Recovering Risk Aversion from Option Prices and Realized Returns

by

Jens Carsten Jackwerth*

First draft: March 25, 1996

This version: October 31, 1997

C:\main\PAPER4\PAPER7.doc

Abstract

A relationship exists between aggregate risk-neutral and subjective probability distributions and risk aversion functions. Using a variation of the method developed by Jackwerth and Rubinstein (1996), we estimate risk-neutral probabilities reliably from option prices. Subjective probabilities are estimated from realized returns. This paper then introduces a technique to empirically derive risk aversion functions implied by option prices and realized returns simultaneously. These risk aversion functions dramatically change shapes around the 1987 crash: Precrash, they are positive and decreasing in wealth and thus consistent with standard economic theory. Postcrash, they are partially negative and increasing and irreconcilable with the theory. Overpricing of out-of-the-money puts is the most likely cause. A simulated trading strategy exploiting this overpricing shows excess returns even after accounting for the possibility of further crashes and transaction costs.

---

* Jens Carsten Jackwerth is a visiting assistant professor at the London Business School. For helpful discussions I would like to thank Ian Cooper, Nils Hakansson, Hayne Leland, Anthony Neuberger, Mark Rubinstein, Stephen Schaefer, and seminar participants at Berkeley, Michigan, McGill, Yale, Harvard, London Business School, Long Term Capital Management (London), University of Southern California, Inquire (UK) at Newcastle, Davis, City University London, Warwick, Bank of England (London), and the Computational Finance Conference at London.

Recovering Risk Aversion from Option Prices and Realized Returns

A relationship exists between aggregate (that is market-wide) risk-neutral and subjective probability distributions and risk aversion functions across wealth. In each state of the world the following relationship holds:

subjective probability = risk-neutral probability $\times$ risk aversion adjustment

The risk-neutral probability is the price that an "average investor" would pay for receiving one dollar in that state multiplied by the riskfree return[1]. The subjective probability is simply the assessment of the "average investor" of how likely a state is to occur. These two probabilities would be identical if the investor was indifferent to risk. However, the investor might value a dollar more highly in certain states, namely ones where wealth is low. The risk aversion adjustment indicates these preferences of the investor. Once we know both risk-neutral and subjective probability distributions, we can derive those preferences empirically and avoid specifying them explicitly a priori which was needed in the past.

Estimating risk aversion functions directly is still notoriously difficult and, until recently, the same was true for risk-neutral distributions despite some theoretical papers on the problem[2]. Jackwerth and Rubinstein (1996) developed a convenient method to recover risk-neutral distributions from option prices[3]. For the subjective distributions, we assume that they are approximated by the historical (actual) return distribution of a broad index such as the S&P500. We then empirically derive risk aversion functions implied by option prices and realized returns[4].

---

[1] We need to multiply by the riskfree return in order to have a proper risk-neutral probability distribution which sums to one.

[2] See Banz and Miller (1978), Breeden and Litzenberger (1978), and Ross (1976).

[3] For alternative methods see Shimko (1993) and the extensions in Brown and Toft (1997) for polynomial interpolation of the smile. Melick and Thomas (1997) use lognormal mixtures and Aït-Sahalia and Lo (1998) use kernels. Derman and Kani (1994) use implied trees and Dumas, Fleming, and Whaley (1997) use deterministic volatility functions.

[4] This methodology was also used in Aït-Sahalia and Lo (1997). Other, less closely, related studies are Rosenberg and Engle (1997) and Derman et al. (1997).

At this point, we are confronted with a potential puzzle which is depicted in the following two figures. During the precrash period both the risk-neutral and the actual distributions look about lognormal. Figure 1 depicts the distributions on October 22, 1986, and is typical of the precrash period. The actual distribution is located slightly to the right due to the risk-premium. We know that in this case a power utility function of moderate risk aversion is likely to provide a reasonable risk aversion adjustment for translating one distribution into the other.

Figure 1. Precrash risk-neutral and actual distributions. The risk-neutral distribution is the 30-day option implied distribution on October 22, 1986. The actual distribution is a kernel density of the 30-day non-overlapping returns over the 4 years prior to October 22, 1986.

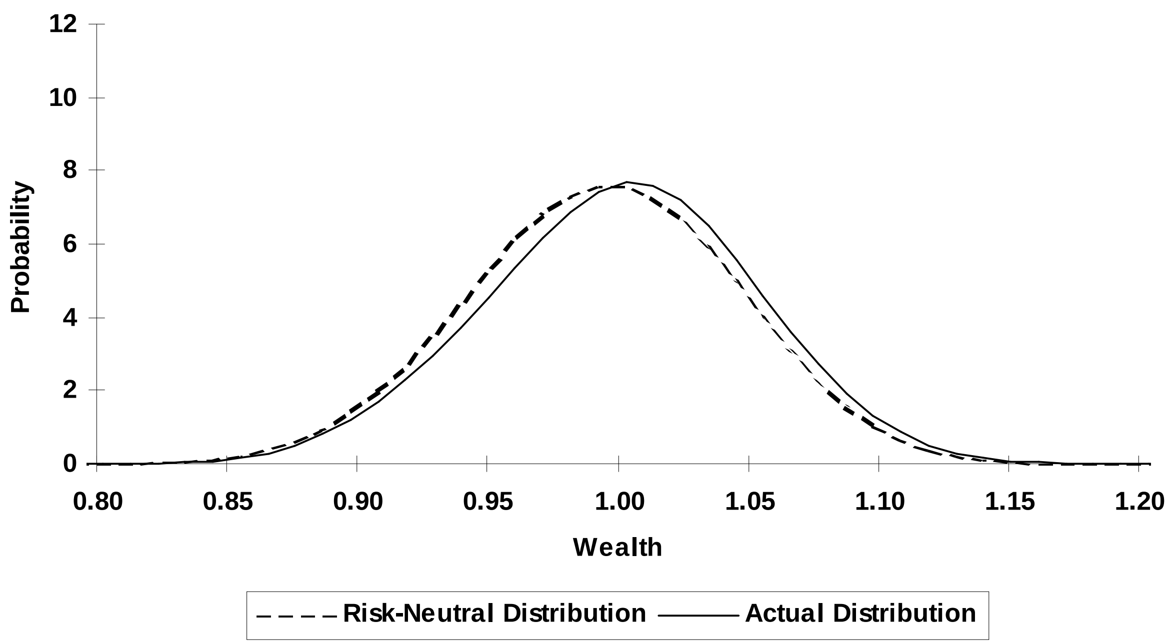

Next, we turn to the postcrash period where the actual distribution looks about lognormal again, but the risk-neutral distribution is now left-skewed and leptokurtic (more peaked).[5] Figure 2 depicts the distributions on April 15, 1992, and is typical of the postcrash period. If we conclude that the risk-neutral distributions changed in shape around the crash and that the actual distributions (which proxy for the subjective distributions) did not, then we are confronted with a relationship between risk-neutral and subjective distributions and risk aversion functions where the first component (risk-neutral distributions) changed and the second (subjective distributions) remained the same. That in turn could imply that the third component (risk aversion functions) changed, too, around the crash. This paper sets out to empirically investigate this possibility.

Figure 2. Postcrash risk-neutral and actual distributions. The risk-neutral distribution is the 30-day option implied distribution on April 15, 1992. The actual distribution is a kernel density of the 30-day non-overlapping returns over the 4 years prior to April 15, 1992.

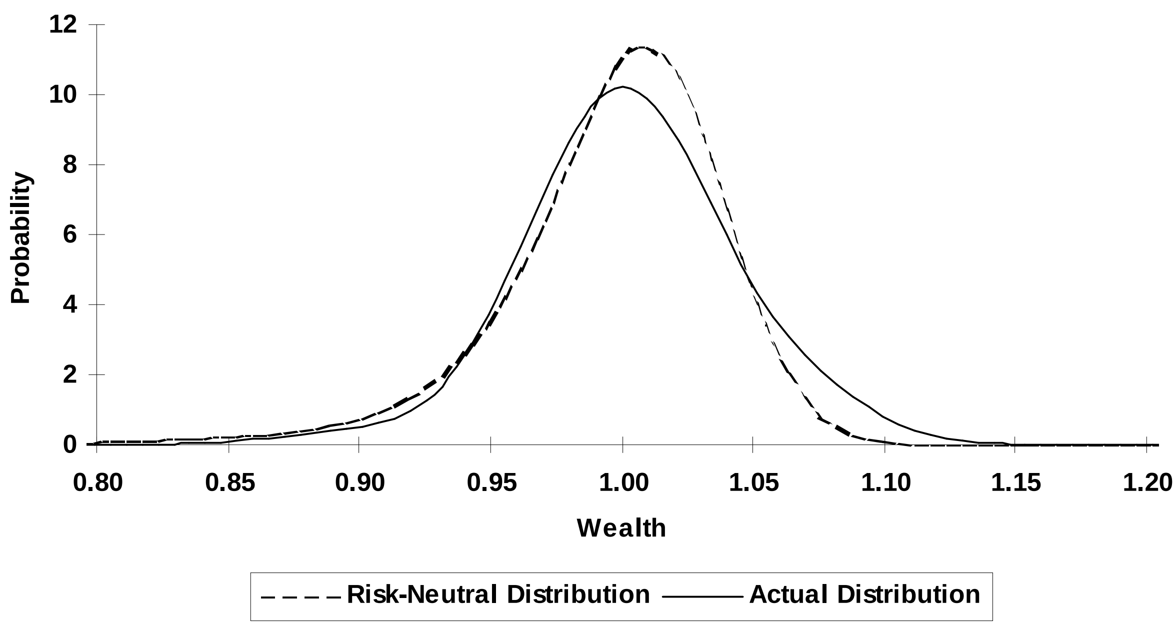

[5] Compare Jackwerth and Rubinstein (1996) for more extensive results with respect to the implied risk-neutral distributions.

We find indeed that implied risk aversion functions dramatically change shapes around the 1987 crash and thereafter exhibit negative and increasing risk aversion across parts of the wealth dimension. These findings are remarkably stable with respect to changes of almost all model parameters. Moreover, a number of potential criticisms of our methodology are discussed and rejected. One potential explanation is that the postcrash options market overvalues deep-out-of-the-money puts. A simulated trading strategy based on selling these puts mean-variance dominates holding the market. This result would not surprise us since there were no more significant crashes between 1988 and 1995. However, we simulate crashes as well and still find put strategies which strictly mean-variance dominate holding the market even in the presence of transaction costs.

Section I presents the model for obtaining the implied risk aversion functions. Section II describes the data and procedures used for the empirical work. Section III presents and analyzes the empirical results. The implied risk aversion functions are checked against standard assumptions such as positivity and potential explanations for empirical violations are provided. The most likely explanation is mispricing of out-of-the-money puts. A trading strategy is simulated which exploits these inconsistencies and mean-variance dominates holding the market. Section IV concludes.

I. Implied Risk Aversion Functions

In this section, we introduce a model of the economy within which we derive implied risk aversion functions across wealth. We consider a complete market economy with a representative investor with an endowed wealth of 1 unit and a fixed time-horizon t.[6] The problem is to maximize utility across future wealth given that all wealth has to be invested in the market[7]:

$$\underset{W}{\text{Max}} \int QU(W)dW - \lambda\left(\frac{1}{r^t}\int PWdW - 1\right) \qquad (1)$$

where: W = future wealth

W	=	future wealth
Q	=	subjective probability across states
U	=	state-independent utility function across states
λ	=	shadow price of the budget constraint (incremental utility associated with increasing the available investment marginally)
r	=	1 + riskfree interest rate
t	=	time horizon
P	=	risk-neutral probability across states

We can then differentiate with respect to the shadow price λ in order to obtain our original wealth constraint again. By differentiating with respect to wealth, we find the first order conditions which have to hold in equilibrium. In equilibrium we know that the aggregate investor has to hold the market portfolio. Thus, if S is the return (cum dividends) on the market portfolio across states then the following relation has to hold in equilibrium[8]:

---

[6] See Constantinides (1982) for the argument that the economy has a representative investor under completeness.

[7] Alternatively, we can think in terms of Arrow-Debreu securities: W is then the number of elementary securities optimally held in a given state.

[8] U′ is guaranteed to be positive since all terms on the right-hand side are positive.

$$U'(S) = \frac{\lambda}{r^t}\frac{P}{Q} \tag{2}$$

The approach to find the implied utility functions by using equation 2 directly suffers from the fact that we cannot measure the value of the shadow price λ accurately. Its formula involves the probability in the tails of the subjective distribution where we have little confidence in our estimates. But there exists an interesting way around this problem. Many utility functions have simple expressions for the absolute risk aversion which we can rewrite in terms of the subjective and risk-neutral distributions after differentiating equation 2 with respect to S once more:

$$U''(S) = \frac{\lambda}{r^t}\frac{P'Q - PQ'}{Q^2} \tag{3}$$

Here we would like to assure that U" < 0. However, the concavity of the utility function is not guaranteed from the data and we will check for violations in the empirical section. Finally, we can write an expression for the absolute risk aversion in terms of the subjective and risk-neutral distributions[9]:

$$RA = -\frac{U''(S)}{U'(S)} = -\frac{\frac{\lambda}{r^t}\frac{P'Q - PQ'}{Q^2}}{\frac{\lambda}{r^t}\frac{P}{Q}} = \frac{Q'}{Q} - \frac{P'}{P} \tag{4}$$

The expression for the risk aversion functions is no longer dependent on the shadow price λ and thus a lot more tractable. In particular, with equation 4 we are able to obtain the risk aversion for any state by using information about that state only. Thus, we can concentrate our attention to the center states with the highest probability of occurring and neglect the notoriously

[9] See Leland (1980).

unreliable tails of the distributions[10]. Evaluation of the risk aversion functions is rather simple since empirically we have to discretize our states and then we can use finite difference approximations to the first derivatives.

II. Data

For our economy with a representative investor, we would like to use a broad index resembling the market in our empirical work. Thus, the empirical research in this paper is based on a database which contains minute by minute trades and quotes covering S&P500 European index options, S&P500 index futures, and S&P500 index levels from April 2, 1986 through December 29, 1995.[11] Moreover, we use daily closing prices and dividends on the S&P500 from January 2, 1928 through December 29, 1995.

A. Risk-Neutral Probability Distribution

In order to find the risk-neutral probability distribution, we use a variation of the method proposed in Jackwerth and Rubinstein (1996) for recovering the risk-neutral distribution from option prices. The implied probabilities derived with the maximum smoothness criterion are given on equally spaced returns S_j. All options used to obtain risk-neutral distributions have times-to-expiration of 30 days. In order to obtain sets of option prices across several striking prices we aggregate all daily quotes into one volatility smile. We calculate the implied volatilities for all options with the same striking price. We compute the mean implied volatilities for each striking price and treat this set as our representative daily volatility smile for the 30-day options. We then deviate slightly from the original method by finding the smoothest implied volatility function and not the smoothest probability distribution consistent with the observed data. We

[10] In particular, relative errors around the probability distributions are lower in the center than in the tails.
[11] See Jackwerth and Rubinstein (1996) for a detailed description of the data.

also measure the standard deviation of those implied volatilities for each striking price. This information is then used to force a tighter fit around at-the-money options (where standard deviations are lower) and allow for a looser fit away from the money. We achieve this by using the following objective function for which the first order conditions can be solved in closed form[12]:

$$(1-p)\int(\sigma'')^2\,dK + p\sum_i\left(\frac{\sigma_i-\overline{\sigma}_i}{STD_i}\right)^2 \qquad (7)$$

where: p = trade-off parameter for balancing smoothness (integral) versus fit (sum)

σ_i = implied volatility associated with a given striking price K_i

$\overline{\sigma}_i$ = implied volatility of observed option with striking price K_i

STD_i = standard deviation of $\overline{\sigma}_i$

From equation 7, we can calculate the Black-Scholes option prices and differentiate twice. After multiplying by r^t, we can write down the risk-neutral probability distribution directly:

$$P = r^t\left[\frac{r^{-t}n(d_2)}{S\sigma\sqrt{t}}[1+2S\sqrt{t}\,d_1\sigma'] + S_0 d^{-t}\sqrt{t}\,n(d_1)[\sigma''+\frac{d_1 d_2}{\sigma}(\sigma')^2]\right] \qquad (8)$$

where: n(x) = normal density function = $\frac{1}{\sqrt{2\pi}}e^{-\frac{1}{2}x^2}$

d_1 = $\dfrac{\ln\frac{S_0 d^{-t}}{Sr^{-t}}}{\sigma\sqrt{t}} + \frac{1}{2}\sigma\sqrt{t}$

d_2 = $d_1 - \sigma\sqrt{t}$

S_0 = index value today

[12] The related method of kernel regression (Härdle (1990)) yields very similar volatility functions.

The dividend is assumed to be a dividend yield based on the actual payments throughout the life of the option. The interest rate is the average of the median implied borrowing and lending rates from put-call parity of all feasible pairs of 30-day options on that day. The index level for our representative daily sets of option prices is the median of the intraday levels of a futures-based index. The futures-based index is obtained by deflating all futures quotes and trades by the median daily implied repo-rate corresponding to the time-to-maturity of the future. For each minute, the median of all deflated quotes and trades is computed and used as the futures-based index for that minute. This futures-based index does not suffer from stale prices as the index does.[13]

Out of our total sample, we are left with 116 days for which we have sufficient data since we only observe the 30-day option once every month.

B. Subjective Probability Distributions[14]

In order to come up with the subjective distribution, we use the historical time series of the index. The length of the sample is an issue and we select 4 year time frames. This particular choice and a number of parameter choices below can be criticized as somewhat arbitrary. However, we show that varying these initial choices does not change the basic findings.

We calculate 30-day, non-overlapping returns from our 4 year sample and derive the kernel density under a Gaussian kernel method for smoothing the raw returns with a bandwidth:[15]

---

[13] See Jackwerth and Rubinstein (1996) for a detailed description of the data.

[14] In an earlier version of this paper, we used a bootstrap method which is susceptible to destroying serial dependence. However, that method seems to be more stable with regard to higher order moments which are hard to estimate but can affect the kernel method below. In particular, Jackwerth (1997a) argues that the loss of serial dependence might only be a minor effect which does not distort the shape of the actual distribution too much.

[15] See Silverman (1986), p. 15. This choice of a window-width leads to slight over-smoothing but removes spurious multi-modalities. The volatility of the kernel density increases in the process but is reset to the sample volatility. Solving an optimization problem which searches for the smoothest distribution consistent with the actual returns gives very similar results but is significantly slower.

$$h = \frac{1.8\sigma}{\sqrt[5]{n}} \qquad (9)$$

where: h = bandwidth for the kernel

σ = standard deviation of the sample returns

n = number of observations

Finally, we resample the smoothed returns onto an equal spacing S_j which we choose to be compatible with the sample intervals of the risk-neutral probability distribution. The resampling is conducted in the cumulative probabilities and with adjustments for the discreteness of the sample space.[16]

III. Empirical Results

At this point, we have to make a few choices to obtain our risk aversion functions. However, we will show shortly that our empirical results are not sensitive to reasonable changes in any of these choices. We measure the actual distribution with a historical mean and volatility which can be rather different from those moments for the future period associated with the risk-neutral probability distribution. We know that the historical mean is very difficult to estimate[17] and that the historical volatility is not necessarily indicative of the volatility which will be realized in the future. Thus, we normalize the actual distribution to a distribution with mean 0 and volatility 1 and then set the mean to the riskfree interest rate + 8% per year for the risk premium[18] and the standard deviation to the standard deviation of the risk-neutral probability distribution. Resetting the volatilities has an additional benefit. If we believe that the subjective

---

[16] See Jackwerth (1996) for a more detailed discussion of the procedure.
[17] See Merton (1980) who also argues that one should explicitly model positive risk-premia.
[18] However, a better assumption might actually be that the risk-premium per unit variance $(\mu - r) / \sigma^2$ is constant rather the risk-premium itself.

and the risk-neutral probability distributions vary only in the drift term due to the risk premium then it is reasonable to require the volatilities to be identical.

A. Empirical Risk Aversion Functions

We are now in the position to calculate the risk aversion functions across wealth. However, we are also interested in error-bounds around our estimates. The following analysis is based on bootstrapped error-bounds[19] where the basic idea is to resample from the inputs (historical returns and option quotes) with replacement and recalculate the risk aversion functions. The process is repeated 100 times and 95% percent error-bounds are calculated. For this procedure to be valid, we have to assure that the inputs, that is the historical returns and the option quotes, are independent from each other. For the historical returns, we assume stationarity and the non-overlapping returns are then independent draws from the true, actual probability distribution. For the option quotes, we use the partial smiles for each minute as independent input variables. They appear to move around some true smile as the day progresses. This way we avoid dependencies across striking prices.

The empirically observed risk aversion functions across wealth are unfortunately not nearly as well behaved as the standard economic theory suggests. Figure 3 shows our empirical results which are aggregated but still indicative of the respective periods covered. The error-bounds are surprisingly tight. This has been known for the error-bounds around the risk-neutral distributions[20] but is a new result for risk aversion functions.

[19] See Press et al. (1988), pp. 548, for a very clear introduction and Härdle (1990), p. 105, for further details on the bootstrap method in particular.
[20] Jackwerth (1997b) and Aït-Sahalia and Lo (1998).

Figure 3. Risk aversion functions across wealth. For 116 days from April 2, 1986 through December 29, 1995 we calculate the absolute risk aversion functions across wealth where a wealth level of one corresponds to the forward price of the index at maturity of the options used for obtaining the risk-neutral probability distribution. For each wealth level we calculate the mean risk aversion across the periods 860402-871018 (○), 881019-910318 (Δ), 910319-930818 (◆), 930819-951229 (□). In addition, we calculate the 95% error bounds for each period through a bootstrap procedure.

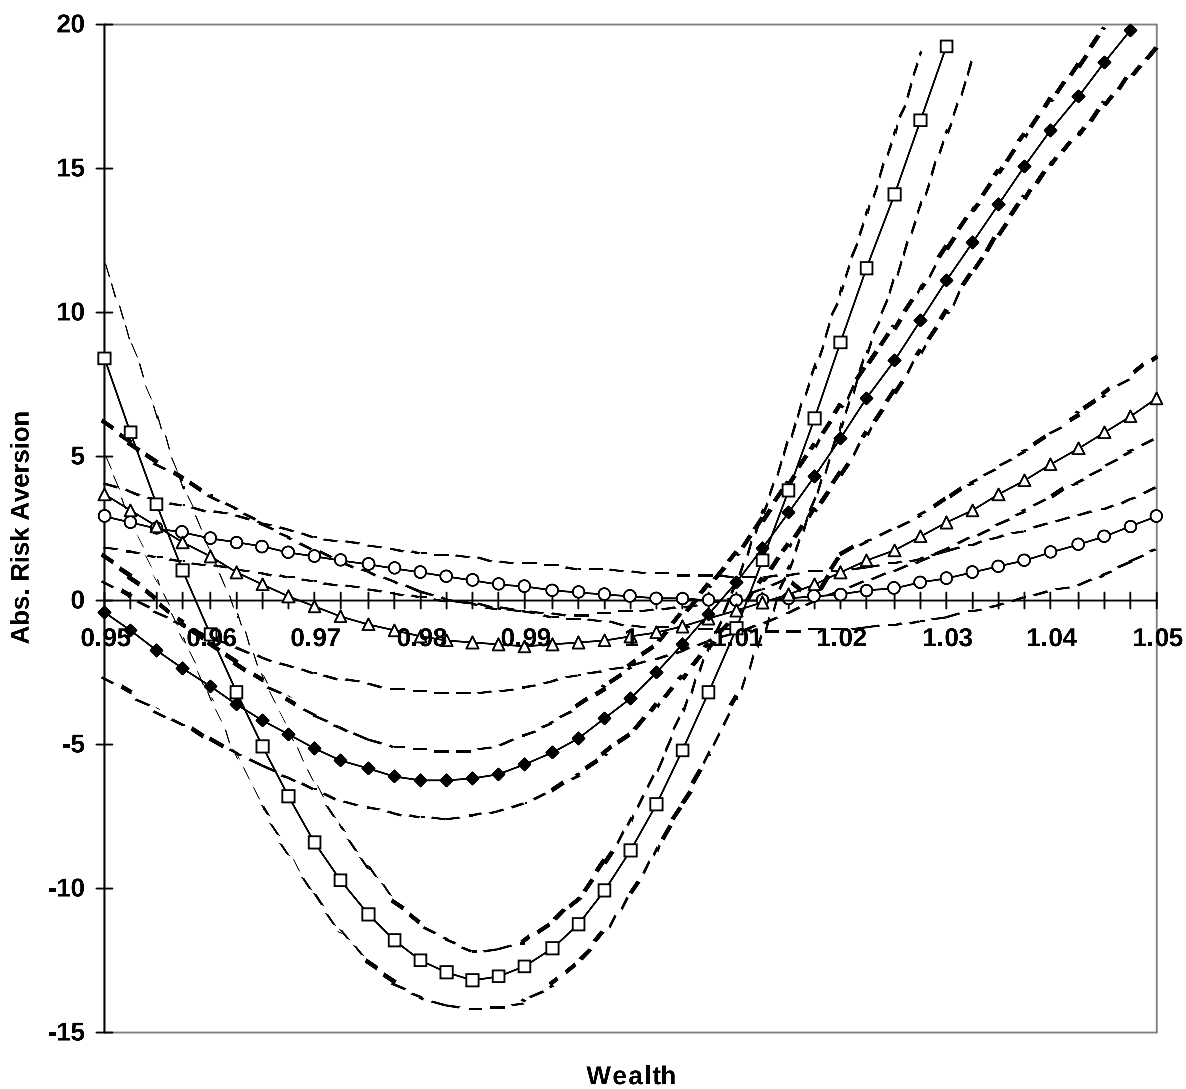

Precrash, the risk aversion function is consistent with standard economic theory. It is positive and decreasing across wealth. However, in the postcrash period two disconcerting features arise and become progressively worse as we move through time[21]. First, the postcrash risk aversion functions are negative around the center. Second, these risk aversion functions rise as wealth increases beyond .99. Before we analyze these features in detail, we want to address the potential concern that our particular choices for the inputs caused them.

B. Stability of the Empirical Risk Aversion Functions

We find that the general shape of the functions in figure 3 is remarkably stable with respect to different input choices. Using a different length of the historical sample of 2 or 10 years instead of 4 years does not change figure 3 significantly. Small changes in the risk premium in the interval where most financial economists would expect it to lie over the long run (5 to 10% per year) does not change figure 3 other than through a slight shift (downward and upward, respectively). However, very large risk premia cause all risk aversion functions to shift upward while the shape of the functions is preserved. A risk premium of more than 25% per year would finally cause all functions to be non-negative throughout. Also, we changed the bandwidth for the kernel density from .5 to 2 times its value in equation 9. Small bandwidths lead to spurious multi-modalities which causes all risk aversion functions to fluctuate more. However, the basic features of figure 3 are still preserved in that the precrash risk aversion function is mostly positive and the postcrash ones are mainly negative. Finally, we use the actual distributions straight with their means and volatilities as calculated from the historical data. Again, figure 3 does not change significantly except that the precrash risk aversion functions are

---

[21] Interestingly, neither Derman et al. (1997) nor Aït-Sahalia and Lo (1997) find negative risk aversion functions postcrash. However, their studies provide only one risk aversion function each and not time series of the functions. It could well be the case that we observe "well-behaved" risk aversion functions on individual days postcrash. Moreover, they both use overlapping returns to estimate the actual distribution. For both papers, I recalculated their actual distribution with and without overlapping returns and the two distributions look rather different.

now even higher and the postcrash ones even lower. Adjusting the means only and using the volatilities straight does not change figure 3 either.

C. Discussion of the Empirical Risk Aversion Functions

In our discussion of figure 3 we want to point out that the results for the precrash period are rather encouraging. The implied risk aversion function is positive throughout and both absolute and relative risk aversion functions are decreasing with increasing wealth as one might deem reasonable behavior for the economy.

The problems occur postcrash where risk aversion functions are now negative throughout wealth levels from about .96 through 1.01. Moreover, risk aversion is increasing for wealth levels greater than about .99, that is for ending up wealthier over the next month than one started out.

Inspection of equation 4 yields insight into this behavior. The risk aversion function is the vertical distance of two downward sloping functions (Q' \ Q and P' \ P). The respective intersections with the wealth axis occur at the modes (that is the maxima) of the probability distributions. In the precrash period, depicted in figure 1, with about lognormal distributions for the subjective and risk-neutral probabilities, the mean, median, and mode tend to lie rather close together for each distribution compared to the risk premium. This changes in the postcrash period, depicted in figure 2, where the risk-neutral distribution suddenly exhibits a mode to the right of the mean. The mode of the risk-neutral distribution often appears even to the right of the mean (and mode) of the subjective distribution, causing negative risk aversion functions. The steep drop in probability to the right of the mode of the risk-neutral distribution causes the P' \ P function to fall rapidly after intersecting the wealth axis. This leads to increasing absolute risk aversion for wealth levels greater than about .99.

One way to interpret figure 3 would simply be to believe what it suggests, namely that the risk aversion functions changed dramatically around the crash. If all the assumptions for the representative agent model hold true, then concavity in the utility functions implies all positive risk aversion functions. In turn, since we observe negative risk aversion functions, we would now have to believe that the utility functions are convex. This would lead to corner solutions where the representative agent with a convex utility function would avoid returns in all states of the world where the utility function is convex and would seek returns in the adjacent states where the utility function is concave. This cannot be the case since the representative agent has to hold the market. If we are not willing to accept the implications of figure 3, then we might suspect that one or more of the assumptions are violated. The main concern is here that the subjective distribution might not be well approximated by the historical distribution. Or, the S&P500 could be a poor proxy for the market as a whole. The assumptions for a representative agent, namely a complete market and state-independent utility function could be violated. Or, transaction costs could increase the prices of deep-out-of-the-money puts. Finally, the options could be mispriced. We will now look into each of these objections in turn.

An explanation might be that the subjective distribution is not well approximated by the historical distribution and that the subjective distribution is much more similar in shape to the risk-neutral distribution than the historical distribution would suggest. This is to argue that we might only sample historically from the center of the true subjective distribution and hardly ever from the far left tail (the so-called peso problem). Then the sample distribution would be closer to lognormal than the true subjective distribution. However, the period after the crash of 1987 provides historical samples which include the crash and therefore the far left tail. This period should then give more well-behaved implied risk aversion functions and figure 3 seems to suggest this since the risk aversion function for the period from October 1988 through March 1991 is only slightly negative and U-shaped. However, this conclusion is misleading since almost the same shapes of the functions are obtained with historical samples spanning 2 or 10 years instead of 4 years. These samples weigh the crash differently and therefore should not all yield that same shape of the risk aversion functions.

The S&P500 index was used as a market proxy and this could be a problem. However, Ibbotson, Siegel, and Love (1985) measure the world wealth portfolio and the US wealth portfolio. They report correlations of the wealth portfolios with the S&P500 of .7 and .9 respectively. Moreover, the less an asset is correlated with the market, the more risk-neutral should the pricing of that asset be. Thus, if the S&P500 was not highly correlated with the market, then our findings would be even more surprising as the risk-neutral distributions do not look like the subjective distributions at all in the postcrash period. Finally, it is not clear why the S&P500 should have been a good proxy for the market before the crash but not after the crash.

In an incomplete market, not all states would be spanned by the observed securities. Since we normally have only about 10-15 option striking prices, we have to be concerned that we are looking at an incomplete market. The question is then, how far away could the true risk-neutral distribution (under a complete market) be from the implied risk-neutral distribution which we measure. Arbitrage bounds help us by limiting the possible probability mass associated with states which are not spanned. But these bounds tend to be rather loose. A better restriction is to impose some reasonable constraints on the smoothness of the risk-neutral distribution since we have the strong prior that the probability of one state should not be vastly different from the next state. Jackwerth (1997b) reports such results and finds rather tight bounds on the probabilities in states which are not spanned.[22] These findings severely limit the extent to which incomplete markets could change our results. We also note that the particular shapes of the implied risk-neutral distributions we use have been confirmed through related but different methodologies which are listed in footnote 2. We are inclined to believe that we are using empirical risk-neutral distributions which approximate the true risk-neutral distributions rather well.

With regard to the assumption of <u>state-independent</u> utility functions, one has to ask which <u>state-dependent</u> utility functions are viable candidates. It is not clear why any of the additional state factors (common candidates would be interest rates, credit spread, national income, inflation, and other macro variables) would have changed dramatically around the crash.

Finally, option traders have pointed out that transaction costs, especially for the deep out-of-the-money put options, are high due to difficulties in hedging these options. An example are

[22] Also, compare Aït-Sahalia and Lo (1998).

trading halts after large down-moves which prevent dynamic hedging in a timely fashion. That would drive up their prices and the prices of the corresponding deep in-the-money calls. The volatility smile would then result from these costs only and the underlying implied risk-neutral distribution without transaction costs could well be approximately lognormal. As a test for the impact of these option prices with low striking prices, we first calculate the implied risk-neutral distributions including puts with striking price / index levels greater than .79 which corresponds to our standard methodology. Next, we calculate the distribution by only using ratios greater than .89. Comparing these distributions does not show significant differences other than in the lower left tail. Moreover, the mode of the distributions and the right tail are virtually identical and these areas are the driving forces behind the negativity of the implied risk aversion functions and the increasing risk aversion for wealth levels greater than about .99.

That leaves one more unpalatable possibility. Maybe the market is pricing the deep-out-of-the-money puts just too high. A reason might be that mutual fund companies desire insurance against crashes through purchasing puts since they do not trust dynamic portfolio insurance any longer. The increased demand could then lead to higher prices for these puts. Then we should be able to make money from this mispricing. We use the following methodology. During the 7 years starting in October 1988 and ending in December 1995[23], we select 116 trading dates one month apart. Each trading period, we start out with $100 to invest and we can invest into the riskfree asset or into the S&P500 or we can select one of four option strategies. These involve selling out-of-the-money puts with either of four striking price / index level ratios: .9, .925, .95, and .975.

We can only sell as many puts as we can margin fully with $100. We use a margin requirement for writing uncovered options of 30% of the stock price less the out-of-the-money amount. All positions are closed out one month later and we earn the riskfree return on our margin. For each of the 116 one month periods we bootstrap 1,000 sample returns from the

[23] The sample is extended from the earlier version of this paper. The additional data for 1995 provides for a true out-of-sample test which fully supports all our earlier findings.

observed monthly returns throughout the whole 7 year sample. Our sample does not include the crash but we suspect that the options are priced to reflect that possibility. We therefore compute 5 different runs where we introduce artificial crashes with a probability of one crash every 4, 8, 16, and 32 years and one run without any crashes[24]. A crash is created by replacing one monthly return during the specified period with that return times .8. We finally compute the average of the 116 means and variances of annualized returns for each strategy. We show in figure 4 the expected excess returns (alphas) across betas[25].

---

[24] With a crash every 4 years, we lose about 5% risk premium per year due to crashes. Therefore, we did not model more frequent crashes since the security market line would then slope downward.

[25] One objection to the above analysis might be that we do not properly adjust for the risk of higher order moments beyond mean and variance such as the negative skewness which is created by selling naked puts. Recent work by Leland (1996), based on Rubinstein (1976), allows us to remedy this problem. If the market exhibits lognormal returns (which approximately holds true for the actual returns), then the usual beta measure should be replaced with an adjusted beta measure. Next, we can replace the usual alpha measure with an adjusted alpha to measure the adjusted expected excess return. However, figure 4 does not change much when we use the adjusted measure.

Figure 4. Risk-adjusted expected excess returns of selling out-of-the-money puts. We depict the strategy of investing $100 in the riskfree asset and in the S&P500. Returns are calculated for 116 one month periods starting in October 1988 until December 1995. The market return for each one month period is bootstrapped from the complete 7 year sample. In four runs, artificial crashes of -20% during one month are added every 4, 8, 16, and 32 years on average and in a fifth run, no crashes are added. We report the average of the 116 expected excess returns (alphas) across betas. The expected excess returns for the riskless asset and the market are zero with betas of 0 and 1 respectively. We similarly report four strategies of selling as many puts as we can margin with $100 with striking price / index level ratios of .9, .925, .95, and .975.

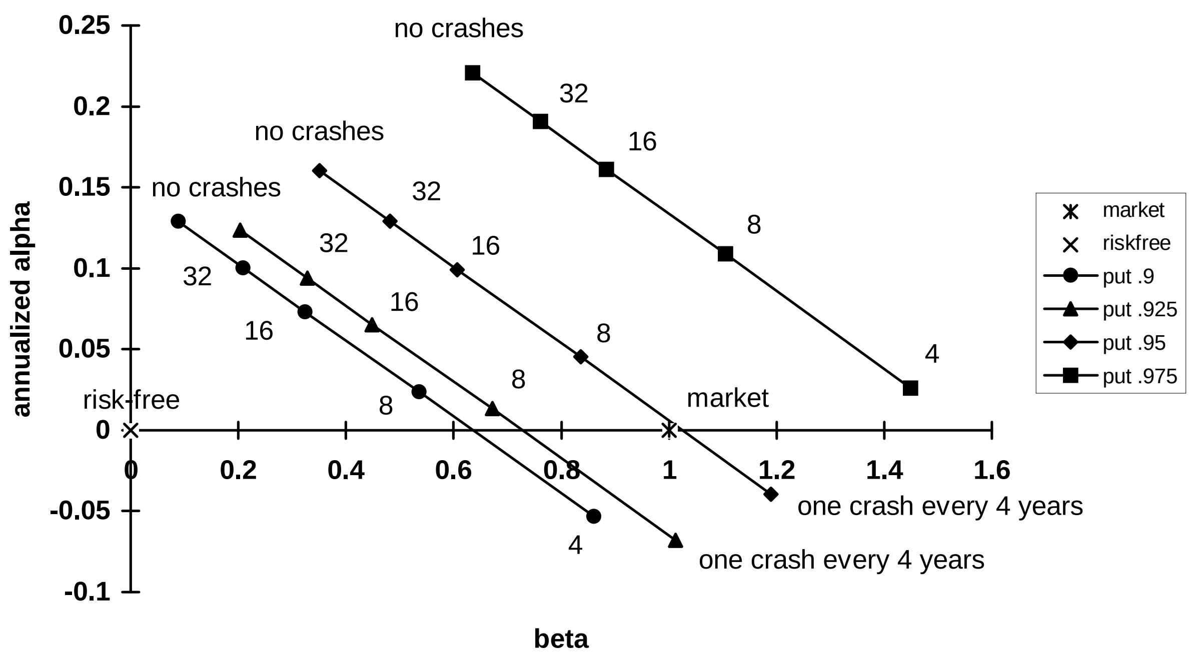

We find excess returns for all option trading strategies as long as the probability of a crash does not exceed one in 8 years and even a 20% crash every 8 years seems to be a rather pessimistic outlook[26]. For the strategy where we sell the puts with a striking price / index level ratio of .975, we would even come out ahead with crashes as frequent as one every four years. Also, if the same analysis is conducted for longer term options (say 180 days-to-expiration) then the strategy is significantly more profitable for crashes occurring as frequently as one every four years with excess returns of about 6-8% per year. The reason for this finding is that over the next half a year, it is more likely for the market to recover from a crash then over the next month only.

An objection to this analysis might be that the bootstrap approach changes the original sequence of returns. We therefore use the above analysis, including the artificial crashes, in combination with the observed returns in their historical sequence for each trading period. Figure 4 changes only insignificantly.

Another concern is that above we left out the deep-out-of-the-money puts but still get the same shape of the risk aversion functions. Here, we only trade in the closer to at-the-money puts (with striking price / index ratios of .9 to .975) and we still achieve excess returns.

The above strategy of selling naked puts can lose all the money invested if there are a sufficient number of crashes and the margin is exhausted. But there are two interesting hedging strategies one could implement. For one, if one sells the out-of-the-money puts, buys the out-of-the-money call, and shorts the index (or sells a future on the index) then the payout pattern is a lot more favorable with limited gains on the downside and limited losses on the upside. More interestingly, Toft and Prucyk (1996) reported that individual stock option smiles are significantly flatter than the index option smiles in the postcrash period. A hedging strategy is therefore to sell out-of-the-money puts on the index and to buy out-of-the-money puts on (only a few) individual stocks. The hedge will perform the better the more we need it since in the case of a crash,

[26] We recalculated figure 4 with transaction costs and the expected excess returns per year are 3-4% lower. With transaction costs, all positions are entered at the current bid prices observed on the trading date and charged a 2% commission. The commission is mainly relevant for the initial sale since the puts are rarely exercised and we therefore normally do not pay commission at expiry.

correlations across different stocks (and the index) tend to go up.[27] Further research will investigate the economic causes of negative skewness as the smile for index options is very steep compared to the smiles for individual stock options.

IV. Conclusion

The present paper models a representative agent economy in which we can empirically determine implied risk aversion functions through simple expressions of risk aversion across wealth in terms of the subjective and risk-neutral probability distributions, only.

The resulting risk aversion functions are consistent with standard economic theory during the precrash period but we find partially negative and increasing risk aversion functions during the postcrash period. These findings are remarkably stable with respect to varying input parameters. Changes in the empirical methodology do not seem to affect the results.

Thus, we are left with either believing that risk aversion functions did indeed change dramatically around the crash or questioning the assumptions of the model in order to explain their radical shapes or concluding that the market misprices some options. The most likely explanation is mispricing of options in the market since negative risk aversion functions are irreconcilable with a representative agent and since we can defend the other assumptions of the model rather well. A simulated trading strategy exploiting such mispricing yields risk-adjusted expected excess returns during the postcrash period.

[27] Kelly (1994).

References

Aït-Sahalia, Yacine and Andrew W. Lo, 1998, Nonparametric estimation of state-price densities implicit in financial asset prices, *Journal of Finance*, forthcoming.

Aït-Sahalia, Yacine and Andrew W. Lo, 1997, Nonparametric risk management and implied risk aversion, working paper, Chicago.

Banz, Rolf and Merton Miller, 1978, Prices for state-contingent claims: Some estimates and applications, *Journal of Business* 51, 653-672.

Breeden, Douglas and Robert Litzenberger, 1978, Prices of state-contingent claims implicit in options prices, *Journal of Business* 51, 621-651.

Brennan, Michael J. and R. Solanki, 1981, Optimal portfolio insurance, *Journal of Financial and Quantitative Analysis* 16, 279-300.

Brown, Gregory and Klaus Bjerre Toft, 1997, Constructing binomial trees from multiple implied probability distributions, working paper, University of Texas at Austin.

Constantinides, George M., 1982, Intertemporal asset pricing with heterogeneous consumers and without demand aggregation, *Journal of Business* 55, 253-267.

Derman, Emanuel, Michael Kamal, Iraj Kani, and Joseph Zou, 1997, Is the volatility skew fair?, Quantitative Strategies Research Notes, Goldman Sachs, New York.

Derman, Emanuel and Iraj Kani, 1994, Riding on a smile, *RISK* 7, No. 2, 32-39.

Dumas, Bernard, Jeff Fleming, and Robert E. Whaley, 1997, Implied volatility functions: empirical tests, working paper, HEC.

Härdle, Wolfgang, (1990), Applied Nonparametric Regression (Cambridge University Press, Cambridge).

Ibbotson, Roger G., Laurence B. Siegel, and Kathryn S. Love, 1985, World wealth: market values and returns, *Journal of Portfolio Management* 12, No. 1, 4-23.

Jackwerth, Jens Carsten, 1996, Implied binomial trees: generalizations and empirical tests, working paper, Berkeley.

Jackwerth, Jens Carsten and Mark Rubinstein, 1996, Recovering probability distributions from option prices, *Journal of Finance* 51, No. 5, 1611-1631.

Jackwerth, Jens Carsten, 1997a, Do we live in a lognormal world?, working paper, London Business School.

Jackwerth, Jens Carsten, 1997b, Geometric and probabilistic bounds on option prices, working paper, London Business School.

Kelly, Michael A., 1994, The U.S. equity market has become more diversified, J P Morgan Equity Derivative Strategies, New York.

Leland, Hayne E., 1980, Who should buy portfolio insurance?, *Journal of Finance* 35, 581-594.

Leland, Hayne E., 1996, Beyond mean-variance: performance measurement of portfolio using options or dynamic strategies, working paper, Berkeley.

Melick, William R. and Charles P. Thomas, 1997, Recovering an asset's implied PDF from option prices: an application to crude oil during the gulf crisis, *Journal of Financial and Quantitative Analysis* 32, No.1, 91-115.

Merton, Robert C., (1980), On estimating the expected return on the market: An exploratory investigation, *Journal of Financial Economics* 8, 323-361.

Press, William H., Brian P. Flannery, Saul A. Teukolsky, and William T. Vetterling, 1988, Numerical Recipes in C (Cambridge University Press, Cambridge).

Rosenberg, Joshua V. and Robert F. Engle, 1997, Option hedging using empirical pricing kernels, working paper, New York University.

Ross, Stephen, 1976, Options and efficiency, *Quarterly Journal of Economics* 90, 75-89.

Rubinstein, Mark, 1973, A comparative statics analysis of risk premiums, *Journal of Business* 46, 605-615.

Rubinstein, Mark, 1976, The valuation of uncertain income streams and the pricing of options, *Bell Journal of Economics and Management Science* 7, 407-425.

Shimko, David, 1993, Bounds on probability, *RISK* 6, No. 4, 33-37.

Silverman, B. W., 1986, Density Estimation for Statistics and Data Analysis (Chapman and Hall, London).

Toft, Klaus Bjerre and Brian Prucyk, 1996, Options on leveraged equity: theory and empirical tests, working paper, University of Texas, Austin.