UNIL | Université de Lausanne
Département de finance

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

FINANCE RESEARCH SEMINAR SUPPORTED BY UNIGESTION

"Sentiment, Risk Aversion, and Time Preference"

Prof. Giovanni BARONE ADESI
University of Lugano, Dept. of Economics

Abstract

This paper provides estimates of aggregate preferences, beliefs, and sentiment from option prices and historical returns. Our market-based estimates correlate well with independent survey-based estimates, and yet deliver several novel insights. Our analysis points out two significant issues related to overconfidence. First, the Baker–Wurgler index strongly reflects excessive optimism but not overconfidence. Second, optimism and overconfidence comove over time and generate a perceived negative risk-return relationship, while objectively the relationship is positive

Friday, October 16, 2015, 10:30-12:00
Room 126, Extranef building at the University of Lausanne

Sentiment, Risk Aversion, and Time Preference*

Giovanni Barone-Adesi† Swiss Finance Institute and University of Lugano

Loriano Mancini‡ Swiss Finance Institute and EPFL

Hersh Shefrin§ Leavey School of Business Santa Clara University

August 28, 2015

Separate Appendix Available Online

*For helpful comments we thank Malcolm Baker, Kelley Bergsma, Menachem Brenner, Pierre Collin-Dufresne, Kent Daniel, Alex Edmans, Rob Engle, Stephen Figlewski, Simon Gilchrist, Paul Glasserman, Cam Harvey, Harrison Hong, Tom Howard, Danling Jiang, Elyès Jouini, Bryan Kelly, Fabio Maccheroni, Roberto Marfè, Veronika Pool (discussant), Norman Schuerhoff, Raghu Sundaram, Fabio Trojani, Raman Uppal, Wei Xiong, Jianfeng Yu, Bill Zame, Alexandre Ziegler and seminar participants at the 2013 European Finance Association meetings in Cambridge, New York University, Santa Clara University and Princeton-Lausanne workshop. Financial support from the Swiss National Science Foundation NCCR-FinRisk (Barone-Adesi and Mancini) and the Sinergia grant "Empirics of Financial Stability" [154445] (Mancini) is gratefully acknowledged. Shefrin acknowledges a course release grant from Santa Clara University.

†Giovanni Barone-Adesi, Institute of Finance, University of Lugano, Via G. Buffi 13, CH-6900 Lugano, Switzerland. E-mail: giovanni.baroneadesi@usi.ch.

‡Loriano Mancini, Swiss Finance Institute at EPFL, Quartier UNIL-Dorigny, CH-1015 Lausanne, Switzerland. E-mail: loriano.mancini@epfl.ch.

§Hersh Shefrin, Department of Finance, Santa Clara University, 500 El Camino Real, Santa Clara, CA 95053, USA. E-mail: hshefrin@scu.edu.

Sentiment, Risk Aversion, and Time Preference

Abstract

This paper provides estimates of aggregate preferences, beliefs, and sentiment from option prices and historical returns. Our market-based estimates correlate well with independent survey-based estimates, and yet deliver several novel insights. Our analysis points out two significant issues related to overconfidence. First, the Baker–Wurgler index strongly reflects excessive optimism but not overconfidence. Second, optimism and overconfidence comove over time and generate a perceived negative risk-return relationship, while objectively the relationship is positive.

Keywords: Sentiment, Pricing Kernel, Optimism, Overconfidence, Option Data

JEL Codes: G02, G12

1. Introduction

This paper proposes a behavioral approach for recovering *aggregate* preferences, beliefs, and sentiment from index option prices and historical returns. As it turns out, options data provide the richest source to identify variations in investors' sentiment. The recovery method is parsimonious and provides estimates of stochastic risk aversion, time preference, and return distributions. Our market-based estimates correlate well with independent survey-based estimates, and yet deliver several novel insights on aggregate investors' beliefs and sentiment.

Our empirical approach is based on Shefrin's theoretical framework (Shefrin, 2005, 2008) which focuses on the manner in which the pricing kernel reflects the market's aggregation of heterogeneous beliefs about returns, preferences about risk taking, and rates of impatience. Jouini and Napp (2006, 2007) and Dumas, Kurshev, and Uppal (2009) provide related general equilibrium approaches involving heterogeneity and sentiment. With this framework as a base, we use the approach developed in Barone-Adesi, Engle, and Mancini (2008) to estimate empirical pricing kernels and draw the relevant inferences about biases, stochastic risk aversion, and stochastic time preference.

Our main finding is that empirical pricing kernels strongly reflect behavioral elements. Our estimates of excessive optimism, overconfidence, risk aversion, and time preference, extracted from empirical pricing kernels, line up with independent estimates reported in the empirical behavioral literature. Specifically, our estimates of excessive optimism are highly correlated with the Baker and Wurgler (2006) sentiment series. Our estimates of overconfidence are related to a variety of survey-based volatility measures such as the Yale/Shiller crash confidence index, the financial executives data analyzed in Ben-David, Graham, and Harvey (2013),[1] as well as the bond default premium developed by Gilchrist and Zakrajšek (2012). Our estimates of risk aversion conform to the general pattern predicted by prospect theory, which suggests that risk aversion will be lower after market losses, than after market gains. Our estimates of time preference are consistent with the survey results reported by Barsky, Juster, Kimball, and Shapiro (1997) which feature negative time preference.

Several new insights arise from our analysis. The most important of these involves the Baker–Wurgler series for sentiment (Baker and Wurgler, 2006), which is the most widely used general measure of sentiment in the academic literature. As conjectured by Baker and Wurgler (2007), we

[1] Although financial executives' overconfidence from survey data is much larger than our estimates of the representative investor's overconfidence, the general patterns are similar.

find that the Baker–Wurgler series primarily captures excessive optimism about stock returns. However, we also show that the Baker–Wurgler series fails to capture the component of overconfidence associated with volatility which is uncorrelated with optimism, as well as sentiment associated with left tail events. We confirm this finding using independent series associated with overconfidence and tail events, namely the Yale/Shiller crash confidence series and the bond default series developed by Gilchrist and Zakrajšek (2012).

Our findings about the Baker–Wurgler series have implications for empirical asset pricing studies of investors' sentiment. For example, Yu and Yuan (2011) report that the relationship between risk and return, while positive when the Baker–Wurgler index is low, weakens when the Baker–Wurgler index is high to the point where it becomes insignificant. Although their informal argument is framed using overconfidence, our findings indicate that their empirical analysis, which relies on Baker–Wurgler as the sole measure of sentiment, fails to capture the component of overconfidence that is uncorrelated with excessive optimism.

Our results indicate that the relationship between risk and return is effectively driven by the comovements of excessive optimism and overconfidence over time, not just the level of excessive optimism. Specifically, we find that investors' excessive optimism and overconfidence comove over time, i.e., investors tend to overestimate (underestimate) future returns and underestimate (overestimate) future return volatility at the same time. This induces a perceived negative risk-return relationship, while objectively the relationship is positive.

Recovery and aggregation provide the backdrop for this paper, with both themes being prominent in the recent literature. Ross (2015) presents a new technique for recovering the constituent components of the pricing kernel, such as subjective beliefs, from a limited set of information. Borovička, Hansen, and Scheinkman (2015) point out that the probability distribution recovered by the Ross procedure typically does not correspond to the representative investor's subjective beliefs. The literature on recovery tends to assume the existence of a representative investor with a specific utility function such as constant relative risk aversion or mean-variance. We do not make this assumption.

The literature on aggregation goes back many decades. Recent studies describe how the preferences and beliefs of a representative investor reflect the underlying preferences and beliefs of heterogeneous investors. Shefrin (2005, 2008) provides an aggregation result (Theorem 14.1) for the case when all investors have time invariant CRRA preferences, with heterogeneity across coefficients of risk aversion, constant time preference parameters, and subjective stochastic processes.

His aggregation result establishes that the representative investor's conditional probability density functions are a generalized Hölder mixture of the individual investors' conditional probability density functions, a feature which can lead to non-monotonic time varying pricing kernels. Bhamra and Uppal (2014) effectively extend Shefrin's result to allow for habit formation, and Bhamra and Uppal (2015) discuss aggregation when individual investors have Epstein–Zin recursive utility.

Finally, in respect to sentiment there is a general understanding that the representative investor's subjective beliefs might differ from the underlying objective process governing the market dynamics, or at least historical probabilities. Ross uses the term "dark matter of finance" when discussing this point, especially when it comes to catastrophic events.[2] Nevertheless, his framework assumes homogenous beliefs, with no disagreement about investors' views about dark matter. Estimating how these dark matter issues are reflected in empirical pricing kernel is of critical importance, and is an important contribution of this paper.

The remainder of the paper is organized as follows. Section 2 presents the intuition underlying our approach. Section 3 describes our methodology for estimating the empirical pricing kernel. Section 4 reviews the theoretical framework for analyzing investors' sentiment. Section 5 presents our estimates of sentiment. Section 6 relates our findings to external measures of sentiment, risk aversion, and time preference. Section 7 concludes.

2. Intuition Underlying Our Approach

We define investors' sentiment in terms of a change of measure that links subjective and objective beliefs. To develop the intuition underlying our approach, we provide a brief nontechnical introduction. Our starting point is the standard neoclassical framework in which equilibrium prices are set as if by a representative investor holding correct beliefs. The objective probability density function (pdf) associated with correct beliefs, is depicted in the top panel of Figure 1, and is labeled Pobj.

In a behavioral framework, equilibrium prices are also set as if by a representative investor, but one whose beliefs possibly reflect biases in the investor population. Because of limits to arbitrage, investor biases are not necessarily eliminated in equilibrium. In the top panel of Figure 1, the function Prep denotes the pdf of the representative investor exhibiting two biases, excessive optimism and overconfidence. Relative to the objective pdf Pobj, excessive optimism means that the rep-

[2]Ross (2015, Page 616) states: The dark matter of finance is the very low probability of a catastrophic event and the impact that changes in that perceived probability can have on asset prices [. . .] Apparently, however, such events are not all that remote and "five sigma events" seem to occur with a frequency that belies their supposed low probability.

resentative investor overestimates expected return. Overconfidence means that the representative investor underestimates return standard deviation. In Figure 1, notice that the mode of Prep is to the right of the mode of Pobj, and Prep attaches much less weight to tail events than Pobj.

Formally, excessive optimism is defined as expected return under Prep minus expected return under Pobj. Overconfidence is defined as return standard deviation under Pobj minus return standard deviation under Prep. Operationally, we estimate Pobj and Prep and then compute excessive optimism and overconfidence from their first and second moments. To estimate Pobj, we use a dynamic model for S&P 500 returns. To estimate Prep, we use S&P 500 index option prices (SPX) and the risk free rate to infer the risk neutral pdf, and then apply a pricing kernel-based change of measure.

The pricing kernel lies at the heart of our process for inferring Prep. The pricing kernel is a function whose values are ratios of state prices to probabilities, which in this case we take to be objective probabilities Pobj. The bottom panel of Figure 1 displays three functions. The function CRRAKernel is the pricing kernel from a neoclassical representative investor model with CRRA preferences. As usual, the function is monotonically decreasing, reflects intertemporal marginal rate of substitution (through function slope), and measures time preference (through function height).

In contrast to CRRAKernel, the function BehavKernel in Figure 1 depicts a pricing kernel associated with a representative investor whose beliefs exhibit excessive optimism and overconfidence. Notice how overconfidence manifests itself in tail events where the BehavKernel function lies below CRRAKernel, as the behavioral representative investor underestimates tail event probabilities. Notably, in this example, the degree of overconfidence leads BehavKernel to feature an upward sloping portion in the left region of the figure. For the middle range, the combination of biases leads BehavKernel to lie above CRRAKernel, so that BehavKernel has the shape of an inverted-U.

We use estimates of BehavKernel, CRRAKernel and their difference to infer values for excessive optimism, overconfidence and other biases, and to disentangle their manifestation within prices. To estimate BehavKernel we use the ratio of the risk neutral pdf to the objective pdf. To estimate CRRAKernel we use a method described later in the paper. To capture the differences between the two pricing kernels, we use the log of BehavKernel minus the log of CRRAKernel, which is displayed as the function LogDiff in the bottom panel of Figure 1. We provide an exact interpretation of LogDiff later in the paper.

Our empirical measures of excessive optimism, overconfidence and other biases are computed relative to a process estimated from historical returns. We do not contend that historical returns are completely free of investor bias. Instead we investigate the extent to which market prices inform an econometrician's best estimate of future returns. Our approach is in the same vein as, e.g., Xiong and Yan (2010) who theoretically analyze the econometrician's prediction of bond premia when investors have heterogenous beliefs.

The representative investor always holds the market portfolio, and therefore does not "lose money" because of biases. To the extent that the representative investor corresponds to a real investor, the biases cause the representative investor to be disappointed and surprised. Optimism leads to disappointment in the realized risk premium, and overconfidence leads to surprise about the amount of volatility.

3. Method to Estimate the Empirical Pricing Kernel

By a pricing kernel we mean a stochastic discount factor (SDF) defined as state price per unit objective probability. Let $M_{t,T}$ denote the empirical SDF associated with returns between date t and date T, *conditional* on the information available at date $t \leq T$. Throughout the paper, $(T-t)$ is fixed and equal to one year. The empirical SDF is given by

$$M_{t,T} = e^{-r_f\,(T-t)} \frac{q(S_T/S_t)}{p(S_T/S_t)} \tag{1}$$

where q is the risk neutral density, p the objective or historical density, r_f the instantaneous risk free rate, and S_t the S&P 500 index at date t, which is a proxy for the market portfolio.[3] The densities q and p are conditional on the information available at date t, but for ease of notation we omit such a dependence. The risk free rate r_f depends on t and T, and such a dependence is omitted as well. Once the conditional densities q and p are estimated, we can recover the SDF by simply taking their discounted ratio, (1). The advantage of this procedure is that no constraint is imposed on the functional form of the SDF.

To estimate the empirical SDF we use the empirical approach in Barone-Adesi, Engle, and Mancini (2008) that we briefly review here. For each date t, we fit an asymmetric Glosten, Jagannathan, and Runkle (1993) GARCH model to historical daily log-returns of the S&P 500 to

[3]Multiple state variables can potentially enter the SDF and there is considerable debate among researchers over the relevant state variables. As in Aït-Sahalia and Lo (2000), Jackwerth (2000), Rosenberg and Engle (2002), and others, we consider the *projection* of the SDF into S&P 500 returns. As discussed in Cochrane (2005), this projected SDF has the same pricing implications as the original SDF for assets whose payoffs depend on S&P 500 returns, which is the relevant case in our setting, as we only consider call and put options on the S&P 500 index.

estimate the index dynamic under the objective pdf p. The model has the form

$$\log(S_u/S_{u-1}) = \mu_u + \epsilon_u \tag{2}$$

$$\sigma_u^2 = \omega + \beta\sigma_{u-1}^2 + \alpha\epsilon_{u-1}^2 + \gamma I_{u-1}\epsilon_{u-1}^2 \tag{3}$$

where $\epsilon_u = \sigma_u\, z_u$, z_u is the standardized historical innovation at day u, $I_{u-1} = 1$ when $\epsilon_{u-1} < 0$ and $I_{u-1} = 0$ otherwise, and $u = t_0, \ldots, t$, with $(t - t_0)$ being the time span of the daily sample starting at day t_0 and ending at day t. When $\gamma > 0$, the model features the so-called leverage effect, namely bad news ($\epsilon_{u-1} < 0$) raises future volatility more than good news ($\epsilon_{u-1} \geq 0$) of the same absolute magnitude. The scaled return innovation, z_u, is from its empirical density function, which is obtained by dividing each estimated return innovations, $\hat{\epsilon}_u$, by its estimated conditional volatility $\hat{\sigma}_u$. This set of estimated scaled innovations gives an empirical density function that incorporates excess skewness, kurtosis, and other extreme return behaviors that are not captured in a normal density. This approach is called filtered historical simulation (FHS). The drift term is specified as $\mu_u = 0.012 + 0.76\,(\mathrm{E/P})_u$, where E/P is the inverse of the price-earnings ratio, adjusted for inflation, developed by Campbell and Shiller (1998).[4] The online appendix shows that our subsequent results remain virtually unchanged when the excess risk premium (in excess of the risk free rate) is set to a constant value of 4%, rather than based on E/P. The GARCH parameter estimates are obtained by maximizing the Pseudo Maximum Likelihood, under the nominal, not necessarily true, assumption of normal innovations. This technique provides consistent parameter estimates even when the true innovation density is not normal, e.g., White (1982). Rosenberg and Engle (2002) use the same approach to estimate the objective distribution of S&P 500 returns in their analysis.

For each date t, a GARCH model (2)–(3) is calibrated to the cross section of out-of-the-money call and put options on the S&P 500 to estimate the index dynamic under the risk neutral pdf q. Given a set of risk neutral GARCH parameters $\{\omega^*, \beta^*, \alpha^*, \gamma^*\}$, a return path is simulated by drawing an estimated past innovation, say, $z_{[1]}$, updating the conditional variance σ_{t+1}^2, drawing a second innovation $z_{[2]}$, updating the conditional variance σ_{t+2}^2, and so on up to day T. Let $\tau = T - t$. The τ-period simulated gross return is $S_{t+\tau}/S_t = \exp(\tau\mu^* + \sum_{i=1}^{\tau} \sigma_{t+i}\, z_{[i]})$, where the drift μ^* is such that the average gross return equals the risk free gross rate e^{r_f}, and is determined

[4]The coefficients in μ_u are obtained by regressing subsequent annualized ten-year returns for the Campbell–Shiller series on a constant and E/P. Campbell and Shiller (1998) show that subsequent ten-year returns to stocks are negatively and statistically related to the price-earnings ratio. Thus, the specification $\mu_u = 0.012 + 0.76\,(\mathrm{E/P})_u$ is forward looking. Updated data series of the price-earnings ratio are available from Robert Shiller's website, http://www.econ.yale.edu/~shiller/.

using the Empirical Martingale Simulation method of Duan and Simonato (1998). We simulate L = 20,000 return paths from t to $t+\tau$. The GARCH call option price at time t with strike price K and time to maturity τ is given by $e^{-r_f \tau} \sum_{l=1}^{L} \max(S_{t+\tau}^{(l)} - K, 0)/L$, where $S_{t+\tau}^{(l)}$ is the simulated index price at time $t+\tau$ in the l-th sample path. Put prices are computed similarly. The risk neutral GARCH parameters $\{\omega^*, \beta^*, \alpha^*, \gamma^*\}$ are varied, which changes the simulated return paths, so as to best fit the cross-section of option prices at date t, minimizing the mean square pricing error $\sum_{j=1}^{N_t} e_t(K_j, \tau_j)^2$, where $e_t(K_j, \tau_j)$ is the difference between the GARCH option price and the market price of the option with strike K_j and time to maturity τ_j, and N_t is the number of options available at day t. The calibration is achieved when, varying the risk neutral GARCH parameters, the reduction in the mean square pricing error is negligible or below a given threshold.[5]

Having estimated objective and risk neutral GARCH parameters on a given date t, the next step to recover the SDF is the estimation of the conditional densities $p(S_T/S_t)$ and $q(S_T/S_t)$. For each date t, these conditional densities are estimated by Monte Carlo Simulation. Given the objective GARCH parameters, we simulate 50,000 return paths of the index at a daily frequency from t to T using the simulation method above. We also simulate 50,000 return paths using the risk neutral GARCH parameters. The conditional densities p and q are obtained by nonparametric kernel density estimation, i.e., smoothing the corresponding simulated distribution of S_T/S_t. Finally, the empirical SDF, $M_{t,T}$, is estimated by computing the discounted ratio of the two densities, as in (1). The whole procedure is repeated for each day t in our sample, producing a time series of functions $M_{t,T}$.

We consider two GARCH models under the risk neutral density q that lead to two estimates of the empirical SDF. One we call Gauss and the other we call FHS. The Gaussian model uses randomly drawn Gaussian innovations for the simulation of the return paths, whereas the FHS model uses the historical, nonparametric innovations z_u, as described above. We use both models in order to contrast the difference that FHS makes.

Risk neutral densities are often estimated by differentiating twice the call pricing function, which can be estimated by interpolating the implied volatility smile. This method could work well for fixed maturities and when many option prices are available, namely for relatively short maturities and when S_T/S_t is approximately one. This is not case in our setting. We are interested in a long and fixed time horizon of one year, and there are essentially no options with time to maturity of

[5]To ensure the convergence of the calibration algorithm, all the FHS innovations, $z_{[i]}$, used to simulate the L return paths are kept fix across all the iterations of the algorithm. Starting values for the risk neutral parameters are the GARCH parameters estimated under the objective measure.

one year for each day t. This motivates our approach of calibrating risk neutral GARCH models to the cross-section of options.

4. Framework for Sentiment

In this section, we define sentiment and discuss its estimation.

4.1. *Definition of Sentiment*

Sentiment impacts the SDF by distorting state prices relative to a neoclassical counterpart. Therefore, a behavioral SDF effectively decomposes into a neoclassical component and a sentiment distortion. As for the neoclassical component, we start with a constant relative risk aversion (CRRA) SDF, which is the standard benchmark in neoclassical theory. Importantly, all our subsequent results about investors' sentiment are virtually unchanged when the CRRA SDF is replaced by simply a *monotonic* non-increasing marginal utility function.

The CRRA SDF has the following form:

$$M_{t,T}(\theta) = \theta_0 \left(S_T/S_t\right)^{-\theta_1} \tag{4}$$

where θ_0 is a discount factor measuring the degree of impatience, θ_1 is the coefficient of relative risk aversion, and $\theta = (\theta_0, \theta_1)$. The logarithmic version of (4) is

$$\log(M_{t,T}(\theta)) = \log(\theta_0) - \theta_1 \log(S_T/S_t). \tag{5}$$

In Shefrin (2008), (5) generalizes to include an additional term $\Lambda_{t,T}$ to reflect the impact of sentiment. The equation for the log-SDF becomes

$$\log(M_{t,T}) = \Lambda_{t,T} + \log(\theta_{0,t}) - \theta_{1,t} \log(S_T/S_t) \tag{6}$$

where the parameter θ is now time varying.[6] Appendix A provides a derivation of (6). We define sentiment as the function $\Lambda_{t,T}$. This function is a scaled log-change of measure, where the change of measure transforms the objective pdf p into the representative investor's pdf p_R. In other words,

[6] $\Lambda_{t,T}$ is a function of S_T/S_t, $\theta_{1,t}$ is a function of T and S_T/S_t, and $\theta_{0,t}$ is a function of T. For simplicity, we omit such dependencies. Shefrin (2008, Theorem 14.1) implies that θ is not just time varying, but stochastic as well.

the function $e^{\Lambda_{t,T}}$ is proportional to the change of measure p_R/p so that

$$p_R = p\, e^{\Lambda_{t,T}}\, \theta_{0,t,p}/\theta_{0,t} \tag{7}$$

where $\theta_{0,t,p}$ is a rescaling of $\theta_{0,t}$ whose purpose is to ensure that p_R integrates to one.

The log-change of measure $\log(p_R/p)$ specifies the percentage error in probability density which the representative investor assigns to the occurrence of a specific return. For example, suppose that the representative investor underestimates by 2% the probability that the market return will be 1%. In this case, the log-change of measure at 1% will be -2%.

In a Gaussian setting, a log-linear change of measure generates a variance preserving shift in mean. If the mean shifts to the right, the log-change of measure is a positively sloped linear function which, when applied to p, shifts probability mass from low values to high values. If the mean shifts to the left, the log-change of measure is a negatively sloped linear function. To put it another way, a positively sloped log-linear change of measure gives rise to excessive optimism, while a negatively sloped log-linear change of measure gives rise to excessive pessimism.

If the log-change of measure is non-linear, then applying the change of measure impacts the second moment. A log-change of measure with a U-shape shifts probability mass from the center to the tails, thereby increasing the variance return. A log-change of measure with an inverted U-shape shifts probability mass from the tails into the center, thereby lowering the variance return. To put it another way, a U-shape gives rise to underconfidence, whereas an inverted U-shape gives rise to overconfidence. With respect to (6), if $\Lambda_{t,T}$ is large enough, then the shape of the sentiment function will dominate the shape of the fundamental component. For example, if the log-change of measure has an inverted U-shape which is sufficiently large, then $\Lambda_{t,T}$ will overpower the other terms in (6), and the log-SDF will also have an inverted U-shape.

If the market reflects a mix of optimists and pessimists with optimism and overconfidence being positively correlated, then log-sentiment can feature an oscillating pattern which is sharply downward sloping in the left tail, upward sloping in the middle region, and downward sloping in the right tail. It is this shape which characterizes the empirical findings for the shape of the pricing kernel in Aït-Sahalia and Lo (2000) and Rosenberg and Engle (2002).

In neoclassical pricing theory, the risk neutral pdf q can be obtained from the objective pdf p by applying a change of measure using the normalized pricing kernel; e.g., Cochrane (2005, Page 51). Of course, this relationship can be inverted to express p as a function of q. In the behavioral framework, an analogous relationship holds between the representative investor's pdf p_R and q, rather than between p and q. The expression for p_R as a function of q is

$$p_R(S_T/S_t) = q(S_T/S_t)\ (S_T/S_t)^{\theta_{1,t}}\ E_t^{p_R}[(S_T/S_t)^{-\theta_{1,t}}] \tag{8}$$

where $E_t^{p_R}$ is the time-t conditional expectation with respect to p_R.

4.2. *Estimation of Sentiment*

We decompose the log-SDF in (6) into its constituent components, the sentiment function $\Lambda_{t,T}$ and a neoclassical fundamental component. For each day t, using the procedure described in Section 3, we obtain a grid of 100 values of gross returns, $S_T^{(i)}/S_t$, $i = 1, \ldots, 100$, spanning the support of the empirical SDF. Then we regress the empirical log-SDF, $\log(M_{t,T}^{(i)})$, on a constant and the log gross return, $\log(S_T^{(i)}/S_t)$.[7] Intercept and slope provide estimates of $\log(\theta_{0,t})$ and $-\theta_{1,t}$, respectively, allowing us to estimate $\log(M_{t,T}^{(i)}(\theta))$. For each gross return $S_T^{(i)}/S_t$, we compute the pointwise difference

$$d_{t,T}^{(i)} = \log(M_{t,T}^{(i)}) - \log(M_{t,T}^{(i)}(\theta)). \tag{9}$$

The differences, $d_{t,T}^{(i)}$, $i = 1, \ldots, 100$, provide an estimate of the sentiment function $\Lambda_{t,T}$ over the support of gross returns, $S_T^{(i)}/S_t$, $i = 1, \ldots, 100$. We repeat this procedure for each day t, and obtain a time series of $\theta_{0,t}$, $\theta_{1,t}$ and the sentiment functions $\Lambda_{t,T}$.

As a major robustness check of our results, we replace the CRRA SDF by a monotonic non-increasing function of $\log(S_T/S_t)$, simply ensuring non-increasing marginal utility of the representative investor. Then we re-estimate sentiment (and all other related variables) for each date t in our sample. For each date t we fit the empirical log-SDF using monotonic regressions and then take the residuals as the estimate of sentiment. The online appendix describes the procedure in detail and shows that all our results below remain largely intact when sentiment is derived from a neoclassical component only featuring non-increasing marginal utility.

In a neoclassical setting, a monotonic pricing kernel only presumes a non-increasing marginal utility and does not impose any further restriction on the representative investor's utility function. Thus, non-increasing marginal utility is a minimal requirement that the representative investor's utility function must satisfy. Indeed, Dybvig (1988) show that if the SDF *projected* on the market

[7] This method is different than Rosenberg and Engle (2002). They calibrated the constrained CRRA SDF directly to option prices, whereas we fit the constrained CRRA SDF to the unconstrained empirical SDF. The two methods give the same result when the empirical SDF conforms to the CRRA pricing kernel.

return (as it is the case in our setting) is not non-increasing, then it is possible to construct a contingent claim with payoff function $f_t(S_T)$, necessarily decreasing over some region, such that $f_t(S_T)$ and S_T have the same objective distribution conditional on information available at day t, and yet $E_t^p[M_{t,T}\, f_t(S_T)] < E_t^p[M_{t,T}\, S_T]$, i.e., such a contingent claim is less expensive than the market index, which is indeed puzzling, and calls for explanations of non-monotonic pricing kernels. Because our results about investors' sentiment remain are largely unchanged when imposing only non-increasing marginal utility of the representative investor, this strongly indicates that our findings are not driven by important misspecifications of the neoclassical pricing kernel.

5. Empirical Results: Pricing Kernels, Sentiment, and Beliefs

Output from the estimation procedure in Sections 3–4 consists of a series of estimates for the objective and risk neutral GARCH parameters, the SDF ($M_{t,T}$), CRRA ($\theta_{1,t}$), time preference ($\theta_{0,t}$), the objective return pdf ($p(S_T/S_t)$), the risk neutral pdf ($q(S_T/S_t)$), the representative investor's pdf ($p_R(S_T/S_t)$), and sentiment ($\Lambda_{t,T}$). We first describe our dataset and then discuss the main features of the estimation results.

5.1. Dataset

We use European options on the S&P 500 index (symbol: SPX) to calibrate the risk neutral GARCH models. SPX options are among the most actively traded index options in the world, have no wild card features, and can be hedged using S&P 500 futures.

We use closing prices of out-of-the-money (OTM) put and call options on Wednesdays from January 2, 2002 to October 28, 2009. It is known that OTM options are more actively traded than in-the-money options. Option data and all the other necessary data are downloaded from OptionMetrics. The average of bid and ask prices are taken as option prices. Options with time to maturity less than 10 days or more than 360 days, or prices less than \$0.05 are discarded. From January 2, 2002 to December 29, 2004, a relatively low volatility period, we also discard options with implied volatility larger than 70%, as in Barone-Adesi, Engle, and Mancini (2008). From January 5, 2005 to October 28, 2009, a relatively high volatility period, we only discard options with implied volatility larger than 150%. This procedure yields a sample of 121,243 options, which are roughly split in calls (45.5%) and puts (54.5%).

Using the term structure of zero coupon rates, the risk free rate for each option maturity is obtained by linearly interpolating the two interest rates whose maturities straddle the given maturity. This procedure is repeated for each contract and each day in the sample.

We divide the option data into several categories according to time to maturity and moneyness, m, which is defined as the ratio of the strike price over the S&P 500 index. A put option is said to be deep OTM if its moneyness $m < 0.85$, or OTM if $0.85 \leq m < 1$. A call option is said to be OTM if $1 \leq m < 1.15$, or deep OTM if $m \geq 1.15$. We also classify option contracts according to the time to maturity: short maturity (< 60 days), medium maturity (60–160 days), or long maturity (> 160 days).

Table 1 describes the 121,243 option prices, and their implied volatilities. The average put (call) prices range from \$1.31 (\$0.67) for short maturity, deep OTM options to \$43.95 (\$44.83) for long maturity, OTM options. OTM put and call options account for 28% and 25%, respectively, of the total sample. Short and long maturity options account for 40% and 29%, respectively, of the total sample. The table also shows the familiar volatility smile and the corresponding term structure. The smile across moneyness is evident for each set of maturities. When the time to maturity increases, the smile tends to become flatter. The number of options on each Wednesday is on average 296.4, with a standard deviation of 127.8, a minimum of 142, and a maximum of 726 option contracts. The average moneyness of OTM put is 0.81, with standard deviation of 0.16, and minimum value of 0.18. The average moneyness of OTM call is 1.21, with standard deviation of 0.24, and maximum value of 3.51. Importantly, our estimates of the empirical pricing kernel pertain to the range of gross returns of about 0.69 to 1.35 and time horizon of one year. The range of gross returns and time horizon are well *within* the span of option moneyness and time to maturities, respectively. Thus there is no extrapolation bias in our estimates.

During our sample period, the S&P 500 ranges from a minimum of \$676.5 to a maximum of \$1,565.2, with an average level of \$1,157.7. The average daily log-return is close to zero (-5.2×10^{-5}), the standard deviation is 22.4% on an annual base, and skewness and kurtosis are -0.13 and 12, respectively. In particular, the high kurtosis of S&P 500 returns appears to be due to the large market swings in the fall 2008.

5.2. GARCH Estimation and Calibration

Table 2 shows estimates of the objective and risk neutral GARCH models (2)–(3). Objective GARCH parameters are estimated rather precisely and exhibit little variation over time. For each day t in our sample, the online appendix reports Ljung–Box and Lagrange Multiplier ARCH tests for squared daily returns and squared standardized historical innovations. These tests show that the GARCH model is highly effective in removing volatility clustering in S&P 500 returns, which is well-known from the GARCH literature.

Risk neutral GARCH parameters exhibit more time variation, but the persistency and long-run mean of the GARCH volatility are estimated quite precisely. FHS GARCH parameters are generally less volatile than Gauss GARCH parameters, especially for the long-run mean volatility. Risk neutral GARCH volatilities appear to be larger and less persistent on average than objective GARCH volatilities.[8] These findings are in line with a recent literature investigating variance risk premiums, e.g., Carr and Wu (2009), Bollerslev and Todorov (2011), and Aït-Sahalia, Karaman, and Mancini (2014).

Table 3 shows mean and root mean square error of option price errors of the risk neutral GARCH model based on the FHS method. The price error is defined as model-based option price minus market option price. Average price errors tend to be positive, but root mean square errors across all moneyness/maturity categories are small and in line with those reported in Barone-Adesi, Engle, and Mancini (2008). The online appendix shows the fitting of the GARCH model to SPX options and visually confirms the good fit of the model.

5.3. *Pricing Kernel Over Time*

Figure 2 displays the empirical SDF estimated on each Wednesday from January 2002 to October 2009 using the FHS method. At the beginning of our sample period, the pricing kernel featured a declining pattern. By December of 2003, the pricing kernel featured a U-shape. During 2005, the shape of the pricing kernel had changed to an inverted-U. In 2009, the pricing kernel became steeper, similar to what it had been at the beginning of the sample period. The empirical SDF based on the Gauss method features a similar evolution over time, and is reported in the online appendix. However, it is significantly steeper to the left, which is in line with the findings in Barone-Adesi, Engle, and Mancini (2008).

5.4. *Representative Investor's Beliefs, Optimism and Overconfidence*

Equations (7) and (8) provide the basis for estimating the beliefs p_R of the representative investor. Equation (7) shows that $d_{t,T}$ is a scaled estimate for the sentiment function $\Lambda_{t,T}$. Therefore $e^{d_{t,T}}$

[8] We compared our risk neutral pdf estimates with Birru and Figlewski (2012), who use a shorter time to expiration than we do. Notably, the general patterns we find appear to be similar to those in Birru and Figlewski (2012).

can be interpreted as being proportional to a change of measure which transforms the objective density p into the representative investor's density p_R.[9]

Figure 3 displays our estimates of optimism and overconfidence. Optimism is defined as the difference between the expected market return under the representative investor's and objective pdfs, i.e., $E_t^{p_R}[S_T/S_t] - E_t^p[S_T/S_t]$, where E_t^p is the conditional expectation under the objective pdf, computed by numerically integrating the gross return against p, and similarly for $E_t^{p_R}$. Overconfidence is defined as the difference between the expected volatility of the market return under objective and representative investor's pdfs, i.e., $\sqrt{\mathrm{Var}_t^p[S_T/S_t]} - \sqrt{\mathrm{Var}_t^{p_R}[S_T/S_t]}$. With the exception of the period following the Lehman bankruptcy in September 2008, both optimism and overconfidence generally rose and fell with the market, exhibiting procyclical behavior. The correlation coefficient for the two variables is 0.5. In the middle of the sample, which is a relatively low volatility period of stable market growth, the representative investor is excessively optimistic and overconfident, judging the expected return as too high and the future volatility as too low. Notably, this pattern is reversed at the beginning and end of the sample period, which are more turbulent periods, when the representative investor is pessimistic and at times underconfident, especially after fall 2008.

The online appendix shows that estimates of optimism and overconfidence are nearly the same as those in Figure 3, when the CRRA SDF is replaced by a monotonic non-increasing SDF. The online appendix also reports estimates of optimism and overconfidence (and other quantities) when the expected return, in excess of the risk free rate, is set to a constant value of 4%. Optimism is more stable over time, but is still economically important. Overconfidence is virtually unaffected by the alternative specification of the expected return. The correlation between optimism and overconfidence is positive and high, 0.6.

5.5. *Impact of Sentiment on Equity and Variance Risk Premiums*

We now discuss the impact of excessive optimism and overconfidence on equity and variance risk premiums. The equity risk premium is the difference between the expected return and the risk free rate. There are two equity risk premiums, one associated with the objective pdf and the other associated with the representative investor's pdf. The objective equity risk premium is negatively correlated with both excessive optimism (-0.9) and overconfidence (-0.5). The signs are consistent

[9]As an example, Figure 1 shows the typical shape of the $d_{t,T}$ function during the middle portion of our sample period, for December 21, 2005. For this day, the $d_{t,T}$ function is positive between 0.99 and 1.16, and negative outside this interval. This means that a change of measure based on $d_{t,T}$, when applied to p, will shift probability mass to the region [0.99, 1.16] from the tails. The modes of p and p_R are about the same, but the probability mass of p is more spread out than the mass of p_R.

with the intuition that increases in excessive optimism and overconfidence drive up prices, thereby reducing the risk premium. Table 4 shows the regression results of the objective risk premium on its most recent lagged value, excessive optimism, and overconfidence. The most recent lag is included as a regressor to control for the autocorrelation of the risk premium. The coefficient on excessive optimism is negative (t-statistic = -2.48), and the coefficient on overconfidence is not statistically significant (t-statistic = 1.46). Thus, the equity risk premium appears to be more affected by optimism than overconfidence.

In a similar regression for the representative investor's equity risk premium, optimism and overconfidence have the opposite signs than the regression above and are both statistically significant, see Table 4. When optimism increases and overconfidence decreases, the representative investor perceives (incorrectly) that the equity risk premium increases. From the representative investor's viewpoint, this is an obvious consequence, given the perceived risk-return trade-off.

The variance risk premium is the difference between the return variance under the objective and risk neutral distributions, and arises as soon as investors require compensation for volatility risk.[10] Based on our estimates with the FHS method, the average variance risk premium is negative and around -1.4% ($= 0.198 - 0.212$, see Table 2) in volatility units, which is in line with the literature. Similarly to the equity risk premium, there are two variance risk premiums, one objective and one perceived by the representative investor. Table 4 shows the regression results of the variance risk premium on its most recent lagged value, excessive optimism, and overconfidence. The objective variance risk premium is significantly affected by overconfidence, but not by excessive optimism. An increase in overconfidence induces a less negative variance risk premium, reducing the risk premium in absolute value. This finding suggests that when investors underestimate return volatility, they also require a lower risk premium for volatility risk. The variance risk premium perceived by the representative investor appears to follow a very persistent dynamic and is unaffected by biases, once we control for its own autocorrelation.

Turning to the dynamic of excessive optimism and overconfidence, we regress these variables on its own lagged values, past one year S&P 500 return and past one year S&P 500 volatility computed using the standard deviation of daily log-returns during that year. Table 5 shows the regression results. Excessive optimism is positively related to past one year returns (t-statistic = 2.17). Both

[10] A fast growing literature studies variance risk premiums, e.g., Jiang and Tian (2005), Carr and Wu (2009), Bollerslev, Tauchen, and Zhou (2009), Todorov (2010), Bollerslev and Todorov (2011), Drechsler and Yaron (2011), Mueller, Vedolin, and Yen (2011), Aït-Sahalia, Karaman, and Mancini (2014), Buraschi, Trojani, and Vedolin (2014), and Filipović, Gourier, and Mancini (2015).

excessive optimism and overconfidence are negatively related to past volatility (t-statistics are −4.22 and −2.95, respectively).[11]

Chaining these relationships together, we have the following: High past returns and low volatility lead to high excessive optimism and high overconfidence. In turn, high optimism, which is likely to grow when the S&P 500 increases smoothly, leads to a low equity risk premium. High overconfidence, which is likely to grow during low volatility periods, leads to a low variance risk premium.

6. Treatment of Sentiment, Risk Aversion, and Time Preference

In this section we demonstrate that our measure of sentiment is parsimonious, strongly reflects a disparate collection of other sentiment measures and yet contains additional information, and yields estimates for risk aversion and time preference that are consistent with findings in the behavioral literature.

To demonstrate parsimony we compare our estimates of sentiment with four independent measures, namely the Baker–Wurgler series, the Duke/CFO survey responses, the Yale/Shiller crash confidence index, and the excess bond premium of Gilchrist and Zakrajšek (2012). We also compare our results to the analysis of Yu and Yuan (2011) who use the Baker–Wurgler series to study how risk and return are related over time. In addition, we examine other aspects of sentiment, besides optimism and overconfidence, such as biases associated with skewness, kurtosis, and left tail events (crashes). In respect to risk aversion and time preference, we relate our estimates of the time series for θ in (6) to the survey evidence presented by Meyer and Meyer (2005) and Barsky, Juster, Kimball, and Shapiro (1997).

6.1. Relationship of Biases to the Baker–Wurgler Series

We analyze the relationship between the Baker and Wurgler (2006) series (BW) and variables based on our estimates of sentiment.[12] Baker and Wurgler do not provide a precise interpretation of what

[11] As robustness checks we re-run these regressions using end-of-month rather than weekly observations, using six months rather than one year returns, as well as other measures of past volatility, such as six months squared returns, standard deviations of monthly returns and high minus low values of returns during the prior twelve months. Regression results remain largely unchanged.

[12] Baker and Wurgler develop two series, one which reflects economic fundamentals and a second which removes the effect of economic fundamentals. We analyze both series and the results are quite similar for both. For this reason we only report findings for the first series. The Baker–Wurgler monthly series for sentiment is available at Jeff Wurgler's website, http://people.stern.nyu.edu/jwurgler/.

their series exactly measures, although they do suggest thinking about the series as if it measures excessive optimism for stocks.

We find that BW strongly reflects excessive optimism. Table 6 shows a regression of BW on its two most recent lagged values, excessive optimism and overconfidence. The t-statistic for excessive optimism is 5.32. Including the S&P 500 monthly returns and the VIX volatility index as regressors, the t-statistic remains high at 3.75. The online appendix shows that optimism has a significant and strong impact on BW when estimating sentiment using a monotonic non-increasing SDF (rather than a CRRA SDF), and when setting the excess expected return of the S&P 500 to 4% (rather than specifying the expected return as the inverse of the price-earnings ratio). Also, allowing for an AR(2) error term in the regression, the t-statistic of optimism remains high at 3.54. Figure 4 (top panel) visually confirms that BW and optimism comove significantly during our sample period.

Although the coefficient of overconfidence is significant in Table 6, this finding is not robust. Controlling for S&P 500 returns and VIX index, the t-statistic of overconfidence drops to -1.78 when sentiment is measured using a monotonic non-increasing SDF, and to -1.07 when the excess expected return of the S&P 500 is set to 4%; see the online appendix. Also, allowing for an AR(2) error term, the t-statistic of overconfidence is only -1.48. Thus, the statistical significance for overconfidence is not robust to alternative specifications.

Table 6 also shows that the VIX index has a significant negative impact on the BW series, when *not* controlling for optimism and overconfidence. Including our sentiment variables as regressors makes the impact of the VIX index on the BW series disappear. This indicates that our sentiment variables subsume the information in the VIX index which is related to the BW series dynamics. The online appendix shows that this finding holds true also when sentiment is estimated using a monotonic non-increasing SDF, and the excess expected return of the S&P 500 is set to 4%.

Although the BW series weakly and negatively reflects overconfidence, our estimated sentiment functions indicate significant overconfidence in much of our sample period. Recall that overconfidence is associated with a sentiment function, or log-change of measure, that has the shape of an inverted U. Figure 5 illustrates several sentiment functions for the first nine months of 2002. Notice the pronounced inverted U-shapes. We conclude that the BW series fails to capture an important aspect of sentiment, namely the overconfidence component that is independent of excessive optimism.

6.2. *Risk, Return, and Sentiment*

The existence of a positive relationship between risk and return is a cornerstone concept of academic finance.[13] Yu and Yuan (2011) report that the relationship between risk and return, while positive when the Baker–Wurgler index is low, weakens when the Baker–Wurgler index is high to the point where it becomes insignificant. Yu and Yuan suggest that *overconfidence* plays a role in the risk-return dynamics, but do not include a measure of overconfidence in their formal analysis. Because our approach to sentiment is quite general, we are able to extend the analysis of the risk-return trade-off to incorporate overconfidence.

We regress (ex-ante) expected return on (ex-ante) return standard deviation, and a constant, under the conditional objective pdf p. Table 7 shows the regression results. The slope coefficient is 0.12 and the intercept is 0.02, and both estimates are statistically significant. Using end-of-month observations, regression estimates are nearly the same as in Table 7, and statistically significant. These parameter values are generally consistent with neoclassical theory, i.e., a positive risk-return trade-off.

In the behavioral approach, prices reflect not the objective pdf p but the representative investor's pdf p_R. A regression of (ex-ante) expected return on (ex-ante) return standard deviation under the representative investor's pdf p_R has a slope coefficient of -0.13 and an intercept of 0.07, and both estimates are statistically significant. Using end-of-month observations, the slope coefficient is -0.11 and the intercept is 0.06, again both statistically significant. The online appendix provides additional robustness checks, e.g., when sentiment is derived from a monotonic non-increasing SDF. The negative slope coefficient reflects the perspective that risk and return are negatively related. Shefrin (2008) discusses several studies about the *perception* that risk and return are negatively related.[14] One key behavioral feature involves excessive optimism and overconfidence being positively correlated. In our data, the correlation between the two series is 0.5. A positive correlation

[13]A large empirical literature studies the risk-return trade-off. After two decades of empirical research, there is little consensus on the basic properties of the relation between the expected market return and volatility. Studies such as Goyal and Santa-Clara (2003), Ghysels, Santa-Clara, and Valkanov (2005), Guo and Whitelaw (2006), and Ludvigson and Ng (2007) find a positive trade-off, while conversely Nelson (1991), Glosten, Jagannathan, and Runkle (1993), Brandt and Kang (2004), and Conrad, Dittmar, and Ghysels (2013) find a negative trade-off. Harvey (1989, 2001) shows that the risk-return relation changes over time. Rossi and Timmermann (2015) provide empirical evidence that the relation may be not linear.

[14]These studies focus on behavior at the level of the individual, and suggest that excessive optimism and overconfidence are positively correlated across the population. From a dynamic perspective, wealth transfers resulting from trading will induce a time series correlation as well. This occurs as wealth shifts from, say, less optimistic, less confident investors to more optimistic, more confident investors, inducing an increase over time in both the representative investor's degree of excessive optimism and overconfidence. Recently, Marfè (2015) developed an equilibrium model to account for a negative risk-return trade-off.

implies that whenever the representative investor is excessively optimistic and overestimates expected return, he tends to be overconfident and underestimates future volatility. Associating high returns to low risk is the hallmark of a negative perceived relationship.

We hasten to add that a positive correlation between excessive optimism and overconfidence does not necessarily induce a negative perceived relationship between risk and return. This is because if sentiment is small, then it will not override the fundamental component. For example, during the period September 2008 through the end of our sample period, the representative investor's perceived risk and return were positively correlated (with a regression slope coefficient of 0.07). At the same time, excessive optimism and overconfidence were still positively related, with a correlation coefficient of 0.8, but sentiment was small during this period.[15] Figure 6 shows the risk-return relationships under the objective and representative investor's pdfs.

Like Yu and Yuan (2011), we find that for the objective pdf, the relationship between risk and return is weaker when the Baker–Wurgler sentiment is positive than when it is negative.[16] Yu and Yuan suggest that this is because when sentiment is high, constraints on short sales magnify the impact of investor errors.

When we perform the same analysis for the representative investor's pdf encapsulating the "market's perception" of risk and return, which provides the basis for pricing assets, we find no statistically discernable difference between periods of high sentiment and periods of low sentiment. Rather, our analysis indicates that the negative risk-return relationship stems from excessive optimism and overconfidence being positively correlated and strong. Yu and Yuan (2011) base their regression analysis on the Baker–Wurgler index, which is effectively a measure only of excessive optimism. Our findings indicate that the perceived risk-return trade-off is driven by the co-movements of excessive optimism and overconfidence over time, not just the level of excessive optimism.

6.3. Duke/CFO Survey Responses

To provide another external check on our sentiment estimates, we use the Duke/CFO survey data. The questions in the Duke/CFO survey that are most relevant to our study pertain to expected S&P 500 return, volatility, and skewness, for a one year horizon. Graham and Harvey (2012)

[15] Distance measures between the log empirical SDF and log CRRA SDF (labeled RMSE and MAE) in the online appendix provide an assessment for the level of sentiment at each date t.

[16] In going from negative to positive sentiment in our analysis, the coefficient of expected return on standard deviation drops by about half, from 0.12 to 0.07. Both coefficients are statistically significant, implying that we do not find the relationship to become flat when sentiment is high.

describe how the survey is conducted and provide an overview of the survey results.[17]

The estimates for expected return, volatility, and skewness that are derived form the Duke/CFO survey responses provide an interesting contrast to our estimates from the representative investor's pdf. Figure 7 (top panel) shows that the Duke/CFO expected return and representative investor's expected return are highly correlated after 2005, with a correlation coefficient of 0.6.[18] For the entire sample period, the correlation coefficient is 0.2.

As discussed in Ben-David, Graham, and Harvey (2013), the Duke/CFO series exhibits very large overconfidence, with an average one year return volatility around 5%. The representative investor's conditional return volatility is around 20%, and thus more in line with historical levels. Interestingly, the correlation between the Duke/CFO volatility series and the representative investor's return volatility is a very high 0.8; see Figure 7 (bottom panel). Although the two volatility predictions are an order of magnitude different, the two measures comove strongly.

As for skewness, the correlation between the Duke/CFO values and the representative investor values is negative (-0.4). The former features an inverted-U shape over time, while the latter is U-shaped over time, as shown in the online appendix. This suggests that when volatility increases at the beginning and end of our sample, the respondents to the Duke/CFO survey overfocus on volatility associated with negative returns. In contrast, the representative investor focuses on high positive returns, as well as negative returns, during periods of heightened volatility.

6.4. Yale/Shiller Crash Confidence Indexes

Next we turn our attention to left tail events or crashes. To do so, we compare the probability of a left tail event under the representative investor's pdf with two independent survey-based counterparts, the Yale/Shiller crash confidence indexes for professional investors (CP) and for individual investors (CI). Each crash confidence index is the percent of respondents who attach little probability to a stock market crash in the next six months.[19] Thus, a high value of the index means a low probability of a market crash, according to the respondents.

To compare the crash confidence indexes CP and CI with our representative investor approach, we consider left tail probabilities under the representative investor's and objective pdfs. For each

[17] An archive of past surveys is available under the "Past Results" tab at http://www.cfosurvey.org.

[18] The sample of CFOs changed in 2004 when Duke changed survey partners from Financial Executives International to CFO magazine. For this reason, the data from 2005 on appears to be more consistent than the data from the earlier period.

[19] A detail description of the index and corresponding data are available at http://icf.som.yale.edu/stock-market-confidence-indices-explanation.

day t, we compute the conditional probabilities of a one year market return being less than -20%, denoted by $p_R\{S_T/S_t < 0.8\}$ and $p\{S_T/S_t < 0.8\}$. Then we define the left tail sentiment bias as $\log(p_R\{S_T/S_t < 0.8\}/p\{S_T/S_t < 0.8\})$, which resembles the sentiment function $\Lambda_{t,T}$ in (7).

We find that the correlation coefficient between the representative investor's left tail probability $p_R\{S_T/S_t < 0.8\}$ and CP is -0.8, and for CI is -0.6. In and of itself, there is no prior stipulation that CP need reflect investors' bias. However, an AR(2) regression of CP on $p\{S_T/S_t < 0.8\}$ and $\log(p_R\{S_T/S_t < 0.8\}/p\{S_T/S_t < 0.8\})$ has only the left tail sentiment bias term being statistically significant, with a t-statistic of -2.0.

Figure 4 (bottom panel) shows CP and $p_R\{S_T/S_t > 0.8\}$, where the latter is the probability of not having a crash under the representative investors' pdf. The comovements between the two series are evident. For example, both series reach lowest levels at the end of 2002 and 2008, i.e., periods of large market turmoil. Figure 4 (bottom panel) also suggests that the fear of a market crash, as measured by $p_R\{S_T/S_t > 0.8\}$ and CP, fell in the middle part of our sample period. This period is characterized by relatively stable market growth and low volatility, as well as high excessive optimism and overconfidence; see Figure 3.

6.5. *Corporate Bond Default Premiums*

Given that left tail events give rise to corporate bond defaults, there is reason to expect that left tail probabilities under the representative investor's distribution impact credit spreads. We now investigate whether this is the case.

In recent work, Gilchrist and Zakrajšek (2012) develop an excess bond premium measuring the component of corporate bond spreads that is not related to firm-specific information on expected defaults. Gilchrist and Zakrajšek contend that a rise in the excess bond premium represents a reduction in the effective risk-bearing capacity of the financial sector and, as a result, a contraction in the supply of credit with adverse consequences for the macroeconomy. Their credit spread index decomposes into a predictable component that captures the available firm-specific information on expected defaults and a residual component – the excess bond premium.

We find a strong positive impact of left tail probabilities under the representative investor's pdf on the excess bond premium. The correlation between the excess bond premium and $p_R\{S_T/S_t < 0.8\}$ is 0.9.[20] In addition, when regressing the excess bond premium on $p_R\{S_T/S_t < 0.8\}$ and $p\{S_T/S_t < 0.8\}$, the t-statistics are 5.25 and 1.09, respectively, with an R-squared of 82%. Although

[20]We thank Simon Gilchrist for providing us with the data for this series.

$p_R\{S_T/S_t < 0.8\}$ and $p\{S_T/S_t < 0.8\}$ are highly correlated, 0.8, the regression indicates that the representative investor's left tail probability has a large impact on the excess bond premium, and subsumes the information in the objective left tail probability. Unreported regression analysis reveals that this result holds true when controlling for excessive optimism and overconfidence. This suggests that large sentiment concentrated on left tail events could impair the intermediation capacity of the financial sector.

6.6. *Risk Aversion and Time Preference*

Meyer and Meyer (2005) survey some of the key studies by economists of how the coefficient of relative risk aversion varies across the population. Most of the survey data suggests values that lie between 0.23 and 8. Meyer and Meyer perform an adjustment to reconcile the scales used by the various studies and suggest an adjusted range of 0.8 to 4.72. We view this as a plausible range within which to evaluate our estimates.

Figure 8 displays the time series estimates of θ_0 and θ_1 when the pricing kernel is constrained to conform to the case of a representative investor with correct beliefs and CRRA utility. In standard theory, θ_0 is the discount factor for time preference and θ_1 is the coefficient of relative risk aversion. Our estimates for the time series θ_1, mostly vary between 0 and 3.1, with a mean of 1.14. During the middle of the sample period, θ_1 lies between 1 and 3.1, thereby falling in the range described in Meyer and Meyer (2005).

In our framework risk levels and risk appetites are not constant over time. At the beginning and end of the sample period, which correspond to recessions, θ_1 falls between 0 and 1, and even dips below 0 in November 2007. This finding is consistent with prospect theory (Kahneman and Tversky, 1979), which posits that after unfavorable events place agents into the domain of losses, risk aversion declines, even to the point where investors might be risk seeking.

Our finding about θ_1 declining during recessions does not imply that investors will take on greater risks. It is so because the risk-return profile and the investment opportunity set are also changing. In general investors' holding of risk is driven both by preferences and beliefs. Our analysis finds that during the financial crisis investors became pessimistic and underconfident (see Figure 3), which may well have induced investors to rebalance their portfolios from risky to less risky assets.

Formal analysis of how risk aversion covaries with excessive optimism and overconfidence reveals that although θ_1 is strongly correlated both with excessive optimism (0.68) and overconfidence

(0.59), the correlation with excessive optimism is dominant.[21] Thus during recessions risk aversion declines and investors become particularly pessimistic about future returns.

One of the papers surveyed by Meyer and Meyer is Barsky, Juster, Kimball, and Shapiro (1997), hereafter BJKS. In addition to investigating risk aversion, BJKS also conduct surveys to identify time preference. They find considerable variation, but point out at zero interest rates, the modal household expresses a preference for flat consumption over time, and the mean household expresses a preference for increasing consumption over time. BJKS report a consumption growth range of 0.28% to 1.28% per year, which they interpret as negative time discounting. Their study provides the backdrop for discussing our estimates of time preference.

The time series for the time preference variable θ_0 tends to lie between 0.99 and 1.04 during the early portion of the sample period, but in late 2004 gravitated to the range between 1.05 and 1.17, peaking in February 2005. The θ_0-series then declined back to the range 1.05 to 1.11 until November 2007, when it declined sharply to the range 0.97 to 1.1. After the Lehman bankruptcy in September 2008, θ_0 rose sharply to 1.3 in October, and then declined back to the region around 1.0 from December on. These findings are consistent with the general negative time preference pattern reported by BJKS, but clearly stronger during a portion of the sample period.

The findings above are based on the CRRA SDF, with potential time varying risk aversion and time preference, as the fundamental component of the empirical pricing kernel. In the literature one common parametric specification of the SDF is the Epstein and Zin (1989) (EZ) SDF. What EZ does, besides separating risk aversion and elasticity of substitution, is to provide a theoretical framework to replace consumption growth as the only state variable with stock return and consumption growth. How would the empirical findings above change if the CRRA SDF would be replaced by the EZ SDF? There are a number of reasons to believe that the findings above would only change marginally. First, consumption growth is notoriously smooth over time and has low correlation with market returns. Including such a smooth variable in the SDF would essentially only impact the constant part of the SDF. Indeed, the EZ log-kernel is linear in consumption growth and market return, and thus consumption growth would theoretically only impact the intercept of the pricing kernel (θ_0), not its slope (θ_1). Second, consumption data are only available at low frequency

[21] In a regression of θ_1 on its own lagged value and contemporaneous values of excessive optimism and overconfidence, the coefficient for optimism has a t-statistic of 1.9, whereas the coefficient for overconfidence has a t-statistic of 0.72; the adjusted R-squared is 81%. Moreover, the relationship is especially strong in first differences: A regression of the first difference of θ_1 on the first differences of excessive optimism and overconfidence respectively features a t-statistic of 6.3 for excessive optimism and -0.44 for overconfidence, with an associated adjusted R-squared of 31%. We interpret these results as descriptive of how our estimates of risk aversion and sentiment covary over time.

(monthly at best) and affected by well-known measurement problems. Our analysis is carried out at a relatively high frequency (weekly), and uses option and price data that are expected to readily incorporate newly available information. Third, for the most part of our sample period, estimates of the relative risk aversion (θ_1) are close to 1, with the mean of 1.16 and standard deviation 0.67. In the EZ SDF, when the coefficient of relative risk aversion is 1, consumption growth is irrelevant to the pricing kernel.

7. Conclusion

We provide empirical estimates of aggregate investors' sentiment, risk aversion, and time preference. Our estimates are extracted from empirical pricing kernels, are largely consistent with independent measures reported in the empirical literature, and yet provide several novel insights about investors' beliefs and sentiment. Our estimates indicate that investors' excessive optimism is highly correlated with the Baker and Wurgler (2006) sentiment index; overconfidence is highly correlated with the volatility predictions in the Duke/CFO survey data; tail risk is highly correlated with both the survey-based Yale/Shiller crash confidence series and the bond premium developed in Gilchrist and Zakrajšek (2012); aggregate risk aversion is in line with the general findings reported in Meyer and Meyer (2005), and varies over time in a way that is consistent with prospect theory (Kahneman and Tversky, 1979); time preference is consistent with the negative discounting, as reported by Barsky, Juster, Kimball, and Shapiro (1997).

Our main finding is that empirical pricing kernels strongly reflect behavioral elements. Our analysis provides a number of insights. For example, the Baker–Wurgler index strongly reflects investors' excessive optimism, but not overconfidence and misjudgments about tail events. The Baker–Wurgler index is thus an incomplete measure of investors' sentiment. Excessive optimism and overconfidence comove over time, i.e., investors tend to overestimate (underestimate) future returns and underestimate (overestimate) future return volatility at the same time. This generates a perceived negative risk-return trade-off, while objectively the relationship is positive. Large amount of sentiment concentrated on left tail events can impair the intermediation capacity of the financial sector with adverse consequences for the macroeconomy. Our empirical findings are robust to the choice of the pricing kernel, and remain largely intact when only minimal and necessary assumptions are imposed on the marginal utility of the representative investor.

A. Derivation of the Sentiment Function Λ

Shefrin (2008) provides a formal derivation of the sentiment function $\Lambda_{t,T}$. For completeness we briefly recall its derivation.

The sentiment function $\Lambda_{t,T}$ in (6) encapsulates the representative investor's biases. In this section, we briefly describe the structure of the sentiment function, and its manifestation within the SDF. To simplify notation, we drop t-subscripts and the argument of the pdf.

Let ξ denote state price. Then the SDF is given by $M = \xi/p$, which in a representative investor CRRA-framework has the form $\xi = p_R\, \theta_0\, (S_T/S_t)^{-\theta_1}$. This last relationship follows from the optimizing condition in which marginal rate of substitution (for expected utility) is set equal to relative state prices, with consumption at $t = 0$ serving as numeraire.

Divide both sides of the previous equation for ξ by $p\, \theta_{0,e}$, where $\theta_{0,e}$ corresponds to the value of θ_0 that would prevail if all investors held correct beliefs. Here, the subscript e denotes efficiency. This last operation leads to the expression $\xi/p = (\theta_0/\theta_{0,e})\, (p_R/p)\, \theta_{0,e}\, (S_T/S_t)^{-\theta_1}$. Define $e^{\Lambda} = (\theta_0/\theta_{0,e})\, (p_R/p)$, which is a scaled change of measure and corresponds to (7).

The change of measure (p_R/p) associated with Λ exactly specifies the transformation of the objective pdf p into the representative investor's pdf p_R. Therefore, Λ encapsulates the representative investor's biases.

Shefrin (2008) establishes that θ_1 does not vary as investors' beliefs change. Then, in the preceding expression for the SDF, e^{Λ} multiplies the term $\theta_{0,e}\, (S_T/S_t)^{-\theta_1}$, and the latter is the SDF M_e that would prevail if all investors held correct beliefs. Therefore $M = e^{\Lambda} M_e$. Taking logs, obtain $\log(M) = \Lambda + \log(M_e)$. This expression stipulates that the log-SDF can be decomposed into two components, one being the sentiment function and the other being the neoclassical log-SDF that would prevail if all investors held correct beliefs.

Rearranging the decomposition of the log-SDF yields $\Lambda = \log(M) - \log(M_e)$. Notably, the last relationship corresponds to (9) and explains why d serves as our estimate of the sentiment function Λ.

Moneyness		Maturity: Less than 60 Mean	Std.	60 to 160 Mean	Std.	More than 160 Mean	Std.
< 0.85	Put price \$	1.31	2.74	3.89	5.61	10.20	11.63
	σ^{bs}%	49.44	20.16	39.45	14.83	31.26	10.18
	Observations	11,849		10,359		10,292	
0.85–1.00	Put price \$	9.94	10.67	24.37	16.55	43.95	23.18
	σ^{bs}%	23.74	10.21	23.52	8.64	22.20	6.86
	Observations	15,876		9,510		8,152	
1.00–1.15	Call price \$	8.58	10.76	21.07	17.48	44.83	25.49
	σ^{bs}%	17.76	8.86	18.57	7.99	18.22	6.46
	Observations	13,807		9,201		7,538	
> 1.15	Call price \$	0.67	1.67	1.74	3.65	6.38	9.33
	σ^{bs}%	37.71	16.05	26.79	9.62	21.61	6.53
	Observations	7,136		7,977		9,530	

Table 1. Options dataset. For each moneyness/maturity category, entries show mean and standard deviation (Std.) of out-of-the-money call and put option prices on the S&P 500 index, as well as of Black–Scholes implied volatility (σ^{bs}) in percentage. Sample data are options observed on Wednesdays from January 2002 to October 2009. Observations are the number of options for each moneyness/maturity category. Filtering criteria of options are described in Section 5. Moneyness is strike price divided by S&P 500 index. Maturity is in calendar days.

	$\omega \times 10^6$	β	$\alpha \times 10^3$	γ	Persist.	Ann. vol.
	Objective GARCH parameters					
Mean	1.215	0.926	3.473	0.117	0.989	0.198
Std.	0.207	0.005	4.141	0.013	0.002	0.016
	Risk Neutral FHS GARCH parameters					
Mean	4.153	0.789	2.169	0.358	0.970	0.212
Std.	5.600	0.208	9.366	0.360	0.033	0.074
	Risk Neutral GAUSS GARCH parameters					
Mean	3.987	0.756	3.479	0.448	0.983	0.252
Std.	5.575	0.201	12.280	0.371	0.021	0.112

Table 2. Objective and risk neutral GARCH parameters. The GARCH model is $\log(S_u/S_{u-1}) = \mu_u + \epsilon_u$, where S_u is the S&P 500 index at day u, μ_u is the drift, and the conditional variance $\sigma_u^2 = \omega + \beta\sigma_{u-1}^2 + \alpha\epsilon_{u-1}^2 + \gamma I_{u-1}\epsilon_{u-1}^2$, where $\epsilon_u = \sigma_u z_u$, z_u is a standardized innovation and $I_{u-1} = 1$ when $\epsilon_{u-1} < 0$, and $I_{u-1} = 0$ otherwise. For each Wednesday from January 2002 to October 2009, a GARCH model is estimated using historical daily S&P 500 returns by maximizing a Pseudo Maximum Likelihood, a GARCH model driven by Gaussian innovations is calibrated to out-of-the-money options on the S&P 500 index by minimizing the sum of squared pricing errors, a GARCH model driven by filtered historical innovations is similarly calibrated to options on the S&P 500 index. Persist. is the persistency of the GARCH volatility and given by $\beta+\alpha+\gamma/2$. Ann. vol. is the annualized long-run mean of the GARCH volatility.

	Maturity					
	Less than 60		60 to 160		More than 160	
Moneyness	Mean	RMSE	Mean	RMSE	Mean	RMSE
< 0.85	0.49	1.20	0.55	1.22	0.56	1.46
0.85–1.00	0.41	1.31	−0.41	1.43	−0.95	1.70
1.00–1.15	0.44	1.29	0.03	1.19	0.42	1.55
> 1.15	0.08	0.46	0.12	0.57	0.94	1.75

Table 3. Option pricing errors. For each moneyness/maturity category, entries show mean and root mean square error (RMSE) of option price errors of the risk neutral FHS GARCH model. Price error is defined as model-based option price minus market option price. Using the FHS method, each Wednesday from January 2002 to October 2009, the GARCH model is calibrated to out-of-the-money call and put options on the S&P 500 index. Calibration procedure is described in Section 3. Filtering criteria of options are described in Section 5. Moneyness is strike price divided by S&P 500 index. Maturity is in calendar days.

Intercept	Lag1	Optimism	Overconf.	R^2
Panel A: Objective Equity Risk Premium				
0.18	0.92	−0.09	0.01	0.99
(1.98)	(7.40)	(−2.48)	(1.46)	
Panel B: Representative Investor's Equity Risk Premium				
0.79	0.71	0.07	−0.02	0.56
(4.79)	(13.26)	(3.96)	(−2.02)	
Panel C: Objective Variance Risk Premium				
−1.57	0.33	0.05	0.38	0.92
(−5.40)	(3.15)	(0.81)	(4.85)	
Panel D: Representative Investor's Variance Risk Premium				
−0.04	0.96	0.01	−0.01	0.94
(−1.03)	(5.95)	(0.81)	(−0.85)	

Table 4. Impact of Sentiment on Equity and Variance risk premiums. Panel A: Time series regression of objective equity risk premium on a constant (Intercept), its two most recent lagged value (Lag1), optimism and overconfidence; t-statistics in parentheses. Objective equity risk premium is $(E_t^p[S_T/S_t] - E_t^q[S_T/S_t]) \times 100$, where E_t^p is the conditional expectation at date t under the objective pdf p, E_t^q is the conditional expectation at date t under the risk neutral pdf q, S_t is the S&P 500 index at date t, and $(T-t)$ is one year. Optimism is $(E_t^{p_R}[S_T/S_t] - E_t^p[S_T/S_t]) \times 100$, where $E_t^{p_R}$ is the conditional expectation at date t under the representative investor's pdf p_R. Overconfidence is $(\sqrt{\mathrm{Var}_t^p[S_T/S_t]} - \sqrt{\mathrm{Var}_t^{p_R}[S_T/S_t]}) \times 100$. Panel B: Same time series regression as in Panel A for the equity risk premium perceived by the representative investor, defined as $(E_t^{p_R}[S_T/S_t] - E_t^q[S_T/S_t]) \times 100$. Panel C: Time series regression of objective variance risk premium on a constant (Intercept), its two most recent lagged value (Lag1), optimism and overconfidence. Objective variance risk premium is $(\mathrm{Var}_t^p[S_T/S_t] - \mathrm{Var}_t^q[S_T/S_t]) \times 100$. Panel D: Same time series regression as in Panel C for the variance risk premium perceived by the representative investor, defined as $(\mathrm{Var}_t^{p_R}[S_T/S_t] - \mathrm{Var}_t^q[S_T/S_t]) \times 100$. R^2 is the adjusted R-squared. Robust standard errors are computed using the Newey and West (1987) covariance matrix estimator with the number of lags optimally chosen according to Andrews (1991). Weekly observations from January 2002 to October 2009.

Intercept	Lag1	Ret	Stdv	R^2
		Panel A: Optimism		
0.70	0.67	0.64	−3.50	0.91
(4.55)	(10.55)	(2.17)	(−4.22)	
		Panel B: Overconfidence		
0.87	0.86	−0.22	−2.48	0.83
(3.93)	(10.54)	(−0.40)	(−2.95)	

Table 5. Optimism and overconfidence. Panel A: Time series regression of optimism on a constant (Intercept), its most recent lagged value (Lag1), past one year S&P 500 return (Ret) and past one year S&P 500 volatility (Stdv) namely the standard deviation of daily S&P 500 log-returns. Optimism is $(E_t^{p_R}[S_T/S_t] - E_t^p[S_T/S_t]) \times 100$, where $E_t^{p_R}$ is the conditional expectation at date t under the representative investor's pdf p_R, E_t^p is the conditional expectation at date t under the objective pdf p, S_t is the S&P 500 index at date t, and $(T-t)$ is one year. Panel B: Same time series regression for overconfidence, defined as $(\sqrt{\mathrm{Var}_t^p[S_T/S_t]} - \sqrt{\mathrm{Var}_t^{p_R}[S_T/S_t]}) \times 100$. R^2 is the adjusted R-squared. Robust standard errors are computed using the Newey and West (1987) covariance matrix estimator with the number of lags optimally chosen according to Andrews (1991). Weekly observations from January 2002 to October 2009.

			Baker–Wurgler Series				
Intercept	Lag1	Lag2	Optimism	Overconf.	S&P	VIX	R^2
0.01	0.92	−0.06	4.35	−1.00			0.93
(0.49)	(10.21)	(−0.75)	(5.32)	(−2.22)			
0.00	0.92	−0.06	4.65	−1.09	−0.06	0.03	0.93
(0.11)	(10.23)	(−0.76)	(3.75)	(−2.01)	(−0.15)	(0.16)	
0.08	0.94	−0.03			−0.27	−0.45	0.92
(3.45)	(8.86)	(−0.29)			(−0.71)	(−4.56)	

Table 6. Baker–Wurgler series and sentiment. Time series regression of monthly Baker–Wurgler series on a constant (Intercept), its two most recent lagged values (Lag1, Lag2), optimism, overconfidence, S&P 500 monthly return (S&P), and VIX index; t-statistics in parentheses. Optimism is $E_t^{p_R}[S_T/S_t] - E_t^p[S_T/S_t]$, where $E_t^{p_R}$ is the conditional expectation at date t under the representative investor's pdf p_R, E_t^p is the conditional expectation at date t under the objective pdf p, S_t is the S&P 500 index at date t, and $(T-t)$ is one year. Overconfidence is $\sqrt{\mathrm{Var}_t^p[S_T/S_t]} - \sqrt{\mathrm{Var}_t^{p_R}[S_T/S_t]}$. Robust standard errors are computed using the Newey and West (1987) covariance matrix estimator with the number of lags optimally chosen according to Andrews (1991). R^2 is the adjusted R-squared. Observations are end-of-month from January 2002 to October 2009.

	Intercept	Slope	R^2
1. Objective expected return vs. volatility	0.02	0.12	0.67
	(7.73)	(6.26)	
2. Rep. investor expected return vs. volatility	0.07	-0.13	0.50
	(15.20)	(-5.03)	

Table 7. Risk and return. Regression 1: Time series regression of objective expected market return on a constant (Intercept) and expected objective volatility (Slope); t-statistics in parentheses. Objective expected return is $E_t^p[S_T/S_t - 1]$, where E_t^p is the conditional expectation at date t under the objective pdf p, S_t is the S&P 500 index at date t, and $(T-t)$ is one year; expected objective volatility is $\sqrt{\mathrm{Var}_t^p[S_T/S_t]}$. Regression 2: Same regression as Regression 1 for expected return and volatility under the representative investor's pdf, p_R. Robust standard errors are computed using the Newey and West (1987) covariance matrix estimator with the number of lags optimally chosen according to Andrews (1991). R^2 is the adjusted R-squared. Observations are weekly from January 2002 to October 2009.

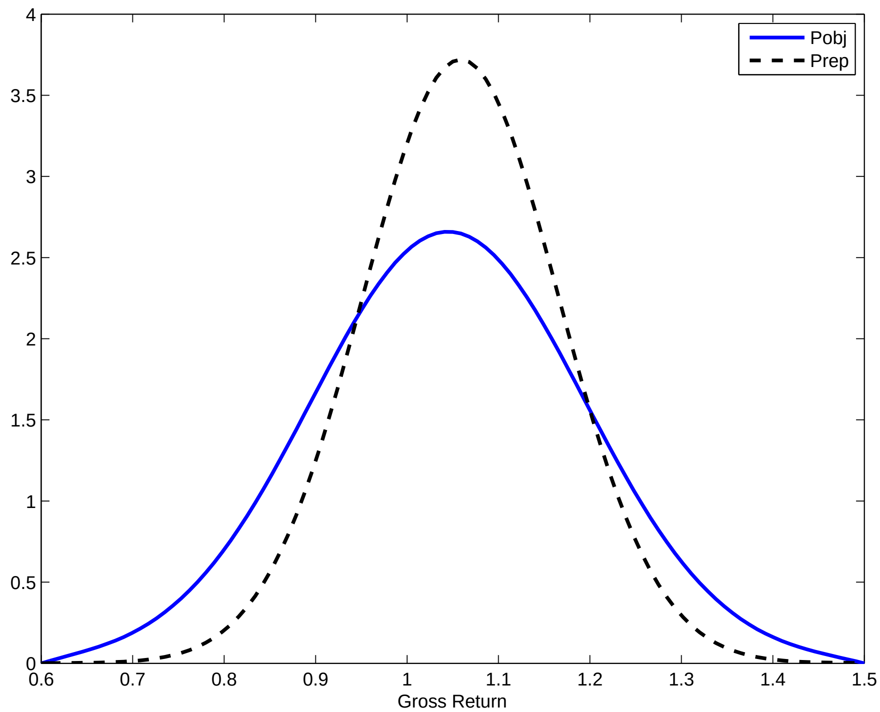

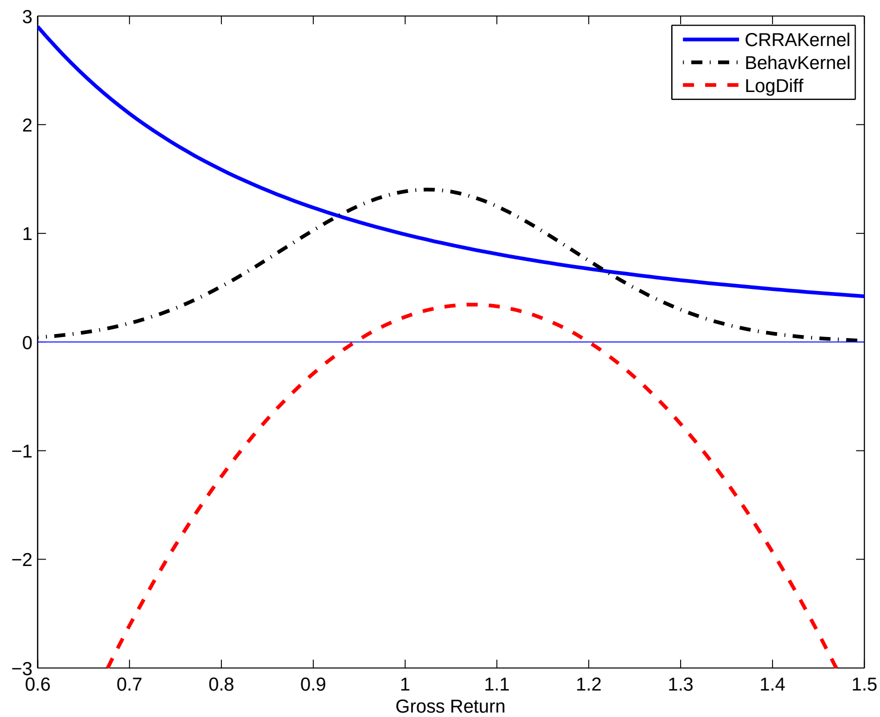

Figure 1. Upper graph: Objective (Pobj) and representative investor's (Prep) probability density functions for December 21, 2005. Lower graph: Behavioral unconstrained SDF (BehavKernel), CRRA-constrained SDF (CRRAKernel), and the LogDiff function, i.e., log BehavKernel minus log CRRAKernel. The latter difference is the function $d_{t,T}$ in (9), for December 21, 2005.

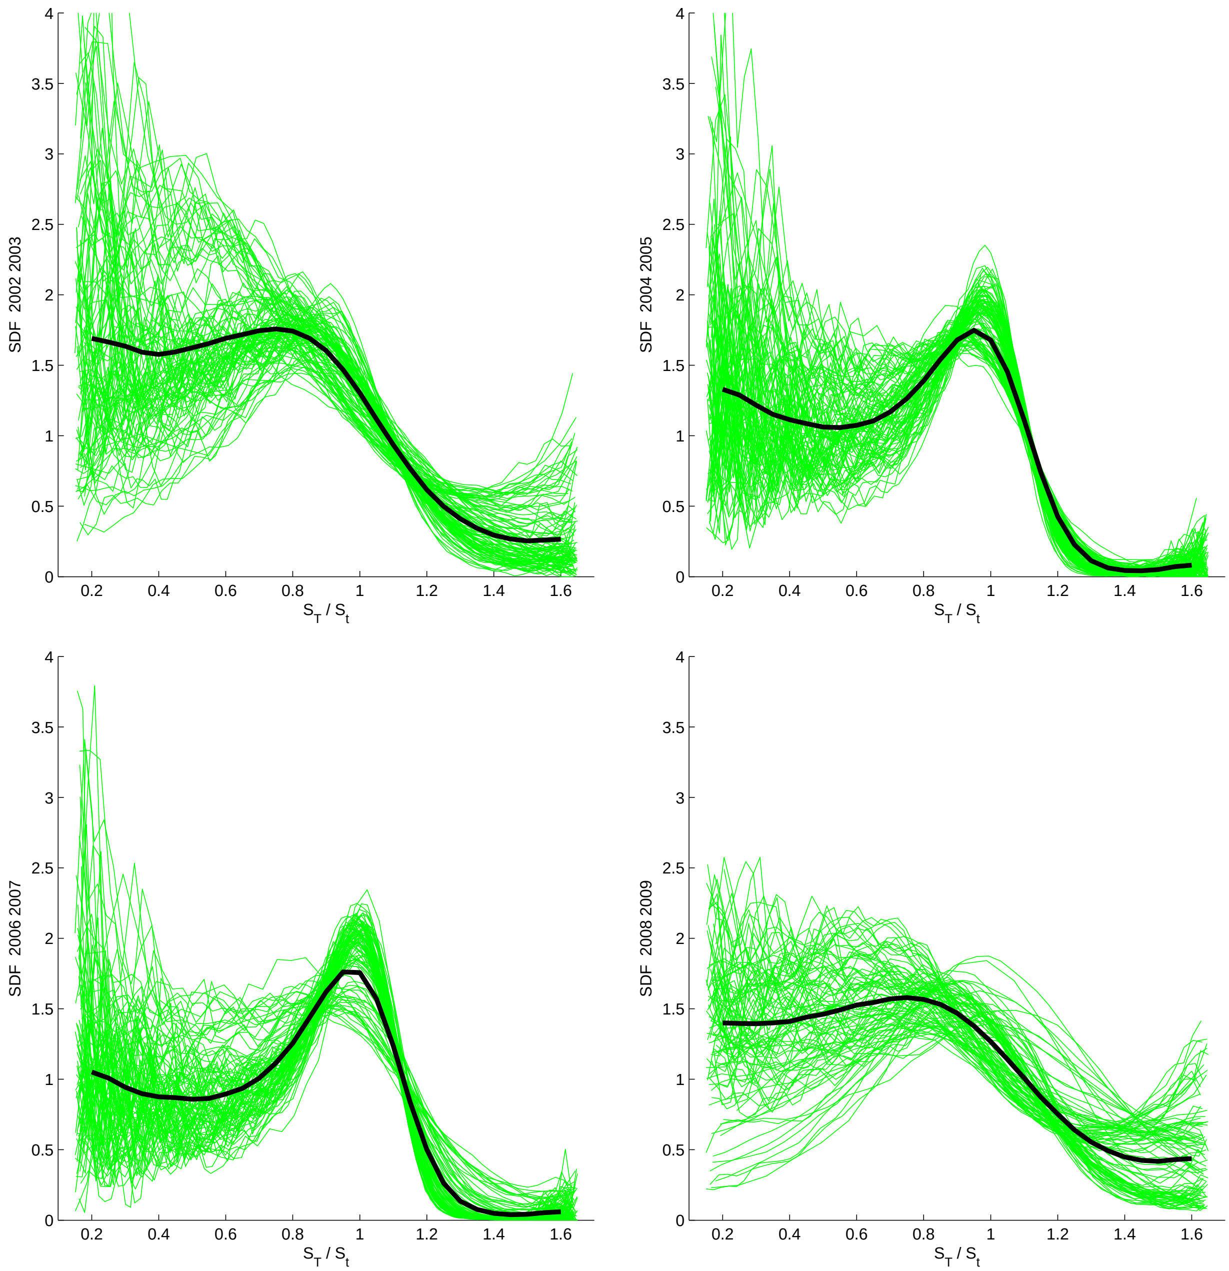

Figure 2. Empirical SDF. For each Wednesday t in our sample, the empirical stochastic discount factor (SDF), $M_{t,T}$, is estimated as $M_{t,T} = e^{-r_f\,(T-t)}\,q(S_T/S_t)/p(S_T/S_t)$, where q is the conditional risk neutral density of S_T/S_t, p the conditional objective density of S_T/S_t, r_f is the risk free rate, S_t the S&P 500 index at date t, and $(T-t)$ is one year. The densities p and q are conditional on the information available at date t and based on GARCH models with FHS innovations estimated using historical S&P 500 returns and SPX options, respectively. Each graph shows the empirical SDF over the corresponding two years period. Superimposed (solid thick line) is the average SDF.

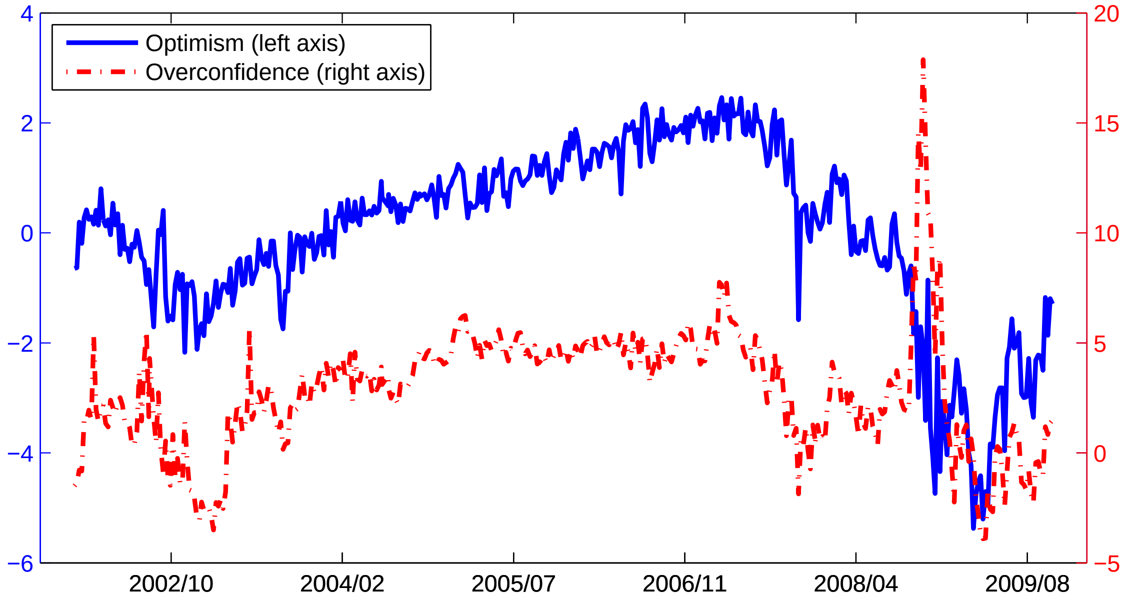

Figure 3. Time series for optimism and overconfidence. Optimism is $(E_t^{p_R}[S_T/S_t] - E_t^{p}[S_T/S_t]) \times 100$, where $E_t^{p_R}$ is the time-t conditional expectation under the representative investor's pdf p_R, S_t is the S&P 500 index at date t, $(T-t)$ is one year, and similarly for E_t^p. Overconfidence is $(\sqrt{\mathrm{Var}_t^p[S_T/S_t]} - \sqrt{\mathrm{Var}_t^{p_R}[S_T/S_t]}) \times 100$. Density estimates are obtained using the FHS method.

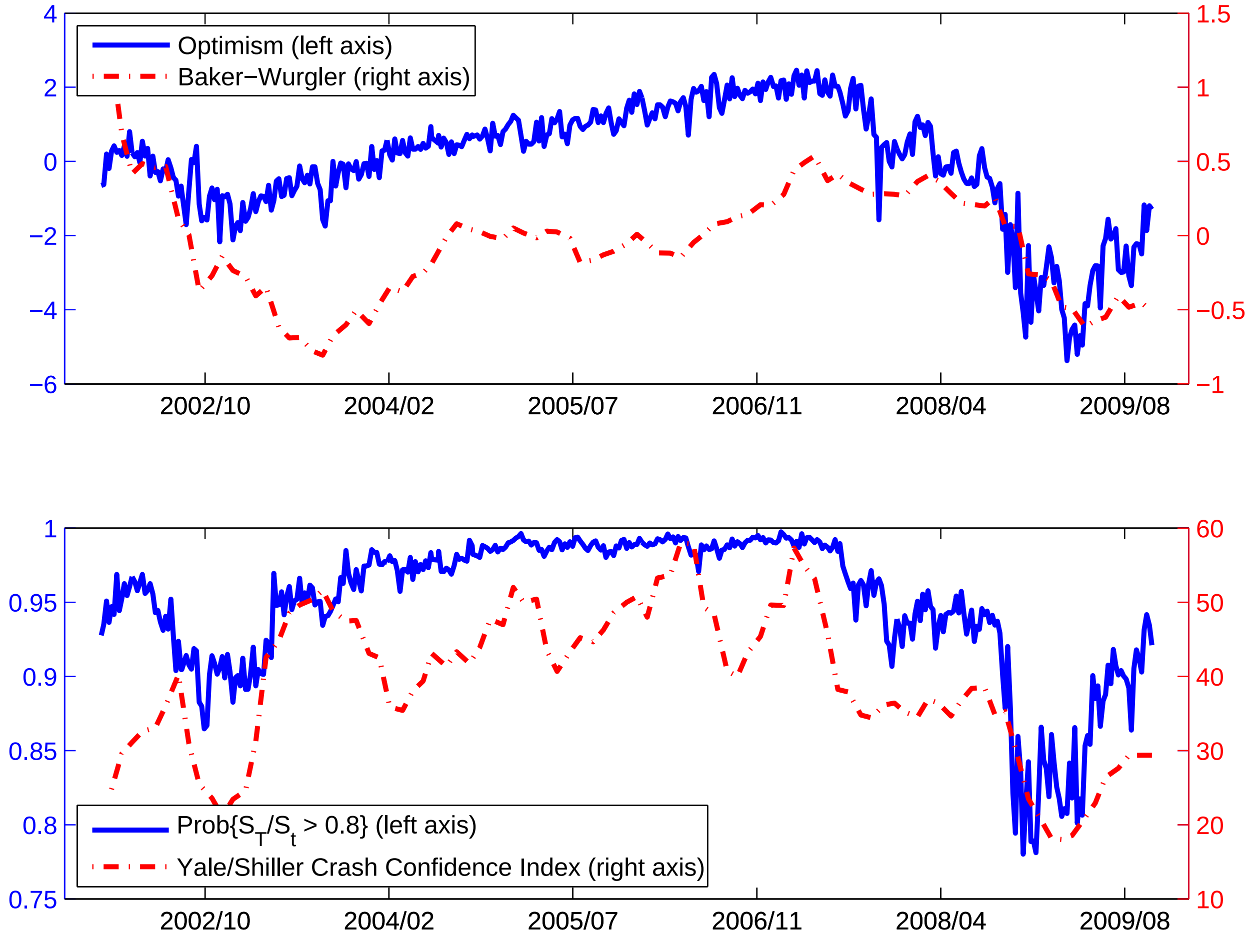

Figure 4. Upper graph: Baker–Wurgler sentiment series and optimism. Baker and Wurgler (2006) monthly series of sentiment extracted using Principal Component Analysis of six specific sentiment proxies, i.e., turnover on the New York Stock Exchange (NYSE), dividend premium, closed-end fund discount, number and first-day returns on IPOs, and the equity share in new issues. Optimism is $(E_t^{p_R}[S_T/S_t] - E_t^{p}[S_T/S_t]) \times 100$, where $E_t^{p_R}$ is the conditional expectation at date t under the representative investor's pdf p_R, S_t is the S&P 500 index at date t, $(T-t)$ is one year, and similarly E_t^{p} is the conditional expectation under the objective pdf p. Lower graph: Yale/Shiller crash confidence index (CP) and "probability of no crash" under the representative investor's pdf. The latter is $Prob\{S_T/S_t > 0.8\}$ under the representative investor's pdf p_R. For each Wednesday t, from January 2002 to October 2009, the conditional probability $Prob\{S_T/S_t > 0.8\}$ is computed numerically integrating the conditional density p_R of the gross return S_T/S_t, given the information available at date t.

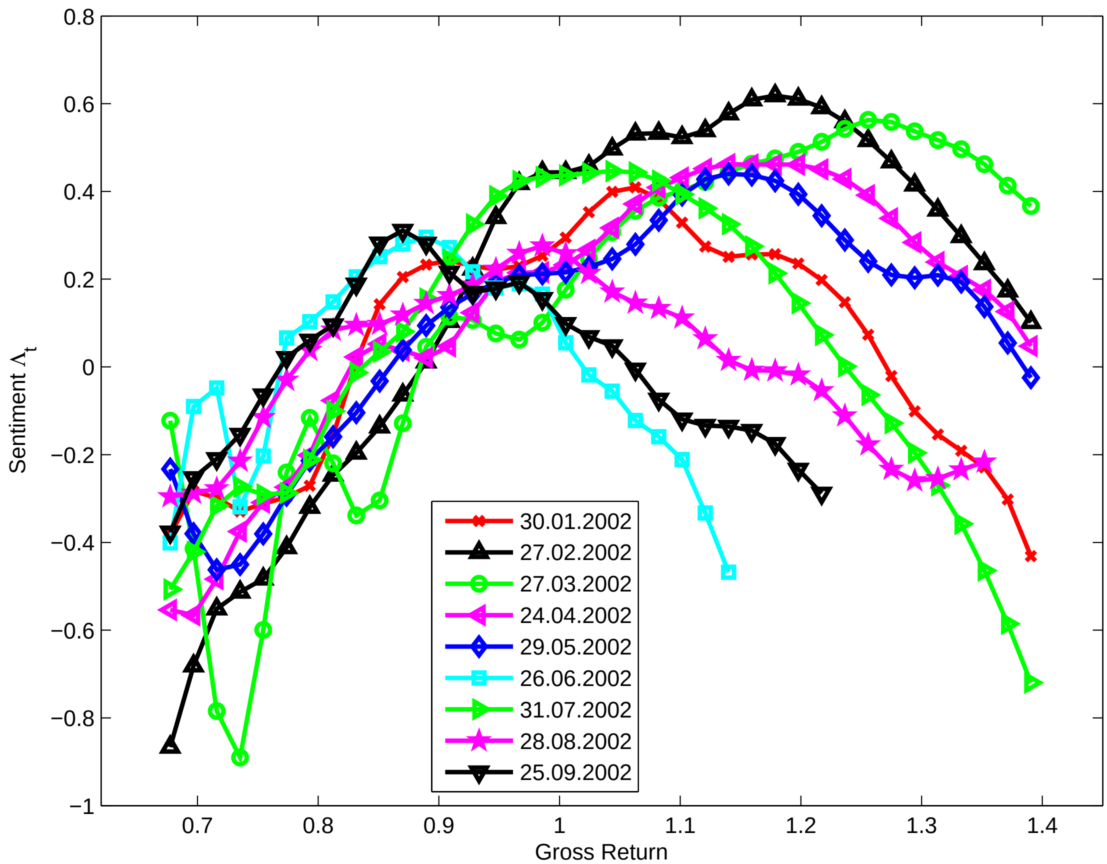

Figure 5. Sentiment functions plotted for several days in 2002. The sentiment function at date t is $\Lambda_{t,T} = \log(M_{t,T}) - \log(M_{t,T}(\theta))$, where $M_{t,T} = e^{-r_f\,(T-t)}\, q(S_T/S_t)/p(S_T/S_t)$ is the unconstrained SDF and $M_{t,T}(\theta) = \theta_{0,t}\,(S_T/S_t)^{-\theta_{1,t}}$ is the CRRA-constrained SDF. q is the conditional risk neutral density of S_T/S_t, p is the conditional objective (i.e., historical) density of S_T/S_t, r_f is the instantaneous risk free rate, $\theta_{0,t}$ is the time discount factor, $\theta_{1,t}$ is the coefficient of relative risk aversion, S_t is the S&P 500 index at date t, and $(T-t)$ is one year. On the x-axis, gross return is S_T/S_t.

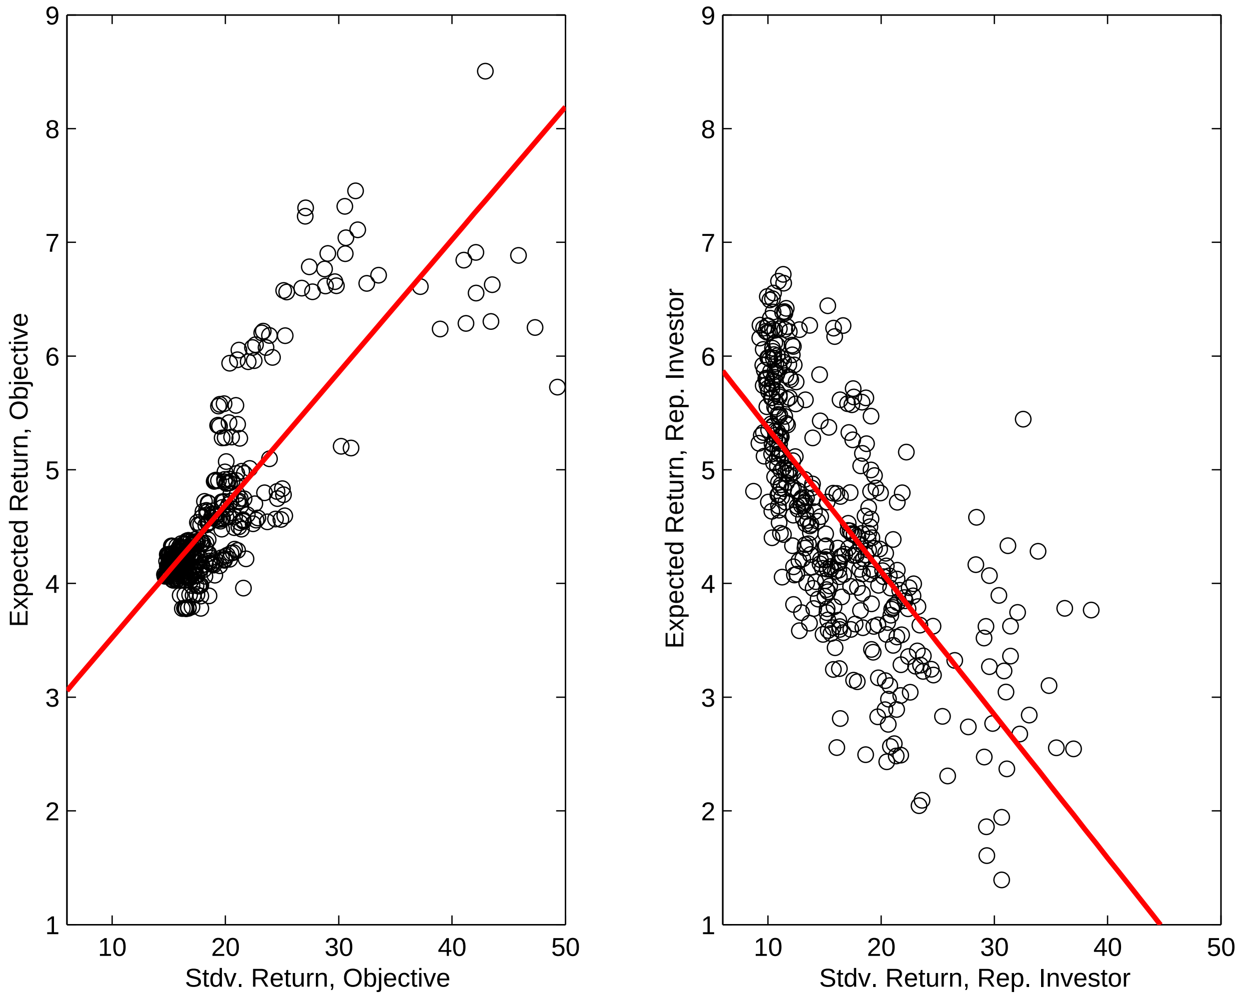

Figure 6. Risk and return. For each Wednesday t from January 2002 to October 2009, "Expected Return, Objective" is the time-t conditional expected market return under the objective pdf p, i.e., $E_t^p[S_T/S_t - 1] \times 100$, where S_t is the S&P 500 index at date t, and $(T-t)$ is one year; "Stdv. Return, Objective" is the time-t conditional expected volatility of market return under the objective pdf p, i.e., $\sqrt{\mathrm{Var}_t^p[S_T/S_t]} \times 100$. "Expected Return, Rep. Investor" and "Stdv. Return, Rep. Investor" are representative investor's expected return and volatility, respectively, computed using time-t conditional representative investor's pdf, p_R. In each graph, superimposed is the regression line.

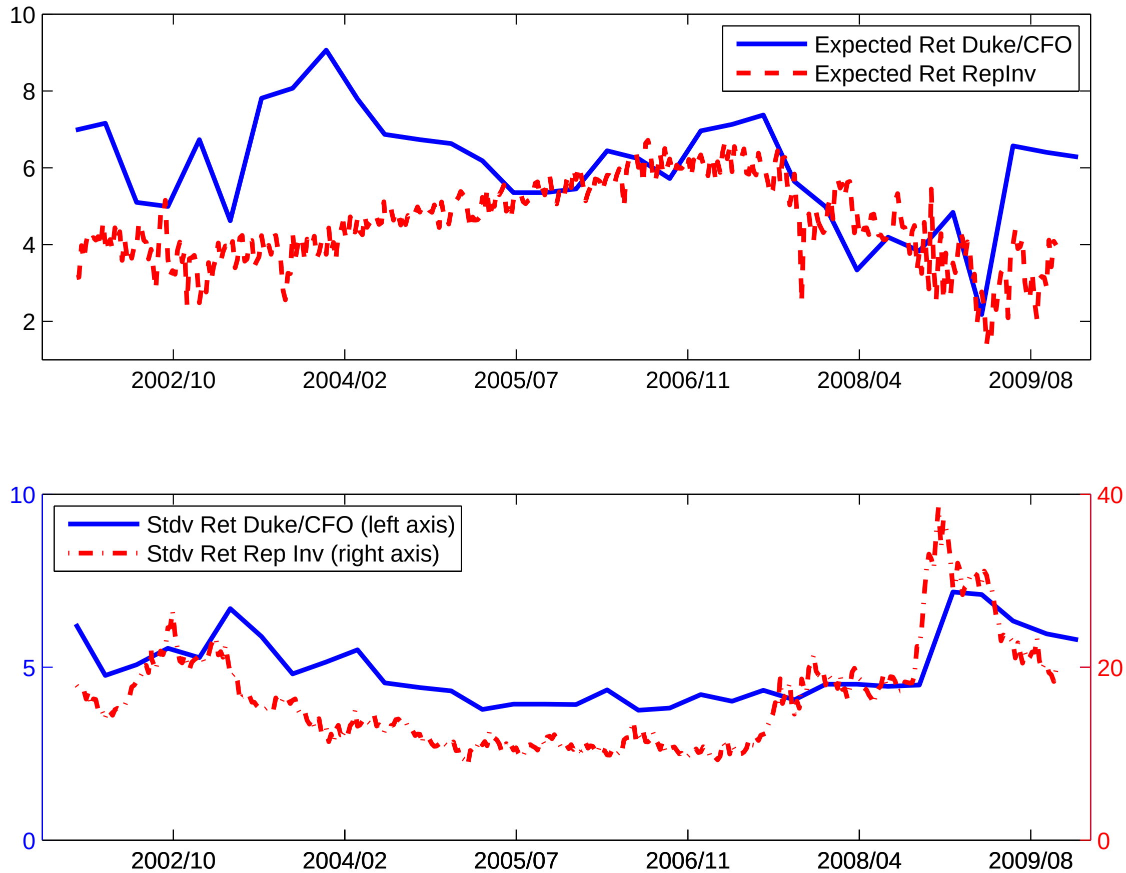

Figure 7. Upper graph: Time series of one year S&P 500 expected return based on Duke/CFO survey responses and the representative investor's distribution. Duke/CFO survey data are described in Graham and Harvey (2012), quarterly frequency. The representative investor one year S&P 500 expected return is given by $E_t^{p_R}[S_T/S_t - 1] \times 100$, where $E_t^{p_R}$ is the conditional expectation at each Wednesday t in our sample under the representative investor's pdf p_R, S_t is the S&P 500 index at date t, and $(T-t)$ is one year; weekly frequency. Lower graph: Time series of one year S&P 500 return standard deviation based on Duke/CFO survey responses and the representative investor's distribution. For each Wednesday t in our sample, return standard deviation under the representative investor's pdf is $\sqrt{\mathrm{Var}_t^{p_R}[S_T/S_t]} \times 100$.

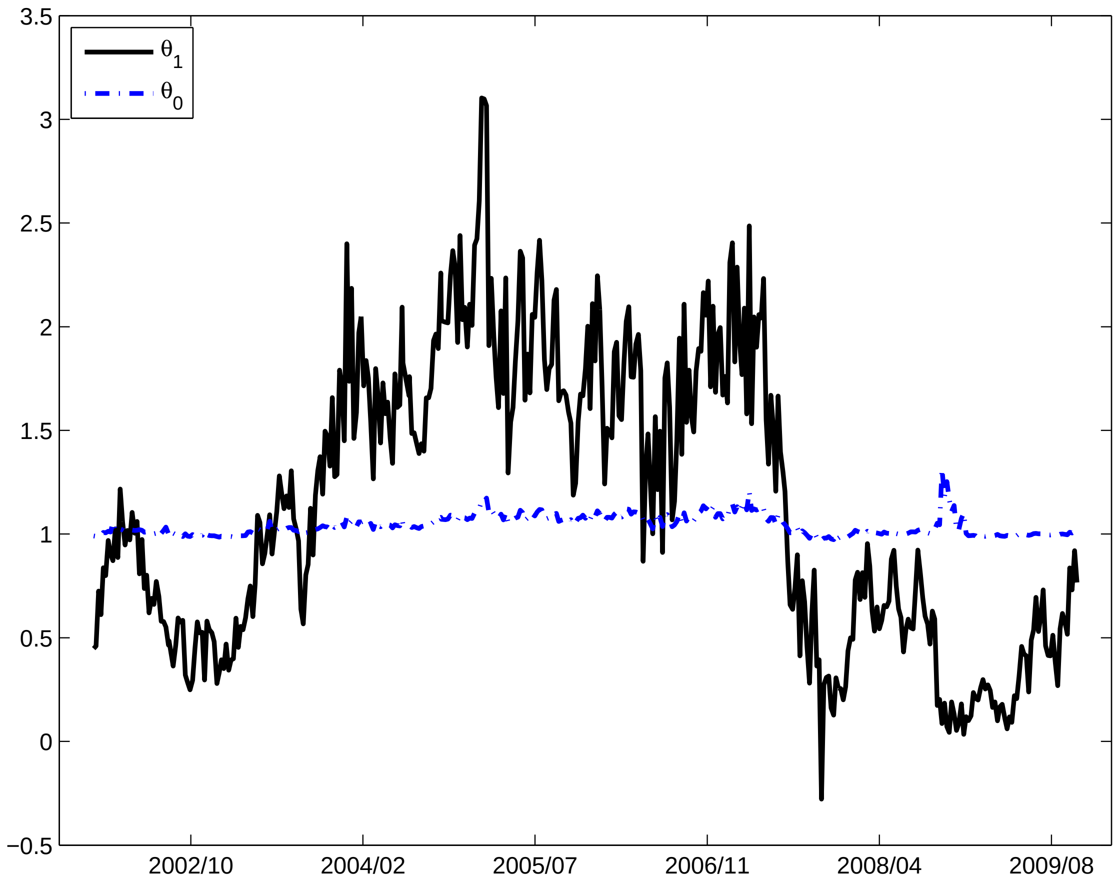

Figure 8. Time series estimates of $\theta_{1,t}$ and $\theta_{0,t}$ in the CRRA SDF. $\theta_{1,t}$ is the coefficient of relative risk aversion and $\theta_{0,t}$ is the discount factor measuring the degree of impatience at date t. The constat relative risk aversion (CRRA) SDF is $M_{t,T}(\theta) = \theta_{0,t}\,(S_T/S_t)^{-\theta_{1,t}}$, where S_t is the S&P 500 index at date t, and $(T-t)$ is one year. For each Wednesday t in our sample, $\theta_{0,t}$ and $\theta_{1,t}$ are estimated fitting the CRRA-constrained SDF, $M_{t,T}(\theta)$, to the unconstrained SDF, $M_{t,T} = e^{-r_f\,(T-t)}\,q(S_T/S_t)/p(S_T/S_t)$, where q is the conditional risk neutral density, p is the conditional objective (i.e., historical) density, S_t is the S&P 500 index at date t, and r_f is the instantaneous risk free rate.

References

Aït-Sahalia, Y., M. Karaman, and L. Mancini, 2014, "The Term Structure of Variance Swaps and Risk Premia," working paper, Princeton University and EPFL.

Aït-Sahalia, Y., and A. Lo, 2000, "Nonparametric Risk Management and Implied Risk Aversion," *Journal of Econometrics*, 94, 9–51.

Andrews, D., 1991, "Heteroskedasticity and Autocorrelation Consistent Covariance Matrix Estimation," *Econometrica*, 59, 817–858.

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

——— , 2007, "Investor Sentiment in the Stock Market," *Journal of Economic Perspectives*, 21, 129–151.

Barone-Adesi, G., R. Engle, and L. Mancini, 2008, "A GARCH Option Pricing Model in Incomplete Markets," *Review of Financial Studies*, 21, 1223–1258.

Barsky, R., F. Juster, M. Kimball, and M. Shapiro, 1997, "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Survey," *Quarterly Journal of Economics*, 107, 537–579.

Ben-David, I., J. R. Graham, and C. R. Harvey, 2013, "Managerial Miscalibration," *Quarterly Journal of Economics*, 128, 1547–1584.

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

——— , 2015, "Do Individual Behavioral Biases Affect Financial Markets and the Macroeconomy," working paper, Imperial College London.

Birru, J., and S. Figlewski, 2012, "Anatomy of a Meltdown: The Risk Neutral Density for the S&P 500 in the Fall of 2008," *Journal of Financial Markets*, 15, 151–180.

Bollerslev, T., G. Tauchen, and H. Zhou, 2009, "Expected Stock Returns and Variance Risk Premia," *Review of Financial Studies*, 22, 4463–4492.

Bollerslev, T., and V. Todorov, 2011, "Tails, Fears and Risk Premia," *Journal of Finance*, 66, 2165–2211.

Borovička, J., L. P. Hansen, and J. Scheinkman, 2015, "Misspecified Recovery," working paper, New York University, University of Chicago and Princeton University.

Brandt, M., and Q. Kang, 2004, "On the Relationship Between the Conditional Mean and Volatility of Stock Returns: A Latent VAR Approach," *Journal of Financial Economics*, 72, 217–257.

Buraschi, A., F. Trojani, and A. Vedolin, 2014, "When Uncertainty Blows in the Orchard: Comovement and Equilibrium Volatility Risk Premia," *Journal of Finance*, 69, 101–137.

Campbell, J. Y., and R. J. Shiller, 1998, "Valuation Ratios and the Long Run Market Outlook," *Journal of Portfolio Management*, 24, 11–26.

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

Cochrane, J., 2005, *Asset Pricing*. Princeton University Press, Princeton, Second edition.

Conrad, J., R. F. Dittmar, and E. Ghysels, 2013, "Ex Ante Skewness and Expected Stock Returns," *Journal of Finance*, 68, 85–124.

Drechsler, I., and A. Yaron, 2011, "What's Vol Got to Do With It," *Review of Financial Studies*, 24, 1–45.

Duan, J.-C., and J. Simonato, 1998, "Empirical Martingale Simulation for Asset Prices," *Management Science*, 44, 1218–1233.

Dumas, B., A. Kurshev, and R. Uppal, 2009, "Equilibrium Portfolio Strategies in the Presence of Sentiment Risk and Excess Volatility," *Journal of Finance*, 64, 579–629.

Dybvig, P. H., 1988, "Distributional Analysis of Portfolio Choice," *Journal of Business*, 63, 369–393.

Epstein, L. G., and S. E. Zin, 1989, "Substitution, Risk Aversion, and the Temporal Behavior of Consumption Growth and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937–969.

Filipović, D., E. Gourier, and L. Mancini, 2015, "Quadratic Variance Swap Models," *Journal of Financial Economics*, forthcoming.

Ghysels, E., P. Santa-Clara, and R. Valkanov, 2005, "There is a Risk-Return Trade-off After All," *Journal of Financial Economics*, 76, 509–548.

Gilchrist, S., and E. Zakrajšek, 2012, "Credit Spreads and Business Cycle Fluctuations," *American Economic Review*, 102, 1692–1720.

Glosten, L., R. Jagannathan, and D. Runkle, 1993, "On the Relation Between the Expected Value and the Volatility of the Nominal Excess Return on Stocks," *Journal of Finance*, 48, 1779–1801.

Goyal, A., and P. Santa-Clara, 2003, "Idiosyncratic Risk Matters!," *Journal of Finance*, 58, 975–1007.

Graham, J. R., and C. R. Harvey, 2012, "The Equity Risk Premium in 2012," working paper, Duke University.

Guo, H., and R. Whitelaw, 2006, "Uncovering the Risk-Return Relation in the Stock Market," *Journal of Finance*, 61, 1433–1463.

Harvey, C. R., 1989, "Time-Varying Conditional Covariances in Tests of Asset Pricing Models," *Journal of Financial Economics*, 24, 289–317.

——— , 2001, "The Specification of Conditional Expectations," *Journal of Empirical Finance*, 8, 573–637.

Jackwerth, J. C., 2000, "Recovering Risk Aversion from Options Prices and Realized Returns," *Review of Financial Studies*, 13, 433–451.

Jiang, G. J., and Y. S. Tian, 2005, "The Model-Free Implied Volatility and its Information Content," *Review of Financial Studies*, 18, 1305–1342.

Jouini, E., and C. Napp, 2006, “Heterogeneous Beliefs and Asset Pricing in Discrete Time,” *Journal of Economic Dynamics and Control*, 30, 1233–1260.

——— , 2007, “Consensus Consumer and Intertemporal Asset Pricing with Heterogeneous Beliefs,” *Review of Economic Studies*, 74, 1149–1174.

Kahneman, D., and A. Tversky, 1979, “Prospect Theory: An Analysis of Decision Making Under Risk,” *Econometrica*, 47, 263–291.

Ludvigson, S., and S. Ng, 2007, “The Empirical Risk-Return Relation: A Factor Analysis Approach,” *Journal of Financial Economics*, 83, 171–222.

Marfè, R., 2015, “Survey Expectations and the Equilibrium Risk-Return Trade Off,” working paper, Swiss Finance Institute.

Meyer, D., and J. Meyer, 2005, “Relative Risk Aversion: What Do We Know?,” *Journal of Risk and Uncertainty*, 31, 243–262.

Mueller, P., A. Vedolin, and Y.-M. Yen, 2011, “Bond Variance Risk Premia,” working paper, London School of Economics.

Nelson, D., 1991, “Conditional Heteroskedasticity in Asset Returns: A New Approach,” *Econometrics*, 59, 347–370.

Newey, W. K., and K. D. West, 1987, “A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix,” *Econometrica*, 55, 703–708.

Rosenberg, J., and R. Engle, 2002, “Empirical Pricing Kernels,” *Journal of Financial Economics*, 64, 341–372.

Ross, S., 2015, “The Recovery Theorem,” *Journal of Finance*, 70, 615–648.

Rossi, A., and A. Timmermann, 2015, “Modeling Covariance Risk in Merton’s ICAPM,” *Review of Financial Studies*, 28, 1428–1461.

Shefrin, H., 2005, *A Behavioral Approach to Asset Pricing*. Elsevier Academic Press, Boston, First edition.

——— , 2008, *A Behavioral Approach to Asset Pricing*. Elsevier Academic Press, Boston, Second edition.

Todorov, V., 2010, “Variance Risk Premium Dynamics: The Role of Jumps,” *Review of Financial Studies*, 23, 345–383.

White, H., 1982, “Maximum Likelihood Estimation of Misspecified Models,” *Econometrica*, 50, 1–25.

Xiong, W., and H. Yan, 2010, “Heterogeneous Expectations and Bond Markets,” *Review of Financial Studies*, 23, 1433–1466.

Yu, J., and Y. Yuan, 2011, “Investor Sentiment and the Mean-Variance Relation,” *Journal of Financial Economics*, 100, 367–381.