The Smirk in the S&P500 Futures Options Prices: a Linearized Factor Analysis

Andrew Carverhill,* Terry H. F. Cheuk and Sigurd Dyrting

July 2004

Abstract

We construct portfolios of S&P500 futures and their associated options, which are Delta (price) and Vega (volatility) neutral. These systematically earn negative abnormal returns, and suggest that out of the money puts are too expensive, relative to out of the money calls. We give evidence that these negative returns are not a payment for insurance against a market crash.

We then do a factor analysis on the Delta hedged option price innovations. Including a 'smirk' factor, there is no evidence of arbitrage opportunities. However, the smirk seems unable to predict the skew in the underlying return, though is useful for hedging portfolios of options. We finally conclude that the smirk represents the risk premium of a dynamic aversion to market falls, which seems unrelated to the underlying futures index.

Addresses:

Andrew Carverhill, School of Business, University of Hong Kong, Pokfulam Road, Hong Kong. Email: carverhill@business.hku.hk. Tel (852) 2857 8358. Fax (852) 2858 5614.

Terry H.F. Cheuk, School of Business, University of Hong Kong, Pokfulam Road, Hong Kong. Email: terrycheuk@business.hku.hk.

Sigurd Dyrting, Deutsche Bank, 20 Finsbury Circus, London EC2M 1NB, UK.

Acknowledgment: The work described in this paper has been partially supported by the University Grants Council of Hong Kong, China, grant number HKUST6229/97H.
JEL classifications: G12, G13.
Keywords: S&P500 index, futures, options, smirk, skew.

*Corresponding author

The Smirk in the S&P500 Futures Options Prices: a Linearized Factor Analysis

Abstract

We construct portfolios of S&P500 futures and their associated options, which are Delta (price) and Vega (volatility) neutral. These systematically earn negative abnormal returns, and suggest that out of the money puts are too expensive, relative to out of the money calls. We give evidence that these negative returns are not a payment for insurance against a market crash.

We then do a factor analysis on the Delta hedged option price innovations. Including a 'smirk' factor, there is no evidence of arbitrage opportunities. However, the smirk seems unable to predict the skew in the underlying return, though is useful for hedging portfolios of options. We finally conclude that the smirk represents the risk premium of a dynamic aversion to market falls, which seems unrelated to the underlying futures index.

Introduction, Aims, and Conclusions

This paper studies the 'smirk' in US equity index options, which is the ubiquitous stylized fact that options with higher strikes have lower implied volatilities, when the implied volatilities are calculated in the standard Black-Scholes Model. Thus, a graph of these implied volatilities looks like a smirk (see Figure 1). The options 'smirk' is a variant of the 'smile', under which options with strikes further away from the money have higher Black-Scholes implied volatilities[1]. Bates (1991) noted that the various S&P index options developed their smirk after the 1987 market crash, and interpreted it in terms of a premium that the market is willing to pay for insurance against a further crash.

Until recently, attempts to explain the smirk in equity index options have focussed on refining the model of the underlying index returns and the associated risk premia. This approach has not been completely successful. For example, stochastic volatility models have the potential to capture the smirk, via a negative correlation between the return and the volatility, but many authors, notably Bakshi, Cao and Chen (1997) have noted that

[1]Bollen and Whaley (2003) document that options on individual stocks, rather than the index, tend to exhibit smiles rather than smirks.

an unreasonably high negative correlation seems necessary to do this. These authors, and later Bates (2000), and Pan (2001) are thus lead to augment the stochastic volatility model with random jumps. The model can then fit the smirk. However, Bates, concludes that the high crash probability implicit in the options prices, under this model, seems inconsistent with the absence of jumps in his data period, which is 1988 - 1993. Pan estimates the parameters of the jump diffusion dynamic for the underlying, simultaneously with the parameters for the corresponding “risk neutral” dynamic, which is appropriate for the option pricing. To do this she takes the underlying and option prices to be a joint process, and applies a GMM technique. However, somewhat similar to Bates, she estimates an extremely high aversion to jumps, even though the jumps themselves might not be very large, and might be indistinguishable from diffusive returns[2].

More recent work has moved beyond the paradigm of perfect, efficient markets. Jones (2001) fits a very flexible semi-parametric model to S&P500 (spot) index put options, which explicitly allows for mis-pricing, in the sense of an expected return in some region of the moneyness-maturity space, which is not a reward for a risk identified in the model. He fits his model in a number of specifications, and concludes that short dated out-of-the-money puts are mis-priced and too expensive, thus giving an interpretation of the smirk. In a similar spirit, Bondarenko (2002) sets up a very general test of rationality, which is immune from peso problems, biased investor beliefs, and does not require knowledge of investor preferences, and shows that the S&P500 futures options fail this test[3]. Also in a similar spirit, but via a different approach, Bollen and Whaley (2003) show that equity option prices react to buying pressure, and suggest that the index option smirk is explained by the fact that there is a particular demand for out-of-the-money index puts, for the purpose of portfolio insurance.

This present paper falls into two parts: first we study the options smirk, using the approach of looking at hedged option portfolio returns, similar to Bakshi and Kapadia (2003), and Coval and Shumway (2001). As Coval and Shumway emphasize, this ap-

[2]See her Table 3. The mean objectively realized jump size is -0.8% of the index, but risk neutral jumps size is -19.2% of the index.

[3]His test does require that the pricing kernel is path independent, and he notes that this puts the popular stochastic volatility models of Heston (1993) etc., and thus also Bakshi, Cao and Chen (1997), Bates (2000), and Pan (2001), and the framework of this present paper, beyond his analysis. However, he he also notes that these models themselves are problematic, in the sense that it seems unclear how they could arise in a general Equilibrium framework.

proach can tell us about the risk premia and efficiency of the market, even if we are unsure whether our model is perfectly specified. As do these authors, we construct our hedges using simple Black Scholes hedge ratios, but we also include checks to ensure that our results are not due to the mis-secification of these hedges. These authors' motivation is to show that the volatility is a priced factor in the market. We should also mention Buraschi and Jackwerth (2001), who use a flexible no-arbitrage framework to draw a similar conclusion, namely that options are not redundant for spanning the pricing kernel. Our motivation is to go a step further, to explore the use of an extra priced factor, beyond the volatility, for spanning the pricing kernel.

Specifically, we construct portfolios which are long the underlying future, short an out of the money (OTM) call, and long an OTM put, and which is Delta and Vega hedged. By regressing the returns of these against the unit constant, we show that they systematically yield less than the riskless return on the investment required to establish the portfolio. Thus, the OTM puts seem too expensive, relative to the OTM calls. We also include the underlying futures return and the at the money (ATM) straddle return in the regressions, representing the price and volatility factors, to ensure that the abnormal negative returns are not a consequence of residual exposure to these factors, caused by the Delta-Vega hedges being imperfect.

This result suggests either that more factors are required to model the data in a no-arbitrage framework, or that there really are arbitrage opportunities. A candidate for an extra factor might be an infrequent jump factor, for which the negative return reflects the risk premium of this factor (i.e. "crash fears"), since our portfolio has a large payoff under such a scenario. To investigate this, we then adapt an idea in Coval and Shumway, and include an extra, deep OTM put in the portfolio, making it "crash neutral", as well as Delta and Vega neutral. By "crash neutral", we mean that it does not give a large positive or negative return, in the event of a large market fall. This extra option does not greatly affect the amounts of the other options in the portfolio, since the extra option has a small Delta and Vega, and the result is substantially unchanged: the expected return of the hedged portfolio is still less than riskless interest rate.

In the second part of this paper, we show that the smirk itself can be taken as a third diffusive, dynamic factor, and the abnormal returns can be explained in terms of the risk premia on this factor and the volatility factor. Our technique for this is essentially to implement a linear priced factor analysis on the option price innovations, having first stripped out the nonlinearity in these price innovations as much a possible using the Black Scholes framework. Our technique includes a test of whether there are arbitrage opportunities, and we find that there are not, so long as we include the smirk factor.

Thus, the smirk represents a third dynamic factor in the options prices, beyond the price and (implied) volatility factors, and its magnitude can be explained in terms of the associated risk premium. Assuming there are no arbitrage opportunities, then options prices should be given as the risk neutral expectation of their payoff, and so must be dependent only on the risk neutral dynamic of the underlying price. It is pertinent therefore to ask whether this factor is connected to the dynamic of the underlying index. One would conjecture that the smirk factor can predict the underlying return skew, just as the implied volatility can predict to the underlying return variance. However, we will see that this is not the case, and the options smirk seems to represent a state variable which resides only in the risk premium itself.

We finally show that the smirk factor is useful for hedging portfolios of options, and we discuss our results in relation to previous literature, including the papers mentioned above, and the very recent papers of Liu and Pan (2003), Eraker, Johannes and Polson (2003), and Eraker (2004).

To summarize our salient conclusions: The smirk in the S&P500 futures options does not represent an arbitrage opportunity or a market imperfection, but can be explained in terms of the risk premium on a third dynamic factor. This factor does not seem to reflect an aspect of the underlying dynamic, but to reside only in the options prices, and it presumably represents a dynamic aversion to market falls.

For our data, we take the CME futures and associated futures options on the S&P500 index. Our data period is the decade of the 1990s, and we concentrate on weekly returns, which seem to yield stronger results than daily returns. Our choice of data contrasts with many of the papers mentioned above, which use the S&P100 or S&P500 spot options. The futures options are American, and this makes the calculations much more computationally intensive, and precludes an easy construction of the risk neutral underlying distribution at maturity time. On the other hand, one expects the futures options market to be very efficient, because hedging with long or especially short positions in the underlying, is much easier with futures, rather than the spot index[4]. Also, with futures

[4]For the same reason, Jorion (1995) uses futures options on foreign exchange, and Amin and Ng (1997) look at Eurodollar futures options.

options, one does not require the dividend yield, to value the option. This seems to be a significant random factor, in the context of our analysis below.

The following is a plan of this paper. We introduce our data in Section 1. In Section 2 we first replicate the results of Coval and Shumway on Delta hedged portfolios, adapted to our data and testing framework. Then we extend this work to Delta and Vega hedged portfolios. In Section 3 we first present a Principal Components Analysis of the options Black-Scholes implied volatilities, to show that their evolution, as a function of moneyness, is overwhelmingly dominated by parallel shifts, and changes in slope, which corresponds to changes in the implied volatility smirk itself. We then adapt this to a formal test on the innovations in the option prices, which does not reject the no-arbitrage hypothesis, when the smirk dynamic is incorporated. We then present some diagnostic tests on the smirk factor, show that it is useful for hedging, and then discuss our results in the context of other recent work on the smirk, and the question of the rationality of equity index option markets. Finally, in section 4 we summarize our results and conclusions.

1 Our Data and its Summary Features

As we have mentioned already, we use the CME futures and associated options on the S&P500 index. These futures contracts trade on a cycle with maturities in March, June, September and December, and each contract matures on the 3rd Friday of the month. Each futures contract is associated with an option, which matures on the same day as the future, and 2 "serial" options, which mature 1 and 2 months earlier than the futures contract. We will restrict our attention to the non-serial options, for simplicity, and because these options are much more heavily traded, and at longer maturities, than the serial options. The options are paid for when they are purchased, i.e. they are not LIFFE style. To value these options, one should replace the dividend yield by the interest rate; also both puts and calls are sometimes optimally exercised early[5].

Our data is purchased from the Futures Industry Institute, of Washington DC. We will work with data at weekly intervals, from Wednesday 01/03/1990 to Wednesday 01/05/2000 (format MM/DD/YYYY), which covers 522 weeks. We will work with set-

[5]This arrangement applies for retail investors. For market makers, there are more complicated margining rules, which are closer to the LIFFE procedures, and are described in Duffie (1989).

tlement prices, which are based on the option prices during the closing period of the day's trading. Also, we will only use option prices for which some trading has been recorded on the appropriate day. These prices are likely to be very reliable, because the daily margining is based on them, and so they are scrutinized closely by the market participants. On the other hand, we do no know whether the prices of the various portfolios that we study are synchronized; however synchroneity problems will tend to erode strength of the results that we will present, and so they will still be valid, to the extent that they are positive.

We also need US$ interest rates. We calculate these from the settlement prices of the Eurodollar futures, which also trade on the CME (and for which the data was also purchased from the FII). In detail, we aggregate and interpolate from the forward rates associated with these futures prices, to get the corresponding interest rates up to the maturity time of the option being valued, and we ignore any convexity correction.

In our tests, we will restrict attention to the non-serial options nearest to maturity, but with maturity beyond a roll-over period of 30 days. Thus, all options considered will have maturity between 30 and 30 days + 3 months $\approx$ 120 days. Figure 1 gives the Black Scholes implied volatilities for all traded puts and calls on the date 06/25/97. The smirk is striking in this figure. Also striking is the number of puts and calls available in this market.

Figure 2 gives the moneyness levels (strike/futures price) available for these call options, over our 522 weeks. At each date, this figure includes a point at the appropriate moneyness level, for each call option for which there has been some trading on that date. We see from this that traded calls are consistently available at moneyness levels between about 96% (4% in the money - 'ITM') and 106% (6% out of the money - 'OTM'), but sometimes not beyond these levels. In our tests, we will restrict attention to this moneyness interval for calls. The figure also includes the total open interest on each date, for moneyness levels below 100% (i.e. ITM), between 100% and 110%, and above 110%. We see from this that there is more open interest in the OTM call options, but there is little open interest beyond 10% OTM. Also, the open interest increases as each maturity approaches, up until the roll-over.

Figure 3 is the same as Figure 2, but for puts. Traded puts are consistently available at moneyness levels below 102% (2% ITM), and availability extends much further OTM than for the calls; up till 1995, traded puts are available to moneyness 88% (12% OTM), and after 1995, to moneyness 80% (20% OTM), and sometimes much further. Our test will also concentrate on these moneyness levels, for the puts. From the open interest graphs, we see that there is more open interest for puts than for calls, and particularly for puts more than 10% OTM.

Many of our tests will involve options with many strikes taken together, interpolated at moneyness intervals at 2% of the underlying price. From Figures 1, 2 and 3, it is clear that this spacing is amply wide enough to ensure that adjacent moneyness values correspond to distinct option prices.

2 Returns to Delta and Vega Hedged Options Portfolios

2.1 An Analytic Framework:

In Subsection 2.3 below, we will give our results on Delta and Vega hedged portfolios. These results can be viewed as an extension of the results of Coval and Shumway (2001), who show that option straddles have negative returns on average, thus providing evidence that volatility risk has a negative price. In Subsection 2.2 we will adapt the Coval and Shumway results themselves to our data and testing framework, and these results will serve to support and strengthen our results of Subsection 2.3.

Specifically, Coval and Shumway construct straddle portfolios, with strikes at various fixed moneyness values, rebalanced on a daily or weekly basis, and constructed to be Delta neutral in the Black-Scholes framework. They show that these straddles have significantly negative return, in an OLS framework, and then argue that this return represents a (negative) payment for the exposure to the volatility risk, that the straddle represents. However, this argument is rather informal, and the conclusion may be spurious, and due to the imperfection of the Black-Scholes hedge. Thus, they then put it on a firmer foundation, by implementing a GMM test of the null hypothesis that their returns obey a specific equilibrium model, which is consistent with Black-Scholes pricing. They show that the null hypothesis is rejected, but then reinstated, when the averages are stripped out of the straddle returns.

We will adopt a different approach to dealing with the imperfection of the Black Scholes hedge, which will involve including extra returns in the regression, to control for residual exposure to the known priced factors. For this, we will work with the following very general assumptions and associated propositions:

Assumption 1: All prices under consideration form a complete market, in which there are no frictions and no arbitrage opportunities. Also, these prices can be determined in terms of state variables (or "factors") $(X_t^1, ..., X_t^n) \equiv \mathbf{X}_t$, which satisfy a joint Ito equation $d\mathbf{X}_t = \sigma(\mathbf{X}_t)(d\mathbf{W}_t + \mu(\mathbf{X}_t)dt)$, in which $d\mathbf{W}_t$ is the differential increment of Brownian motion in n dimensions. Under these assumptions, and some general technical conditions see eg. Duffie (2001), there exist associated risk premia denoted say $\lambda^1, ..., \lambda^n$, which can be functions of $\mathbf{X}_t$, and and are defined as the expected return that the market requires per unit of volatility, for exposure to the risk represented by the state variable X_t^i. The expected rate of return associated with taking the risk represented by dX_t^i is thus $\sigma^i(\mathbf{X}_t)\lambda^i dt$, and the equation for this state variable with respect to risk neutral probabilities is $dX_t^i = \sigma^i [d\mathbf{W}_t + (\mu^i - \lambda^i)dt]$.

Proposition 1: Suppose there is a non-dividend paying asset in this market with price Q_t. Then under Assumption 1, we must have

$$dQ_t - r_t Q_t dt \;=\; \sum_{j=1}^{n} \beta^j \left[d\tilde{X}_t^j + \lambda^i \sigma dt \right], \tag{1}$$

for some $\beta^1, ..., \beta^n$, which can also be functions of $\mathbf{X_t}$, and in which[6] $d\tilde{X}_t^j := dX_t^j - \sigma^j \mu dt \equiv \sigma d\mathbf{W}_t$, so that this is a martingale difference.

Proof: First, Q_t must be a function of $X_t^1, ..., X_t^n$, since these are a complete set of state variables. Thus, applying the Ito Formula, we can write dQ_t in terms of $dX_t^1, ..., dX_t^n$, together with a drift. Taking $\beta^1, ..., \beta^n$ to be the resulting noise coefficients, the drift must be as in Equation (1): to see this, note that this LHS must be a martingale difference with respect to risk neutral probabilities, since it is a cash flow which can be obtained net of investment (See Duffie (2001)); also, to transform the RHS to RNPs corresponds to subtracting $\sum_{j=1}^{n} \beta^j \lambda^j \sigma^j dt$, to leave terms involving $d\tilde{X}_t^j$, making it a martingale difference. *QED*

Assumption 2: The index level S_t and its volatility σ_t are sufficient state variables for the prices under consideration, and the dynamics are homogeneous of degree 1 with respect to the index level. The dividend yield rate d_t and instantaneous interest rate r_t

[6]Here and throughout, the notation ":=" in an equation means that the equation serves to define the term on the left of the notation.

are known up to the maturities of the futures and options prices being considered. Thus we can write the state equations as $dS_t/S_t = \mu^S dt + \sigma_t dW_t^S$ and $d\sigma_t = \mu^\sigma dt + \sigma^\sigma dW_t^\sigma$, and these coefficients, and the risk premia λ^S and λ^σ, can depend on σ_t but not S_t.

Proposition 2: Take a non-dividend paying asset as in Proposition 1, with price Q_t. Also, take a fixed futures maturity T, later than t, and denote the corresponding futures price by f_t^T. Then under Assumptions 1 and 2, we must have

$$dQ_t - r_t Q_t dt \quad = \quad \beta^f df_t^T + \beta^\sigma (d\tilde{\sigma}_t + \lambda^\sigma \sigma^\sigma dt), \tag{2}$$

for some coefficients β^f, β^σ, which can depend on σ_t, and in which $d\tilde{\sigma}_t := d\sigma_t - \mu^\sigma dt \equiv \sigma^\sigma dW_t^\sigma$.

Proof: This follows from Proposition 1, noting that the futures price must be a martingale, with respect to RNPs, since it is a cash flow which can be obtained zero investment. *QED*

In Subsection 2.2, we will adapt the tests of Coval and Shumway, arguing that $\lambda^\sigma < 0$, under Assumptions 1 and 2. In Subsection 2.3 we will argue that we cannot account for the Delta-Vega hedged returns under these assumptions, and in Section 3 we will extend our set of state variables under Assumption 1, to including the smirk factor.

2.2 Analysing Straddle Returns:

The work of this section is closely related to that of Coval and Shumway (2001) (see also Bakshi and Kapadia (2003)), in that we isolate the volatility risk premium by looking at the returns of straddles, in which the price factor should be hedged away. Coval and Shumway show that these consistently tend to have negative returns, after subtracting their financing costs, and they conclude from this that the volatility factor has a negative risk premium.

Coval and Shumway construct straddles, comprising put and call with the same strike. In order to avoid options for which there are no trades, we form our straddles from ATM calls and OTM puts, or ATM puts and OTM calls, as detailed in Tables 1 and 2, which contain the empirical results of this subsection. Given a choice of moneyness levels for the call and put, then we form our straddle portfolios comprising amounts h_t^c of the call and h_t^p of the put with strikes nearest to these levels, and such that

$$h_t^c \Delta_t^c + h_t^p \Delta_t^p \quad = \quad 0, \tag{3}$$

$$h_t^c V_t^c + h_t^p V_t^p \quad = \quad f_t, \tag{4}$$

where f_t is the futures price on which the options are written, and the Δ and V values are the appropriate Delta and Vegas, calculated in the Black-Scholes Model for the underlying returns[7]. The first equation here corresponds to the straddle being Delta neutral, and the other equation is strictly speaking unnecessary, but is included so that the volatility exposure for the straddle is constant unity, in the Black-Scholes Model, when normalized[8] by the futures price.

Table 1 gives the results of 2 regressions, with dependent variable the weekly price innovations in the straddle, for a number of put and call moneyness levels as indicated, minus the financing costs[9], and normalized by the underlying futures price. In the first regression, the independent variable is just the unit constant. The coefficient is negative for every x, usually significant at a level of 0.1%, agreeing with the Coval and Shumway result. The t-statistics in this table are all calculated using the Newey-West information matrix taking 4 lags, and are thus robust to non-normality and heteroskedasticity of the residuals.

However, we cannot rule out the possibility that the straddle is not perfectly Delta hedged, and that the consistently negative returns reflect the risk premium on the residual exposure to the underlying. To address this possibility, we appeal to the above Proposition 2, and include the futures return[10] into the regression. The second regression equation for Table 1 is thus

$$((Str_{t+1} - Str_t) - r_t Str_t \delta t)/f_t = \alpha + \beta(f_{t+1} - f_t)/f_t + (mean\, zero\, residual), \tag{5}$$

[7] By "Black Scholes Model for the underlying returns", we mean that the interest rate is consant, and the underlying index with dividends reinvested, whose rate of return will be the same as that of the futures conract, plus the riskless rate, follows a Geometric Brownian Motion. To calculate the options prices, etc, we use the Binomial method, taking account of early exercise, opportunities.

[8] This normalization is appropriate, and prevents h_t^c, h_t^p from being homogeneous with respect to the futures price, because Vega is homogeneous, of degree 1, with respect to the underlying futures price.

[9] In all our tests, we will take the interest rate to be that derived from the Eurodollar futures prices, and appropriate for the option maturity under consideration. This is rather crude. However, all our empirical results will be essentially unchanged if we double this interest rate, or take it to be zero.

[10] By "futures return" we mean the futures price innovation, divided by the initial futures price. This is a misnomer, because taking a futures position does not entail investment; but it is useful, because the futures return only trivially differs, by the riskless return minus the dividend yield, from the return on the index.

where Str_t is the time t price of the straddle. Comparing this equation with Equation (2), then $\beta^\sigma d\tilde{\sigma}_t$ in Equation (2) (which is a Martingale difference) corresponds to the (mean zero residual) term in (5), and $\beta^\sigma \lambda^\sigma \sigma^\sigma dt$ in (2) corresponds to α in (5). Since the straddles are positively exposed to the volatility ($\beta^\sigma > 0$), then $\alpha < 0$ corresponds to $\lambda^\sigma < 0$.

As we have said, in the first regression of Table 1, for all call and put moneyness combinations, the coefficient on the unit constant is negative, and usually at a significance level of 0.1%. However, in the second regression, the coefficient on the futures return is also always negative. This indicates that the hedge might be mis-specified; and the negative straddle return might be attributed at least partially to the risk premium on the price factor, which is well established to be positive. But as argued above, under Assumptions 1 and 2, the coefficient on the constant in the second regression corresponds to the volatility risk premium having stripped out the price factor, and this coefficient is also negative for each strike combination, and usually significant at 0.1%. The conclusion that $\lambda^\sigma < 0$ thus still stands.

Table 1 also gives the residual Sharpe ratio for the second regression for each strike combination, i.e. the mean of the residual returns, divided by its standard deviation. Identifying Equations (2) and (5), then this mean is $\alpha \equiv \beta^\sigma \lambda^\sigma \sigma^\sigma dt$, and this residual standard deviation is $\beta^\sigma \sigma^\sigma \sqrt{dt}$. Thus, under Assumptions 1 and 2, this Sharpe ratio is an estimate of $\lambda \sqrt{dt}$, i.e. the volatility risk premium itself.

Looking more closely at Table 1, the evidence for $\lambda^\sigma < 0$ tends to be weaker for higher strikes (calls further OTM and puts less OTM). Correspondingly, the Sharpe ratios decline for higher strikes. Since a higher option price corresponds to a lower return, the maturity price being fixed in terms of the underlying, this is consistent with our interpretation of the smirk, that prices are “too high” for lower strikes, under Assumptions 1 and 2 above.

Coval and Shumway also show how to make their straddles “crash neutral”, by including an extra, deep out of the money put in the Delta hedged portfolio, in an amount such that if the market falls so that both puts are exercised, then the value of the portfolio is constant. The motivation for this is to see whether the negative return that they isolate, is really a jump risk premium, and the point about these crash neutral straddles is that they are not vulnerable, and neither do they provide insurance, against a market crash. Following this idea, we add to our portfolio an amount $h_t^{p,Crash}$ of the traded put, which is nearest to being of the money by 15%, and taking the amounts $h_t^c, h_t^p, h_t^{p,Crash}$ to solve

$$h_t^c \Delta_t^c + h_t^p \Delta_t^p + h_t^{Crash} \Delta_t^{p,Crash} = 0, \quad (6)$$

$$h_t^c V_t^c + h_t^p V_t^p + h_t^{p,Crash} V_t^{p,Crash} = f_t, \quad (7)$$

$$h_t^p + h_t^{p,Crash} = 0. \quad (8)$$

The regression results, replacing the straddles of Table 1 by such "crash neutral" straddles, are given in Table 2. Making the straddles crash neutral further erodes the t-statistic on the unit constant coefficient, in each regression, but it is usually still highly significant.

2.3 Delta-Vega Hedged Returns:

Our Delta-Vega hedged portfolios comprise 1 long futures contract, and an amount h_t^c of the call with strike nearest to being out of the money by a proportion x of the underlying futures price, and an amount h_t^p of the put, also with strike nearest to being out of the money by a proportion x of the underlying futures price. We use only options which are traded on each date, and so we take $x = -2\%, 0, +2\%, +4\%$ and $+6\%$. These amounts h_t^c and h_t^p are chosen such that

$$1 + h_t^c \Delta_t^c + h_t^p \Delta_t^p = 0, \quad (9)$$

$$h_t^c V_t^c + h_t^p V_t^p = 0, \quad (10)$$

where Δ_t^c and Δ_t^p are the Deltas, and V_t^c and V_t^p are the Vegas of these options, calculated in the Black-Scholes Model. (Note that the underlying future itself has Delta = 1, and Vega = 0.)

Now, since the out of the money amounts x are always approximately the same for the call and the put in the portfolio, then we have $\Delta_t^c \approx -\Delta_t^p > 0$ and $V_t^c \approx V_t^p > 0$. From these, and Equations (9) and (10), it follows that $h_t^c \approx -h_t^p < 0$. Thus, for $x > 0$ ($x < 0$) we expect this Delta Vega hedged portfolio to lose (gain) money. For $x = 0$, we expect $h_t^c \approx -1, h_t^p \approx +1$, and we expect the portfolio to be close to zero, on the basis of put-call parity, ignoring the effects of early-exercise on the option prices.

In Table 3 gives the results of 3 regressions, for each x above. The first regression includes only the unit constant as independent variable, and it confirms our conjecture, in that the coefficient is positive for $x < 0$, negative for $x > 0$, and indistinguishable from zero for[11] $x = 0$. Motivated as in the previous subsection, we also give a second regression, which includes the futures return as an extra independent variable. This is sometimes highly significant, but does not substantially alter the pattern of significance of the constant coefficient.

The third regression of Table 3 includes the return of the at the money straddle[12] minus financing costs, as another independent variable. Under Assumptions 1 and 2, then this return, taken together with the futures return, will account for the price and volatility risk factors, and will also account for the corresponding risk premia, since these returns are available in the market, with zero investment. This extra independent variable is never significant, and does not alter the pattern of significance of the constant coefficient.

Following Coval and Shumway's idea as in the previous subsection, Table 4 gives the same results as Table 3, but for the Crash neutral Delta-Vega hedged portfolio returns, which include an amount $h_t^{p,Crash}$ of the traded put, which is nearest to being OTM by 15%, and with amounts satisfying

$$1 + h_t^c \Delta_t^c + h_t^p \Delta_t^p + h_t^{p,Crash} \Delta_t^{p,Crash} = 0, \quad (11)$$

$$h_t^c V_t^c + h_t^p V_t^p + h_t^{Crash} V_t^{p,Crash} = 0, \quad (12)$$

$$1 + h_t^p + h_t^{p,Crash} = 0. \quad (13)$$

As in the previous subsection, this extra ingredient slightly erodes the coefficient on the

[11]Bollen and Whaley present a similar regression to this one, but with a different result. See Strategy 5, in their Table 9. This strategy is to sell puts in each of their moneyness categories, and to Delta-Vega hedge by buying an ATM call, and selling the underlying index. This is essentially the reverse of our strategy, and to be consistent with our results, it should make a profit when the put is OTM, and be riskless when the put is ATM. In fact they lose money both for OTM and ATM puts. As an explanation, they suggest that "the market maker is not charging a high enough volatility risk premium". We suggest that this discrepancy between their results and ours is driven by a fundamental difference between our futures options and their spot options, namely that for futures options we do not have to deal explicitly with the dividend yield factor. In fact an 'implied dividend yield' dc_t can be extracted from futures prices via $dc_t = df_t/f_t - ds_t/s_t + rdt$, where s_t is the spot index price. In unreported work, we have compared this with the actual, realized dividend yield, which is available from Datastream. The difference is small, but might be enough to swamp the very small abnormal residual returns in our Table 3.

[12]By "return to the straddle", and any other "return", we mean the price innovation, divided by the initial associated futures price. This again is a misnomer, since the ivestment required to establish the straddle might not be equal to the futures price; but it is useful from the point of view of homogeneity in the regressions.

constant in the regressions, but they are still significant, and we can conclude that the returns on the Delta-Vega hedged portfolios cannot entirely be explained in terms of the risk premium on the price or the volatility, or in terms of a jump risk ("crash") premium.

3 The Smirk Factor

In this section we analyze the abnormal returns of the Delta-Vega hedged portfolios, in terms of what we call the smirk factor, and its risk premium. In Subsection 3.1 we use a Principal Components Analysis to identify 2 factors in the dynamic of the implied volatility structure, which account for the overwhelming part of this dynamic. These factors can be interpreted as a parallel shift and a twist of the implied volatility, thus corresponding to changes in the implied volatility, and in the smirk itself.

In Subsection 3.2 we test a linear priced factor (no-arbitrage) structure, with 1 and 2 factors corresponding to these components, against an alternative, which allows returns which are not accounted for in terms of exposure to these factors. The input data to this test is a vector of option price innovations, in which the futures price factor has already been hedged out. The result is that with 1 factor, the priced factor structure is rejected, consistent with the results of Section 2, but with 2 factors, the priced factor structure is not rejected. Thus, the smirk factor accounts for the apparent arbitrage opportunities in the Delta-Vega hedged portfolios.

In Subsection 3.3 we present some diagnostic tests on the smirk factor, and finally in Subsection 3.4 we present a hedging exercise, in which a single option is hedged with other options. Hedging the smirk factor does lead to a significant reduction of the hedged volatility. In this hedging exercise, we use the Principal Components approach to construct the hedges, and we are careful to do this out of sample, i.e. the hedge on each date is constructed using only data available before that date.

3.1 A Principal Components Analysis (PCA) of the Implied Volatility Innovation Structure:

Denote by $\Phi_t^{x,T}$, the time t price of the OTM option with moneyness $x := X/f_t^T$ and maturity T. (Thus, the option is a put if $x < 1$ and a call if $x \geq 1$.) Also define the corresponding implied volatility $\sigma_t^{x,T}$ to be such that

$$\Phi_t^{x,T} = BS(f_t^T, x, r_t, \sigma_t^{x,T}, T-t), \tag{14}$$

in which BS is the Black Scholes valuation function of the option, taking account of the fact that it is American, and that the dividend yield should be taken to be equal to the interest rate, since the underlying is the futures contract.

In this subsection, we will present a PCA of the weekly implied volatility innovations $\{\delta\sigma_t^{x,T_t})_{\{x=x_1,x_2,\ldots,x_n\}}\}_t$, linearly interpolated at moneyness points x_i, and for each t, taking T_t to be the nearest maturity beyond the roll-over time[13]. A similar exercise has been carried out by Skiadopoulos, Hodges and Clewlow (2000). Principal Components Analysis simply takes an orthogonal basis of eigenvectors $\mathbf{e}^1, \mathbf{e}^2, ..., \mathbf{e}^n$ of the covariance matrix of these volatility innovations, with corresponding eigenvalues $\rho^1 \geq \rho^2 \geq ... \geq \rho^n \geq 0$. The innovations vector can then be expressed as

$$\delta\sigma_t = \mathbf{g}^1\delta s_t^1 + \mathbf{g}^2\delta s_t^2 + ... + \mathbf{g}^1\delta s_t^n, \tag{15}$$

in which $\mathbf{g}^j = \sqrt{\rho^j}\mathbf{e}^j$, and we refer to this as the jth principal component, and $\delta s_t^j := \frac{1}{\sqrt{\rho^j}}\sum_{i=1}^n e_i^j \delta\sigma_t^{x_i}$ are uncorrelated dynamic factors with unit variance. The most efficient way to summarize the innovations in terms of k independent factors, is then to take only the first k components on the RHS of Equation (15).

In Figure 4 we graph the first 2 principal components, taking moneyness points 0.90, 0.92, 0.94, 0.96, 0.98, 1.00, 1.02, 1.04, 1.06. We see from this graph, as one might expect, that the most important aspect of movement in the implied volatility vector is (roughly) parallel shift, and the second is a twist. The corresponding eigenvalues make up respectively 91.3% and 6.8% of the sum of all the eigenvalues, which tells us that these 2 components account for 98.1% of the part of the implied volatility dynamic, that has been accounted for in the linear structure we have imposed by using the covariance structure.

[13] Thus, our implied volatilities are not time homogeneous with respect to maturity. We agree, in unreported work, with the conclusions of Skiadopoulos, Hodges and Clewlow (2000), that the implied volatility structure is not sensitive to the time to maturity.

3.2 Fitting a Linear Priced Factor Structure to the Option Price Innovations:

Applying the Ito formula to Equation (14), we have

$$\delta\Phi_t^{x,T} = \Delta_t^{x,T}(f_t^T, \sigma_t^{x,T})\delta f_t^T + V_t^{x,T}(f_t^T, \sigma_t^{x,T})\delta\sigma_t^{x,T} + O(\delta t), \tag{16}$$

where Δ and V are the corresponding derivatives of the Black Scholes function, i.e. the Delta and Vega, and we have subsumed the Ito correction terms, and terms involving derivatives with respect to x, into the term[14] $O(\delta t)$.

Rearranging this equation, we can then deduce

$$\delta\Psi_t^x := \left[(\delta\Phi_t^{x,T_t} - r\Phi_t^{x,T_t}\delta t) - \Delta_t^{x,T_t}\delta f_t^{x,T_t}\right]/V_t^{x,T_t} = \delta\sigma_t^{x,T_t} + O(\delta t). \tag{17}$$

Note that $\delta\Psi_t^{x,T_t}$ here is the price innovation, net of financing costs, of a portfolio of an option and the futures contract. In fact replacing $(\delta\sigma_t^{x_i,T_t})_i$ by $(\delta\Psi_t^{x_i,T_t})_i$ in the PCA of the previous subsection does not significantly alter the result there. (This is not shown.) Moreover, the corresponding dynamic factors $\delta s_t^j := \frac{1}{\sqrt{\rho^j}}\sum_{i=1}^n e_i^j \delta\Psi_t^{x_i}$ can also be realized in terms of actual portfolios of options and the future.

In this subsection, we will apply a linear factor analysis to the vector of innovations $\delta\boldsymbol{\Psi}_t := (\delta\Psi_t^{x_1}, \delta\Psi_t^{x_2}, ..., \delta\Psi_t^{x_n})$. Our aim is to test for arbitrage opportunities. Our approach is conceptually parallel to that of Jones (2001), except that he attempts to model explicitly the non-linear nature of options returns, whereas we are relying on the Black Scholes Model to account sufficiently for the nonlinearity, so that we can apply linear techniques to the pricing dynamics of $\delta\boldsymbol{\Psi}_\mathbf{t}$.

In detail, we fit the linear model

$$\delta\boldsymbol{\Psi}_t = \sum_{j=1}^k \mathbf{g}^j(\delta s_t^j + \lambda^j \delta t) + \epsilon_t, \tag{18}$$

in which each $\mathbf{g}^j$ is an n vector of parameters, constrained such that each is orthogonal to its predecessors for lower j; the δs_t^j's are jointly normal, with mean zero, standard deviation $\sqrt{\delta t}$, and independent with respect to t and j; and the ϵ_t^i's making up ϵ_t are normally distributed, with mean zero and variance w^i, and independent with respect to

[14] Our numerical results will be presented on a weekly basis, i.e. taking $\delta t = 1$. Thus, we could delete the factor δt from our equations, but we keep it for the sake of clarity.

t and i. Under our linearity assumptions, and applying Proposition 1 of Section 2, this equation will hold if there are no arbitrage opportunities, and then λ^j represents the price of risk associated with the factor s_t^j. This equation represents our null hypothesis, and the alternative hypothesis is that the drifts are not so constrained, so that we can write

$$\delta\mathbf{\Psi}_t = \sum_{j=1}^{k} \mathbf{g}^j \delta s_t^j + \mu\delta t + \epsilon_t, \tag{19}$$

in which the n parameters $\mu := (\mu^1, ..., \mu^n)$ replace the k parameters $\lambda^1, ..., \lambda^k$. If Equation (19) is favored over Equation (18), then this is evidence either that there are arbitrage opportunities, or that the model is not adequate to represent the dynamic of the options prices.

Our technique for fitting the null and its alternative, is simply MLE, using the fact that under the null, then $d\mathbf{\Psi}_t \sim N(\sum_{j=1}^{k} \lambda^j \mathbf{g}^j \delta t, BB^T + W)$, where $B = [\mathbf{g}^1, ...\mathbf{g}^n]$ and $W = diag(w^1, ..., w^n)$, and under the alternative, then the mean $\sum_{j=1}^{k} \lambda^j \mathbf{g}^j$ should be replaced by μ.

We have implemented this test taking 9 moneyness values x^i, ranging over 0.90, 0.92, 0.94, 0.96, 0.98, 1.00, 1.02, 1.04, 1.06, and taking $k = 1$ and $k = 2$. The results are given in Table 5. With $k = 1$, i.e. the 1 factor model, the null hypothesis is rejected strongly in favor of the alternative: the difference in the log-likelihoods is 65.12, and under the null hypothesis, this should be distributed as $\frac{1}{2}$ times χ^2, with $n - k = 8$ degrees of freedom, and so the probability of the null is virtually zero. This result is in tune with those of Section 3 relating to the Delta-Vega hedged portfolios: First, note that our portfolio innovations $\delta\Psi_t^{x_i}$ represent the price innovations of Delta hedged options, and the factor δs_t^1 can be identified with the innovation of implied volatility. If price and volatility (corresponding to Delta and Vega) are the only random factors, and Equation (18) holds, then hedging the volatility between any pair of innovations will yield a portfolio whose return is essentially just the riskless return. The portfolios of Section 3 are Delta and Vega hedged, but their returns are not just the riskless return.

Turning to $k = 2$, i.e. the 2 factor model, the difference in the log-likelihoods tells us that this is overwhelmingly favored over the 1 factor model, both in null and alternative forms. But more striking, the alternative is now not favored over the null: the difference in the log-likelihoods is only 3.96, which under the null is distributed as $\frac{1}{2}$ times χ^2, with $n - k = 7$ degrees of freedom. This has a p-value of about 20%. Moreover, the factor loadings $\mathbf{g}^1$ and $\mathbf{g}^2$ are virtually the same as the principal components of the previous subsection, and so we can identify the factors as the "(implied) volatility", and the "smirk". The associated risk premia should be compared with the Sharpe ratio for investing in the underlying futures themselves, which we estimate at 0.11 (t-stat. 3.17) for our data. We see that the volatility premium is comparable in magnitude to this ratio, but negative, and the smirk premium is about twice the magnitude of this ratio, and positive. These signs are consistent with option prices being generally "high", and to the prices of options with low strikes, i.e out of the money puts, being "high" relative to the rest[15]. Thus, consistent with Backshi and Kapadia, etc, the market is prepared to pay a premium for options, since they provide insurance against volatility. The market is also prepared to pay a premium for OTM puts, causing the smirk, presumably to insure against market declines, but this does not represent an arbitrage opportunity, and can be explained in terms of the risk premium of the dynamic smirk factor. The risk premium on the smirk factor is rather high, relative to that on the index itself. Writing OTM puts is thus good business; but we suggest that this premium represents a reasonable rent for the put writers' expertise in hedging, and so it does not represent a market anomaly or disequilibrium.

3.3 Some Diagnostics on the Smirk Factor:

The factor innovations $\delta s_t^1, \delta s_t^2$ can be extracted from the data in the MLE framework, by minimizing the residual $||\delta \mathbf{\Psi}_t - \sum_{j=1}^k \mathbf{e}_t^j \delta x_t^j||^2$ for each t. Graphs of our 3 factor innovations $\delta f_t^{T_t}/f_t^{T_t}, \delta s_t^1, \delta s_t^2$ are given in Figure 5.

Table 6 presents some diagnostic statistics relating to these innovations. For clarity, in this table we write the implied volatility factor innovation ds_t^1 as $dimpl_t$, and the smirk factor innovation ds_t^2 as $dsmirk_t$. From Panel A, we see that at each moneyness level, the residuals from the 2 factor fit have standards deviation of about 12% of the standard deviation of the innovations $\delta \mathbf{\Psi}_t$ themselves. This is consistent with an R^2 for the model of about $1-(12\%)^2 = 98.5\%$, i.e. the 2 factors account for 98.5% of the innovations. Also, these innovations are not much correlated with each other, as one should expect from their construction, but they are significantly correlated with the futures return $df_t^{T_t}/f_t^{T_t}$.

[15]Note that the smirk factor, as given in Table 5 is negative for low strikes. This is arbitrary: If the signs on this factor were reversed, and the risk premium were reversed, the model would not be altered.

This is consistent with the results of Section 2, that the Black Scholes Delta hedge, which has been used to construct the innovation $\delta\mathbf{\Psi_t}$, is not perfect.

We see in Panel B, that the factor innovations are not autocorrelated. This indicates that they are distinct from the "buying pressure" factor identified by Bollen and Whaley (2003), which corresponds to option price increases, when there is more demand to buy them. Bollen and Whaley show that this price distortion is temporary, and in fact its innovation will be highly negatively autocorrelated, and self correcting within a few days, as more options are written to satisfy the demand.

Panel C presents a regression of the underlying return cubed, and normalized by the implied volatility at the beginning of the time step, against the contemporaneous innovation in the smirk factor, and its lags. The significant result shows that the smirk factor can be used to hedge the skew in the underlying return. This is not surprising, since the smirk innovation itself is skewed, since it corresponds to short puts and long calls. This panel also presents a regression of the underlying returns squared and normalized by the implied volatility at the beginning of the time step, against the implied volatility factor innovation. The result is again significant, but not surprising, since this return squared is a proxy for the innovation in the underlying volatility.

In Panel D we investigate whether the implied volatility and smirk factors themselves (not their innovations) can predict respectively the underlying volatility and skew, proxied respectively by the underlying return squared, and the underlying return cubed and normalized by the implied volatility. We characterize the implied volatility and smirk factors by projecting the implied volatility vector onto the corresponding eigenvector. Also, we include in the regression the time to maturity of the option at each date, to obviate any maturity effects. The implied volatility factor can, as expected, predict the underlying volatility. This result would seem necessary, under the assumption that volatility is a dynamic factor, which is reflected in option prices. However, it seems that the smirk factor cannot predict the underlying skew. It seems that the options smirk, unlike the implied volatility, is not determined in relation to the underlying dynamic. This negative result has consequences for constructing a no-arbitrage option pricing model incorporating our dynamic smirk factor. Any such model could be characterized as an Equivalent Martingale measure on the underlying dynamic, which differs from the objectively realized measure by a Radon Nykodym derivative representing a risk premium. Our result suggests that in a no-arbitrage model, the smirk dynamic must reside purely in this risk premium, and not in the objectively realized dynamic. It seems that the options smirk represents a dynamic "crash fear", that cannot be detected in the underlying returns.

3.4 Hedging with the Smirk Factor:

We now present a Delta-Vega-Smirk hedging exercise, based on the factors isolated above. We will hedge the option at each moneyness level 0.90, 092, ...1.06, with all the options together, at the other moneyness levels. Taking the the first k components, using Equation (18) above, and recalling that $\delta s_t^j := \frac{1}{\sqrt{\rho^j}} \sum_{\ell=1}^n e_\ell^j \delta \Psi_t^{x_\ell}$ and $\mathbf{g}^j = \sqrt{\rho^j}\mathbf{e}^j$, then for each i, the portfolio $\delta\Psi_t^{x_i} - \sum_{j=1}^k \sum_{l=1}^n e_i^j e_\ell^j \delta\Psi_t^{x_\ell}$ is hedged against the first k factors. This portfolio has weight $1 - \sum_{j=1}^k (e_i^j)^2$ in $\delta\Psi_t^{x_i}$, and so the portfolio in which 1 unit of the option with moneyness i is hedged with the other options, has price innovation

$$\frac{\left(\delta\Psi_t^{x_i} - \sum_{j=1}^k \sum_{l=1}^n e_i^j e_\ell^j \delta\Psi_t^{x_\ell}\right)}{\left(1 - \sum_{j=1}^k (e_i^j)^2\right)}. \tag{20}$$

Table 7 shows that standard deviations of these price innovations, for values $k = 0, 1, 2$, which correspond to the Delta hedge, the Delta-Vega hedge, and the Delta-Vega-Smirk hedge, respectively. The e_i^j's used to construct these portfolios are calculated separately for each date, via the Principal Components Analysis of Section 3.1, and using only data for previous dates, i.e. out of sample. This ensures that the hedging could have been done in real-time. In detail, for each date we use the covariance matrix, calculated with previous data, but with exponentially decaying weight factor, such that the data 0.25 years before the date being considered, is weighted half as much as the most recent data. For this reason, we have started the hedging from 1991, instead of 1990.

We see from table 7, that the Delta-Vega hedge is always much better than the Delta hedge, and the Delta-Vega-Smirk is significantly better that the Delta-Vega hedge, except for middle moneyness levels, where the influence of the smirk factor is relatively small. For $k = 0$ and 1, the Sharpe ratios of the corresponding returns reflect the risk premia of the unhedged exposures.

3.5 Discussion of our results:

We now discuss our analysis of the smirk, and how it relates to some of the previous literature. First, like us, many authors, notably Bakshi, Cao and Chen (1997), Bates

(2000) and Pan (2001), and very recently Liu and Pan (2003), Eraker, Johannes and Polson (2003) and Eraker (2004) have studied index options pricing and hedging, taking the price and the volatility to be state variables, as we have done, but with Poisson jumps as a third source of risk, in place of our dynamic smirk factor. It is pertinent therefore to ask whether our smirk factor can be interpreted in the context of these models. We note that in these models, the Poisson jump is not a sate variable, since the propensity for the market to jump is dependent only on the volatility. Therefore the jump component cannot make the smirk dynamic, which seems at odds with the behavior that we have documented. On the other hand, Liu and Pan show that the jump component in their model can be hedged in a portfolio containing at least 2 options[16]. From our Table 7, we see that the option portfolios of our Section 3.4 are very effective in hedging our price, volatility and smirk factors, and so it is natural to ask whether our smirk factor is the same as the Poisson jump component in these papers, which is isolated by Eraker, Johannes and Polson, in their Figure 3. Comparing this Figure with Figure 5C of the present paper, there seems to be no resemblance between our smirk innovations and their jumps: our smirk innovations seem to be white noise, whereas the jumps are quite infrequent impulses, occurring about once per year. In fact Bakshi, Cao and Chen find that their third jump factor can improve the pricing of their model, by better representing the smirk, but not the hedging ability. We conclude that our smirk factor is not accounted for by the jump component of these papers.

In fact our approach is different from that of the papers mentioned above, in that these concentrate on modelling the underlying dynamic and associated risk premia, such that options can be accurately priced in terms of the Equivalent Martingale Measure (EMM). By contrast, and consistent with Coval and Shumway (2001) and Bakshi and Kapadia (2003), we have taken the option prices as given by the market, and concentrated on directly modelling their returns. Thus, we cannot speak directly of the "pricing ability" of our approach, but only whether the market option prices are arbitrage free and consistent with reasonable levels of risk premia, and we find that they are.

In contrast to our smirk risk premium estimation, we have mentioned above that in her GMM fit of the jump diffusion model to index option prices, Pan (2001) estimates a mean jump size of about -0.9% but a risk neutral mean jump size of -19.2%. This represents an

[16]This paper works with the relatively simple model, in which the jump size is constant, but intuitively, their procedure will hedge most of the risk, if the jump size is random.

unreasonably large premium of 18.3%. Using more recent MCMC econometric methods, Eraker gets a lower premium of $\mu_y^Q = 7.9\%$ (see his Table III), but this still seems very high[17].

Our diagnostic tests failed to show that the smirk can predict the underlying returns skew, and so we suspect that it might be problematic to identify our dynamic smirk with some aspect of the underlying return. The EMM must exist, if there are no arbitrage opportunities, but we suspect that the smirk corresponds to a dynamic risk premium, which is not tied to the underlying returns.

We have mentioned the conceptual parallels and contrasts between our approach and that of Jones (2001). The basic difference is that he uses a semi-parametric framework to accommodate the nonlinearity of the problem, whereas we work with the residuals from a Black Scholes approximation, and rely on a linear analysis of these. We agree with Jones, in favoring a 3 factor diffusive model, but it seems unclear whether Jones' factors can encompass our smirk factor. With his factors, Jones is not able to account for all option returns in terms of risk premia associated with his factors.

Finally, our results are distinct from, but not in disagreement with, those of Bollen and Whaley (2003), who identify a "price pressure" factor for options, which is essentially the excess of demand interest over supply interest. They show that this factor influences the price. This is perhaps not surprising, but is problematic for no-arbitrage models, under which options prices should be completely determined by the underlying factors, with supply/demand imbalances being immediately obviated. However, they show that the effect on the option price is temporary, lasting only a few days, and this means that the effect will not be very strong in our data, because we are working with a weekly time step. Also, our smirk factor is not temporary, since its increments are not negatively autocorrelated.

4 Summary and Conclusions

Our aim in this paper has been to analyze the "smirk" in S&P500 futures options prices, which is the stylized fact that options with lower strikes have a higher Black-Scholes

[17]Eraker, Johannes and Polson (2003) note, on page 1291, that a more modest risk premium of 2% seems to account for the smirk in the jump diffusion model. However, this result is informal and predates the estimate $\mu_y^Q = 7.9\%$ of Eraker.

implied volatility, so that the option pricing "smile" is lop-sided.

Our approach has been first to construct portfolios of options and the underlying future, which are Delta hedged, i.e. hedged against movements in the underlying price, and Vega hedged, i.e. hedged against movements in the options implied volatility. These earn excess returns, consistent with a "market imperfection" interpretation of the smirk, i.e. that the out of the money puts are "too expensive", relative to the out of the money calls. Our hedges have been constructed using the Black Scholes Model, but we have also included checks to account for mis-specification of this model. We conclude from this exercise that either there really are market imperfections, of that we must search for factors beyond price and volatility, to explain the smirk. By including an extra, deep out of the money put in the Delta-Vega hedged portfolio, to make it "crash neutral", we can conclude that the smirk does not represent an insurance premium against a large market crash. These tests are extensions of the tests of Coval and Shumway (2001) and Bakshi and Kapandia (2003), who deal with Delta (but not Vega) hedged portfolios.

Second, we have isolated the smirk as a third, diffusive factor, and shown that the returns to the Delta-Vega hedged portfolios do not represent arbitrage opportunities, but can be explained in terms of the risk premium of this factor. Our technique for this has been Principal Components Analysis, applied to the Delta hedged option price innovations, and an estimation of a linear priced factor model for these innovations. This risk premium is about twice the risk premium for investing in the underlying S&P500 index itself, and we suggest that this is not excessive, but represents a reasonable rent for the expertise of writing and hedging these options. We show that this dynamic smirk factor is useful for hedging options portfolios, but in contrast to the volatility factor, it does not seem able to predict any aspect of the underlying return. The smirk factor seems to represent a dynamic aversion to market falls, which is not reflected in the underlying index futures prices.

Finally, we have discussed our results in connection with other recent work on the smirk. Most authors have modelled the smirk in terms of Poisson jumps in the underlying price. However, this is not able to make the smirk dynamic, and fails to account adequately for its magnitude. Also, including the Poisson jump component is not able to improve the hedging performance of the resulting models.

TABLE 1

Regressions of the Weekly Returns to Delta Neutral Straddles

(t-statistics in brackets)

Call mon.	Put mon.	Av. call position	Av. put position	Regressor	1st regr.		2nd regr.		Residual Sharpe ratio
1.00	0.92	2.015	7.792	Unit const.	-0.00322	(-5.07)***	-0.00302	(-4.29)***	
				Futures ret.			-0.08879	(-1.17)	-0.203
1.00	0.94	2.165	5.914	Unit const.	-0.00321	(-5.08)***	-0.00299	(-4.31)***	
				Futures ret.			-0.09627	(-1.29)	-0.201
1.00	0.96	2.351	4.589	Unit const.	-0.00301	(-4.67)***	-0.00276	(-3.89)***	
				Futures ret.			-0.11614	(-1.50)	-0.182
1.00	0.98	2.602	3.682	Unit const.	-0.00280	(-4.31)***	-0.00251	(-3.55)***	
				Futures ret.			-0.12843	(-1.63)	-0.163
1.00	1.00	2.964	3.089	Unit const.	-0.00292	(-3.89)***	-0.00267	(-3.22)***	
				Futures ret.			-0.11204	(-1.35)	-0.159
1.02	1.00	3.595	2.696	Unit const.	-0.00255	(-3.62)***	-0.00225	(-2.95)**	
				Futures ret.			-0.13518	(-1.66)*	-0.138
1.04	1.00	5.245	2.355	Unit const.	-0.00223	(-3.25)***	-0.00189	(-2.62)**	
				Futures ret.			-0.15513	(-1.93)*	-0.118
1.06	1.00	9.129	2.100	Unit const.	-0.00197	(-2.86)**	-0.00160	(-2.33)**	
				Futures ret.			-0.16661	(-2.12)*	-0.100

Notes:

- This table gives the results of regressions, in which the dependent variable is the weekly return to the straddle, comprising puts and calls at the moneyness levels indicated, and with amounts, so that the straddle is Delta neutral, in the Black Scholes framework.
- In the first regression, the regressor is the unit constant only, and the salient point is that the coefficient is always negative.
- In the second regression, then futures return is also included as a regressor, to compensate for any mis-specification in the hedge. The corresponding coefficient are sometimes significantly negative, indicating that there is some mis-specification, but the unit constant is usually still highly significant.
- The t-statistics are obtained using the Newey-West procedure, with 4 lags. Single, double and triple asterisk indicates significance at 5%, 1%, 0.1%, in a 2 tailed test.

TABLE 2

Regressions of the Weekly Returns to Delta and Crash Neutral Straddles

(t-statistics in brackets)

Call mon.	Put mon.	Av. call position	Av. put position	Regressor	1st regr.		2nd regr.		Residual Sharpe ratio
1.00	0.92	2.315	14.307	Unit const.	-0.00289	(-4.53)***	-0.00268	(-3.91)***	
				Futures ret.			-0.09319	(-1.23)	-0.170
1.00	0.94	2.467	9.216	Unit const.	-0.00289	(-4.53)***	-0.00268	(-3.91)***	
				Futures ret.			-0.09580	(-1.31)	-0.173
1.00	0.96	2.659	6.572	Unit const.	-0.00271	(-4.11)***	-0.00243	(-3.41)***	
				Futures ret.			-0.12280	(-1.59)	-0.155
1.00	0.98	2.926	5.007	Unit const.	-0.00249	(-3.78)**	-0.00219	(-3.10)**	
				Futures ret.			-0.13752	(-1.72)*	-0.136
1.00	1.00	3.315	4.056	Unit const.	-0.00262	(-3.43)**	-0.00235	(-2.82)**	
				Futures ret.			-0.12388	(-1.49)	-0.136
1.02	1.00	3.951	3.524	Unit const.	-0.00226	(-3.15)**	-0.00193	(-2.53)**	
				Futures ret.			-0.14689	(-1.79)*	-0.115
1.04	1.00	5.654	3.063	Unit const.	-0.00196	(-2.79)**	-0.00158	(-2.19)**	
				Futures ret.			-0.16747	(-2.08)*	-0.096
1.06	1.00	9.667	2.721	Unit const.	-0.00169	(-2.39)**	-0.00129	(-1.86)*	
				Futures ret.			-0.17936	(-2.27)**	-0.078

Notes:

- This table gives the results of regressions, in which the dependent variable is the weekly return to the straddle, as in Table 1, comprising puts and calls at the moneyness levels indicated, together with a deep out of the money put, in an amount to make the portfolio crash neutral.
- In the first regression, the regressor is the unit constant only, and the salient point is that the coefficient is always negative.
- In the second regression, then futures return is also included as a regressor, to compensate for any mis-specification in the hedge. The corresponding coefficients are usually not significant, and do not greatly effect on the coefficient of the constant.
- The t-statistics are obtained using the Newey-West procedure, with 4 lags. Single, double and triple asterisk indicates significance at 5%, 1%, 0.1%, in a 2 tailed test.

TABLE 3
Regressions of the Weekly Returns to Delta-Vega Neutral Option Portfolios

(t-statistics in brackets)

Call mon.	Put mon.	Av. call position	Av. put position	Regressor	1st regr.		2nd regr.		3nd regr.		Residual Sharpe ratio
0.98	1.02	-0.80855	0.80144	Unit const.	0.00021	(1.91)*	0.00018	(1.63)	0.00029	(2.96)**	
				Futures ret.			0.01407	(1.82)*	0.01862	(1.90)*	
				ATM Straddle ret.					0.04057	(1.42)	0.120
1.00	1.00	-1.00297	1.00330	Unit const.	-0.00005	(-0.78)	-0.00006	(-0.89)	0.00002	(0.73)	
				Futures ret.			0.00811	(1.37)	0.01211	(1.45)	
				ATM Straddle ret.					0.03567	(1.23)	0.019
1.02	0.98	-1.36241	1.36189	Unit const.	-0.00015	(-4.39)***	-0.00017	(-4.86)***	-0.00017	(-3.92)***	
				Futures ret.			0.01100	(3.81)***	0.01115	(3.73)***	
				ATM Straddle ret.					0.00133	(0.23)	-0.226
1.04	0.96	-2.18152	1.93301	Unit const.	-0.00036	(-4.64)***	-0.00041	(-5.09)***	-0.00042	(-4.14)***	
				Futures ret.			0.02224	(3.29)***	0.02191	(3.01)**	
				ATM Straddle ret.					-0.00296	(-0.21)	-0.245
1.06	0.94	-4.13131	2.77609	Unit const.	-0.00060	(-4.36***)	-0.00068	(-4.95)***	-0.00068	(-3.93)**	
				Futures ret.			0.03543	(3.18)**	0.03551	(2.86)**	
				ATM Straddle ret.					0.00069	(0.03)	-0.232

Notes:

- This table gives the results of regressions, in which the dependent variable is the weekly return to a Delta Vega neutral portfolio, which is long the futures, long the put, and short the call, at the moneyness levels indicated.
- When the put and call are out of the money (put mon. < 1.00, call mon. > 1.00), then the portfolio should lose money, on the basis that out of the money puts are too dear relative to the calls. We see that it does, since the coefficient of the first regression is negative. The results for the other strikes are also consistent with this notion.
- In the second regression, the futures return is also included as a regressor, to compensate for any residual exposure to the price factor in the hedge. The presence of the futures return does not make a significant difference to the coefficient on the constant.
- In the third regression, the return to the at the money straddle is also included, to account for any residual exposure to the volatility factor. This also does not make a significant difference to the coefficient on the constant.
- The t-statistics are obtained using the Newey-West procedure, with 4 lags. Single, double and triple asterisk indicates significance at 5%, 1%, 0.1%, in a 2 tailed test.

TABLE 4

Regressions of the Weekly Returns to Delta-Vega Neutral Crash Proof Portfolios

(t-statistics in brackets)

Call mon.	Put mon.	Av. call position	Av. put position	Regressor	1st regr.		2nd regr.		3nd regr.		Residual Sharpe ratio
0.98	1.02	-0.835	0.752	Unit const.	0.00018	(2.01)*	0.00016	(1.72)*	0.00024	(3.01)***	
				Futures ret.			0.01166	(1.91)*	0.01522	(1.94)*	
				ATM Straddle ret.					0.03170	(1.34)	0.120
1.00	1.00	-1.003	1.007	Unit const.	-0.00004	(-0.76)	-0.00005	(-0.88)	0.00002	(0.76)	
				Futures ret.			0.00722	(1.41)	0.01067	(1.48)	
				ATM Straddle ret.					0.03085	(1.23)	0.020
1.02	0.98	-1.328	1.466	Unit const.	-0.00012	(-3.89)***	-0.00014	(-4.37)***	-0.00013	(-3.35)***	
				Futures ret.			0.00997	(3.57)***	0.01035	(3.68)***	
				ATM Straddle ret.					0.00336	(0.62)	-0.188
1.04	0.96	-2.086	2.251	Unit const.	-0.00029	(-4.09)***	-0.00034	(-4.52)***	-0.00034	(-3.57)***	
				Futures ret.			0.02005	(3.05)**	0.02017	(2.91)**	
				ATM Straddle ret.					0.00101	(0.08)	-0.209
1.06	0.94	-3.882	3.574	Unit const.	-0.00050	(-3.85)***	-0.00057	(-4.47)***	-0.00055	(-3.46)***	
				Futures ret.			0.03329	(3.22)***	0.03402	(2.97)**	
				ATM Straddle ret.					0.00651	(0.32)	-0.201

Notes:

- This table gives the results of regressions, in which the dependent variable is the weekly return to a Delta Vega neutral portfolio, which is long the futures, long the put, and short the call, at the moneyness levels indicated and as in Table 3, together with a deep out of the money put, to make the portfolio crash-neutral.
- Apart from the extra deep out of the money out, this regressions of this table are the same as in Table 3, and the results are consistent, but somewhat weaker.
- The t-statistics are obtained using the Newey-West procedure, with 4 lags. Single, double and triple asterisk indicates significance at 5%, 1%, 0.1%, in a 2 tailed test.

TABLE 5

Testing the Linearized Factor Model

Panel A - log-likelihood values:

1 factor NULL Model	18119.91
1 factor ALTERNATIVE Model	18185.03
2 factor NULL Model	19844.41
2 factor ALTERNATIVE Model	19848.37

Panel B - Parameter Estimates for the 2 factor NULL Model:

Moneyness	0.90	0.92	0.94	0.96	0.98	1.00	1.02	1.04	1.06
Factor loadings $\mathbf{g}^1$	0.0135	0.0141	0.0143	0.0149	0.0155	0.0154	0.0158	0.0159	0.0153
(Std errors)	0.0006	0.00057	0.0005	0.0005	0.0005	0.0006	0.0006	0.0006	0.0007
Factor loadings $\mathbf{g}^2$	-0.0050	-0.0042	-0.0034	-0.0023	-0.0007	0.0008	0.0028	0.0045	0.0060
(Std errors)	0.0002	0.0001	0.00011	0.00011	0.00009	0.00008	0.00009	0.00023	N/A
Risk premium λ^1 (and its standard error)			-0.181			(0.0613)			
Risk premium λ^2 (and its standard error)			0.250			(0.0632)			
SD model errors $\sqrt{\mathbf{w}}$	0.00338	0.00217	0.00121	0.00146	0.00178	0.00170	0.00105	0.00264	0.00529
(Std errors)	0.00011	0.00008	0.00005	0.00005	0.00007	0.00005	0.00006	0.00008	0.00016

Notes:

- This table gives the results of fitting the null model represented by $\delta\mathbf{\Psi}_t = \sum_{j=1}^{k} \mathbf{g}^j(\delta s_t^j + \lambda^j \delta t) + \epsilon_t$, against the alternative model $\delta\mathbf{\Psi}_t = \sum_{j=1}^{k} \mathbf{g}^j \delta s_t^j + \mu\delta t + \epsilon_t$.
- In these equations, $\delta\mathbf{\Psi}_t$ is the vector of weekly innovations in the option prices for the moneyness values given, after these innovations have been delta hedged, and their financing costs have been included.
- The number of factors in these models is represented by k in these equations.

TABLE 6

Diagnostic Tests on the Implied Volatility and Smirk Factors

PANEL A - Correlations and residual standard deviations

Moneyness x_i	0.90	0.92	0.94	0.96	0.98	1.00	1.02	1.04	1.06
SD of $d\Psi_t^{x_i}$	0.01496	0.01510	0.01492	0.01532	0.01573	0.01556	0.01610	0.01667	0.01717
Residual SD	0.00274	0.00159	0.00141	0.00171	0.00218	0.00223	0.00195	0.00141	0.00303
Correlations		$df_t^{T_t}/df_t^{T_t}$	$dimpl_t$	$dsmirk_t$					
	$dimpl_t$	-0.16069	1						
	$dsmirk_t^2$	-0.23729	-0.01767	1					
SDs		0.01954	1.00352	1.06017					

PANEL B - Autocorrelation tests

Regress	$dimpl_t$	versus	$dimpl_{t-1}$	$dimpl_{t-2}$	$dimpl_{t-3}$	1
Coefficient			-0.0448	-0.00136	-0.0645	0.00925
t-statistic			(-0.73)	(-0.02)	(-1.64)	(0.21)
Regress	$dsmirk_t$	versus	$dsmirk_{t-1}$	$dsmirk_{t-2}$	$dsmirk_{t-3}$	1
Coefficient			0.0705	-0.0370	0.0022	-0.0200
t-statistic			(1.21)	(-0.84)	(0.06)	(-0.43)

TABLE 6 continued

PANEL C - Test for correlation between factor innovations and return moments

Regress	$\left(\frac{df_t^{T_t}}{f_t^{T_t}\times impl_t}\right)^3$	versus	$dsmirk_t$	$dsmirk_{t-1}$	$\left(\frac{df_t^{T_t}}{f_t^{T_t}\times impl_t}\right)$	$\left(\frac{df_t^{T_t}}{f_t^{T_t}\times impl_t}\right)^2$	1	
Coefficient			0.002227	-0.00007204	0.05171	-0.1.332	0.0007898	
t-statistic			(5.00)***	(-0.43)	(9.92)***	(-2.87)	(1.68)	
Regress	$\left(\frac{df_t^{T_t}}{f_t^{T_t}\times impl_t}\right)^2$	versus	$dimpl_t$	$dimpl_{t-1}$	$\left(\frac{df_t^{T_t}}{f_t^{T_t}\times impl_t}\right)$	1		
Coefficient			0.01609	-0.0005297	-0.01854	0.01352		
t-statistc			(10.76)***	(-0.71)	(-1.39)	(18.41)***		

PANEL D - Test for prediction power of the factors

Regress	$\left(\frac{df_t^{T_t}}{f_t^{T_t}\times impl_t}\right)^3$	versus	$smirk_t$	$\left(\frac{df_t^{T_t}}{f_t^{T_t}\times impl_t}\right)^2$	$\left(\frac{df_t^{T_t}}{f_t^{T_t}\times impl_t}\right)$	$T_t - t$	$(T_t - t)^2$	1
Coefficient			0.03745	-0.01971	0.001583	0.000223	-0.0006065	-0.00002068
t-statistic			(1.95)	(-1.20)	(6.96)***	(1.26)	(-1.22)	(-1.75)
Regress	$\left(\frac{df_t^{T_t}}{f_t^{T_t}}\right)^2$	versus	$impl_t$	$\frac{df_t^{T_t}}{f_t^{T_t}}$	$T_t - t$	1		
Coefficient			2.388	-0.003702	-0.0007278	0.0001063		
t-statistic			(6.78)***	(-0.76)	(-2.13)	(1.12)		

Notes:

- In the first regression of Panel C, the dependent variable is the underlying returns cubed, and normalized by the implied volatility. This is chosen as a proxy for the underlying skew. We see that it is strongly related to the contemporaneous innovation in the smirk factor.
- The second regression of Panel C shows similarly, that a proxy for the innovation in the underlying volatility is strongly related to the innovation in the implied volatility.
- The first regression of Panel D take the same proxy for the underlying skew, and asks if it can be predicted by the smirk factor. The answer is negative.
- The second regression of Panel D asks the same question for the second moment, rather than the third, i.e. whether the implied volatility factor can predict the underlying volatility, and the answer is affirmative.

TABLE 7

Residual Standard Deviations of Hedged Returns

		Residual SD	Sharpe Ratio	(t-statistic)
Moneyness 0.90		Residual SD	Sharpe Ratio	(t-statistic)
	$k = 0$ (Delta hedge)	0.01413	-0.249	$(-5.49)^{***}$
	$k = 1$ (Delta-vega hedge)	0.00682	-0.177	$(-3.76)^{***}$
	$k = 2$ (Delta-vega-Smirk hedge)	0.00387	-0.051	(-1.17)
Moneyness 0.92		Residual SD	Sharpe Ratio	(t-statistic)
	$k = 0$ (Delta hedge)	0.01420	-0.231	$(-5.18)^{***}$
	$k = 1$ (Delta-vega hedge)	0.00559	-0.162	$(-3.29)^{***}$
	$k = 2$ (Delta-vega-Smirk hedge)	0.00237	0.034	(0.84)
Moneyness 0.94		Residual SD	Sharpe Ratio	(t-statistic)
	$k = 0$ (Delta hedge)	0.01397	-0.236	$(-5.30)^{***}$
	$k = 1$ (Delta-vega hedge)	0.00424	-0.227	$(-4.25)^{***}$
	$k = 2$ (Delta-vega-Smirk hedge)	0.00176	-0.067	(-1.53)
Moneyness 0.96		Residual SD	Sharpe Ratio	(t-statistic)
	$k = 0$ (Delta hedge)	0.0141	-0.215	$(-4.89)^{***}$
	$k = 1$ (Delta-vega hedge)	0.0034	-0.195	$(-3.98)^{***}$
	$k = 2$ (Delta-vega-Smirk hedge)	0.0019	0.002	(0.06)
Moneyness 0.98		Residual SD	Sharpe Ratio	(t-statistic)
	$k = 0$ (Delta hedge)	0.0144	-0.192	$(-4.42)^{***}$
	$k = 1$ (Delta-vega hedge)	0.0028	-0.098	$(-2.07)^{*}$
	$k = 2$ (Delta-vega-Smirk hedge)	0.0026	0.023	(0.56)
Moneyness 1.00		Residual SD	Sharpe Ratio	(t-statistic)
	$k = 0$ (Delta hedge)	0.0145	-0.175	$(-4.00)^{***}$
	$k = 1$ (Delta-vega hedge)	0.0027	0.013	(0.31)
	$k = 2$ (Delta-vega-Smirk hedge)	0.0025	0.024	(0.53)
Moneyness 1.02		Residual SD	Sharpe Ratio	(t-statistic)
	$k = 0$ (Delta hedge)	0.01501	-0.140	$(-3.30)^{***}$
	$k = 1$ (Delta-vega hedge)	0.00383	0.143	$(3.04)^{**}$
	$k = 2$ (Delta-vega-Smirk hedge)	0.00245	0.042	(0.93)
Moneyness 1.04		Residual SD	Sharpe Ratio	(t-statistic)
	$k = 0$ (Delta hedge)	0.01587	-0.105	$(-2.45)^{**}$
	$k = 1$ (Delta-vega hedge)	0.00592	0.186	$(3.91)^{***}$
	$k = 2$ (Delta-vega-Smirk hedge)	0.00221	0.032	(0.71)
Moneyness 1.06		Residual SD	Sharpe Ratio	(t-statistic)
	$k = 0$ (Delta hedge)	0.01669	-0.076	$(-1.73)^{*}$
	$k = 1$ (Delta-vega hedge)	0.00902	0.173	$(3.40)^{***}$
	$k = 2$ (Delta-vega-Smirk hedge)	0.00529	-0.038	(-0.86)

Notes:

- This table gives the residual standard deviations for each option price innovation, hedged with the other options, and with financing costs included.
- The t-statistics correspond to the hypothesis that the corresponding Sharpe Ratio is zero.

Figure 1: The smirk for put and calls, on the date 06/25/97:

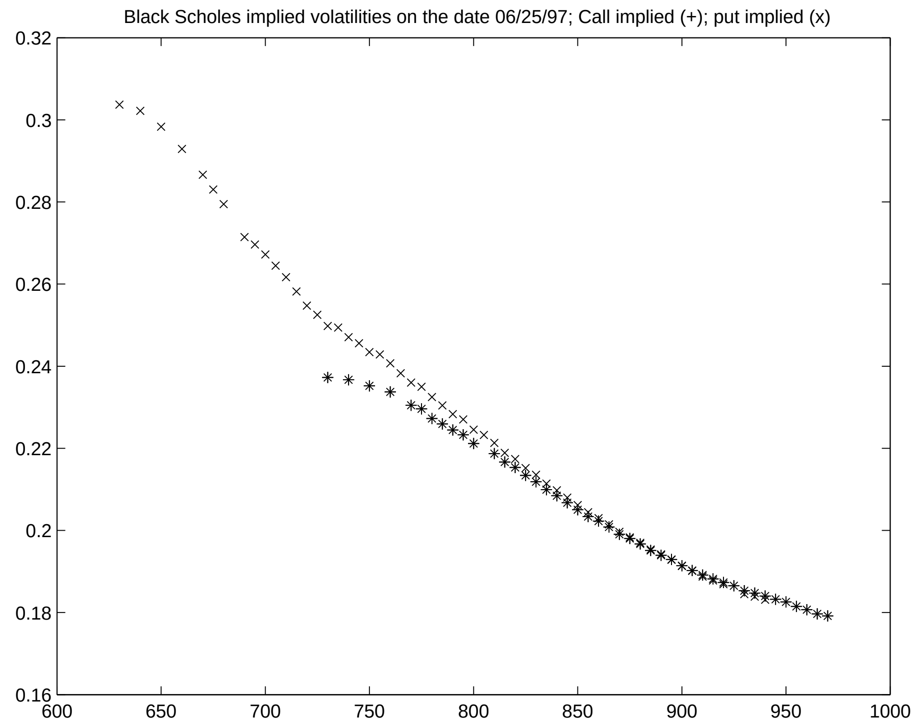

Figure 2: Traded Strikes and Open Interest for our Call Data, in terms of Moneyness:

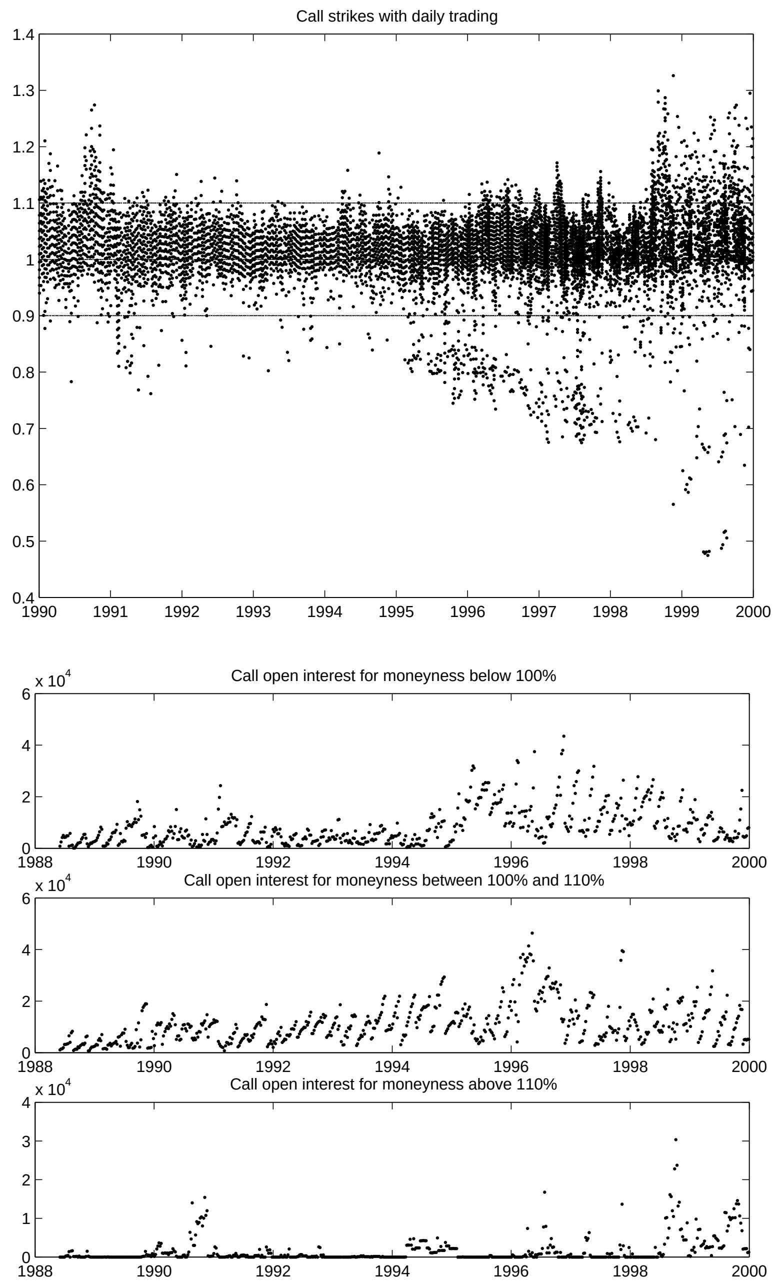

Figure 3: Traded Strikes and Open Interest for our Put Data, in terms of Moneyness:

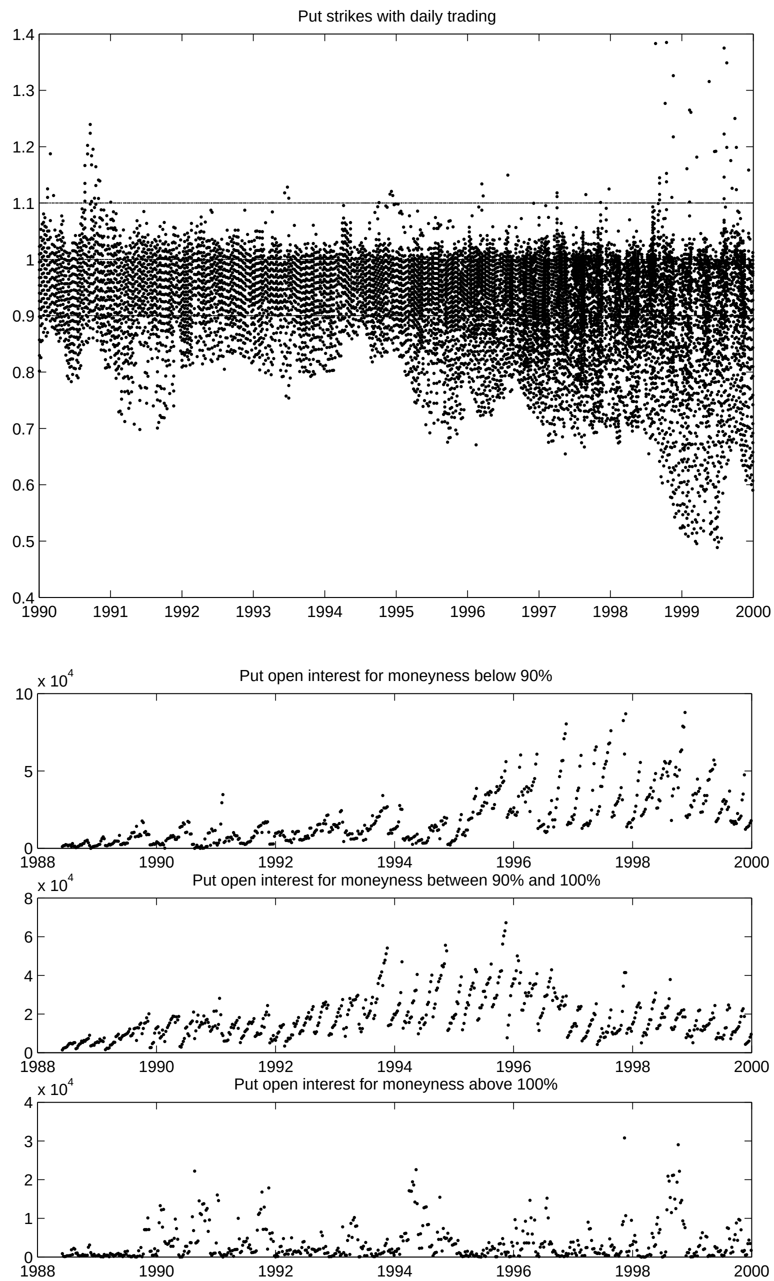

Figure 4: First and Second Principal Components of the Implied Volatility Innovations:

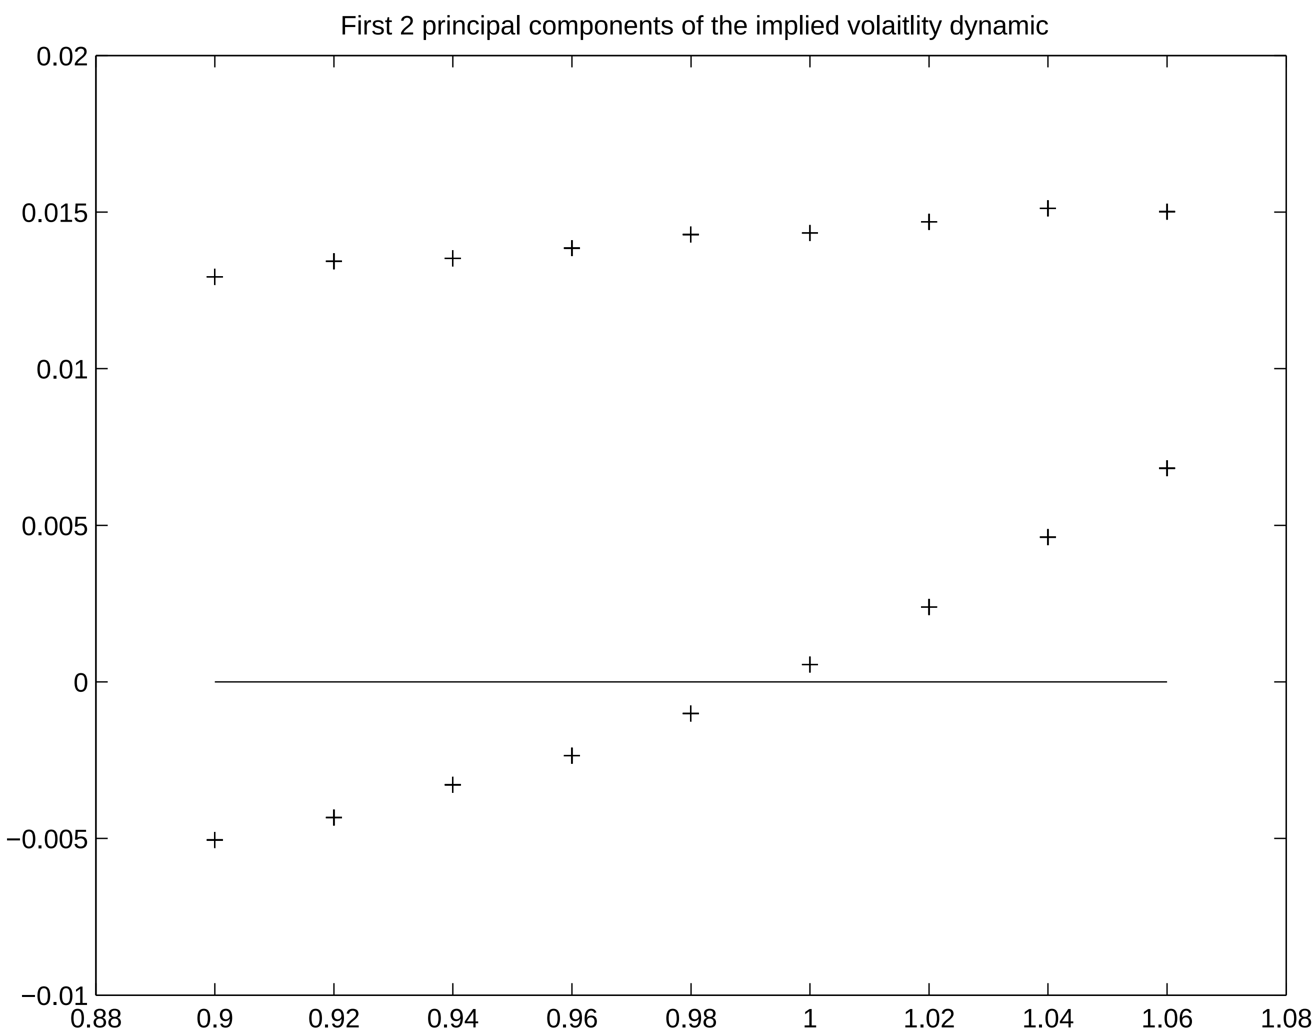

Figure 5: Time Series of our 3 Factor Innovations:

Futures return:

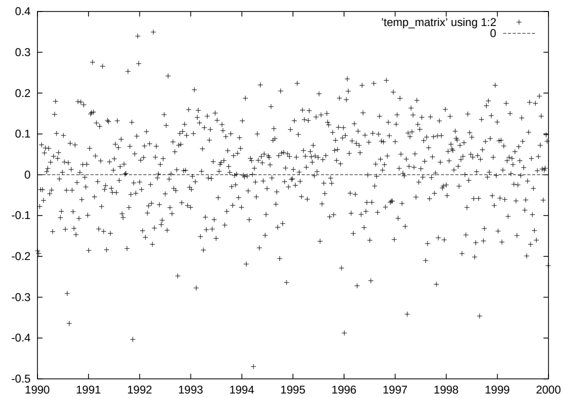

Implied volatility factor innovation:

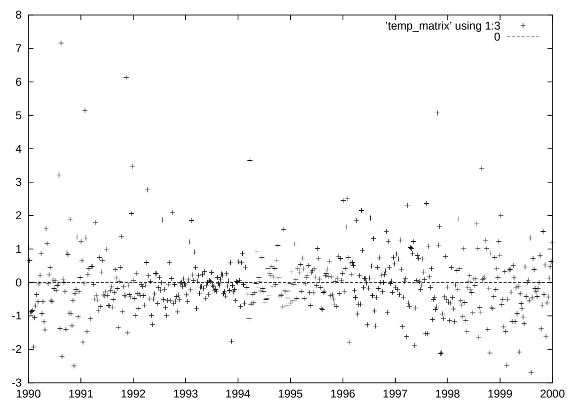

Smirk innovation:

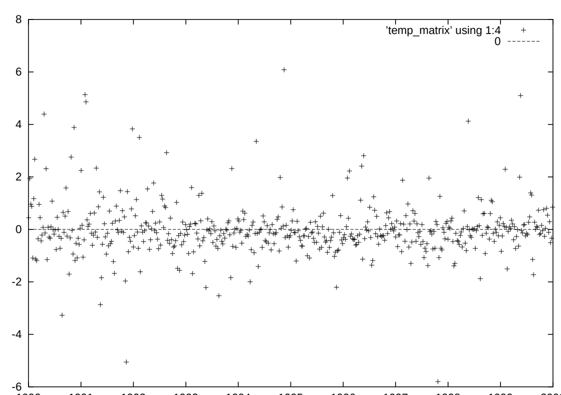

References

K. Amin and V. K. Ng, 1997, *Inferring Future Volatility from the Information in Implied Volatility in Eurodollar Options: A new Approach*, Review of Financial Studies, Vol 10, No 2, pp 333 - 367.

A. Buraschi and J. Jackwerth, 2001, *The Price of a Smile: Hedging and Spanning in Options Markets*, Review of Financial Studies, Vol 14, No 2, pp 495 - 527.

G. Bakshi, C. Cao and Zhiwu Chen, 1997, *Empirical performance of Alternative Option Pricing Models*, Journal of Finance, Vol LII, No 5, pp 2003 - 2049.

G. Bakshi and N. Kapadia, 2003, *Delta-Hedged Gains and the Negative Market Volatility Risk Premium*, Review of Financial Studies, Vol 16, pp 527-566.

D. Bates, 1991, *The Crash of '87: was it Expected? The Evidence from Options Markets*, Journal of Finance, Vol 46, pp 1009 - 1044.

D. Bates, 2000, *Post-'87 Crash Fears in the S&P500 Futures Options Market*, Journal of Econometrics, Vol 94, pp 181 - 238.

N. P. B. Bollen and R. E. Whaley, 2003, *Does Net Buying Pressure Affect the Shape of Implied Volatility Functions?*, Journal of Finance, forthcoming.

O. Bondarenko, 2002, *Why are Put Options So Expensive?*, Working Paper, University of Illinois at Chicago, USA.

J. D. Coval and T. Shumway, 2001, *Expected Option returns*, Journal of Finance, Vol LVI, No 3, pp 983 - 1009.

D. Duffie, 1989, "Futures Markets", Prentice Hall.

D. Duffie, 2001, *Dynamic Asset Pricing Theory*, Princeton University Press, 3rd Edition.

B. Eraker, 2004, *Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices*, Journal of Finance, Vol 59, pp 1367 - 1403.

B. Eraker, M. Johannes and N. Polson, 2003, *The Impact of Jumps in Volatility and Returns*, Journal of Finance, Vol 58, pp 1269 - 1300.

S. Heston, 1993, *A Closed Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options*, Review of Financial Studies, Vol 6, pp 327 - 343.

P. Jorion, 1995, *Predicting Volatility in the Foreign Exchange market*, Journal of Finance, Vol L, No. 2, pp 507 - 528.

C. S. Jones, *A Nonlinear Factor Analysis of S&P500 Index Options Returns*, Working Paper, Simon School of Business, University of Rochester, NY 14627, USA.

J. Liu and J. Pan, 2003, *Dynamic Derivatives Strategies*, Journal of Financial Economics, Vol 69, pp 401 - 430.

D. B. Nelson, 1990, *ARCH Models as Diffusion Approximations*, Journal of Econometrics, Vol 45, pp 7 - 38.

J. Pan, 2002, *The Jump-Risk premium Implicit in Option Prices: Evidence from an Integrated Times-Series Study*, Journal of Financial Economics, Vol 68, pp 3 - 50.

G. Skiadopoulos, S. Hodges and L. Clewlow, 2000, *The Dynamics of the S&P 500 Implied Volatility Surface* Review of Derivatives Research Vol 3, No 3, pp 263 - 282.