



Working Paper no. 283

Measuring investors' risk appetite

Prasanna Gai and Nicholas Vause

November 2005

Bank of England

Measuring investors' risk appetite

*Prasanna Gai**

and

*Nicholas Vause***

Working Paper no. 283

* Bank of England.

Email: prasanna.gai@bankofengland.co.uk

** Bank of England.

Email: nicholas.vause@bankofengland.co.uk

The views expressed in this paper are those of the authors, and not necessarily those of the Bank of England. We would like to thank Alex Bowen, Damien Lynch, Paul Robinson, Hyun Shin and Peter Westaway for helpful comments and encouragement. This paper was finalised on 20 October 2005.

The Bank of England's working paper series is externally refereed.

Information on the Bank's working paper series can be found at www.bankofengland.co.uk/publications/workingpapers/index.htm.

Publications Group, Bank of England, Threadneedle Street, London, EC2R 8AH; telephone +44 (0)20 7601 4030, fax +44 (0)20 7601 3298, email mapublications@bankofengland.co.uk.

Contents

Abstract	3
Summary	4
1 Introduction	5
2 The concept of risk appetite	6
3 Estimating risk appetite	9
4 Comparison with existing approaches	13
5 Conclusion	16
Appendix: Robustness of risk-appetite measure	17
Charts and tables	19
References	25

Abstract

This paper proposes a new method for measuring investor ‘risk appetite’. Like other indicators in the literature, it is based on a comparison of risk-neutral probabilities of future returns with the corresponding subjective probabilities. The precise nature of the comparison is novel, however, and involves comparing probabilities across the full range of potential returns. Unlike other indicators, 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.

Summary

Financial market participants and policymakers frequently cite increased ‘risk appetite’ as an important driver of the recent downward trend in risk premiums and yield spreads. Risk appetite reflects investors’ willingness to hold risky assets and, as such, depends on their attitudes to risk as well as the size of other risks they carry on their balance sheets, such as that relating to employment. As a general determinant of asset prices, changes in risk appetite can generate correlation among the returns of otherwise unrelated assets. For example, a decline in risk appetite may help explain financial market contagion during the 1997-98 East Asian crisis. More recently, there have been concerns among policymakers that a build-up of risk appetite may have led to a ‘search for yield’, whereby investors demand more risky assets in pursuit of higher yields.

This paper reviews the concept of risk appetite, distinguishing it from the related notions of risk aversion and the risk premium. It suggests a precise definition of risk appetite, relating it to the price of risk – the compensation that investors require to hold a given amount of risk. This definition can be related to differences between investors’ expectations about future asset returns and those implicit in options prices. Calibration of our measure of risk appetite suggests that it fluctuates within a relatively narrow range during ‘normal’ times, but falls sharply during crises.

1 Introduction

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 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 domestic fundamentals alone and coincided with claims that a decline in ‘risk appetite’ was an underlying reason for contagion and financial instability.

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 make-up of the investor. It is a parameter that our theoretical priors suggest is unlikely to change markedly, or frequently, over time.⁽¹⁾

Risk appetite, by contrast, is likely to shift periodically as investors respond to episodes of financial distress and macroeconomic uncertainty. In adverse circumstances, an investor 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 the investor for holding that asset is the *risk premium*. Chart 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.

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

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 pay-off 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 *et al* (2003), Tarashev *et al* (2003), and Bollerslev *et al* (2004). But unlike these papers, however, we are able to extract an indicator of market sentiment that is quite distinct from risk aversion. Moreover, the 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 favourably with other measures advanced in the literature.

The paper is organised as follows. Section 2 sets out the theoretical basis of our risk-appetite measure. Section 3 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 4, and a final section concludes.

2 The concept of risk appetite

The standard treatment of asset pricing theory (eg 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 pay-offs, x_{t+1} . These pay-offs 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,

$$p_t = E_t(m_{t+1} \cdot x_{t+1}) \quad (1)$$

where x_{t+1} denotes the pay-off 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.

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,

$$1 = E_t(m_{t+1} \cdot R_{t+1}) \quad (2)$$

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

$$1 = \underbrace{E_t(m_{t+1}) \cdot E_t(R_{t+1})}_{\text{risk-neutral component}} + \underbrace{cov_t(m_{t+1}, R_{t+1})}_{\text{risk adjustment}} \quad (3)$$

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 $R_{t+1}^f = 1/E_t(m_{t+1})$, we can re-arrange to obtain the familiar expression

$$\underbrace{E_t(R_{t+1}) - R_{t+1}^f}_{\text{risk premium}} = -R_{t+1}^f cov_t(m_{t+1}, R_{t+1}) \quad (4)$$

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 discount factor. Intuitively, an asset that pays a high return in good times when the investor has a high level of consumption, but fails to pay out in bad times when the 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, ie the asset must offer a *risk premium*.

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

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

The price of risk, λ_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 an investor's risk appetite falls, they require larger expected excess returns 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 an individual’s utility function – 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 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 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 pay-offs in bad states – means that the expected pay-off of an asset can also be evaluated using a set of adjusted probabilities. These adjusted probabilities are risk-neutral in the sense that they assign greater weight to undesirable states in order to keep the expected utility of the investor from holding the asset unchanged as his risk preferences are hypothetically changed. As discussed in Section 3 below, these adjusted probabilities can be inferred from the prices of options contracts on the underlying asset.

Assets can, therefore, be priced by (a) trying to estimate investors’ best guesses of probabilities in order to compute expected returns; or, equivalently, (b) by considering the behaviour of a risk-neutral agent, discounting by the risk-free rate, and evaluating the expected pay-offs of an asset 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

(2) 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 *et al* (2001)).

state occurring,

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

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

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

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, ie

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

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. 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. The mean of the risk-neutral density is given by $R_{t+1}^f = 1/E_t(m_{t+1})$, whereas the mean of the subjective density is given by equation (2). The difference between the two means is therefore the risk premium. Chart 2 provides a stylised illustration of the two probability distributions. It then follows from equations (5) and (9) that

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

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

3 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 pay-off 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.

3.1 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 an investor will pay for an option depends both on his subjective beliefs about the relative likelihoods of returns having particular future values and on the investor's attitude to risk. If the investor were neutral towards 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 *et al* (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.

3.2 Subjective densities

Estimation of the subjective probability distributions of returns follows the approach of Hayes *et al* (2003).⁽³⁾ It is based on the following threshold-GARCH model of returns on the S&P 500 index, r_t :

$$\begin{aligned}r_t &= \beta_0 + x'_{t-1}\beta_1 + \varepsilon_t \\ \sigma_t^2 &= \beta_2 + y'_{t-1}\beta_3 + \beta_4\varepsilon_{t-1}^2 + \beta_5\varepsilon_{t-1}^2d_{t-1} + \beta_6\sigma_{t-1}^2 \\ d_t &= 1 \text{ if } \varepsilon_t > 0, \text{ and } 0 \text{ otherwise.}\end{aligned}$$

where x_t and y_t are vectors of explanatory variables and σ_t^2 is the variance of the residuals, ε_t .⁽⁴⁾

The shape of the subjective density of returns is equated to the shape of the density of the standardised residuals, ε_t/σ_t . But to construct the precise subjective density of one-quarter ahead returns, the variance of the density of the standardised residuals is multiplied by the forecast conditional variance, σ_{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, ie 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 3-stage Dividend Discount Model.⁽⁵⁾

The threshold-GARCH model is initially estimated using quarterly data from 1920 Q1 to 1983

(3) An alternative approach is suggested by Bliss and Panigirtzoglou (2004), who estimate the subjective probability by hypothesising 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 *et al*.

(4) Hence, the modelled variance of returns depends on previous errors in modelling 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 modelled variance of return via the parameter β_4 . Furthermore, extreme returns can have differential effects on the modelled 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 β_5 . Finally, the variance of returns is postulated to exhibit some persistence according to β_6 .

(5) See Panigirtzoglou and Scammell (2002).

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 re-estimated 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. Variables included in x_t and y_t were selected by adopting a general-to-specific modelling 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-rated and AAA-rated US corporate bonds; the yield on three-month US Treasury bills; the term spread between the yields on ten-year US government bonds and three-month US Treasury bills; the rate of commodity price growth according to the Commodity Research Bureau; US consumer price inflation and the rate of unemployment in the United States,⁽⁶⁾ 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 re-estimating the model until only statistically significant variables remained. This choice of variables also optimised 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 A. 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. Consistent with the findings of empirical finance, this generates mean reversion in the dividend yield.

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 (ε_{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 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

(6) The choice of variables was motivated by the literature on equity return predictability, eg 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 *et al* (1986)).

generates persistence in volatility, resulting in clusters of high and low volatility.

3.3 *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 π^*/π , 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 US Treasury bills serving as a proxy for the risk-free rate.

3.4 *The ‘variance’ measure*

Chart 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 behaviour of the measure during the period in question (1983-2005) seems plausible.⁽⁷⁾

4 **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 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,

(7) The appendix investigates the robustness of the risk appetite measure to changes in the assumptions made in its derivation.

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.⁽⁸⁾

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, emphasised by Tarashev *et al* (2003) and Hayes *et al* (2003), focuses on a comparison of the risk-neutral and subjective probability densities.⁽⁹⁾ 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. Recognising this shortcoming, Hayes *et al* 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 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. 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.

(8) 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.

(9) 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.

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 Chart 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 *et al* (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 is risk appetite, the smaller the degree to which implied (risk-neutral) volatilities derived from option prices will exceed realised (subjective) volatilities. An advantage of their approach is in its use of model-free volatilities – neither the implied nor realised volatilities rely on the accuracy of a model such as the Black-Scholes option pricing model. A potential disadvantage, however, is that by focusing only on the standard deviations of the risk-neutral and subjective distributions, 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 his 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.

5 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, appear 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 cut-off 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 (ie setting $\beta_1 = 0$) and the threshold ARCH effect (ie setting $\beta_5 = 0$) from our preferred specification. Regression results for these two modifications are respectively displayed in columns 2 and 3 of Table A. Coefficient estimates are quite similar to those of our preferred specification. As a result, the risk-appetite measure derived using these alternative GARCH specifications are also quite similar to our preferred measure of risk appetite, as can be seen in Chart 5.

Second, we change the sample period over which we estimate the coefficients of our GARCH model to the post-war interval of 1945 Q4 to 1983 Q1. Regression results are displayed in the final column of Table A. 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 Chart 6). The gap between the measure based only on post-war 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%, 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. As Chart 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 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). Chart 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.

Charts and tables

Chart 1: Relationship between risk concepts

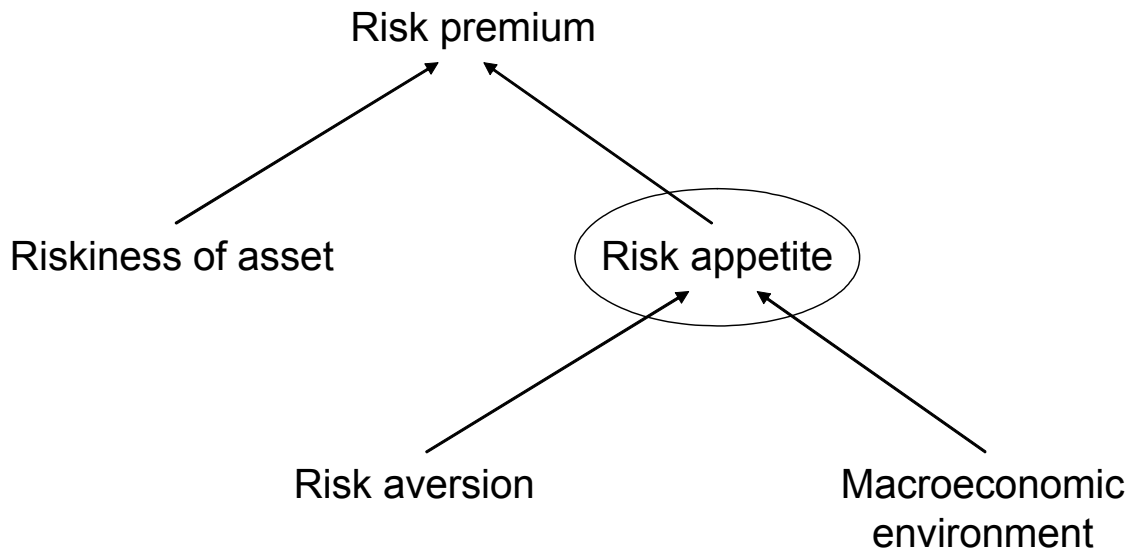


Chart 2: Risk-neutral and subjective PDFs

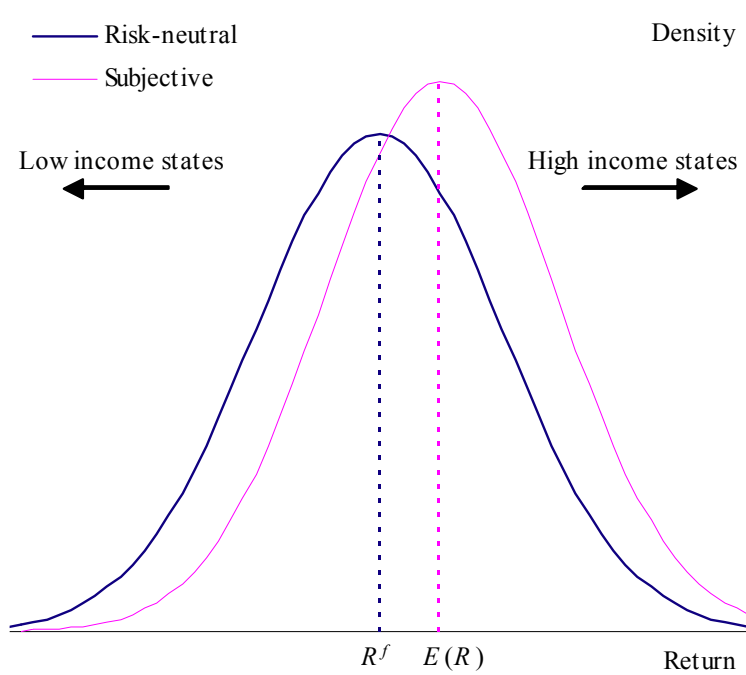


Chart 3: Estimated risk appetite




Chart 4: Importance of using whole densities




Chart 5: Robustness of risk-appetite measure to alternative GARCH specifications

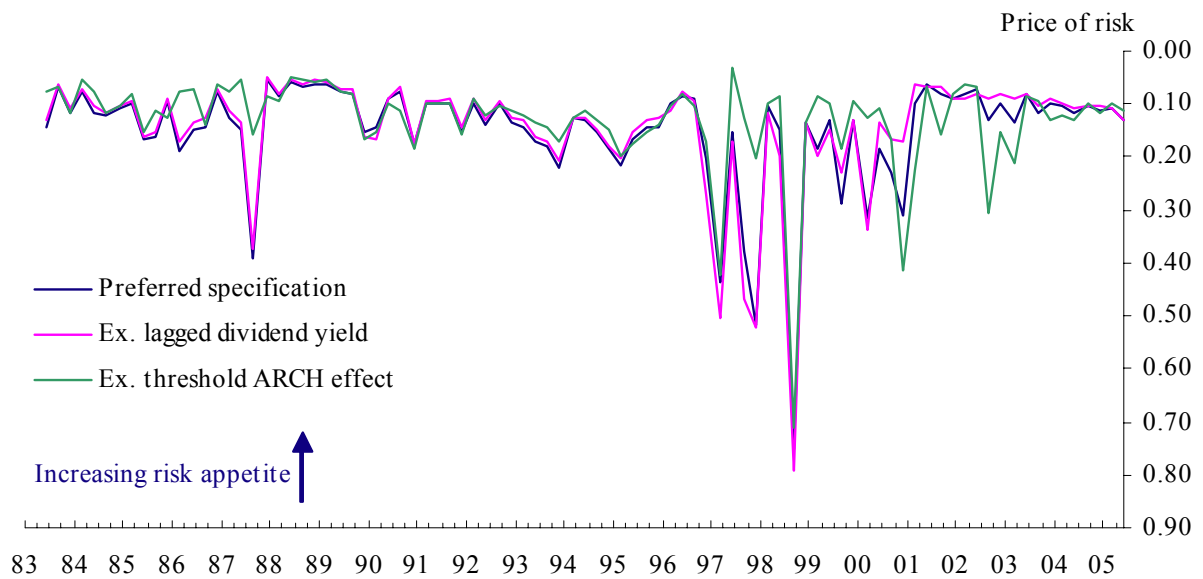


Chart 6: Robustness of risk-appetite measure to alternative GARCH estimation periods

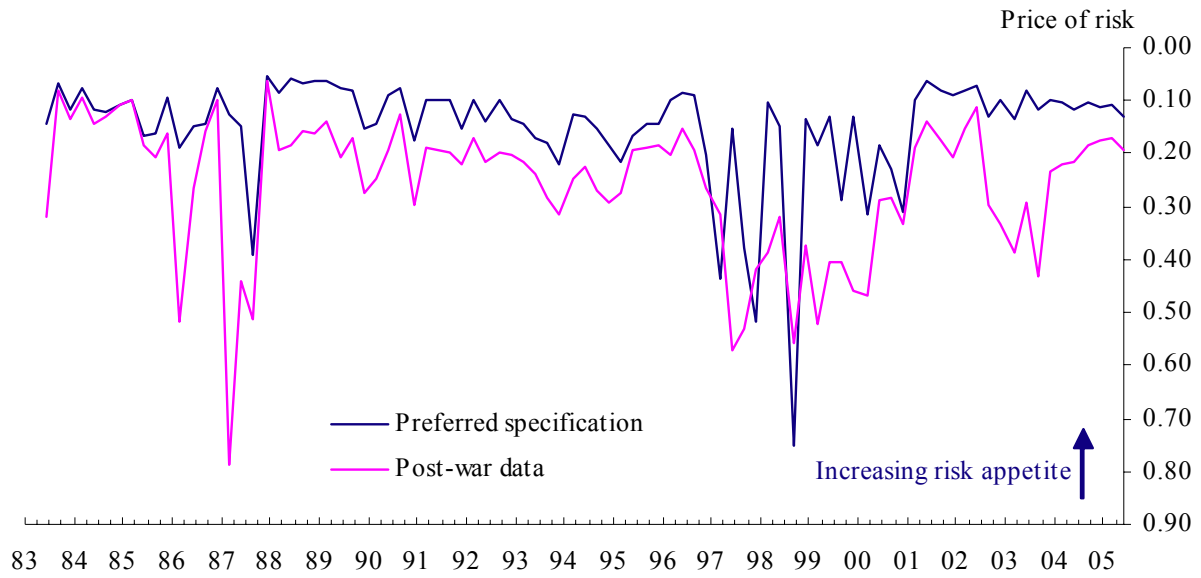


Chart 7: Robustness of risk-appetite measure to alternative equity risk premium estimates

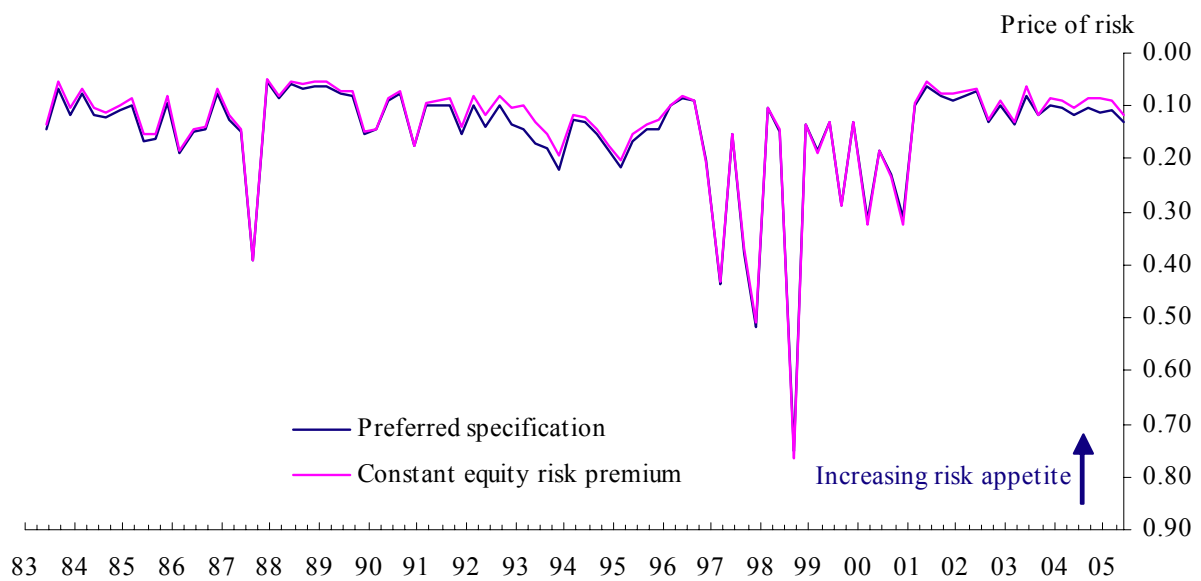


Chart 8: Robustness of risk-appetite measure to alternative threshold ratios




Table A: Coefficient estimates (* = significant at 5%, ** = 10%)

	1920:1 to 1983:1	1920:1 to 1983:1	1920:1 to 1983:1	1945:4 to 1983:1
β_0	-0.048908	0.015267**	-0.050217**	-0.048707
	-0.031992	-0.006122	-0.025616	-0.037335
β_1	0.043624**		0.046504**	0.046079*
	-0.021875		-0.018178	-0.026267
β_2	0.001428**	0.001569**	0.001116**	0.001496**
	-0.000485	-0.000502	-0.000423	-0.00079
β_3	0.099892	0.080285	0.267145**	-0.130369
	-0.064973	-0.058337	-0.073701	-0.090571
β_4	0.278501**	0.325949**		0.341162**
	-0.114845	-0.12111		-0.152778
β_5	0.614014**	0.606103**	0.634145**	0.654751**
	-0.106738	-0.112089	-0.091232	-0.232843

References

Ait-Sahalia, Y and Lo, A (2000), ‘Non-parametric risk management and implied risk aversion’, *Journal of Econometrics*, Vol. 94, pages 9-51.

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

Bliss, R and Panigirtzoglou, N (2004), ‘Option-implied risk aversion estimates’, *Journal of Finance*, Vol. 59, pages 407-46.

Bollerslev, T, Gibson, M and Zhou, H (2004), ‘Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realised volatilities’, *Finance and Economics Discussion Series Working Paper #2004-56*, Federal Reserve Board.

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

Campbell, J and Shiller, R (1988), ‘Stock prices, earnings and expected dividends’, *Journal of Finance*, July, pages 661-76.

Chen, N, Roll, R and Ross, S (1986), ‘Economic forces and the stock market’, *Journal of Business*, pages 383-403.

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

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

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

Glosten, L, Jagannathan, R and Runkle, D (1993), ‘On the relation between expected value and the volatility of the nominal excess return on stocks’, *Journal of Finance*, Vol. 48, pages 1,779-801.

Hayes, S, Panigirtzoglou, N and Shin, H S (2003), ‘Liquidity and risk appetite: evidence from equity index option prices’, *mimeo*, London School of Economics.

Jackwerth, J (2000), ‘Recovering risk aversion from option prices and realised returns’, *Review of Financial Studies*, Vol. 13, pages 433-51.

Kothari, S and Shanken, J (1992), ‘Stock return variation and expected dividends: a time-series and cross-section analysis’, *Journal of Financial Economics*, pages 117-210.

Kumar, M and Persaud, A (2002), ‘Pure contagion and investors’ shifting risk appetite: analytical issues and empirical evidence’, *International Finance*, Vol. 5, pages 401-26.

Lamont, O (1998), 'Earnings and expected returns', *Journal of Finance*, pages 1,563-87.

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

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

Pericoli, M and Sbracia, M (2004), 'The CAPM and the risk appetite index: theoretical differences and empirical similarities?', *mimeo*, Bank of Italy.

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

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

Taylor, B (2005), 'The equity risk premium', *mimeo*, Global Financial Data.