that exploit such information. Schneider (2015) shows how the subjective Sharpe ratio-maximizing strategy for an agent with particular physical and risk-neutral measures can be projected on the space of traded asset returns.

We analyze the effects of (1) the source of information for extracting the subjective physical and risk-neutral measures (that is, 1-D vs. 2-D problems), and, (2) the payoff space, which is used in the trading strategy (only S&P500, or both S&P500 and VIX futures).

Projecting a given likelihood ratio on a particular payoff space defined by payoffs to either S&P500 forwards, or both S&P500 and VIX futures, as formalized by equation (8) in Subsection 1.4, we obtain on each calibration date the optimal investments (θ_j number of notional units) in the relevant instruments. We create a portfolio such that on each rebalancing (calibration) date its initial value is set to 100, and we invest $10 \times \theta_j$ notional in each of the risky instruments j. Table 5 shows the performance of such a portfolio for all agents for the 1-D problem (with projection on the S&P500 only), and for the 2-D problem (with projection on just the S&P500, and on both the S&P500 and the VIX). Several observations emerge. First, quite

		1-D	2-D	2-D, S&P500 and VIX		
Agent label	Metric	S&P500	S&P500	S&P500	VIX	Portfolio
Extreme Bear	Return, % p.a.	-23.22	-16.56	-10.91	-7.82	-18.67
		(0.000)	(0.000)	(0.000)	(0.181)	(0.005)
	Volatility, % p.a.	15.48	11.11	6.77	15.24	19.12
	Sharpe Ratio	-1.50	-1.49	-1.61	-0.51	-0.98
	Investment, # units	-1.32	-0.95	-0.59	0.16	-
		(0.000)	(0.000)	(0.000)	(0.000)	-
Moderate Bear	Return, % p.a.	5.49	-1.36	0.59	-9.64	-8.98
		(0.000)	(0.003)	(0.110)	(0.045)	(0.053)
	Volatility, % p.a.	4.58	1.26	1.44	12.02	11.62
	Sharpe Ratio	1.20	-1.08	0.41	-0.80	-0.77
	Investment, # units	0.26	-0.03	0.08	0.15	-
		(0.006)	(0.016)	(0.000)	(0.000)	-
Flat	Return, % p.a.	10.40	6.67	6.14	0.56	6.70
		(0.000)	(0.000)	(0.000)	(0.611)	(0.000)
	Volatility, % p.a.	6.24	4.20	3.88	3.22	5.65
	Sharpe Ratio	1.67	1.59	1.58	0.18	1.19
	Investment, # units	0.52	0.35	0.33	-0.02	-
		(0.000)	(0.000)	(0.000)	(0.000)	-
Moderate Bull	Return, % p.a.	22.08	8.72	6.04	9.02	14.99
		(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
	Volatility, % p.a.	16.13	6.58	5.34	6.44	10.40
	Sharpe Ratio	1.37	1.32	1.13	1.40	1.44
	Investment, # units	1.38	0.57	0.44	-0.10	-
		(0.000	(0.000)	(0.000)	(0.000)	-
Extreme Bull	Return, % p.a.	28.00	16.64	10.25	12.53	22.69
		(0.000)	(0.000)	(0.000)	(0.001)	(0.000)
	Volatility, % p.a.	19.96	11.88	8.33	10.12	16.68
	Sharpe Ratio	1.40	1.40	1.23	1.24	1.36
	Investment, # units	1.67	1.02	0.71	-0.15	-
		(0.000)	(0.000)	(0.000)	(0.000)	-

Table 5: This table reports realized performance measures for the portfolio (and its parts) resulting from the projection of likelihood ratios of agents defined in Table 2 on the space of S&P500 futures returns (for both 1-D and 2-D problems), and on the space of S&P500 and VIX futures returns (for 2-D problem). The portfolio is rebalanced on each calibration date so that the starting value of the portfolio is always 100, and the notional of each risky portfolio part is equal to 10× Investment, # units.

naturally with the bullish market from 2006 to 2018, the agents with flat or bullish sentiments outperform the bearish ones. Thus, comparing average returns or even Sharpe ratios is not very useful. It is worth noting however that for all agent types the 1-D problem solution produces far more aggressive trading strategies compared to the ones based on the 2-D state space. The investments in the market index are larger and returns are larger in magnitude as well. Second, the 2-D problem solution produces more balanced and less risky portfolios. Comparing strategies that invest only in the market index, for example, we observe comparable Sharpe ratios, and volatilities in the 2-D case of 1.5 to 3 times lower compared to the ones in the 1-D case. For the *Moderate Bear* we also observe a "reversion" of strategy, so that she goes long the market index on average in the 1-D case, and goes short the index in the 2-D case when projected on the S&P500 space alone. The portfolio valuation over time is depicted in Figure 6. Visually, the

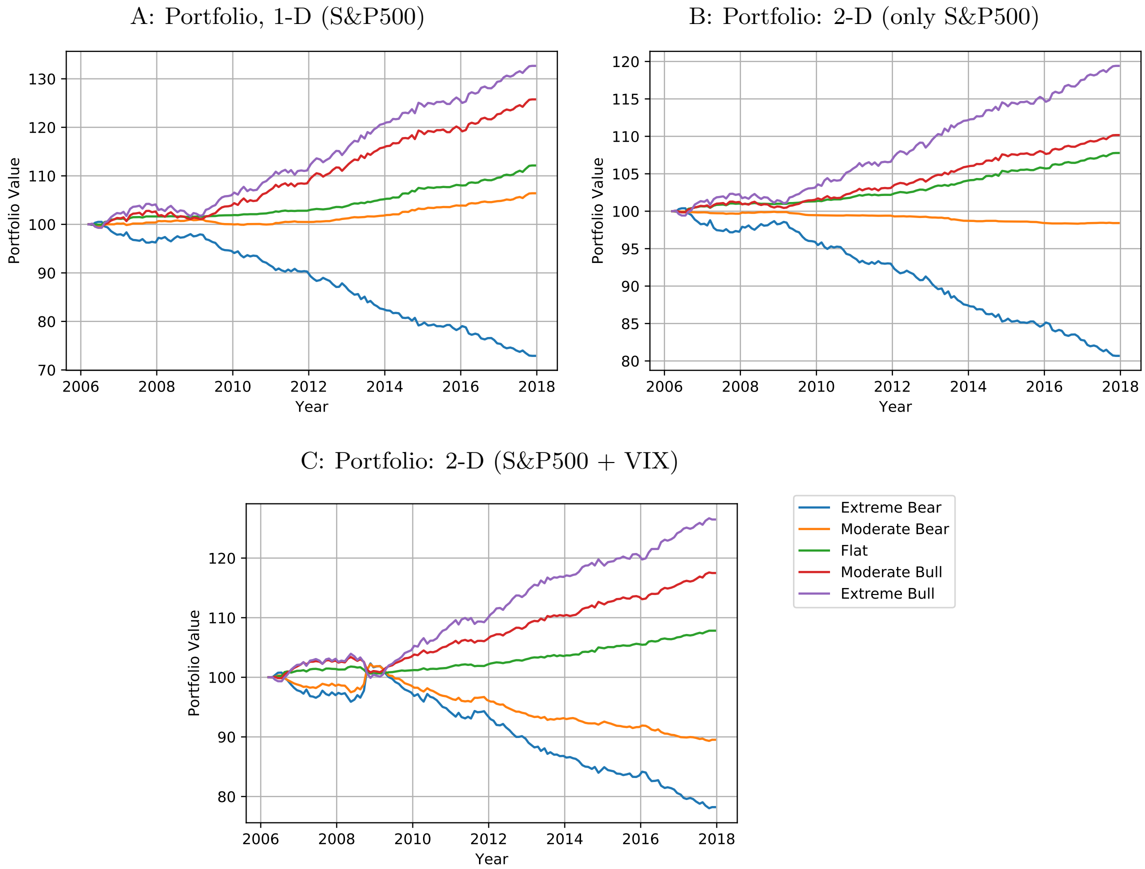

Figure 6: This figure depicts the valuation of the portfolios resulting from the projection of likelihood ratios of agents defined in Table 2 on the space of S&P500 futures returns (for both 1-D and 2-D problems) and on the space of S&P500 and VIX futures returns (for 2-D problem). The portfolio is rebalanced on each calibration date so that the starting value of the portfolio is always 100, and notional of each risky portfolio part is equal to 10× Investment, # units. The accumulated P&L is added to the portfolio value.

performance is extremely smooth for all agents, positive for agents with bullish sentiment, and negative—for bearish sentiment. During the 2007-2009 downmarket, bullish portfolios slightly underperform, while bearish portfolios recover, especially the ones invested in both S&P500 and VIX futures. The dynamics of the portfolio composition in Figure 7 reveal that in general

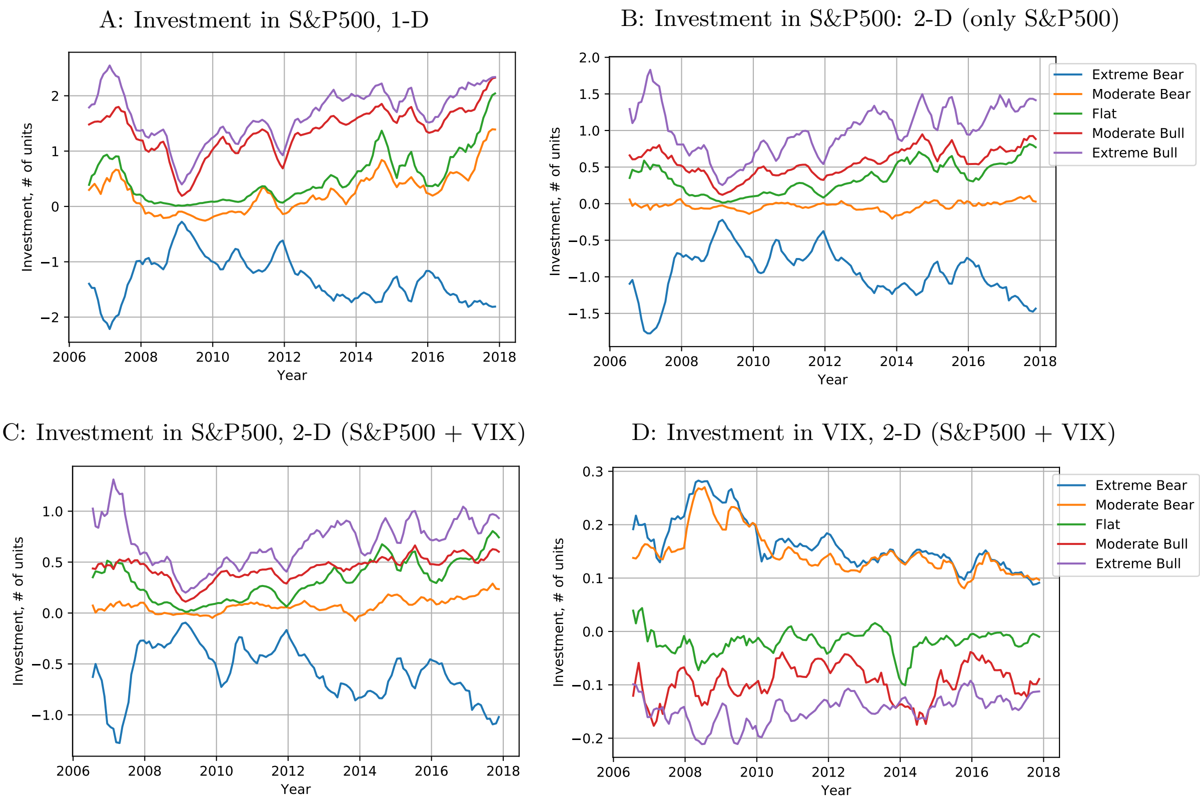

Figure 7: This figure depicts the time series of risky asset investments in the portfolio resulting from the projection of likelihood ratios of agents defined in Table 2 on the space of S&P500 futures returns (for both 1-D and 2-D problems) and on the space of S&P500 and VIX futures returns (for 2-D problem). The portfolio is rebalanced on each calibration date so that the starting value of the portfolio is always 100, and notional of each risky portfolio part is equal to 10× Investment, # units. Investments are smoothed using MA(5).

the weights of agents of the same type are correlated across different setups, and that bullish agents correctly reduce market exposure starting in 2007 with the lowest levels reached around beginning of 2009, after which the market investment again increases. Very intriguingly, the comparison of market index weights and model-implied expected returns in Figure 3 shows: in 2009 when the market investments for bullish agents are at the lows, the expected returns for *Extreme Bull* are the highest, and for *Moderate Bull* close to the highest levels. *Extreme Bear* starts covering market shorts at the same time while her expectation of market return turns more negative. Effectively, agents balance risk and return, and while the expected return has a first-order effect on the portfolio weights, increasing the amount of market risk captured by the projection of likelihood ratio leads to reduced exposure to the S&P500 index.

By executing trading on both S&P500 and VIX futures markets, the optimal strategy does not improve, and in some cases suffers from slightly higher volatility. Investments are typically positive for bearish, and negative for bullish agents. Especially long VIX futures positions of bearish agents in 2009 result in gains during the market crash and the subsequent high volatility regime.

Thus, for proper recovery it is important not only to specify sentiment of an agent correctly, but also to use a well-specified state space, augmenting the expected volatility dimension in addition to traditionally used market index states. Information from volatility markets leads to less noisy and more balanced parameters of subjective physical and risk-neutral distributions, likelihood ratios, and the resulting optimal strategy. Subsequent execution of the optimal strategy can be carried out using only market instruments. It is not our goal to produce a functional and profitable trading strategy, and we do not allow for state-dependent sentiment. However, a logical continuation of our analysis in the direction of trading would be to mix our agents in state-dependent proportions depending, for example, on any existing sentiment measure, and analyze the resulting trading strategy. Below we discuss one of the ways to aggregate agents into a aggregate agent.

3.4 Aggregation of Agents and Dispersion of Beliefs

As we admit freely, we define agents to cover a relatively wide range of sentiments towards market conditions. Each agent's expectations and the resulting behavior are not that interesting per se; in relative terms, however, the degree of heterogeneity of beliefs, and the market-implied average sentiment can provide us with additional insights.

First, we split all market participants into groups according to their sentiment, or, in other words, we find the composition of an aggregate agent in terms of five agent types with sentiments specified in Table 2. Second, we analyze the dynamics of weights, of average beliefs, and of the dispersion of beliefs defined as the weighted standard deviation of expected returns as seen by our agents on each calibration date.

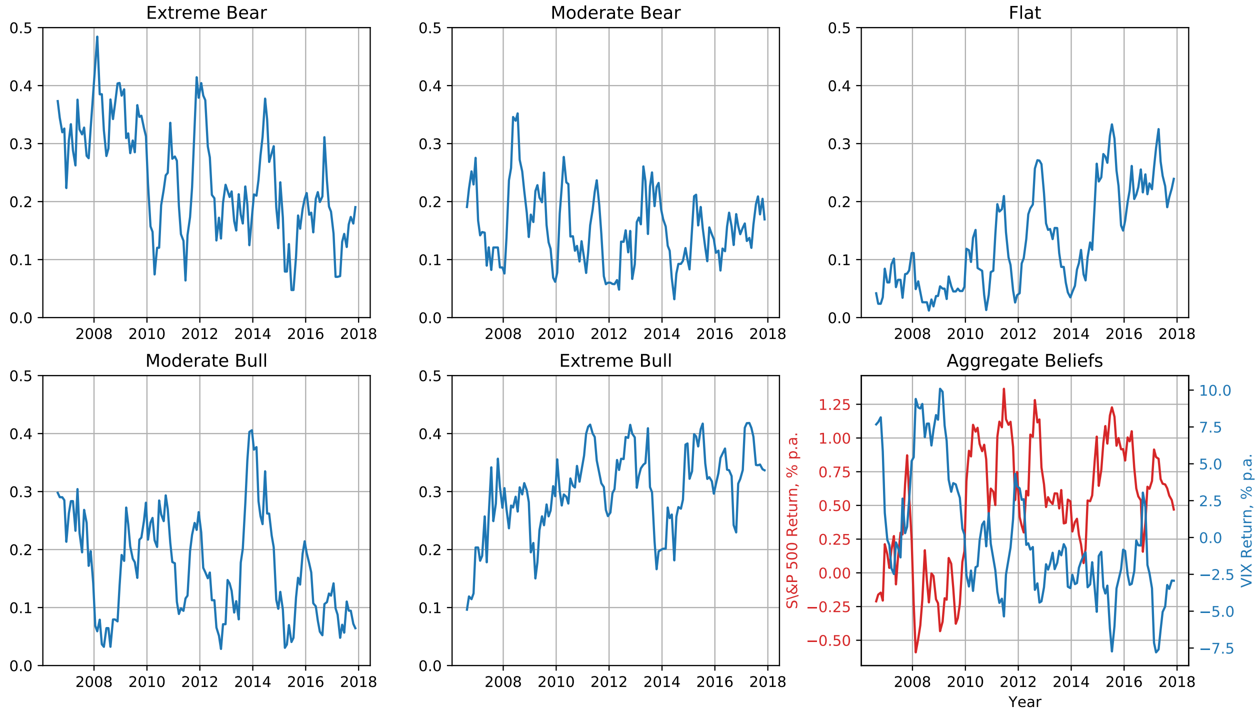

Figure 8: This figure depicts the time series of aggregate agent weights defined through equation (9). The last panel shows the time-series of aggregate agent beliefs with respect to S&P500 and VIX future returns, computed on each date as the weighted expected return for a given asset, using aggregate agent weights and agent-specific expected returns for each asset. All series are smoothed using MA(5).

To construct an aggregate agent from our agents we follow the procedure outlined in Section 1.2 and solve problem in equation (6) to find the weights w_t of the agents on each day t.[9] Because agents by construction agree on the current prices of S&P500 and VIX futures, we fit only the agent-weighted prices of observed options on S&P500 and VIX futures to the true prices, assumed to be equal to the mid-point quotes observed on a given day. To avoid the effect of small prices "unimportance," instead of using prices, we minimize the mean squared pricing error in terms of relative implied volatility deviations:

$$w = \arg\min_{w} \frac{1}{N} \sum_{i=1\dots N} \left(\frac{IV_{i,SP}^{mid}}{w^\top IV_{i,SP}} - 1 \right)^2 + \frac{1}{M} \sum_{i=1\dots M} \left(\frac{IV_{i,VIX}^{mid}}{w^\top IV_{i,VIX}} - 1 \right)^2, \tag{9}$$

$$s.t. \quad |w|^\top \times \mathbf{1} = w^\top \times \mathbf{1} = 1.0,$$

[9]We omit the time subscript for brevity in the following.

where $IV_{i,j}^{mid}$ is the implied volatility of an option i on asset j computed from its mid-price, $IV_{i,j}$ is the vector of implied volatilities computed from agent-specific option prices, and N and M are the number of available options on a particular date for each underlying.

The resulting aggregate agent weights are depicted in Figure 8, where the last panel also shows the aggregate expected return on each asset computed from the aggregate agent weights and agent-specific expected returns. Over the last 12 years, the weights of agents with bearish views decline, moderate bullish sentiment also plays less of a role, while agents with flat and extreme bullish sentiments have considerably gained in importance. The last panel of the figure shows the dynamics of aggregate agent expected returns, with an increase on average for the S&P500 and a decrease for VIX futures. Interestingly, during the downmarket period 2008-2009 the expected market return has decreased rapidly at the beginning of 2009 and stayed at these low levels until the end of 2010. At the same time the volatility index VIX was at the highest expected return levels. Aggregate expectations of market and its expected volatility are on average negatively correlated, consistent with the asymmetric volatility effect.

Aggregate levels of expected returns appear plausible, and, hence, serve as additional support for our idea to construct a trading strategy using aggregated beliefs. We form an aggregate agent by weighting respective probability measures of individual agents by the weights derived in (9):

$$dP^{agg} = \sum_L w^l \times dP^l, \quad dQ^{agg} = \sum_L w^l \times dQ^l, \tag{10}$$

and then get her optimal portfolio weights by projecting the resulting likelihood ratio $\frac{dQ^{agg}}{dP^{agg}}$ onto the selected payoff space as we did in the previous section for individual agents. The resulting investments and the portfolio performance are given in Figure 9. Visually, the performance in all cases matches the performance of the "most lucky" individual agent (see Figure 6), and the aggregate agent definitely behaves very intelligently in terms of portfolio composition. In the 1-D case the trading is far more aggressive than in both 2-D projections, and in the 2-D case with two instruments she correctly goes long volatility in the turbulent market periods. Being mostly long in VIX futures helps aggregate agent during 2008-2010 period, and slightly lowers absolute returns in the later bullish markets. The Sharpe ratios are around 1.50 for the 1-D case and 1.15 for both 2-D cases.

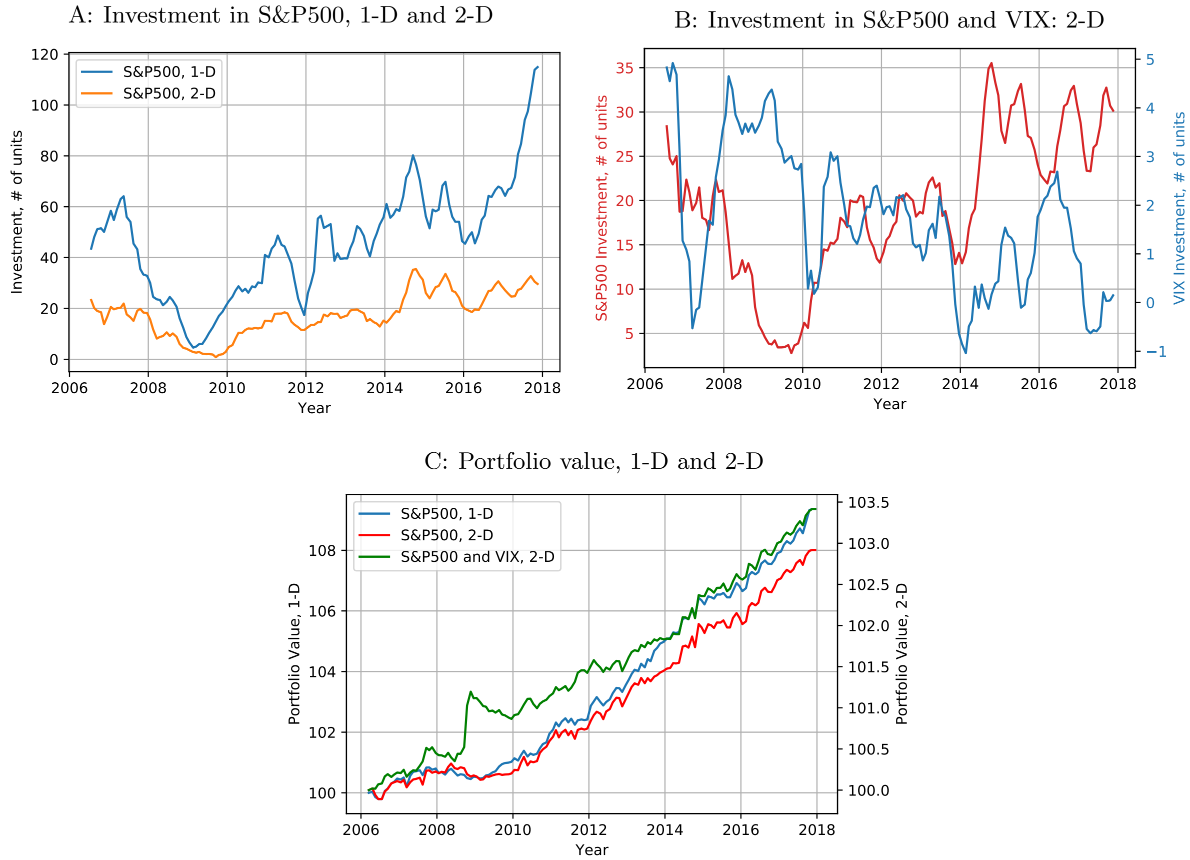

Figure 9: This figure depicts the time series of risky asset investments and the portfolios performance resulting from the projection of the aggregate agent's likelihood ratio on the space of S&P500 futures returns (for both 1-D and 2-D problems) and on the space of S&P500 and VIX futures returns (for 2-D problem). The portfolio is rebalanced on each calibration date so that the starting value of the portfolio is always 100, and notional of each risky portfolio part is equal to 10× Investment, # units. Investments are smoothed using MA(5).

Now, we are interested in the how agents agree or disagree on their subjective expectations, and we compute a dispersion of beliefs measure (DOB), applying equation (7) to both S&P500 and VIX futures.

The dynamics of both measures are provided in Figure 10, and the contemporaneous correlation with a number of financial, macroeconomic, and sentiment-type variables is provided in Table 6. Intuitively, dispersion of beliefs reflects the degree of disagreement about the future economic state, which can come along due to differential information, uncertainty, or sentiment. As we discover already above by analyzing 2-D likelihood ratios, expected volatility changes are

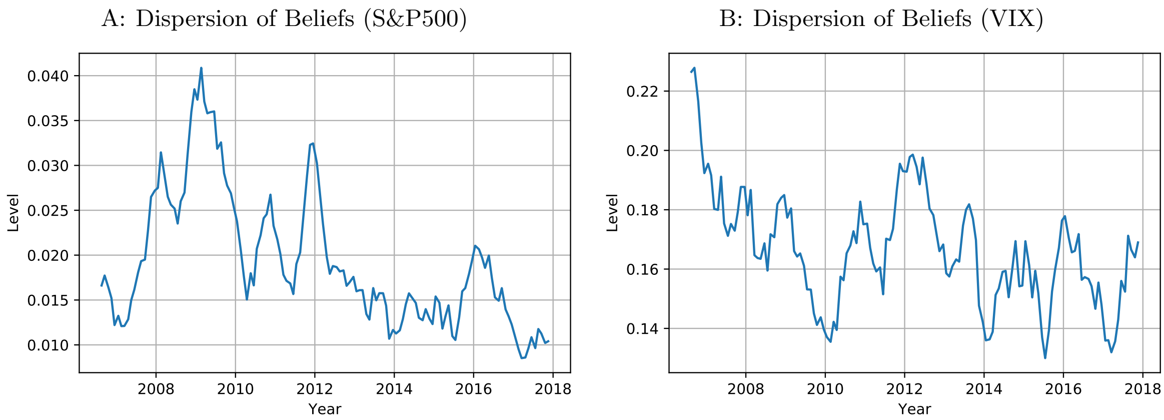

Figure 10: This figure depicts the time series of dispersion of expectations of S&P500 and VIX futures returns, defined at each point in time as the aggregate agent-weighted standard deviation of expected returns from their weighted mean. Both series are smoothed using MA(5).

priced much stronger than market index states. Thus, potentially, we expect to have stronger implications for differences of beliefs regarding VIX for the future economic state compared to the difference of beliefs about the market index. First, both DOB's are positively correlated at

	Dispersion of Beliefs (S&P500)	Dispersion of Beliefs (VIX)
Financial Variables		
Riskfree Rate	-0.01	0.35
Yield, 10 years	0.12	0.21
Term Spread	0.19	-0.34
Default Spread	0.63	-0.01
Market Factor	-0.31	-0.17
Dividend Yield	0.56	-0.05
VIX	0.75	0.12
Macroeconomic Variables		
Industrial Production Index	-0.50	0.06
Economic Activity	-0.61	-0.09
Nonfarm Payrolls	-0.42	-0.10
Leading Indicator	-0.62	-0.08
Recession Probability (smoothed)	0.54	0.04
Sentiment Variables		
Sentiment	-0.22	0.19
Economic Policy Uncertainty	0.36	-0.02
Dispersion of Beliefs (S&P500P)	1.00	0.45
Dispersion of Beliefs (VIX)	0.45	1.00

Table 6: This table shows the contemporaneous correlations between Dispersion of Beliefs measures defined in equation (7), for S&P500 and VIX futures markets, and a number of indicators and economic variables.

a level of 0.45, even though only the DOB(S&P500) is positively correlated to Economic Policy

Uncertainty. Both DOB's are correlated with Baker and Wurgler (2006) Sentiment, though the correlation is negative for DOB(S&P500) and flips sign for DOB(VIX). The Dispersion of Beliefs regarding the market index clearly goes up in bad economic conditions and market downturns: it has very negative correlations with most macroeconomic variables and with the market factor, and a positive one with the recession probability and default spread. Thus, increasing DOB(S&P500) is clearly considered bad news. Dispersion of Beliefs regarding VIX has a different connotation: it has very low correlations with macro variables. Aside from a negative correlation with inflation, it has a negative correlation with the term spread, but positive with both short- and long-term interest rates. A positive correlation with interest rate levels indicates that high DOB(VIX) reduces demand for precautionary savings, thus expecting better economic conditions in the future.

To understand if the difference of beliefs in S&P500 and VIX markets anticipates changes in economic conditions, we also perform a number of predictive regressions as follows:

$$Var_{t+m} = DOB_{j,t} + Var_t + Controls_t, \tag{11}$$

where Var_{t+m} and Var_t are the future and current values of a variable of interest (we exclude the current value of the variable for cumulative market returns), $DOB_{j,t}$ is the current value of dispersion of beliefs about either the market index or its future volatility, and $Controls_t$ are the controls including VIX, the current level of the riskfree rate, the last month realization of market factor, default and term spreads, and the Sentiment indicator by Baker and Wurgler (2006).[10] The results are provided for horizons from one to 24 months for a a number of variables (cumulative market factor over the horizon, future dividend yield, future industrial production, and economic activity) only in terms of significance of coefficients for DOB(S&P500) and DOB(VIX) in Figure 11. We observe that DOB(S&P500) works at intermediate horizon, while DOB(VIX) effect is rather long-run. A high level of dispersion of beliefs regarding market negatively predicts market returns at around annual horizon, and positively (due to lower market valuation) – dividend yield. DOB(VIX) has no short-term predictability for the market, but

[10] We do not include both DOB variables into one regression due to potential multicollinearity issues; moreover, we exclude VIX as control for regression with DOB(S&P500) due to their high correlation (0.75).

Figure 11: This figure depicts the t-statistic for one of of the Dispersion of Beliefs from regressing future value of variable of interest (or cumulative value over time for market factor) on DOB(S&P500), DOB(VIX), lagged variable, and controls, as specified in equation (11). Standard errors are corrected using Newey and West (1987) with number of lags equal to horizon (months) -1.

predicts it positively for the long horizons (18 months and beyond). Market DOB negatively predicts industrial production (mostly beyond one year horizon) and economic activity (six to 17 months) with borderline significance. DOB(VIX), positively predicts industrial production for longer horizons and has mostly insignificant (though positive) relation to economic activity. Thus, even though contemporaneously both DOB variables are positively correlated, they posses quite distinct relation to the future economy: DOB(S&P500) has medium-term negative predictability, while DOB(VIX) shows longer-term positive predictability.

4 Conclusion

In this paper, we recover subjective physical and risk-neutral distributions using only observed option quotes on the S&P 500 and VIX futures. We make assumptions about agent's sentiment towards market and its volatility, and a good-deal bound. Our assumptions concern the location in the state space where the agents concentrate probability mass. None of our assumptions could ex-ante be verified to be unreasonable. In particular they yield distributions that price all options into their respective bid-ask spreads.

The recovered subjective distributions show rich heterogeneity, stressing the importance of sentiments towards the economy in the recovery procedures. While existing literature concentrates on using the market index as the only variable describing the consumption state, we provide evidence that expected index volatility is an important information source that is not spanned by the market index, and should not be omitted. It results in lower dispersion of beliefs about the future market growth, and in a less extreme optimal trading strategies.

Both, the physical, as well as the risk-neutral distributions are subjective to the agents. Their ratio is a change of measure and informative about risk premiums the agents attach to the states. Analyzing agent-specific changes of measure for the case of using only market index options for recovery (1-D problem), we document very disperse (depending on initial sentiments), but always U-shaped pricing kernels in the market index dimension. In the 2-D case (using both market index and volatility) we observe almost flat pricing kernels in the market index dimension for volatility states different from the current futures level; in general most heterogeneity in pricing of economic states coming from future volatility dimension. It stresses the importance of expected volatility as a priced state variable in the ICAPM (Merton, 1973) sense.

The dispersion of beliefs about future market returns and expected volatility between the agents is related to the current state of the economy. It is also able to predict future macroeconomic variables. High dispersion of beliefs regarding the market indicates bad news contemporaneously, and also predicts worsening economic conditions in the medium term. High dispersion of beliefs regarding volatility brings more positive news, coincides with lower interest rates and predicts market and industrial production growth in the long run. Thus, our study emphasizes a role of option-implied information in macro-finance research and underscores the potential for incorporating volatility sentiment to enhance asset pricing models/theories.

References

Andersen, T., N. Fusari, and V. Todorov, 2015, "The risk premia embedded in index options," *Journal of Financial Economics*, 117(3), 558–584.

Andersen, T. G., T. Bollerslev, and F. X. Diebold, 2007, "Roughing It Up: Including Jump Components in the Measurement, Modeling, and Forecasting of Return Volatility," *The Review of Economics and Statistics*, 89(4), 701–720.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and P. Labys, 2003, "Modeling and Forecasting Realized Volatility," *Econometrica*, 71(2), pp. 579–625.

Baker, M., and J. Wurgler, 2006, "Investor Sentiment and the Cross-Section of Stock Returns," *Journal of Finance*, 61, 1645–1680.

Baker, S. R., N. Bloom, and S. J. Davis, 2016, "Measuring Economic Policy Uncertainty," *The Quarterly Journal of Economics*, 131(4), 1593–1636.

Bakshi, G. S., N. Kapadia, and D. B. Madan, 2003, "Stock Return Characteristics, Skew Laws, and the Differential Pricing of Individual Equity Options," *The Review of Financial Studies*, 16(1), 101–143.

Bardgett, C., E. Gourier, and M. Leippold, 2019, "Inferring volatility dynamics and risk premia from the S&P 500 and VIX markets," *Journal of Financial Economics*, 131, 593 – 618.

Bates, D. S., 2019, "How Crashes Develop: Intradaily Volatility and Crash Evolution," *The Journal of Finance*, 74, 193–238.

Bekaert, G., and G. Wu, 2000, "Asymmetric Volatility and Risk in Equity Markets," *Review of Financial Studies*, 13(1), 1–42.

Bhamra, H. S., and R. Uppal, 2014, "Asset Prices with Heterogeneity in Preferences and Beliefs," *Review of Financial Studies*, 27(2), 519–580.

Bollerslev, T., V. Todorov, and L. Xu, 2015, "Tail risk premia and return predictability," *Journal of Financial Economics*, 118(1), 113–134.

Borovička, J., L. P. Hansen, and J. A. Scheinkman, 2016, "Misspecified Recovery," *Journal of Finance*, pp. 2493–2544.

Britten-Jones, M., and A. Neuberger, 2000, "Option Prices, Implied Price Processes, and Stochastic Volatility," *Journal of Finance*, 55, 839–866.

Buss, A., and G. Vilkov, 2012, "Measuring Equity Risk with Option-implied Correlations," *Review of Financial Studies*, 25(10), 3113–3140.

Carr, P., and L. Wu, 2009, "Variance Risk Premiums," *Review of Financial Studies*, 22(3), 1311–1341.

——— , 2017, "Leverage Effect, Volatility Feedback, and Self-Exciting Market Disruptions," *Journal of Financial and Quantitative Analysis*, 52, 2119–2156.

CBOE , 2018, "Cboe Volatility Index," White Paper.

Chabakauri, G., 2013, "Dynamic Equilibrium with Two Stocks, Heterogeneous Investors, and Portfolio Constraints," *Review of Financial Studies*, 26(1), 3104–3141.

Chabi-Yo, F., R. Garcia, and E. Renault, 2008, "State Dependence Can Explain the Risk Aversion Puzzle," *Review of Financial Studies*, 21, 973–1011.

Chabi-Yo, F., and J. Loudis, 2019, "The Conditional Expected Market Return," *Journal of Financial Economics*.

Driessen, J., P. Maenhout, and G. Vilkov, 2005, "Option-Implied Correlations and the Price of Correlation Risk," Working paper, INSEAD.

Driessen, J., P. J. Maenhout, and G. Vilkov, 2009, "The Price of Correlation Risk: Evidence from Equity Options," *Journal of Finance*, 64(3), 1377–1406.

Dumas, B., 1995, "The Meaning of the Implicit Volatility Function in Case of Stochastic Volatility," *HEC Working Paper*.

Gallmeyer, M., and B. Hollifield, 2008, "An Examination of Heterogeneous Beliefs with a Short-Sale Constraint," *Review of Finance*, 12(2), 323–364.

Hansen, L. P., and R. Jagannathan, 1991, "Implications of Security Market Data for Models of Dynamic Economies," *Journal of Political Economy*, 99(2), 225–262.

Hansen, L. P., and J. A. Scheinkman, 2009, "Long-Term Risk: An Operator Approach," *Econometrica*, 77, 177–234.

Jackwerth, J., and G. Vilkov, 2018, "Asymmetric Volatility Risk: Evidence from Option Markets," *Review of Finance*, 23(4), 777–799.

Jensen, C. S., D. Lando, and L. H. Pedersen, 2019, "Generalized recovery," *Journal of Financial Economics*, 133, 154 – 174.

Kadan, O., and X. Tang, forthcoming, "A Bound on Expected Stock Returns," *The Review of Financial Studies*, hhz075.

Kozhan, R., A. Neuberger, and P. Schneider, 2012, "Understanding the Skew in Index Option Prices," *Review of Financial Studies*, 26(9), 2174–2203.

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

Martin, I., and D. Papadimitriou, 2018, "Sentiment and speculation in a market with heterogeneous beliefs," Working Paper LSE.

Martin, I., and C. Wagner, forthcoming, "What is the Expected Return on a Stock?," *Journal of Finance*.

Merton, R. C., 1973, "An Intertemporal Capital Asset Pricing Model," *Econometrica*, 41, 867–888.

Newey, W. K., and K. D. West, 1987, "A Simple, Positive-semidefinite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," *Econometrica*, 55(3), 703–708.

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

Scheinkman, J. A., and W. Xiong, 2003, "Overconfidence and Speculative Bubbles," *Journal of Political Economy*, 111, 1183–1219.

Schneider, P., 2015, "Generalized risk premia," *Journal of Financial Economics*, 116(3), 487 – 504.

Schneider, P., and F. Trojani, 2018, "Divergence and the Price of Uncertainty," *Journal of Financial Econometrics*.

——— , 2019, "(Almost) Model-Free Recovery," *Journal of Finance*, 74, 323–370.

Song, Z., and D. Xiu, 2016, "A Tale of Two Option Markets: State-Price Densities and Volatility Risk," *Journal of Econometrics*, 190(1), 176–196.

Xiong, W., and H. Yan, 2010, "Heterogeneous Expectations and Bond Markets," *Review of Financial Studies*, 23(4), 1433–1466.