Systemic Risk and Sentiment
Chapter Contribution to Handbook on Systemic Risk*
Edited by Jean-Pierre Fouque and Joe Langsam

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

This version: October 30, 2011

Abstract

Regulators charged with monitoring systemic risk need to focus on sentiment as well as narrowly defined measures of systemic risk. This chapter describes techniques for jointly monitoring the co-evolution of sentiment and systemic risk. To measure systemic risk, we use Marginal Expected Shortfall. To measure sentiment, we apply a behavioral extension of traditional pricing kernel theory, which we supplement with external proxies. We illustrate the technique by analyzing the dynamics of sentiment before, during, and after the global financial crisis which erupted in September 2008. Using stock and options data for the S&P 500 during the period 2002–2009, our analysis documents the statistical relationship between sentiment and systemic risk.

*Financial support from the Swiss National Science Foundation NCCR-FinRisk (Barone-Adesi and Mancini) is gratefully acknowledged. Shefrin ackowledges a course release grant from Santa Clara University. We thank Rob Engle for his insightful comments during the development of this chapter.

[†]Giovanni Barone-Adesi, Faculty of Economics, University of Lugano, Via Buffi 13, CH-6900 Lugano, Switzerland. E-mail: `giovanni.baroneadesi@usi.ch`

[‡]Loriano Mancini, Swiss Finance Institute at EPFL, Quartier UNIL-Dorigny, Extranef 217, CH-1015 Lausanne, Switzerland. E-mail: `loriano.mancini@epfl.ch`

[§]**Corresponding author**: Hersh Shefrin, Leavey School of Business, Santa Clara University, Santa Clara, California, CA 95053. E-mail: `hshefrin@scu.edu`

1 Introduction

The report of the Financial Crisis Inquiry Commission (FCIC, 2011) emphasizes the importance of systemic risk and sentiment. These two concepts, and the relationship between them, are important for regulatory bodies such as the Financial Stability Oversight Council (FSOC) who, with the support of the Office of Financial Research (OFR), is charged with the responsibility for monitoring systemic risk throughout the financial system. This chapter describes tools regulators can use to monitor sentiment and its impact on systemic risk.

To measure systemic risk we use Marginal Expected Shortfall (MES), defined for a firm as the expected equity loss per dollar conditional on the occurrence of a systemic event. See [Acharya et al. (2010)]. An example of a systemic event is a decline in the value of the market portfolio on a given day by 2% or more. Values for MES are reported at New York University's Volatility Lab website. These values are computed using the methodology developed in [Brownlees and Engle (2010)], where MES is computed as a function of volatility, correlation with the market return, and tail expectations of the standardized innovations distribution.

To measure sentiment we employ several sources, some derived from market data and some based on survey evidence. Much of the chapter describes a technique for estimating sentiment from market prices, which was developed by the authors of this chapter in [Barone-Adesi, Mancini, and Shefrin (2011)]. Our discussion deals with estimates of optimism and overconfidence, Campbell-Shiller P/E, the crash confidence indices created by economist Robert Shiller and managed by Yale University, and a variable developed by economists Malcolm Baker and Jeffrey Wurgler, which is based on variables studied in the behavioral finance literature.

Consider what the FCIC states about systemic risk and sentiment. The FCIC describes systemic risk as "a precipitous drop in asset prices, resulting in collateral calls and reduced liquidity." (p. 334) In its report, the FCIC criticized regulators for viewing "the institutions they oversaw as safe and sound even in the face of mounting troubles," and concluded that "dramatic failures of corporate governance and risk management at many systemically important financial institutions were a key cause of this crisis." (p. xviii)

Notably, the members of the FCIC were not unanimous in their conclusions about the key cause of the financial crisis. A dissenting minority pointed to "U.S. government housing policy, which led to the creation of 27 million subprime and other risky loans," noting that "[if] the U.S. government had not chosen this policy path ... the great financial crisis of 2008 would never have occurred." (p. 444) What is important about this perspective for regulators is not so much whether it is true, but whether regulatory measures are able to prevent such policies from leading to a financial crisis.

What is especially important about the minority position's perspective was its emphasis on sentiment. That position highlighted the role of private mortgage-backed securities (PMBS) issued by financial firms such as Countrywide Financial Corp., at the time the nation's largest mortgage lender. The report states that "PMBS, however, are far more vulnerable to swings in sentiment than whole mortgages held on bank balance sheets." (p. 476) In this regard, PMBS comprised about one third of issuance of non-traditional mortgage-backed securities.

As an example, the minority position pointed out that investment banks such as Bear Stearns relied heavily on AAA-rated PMBS as collateral for raising short-term financing through repurchase agreements (repos). PMBS comprised roughly one third of Bear Stearns's collateral. Once concerns about mortgage defaults began to rise, sentiment about PMBS quickly became negative, and investment banks lost much of their ability to borrow short-term. The FCIC staff "showed that the loss of the PMBS market was the single event that was crippling for Bear, because it eliminated a major portion of the firm's liquidity pool ..." (p. 477-8)

The previous discussion about Bear Stearns makes the point that sentiment played a critical role as the financial crisis unfolded. As it turned out, the risks Bear Stearns faced were not unique but systemic. Of course, sentiment did not only materialize when Bear Stearns lost its ability to borrow short-term. Rather, sentiment played a key role in fostering the climate in which systemic risk grew. The unsound mortgage lending practices highlighted by the FCIC accelerated in the period 2003–2006. For example, loan-to-value ratios averaged 80.5% in 2002, but climbed to 89.1% in 2006. The combination of limited documentation and 100% financing climbed from 1% of all mortgages in 2002 to 15% in 2006.

For the purpose of illustration, consider a thumbnail sketch of what our analysis indicates about the co-evolution of sentiment and systemic risk as the financial crisis unfolded. As we shall see, during the deleveraging phase of the financial crisis, leverage and systemic risk increased dramatically. Notably, both were highly negatively correlated with sentiment, suggesting that negative shifts in sentiment exacerbated the impact of declining fundamentals.

Between January 2004 and February 2007, the values of optimism and overconfidence implied by our model both rose. The Yale/Shiller crash confidence index, which measures the confidence investors have that a crash is not imminent, trended upward. The Baker–Wurgler sentiment index, a proxy for optimism, rose. The Campbell-Shiller P/E, based on long-term average earnings, stabilized around 25, a level historically associated with major market declines. The risk premium gradually declined. Housing prices continued to climb at about 10% per year, and peaked in early 2006. Defaults on subprime mortgages increased in the first quarter of 2007. Notably, both optimism and overconfidence were strongly correlated with the level of housing prices.

Between March and September 2007 overconfidence began to decline, but optimism continued to increase. This finding is especially interesting because the FCIC uses the term "madness" to characterize investment decisions during the first part of 2007. They do so because despite housing prices having peaked in early 2006 and subprime mortgage default rates beginning to rise, some financial firms such as Citigroup, Merrill Lynch, and UBS continued to increase their exposure to subprime mortgages. The crash confidence index peaked at the end of February 2007 and subsequently declined sharply. A run on commercial paper took place between August and December of 2007, as bad news about defaults on subprime mortgages intensified. During this time, optimism and overconfidence both declined, temporarily dipping below zero in November.

At the end of September 2007, there was considerable variation in leverage across financial firms, where leverage is defined as assets over equity, measured in terms of market value. AIG's leverage was 6. Citigroup's leverage was 10, as was that of commercial bank Washington Mutual. Goldman Sachs's leverage was 11. Fannie Mae's leverage was 15 and Freddie Mac's leverage was 21. Merrill Lynch and Morgan Stanley were both at 18. However, leverage levels for the two major investment banks who were first to fail were higher: Lehman Brothers's leverage was 21 and Bear Stearns's leverage was 28.

The period between late 2007 and the Lehman Brothers bankruptcy in September 2008 featured considerable volatility. Systemic risk soared in March 2008 for both Bear Stearns and Lehman Brothers. Bear Stearns almost failed in March, but for an emergency loan provided by the Federal Reserve Bank of New York, and Bear Stearns agreeing to be acquired by JP Morgan Chase. Bear Stearns's March MES increased from about 5% in February 2008 to 47% in March. Interestingly, its February value for MES was higher than for some banks, such as Goldman Sachs and JP Morgan whose MES values were in the 3–4% range, but not as high as Lehman Brothers and Merrill Lynch whose values exceeded 5%. However, its March value was the highest by far: Lehman Brothers's MES came in second at 14%.

Bear Stearns's leverage peaked at 254 in March 2008. By this time Fannie Mae's leverage had increased to 32 and Freddie Mac's leverage had increased to 49. Lehman Brothers's leverage stood at 39. After the acquisition of Bear Stearns by JP Morgan Chase was completed in May 2008, leverage and MES for financial institutions began to rise, as optimism and overconfidence fell sharply. At the end of June, Fannie Mae's leverage had increased to 45. Freddie Mac's leverage had increased to 83. Lehman Brothers's leverage had increased to 46. As for MES, by July the MES for both Fannie Mae and Freddie Mac had increased to 15%, from their May values of 5%. Lehman Brothers's MES revisited its March value of 14%, after having declined to 8% in May.

Optimism plummeted after the bankruptcy of Lehman Brothers and rescue of AIG, and became negative, which is to say that the market became excessively pessimistic. Lehman Brothers's MES rose above 12% in the two months before its bankruptcy in September 2008. Within days of the Lehman bankruptcy, insurance firm AIG required a government bailout to survive, because of credit default swap collateral calls triggered by a downgrade to its credit rating. AIG's September MES increased dramatically to 25% from already high values of 10% in July and August. During the same period, Fannie Mae's MES and Freddie Mac's MES both rose above 14% prior to their being taken into conservatorship by the U.S. government. At this time, their leverage had risen to 116 and 298 respectively. Lehman Brothers's leverage stood at 56 at the time it declared bankruptcy. Leverage for Washington Mutual, which failed at the same time, rose to 42. Leverage levels for AIG, Citigroup, Merrill Lynch, and Morgan Stanley, at this point all firms in danger of failing, had risen above 20.

Sentiment played a prominent part as the degree of systemic risk became apparent during the market decline that followed the Lehman Brothers bankruptcy. Between the Lehman bankruptcy in September 2008 and the market bottom in March 2009, optimism and overconfidence declined dramatically, and became quite negative. Optimism bottomed at -5.4%. Overconfidence bottomed at -3.4%. Both optimism and overconfidence subsequently rose after April, although optimism remained negative.

As the financial crisis began to unfold in 2007, for many financial firms MES was negatively correlated with excessive optimism and to a lesser degree with overconfidence. Financial firm leverage was also negatively correlated with excessive optimism. Although we document these relationships in hindsight, the exercise suggests to us that tracking sentiment is an important task for regulators as they monitor systemic risk. In particular, the correlations are stronger for some firms than for others.

The remainder of the chapter is divided into five sections, plus a conclusion. Section 2 is theoretical, and focuses on behavioral asset pricing theory, centered on the concept of a pricing kernel or stochastic discount factor (SDF). Notably, in behavioral asset pricing theory, the SDF can be decomposed into a fundamental component and sentiment. Section 3 is empirical, and focuses on estimating the SDF, or at least its projection onto the S&P 500. We note that SDF-based theory is very general, applying to a multitude of assets such as stocks, bonds, options, and physical assets such as housing. In this chapter, we focus on the U.S. equity market. Section 4 uses the empirical SDF to impute the time series for sentiment between 2002 and 2009. Section 5 describes external measures of sentiment. Section 6 analyzes the variation in systemic risk and leverage as sentiment shifted over time. Section 7 contains concluding comments.

2 Behavioral Asset Pricing Theory and Sentiment

2.1 Sentiment

In SDF-based theories, the price P of an asset with random payoff X is the expected value of its discounted payoff, where the SDF M is the discount factor used to capture the effects of both time value of money and risk. In the equation $P = E(MX)$, both M and X are random variables. That is, the discount factor M typically varies across payoff levels in order to reflect that risk is priced differently across payoff levels.

Sentiment pertains to erroneous beliefs. In this regard, think of X as having a probability density function (pdf) p which is objectively correct, but about which individual investors only possess subjective beliefs. The beliefs of an investor whose subjective beliefs are correct are said to feature zero sentiment. The beliefs of an investor whose subjective beliefs are incorrect are said to feature nonzero sentiment.

In the neoclassical SDF framework, investors' beliefs refer to an underlying state variable such as aggregate consumption growth. In this case, the sentiment of an individual investor i can be described as the "difference" between two probability density functions: the objective pdf p and the individual investor's subjective pdf p_i. We measure this difference as $ln(p_i/p)$, which is technically a log-change of measure. The log-change of measure is a function of the underlying state variable. It specifies the percentage error in probability density which investor i assigns to the occurrence of a specific value for consumption growth. For example, suppose that investor i underestimates by 2% the probability that consumption growth will be 1%. In this case, i's log-change of measure at 1% will be -2%.

Because the log-change of measure completely captures the error in an investor's beliefs, we use it as our measure of sentiment. The shape of the log-change of measure function captures many of the essential characteristics of sentiment. In a Gaussian framework, a log-linear change of measure generates a variance preserving shift in mean (with the form $x\mu - \frac{1}{2}\mu^2$). If the mean shifts to the right by μ, 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. A log-change of measure with an inverted U-shape shifts probability mass from the tails into the center, thereby lowering the variance. To put it another way, a U-shape gives rise to underconfidence, whereas an inverted U-shape gives rise to overconfidence.

2.2 CRRA: Equilibrium Aggregation

A central issue in behavioral asset pricing theory is how in equilibrium the market aggregates the probability density functions of the individual investors to arrive at a probability density function for the market as a whole. It is this "market pdf" or representative investor's pdf which underlies the pricing equation $P = E(MX)$. And of course, the same notion of sentiment that applies to the probability density functions of individual investors also applies to the market pdf.

[Shefrin (2008)] develops a behavioral extension of the CRRA-based pricing kernel. The equation for the CRRA pricing kernel traditionally used to estimate the empirical pricing kernel has the form

$$M_{t,t+\tau}(\theta) = \theta_0 (S_{t+\tau}/S_t)^{-\theta_1} \tag{1}$$

In (1), M is the pricing kernel, t and τ are indexes for time, S is a proxy for the value of the market portfolio, θ_0 is a discount factor measuring the degree of impatience and θ_1 is the coefficient of relative risk aversion, and $\theta = (\theta_0, \theta_1)$. In empirical analysis, $S_{t+\tau}/S_t$ plays the role of aggregate consumption growth.

In the behavioral framework, every individual investor's utility function conforms to CRRA, with investors' coefficients of CRRA allowed to vary from investor to investor. Notably, the behavioral extension allows for inter-investor variation in respect to time preference parameters and most importantly, beliefs about the stochastic process p governing the evolution of S_t. Therefore, if investors are indexed by i, then heterogeneity implies that θ will be indexed by i, as is already the case for beliefs p_i.

The heterogeneity in respect to beliefs, risk attitude, and time preference is sufficient to accommodate a wide spectrum of psychological features that feature in the behavioral finance literature. Examples include biases such as excessive optimism and overconfidence, prospect theoretic features such as probability weighting and asymmetric treatment of gains and losses, and a nonexponential time preference extension to accommodate hyperbolic discounting.

The logarithmic version of (1) is

$$ln(M_{t,t+\tau}(\theta)) = ln(\theta_0) - \theta_1 ln(S_{t+\tau}/S_t) \tag{2}$$

In the behavioral pricing theory extension, (2) generalizes to include an additional sentiment term for the market, Λ. The generalization of the log-pricing kernel (2) has the following form:

$$ln(M_{t,t+\tau}(\theta)) = \Lambda(S_{t+\tau}/S_t) + ln(\theta_{0,t}) - \theta_{1,t} ln(S_{t+\tau}/S_t) \tag{3}$$

Notice that θ is now indexed by t and so becomes time varying. This time variation is generated by wealth transfers among investors that result from trading.

The function e^{Λ} is the product of a change of measure in respect to p and a proportional rescaling of θ_0. Formally,

$$p_R = pe^{\Lambda}\theta_{0,t,p}/\theta_{0,t} \tag{4}$$

Here $\theta_{0,t,p}$ connotes the time preference variable that would apply were all investors to hold correct beliefs p, whereas $\theta_{0,t}$ is the time preference variable of the representative investor. The change of measure, when applied to p, gives rise to the beliefs p_R of a representative investor R. Likewise the time preference rescaling term, when applied to the objective time preference variable, gives rise to the time preference variable of the representative investor. As for the exponent term $\theta_{1,t}$, it is determined as a consumption-weighted harmonic mean of the individual investor coefficients ($\theta_{1,i}$).

By representative investor, we mean that equilibrium prices are determined as if there were a single investor in the market, or equivalently as if all investors were homogeneous and identical to the representative investor. In this sense, the representative investor aggregates the risk tolerance, time preference, and beliefs of the individual investors into a composite. Time variation in the representative investor's features occurs as a result of wealth transfers over the course of trading, in that investor wealth is a key contribution to the weight with which the representative investor reflects the traits of a specific individual investor in the population. Notably, when all investors hold correct beliefs, then so too will the representative investor, meaning that $p_R = p$.

[Shefrin (2008)] points out that equation (3) stipulates that the log-pricing kernel can be interpreted as the sum of a sentiment function and a function that corresponds to what the log-pricing kernel would be if investors all held correct beliefs ($p_i = p$ for all i). Restated, the equilibrium log-pricing kernel is the sum of sentiment and the fundamental log-pricing kernel. In our empirical analysis, we use this equation to solve for the sentiment function, by taking the difference between the empirical SDF and an estimate of the sum of the second and third terms in (3).

Imbedded within the sentiment function is the log-change of measure $ln(p_R/p)$. Suppose that $ln(p_R/p)$ is linear, and positively sloped to reflect excessive optimism. Then if the slope is steep enough, the sum of sentiment and the fundamental log-pricing kernel will also be positive. That is, nonzero sentiment can induce the pricing kernel to be positively sloped, in line with the pricing kernel puzzle.

If sentiment is strong enough, then the shape of the sentiment function will dominate the shape of the fundamental log-pricing kernel in the determination of the equilibrium pricing kernel. In this regard, overconfidence reflects investors' underestimate of return standard deviation. The log-change of measure has an inverted U-shape, and if strong enough will carry over to the shape of the log-pricing kernel.

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 characterized the empirical findings for the shape of the pricing kernel in the work of Aït-Sahalia and Lo (2000) and Rosenberg and Engle (2002).

2.3 Risk Neutral pdf

Closely related to the pricing kernel M is the risk-neutral density q. In the behavioral framework, the risk neutral measure can be derived from the representative investor's beliefs through the following change of measure:

$$q(S_{t+\tau}/S_t) = \frac{p_R(S_{t+\tau}/S_t)(S_{t+\tau}/S_t)^{-\theta_{1,t+\tau}}}{E(S_{t+\tau}/S_t)^{-\theta_{1,t+\tau}}} \tag{5}$$

where the expectation in the denominator of the right-hand-side is with respect to p_R. Notice that (5) implies that the risk neutral density q is determined in accordance with the beliefs p_R of the representative investor and not necessarily the objective density p. Equation (5), with p in place of p_R, is a key relationship in traditional asset pricing theory.

Equation (5) can be inverted to obtain p_R by applying a change of measure to the risk neutral density q. However, the expectation $E(S_{t+\tau}/S_t)^{-\theta_{1,t+\tau}}$ is taken with respect to p_R, which is what we are solving for. To eliminate the appearance of circular dependence, we therefore use the equivalent term $E_p(M)/\theta_{0,t+\tau}$. Here $E_p(M)$ is the expected value of the SDF M with respect to the objective pdf p, or equivalently the inverse of the gross risk-free rate. Making the substitution, and for simplification omitting arguments for p_R and q as well as time subscripts on θ_0 and θ_1, yields the expression:

$$p_R = \frac{E_p(M)(S_{t+\tau}/S_t)^{\theta_1}}{\theta_0} q \tag{6}$$

Notably, (4) and (6) provide alternative methods for estimating the representative investor's beliefs p_R. The first involves transforming the objective density p, and the second involves transforming the risk neutral density q.

Equations (2), (3), (4) and (6) serve as the basis for applying behavioral asset pricing theory to analyze the empirical pricing kernel. In applying these equations to the data, our interest is in using the empirical pricing kernel to estimate sentiment, meaning behavioral biases on the part of the market (understood as the representative investor). For example, at any point in time, is the representative investor excessively optimistic, overconfident, or possibly both? What are the time series properties of excessive optimism and overconfidence? Are behavioral biases related to broader asset pricing issues such as the equity premium puzzle?

3 Estimating the Empirical SDF

We use the same approach as [Barone-Adesi, Engle, and Mancini (2008)] (BEM) to estimate the SDF (or pricing kernel). Options and stock data are from OptionMetrics, filtered as in BEM, and cover the period January 2002 to October 2009. Options data are used to estimate the risk neutral density q and stock data are used to estimate the objective density p. We briefly recall the method. For a detailed description of it we refer the reader to BEM.

3.1 Estimating p and q

For each Wednesday t in our sample, from 2002 to 2009, we estimate two asymmetric [Glosten, Jagannathan, and Runkle (1993)] (GJR) GARCH models. A GJR GARCH model is fitted to historical daily returns of the S&P 500 to describe the index dynamic under the objective or historical distribution with pdf p. The estimation is obtained via Gaussian Pseudo Maximum Likelihood. Another GJR GARCH model is calibrated to the cross section of out-of-the-money (OTM) options on S&P 500 capturing the index dynamic under the risk neutral or pricing distribution with pdf q. The calibration is achieved via non-linear least squares, i.e. minimizing the sum of squared pricing errors with respect to the GARCH parameters. Then, for each Wednesday t, transition densities of the S&P 500 under p and q are estimated by Monte Carlo Simulation. Using [Duan and Simonato (1998)] Empirical Martingale Simulation method, we simulate 50,000 trajectories of the S&P 500 from t to $t+\tau$, where e.g. τ corresponds to one year. Transition densities p and q are obtained by smoothing the corresponding simulated distribution, as in nonparametric kernel density estimation. In this approach, first and second moments for annual returns under p and q respectively are based on GARCH models estimated at daily frequencies.

To estimate the process for p, we need to make an assumption about the evolution of its mean. To this end, we use a measure of earnings-to-price ratio (E/P) based on the work of [Campbell and Shiller (1998)]. They develop a P/E ratio for U.S. stocks in which a stock price index P is divided by an average of earnings over the prior ten years. They call this the cyclically adjusted price-earnings ratio (CAPE). Here both price and earnings are adjusted for inflation. The key result of their analysis is that subsequent ten-year returns to stocks are negatively and statistically related to P/E. Campbell and Shiller suggest that P/E reflects sentiment. When investors become "irrationally exuberant," prices rise relative to earnings in an unwarranted manner. That is, future returns are low because current prices are too high. Updated data series are available from Robert Shiller's website.

For the market as a whole, E/P can be interpreted as expected steady state long-run return, as the present value of growth opportunities for the market is zero. A value of E/P equal to 25 corresponds to an expected long-term return of 4%. As a consistency check, we regress subsequent annualized ten-year returns for the Campbell-Shiller series on E/P. The regression equation for annualized return is 0.012 + 0.76 E/P. We also regressed returns just for the S&P 500 on E/P and obtained the regression equation $-0.022 + 1.21E/P$. In this regard, we ignore the issue of overlapping intervals in the estimation itself, as the bias strikes us as minor.

In our analysis, we used CAPE as the basis for expected return in estimating the objective pdf p. We report the results for the regression equation in which expected return is given by 0.012 + 0.76 E/P. However, results based on the other specifications are similar in most respects.

3.2 Estimating the SDF

The empirical or unconditional SDF $M_{t,t+\tau}$, can then be estimated semiparametrically as

$$M_{t,t+\tau} = e^{-r\tau} \frac{q(S_{t+\tau}/S_t)}{p(S_{t+\tau}/S_t)} \tag{7}$$

where q is the risk-neutral density, p the objective (i.e., historical) density, and S the S&P 500 index.

On each Wednesday t, we estimate $M_{t,t+\tau}$ for the fixed horizon τ = 1 year. We also consider two GJR GARCH models under the risk neutral pdf q. One GJR GARCH model driven by Gaussian innovations, and another GJR GARCH model driven by filtered historical innovations; see BEM. We refer to these GARCH models and the corresponding estimates of SDF simply as Gauss and FHS methods, respectively.

3.3 Jointly Estimating CRRA Parameters and Sentiment

The least square regression of the empirical pricing kernel, $M_{t,t+\tau}$, on the CRRA pricing kernel gives the closest (in least square sense) CRRA pricing kernel to the empirical one. Given (2), the regression is run in log-log space because in this space the CRRA-pricing kernel is linear. We estimate $\theta_{0,t}$ and $\theta_{1,t}$ on each Wednesday t between January 2002 and October 2009 by regressing the log of the unconstrained pricing kernel on the log gross return.

To illustrate the procedure, consider Figure 1. This figure displays the empirical pricing kernel (SDF unconstrained), the best CRRA fit to the points along the empirical pricing kernel (CRRA-constrained SDF), and the difference between the two (whose logarithm is the sentiment function, computed using (3)).

From January 2002 to October 2009, we run the regression above on each Wednesday t and for the time horizon $\tau = 1$ year. For each Wednesday, we obtain a grid of 100 values of $S_{t+\tau}/S_t$ and compute the pointwise difference, d_t, between the unconstrained and CRRA-constrained pricing kernel. For each gross return, $S^{(j)}_{t+\tau}/S_t$, $j = 1, \ldots, 100$, this difference is defined as

$$d_t^{(j)} = ln(M_{t,t+\tau}) - ln(M_{t,t+\tau}(\theta)) \tag{8}$$

where $M_{t,t+\tau}$ is the unconstrained SDF and $M_{t,t+\tau}(\theta)$ the CRRA-constrained SDF. We use the following distance measures:

$$RMSE_t = 100\sqrt{\frac{1}{n}\sum_{j=1}^{n}(d_t^{(j)})^2} \tag{9}$$

and

$$MAE_t = 100\sum_{j=1}^{n}|d_t^{(j)}| \tag{10}$$

resembling traditional root mean square error and mean absolute error.

4 Sentiment and the Financial Crisis

In this section, we describe our findings within the context of how the S&P 500 evolved between 2002 and 2009. In doing so, we begin with earnings and returns, and then move on to the evolution of sentiment.

4.1 Earnings and Return History

Figure 2 displays the time series for real earnings and the S&P 500 for our sample period. During the period, the S&P 500 fell during 2002, then steadily rose through late 2007, after which it declined in two stages. A minor decline ("correction") in stock prices occurred until the Lehman Brothers bankruptcy on September 15, 2008. Thereafter, a major decline occurred between September 15, 2008 and March 2009, after which the S&P 500 increased rapidly. The earnings trajectory followed a similar pattern. The correlation coefficient for earnings and the S&P 500 is 0.84. Housing prices, which played a key role in the runup to the financial crisis and its aftermath, peaked in late 2006. The eighteen month recession associated with the financial crisis began in late 2007.

4.2 P/E

Over the course of the sample period, CAPE fell from about 30 at the start of our sample period to about 22 at the end of 2002. It then rose and meandered around 26 until the decline that began in late 2007. During the decline it fell to about 14, and then rose to about 18 at the end of our sample period. To place P/E in context, sustained values above 25 are historically rare. The major stock market declines that occurred in 1901, 1929, 1937, and 1966 are cases in point. In each of these cases, crashes occurred within months of P/E rising above 25. Only in the period 1998 through 2000, did P/E rise above 25 without a market crash occurring within months. When the market did peak in March 2000, P/E was about 42. The dramatic decline in earnings during 2008 is evident in Figure 2.

The period between mid-2003 and late 2007 was exceptional, in that P/E remained in the vicinity of 25 for the entire time. This was a period of steady, robust earnings growth. Here is a story about the evolution of events in the equity market which preceded the financial crisis. After the recession of 2001, earnings began to grow in 2002. However, investors were not persuaded that earnings would persist, especially as the growth rate declined. As a result, P/E fell in 2002. However, in early 2003, the growth rate in earnings increased markedly. Investors then became persuaded that the higher rate would persist, and P/E rose, from 22 (which although much less than 42 was still high) to 27. Notably volatility began to decrease in the second half of 2003. Investor overconfidence rose during 2003, as did pessimism. The increased pessimism reflects the fact that investors did not raise their earnings expectations sufficiently high, thereby resulting in positive earnings surprises. Over the next four years, stock prices increased at about the same rate as earnings, thereby responding to the combination of earnings surprises and lower volatility.

4.3 Sentiment

Consider the beliefs and biases encapsulated within the representative investor's pdf p_R. Recall that p_R is computed by transforming the risk neutral pdf q using (6). Figure 3 illustrates the expected return and volatility for the objective stochastic process, along with our estimates of excessive optimism and overconfidence, measured as the difference between the processes for p_R and p.

According to our estimations, the expected objective return fell during 2002 and then remained stable at around 4.5% through 2007. In contrast, the representative investor's expected return rose sharply during this period, peaking at about 7%. Given our assumptions, optimism grew during the middle portion of the sample period, peaking in late 2006 at 2.5%, at which point it generally fell until the market bottom in March 2009. Notably, optimism turned to pessimism in April 2008, and with a couple of exceptions became consistently negative in September 2008, at the time of the Lehman Brothers bankruptcy. We note that for the alternative assumptions we employed, this pattern is quite robust.

Overconfidence generally rose discernably between 2003 and late 2007, falling sharply thereafter in two phases, once in June 2007 and a second time in the month after the Lehman bankruptcy, when overconfidence rose sharply and then plummeted. This second episode, in 2008, is especially interesting, and is related to volatility, which we discuss below.

Given the prominence that housing prices played in the financial crisis, we analyzed the degree to which our measures of sentiment were related to the Shiller housing price index. Using quarterly data, we found strong correlations between the level of housing prices and both optimism ($\rho = 0.88$) and overconfidence ($\rho = 0.69$). A regression of optimism on its own lagged value and the housing price index resulted in t-stats of -2.8 and 4.6 respectively. A similar regression for overconfidence resulted in t-stats of 0.4 and 3.6 respectively. Interestingly, there was no statistical relationship between the quarterly change in housing prices and the two sentiment variables.

Because the objective and risk neutral pdfs are not typically normal, overconfidence is not necessarily uniformly distributed between left and right tails. The left tail is measured by the probability that the gross return is less than 0.8, and the right tail by the probability that the gross return exceeds 1.2. In investigating this issue, we examined the tails under objective and risk neutral densities, as well as under the representative investor density, calculated using both (4) and (6). Overall, we find that overconfidence was manifest in both tails. In the middle of the sample, the market underestimated the size of the right tail by about 10% and the left tail by just under 5%.

4.4 Pricing Kernel

Figure 4 displays the SDF estimated on each Wednesday from 2002 to 2009 using FHS GARCH. This figure provides a three-dimensional bird's eye view of how the shape of the pricing kernel changed with time. In the middle portion of the sample, the FHS-based pricing kernel had the shape of an inverted-U. This shape corresponds to overconfidence. However, in January 2003, the shape of the pricing kernel was much closer to being monotone declining, with a flat region at the left.

Notably, in January 2003, overconfidence was -2.4%. Recall the earlier discussion about how optimism impacts the shape of the pricing kernel when overconfidence is zero. That discussion pointed out that in this case the pricing kernel has a traditional monotone declining shape, but is excessively steep. In the Gaussian case, the magnitude of the pricing kernel slope is flatter by the magnitude of optimism bias. In January 2003, the magnitude of optimism bias was about -1%, so the adjustment to the slope from optimism bias was -0.01. To place this value in context, in our sample, the mean value for optimism bias is 0.2%.

Optimism and overconfidence both declined sharply after the Lehman Brothers bankruptcy. Optimism fell from -1% in October 2008 to -5.4% in March 2009 before climbing back to 1% at the end of the sample period.

In the presence of pessimism (negative optimism), the pricing kernel will most resemble its CRRA-constrained counterpart when overconfidence approaches zero. In this case, RMSE and MAE should be small. In our data, RMSE and MAE achieve their lowest values in 2002 and 2007–2009, as optimism and overconfidence reach their respective local minima.

As overconfidence fell in the aftermath of the Lehman Brothers bankruptcy, the shape of the pricing kernel shifted. It went from being inverse-U with a very high peak in September 2008 to being monotone declining, except for a small hump or flat region at the left, during 2009.

4.5 Volatility

Figure 5 displays the time series for volatility (during the prior twelve months), VIX, and return standard deviations for both the objective density and representative investor's density. Notice that these curves are quite similar to each other in the middle of the sample period, but diverge at the extremes when actual volatility and the VIX exceed the return standard deviations associated with both the objective pdf and the representative investor's pdf. For the most part, the representative investor's return standard deviation lay below its objective counterpart, with the gap representing overconfidence. After the Lehman Brothers bankruptcy, when volatility spiked dramatically and atypically exceeded the VIX, overconfidence fell to near zero.

During the middle portion of the sample, the higher levels of overconfidence, in combination with high P/E, meant that investors underestimated the probability of a long-term tail event associated with a sharp drop in earnings and very low returns. Earnings growth begin to decline in 2007, with a sharp jump at the end of June. Earnings growth subsequently turned negative in September. Earnings continued to decline through 2008, with the next sharp jump occurring in October 2008, shortly after the Lehman bankruptcy.

Pessimism, meaning negative optimism, and overconfidence were highly volatile during the final quarter of 2008, and continued to fall at the end of 2008. They also continued to be volatile, but began to increase after the market bottomed in March 2009. These events underlie the time path of RMSE and MAE, the two measures of the degree to which the pricing kernel deviated from its CRRA-constrained counterpart. Notably, RMSE and MAE both drop sharply during 2007, and generally increase until September 2008, after which they decline as earnings fall dramatically.

4.6 Interest Rates and Risk Premium

Interest rates fell to 1% in 2002, and then rose through early 2008, peaking just below 4%. Thereafter interest rates declined sharply, falling to near zero after the Lehman Brothers bankruptcy. The objective risk premium is the difference between expected return and the risk-free interest rate. The objective risk premium rose sharply from a low of about 2.4% in 2002 to about 4% several months later, after which it followed a downward path, with considerable volatility, dipping close to zero in 2007. It then rose sharply over the next few months, to about 6.5% before falling to about 5% at the end of the sample period. The representative investor's risk premium followed a different pattern, fluctuating between 2% and 4% for most of the sample period.

The regression equation for annualized return is $0.012 + 0.76\ E/P$. As mentioned previously, we also regressed returns just for the S&P 500 on E/P and obtained the regression equation $-0.022 + 1.21E/P$. Although the general character of the results under both specifications is similar, one point of difference is the sign of the risk premium. Under the first specification, the risk premium is always positive. However, under the second specification, the risk premium is negative between the second quarter of 2005 and the second quarter of 2008. The negative risk premium reflects the combination of a low expected return associated with high values of P/E, stemming from the negative intercept in the regression equation, and rising interest rates during this period.

From a behavioral perspective, a negative risk premium is not problematic. Rather, it stems from the combination of optimism and overconfidence. For example, in the first specification, optimism peaks at about 2%, while in the second specification, optimism peaks at just under 4%.

4.7 Risk Aversion

The associated estimates of $\theta_{1,t}$ are generally consistent with the range reported in the survey of [Meyer and Meyer (2005)], 0.8 to 4.72. Meyer and Meyer discuss the survey paper by [Barsky, Juster, Kimball, and Shapiro (1997)] (BJKS). BJKS describe the use of a survey question, designed to elicit θ_1. The question asks people if they would accept a 50-50 chance that their income next year either doubles or decreases by x%. The status quo is to maintain their current income. In a CRRA framework, the value of x that makes a person indifferent to accepting the risk is approximately $1/1+\theta_1$. For logarithmic utility, meaning $\theta_1 = 1$, $x = 50$. This is consistent with the prospect theory property that people experience a loss at twice the intensity of a gain of comparable magnitude. For $\theta_1 = 2$, $x \approx 33$. For $\theta_1 = 3$, $x \approx 25$. These are values that conform to the survey responses reported by BJKS. However, for the values of θ_1 needed to explain the equity premium puzzle, that being $\theta_1 = 20$, $x \approx 5$, a figure well below the estimates from survey evidence.

With respect to our sample, risk aversion generally declined during 2002 from about 1.3 to 0.3, and then rose in a wave pattern through late 2007, first increasing through 2005 above 3.0, falling through mid-2006 to just under 1.0, rising again until 2007, and then following dramatically with volatility, reaching zero in the aftermath of the Lehman Brothers bankruptcy. Risk aversion then rose through the rest of the sample period. Notably, risk aversion was at its lowest during the down markets at the beginning and end of the sample periods and highest in the up market during the middle period.

Interestingly, the time series for risk aversion is consistent with the well known behavioral property in which individual risk attitude features risk aversion in the domain of gains and risk seeking in the domain of losses.

4.8 Time Preference

BJKS report evidence that time preference discount factor lies above 1. We do find that the representative investor's time discount factor, adjusted for estimation bias, remained in a region above 1.0 for most of the sample period. The mean value for time preference was 1.04. Its range was from 0.97 to 1.31, with the peak being associated with the Lehman bankruptcy. The discount factor became quite volatile between 2005 and late 2007, with some sharp spikes, indicating increased patience. The greatest volatility and highest spikes occurred at the time of the Lehman Brothers bankruptcy, but the discount factor reverted to its historical range after the market bottom in March 2009.

5 External Measures of Sentiment

Systemic risk can increase as excessive optimism and overconfidence induce investors to become complacent. In this section, we describe two measures of sentiment that are independent of the estimates we reported earlier. The first is the series of Yale/Shiller confidence indices. The second is the Baker–Wurgler sentiment series.

5.1 Yale/Shiller Confidence Indexes

The Yale/Shiller U.S. data is based on a survey of two samples of investors. The first is a sample of wealthy individual investors, and the second is a sample of institutional investors. For the time period studied in our paper, the first sample consists of a random sample of high-income Americans from Survey Sampling, Inc. The second sample consists of investment managers listed in the *Money Market Directory of Pension Funds and Their Investment Managers.*

The Yale/Shiller confidence indexes consist of monthly six-month averages of monthly survey results. For example, an index value for January 2002 is an average of results from surveys between August 2001 and January 2002. Sample size has averaged a little over one hundred per six-month interval since the beginnings of the surveys. This means that standard errors are typically plus or minus five percentage points.

There are four confidence indexes. These measure confidence that there will be no crash in the next six months (C), confidence that the market will go up in the next twelve months (O), confidence that the market is fairly valued (V), and confidence that the market will reverse in the short-term (buying-ondips, B). We use the symbol P for institutional (professional) investor and I for individual investor. This leads to eight confidence series, CP, CI, OP, OI, VP, VI, and BP, BI. Of these, the crash confidence index for institutional investors turns out to be the most informative. For this reason, we provide more detail about its construction below.

The crash confidence index measures the percentage of the population who attach little probability to a stock market crash in the next-six months. The survey question used to elicit the index is as follows.

> What do you think is the probability of a catastrophic stock market crash in the U. S., like that of October 28, 1929 or October 19, 1987, in the next six months, including the case that a crash occurred in the other countries and spreads to the U. S.? (An answer of 0% means that it cannot happen, an answer of 100% means it is sure to happen.)

The Crash Confidence Index is the percentage of respondents who think that the probability is less than 10%.

5.1.1 Crash Confidence During 2002–2009

At the beginning of the sample period, the crash confidence index CP was low, below 30. During 2002, it rose to about 40 and then fell sharply to about 21. Between 2003 and late 2007, the crash confidence index trended up, peaking just below 60. During the decline, crash confidence fell to the mid 30s, where it remained until the Lehman Brothers bankruptcy. After the bankruptcy, crash confidence fell sharply, bottoming below 20. As with RMSE and MAE, it too dropped sharply in 2007 and the last portion 2008, as earnings declined sharply. In this respect, CP provides external corroboration for the evolution of optimism and overconfidence.

There is reason to suggest that CP serves as an indicator of systemic risk. High values of CP suggest that the majority of institutional investors attached not just low probability, but insufficient probability, to outlying events. Keep in mind that the Yale/Shiller indexes are indicators of the proportion of those holding particular views. Other indexes are also informative. The value confidence index, both for institutional investors and individual investors rose during 2008, suggesting that investors increasingly viewed the market decline during this period to have been an overreaction. A similar statement applies to the one year confidence indexes.

5.2 Baker–Wurgler Sentiment Index

[Baker and Wurgler (2006)] consider six proxies for sentiment suggested in the behavioral finance literature and form a composite sentiment index based on their first principal component. The six proxies are closed-end fund discount, detrended log-turnover, number of IPOs, first-day return on IPOs, dividend premium, and equity share in new issues. The composite series is the Baker–Wurgler sentiment index. To reduce the likelihood that these proxies are connected to systematic risk, they also form an index based on sentiment proxies that have been orthogonalized to several macroeconomic conditions. They point out that the two sentiment indexes visibly line up with historical accounts of bubbles and crashes. This series is currently available through December 2007, from Jeff Wurgler's website.

We find that the Baker-Wurgler sentiment index is positively correlated with optimism and is negatively correlated with OP and OI. Statistically, Baker-Wurgler sentiment is driven by a regression equation with these three variables, together with an AR(2) process for the disturbances. The positive relationship between the Baker-Wurgler index and optimism is as expected. The negative relationship between the Baker-Wurgler index and the Yale/Shiller indexes is something of a surprise, and reflects the fact that during the early portion of our sample, the Baker-Wurgler index fell sharply, whereas the Yale-Shiller indexes rose.

6 Sentiment, Systemic Risk and Leverage

[Acharya et al. (2010)] introduce the concept of Marginal Expected Shortfall (MES) for an individual financial company. [Brownlees and Engle (2010)] define MES as "the expected equity loss per dollar invested in a particular company if the overall market declines by a certain amount." When it comes to systemic risk, Brownlees and Engle indicate that the "companies with the highest MES are the companies that contribute the most to the market decline and are therefore the most important candidates to be systemically risky. Equity holders in a company that is systemically risky will suffer major losses in a financial crisis and consequently will reduce positions if a crisis becomes more likely. MES measures this effect. It clearly relies on investors recognizing which companies will do badly in a crisis."

The Volatility Lab (Vlab), at New York University, computes and reports values of MES and leverage for major financial institutions. Using their data, we analyze how leverage and MES changed as market sentiment began to deteriorate beginning in March 2007, and through the crisis which erupted in 2008. In particular, we consider how MES and leverage changed for particular financial firms as optimism and overconfidence began to decrease. In this regard, we note that sentiment impacts both the magnitude of systemic events as well as the frequency with which they occur.

Leverage is measured by the ratio of assets to market value of equity. We focus attention on the leverage time series for the following firms who were prominent in the financial crisis: American International Group (AIG), Bank of America, Bank of New York Mellon, Bear Stearns, Citigroup, Fannie Mae, Freddie Mac, Goldman Sachs, JP Morgan Chase, Lehman Brothers, Merrill Lynch, Morgan Stanley, and Washington Mutual.

We draw attention to a distinct shift in sentiment that took place between March and August of 2007. Overconfidence steadily decreased from 6.25% to 3.7%, even though the S&P 500 rose by 3.7%, and optimism was quite stable at about 2% until July when it dipped to 1.2%. In respect to external validation, the Yale/Shiller crash confidence index CP plummeted from 55% to 34%, and the Baker-Wurgler index exhibited no discernable trend.

Between March and August, leverage began to increase for almost all of the firms mentioned above. Bear Stearns and Lehman were the first firms to fall during the crisis. For Bear Stearns, the increase was marked. From an initial level of 22, its leverage moved to 31. Lehman Brothers's leverage went from 16 to 20. In contrast, Goldman Sachs's leverage went from 10 to 11. Morgan Stanley and Merrill Lynch saw their leverage go from 14 to 18.

The sharp drop in overconfidence between March and August preceded one of two runs on asset backed commercial paper during the financial crisis. The first took place between August and December of 2007. As subprime mortgage defaults rose, financial firms that issued commercial paper in order to raise funds to purchase mortgage backed securities suddenly found themselves unable to do so. [Covitz et al. (2009] report evidence of panic early on, as the run extended to financial firms issuing commercial paper not backed by mortgage assets. During the run, optimism and overconfidence both fell sharply, falling to about -2% in mid-October before climbing back to positive territory near zero for the remainder of the run.

The movements in sentiment coincided with concerns that were not reflected in the S&P 500, which continued to trend upward until December 2007. In this regard, sentiment might have been related to conditions in the housing market. Between March and December of 2007, the house price index declined at a 10% rate. During our sample period, the house price index was highly correlated with sentiment. The correlation coefficients for optimism, overconfidence, and crash confidence are respectively 0.88, 0.69, and 0.73. However, we note that none of these variables Granger caused any of the others, which leads us to be cautious about making attributions of causality.

December 2007 marks the beginning of an eighteen month recession. Optimism and overconfidence briefly rebounded in January to 1% and 4% respectively before beginning a volatile decline back to negative territory. By mid-March when Bear-Stearns failed, but for agreeing to be acquired by JP Morgan Chase, optimism was near zero and overconfidence was at 2%. At the time Bear Stearns's leverage stood at 254.

By May of 2008, optimism and overconfidence had both fallen to near zero. At this time, JP Morgan Chase completed its acquisition of Bear Stearns. JP Morgan Chase's leverage stood at 11, up from 8 in March 2007. In contrast, the leverage of Washington Mutual, which had been a key player in mortgage origination, was 10 in March 2007, rose to 33 in March, when Bear Stearns failed, and jumped to 55 in June. By August, Washington Mutual had failed.

Turning next to systemic risk, consider MES values during April 2008. MES measures expected equity loss for a financial firm if the market loss exceeds 2% on a daily basis. Bear Stearns, which was about to be acquired, had an MES of 12.4%. Lehman Brothers, which declared bankruptcy in September, had an MES was 8.75%. AIG's MES stood at 6.21%. Merrill Lynch, the dominant underwriter of CDOs, and which needed to be acquired by Bank of America at year-end, was close behind at 8.15%. In March 2007, MES values had been considerably lower. Bear Stearns's MES was at 4.5%, as was the MES of Lehman.

By the end of September when Lehman declared bankruptcy, AIG's MES had soared to 25.8%. Citigroup, Merrill Lynch, Morgan Stanley, and even JP Morgan Chase all had MES values above 8%. Goldman Sachs stood out as an exception, with an MES of 6.5%. By this time, the market had turned from being optimistic to being pessimistic (by over 1%). Notably, overconfidence soared to almost 8%, as investors seriously underestimated future volatility.

During the remainder of our sample period, leverage levels continued to rise. Morgan Stanley's leverage peaked at 52 in October 2008. It was unclear whether Morgan Stanley would survive, and the firm sought protection by registering as a holding company with the Federal Reserve. AIG's leverage peaked at 143 in February 2009. Citigroup's leverage peaked at 122 in March. MES levels were also elevated between September 2008 and March 2009, and for most of the surviving firms began to decline towards the end of the sample period. AIG was an exception, as its MES, which had fallen to 5.53% in May 2009, rose again in the summer to 11%.

The relationship between systemic risk and sentiment is complex. If one asks how sensitive was MES to changes in sentiment, the answer is that it was highly firm dependent. For Bank of America, the correlation between MES and optimism was −84%. For Fannie Mae and Freddie Mac, MES was −72% and −77% respectively. For Citigroup, it was −71%. For AIG, it was −48%. For Bear Stearns, it was −41%. Notably, for Goldman Sachs, it was −36%. Correlations of MES with overconfidence are much lower than for optimism, and also vary in sign.

If one asks about the correlation of leverage with optimism, the answer is that these tended to be large and negative, in the range of −70% and below. Goldman Sachs was an exception, at −35%. However, the correlations of leverage with overconfidence were closer to zero, and mixed in sign.

If one asks about the impact of sentiment on systemic events, the answer is that fluctuations in sentiment tend to increase the frequency of these events, along with their magnitude. This statement is consistent with the events occurring between September 2008 and April 2009.

Finally, if one asks about whether large changes in sentiment can serve as signals in respect to systemic events, the answer is mixed. We find no evidence of Granger causality. At the same time, we note that a dramatic drop in overconfidence during the first half of 2007 preceded the run occurring several months later on asset backed commercial paper which represented the leading edge of the financial crisis. In addition, the sharp declines in both optimism and overconfidence, that began in May 2008 preceded the major downturn of September 2008–March 2009.

7 Conclusion

Regulators charged with monitoring systemic risk need to focus on sentiment as well as narrowly defined measures of systemic risk. This chapter describes techniques for jointly monitoring the co-evolution of sentiment and systemic risk.

Regulators need to be consciously aware that systemic risk builds over time, as excessive optimism and overconfidence induce investors to become complacent. This complacency is typically manifest within unsound lending practices and increased leverage. When sentiment changes direction, often in response to changing fundamentals, the reversal in optimism and confidence can give rise to large increases in systemic risk, high volatility, and large losses.

References

[Acharya et al. (2010)] Acharya, V., L. Pedersen, T. Phillipe, and M. Richardson (2010). *Measuring Systemic Risk*. Technical Report, Department of Finance, New York University.

[Aït-Sahalia and Lo (2000)] Aït-Sahalia, Y., and A. Lo, 2000. "Nonparametric risk management and implied risk aversion," *Journal of Econometrics*, 94, 9–51.

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

[Barone-Adesi, Engle, and Mancini (2008)] Barone-Adesi, G., R. Engle, and L. Mancini, 2008. "A GARCH option pricing model with filtered historical simulation," *Review of Financial Studies*, 21, 1223–1258.

[Barone-Adesi, Mancini, and Shefrin (2011)] Barone-Adesi, G., L. Mancini, and H. Shefrin, 2011. "The Pricing Kernel Puzzle: Risk Aversion, Time Preference, and Sentiment," Working paper: Santa Clara University.

[Barsky, Juster, Kimball, and Shapiro (1997)] Barsky, R., F. T. 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.

[Brownlees and Engle (2010)] Brownlees, C., and R. Engle (2010). "Volatility, Correlation and Tails for Systemic Risk Management," Working paper, New York University.

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

[Campbell and Cochrane (1999)] Campbell, J., and J. Cochrane, 1999. "By force of habit: a consumption-based explanation of aggregate stock market behavior," *Journal of Political Economy*, 107, 205–251.

[Cochrane (2005)] Cochrane, J., 2005. *Asset pricing, second edition*, Princeton: Princeton University Press.

[Covitz et al. (2009] Covitz, D., N. Liang, and G. Suarez, 2009. "The Evolution of a Financial Crisis: Runs in the Asset-Backed Commercial Paper Market," Working paper, Federal Reserve Board.

[Duan and Simonato (1998)] Duan, J.-C., and J.-G. Simonato, 1998. "Empirical martingale simulation for asset prices," *Management Science*, 44, 1218–1233.

[Glosten, Jagannathan, and Runkle (1993)] 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.

[Hens and Reichlin (2011)] Hens, T., and C. Reichlin, 2011. "Three Solutions to the Pricing Kernel Puzzle," Working paper, University of Zurich.

[Jackwerth (2000)] Jackwerth, J. C., 2000. "Recovering risk aversion from options prices and realized returns," *Review of Financial Studies*, 13, 433–451.

[Jouini and Napp (2006)] Jouini, E., and C. Napp, 2006. "Heterogeneous beliefs and asset pricing in discrete time," *Journal of Economic Dynamics and Control*, 30, 1233–1260.

[Jouini and Napp (2007)] Jouini, E., and C. Napp, 2007. "Consensus consumer and intertemporal asset pricing with heterogeneous beliefs," *Review of Economic Studies*, 74, 1149–1174.

[Kahneman and Tversky (1979)] Kahneman, D., and A. Tversky, (1979). "Prospect Theory: An Analysis of Decision Making Under Risk," *Econometrica*, 47, 263–292.

[Lopes and Oden (1999)] Lopes, L. L., and Oden, G. C. (1999). "The role of aspiration level in risk choice: A comparison of Cumulative Prospect Theory and SP/A Theory," *Journal of Mathematical Psychology*, 43, 286–313.

[Meyer and Meyer (2005)] Meyer, D., and J. Meyer, 2005. "Relative Risk Aversion: What Do We Know?," *Journal of Risk and Uncertainty*, 31, 243–262.

[Rosenberg and Engle (2002)] Rosenberg, J., and R. Engle, 2002. "Empirical pricing kernels," *Journal of Financial Economics*, 64, 341–372.

[Shefrin (2008)] Shefrin, H., 2008. *A behavioral approach to asset pricing*, second edition. Boston: Elsevier Academic Press.

[Tversky and Kahneman (1992)] Tversky, A., and D. Kahneman, 1992. "Advances in Prospect Theory: Cumulative Representation of Uncertainty," *Journal of Risk and Uncertainty*, 5, 297–323.

[Ziegler (2007)] Ziegler, A., 2007. "Why Does Implied Risk Aversion Smile?," *Review of Financial Studies*, 20, 859–904.

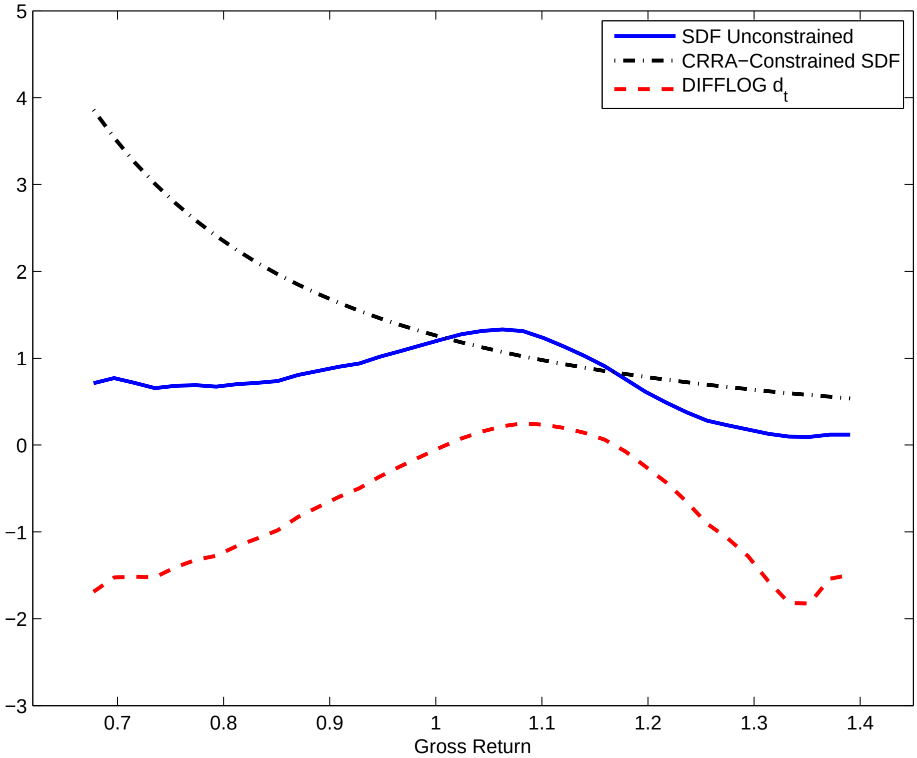

Figure 1: Unconstrained SDF, CRRA-constrained SDF, and d_t function in (8) for the date 21/12/2005.

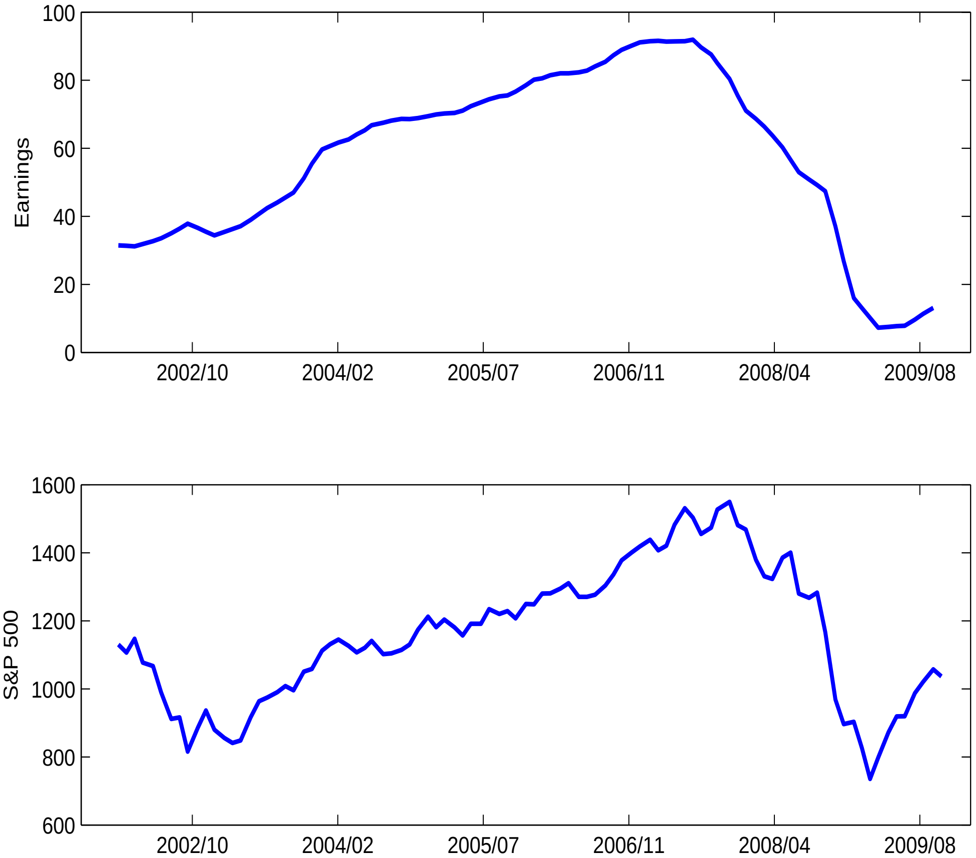

Figure 2: Time series for real earnings and the S&P 500 for our sample period January 2002 – October 2009.

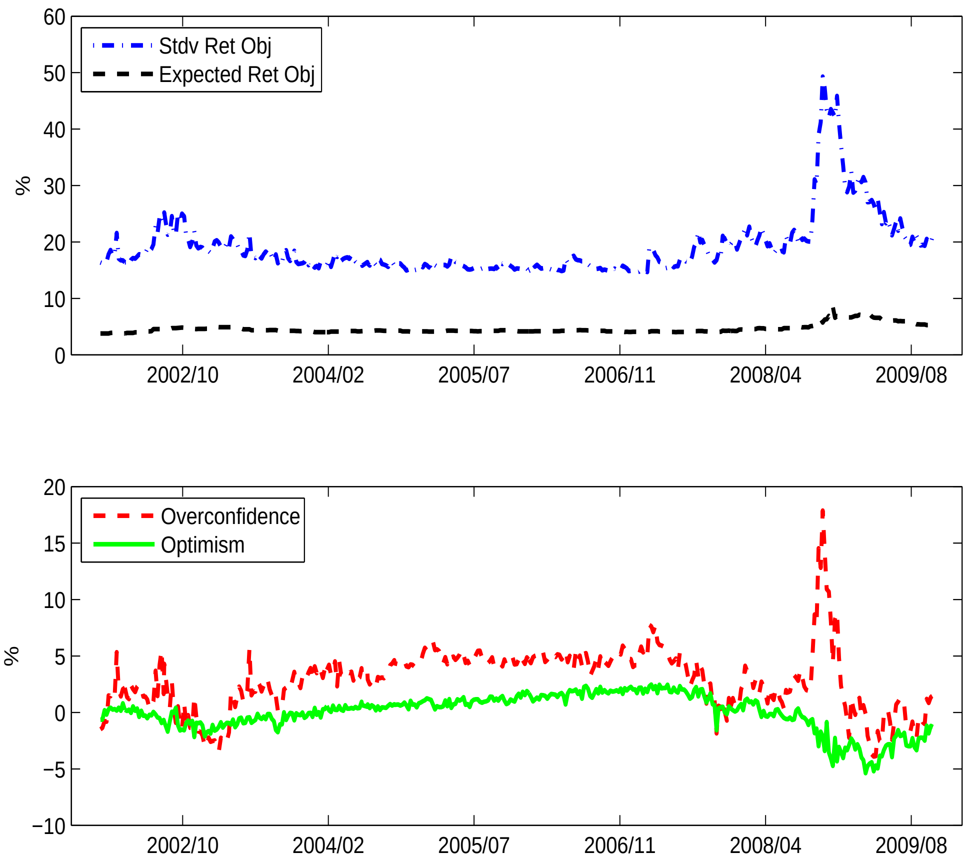

Figure 3: Upper graph: time series for the expected mean and standard deviation of the market return. Lower graph: time series difference between the means and standard deviations of the objective density and the representative investor's density (i.e. optimism and overconfidence, respectively) using the FHS method for our sample period January 2002 – October 2009.

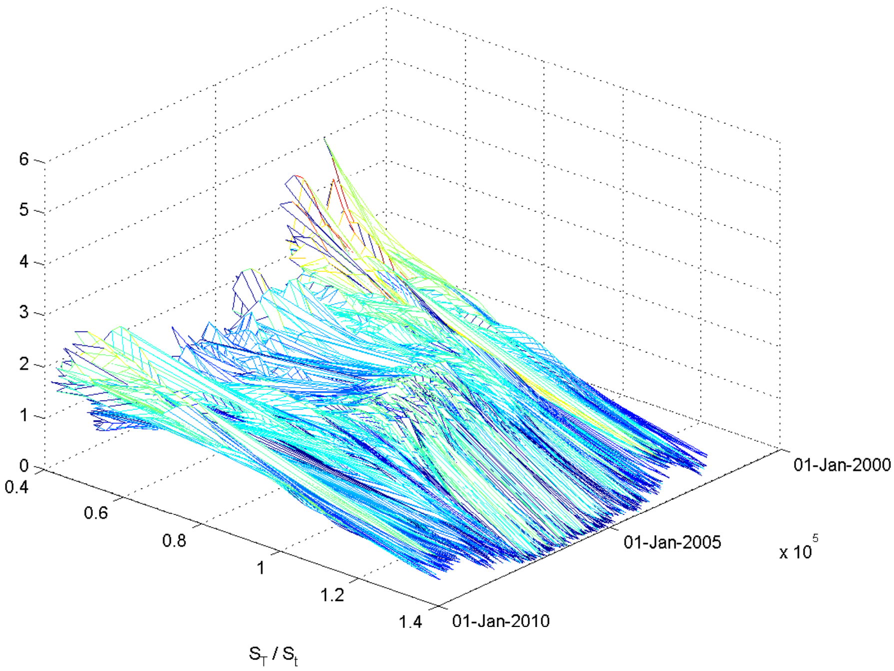

Figure 4: SDF estimated semiparametrically using FHS method for each Wednesday between January 2002 and October 2009 and for the horizon of one year.

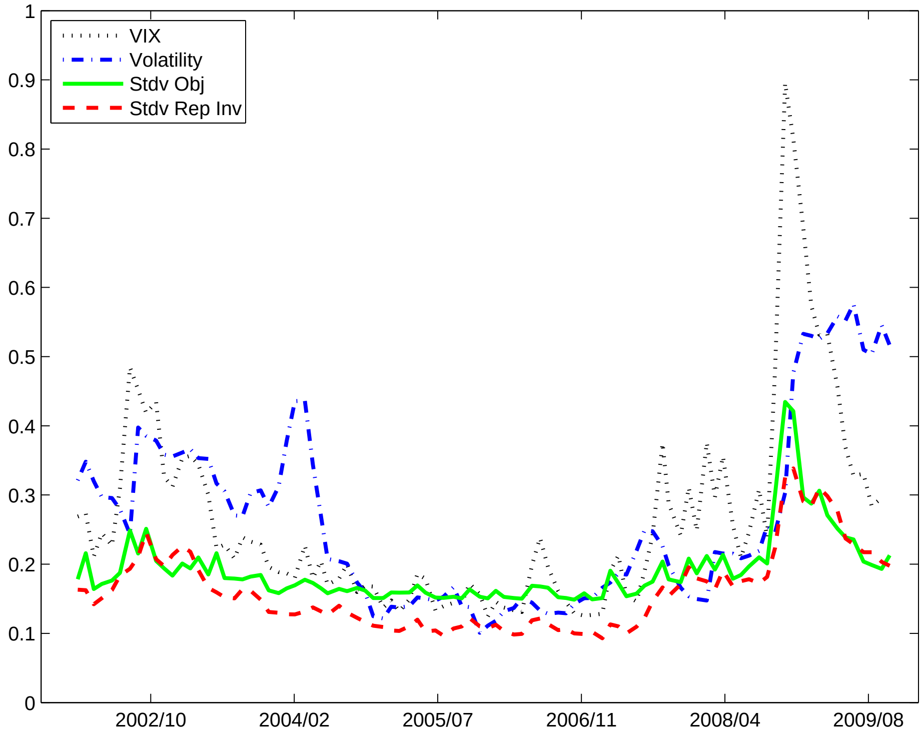

Figure 5: Time series for objective volatility (during the prior twelve months), VIX, and return standard deviations for both the objective density and representative investor's density for our sample period January 2002 – October 2009.