Nikola Tarashev
+41 61 280 9213
nikola.tarashev@bis.org

Kostas Tsatsaronis
+41 61 280 8082
ktsatsaronis@bis.org

Dimitrios Karampatos
+41 61 280 8324
dimitrios.karampatos@bis.org

Investors’ attitude towards risk: what can we learn from options?[1]

Market commentators often cite changes in investors’ attitude towards risk as a possible explanation for swings in asset prices. Indeed, episodes of financial turmoil coincide with anecdotal evidence of abrupt shifts in market sentiment from risk tolerance to risk avoidance. While these shifts may be potentially driven by changes in the fundamental disposition of individual investors towards risk, they are more likely to reflect the *effective risk attitude* as manifested through the behaviour of currently active investors. In particular, behaviour similar to that induced by shifts in the fundamental preferences of investors over risk and return can also reflect changes in the composition of active market players or tactical trading patterns, induced by the interaction of prevailing market conditions with institutional features. Tools that track the dynamics of investors’ willingness to take risks can lead to a better understanding of the functioning of financial markets. In particular, they can contribute not only to more effective risk management from the point of view of individual institutions, but also to improved monitoring of market conditions by policymakers.

This article constructs an indicator of investors’ effective aversion to risk. The indicator is obtained by comparing the statistical likelihood of future asset returns, which is estimated on the basis of historical patterns in spot prices, with an assessment of the same likelihood filtered through market participants’ effective risk preferences, which are derived from option prices. In particular, we argue that the *relative* size of downside risk, as assessed from the preference-weighted and the statistical vantage points, co-moves with the prevailing effective attitude of market participants towards risk. Remarkably, we find that indicators of risk attitude derived from different equity markets have a significant common component, indicating that investor sentiment transcends national boundaries.

In the next two sections we first describe and motivate the methodology and then discuss the time patterns displayed by the indicator of effective risk aversion for three equity market indices. In the last section we analyse the

[1] The views expressed in this article are those of the authors and do not necessarily reflect those of the BIS. The authors would like to thank Marian Micu for his help with computer programming.

statistical behaviour of asset prices, conditional on whether the indicator signals a high or low investor aversion to risk. The observed patterns are consistent with accounts suggesting that periods of investor retrenchment from risk-taking are also characterised by higher equity price volatility and subdued co-movement between bond and equity markets.

An indicator of investors' risk aversion

The price of an asset reflects investors' preferences with regard to possible future payoffs as well as their assessment of the likelihood of those payoffs. The incremental value to an investor of a future payoff decreases with the level of the investor's wealth. Hence, everything else constant, assets that tend to produce higher payoffs in situations when wealth is lower are valued more highly. Based on this premise, modern finance theory models asset prices as the expectation of future payoffs, calculated not on the basis of their objective *statistical* likelihood, but rather on the basis of a *preference-weighted* likelihood measure that filters statistical probabilities by investors' preferences with regard to risk.

An indicator that compares ...

Graph 1 provides an illustration of the difference between the two likelihood measures, taking as an example an investor whose only source of wealth is a single security. The red curve plots the statistical likelihood of the security's possible future payoffs. The blue curve depicts instead the assessment of payoffs from the point of view of the investor and weights the statistical probabilities according to the investor's risk preferences. This probability distribution, which is filtered by the investor's subjective preferences, assigns greater weight to lower payoffs that coincide with low wealth. According to theory, the value of the security to the hypothetical investor equals the average payoff calculated using this preference-weighted probability distribution.

... the statistical likelihood of a payoff ...

... with a preference-filtered likelihood ...

The ratio of downside risk measured under the two probability distributions is related to the investor's risk aversion. In terms of the labelled areas in

Statistical versus risk preference-weighted probability density functions of payoffs

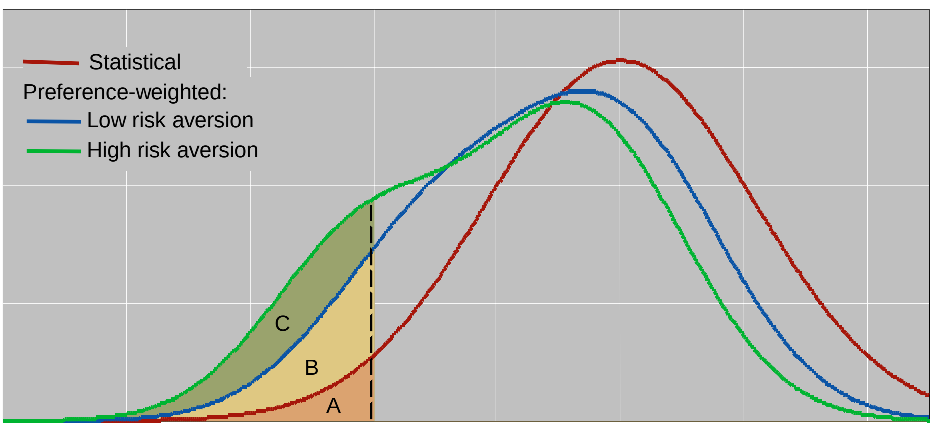

Graph 1

Graph 1, this ratio is equal to $(A+B)/A$. If the hypothetical investor is less willing to bear risk or, in other words, attaches less value to the possibility of receiving high payoffs than to the avoidance of low payoffs, then his valuation will be based on a distribution such as the one depicted in green. Clearly, for such an investor the security is less valuable and the indicator of risk aversion, $(A+B+C)/A$, has a higher score.

... derived from options

Our derivation of the indicator of investors' risk aversion follows the above logic closely and is detailed in the box on the next page. We use option prices for the estimation of market participants' preference-weighted assessment of the likelihood of future returns.[2] Option prices provide a unique insight into investor assessments of future payoffs. This is due to the fact that an array of option contracts, based on different strike prices of the same underlying asset, is observed simultaneously on each trading date. This cross section of option prices makes it possible to estimate the subjective probability that investors ascribe to future payoffs, represented by the option strikes.

Effective risk aversion may reflect preferences ...

There are reasons to believe that the indicator of risk attitude may change over time. For instance, there is the possibility that different periods might be characterised by a different collective disposition of investors vis-à-vis risk-taking. Arguably, the component of our indicator that is based on such fundamental determinants of risk aversion can evolve only gradually, if at all.

... the identity of active players ...

Alternatively, one could argue that the indicator measures the *effective* risk aversion of those investors that actively participate in the market. In this respect, a possible source of time variation can be changes in the composition of the set of active investors. Our calculations are based on observed prices in the cash and derivatives markets and, as such, reflect the collective views of the active participants at the time. For a variety of institutional and regulatory reasons, different types of market participants have a different tolerance of risk. For example, pension funds and foundations are typically more conservative investors that put high priority on capital preservation. In contrast, hedge funds are more aggressive in their pursuit of high returns. Even if neither type of investor changes its attitude towards risk and return, the effective choices between risk and return reflected in the spot and option prices will be sensitive to the identity of the active participants at any given juncture.

... or responses to changing conditions

Finally, the risk aversion indicator we construct might also be viewed as reflecting the insurance value of an option that can also be time-varying. At times, risk management systems may impose mechanical trading behaviour that is effectively similar to that implied by heightened risk aversion.[3] For instance, when the predetermined floor for a portfolio's value is reached, or an operation's risk budget is exceeded, the systems prescribe the sale of risky assets. Thus, the value of an option with a sufficiently high payoff in such situations would offer valuable protection to investors against reaching their

---

[2] Hayes and Shin (2002) construct a similar indicator of risk aversion.

[3] A theoretical treatment of this issue is provided in Danielsson et al (2002).

Calculation details

In this box, we outline the methodology for deriving a numerical indicator of aversion to downside risk. The indicator is calculated as the ratio of two measures of downside risk: one based on the preference-weighted probability density function (PDF), derived from option contracts on a specific security, and the other on the statistical PDF, which is calculated on the basis of the historical behaviour of returns on the asset underlying the option contracts. In the finance literature, the preference-weighted PDF is often referred to as the "risk neutral" PDF.

The derivation of the *preference-weighted* PDF is based on the non-arbitrage argument of Breeden and Litzenberger (1978), who show that such a PDF is equal to the second derivative of the option price with respect to the option strike. Unfortunately, option contracts are traded only for a set of discrete strike prices of the underlying security. To overcome this difficulty, we follow Shimko (1993) by first estimating a continuous, "smooth" implied volatility function that is consistent with the option prices for the range of observed strikes. More specifically, we estimate a quadratic volatility "smile" by minimising the sum of the weighted squared differences between it and the volatility implied by the observed prices for the range of traded contracts. We use option prices from contracts with 45 days to maturity. We then derive the corresponding continuous option price function based on this implied volatility function and calculate its second derivative numerically.

Preference-weighted and statistical PDFs

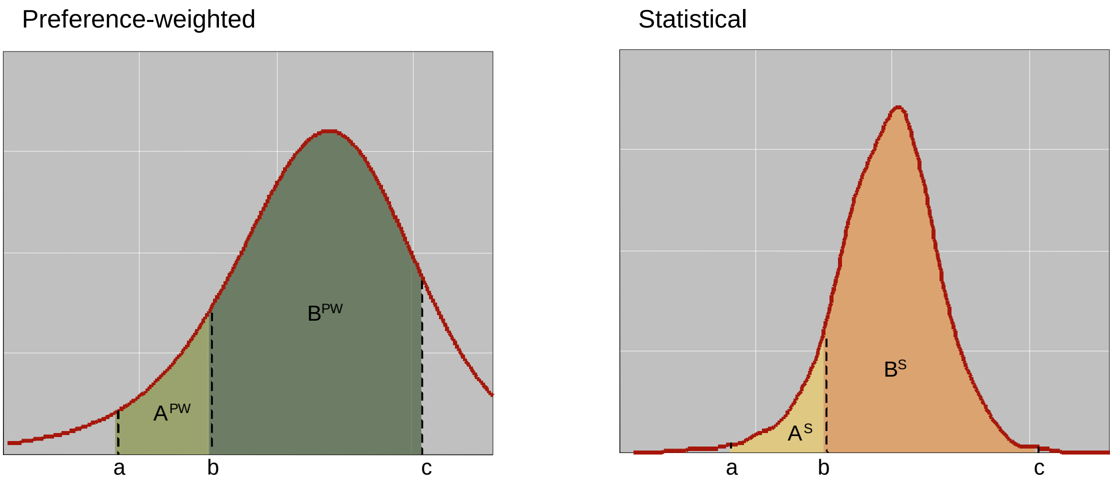

The estimate of the *statistical* distribution is based on an asymmetric GARCH model first suggested by Glosten et al (1993). It incorporates two established characteristics of asset returns: the persistence of volatility, and the tendency of volatility to rise as returns fall. The model is estimated each month on the date we observe the option prices, using information available up to that point in time. We then simulate the estimated model 5,000 times, generating a distribution of the asset's returns 45 days into the future in order to match the date of expiration of the option contracts.

The graph above shows the two distributions for a typical day in our sample. The preference-weighted distribution (left-hand panel) is truncated between points b and c, reflecting the range of strikes for which we observe option prices on that particular day. Because our indicator of risk aversion is sensitive to the probability mass in the left tail of the distribution, we do not extrapolate beyond the bounds of observed strikes. Hence, the indicator is expressed as the ratio of (i) the preference-weighted conditional probability of a 10% or larger decline in the underlying asset to (ii) the corresponding statistical probability. In terms of the labelled areas in the graph above, our indicator is equal to:

$$\left(\frac{A^{PW}}{A^{PW}+B^{PW}}\right)\Big/\left(\frac{A^{S}}{A^{S}+B^{S}}\right)$$

This formulation comes as close to the one presented in Graph 1 as the objective limitations of the data permit.[①] Note that our methodology is qualitatively similar to the estimation of risk premia as the difference between futures prices, which account for investors' risk preferences, and statistical expectations of the same underlying asset's returns.

[①] It is conceivable that the truncation bounds of the preference-weighted PDF change over time for reasons unrelated to our analysis and may, in principle, affect the value of the indicator. Inspection of the movement of these bounds suggests, however, that it cannot be at the root of the empirical regularities found in this special feature.

risk limits and, consequently, against being forced to liquidate positions under stress. The indicator of risk aversion will increase in situations when such considerations are expected to dominate the behaviour of market participants.

Analysis of risk aversion indicators

Indicators for three equity markets ...

We apply the basic idea outlined in the previous section, and detailed in the box, in order to calculate monthly indicators of market participants' effective risk attitude using information from option prices and cash returns on the S&P 500, FTSE 100 and Dax 30 equity indices. The data cover the period from December 1995 to December 2002. We calculate the risk preference-weighted likelihood of index returns as implied by option prices observed 45 days prior to each option contract's expiry date. On average, there are 37 strikes for the S&P 500 options, 25 for the FTSE options and 29 for the Dax options.

Graph 2 plots the derived indicators for the three equity markets. Higher values of the indicators are associated with lower investor tolerance of risk. The three indicators exhibit a fair degree of variation over time. There is an upward shift timed around the second half of 1997, during the period when currency crises spread widely in the Southeast Asian region. This heightened sensitivity to risk is not fully reversed in the subsequent years. In fact, during the market turbulence in the autumn of 1998, our indicator series register the longest sustained rise in investors' reluctance to bear risk. In contrast, the events of 11 September 2001 are marked only by a short-lived jump in the three indicators.[4]

... show a high degree of co-movement

A striking feature of the graph is the degree of co-movement between the three indicator series. Bilateral correlation coefficients ranging between 62% and 78% confirm the visual impression. We interpret this fact as suggesting that integrated financial markets tend to be driven by the actions of investors with similar perceptions and objectives. Furthermore, since we estimate the three indicators independently for each equity index, this co-movement provides a reassuring signal for the validity of our methodology.

[4] These patterns are very similar to those exhibited by the indicator of risk aversion constructed by Hayes and Shin (2002), which is based on the same principles.

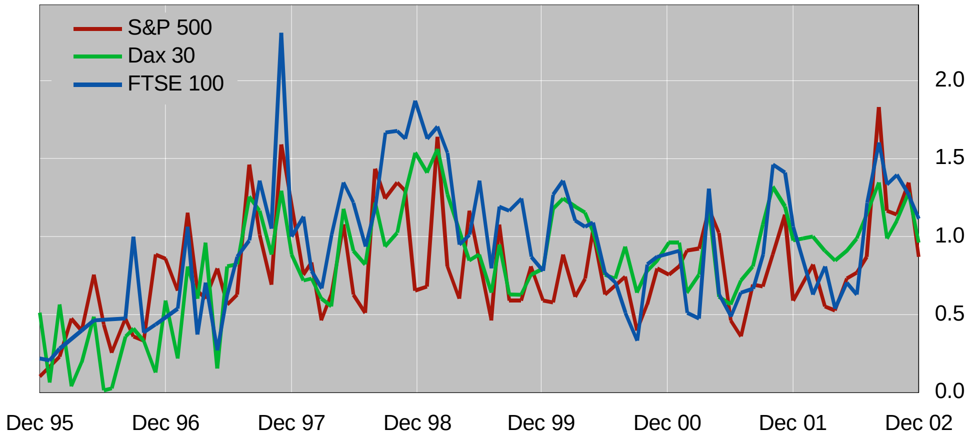

Sources: Bloomberg; CME; Eurex; LIFFE; BIS calculations. Graph 2

The high correlation between the three indicators suggests that there is a strong common factor driving their dynamics. For the rest of this special feature we base the analysis on this common factor, which we derive statistically as the first principal component of the three indicators. The new series accounts for 80% of the overall variation in its constituent series.

Risk aversion and the dynamics of financial markets

In this section, we examine whether the behaviour of asset prices changes systematically with the level of investors' effective risk aversion. To this end, we focus on the three equity indices we used to derive the indicator, and price indices of US, UK and German government bonds with a maturity of seven to 10 years. We have classified each month in our sample as being characterised either by "high" or "low" effective aversion to risk on the basis of the value of the risk aversion indicator.[5] Tables 1 and 2 contain, respectively, univariate and bivariate descriptive statistics on the annualised daily returns on the six assets. The statistics are calculated over the entire period and over each of the two subsamples characterised by different levels of effective risk aversion.

High risk aversion coincides with ...

We first test whether asset return distributions are similar across the two subsamples marked by "high" or "low" risk aversion. The results of a test for equality of the distribution functions are reported in the bottom row of each panel of Table 1. The test concludes that returns on most assets exhibit different statistical behaviour in periods characterised by different levels of risk aversion. The sole exception is the gilt market, where we cannot reject the hypothesis that the returns are drawn from the same distribution.

In order to cast light on what factors drive the outcome of the distribution test, we examine separately the returns' first four moments. More specifically, we calculate for the entire period and for each of the two subsamples the

---

[5] We use the median value over the entire period of our composite indicator as the cutoff point in order to determine the high and low risk aversion subsamples. We have tried a variable trend as the cutoff point with no material impact on the results.

average annualised daily return, the annualised volatility of the daily returns, a measure of the asymmetry in the probability of low and high returns (skewness), and a measure of the likelihood of extreme returns in either direction (kurtosis). Tests of equality of these measures across the two subsamples attribute the differences in the distribution of returns, indicated by the overall test, to differences in the first two moments. Mean returns of equity indices are lower in periods that are characterised by higher effective risk aversion, while the opposite is true for fixed income securities. This finding is consistent with the intuitive argument that investors would tend to withdraw from riskier asset classes as they become less inclined to take on risk. In contrast, at times of increased risk tolerance, the demand for riskier assets would tend to support an increase in their price at the expense of that of bonds.

... below par performance of equity markets ...

... higher volatility ...

Another general pattern that emerges from Table 1 is that higher risk aversion is associated with higher volatility of asset prices. This result holds for both asset classes, but is more pronounced in the case of equities. A possible interpretation of this pattern is consistent with one of the motivations we offered for the time variation in market participants' effective risk aversion. Increased price volatility is tantamount to heightened market risk and is likely to coincide with periods when participants' capital base is stretched to its limits. This, in

Return distributions and risk aversion

December 1995–December 2002

	Equities[1]											
	S&P 500				FTSE 100				Dax 30			
	Whole	High	Low	Test[2]	Whole	High	Low	Test[2]	Whole	High	Low	Test[2]
Mean	0.081	–0.041	0.206	**	0.080	–0.014	0.177	*	0.109	–0.064	0.287	**
Std dev	0.195	0.228	0.153	***	0.191	0.234	0.132	***	0.265	0.331	0.173	***
Skewness	–0.168	–0.197	0.159	.	–0.237	–0.197	–0.026	.	–0.280	–0.171	–0.308	.
Kurtosis	5.868	5.245	4.385	.	5.347	4.196	4.359	.	5.785	4.414	3.913	.
K-S test[3]	0.000***				0.000***				0.000***			
	Fixed income[1]											
	US Treasury notes				UK gilts				German bunds			
	Whole	High	Low	Test[2]	Whole	High	Low	Test[2]	Whole	High	Low	Test[2]
Mean	0.043	0.074	0.012	*	0.024	0.055	–0.007	*	0.000	0.020	–0.020	.
Std dev	0.066	0.069	0.063	**	0.063	0.065	0.062	.	0.057	0.057	0.057	.
Skewness	–1.032	–0.946	–1.160	.	–0.753	–0.947	–0.532	.	–2.177	–2.045	–2.324	.
Kurtosis	7.431	7.151	7.725	**	8.964	8.666	9.382	.	14.941	13.414	16.620	.
K-S test[3]	0.021**				0.025**				0.429			

[1] A normality test rejects the hypothesis that the returns are drawn from normal distributions. [2] Outcome of the test of whether the difference between the moment estimates across the two subsamples is greater than zero. *, ** and *** indicate that the null hypothesis is rejected at the 10%, 5% and 1% significance levels, respectively. [3] Kolmogorov-Smirnov test of whether returns are distributed identically across the two subsamples: p-values indicate the significance level at which one rejects the null hypothesis that the distribution of returns is invariant to the measure of risk aversion.

Sources: Bloomberg; Datastream; BIS. Table 1

Cross-correlations of asset returns[1]

December 1995–December 2002

	S&P 500	FTSE 100	Dax 30	US Treasury notes	UK gilts	German bunds
S&P 500	...	0.448	0.493	–0.229	–0.071	–0.022
FTSE 100	0.420	...	0.741***	–0.197	–0.150	–0.133
Dax 30	0.437	0.587	...	–0.269	–0.178	–0.077
US Treasury notes	0.126***	0.085***	–0.036***	...	0.389	0.183
UK gilts	0.142***	0.149***	0.095***	0.362	...	0.335**
German bunds	0.026	0.045***	0.176***	0.116	0.233	...

[1] The numbers above (below) the main diagonal correspond to correlations in "high"("low") risk aversion periods. *, ** and *** indicate that the hypothesis of equality of correlations between asset returns across "high" and "low" periods of risk aversion is rejected at the 10%, 5% and 1% significance levels, respectively.

Sources: Bloomberg; Datastream; BIS. Table 2

turn, increases the insurance value of options. In terms of Graph 1, the area under the left tail of the preference-weighted likelihood function increases and so does the indicator of risk aversion.

An alternative interpretation would reverse the direction of causality. When traders are more reluctant to expose themselves to risk, they are particularly cautious in managing their portfolios and tend to react more vigorously to news. Furthermore, when the overreaction is market-wide, it would be difficult to find counterparties for investment positions. Large swings in prices would then be a natural consequence.

The correlations of returns across different equity markets appear to increase during periods when investors are more apprehensive about risk (Table 2).[6] The direction of causality is ambiguous. On the one hand, a closer co-movement of stock markets narrows the scope for portfolio diversification, thus increasing the correlation of market returns with investors' wealth. As explained earlier, this would tend to increase the effective risk aversion of investors. On the other hand, it is possible that increased volatility (or, equivalently, measured market risk) might be driving both the higher correlations and the higher values of the indicator of effective risk aversion. Loretan and English (2000) show that higher correlation between asset prices should be expected during periods of increased volatility. As risk management systems typically register higher market risk during these periods, one would expect investor behaviour that is observationally similar to lower tolerance of risk, and similarities in investment strategies might lead to a tighter relationship between stock markets. Despite the fact that the differences in correlations between the two subsamples appear economically significant, formal tests fail to establish statistical significance except for the correlation between the German and British equity markets.

... and decoupling of bond and equity returns

The comparison of tightness of correlation between bond and equity returns exhibits a clearer pattern. The co-movement between the two asset

[6] This can be observed by comparing entries in the table that are symmetrically positioned with respect to its main diagonal.

classes is uniformly higher during periods of lower risk aversion. The differences are not only economically significant (in the range of 10–20 percentage points) but also pass the statistical test of equality. This finding is consistent with the results from the comparison of the univariate statistics reported earlier. According to Table 1, government bond markets are less sensitive to shifts in investors' attitude towards risk than equity markets, the returns on which tend to suffer as investors withdraw from risk-taking during periods of heightened risk aversion. Thus, during those periods the prices in the two asset classes tend to move in opposite directions, leading to lower correlation.

Conclusion

This special feature compares data that can be extracted from option and cash markets in order to derive time series of risk aversion indicators. An encouraging feature of the estimation results is that these indicators co-move closely across market segments.

Furthermore, we find evidence that financial market dynamics tend to change systematically with the level of investors' effective risk aversion. In particular, heightened risk aversion is associated with lower returns and higher volatility, especially for equity markets, and weaker co-movement of asset classes. Our findings thus have a bearing on the interpretation of signals sent by financial markets. Incorporating changes in risk attitudes in such an interpretation adds information relevant for understanding the functioning of financial markets.

References

Breeden, D and R Litzenberger (1978): "Prices of state-contingent claims implicit in option prices", *Journal of Business*, vol 51, pp 621–51.

Danielsson, J, H S Shin and J-P Zigrand (2002): "The impact of risk regulation on price dynamics", London School of Economics working paper.

Glosten, L R, R Jagannathan and D E Runkle (1993): "On the relation between the expected value and the volatility of the nominal excess return on stocks", *Journal of Finance*, vol 48, no 5, pp 1779–801.

Hayes, S and H S Shin (2002): "Liquidity and risk appetite: evidence from equity index option prices", Bank of England working paper.

Loretan, M and W B English (2000): "Evaluating changes in correlations during periods of high market volatility", *BIS Quarterly Review*, June.

Shimko, D C (1993): "Bounds of probability", *Risk*, vol 6, pp 33–7.