

# Measuring Investors' Risk Appetite\*

Prasanna Gai and Nicholas Vause  
Bank of England

This paper proposes a method for measuring investor *risk appetite* based on the variation in the ratio of risk-neutral to subjective probabilities used by investors in evaluating possible future returns to an asset. Unlike other indicators advanced in the literature, our measure of market sentiment distinguishes risk appetite from *risk aversion*, and is reported in levels rather than changes. Implementation of the approach yields results that respond to crises and other major economic events in a plausible manner.

JEL Codes: G10, G12, G13.

Financial market practitioners often cite market sentiment as a key factor driving broad trends in asset prices. The prices of financial assets frequently move together, even though many of the factors affecting valuations in different asset markets can be quite different. The Asian financial crisis illustrates how shifting attitudes toward risk can generate correlation among the prices of seemingly unrelated assets. Following the devaluation of the Thai baht in July 1997, investors reduced their risk exposures across a range of emerging markets, causing a rise in the cost of borrowing beyond Asia, and into Latin America and Emerging Europe. The spillover of financial stress across borders could not be explained by domestic fundamentals alone and coincided with claims that a decline in “risk appetite” was an underlying reason for contagion and financial instability.

---

\*We thank Alex Bowen, Damien Lynch, Paul Robinson, Hyun Shin, and Peter Westaway for helpful comments and encouragement. The usual caveat applies. The views expressed are those of the authors and do not reflect those of the Bank of England. Corresponding author: Vause: Bank of England, Threadneedle Street, London EC2R 8AH, United Kingdom; e-mail: [nicholas.vause@bankofengland.co.uk](mailto:nicholas.vause@bankofengland.co.uk). Other author contact: Gai: Bank of England, Threadneedle Street, London EC2R 8AH, United Kingdom; e-mail: [prasanna.gai@bankofengland.co.uk](mailto:prasanna.gai@bankofengland.co.uk).

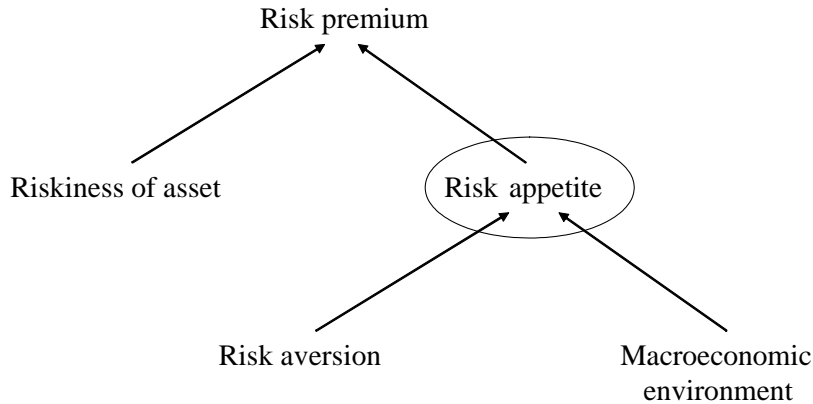
The terms “risk appetite,” “risk aversion,” and “risk premium” are frequently used interchangeably to refer to sentiment in asset markets. But the concepts are very distinct, and inappropriate use makes it difficult to assess and convey the true extent of the willingness to hold risky assets. Investors dislike uncertainty surrounding the future consumption implied by their asset holdings. *Risk appetite*—the willingness of investors to bear risk—depends on both the degree to which investors dislike such uncertainty *and* the level of that uncertainty. The level of uncertainty about consumption prospects depends on the macroeconomic environment. And the degree to which investors dislike uncertainty reflects underlying preferences over lotteries. This *risk aversion* is part of the intrinsic makeup of investors. It is a parameter that our theoretical priors suggest is unlikely to change markedly, or frequently, over time.<sup>1</sup>

Risk appetite, by contrast, is likely to shift periodically as investors respond to episodes of financial distress and macroeconomic uncertainty. In adverse circumstances, investors will require higher excess expected returns to hold each unit of risk and risk appetite will be low—it is the inverse of the price of risk. And when the price of risk is taken together with the quantity of risk inherent in a particular asset, the expected return required to compensate investors for holding that asset is the *risk premium*. Figure 1 illustrates these concepts. It is clearly difficult to disentangle risk appetite from risk aversion and, as Pericoli and Sbracia (2004) note, an increase in either one of them causes asset prices to decline and risk premia to increase.

In what follows, we formally distinguish risk appetite from risk premia and aversion. Specifically, we propose a measure based on the variation in the ratio of risk-neutral to subjective probabilities used by investors in evaluating the expected payoff of an asset. By exploiting the linkages between the risk-neutral and subjective probabilities that can be extracted from financial market prices, we follow Hayes, Panigirtzoglou, and Shin (2003), Tarashev, Tsatsaronis, and Karampatos (2003), and Bollerslev, Gibson, and Zhou (2004). Unlike these papers, however, we are able to extract an indicator of market sentiment that is quite distinct from risk aversion. Moreover, the

---

<sup>1</sup>For recent market-based estimates of risk aversion, see Bliss and Panigirtzoglou (2004).

**Figure 1. Relationship between Risk Concepts**

index of risk appetite based on our approach appears to respond to crises and other economic events in a plausible fashion and, as such, compares favorably with other measures advanced in the literature.

The paper is organized as follows. Section 1 sets out the theoretical basis of our risk appetite measure. Section 2 presents the data and discusses the empirical strategy used to obtain estimates of risk-neutral and subjective probabilities from options prices and presents our measure of risk appetite. We contrast our approach with the recent literature in section 3, and a final section concludes.

## 1. The Concept of Risk Appetite

The standard treatment of asset pricing theory (e.g., Cochrane 2001) states that in an efficient market, with fully rational and informed investors, the current price of an asset,  $p_t$ , should equal the expected discounted value of its possible future payoffs,  $x_{t+1}$ . These payoffs comprise income (such as dividend payments) received over the horizon, plus the ongoing value of the asset as implied by its future price. More formally,

$$p_t = \mathbb{E}_t(m_{t+1} \cdot x_{t+1}), \quad (1)$$

where  $x_{t+1}$  denotes the payoff in period  $t + 1$ , and  $m_{t+1}$  denotes the discount factor—the marginal rate at which the investor is willing to substitute consumption at time  $t + 1$  for consumption at time  $t$ .

Both  $x_{t+1}$  and  $m_{t+1}$  vary across states of the world. Indeed,  $m_{t+1}$  is usually referred to as the *stochastic* discount factor.

We ensure that  $m_{t+1}$  is unique by assuming that the asset market is complete. This means that it is possible to form portfolios as linear combinations of the assets traded in the market that have positive payoffs in a single state of the world, and otherwise pay zero. Furthermore, it is possible to create many of these portfolios, so that there is a positive payoff for every state. So, if  $m_{t+1}$  were not unique, multiple prices would be a possibility for at least one of the portfolios. But this is inconsistent with the absence of arbitrage opportunities that is associated with rational investors. Hence,  $m_{t+1}$  is a unique stochastic discount factor that prices all assets.<sup>2</sup>

The basic asset pricing equation can also be expressed in terms of gross returns,  $R_{t+1}$ , by dividing equation (1) by current prices. Thus,

$$1 = \mathbb{E}_t(m_{t+1} \cdot R_{t+1}). \quad (2)$$

Although, in general, different assets have different expected returns, all assets have the same expected *discounted* return in equilibrium (of unity). Since both the gross return and the stochastic discount factor are random variables, equation (2) can be written as

$$1 = \underbrace{\mathbb{E}_t(m_{t+1}) \cdot \mathbb{E}_t(R_{t+1})}_{\text{risk-neutral component}} + \underbrace{\text{cov}_t(m_{t+1}, R_{t+1})}_{\text{risk adjustment}}. \quad (3)$$

The first term on the right-hand side of equation (3) reflects the mean return required by investors to hold the asset *if* they were indifferent to risk. The second term is a risk correction required by risk-averse investors. Noting that the gross risk-free rate is given by  $R_{t+1}^f = 1/\mathbb{E}_t(m_{t+1})$ , we can rearrange to obtain the familiar expression

$$\underbrace{\mathbb{E}_t(R_{t+1}) - R_{t+1}^f}_{\text{risk premium}} = -R_{t+1}^f \text{cov}_t(m_{t+1}, R_{t+1}). \quad (4)$$

Equation (4) states that the expected return of a risky asset in excess of that available on a risk-free asset is proportional to *minus* the covariance of its state-contingent rate of return and the stochastic

<sup>2</sup>See Danthine and Donaldson (2005, chap. 11) or Milne (1995, chap. 5) for further detail.

discount factor. Intuitively, an asset that pays a high return in good times when investors have a high level of consumption, but fails to pay out in bad times when investors' consumption is lower, has a disadvantageous pattern of returns. So to encourage investors to hold this asset, the expected return must exceed the risk-free rate, i.e., the asset must offer a *risk premium*.

The risk premium can, in turn, be decomposed into the quantity of risk,  $\beta_i$ , inherent in each asset and the unit price of risk that is common across assets,  $\lambda_t$ . In particular,

$$\mathbb{E}_t(R_{t+1}) - R_{t+1}^f = \underbrace{\frac{-cov_t(m_{t+1}, R_{t+1})}{var(m_{t+1})}}_{\beta_i} \cdot \underbrace{var(m_{t+1}) \cdot R_{t+1}^f}_{\lambda_t}. \quad (5)$$

The price of risk,  $\lambda_t$ , is the expected excess return that investors require to hold each unit of risk in equilibrium. *Risk appetite*—the willingness of investors to bear risk—can therefore be defined as the inverse of the price of risk. So when risk appetite falls, larger expected excess returns are required to hold risky assets.

It is apparent from equation (5) that risk appetite reflects variation in the stochastic discount factor,  $var(m_{t+1})$ . Since the stochastic discount factor specifies the marginal rate at which the investor is willing to substitute uncertain future consumption for present consumption, risk appetite depends on the *degree* to which investors dislike uncertainty about their future consumption and on factors that determine the overall *level* of uncertainty surrounding consumption prospects. The degree of such uncertainty corresponds to *risk aversion*, since the more risk averse the investor, the more valuable is additional income in bad states of the world. Accordingly, risk aversion reflects innate preferences over uncertain future consumption prospects—the curvature of individuals' utility functions—that are unlikely to vary significantly over time.

The factors underpinning risk appetite can be seen more clearly by imposing some structure on the stochastic discount factor. In particular, if consumption growth is log-normally distributed with variance,  $\sigma_t^2(c_{t+1})$ , and investors have power utility functions, then the price of risk is

$$\lambda_t = \gamma \sigma_t^2(c_{t+1}), \quad (6)$$

where  $\gamma$  is the coefficient of absolute risk aversion.<sup>3</sup> So a rise in  $\gamma$  would mean a fall in risk appetite. But risk appetite will also fall if uncertainty about future consumption growth increases. The expected volatility of future consumption is likely to depend on factors such as unemployment prospects, the stance of macroeconomic policy, and so on. In general, one would expect that the periodic shifts in market sentiment witnessed over time are more likely to be driven by the macroeconomic environment rather than by changes in the risk aversion of investors.

The analysis of asset pricing above is couched in terms of investors' subjective probabilities about various states of the world. But the risk aversion of investors—their tendency to value more highly assets that produce high payoffs in bad states—means that the expected payoff of an asset can also be evaluated using a set of adjusted probabilities. These adjusted probabilities are risk neutral, as by assigning greater weight to undesirable states they generate the same utility for a risk-neutral investor as for a risk-averse investor with the original probabilities. As discussed in section 2 below, these adjusted probabilities can be inferred from the prices of options contracts on the underlying asset.

Assets can, therefore, be priced by (i) evaluating the expectation of discounted payoffs using investors' best guesses of the probabilities of different states of the world occurring or, equivalently, by (ii) discounting payoffs by the risk-free rate and evaluating expectations using a set of adjusted probabilities. If there are  $S$  possible future states of the world, indexed by  $s = 1, 2, 3, \dots, S$ , then the expected discounted return of an asset can be expressed either as the sum of the discounted returns in each state, weighted by investors' subjective probability of the state occurring,

$$1 = \mathbb{E}_t(m_{t+1} \cdot R_{t+1}) = \sum_{s=1}^S m_{t+1}(s) \cdot R_{t+1}(s) \cdot \pi_{t+1}(s), \quad (7)$$

---

<sup>3</sup>This is a standard result in asset pricing. See Cochrane (2001) for a detailed exposition. Asset pricing models that employ these restrictions do, however, significantly underestimate the risk premia observed in practice due to the low volatility of consumption. Models with less restrictive utility functions and, hence, stochastic discount factors that depend on a broader set of variables may help to reconcile such anomalies (see, for example, Barberis, Huang, and Santos 2001).

or in terms of risk-neutral probabilities ( $\pi_{t+1}^*(s)$ ), discounted with the risk-free interest rate,

$$1 = \mathbb{E}_t(m_{t+1}) \cdot \mathbb{E}_t^*(R_{t+1}) = \sum_{s=1}^S \frac{1}{R_{t+1}^f} \cdot R_{t+1}(s) \cdot \pi_{t+1}^*(s). \quad (8)$$

Taken together, equations (7) and (8) imply that the ratio of the risk-neutral to subjective probabilities is proportional to the stochastic discount factor, where the constant of proportionality is given by the gross risk-free rate of return, i.e.,

$$\frac{\pi_{t+1}^*(s)}{\pi_{t+1}(s)} = m_{t+1}(s) \cdot R_{t+1}^f. \quad (9)$$

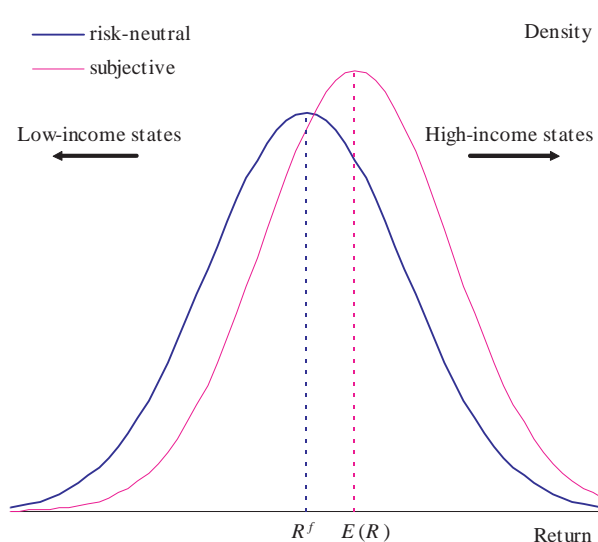
Note that the risk-neutral probability distribution is *pessimistic* in the sense that it assigns excessive probability to low-income states and too little probability to high-income states. The mean of the risk-neutral density is given by  $R_{t+1}^f = 1/\mathbb{E}_t(m_{t+1})$ , whereas the mean of the subjective density is given by equation (2). The difference between the two means is therefore the risk premium.

Investors' risk aversion also enters the risk-neutral probabilities. Since risk-averse investors value additional income more highly in poor states of the world, low-income states receive an increased weight when computing the expected return of an asset using the risk-neutral asset pricing relationship. When the marginal utility of consumption is high in state  $s$ , the risk-neutral probability is greater than the true probability and vice versa. Figure 2 provides a stylized illustration of the two probability distributions.

An increase in the ratio between the risk-neutral and subjective probabilities may therefore reflect either an increase in risk aversion or changes in other state variables that increase the marginal utility of consumption. As we have seen, the willingness of the investor to pay for insurance across such states—the investor's risk appetite—depends on the variance of the stochastic discount factor across states of the world. It follows from equations (5) and (9) that

$$\lambda_t = \frac{1}{R_{t+1}^f} \cdot \text{var} \left( \frac{\pi_{t+1}^*(s)}{\pi_{t+1}(s)} \right) \quad (10)$$

is a measure of risk appetite, once the two probability densities over future returns are derived.

**Figure 2. Risk-Neutral and Subjective PDFs**

## 2. Estimating Risk Appetite

Our analysis suggests that a measure of risk appetite may be derived by computing the variation in the ratio of risk-neutral to subjective probabilities used by investors in evaluating the expected payoff of an asset. This requires estimating two probability density functions over future returns—one risk-neutral distribution and one subjective distribution—on an index such as the S&P 500. To generate a time series for risk appetite, these distributions are estimated every three months, at the end of each quarter. As the return forecasts for the end of a particular quarter are made at the end of the previous quarter, the corresponding estimate of risk appetite would also be for the previous quarter. In what follows, we outline the approach used in estimating these distributions.

### 2.1 Risk-Neutral Densities

Option prices offer a forward-looking guide to the likelihood investors attach to future values of asset prices. But it is only a guide, because the price that investors will pay for an option depends both on their



subjective beliefs about the relative likelihoods of returns having particular future values and on their attitude to risk. If investors were neutral toward risk, however, option prices would only reflect expectations about returns. So, by comparing options with different strike prices, we can infer the risk-neutral probabilities attached by market participants to an asset being within a range of possible prices at some future date. Indeed, the whole risk-neutral density function can be inferred from the prices of marketed options using the no-arbitrage argument of Breeden and Litzenberger (1978), who demonstrate that the density function is the second derivative of the option price with respect to the option strike.

We use risk-neutral density functions for the S&P 500 index constructed by the Bank of England. These are estimated by a two-step procedure. The first step is to estimate a call price function, which shows how the prices of call options with identical maturities vary as strike prices change. This is achieved by applying a cubic-spline interpolation to the available data on pairs of call and strike prices. For more robust results, the interpolation is actually applied to transformations of these prices (see Clews, Panigirtzoglou, and Proudman [2000] for details) and the resulting function is converted back into a smooth and continuous relationship between call and strike prices. The second step is then to twice differentiate the resulting call price function. As demonstrated by Breeden and Litzenberger, this delivers the density function of the underlying asset based on the assumption that investors are risk neutral.

## 2.2 *Subjective Densities*

Estimation of the subjective probability distributions of returns follows the approach of Hayes, Panigirtzoglou, and Shin (2003).<sup>4</sup> It is

---

<sup>4</sup>An alternative approach is suggested by Bliss and Panigirtzoglou (2004), who estimate the subjective probability by hypothesizing a specific utility function for a representative agent and then using it to convert the estimated risk-neutral density function into a subjective density using the method suggested by Ait-Sahalia and Lo (2000). As Bliss and Panigirtzoglou observe, knowledge of any two of three functions—the risk-neutral density, the subjective density, and the utility function—allows the third to be inferred. So it is not immediately obvious whether this alternative is superior to the approach suggested by Hayes, Panigirtzoglou, and Shin (2003).

based on the following threshold-GARCH model of returns on the S&P 500 index,  $r_t$ :

$$\begin{aligned} r_t &= c + x'_{t-1}\theta + \delta\sigma_t + \varepsilon_t, \\ \sigma_t^2 &= \omega + y'_{t-1}\lambda + \alpha\varepsilon_{t-1}^2 + \gamma\varepsilon_{t-1}^2 d_{t-1} + \beta\sigma_{t-1}^2 \\ d_t &= 1 \quad \text{if } \varepsilon_t > 0, \text{ and } 0 \text{ otherwise,} \end{aligned}$$

where  $x_t$  and  $y_t$  are vectors of explanatory variables and  $\sigma_t^2$  is the variance of the residuals,  $\varepsilon_t$ .<sup>5</sup>

The shape of the subjective density of returns is equated to the shape of the density of the standardized residuals,  $\varepsilon_t/\sigma_t$ . But to construct the precise subjective density of one-quarter-ahead returns, the variance of the density of the standardized residuals is multiplied by the forecast conditional variance,  $\sigma_{t+1}$ , and the mean of the resulting density is set equal to a particular value. In principle, this value could be the forecast conditional mean,  $r_{t+1}$ , but in practice this occasionally implies that the mean of the subjective density is smaller than the mean of the risk-neutral density, i.e., that the risk premium is negative, which seems implausible for an equity index. Instead, we locate each subjective density such that the difference between its mean and the mean of the corresponding risk-neutral density is equal to the value of the equity risk premium implied by the Bank of England's three-stage dividend discount model.<sup>6</sup>

The threshold-GARCH model is initially estimated using quarterly data from 1920:Q1 to 1983:Q1. The fitted model is then used to forecast the conditional variance in 1983:Q2, which, as noted above, is used to construct the subjective density of returns in 1983:Q2. The model is then reestimated using data from 1920:Q1 to 1983:Q2 and the new model is used to forecast the conditional variance and hence construct the subjective return density in 1983:Q3, and so on.

<sup>5</sup>Hence, the modeled variance of returns depends on previous errors in modeling the level of returns. So, extreme returns that are not fully captured by the GARCH model would generate large residuals, and these would affect the subsequent modeled variance of return via the parameter  $\beta_4$ . Furthermore, extreme returns can have differential effects on the modeled variance of returns depending on whether they are extremely high or extremely low and hence whether residuals are positive or negative. The scale of the difference is governed by the parameter  $\beta_5$ . Finally, the variance of returns is postulated to exhibit some persistence according to  $\beta_6$ .

<sup>6</sup>See Panigirtzoglou and Scammell (2002).

Variables included in  $x_t$  and  $y_t$  were selected by adopting a general-to-specific modeling approach. From an initial list comprised of the natural logarithm of the dividend yield on the S&P 500, the spread between the yield on BBB- and AAA-rated U.S. corporate bonds, the yield on three-month U.S. Treasury bills, the term spread between the yields on ten-year U.S. government bonds and three-month U.S. Treasury bills, the rate of commodity price growth according to the Commodity Research Bureau, U.S. consumer price inflation, and the rate of unemployment in the United States,<sup>7</sup> the natural logarithm of the dividend yield was selected as the only variable to include in  $x_t$ , while  $y_t$  was selected to be empty. Variables were selected by deleting any found to be insignificant in the most general specification and reestimating the model until only statistically significant variables remained. This choice of variables also optimized the Hannan-Quinn information criterion.

The parameter estimates for the preferred model, estimated over 1920:Q1 to 1983:Q1, are reported in the first column of table 1. These estimates are quite stable as the sample period is lengthened. The positive coefficient associated with the logarithm of the dividend yield implies that returns tend to be low when prices are high relative to dividends. This generates a degree of mean reversion in the dividend yield that is consistent with the findings of empirical finance. Returns are also found to vary positively with their standard deviation. This is consistent with theoretical models in which risk and expected returns are positively associated.

The conditional variance equation generates three further features of equity returns that are commonly found in empirical work: fat tails, negative skewness, and volatility clustering. The ARCH term ( $\varepsilon_{t-1}^2$ ), which has a positive coefficient, means that a significant shock to returns will boost the conditional variance, so that extreme returns are more likely to follow an initial extreme return than an initial moderate return. This increases the thickness of tails in the distribution of returns. The threshold-ARCH term ( $\varepsilon_{t-1}^2 d_{t-1}$ ), which also has a positive coefficient, implies that negative shocks are more

---

<sup>7</sup>The choice of variables was motivated by the literature on equity return predictability, e.g., Lamont (1998) and Kothari and Shanken (1992). Also, variables such as inflation, the unemployment rate, and commodity price growth were included to capture potential business cycle effects on returns and their variability (see, for example, Chen, Roll, and Ross 1986).

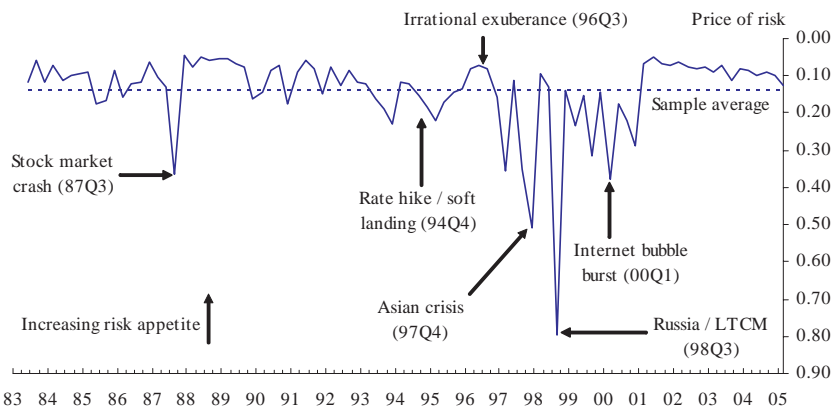
**Table 1. Coefficient Estimates of GARCH Model**

	1920:1 to 1983:1	1920:1 to 1983:1	1920:1 to 1983:1	1945:4 to 1983:1
$c$	-0.0702 (0.0302)	-0.0241 (0.0172)	-0.0782 (0.0282)	-0.0475 (0.0356)
$\theta$	0.0436 (0.0209)		0.0419 (0.0186)	0.0314 (0.0235)
$\delta$	0.3989 (0.2041)	0.4820 (0.2077)	0.4261 (0.1949)	0.2531 (0.4113)
$\omega$	0.0017 (0.0007)	0.0020 (0.0007)	0.0014 (0.0006)	0.0016 (0.0006)
$\alpha$	0.0997 (0.0813)	0.0803 (0.0888)	0.2911 (0.0913)	-0.0481 (0.0955)
$\gamma$	0.2836 (0.1313)	0.3224 (0.1403)		0.3156 (0.1303)
$\beta$	0.5611 (0.1329)	0.5355 (0.1465)	0.5755 (0.1186)	0.6120 (0.1555)
<b>Note:</b> Standard errors are shown in parentheses.				

likely to be followed by high volatility than positive shocks. This generates negative skewness. Finally, the GARCH term ( $\beta_6\sigma_{t-1}^2$ ), with its positive coefficient, generates persistence in volatility, resulting in clusters of high and low volatility.

### 2.3 Comparing the Two Densities

We plot histograms of the estimated risk-neutral and subjective densities. For each bin of the histograms, we compute the ratio of  $\pi^*/\pi$ , as required by equation (10). Due to inaccuracy of estimation, however, this ratio is sometimes spuriously high in the tails of the histograms. Therefore, any bins for which  $\pi^*/\pi > 10$  are dropped from the histograms, which are subsequently rescaled so that the probabilities of the various feasible returns continue to sum to unity. Finally, risk appetite is computed in accordance with equation (10), with the yield on three-month U.S. Treasury bills serving as a proxy for the risk-free rate.

**Figure 3. Estimated Risk Appetite**

#### 2.4 The “Variance” Measure

Figure 3 shows the quarterly time series of risk appetite from our estimation procedure. The illustrated series fluctuates close to its average for most of the time, but has occasional sharp downward movements. The sharp downward movements coincide with the 1987 stock market crash, the Asian financial crisis, the Russian/LTCM crisis, and the Internet stock crash. The series suggests that investors’ risk appetite is likely to be fairly stable during “tranquil” periods, but move sharply in response to exogenous shocks. More recently, investors’ appetite for risk has been strong, above the sample average and at levels comparable to those of 1996 when Alan Greenspan spoke of irrational exuberance. Of course, the true path of investors’ risk appetite remains unobserved, but the behavior of the measure during the period in question (1983–2005) seems plausible.<sup>8</sup>

### 3. Comparison with Existing Approaches

A number of recent papers have also attempted to measure market sentiment. A first approach is based on changes in excess returns. Equation (5) showed how the excess return required by investors to

<sup>8</sup>The appendix investigates the robustness of the risk appetite measure to changes in the assumptions made in its derivation.

hold an asset depends on the level of risk inherent in the asset and the risk appetite of the investor. Kumar and Persaud (2002) propose a measure of risk aversion based on the distribution of excess returns across assets. Their hypothesis is that when risk appetite increases, excess returns of very risky assets increase by more than for less-risky assets. In contrast, changes in the overall level of risk across assets should not have a differential impact on expected returns. Thus, the degree of correlation between changes in excess returns and the level of risk across a number of assets should indicate any change in risk appetite.<sup>9</sup>

There are a number of difficulties with this measure, however. First, the measure only indicates changes in risk aversion and does not suggest what its level might be. Second, the measure does not give an indication of the magnitude of the change in risk aversion. The rank correlation is theoretically unity when risk aversion is driving returns and zero when changing risk is driving returns. And finally, a rank correlation may be detected even when risk aversion is constant, if the level of risk associated with different assets changes to differing degrees. For example, if the volatility of the market return increased, this would increase the risk of some assets more than others and lead to a rank correlation.

A second approach, emphasized by Tarashev, Tsatsaronis, and Karampatos (2003) and Hayes, Panigirtzoglou, and Shin (2003), focuses on a comparison of the risk-neutral and subjective probability densities.<sup>10</sup> They interpret the ratio on the left-hand side of equation (9), evaluated at a particular percentile, as an indicator of risk aversion. As we have argued, however, the stochastic discount factor generally reflects rather more than just investor preferences. So movements in the probability ratio over time are more likely to reflect factors other than risk aversion. Recognizing this shortcoming,

---

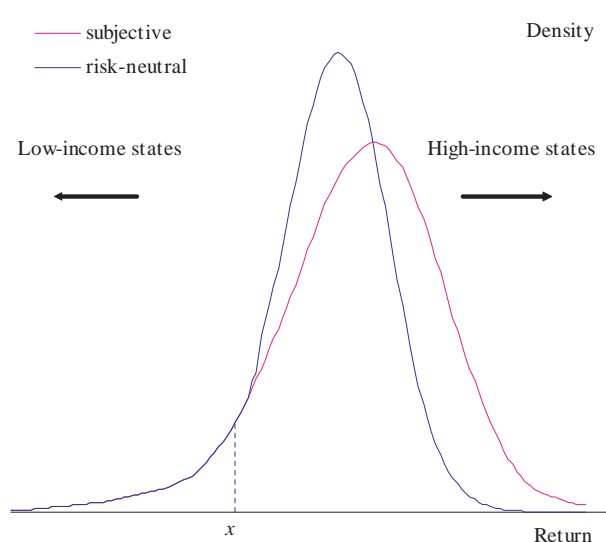
<sup>9</sup>See Misina (2003) and Pericoli and Sbracia (2004) for a reconciliation of the Kumar and Persaud measure with the general asset pricing framework outlined above.

<sup>10</sup>See also Scheicher (2003). Jackwerth (2000) also uses the probability ratio to obtain a function for risk aversion that can be computed from option contracts on the market portfolio. But his approach has two drawbacks. First, the risk aversion function can take on negative values in some states of the world, suggesting that risk aversion may (on occasion) increase with increasing wealth. And second, the risk aversion schedule does not allow a measure of market sentiment to be readily tracked over time.

Hayes, Panigirtzoglou, and Shin suggest that movements in the ratio might reflect investors' concerns about liquidity. Their hypothesis is that investors discount asset returns less heavily when their wealth is illiquid because it is more difficult to support consumption from retained wealth in such circumstances. They suggest that the importance of illiquidity in the stochastic discount factor is greatest in bad states of the world that are characterized by low asset returns. This is supported by the fact that, in such states, there is a positive relationship between implied volatilities (which tend to increase when market liquidity falls) and the estimated probability ratio. But in other states of the world, a better indication of risk aversion may be obtained since the liquidity factor is less likely to be important.

A further drawback of such an approach is that, by estimating the stochastic discount factor at a particular percentile, a "ratio" measure can misrepresent investors' overall attitude to risk. By contrast, our "variance" measure uses estimates of the stochastic discount factor across many states of the world, in which asset returns differ. If the subjective and risk-neutral distributions differ in shape markedly, then using all the information contained in the distributions is likely to offer a more reliable indicator of sentiment. For example, a ratio measure evaluated at a point like  $x$  in figure 4 would suggest that investors were risk neutral, as the tails of the risk-neutral and subjective densities coincide. As the densities diverge away from the left tail, however, the variance measure would suggest that investors disliked risk.

An approach that is very close in spirit to our own is that of Bollerslev, Gibson, and Zhou (2004), who essentially compare estimates of the standard deviations (or volatilities) of the risk-neutral distribution and the subjective distribution, rather than the whole distributions. The difference between the two standard deviations reflects a "volatility risk premium." The higher the risk appetite, the smaller the degree to which implied (risk-neutral) volatilities derived from option prices will exceed realized (subjective) volatilities. An advantage of their approach is in its use of model-free volatilities. The computation of implied volatility does not rely on the accuracy of the Black-Scholes model, for example, while the subjective volatility is computed using high-frequency historical data without imposing a GARCH model. A further advantage is that the authors are able to relate their measure of risk appetite to macroeconomic factors. One

**Figure 4. Importance of Using Whole Densities**

of their key findings is that risk appetite appears to increase with industrial production, which supports the notion advanced earlier that risk appetite is positive related to the macroeconomic environment that investors inhabit. A potential disadvantage of this study, however, is that by focusing only on the standard deviations of the risk-neutral and subjective distributions, and ignoring their higher moments, an incomplete picture of risk appetite may be obtained.

A final approach to measuring risk appetite relies on cross-border portfolio flows (Froot and O'Connell 2003). By assuming that investors have constant absolute risk aversion (CARA) utility functions, the authors show that each investor's demand for a risky asset will depend on the investor's wealth, the variance of the risky asset's excess returns, the covariance of these excess returns with the excess returns to other risky assets, as well as on the risk aversion parameter. Investors are then divided into two categories: international investors, who can purchase all assets, and domestic investors, who can only purchase the asset of the market that they inhabit. Froot and O'Connell show that cross-border portfolio flows will reflect only the risk aversion of international investors *relative to* the risk aversion of domestic investors. Using data on cross-border portfolio flows, they



infer the relative measure of risk aversion in the form of an “investor confidence index.” By contrast, our approach provides an absolute measure of risk appetite, rather than a measure of relative risk aversion that relies on particular utility functions.

#### 4. Conclusion

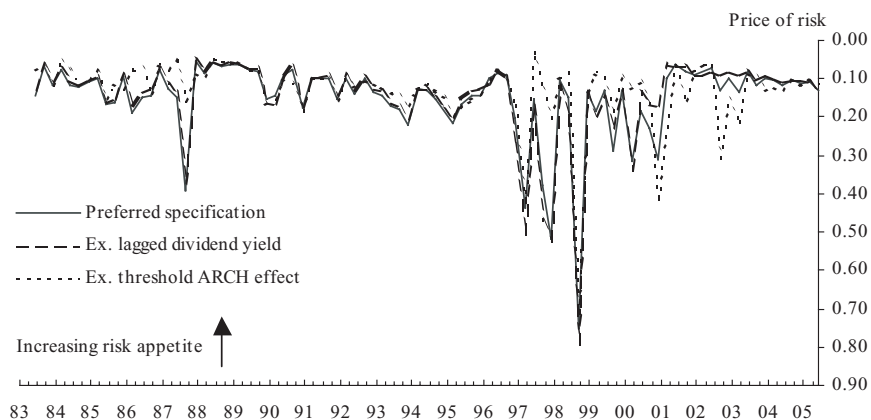
This paper has proposed a measure of market sentiment that is distinct from risk aversion and can be used to gauge how investors' appetite for risk evolves over time. The empirical analysis suggests that a measure based on the variation in the ratio of risk-neutral to subjective probabilities, derived from equity index option prices, appears to generate results that conform to intuition—the measure responds to major financial events in a plausible manner. Our approach has a number of advantages over existing measures of market sentiment. In particular, it does not rely on restrictive assumptions on investor preferences and it uses all the available information in the risk-neutral and subjective probability distributions.

#### Appendix. Robustness of Risk Appetite Measure

This appendix investigates the robustness of the risk appetite measure to changes in the assumptions made in its derivation. In particular, we investigate the effects of changing (i) the specification of the GARCH model used in constructing the subjective density of returns, (ii) the sample period over which the coefficients of the GARCH model are estimated, (iii) the estimates of the risk premium used to separate the means of the risk-neutral and subjective densities, and (iv) the cutoff point at which we reject our estimates of the ratio of the risk-neutral probability to the subjective probability.

First, we adapt our specification of the GARCH model, dropping the lagged dividend yield (i.e., setting  $\theta = 0$ ) and the threshold ARCH effect (i.e., setting  $\gamma = 0$ ) from our preferred specification. Regression results for these two modifications are respectively displayed in columns 2 and 3 of table 1, shown previously in section 2.2. Coefficient estimates are quite similar to those of our preferred specification. As a result, the risk appetite measure derived using these alternative GARCH specifications is also quite similar to our preferred measure of risk appetite, as can be seen in figure 5.

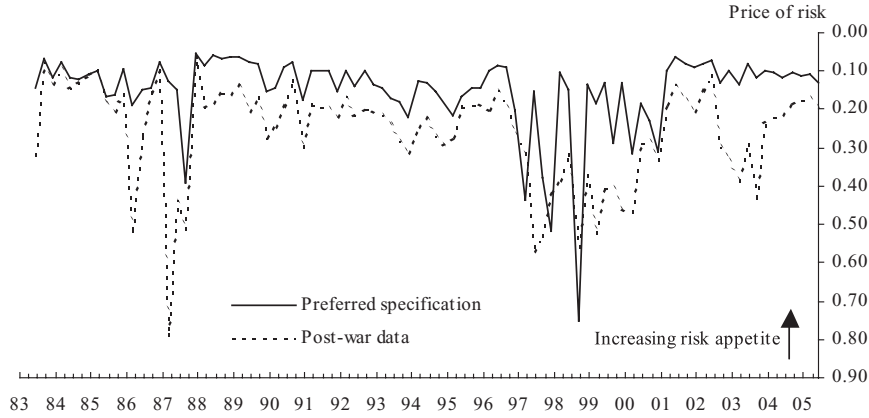
**Figure 5. Robustness of Risk-Appetite Measure to Alternative GARCH Specifications**



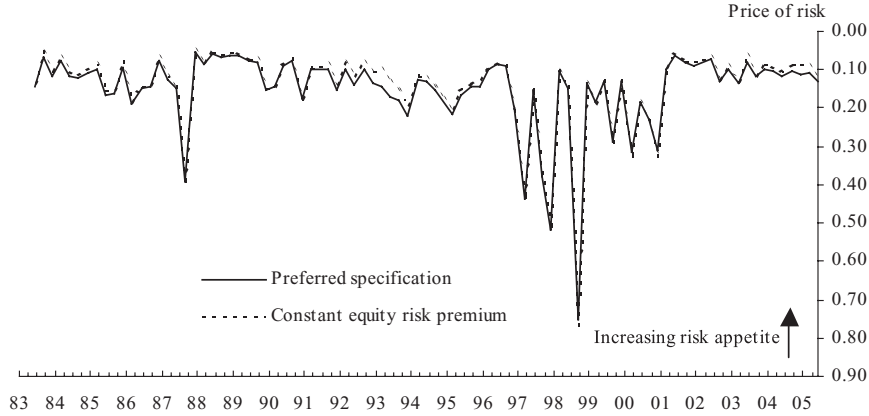
Second, we change the sample period over which we estimate the coefficients of our GARCH model to the postwar interval of 1945:Q4 to 1983:Q1. Regression results are displayed in the final column of table 1. Again, the coefficient estimates are similar to those of our preferred specification and the resulting profile of risk appetite is broadly similar to that of our preferred measure, although the correlation between the two does fall in the last few years of the sample (see figure 6). The gap between the measure based only on postwar data and our preferred measure of risk appetite that emerges at certain times is attributable to a change in the estimated shape of the subjective density. As the shape of the subjective density is constructed from the GARCH residuals, it is affected by the change of sample period. In particular, some probability mass is removed from the left tail of the subjective density due to the exclusion of the 1929 crash from the data sample. This results in higher ratios of risk-neutral probabilities to subjective probabilities and, hence, higher estimates of the price of risk.

Third, we change the estimates of the equity risk premium used to separate the means of the risk-neutral and subjective densities from the time-varying estimates obtained from the Bank of England's discounted dividend model to a constant estimate of 3.3 percent, which is taken from Taylor (2005). The latter is an estimate of the average equity risk premium on the S&P 500 since the beginning of the 1980s.

**Figure 6. Robustness of Risk-Appetite Measure to Alternative GARCH Estimation Periods**



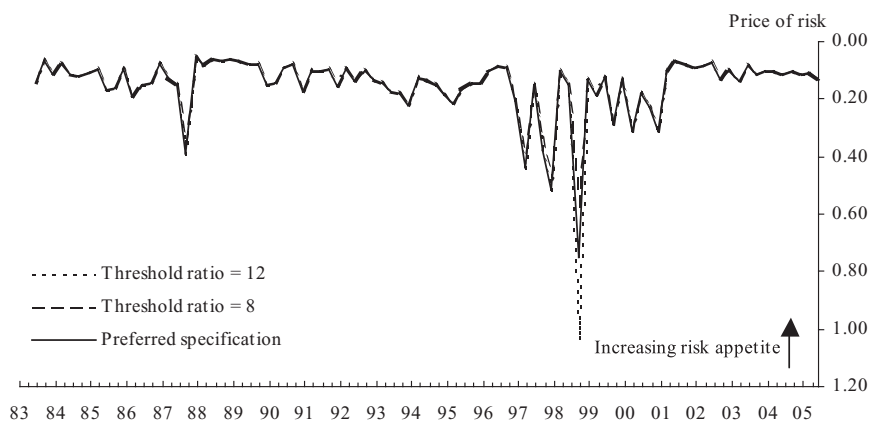
**Figure 7. Robustness of Risk-Appetite Measure to Alternative Estimates of the Equity Risk Premium**



As figure 7 indicates, the risk appetite measure is highly robust to alternative estimates of the equity risk premium, with constant and time-varying estimates producing very similar profiles.

Finally, we experiment by changing the threshold above which we reject our estimates of the ratio of risk-neutral probability to subjective probabilities. These ratios are occasionally found in the

**Figure 8. Robustness of Risk-Appetite Measure to Alternative Threshold Ratios**



tails of return distributions, where errors can result in subjective probability estimates that are very close to zero. This produces very high estimates of the ratio of risk-neutral probability to subjective probability. As the risk appetite measure is derived from variation in this ratio across the estimated probability distributions, it could potentially become driven by spuriously high ratios in the tails of the distributions. Hence, our preferred measure of risk appetite is computed by omitting any ratio estimates greater than ten from the variance calculation of equation (10). Figure 8 shows the effect of varying this threshold. The threshold appears to affect the degree to which crisis periods stand out as episodes of low risk appetite, while leaving the broad profile of the series essentially unchanged.

## References

- Ait-Sahalia, Yacine, and Andrew W. Lo. 2000. "Non-parametric Risk Management and Implied Risk Aversion." *Journal of Econometrics* 94 (1-2): 9-51.
- Barberis, Nicholas, Ming Huang, and Tano Santos. 2001. "Prospect Theory and Asset Prices." *Quarterly Journal of Economics* 116 (1): 1-53.

- Bliss, Robert R., and Nikolaos Panigirtzoglou. 2004. "Option-Implied Risk Aversion Estimates." *Journal of Finance* 59 (1): 407–46.
- Bollerslev, Tim, Michael Gibson, and Hao Zhou. 2004. "Dynamic Estimation of Volatility Risk Premia and Investor Risk Aversion from Option-Implied and Realized Volatilities." Finance and Economics Discussion Series Working Paper No. 2004-56, Board of Governors of the Federal Reserve System.
- Breeden, Douglas T., and Robert H. Litzenberger. 1978. "Prices of State-Contingent Claims Implicit in Options Prices." *Journal of Business* 51:621–51.
- Chen, Nai-Fu, Richard Roll, and Stephen A. Ross. 1986. "Economic Forces and the Stock Market." *Journal of Business* 59:383–403.
- Clews, Roger, Nikolaos Panigirtzoglou, and James Proudman. 2000. "Recent Developments in Extracting Information from Options Markets." *Bank of England Quarterly Bulletin* (February):50–60.
- Cochrane, John H. 2001. *Asset Pricing*. Princeton, NJ: Princeton University Press.
- Danthine, Jean-Pierre, and John B. Donaldson. 2005. *Intermediate Financial Theory*. Elsevier Academic Press.
- Froot, Kenneth, and Paul O'Connell. 2003. "The Risk Tolerance of International Investors." NBER Working Paper No. 10157.
- Hayes, Simon, Nikolaos Panigirtzoglou, and Hyun Song Shin. 2003. "Liquidity and Risk Appetite: Evidence from Equity Index Option Prices." Mimeo, Bank of England.
- Jackwerth, Jens. 2000. "Recovering Risk Aversion from Option Prices and Realized Returns." *Review of Financial Studies* 13 (2): 433–51.
- Kothari, S. P., and Jay Shanken. 1992. "Stock Return Variation and Expected Dividends: A Time-Series and Cross-Section Analysis." *Journal of Financial Economics* 31 (2): 177–210.
- Kumar, Manmohan, and Avinash Persaud. 2002. "Pure Contagion and Investors' Shifting Risk Appetite: Analytical Issues and Empirical Evidence." *International Finance* 5 (3): 401–26.
- Lamont, Owen. 1998. "Earnings and Expected Returns." *Journal of Finance* 53 (5): 1563–87.
- Milne, Frank. 1995. *Finance Theory and Asset Pricing*. Oxford University Press.

- Misina, Miroslav. 2003. "What Does the Risk Appetite Index Measure?" Working Paper No. 2003-23, Bank of Canada.
- Panigirtzoglou, Nikolaos, and Robert Scammell. 2002. "Analysts Earning Forecasts and Equity Valuations." *Bank of England Quarterly Bulletin* (Spring):59–66.
- Pericoli, Marcello, and Massimo Sbracia. 2004. "The CAPM and the Risk Appetite Index: Theoretical Differences and Empirical Similarities?" Mimeo, Bank of Italy.
- Scheicher, Martin. 2003. "What Drives Investor Risk Aversion? Daily Evidence from the German Equity Market." *BIS Quarterly Review* (June):67–74.
- Tarashev, Nikola, Kostas Tsatsaronis, and Dimitrios Karampatos. 2003. "Investors' Attitude Towards Risk: What Can We Learn from Options?" *BIS Quarterly Review* (June):57–66.
- Taylor, Bryan. 2005. "The Equity Risk Premium." Mimeo, Global Financial Data.