Risk appetite:

concept and measurement[1]

Prasanna Gai and Nicholas Vause, Financial Stability Assessment Division, Bank of England

This article critically reviews the analytical underpinning and measurement of investor '*risk appetite*'. We reconcile a number of different approaches with asset pricing theory, and articulate a new 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. The measure distinguishes risk appetite from *risk aversion*, and is reported in *levels* rather than *changes*. A preliminary application of the approach is assessed alongside other indicators of market sentiment and appears to yield generally plausible results.

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 of 1997 illustrates how shifting perceptions of 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 trade links and financial interconnections and coincided with claims that a decline in 'risk appetite' was an underlying reason for this contagion.

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. Fundamentally, investors dislike uncertainty about the level of consumption that will be possible in the future given 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. This *risk aversion* is part of the intrinsic make-up of the investor. It is a parameter that our theoretical priors suggest should not change markedly, or frequently, over time.

Risk appetite thus reflects somewhat more than the notion of risk aversion in microeconomics. It shifts periodically as investors respond to episodes of financial distress and macroeconomic uncertainty. In adverse circumstances, an investor will require higher expected excess returns to bear risk and risk appetite will be low. Conversely, high risk appetite will be associated with low expected excess returns.

The expected excess return required to compensate an investor for holding a risky asset is, in turn, known as the *risk premium*. The risk premium is determined partly by the inherent riskiness of the asset, and partly by the level of risk appetite. The higher the appetite for risk, the lower the risk premium. Figure 1 illustrates how these concepts are linked.

Figure 1
Summary of concepts

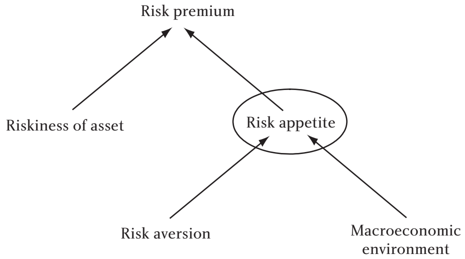

(1) We are grateful to Manmohan Kumar, Michael Metcalfe, and Kostas Tsatsaronis for generously sharing data underpinning the risk appetite measures used in this paper. We also thank Frank Milne and Hyun Shin for helpful comments and encouragement.

This article critically assesses the concept and measurement of risk appetite. It draws on asset pricing theory to distinguish risk appetite from risk aversion and risk premia. It then outlines a preferred measure, before contrasting some preliminary calculations based on this method with other indicators of risk appetite advanced in the literature.

The concept of risk appetite

Textbook treatments of asset pricing theory (eg Cochrane, 2001) state that in an efficient market, with fully rational and informed investors, the current price of an asset, p_t, should equal the expected present 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, p_{t+1}. The payoffs of an asset vary across future states of the world, such as whether a boom or recession develops. The rate at which investors discount in order to relate future purchasing power to present purchasing power varies among states of the world as well. More formally, the price of an asset can be expressed as

$$p_t = E_t\left(m_{t+1}x_{t+1}\right) \qquad [1]$$

where m_{t+1} denotes the *stochastic* discount factor that investors use to translate future payoffs in present value terms. States of the world in which high levels of future consumption are available to investors are states in which future payoffs are discounted at high rates. This is because asset payoffs are not required to support the level of consumption as much as would otherwise have been the case.

The asset pricing relationship described above can be expressed in terms of gross returns,[1] R_{t+1} , by dividing (1) by current prices. Thus

$$1 = E_t\left(m_{t+1}R_{t+1}\right). \qquad [2]$$

All assets have the same expected *discounted* return in equilibrium (of unity), even though different assets generally have different expected returns. Since both the gross return and the stochastic discount factor are random variables that depend on states of the world, we can write (2) as:

$$1 = \underbrace{E_t\left(m_{t+1}\right)E_t\left(R_{t+1}\right)}_{risk\text{-}neutral\,component} + \underbrace{cov_t\left(m_{t+1}, R_{t+1}\right)}_{risk\,adjustment}. \qquad [3]$$

If an asset were completely risk free, returns would not vary, and the gross risk-free return would be given by

$$R^f_{t+1} = 1/E_t\left(m_{t+1}\right). \qquad [4]$$

And if investors were neutral towards risk – so that they were indifferent between the particular states of the world in which asset returns were high or low – the rate of return of an asset would not be correlated with the stochastic discount factor. The covariance term in (3) would then be zero and all assets would offer the same expected rate of return, given by equation (4).

In reality, however, investors prefer to receive higher returns in some states than in others. Most commonly, investors prefer excess returns in those states of the world that deliver low consumption, since the payout in these circumstances is particularly valuable. An asset that pays a high return in good times when consumption is relatively high, but fails to pay out in bad times, has an unfavourable distribution of payoffs. In these situations, the stochastic discount factor and asset returns are negatively correlated. Investors then require a *risk premium* over and above the risk-free rate to compensate them for holding such an asset. Rearranging (3) and making use of the definition of the risk-free rate in equation (4) allows us to write the risk premium as:

$$\underbrace{E_t\left(R_{t+1}\right) - R^f_{t+1}}_{risk\ premium} = -R^f_{t+1}\,cov_t\left(m_{t+1}, R_{t+1}\right). \qquad [5]$$

The risk premium can, in turn, be decomposed into the quantity of risk associated with each asset, β^i, and the unit price of risk, λ, that is common across all assets. In particular, we can re-write (5) as:

$$E_t\left(R_{t-1}\right) - R^f_{t+1} = -\underbrace{\frac{cov_t\left(m_{t+1}, R_{t+1}\right)}{var_t\left(m_{t+1}\right)}}_{\beta_i}\underbrace{var_t\left(m_{t+1}\right)R^f_{t+1}}_{\lambda}. \qquad [6]$$

(1) The gross return on an investment is its current value expressed as a proportion of its initial value. This is equal to one *plus* the net rate of return.

The price of risk is the expected excess return that investors require to hold financial wealth at the margin. We can, therefore, define *risk appetite* – the willingness of investors to bear risk – as the inverse of the price of risk. So when investors' risk appetites fall, they require larger expected excess returns to hold risky assets.

It is immediately apparent from equation (6) that the behaviour of risk appetite hinges on the volatility of the stochastic discount factor. 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. *Risk aversion* reflects the former, since the more risk averse the investor, the more valuable is additional income in bad states of the world relative to good states, as reflected by the curvature of the investor's utility function. These are innate preferences over uncertain future prospects. As such, they 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 σ_t^2, and investors have utility functions that depend only on their consumption and capture impatience and aversion to risk, then the price of risk can be expressed as[1]

$$\lambda_t = \gamma\sigma_t, \qquad [7]$$

where γ is the coefficient of absolute *risk aversion*.[2] So a rise in γ would mean a fall in risk appetite. But risk appetite will also fall if uncertainty about consumption growth, σ_t, 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. So the shifts in market sentiment that are witnessed over time are more likely to be driven by the macroeconomic environment than the risk aversion of investors.

The discussion so far has shown how the price of an asset depends on the subjective probabilities assigned by risk-averse investors assign to future payoffs. The same price may be deduced by treating this risk-averse investor as a risk-neutral investor who attaches increased probability to bad outcomes and reduced probability to good outcomes. These adjusted probabilities are known as 'risk-neutral' probabilities and, importantly, can be readily inferred from the prices of financial options.

So we can either try (i) to estimate investors' best guesses of probabilities in order to compute expected returns, or, equivalently, (ii) consider the behaviour of a typical risk-neutral agent, discount by the risk-free rate, but evaluate the expected payoffs of an asset using a set of adjusted probabilities. Suppose there are S possible future states of the world, indexed $s=1,2,3...S$. 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,

$$1 = E_t\left(m_{t+1}R_{t+1}\right) = \sum_{s=1}^{S} m_{t+1}(s)R_{t+1}(s)\pi_{t+1}(s); \qquad [8]$$

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

$$1 = E_t\left(m_{t+1}\right)E_t\left(R_{t+1}\right) = \sum_{s=1}^{S} \frac{1}{R_{t+1}^f} R_{t+1}(s)R_{t+1}(s)\pi_{t+1}^*(s). \qquad [9]$$

Taken together, equations (8) and (9) 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, ie

$$\frac{\pi_{t+1}^*(s)}{\pi_{t+1}(s)} = m_{t+1}(s)R_{t+1}^f. \qquad [10]$$

(1) More specifically, investors have utility functions that are defined over consumption of the form: $u(c) = (c^{1-\gamma})/(1-\gamma)$. These utility functions have the convenient property that the composition of consumption is not affected by wealth, but depends only on relative prices.

(2) This is a standard result in asset pricing. For a detailed explanation, see Cochrane (2001), page 19. Asset pricing models that employ these restrictions do, however, significantly underestimate the risk premia observed in practice. This is 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 *et al* 2001).

Note that the adjusted, 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 adjusted (risk-neutral) density is given by equation (4), whereas the mean of the subjective density is given by equation (3). So the difference between the two means captures the risk premium.

The risk aversion of the representative investor 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 increasing 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 a poor state, *s*, the adjusted, risk-neutral probability is greater than the subjective probability and *vice versa*. Chart 1 provides a stylised illustration of the two probability distributions.

Chart 1
Risk-neutral and subjective probability densities

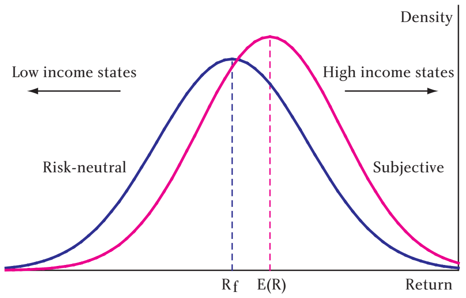

An increase in the ratio between the risk-neutral and subjective probabilities may reflect either an increase in risk aversion, or changes in other variables that increase the marginal utility of consumption. As we have seen, the willingness of the investor to pay for insurance in such states – his risk appetite – depends on the variance of the stochastic discount factor. Since equation (10) provides a measure of that stochastic discount factor across states of the world, it follows from equations (6) and (10) that risk appetite is

$$\lambda_t = \frac{1}{R_{t+1}^f} var\left(\frac{\pi_{t+1}^*(s)}{\pi_{t+1}(s)}\right). \qquad [11]$$

Measures of risk appetite

Existing measures of market sentiment fall into two categories, which can both be nested within the conceptual framework outlined above.[1] The first, typified by Kumar and Persaud (2002), is based on changes in excess returns. Equation (6) shows how the expected excess return required by investors to hold an asset depends on the level of risk inherent in the asset and the risk appetite of the investor. If the level of risk or risk appetite should change, then the required excess return should also change. The second approach, emphasised by Karampatos *et al* (2003) and Hayes *et al* (2003), focuses on a comparison of the risk-neutral and subjective probability densities. They interpret the ratio of the risk-neutral probability of future returns to subjective probability, evaluated within a certain range, as reflecting risk aversion.[2]

Kumar and Persaud (2002) propose a measure based on the distribution of excess returns across assets. Their hypothesis is that when appetite for risk 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 all assets should not have a differential impact on expected returns. So, the degree of correlation between changes in excess returns and the level of risk across a number of assets should indicate any change in the willingness to bear risk.

Kumar and Persaud implement this hypothesis by computing Spearman's rank correlation between the excess returns and volatilities of 17 currencies.[3] Excess returns are defined as the difference between actual returns and those implied by futures contracts.

To reconcile their approach with the general asset pricing framework outlined above, it is necessary to

(1) Index measures, which combine a number of variables thought to correlate with risk appetite into one indicator, are not considered here. Such measures include the Deutsche Bank Currency Risk Appetite Index, Lehman Brothers' Risk Aversion Index and JP Morgan's Liquidity and Credit Premia Index. Another popular index, the VIX, which is a weighted average of several measures of implied volatility of the US stock market, also falls under this category.

(2) See also Scheicher (2003). Karampatos *et al* (2003) argue that changes in the composition of investors or the introduction of mechanised trading, such as stop-loss sales, could affect the probability ratio in the same manner as changes in risk aversion. They therefore suggest that it is interpreted as a measure of *effective* risk aversion. Hayes *et al* note that the ratio measure may also reflect a willingness to provide liquidity to the market.

(3) The currencies are those of Argentina, Australia, Canada, the Czech Republic, the euro area, Hong Kong, Japan, Mexico, New Zealand, Norway, Poland, Singapore, South Africa, Sweden, Switzerland, Taiwan and the United Kingdom vis-à-vis the US dollar.

make the simplifying assumption that asset returns are normally distributed, and that investors have exponential utility.[1] This gives rise to the Capital Asset Pricing Model, or CAPM,

$$E_t\left(R_{t+1}\right)-R_{t+1}^f=\gamma\, cov_t\left(R_{t+1},R_{t+1}^m\right) \qquad [12]$$

where γ is the coefficient of absolute risk aversion in the investors' utility function and R^m is the return on the market portfolio, ie the return on all assets in the market portfolio, weighted according to their importance in the index, α_i. In other words,

$$R_{t+1}^m=\sum_i \alpha_i R_{t+1}^i. \qquad [13]$$

Combining equations (12) and (13), and assuming that asset returns are independently distributed, the changes in excess returns when risk aversion and asset volatility increase are given by the derivatives

$$\frac{\partial\left[E_t\left(R_{t+1}\right)-R_{t+1}^f\right]}{\partial\gamma}=\sigma_i^2\alpha_i, \qquad [14]$$

and

$$\frac{\partial\left[E_t\left(R_{t+1}\right)-R_{t+1}^f\right]}{\partial\sigma_i^2}=\gamma\alpha_i. \qquad [15]$$

Equation (14) shows that an increase in risk aversion will increase expected excess returns according to the volatility of the asset's return. This is higher for riskier assets than for less-risky assets. In contrast, equation (15) shows that changes in asset-specific risk will have a uniform effect on expected excess returns, given by the risk aversion parameter. Both responses are also influenced by the weights of assets in the market portfolio. But, there is no strong reason to expect a relationship between the riskiness of an asset and its weight in the market portfolio. So any correlation between excess returns and asset riskiness can be attributed to changes in risk aversion. The Kumar and Persaud technique therefore detects risk aversion, rather than risk appetite in the sense defined in this paper.

Several other issues also need to be borne in mind when interpreting the Kumar and Persaud indicator. First, it 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, the rank correlation may be non-zero even when risk aversion is constant, if the level of risk associated with different assets changes to differing degrees.

Karampatos *et al* (2003) and Hayes *et al* (2003) interpret the ratio of the risk-neutral to subjective probabilities on the left-hand side of equation (10) as an indicator of risk aversion. As we have argued, however, the stochastic discount factor, which features on the right-hand side of equation (10), generally reflects more than just investor preferences. So movements in the probability ratio over time are more likely to reflect factors other than risk aversion.

Hayes *et al* argue that one such factor may be the liquidity of investors' wealth. 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 an illiquidity factor in the stochastic discount factor is at its greatest in bad states of the world that are characterised 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. In other states of the world, however, a better indication of risk aversion may be obtained, since the liquidity factor is less likely to be important.

Jackwerth (2000) also uses the probability ratio to compute risk aversion. In his model, the representative agent holds the market portfolio. This means that any state of the world that generates a particular market return also generates a particular discount factor and, hence, a particular risk-neutral probability. His analysis points to a risk aversion function, across states, of the form

$$\frac{\pi'(s)}{\pi(s)}=\frac{\pi^{*}{}'(s)}{\pi^{*}(s)}, \qquad [16]$$

(1) The exponential utility function, $u(c) = e^{-\gamma c}$, has the advantage of allowing the demand for risky assets to be linear in expected returns. See Misina (2003) for a detailed attempt to reconcile the Kumar and Persaud measure with asset pricing theory.

which can be computed from option contracts on the market portfolio.

The drawbacks of this approach are twofold. First, the risk aversion function can take on negative values in some states of the world. In other words, investors are, on occasion, risk loving. This stands in contrast to our normal priors, which suggest that investors become more attracted to gambles as their wealth increases rather than less. Second, a risk aversion schedule does not offer a measure of market sentiment that can readily be tracked over time. Parts of the risk aversion schedule can rise, while others can fall. So it is difficult to gauge whether risk appetite has increased or not.

The discussion so far suggests that a measure of risk appetite based on the variance of the probability ratio (equation 11) is appealing, for several reasons. The measure is more commensurate with investors' willingness to pay for risk, rather than their aversion to risk. The focus on the variability of the stochastic discount factor also permits time series estimation in a way that allows changes in the absolute level of risk appetite to be assessed.[1] But the measure does rely on statistical methods to model the risk-neutral and subjective densities, and so may suffer from the same problems of measurement error encountered by studies that estimate a 'ratio' measure (equation 10).[2]

In contrast to 'ratio' measures, however, the 'variance' measure uses estimates of the stochastic discount factor across many different states of the world, in which asset returns differ. Risk appetite is a summary measure of investors' attitudes to payoffs across many different states of the world. By estimating the stochastic discount factor at only one return level, a 'ratio' measure could misrepresent investors' overall attitude to risk. If the two distributions differ in shape markedly, using all the information in the distributions is likely to offer a more reliable indicator of sentiment.

Chart 2 illustrates this point with an example. A ratio measure evaluated at x would suggest that investors were risk neutral, as the left tails of the risk-neutral and subjective probability densities coincide. As the densities diverge away from the left tail, however, the variance measure would suggest that investors dislike risk. As we have seen, risk-neutral distributions assign higher probabilities to lower returns than the subjective beliefs of investors. But the utility of risk-averse investors is driven to a greater extent by low returns relative to high returns than if they were risk neutral. As equations (8) and (9) indicate, this is how the two distributions correspond to the same asset and why both distributions imply exactly the same price.

Chart 2
Importance of using whole distributions

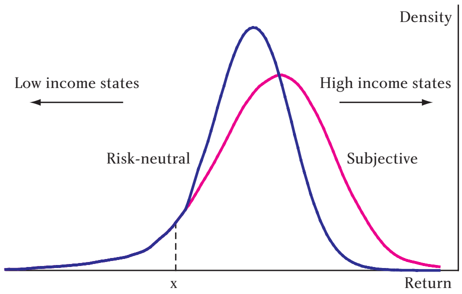

Empirical estimates of risk appetite

We now consider how various measures of risk appetite perform in practice. Specifically, we compare a measure of risk appetite computed according to equation (11) against a 'ratio' measure based on equation (10). The performance of other measures in accurately gauging market sentiment is also reviewed.

The 'variance' measure of equation (11) and the ratio measure of equation (10) are both computed by estimating probability density functions for future returns – one risk-neutral and one subjective distribution – on the S&P 500 index. To produce a time series of 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.

The risk-neutral density function is estimated using three-month option prices (see Breeden and

(1) Froot and O'Connell (2003) present an alternate measure of risk appetite that relies on cross-border portfolio flows. They obtain an index that compares changes in the risk appetite of cross-border investors with changes in the risk appetite of domestic investors. But, unlike the measure implied by equation (11), it is a relative rather than absolute measure.

(2) Note, however, that unlike the ratio measure, the variance-based measure does not distinguish between cases where investors suddenly switch from a 'normal' risk-averse state to a risk-loving state – perhaps due to a marked shift in the composition of investors. Such a situation, however, is extremely unlikely and is not observed in the data.

Litzenberger, 1978; Clews *et al*, 2000). Option prices provide us with a forward-looking guide to the likelihood the market attaches to future values of asset prices. So by comparing options with different state 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. As equation (9) suggests, this can be done taking the option prices as given and applying a known risk-free rate (such as the US Treasury bill yield) as the discount factor.

In order to determine the 'true' subjective density function, we need to estimate its overall shape, along with the mean and variance. We do this by fitting a backward-looking statistical model to historical equity returns and using that model to forecast three months into the future.[1] The essential features of the return distribution of equities are fat tails and negative skewness.[2] In order to capture these features, we adopt the simplest method possible. Specifically, we model the distribution of equity returns, r_t, as an asymmetric GARCH model of the form

$$r_t = c + \varepsilon_t \text{ where } \varepsilon_t \sim N\left(0, \sigma_t^2\right), \quad [17]$$

and

$$\sigma_t^2 = \omega + \alpha\varepsilon_{t-1}^2 + \gamma\varepsilon_{t-1}^2 D_{t-1}. \quad [18]$$

where a dummy variable, D_{t-1}, is included in the final term of equation (18) to generate negative skewness.

Karampatos *et al* (2003) and Hayes *et al* (2003) develop more sophisticated GARCH models that relate both the mean and variance of returns to additional financial and economic variables. But the simple approach outlined here is sufficient to identify key episodes of financial stress and allow consistent comparison across measures. Details of the GARCH model and the method used to generate the subjective density function are provided in the appendix.

Chart 3 illustrates the movement of the variance-based and ratio-based measures during the period 1990–2004. The broad similarity of the two measures reflects the similar shapes of the density functions during the period under consideration. Nevertheless, the variance measure is significantly more volatile, moving in a limited range during 'normal' times and sharply during episodes of financial stress. The large spikes during 1997 Q4 and 1998 Q3 correspond to the Asian and Russian/LTCM crises. Another notable shift in risk appetite in 2000 Q1 coincides with the sharp fall in high-tech stock prices. The pronounced spike in 1997 Q1 is somewhat harder to explain, however.

Chart 3
Variance and ratio-based measures derived from the S&P 500[(a)]

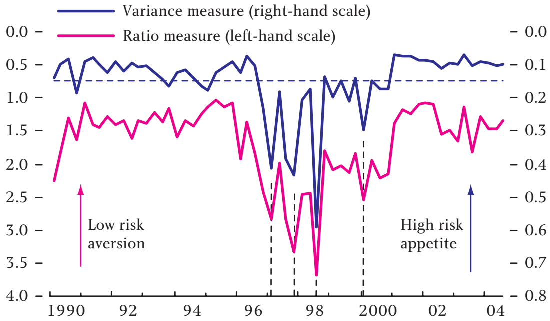

Source: Bank of England.

(a) Dotted line denotes sample average.

The variance-based measure also seems to suggest that risk appetite has been above its long-run average value in recent years. While this appears to be broadly consistent with financial stability surveillance that has pointed to a shift towards more risky investments, the measure should be interpreted cautiously. For example, it does not detect the notable reduction in the willingness to bear risk identified in the Bank's December 2002 *Review*.[3] More generally, difficulties in estimating the probability densities mean that the variance-based measure can be disproportionately sensitive to tail probabilities (see appendix).

The ratio-based measure identifies the same episodes. But since the measure implied by (10) reflects risk aversion, as opposed to risk appetite, we might expect it to be much less volatile than is the case in Chart 3. It suggests that the ratio-based measure may be reflecting other factors in addition to risk aversion,

(1) This assumes no structural change in asset markets, so future asset returns behave in the same way as in the past.

(2) See Alexander (2001) for a detailed discussion of the stylised features of return distributions.

(3) This may, however, equally reflect the localised concentration of risk during the period (within the LCFI sector).

lending credence to the liquidity hypothesis advanced by Hayes *et al* (2003).

Chart 4 compares the variance measure with the more sophisticated ratio measure of Karampatos *et al* (2003). It is immediately apparent that the Karampatos *et al* measure looks quite different from the ratio-based and variance-based measures from our illustrative GARCH. So variations in the construction of the risk-neutral and subjective densities make an important difference to estimates of risk appetite and risk aversion.

unclear whether either investors' risk appetite or risk aversion should fluctuate to such an extent. *A priori*, risk appetite is unlikely to change much during normal periods, but can be expected to shift markedly during financial crises. As discussed above, the measure is likely to reflect changing risk aversion, rather than its level. Chart 5 suggests risk aversion increased sharply in 1995 Q1, with the Mexican crisis, but also in 2000 Q3, which is more difficult to explain. And there appear to have been notable declines in investor sentiment during 1995 Q2 and 2000 Q4. The events to which these movements correspond are not readily evident.

Chart 4
Variance-based measure and the Karampatos et al indicator[(a)]

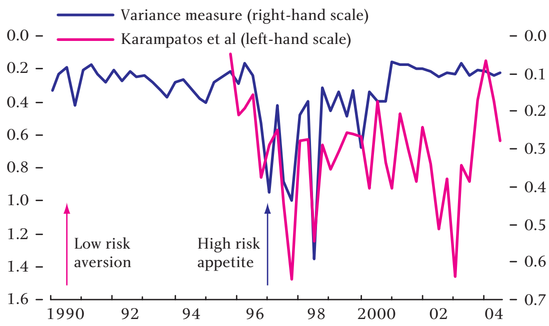

Sources: Bank of England and BIS.
(a) Dotted line denotes sample average.

Chart 5
Variance-based measure and the Kumar and Persaud indicator

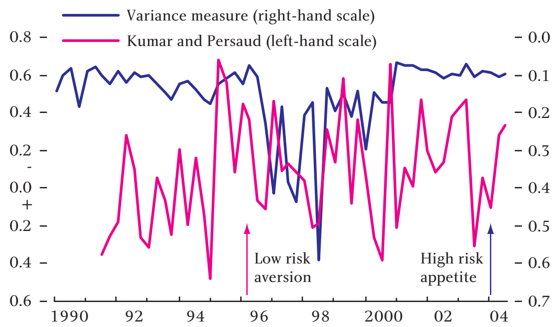

Sources: Bank of England and State Street Bank.

The Karampatos *et al* measure suggests that investor sentiment deteriorated sharply at the time of the Asian and Russian crises. It also appears to pick up the withdrawal from risk since the peak of the equity markets in 2000 reported in various financial stability reports. While the measure has pointed to a marked increase in risk appetite since 2003 Q1, some of this appears to have unwound on the most recent data. By contrast, our simple variance measure suggests that recent appetite for risk has been high but relatively stable. It is difficult to compare any further across indicators since the true subjective probability density of investors is unknown.

Chart 5 plots the variance measure against rank correlation between excess return and risk proposed by Kumar and Persaud. Excess returns are computed as actual returns minus the return implied by forward rates. And risk is taken to be the standard deviation of these excess returns over the preceding year.

The Kumar and Persaud measure is more volatile than either the variance-based or ratio measures, and it is

Chart 6 compares the variance-based measure from the simple GARCH model (used in the previous charts) against an alternative GARCH model that estimates subjective probabilities along the lines adopted by Hayes *et al*. In particular, dividend yields are used to help estimate mean returns in equation (17) and the spread between BAA- and AAA-rated corporate bonds is used to help estimate the variance of returns in equation (18). These variables might be helpful in modelling the distribution of returns as the former is a component of gross returns, while yield spreads and equity returns generally move together as they are both determined by the value of corporate assets.

The additional variables identified by Hayes *et al* are found to be statistically significant in helping to explain the statistical distribution of past returns. This suggests that there is scope for improving the measurement of subjective probabilities. Although we do not report details of the modified GARCH equation here, the similar profiles of risk appetite in Chart 6 suggest that the simple model used for illustration in this paper is reasonably robust.

Chart 6
Measures of risk appetite based on different underlying GARCH models

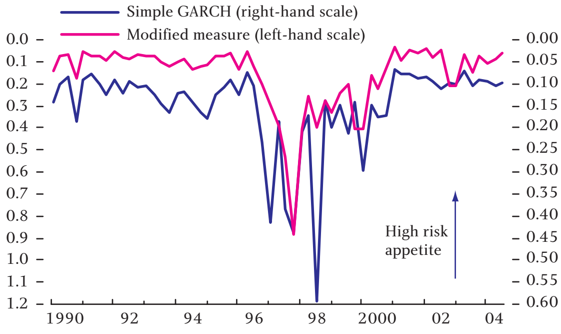

Source: Bank of England.

Conclusion

This article has reviewed the notion of risk appetite in theory and in practice. Unlike existing measures, our approach provides an indicator of market sentiment that is distinct from risk aversion and focuses on levels rather than changes. A further feature of the measure is that it uses all the information in the risk-neutral and subjective probability distributions. This may make it a better gauge of risk appetite.

The preliminary empirical analysis reported here suggests that measures of risk appetite based on this approach seem plausible. But it should be stressed that our findings are a tentative first step in the measurement of perceptions of risk across time. The role of market liquidity and changes in the composition of investors in influencing risk appetite merits further investigation. Further work is also needed to develop better estimates of investors' subjective probabilities over states of the world, and to relate better the mean and variance of asset returns to financial and economic variables.

Appendix

The parameter values of the GARCH model were estimated using quarterly data from 1928–89 (Table 1). The positive value of α implies that large deviations from the average return are more likely to follow previous large deviations. This generates volatility clustering and, hence, fat tails. A positive value for γ generates negative skewness, as D_{t-1} is an indicator variable that equals 1 when the previous quarter's return was below average and otherwise equals 0.

Three steps are required to generate a subjective density function of three-month future returns using this model. First, the shape of this distribution is given by the actual distribution (rather than the assumed Normal distribution) of the standardised residuals, ε_t/σ_t, from equation (17). Second, this distribution is scaled by a forecast of the conditional volatility obtained by rolling equation (18) forward by one quarter, ie

$$\sigma_{t+1}^2 = \omega + \alpha\varepsilon_t^2 + \gamma\varepsilon_t^2 D_{t-1} + \beta\sigma_t^2 . \qquad [19]$$

Finally, the distribution is shifted, so that its mean is equal to the mean of the risk-neutral distribution *plus* a risk premium. The estimated mean from equation (18) is disregarded as this is a very simple equation, which is unlikely to forecast accurately. The GARCH model was employed principally for the volatility forecast used in step 2. The equity risk premium was taken as the residual from a Dividend Discount Model of the S&P index level estimated by Panigirtzoglou and Scammell (2002).

Table 1
GARCH model parameter values and standard errors

	Parameter	Standard Error
χ	0.017	0.0052
ω	0.0015	0.00048
α	0.083	0.062
γ	0.34	0.12
β	0.60	0.11

The construction of the subjective density forecast is illustrated in Charts 7 and 8.

Chart 7
Probability densities of standardised residuals and return forecast errors

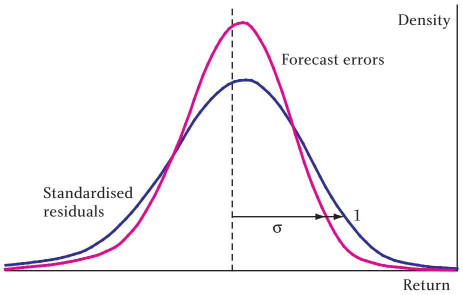

In Chart 7, the standardised residuals from equation (17) are multiplied by the square-root of the variance forecast of equation (19), evaluated at a particular point in time. This has the effect of scaling the distribution of the standardised residuals, as illustrated.

In Chart 8 the estimated subjective distribution (the scaled standardised residuals) is shifted. The estimated mean is disregarded and substituted for the mean of the risk-neutral distribution *plus* the equity risk premium from the dividend discount model. This shift moves the estimated subjective distribution from the dotted pink distribution to the solid pink distribution.

Computing the variance-based and ratio-based measures is then a simple matter of employing equations (10) and (11).[1] Three-month US Treasury bill yields are used as a proxy for the risk-free rate in equation (11).

Chart 8
Correction of the mean of the subjective return density

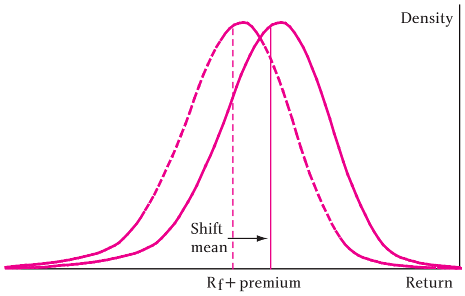

(1) Since errors in estimating both the risk-neutral and subjective probability densities can have disproportionate effects in the tails, the variance measure is computed across a wide range of returns. More specifically, we compute our simple measure from two standard deviations below the mean of the risk-neutral distribution to two standard deviations above it. While this may appear arbitrary, it has the benefit of capturing around 95% of probability mass in these densities.

References

Alexander, C (2001), *Market models*, John Wiley and Sons: Chichester, UK.

Barberis, N, Huang, M and Santos, T (2001), 'Prospect theory and asset prices', *Quarterly Journal of Economics*, 116, pages 1–53.

Breeden, D and Litzenberger, R (1978), 'Prices of state-contingent claims implicit in options prices', *Journal of Business*, 51, pages 621–651.

Clews, R, Panigirtzoglou, N and Proudman, J (2000), 'Recent developments in extracting information from options markets', *Bank of England Quarterly Bulletin*, February, pages 50–60.

Cochrane, J (2001), *Asset pricing*, Princeton University Press.

Froot, K and O'Connell, P (2003), 'The risk tolerance of international investors', *NBER Working Paper, No. 10157*.

Hayes, S, Panigirtzoglou, N and Shin, H (2003), 'Liquidity and risk appetite: evidence from equity index option prices', Bank of England mimeo.

Jackwerth, J (2000), 'Recovering risk aversion from option prices and realised returns', *Review of Financial Studies*, 13, pages 433–451.

Karampatos, D, Tarashev, N and Tsatsaronis, K (2003), 'Investors' attitude towards risk: what can we learn from options?', BIS *Quarterly Review*, June, pages 57–66.

Kumar, M and Persaud, A (2002), 'Pure contagion and investors' shifting risk appetite: analytical issues and empirical evidence', *International Finance*, 5(3), pages 401–436.

Misina, M (2003), 'What does the risk-appetite index measure?', *Bank of Canada Working Paper, no. 23*.

Panigirtzoglou, N and Scammell, R (2002), 'Analysts' earnings forecasts and equity valuations', *Bank of England Quarterly Bulletin*, February, pages 59–66.

Scheicher, M (2003), 'What drives investor risk aversion? Daily evidence from the German equity market', BIS *Quarterly Review*, June, pages 67–74.