A Note on The Empirically Optimal Volatility for Black/Scholes Hedging in Various Market Conditions

David Murphy

Draft of 17th February 2010

Abstract

This note studies the optimal volatility for Black Scholes hedging in practice. The notion of hedging at a volatility which minimises the variance of the daily P/L of the hedged portfolio is introduced, and this measure is christened *prophetic volatility.* The relationships between realised volatility, implied volatility and prophetic volatility are studied for S&P 500 index options during the period 1999-2009.

The main result is that prophetic volatility is close to realised for the whole of the data set; for quiet market conditions (such as pertained from mid 2003 to mid 2005); and for the market turbulence induced by the Credit Crunch. Implied volatility differs from prophetic for the whole period, indicating that it is at best an imperfect estimate of the right hedge volatility.

We review situations when prophetic volatility differs significantly from realised, highlighting the practical importance of gamma weighted realised volatility. Finally, the spread required to at least break even from selling calls is analysed for both prophetic and implied volatility.

1 Introduction

The Black Scholes formula for option pricing focuses attention on the volatility of the underlying. Under the assumptions used to derive the formula, hedging at the correct volatility minimises the variance of the P/L of the portfolio. Two questions therefore arise immediately:

- What is the correct volatility to use in a Black Scholes universe?
- How does the deviation of the real world from the Black Scholes assumptions effect derivatives hedging?

Both questions have received a lot of attention in the literature. To pick two classes of suggestion from many, both historical volatilities[1] and (G)ARCH-derived volatilities have been investigated as candidates for the right hedge volatility [8, 18]. The information about the expected future return distribution of the underlying available from the options market has been elaborated, leading to local volatility models [9, 6]. Furthermore a plethora of more sophisticated models of the evolution of the underlying have been proposed, including a number in which volatility itself is another stochastic variable [3, 12]. See [11] or [13] for a further survey of this large and well developed area of research.

Many of these more sophisticated models give significant insight into the pricing of contingent claims. However, the simple Black Scholes model remains important not least because it is still used for much option pricing and hedging in the industry. Moreover the performance of the more advanced models in real world hedging is often worse than that of Black Scholes [1, 7, 15]. Therefore it is interesting to investigate the optimal volatility for Black Scholes hedging in a

[1] I.e. those calculated using various frequencies of observations, lengths of data, weighting schemes etc.

variety of market conditions. This note is investigates this measure. In particular, we find the optimal volatility for hedging short dated at the money S&P 500 index call options, and compare this empirically optimal volatility with realised and implied volatilities.

2 Black Scholes Prophetic Volatility

The hedging of a contingent claim is intended to reduce P/L variance: dealers want to be able to sell options for more than they are worth, or buy them for less, and then lock in those profits by hedging.

In the Black-Scholes setting, hedge ratios depend on the volatility used, σ. In particular, since the partial derivatives $\frac{\partial \delta}{\partial \sigma}$ and $\frac{\partial \gamma}{\partial \sigma}$ are significant in many practical situations, the performance of Black-Scholes delta hedging materially depends on the volatility used. Therefore an interesting question in the Black Scholes framework is 'which volatility produces hedge ratios which minimise P/L variance?'[2]. We will call this empirically optimal hedge volatility *prophetic volatility* and write σ_P. A dealer wishing to use Black Scholes hedging and able to prophesise the future evolution of the underlying accurately would generate hedge ratios using σ_P.

Prophetic volatility would be the same as realised volatility if the underlying followed a diffusion: since it doesn't [16], these two volatilities differ in general. Moreover, prophetic volatility would be the same as implied volatility if the prices of options was a perfect estimator of the cost of hedging them. It is well known that they aren't [2, 7], so prophetic volatility also differs from implied. In order to know how much this matters, we need to look at how significant these differences are in a variety of market conditions.

2.1 Deriving Prophetic Volatility

Suppose that we sell an at the money call option C expiring at time T. The daily returns of the underlying are $S(0), \ldots S(T)$. The decision to hedge at a volatility σ means that each day t we generate a hedge ratio $\delta = \frac{\partial C}{\partial S}$ using Black-Scholes with volatility σ. The hedged portfolio is then $\delta S - C$. This hedge is rebalanced each day[3] until expiry at T, with cash invested or borrowed at the then current one month Libor, dividends received from the hedge at the appropriate rate, and all trading being assumed to happen at mid market. Each day a P/L $P(t)$ is generated. The prophetic volatility is simply that volatility which minimises the variance of $P(t)$. In other words, hedging with prophetic volatility in the Black Scholes framework minimises the variance of the P/L of the hedged portfolio, and hence is a reasonable candidate for the optimal Black Scholes hedge volatility.

The underlying is governed by $dS = \mu S dt + \sigma_R S dW$ in the Black Scholes framework, where we write σ_R for the realised volatility. Making the volatility dependence explicit, we write $\delta_A S$ for the hedge for some call option at some arbitrary volatility σ_A. After dt, the call changes in value from C_A to $C_A + dC_A$; the hedge by $\delta_A dS$; and the cash position by $r(\delta_A S - C)dt$ ignoring dividends. Itô's lemma gives us

$$dC_A = \theta_A dt + \delta_A dS + \frac{1}{2}\sigma_A^2 S^2 \gamma_A dt$$

The mark to market P/L is $\theta_A dt + \delta_A dS + \frac{1}{2}\sigma_A^2 S^2 \gamma_A dt - \delta_A dS + r(\delta_A S - C)dt$, or simplifying

$$\frac{1}{2}(\sigma_R^2 - \sigma_A^2) S^2 \gamma_A dt$$

[2] In passing it is worth considering whether minimising the second moment of the daily P/L distribution is the right measure of a 'good' hedge. Some market participants consider that a measure that weights the tails of the P/L distribution more heavily might be better. However this view is not standard, and thus we stick to a definition of the risk of hedging based on P/L variance.

[3] In reality, much hedging is performed daily, so while the question of prophetic volatility for different rebalance periods is interesting, daily rebalances reflect typical industry practice.

The P/L is thus quadratic in the difference between the input vol σ_A and the realised volatility σ_R provided that the underlying follows a diffusion. In this case the optimal volatility to hedge at is simply realised. (A much more comprehensive discussion of the issues here can be found in [11].)

In practice for at the money options on the S&P 500 index, we find a similar concave dependence. Figure 1 shows the variation of the standard deviation of P/L of the portfolio as the volatility varies for three days in a volatile period and three days in a quieter market: other days exhibit similar curves.

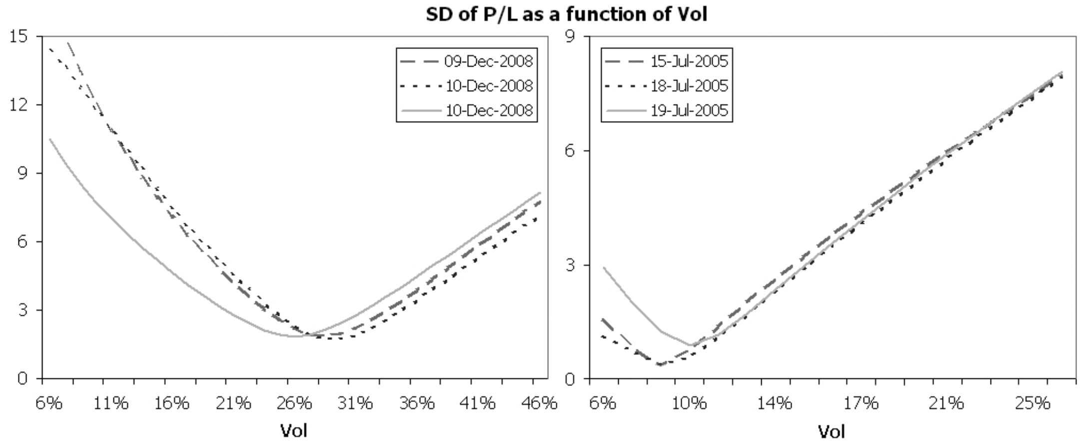

Figure 1: The variation of the P/L of the hedge strategy as the volatility used to generate Black-Scholes hedge ratios changes

In reality, then, while the best volatility to hedge at is a strong function of conditions, the standard deviation of daily P/L for our data shows a consistent, roughly quadratic dependence on hedge volatility, and thus we can estimate the prophetic volatility using a relatively simple minimum finder applied to the standard deviation of daily P/L function.

2.2 Model Design

The S&P 500 index makes an ideal tool for investigating prophetic volatility for several reasons:

- All of the necessary data - including dividend yields, Libor and daily quotes of implied volatility - is readily available;
- Both the underlying and options on it are highly liquid;
- The available data set includes two periods of highly elevated volatility: the last half of 2002; and the credit crunch of 2008-2009. It also includes extended periods of lower volatility such as that that pertained for most of 2003-2005. Therefore we can gain some insight into a variety of volatility conditions. Figure 2 illustrates historic and implied volatility for the period available.

We took the entire period 1999-late 2009, and considered selling a one month at the money index option for each day in the data set. This option was initial priced at the then current at the money volatility[4] σ_I, and delta hedged using the Black Scholes formula for each day in its life.

[4]This was kindly supplied by `ivolatility.com`. The results obtained using this ATM implied volatility also hold if we take the CBOE's VIX (the volatility corresponding to a fair value one month variance swap on the S&P 500) or the VXO (one month S&P 100 at the money implied volatility) as our implied volatility measure [17], and hence are robust under the precise implied volatility data used.

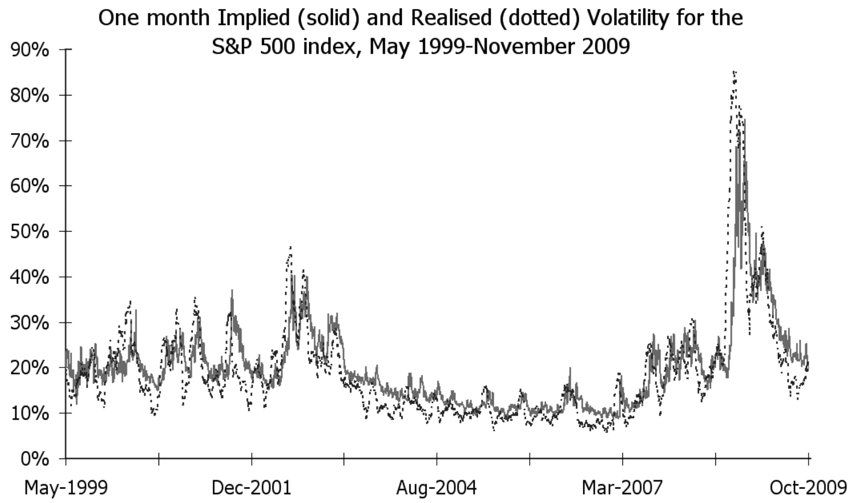

Figure 2: Historical and implied volatility May 1999-November 2009

A P/L was calculated for each day in the option's life, and the variance of those P/Ls was determined. The process was repeated for higher and slightly lower vols, and thus the vol which minimises the variance of P/Ls was found by convergence.

2.3 Market Practice and Accounting

We assume that once an option has been traded, the same volatility is used for marking and hedging during its entire life. This assumption is standard in the literature, and moreover not too unrealistic for the one month options studied. However it is worth bearing in mind that dealers sometimes remark volatilities more frequently than monthly (although they typically do not remark all volatilities daily).

Our assumption that the option is marked to an optimal volatility is more problematic. International accounting standards [14] – and investment bank product control departments – demand that options are marked to the current market price rather than some arbitrary volatility that happens to suit the trader's hedging. However this need not be problematic for two reasons: firstly, P/L from *marking* an option is unrealised, whereas P/L from following a hedge strategy is realised. In the long term, then, it is better to get the hedging right and the marking wrong rather than the converse. The second justification for looking at the optimal volatility to hedge at follows from this: some derivatives trading systems allow both mark vol and hedge vol inputs, allowing the trader to generate hedge ratios based on their view of the optimal volatility while marking at the correct level.

	Whole Period 1999-2009		Quiet Period mid 2003-2005		Crunch Period 2008-2009	
	Average	SD	Average	SD	Average	SD
Realised Vol σ_R	18.9%	11.8%	11.2%	2.5%	41.1%	19.6%
Prophetic Vol σ_P	18.9%	11.4%	11.3%	2.5%	40.2%	18.5%
Implied Vol σ_I	19.8%	8.9%	13.6%	2.7%	36.1%	13.4%
Realised Prophetic spread	0.003%	2.4%	-0.08%	1.0%	0.9%	3.7%
Implied Prophetic spread	0.92%	6.7%	2.3%	2.4%	-4.1%	15.8%

Figure 3: Summary of the three volatilities

3 Summary of Results

Perhaps slightly surprisingly, prophetic volatility tracks realised volatility very accurately, even in the market stress of 2008-2009: the table in figure 3 gives a overview. For the purposes of hedging, at least, estimating what the realised volatility will be is sufficient to hedge the option well on most occasions. Interestingly, implied volatility generates less effective hedge ratios, indicating that the alternation between fear and greed in volatility markets described by Derman [5] amongst others persists into the 21st century.

3.1 Do the volatilities differ?

Figure 4 illustrates the relationships between the three volatilities for the quiet markets of mid 2003-2005 and the Crunch period of 2008-2009.

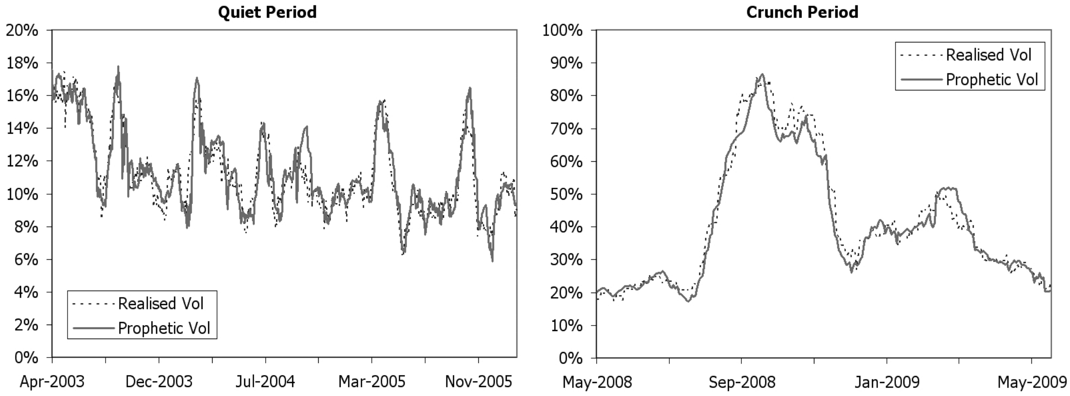

Figure 4: Prophetic vs. realised and prophetic vs. implied volatility during different market conditions

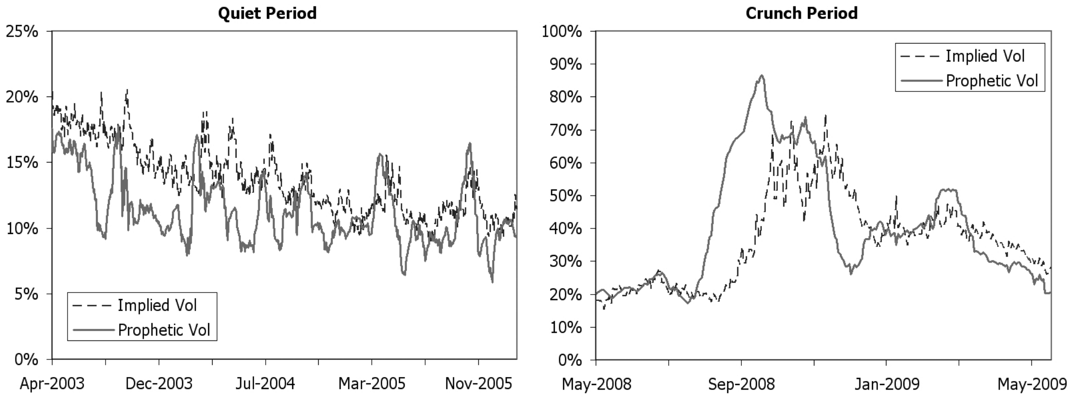

Here we can see a clear tendency for implied volatility to lag: when the crisis hits, realised (and prophetic) vols spike rather quickly, while implied vol is slower to react. This, combined with the slightly higher average implied vol, suggests that while prophetic and realised vol might well come from the same distribution, implied volatility doesn't. A Mann Whitney test bears this out: we cannot reject the hypothesis that prophetic and realised are the same; but we can have some confidence that prophetic differs from implied. The table in figure 5 presents the results for the whole period and the two sub-periods.

	Whole Period 1999-2009		Quiet Period 2003-2005		Crunch Period 2008-2009	
Confidence	99%	95%	99%	95%	99%	95%
Reject hypothesis $\sigma_P = \sigma_R$?	No	No	No	No	No	No
Reject hypothesis $\sigma_P = \sigma_I$?	Yes	Yes	Yes	Yes	No	Yes

Figure 5: Summary of the differences between the three volatilities

3.2 How far has prophetic volatility been from realised?

On average realised volatility – if we knew what it was going to be – would be a good vol to hedge at. But have there been particular paths where this is not true? Unsurprisingly, there have been. For example, a bad day to sell a one month call option and hedge at realised during the data set was 27th March 2000. For the month going forward from this date, realised volatility was 32.3% while prophetic was less than half that, 14.8%. It is worth looking at this path in particular to see how hedging at realised can fail to minimise P/L variance: figure 6 illustrates the level of the S&P 500 and its one week realised volatility during the period.

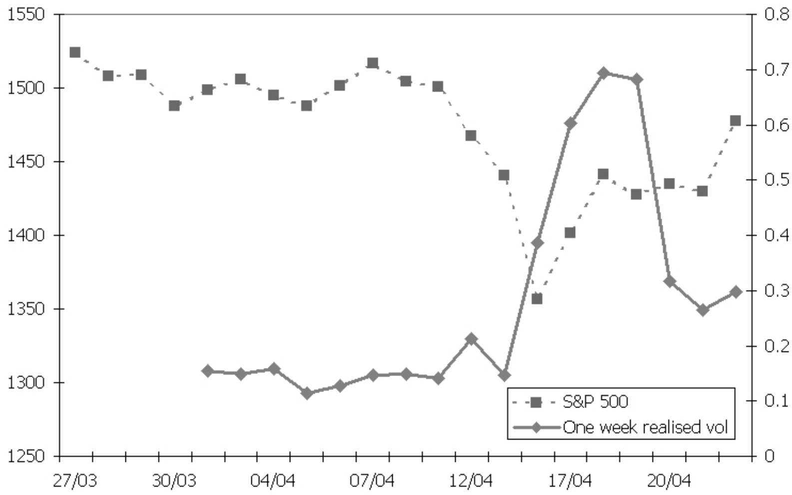

Figure 6: S&P 500 level and one week volatility for the month starting 27th March 2000

Clearly here the delivered volatility is relatively low for the first two weeks of the option's life, and the option has a reasonable amount of gamma as the S&P 500 is not trending away from the strike. The P/L for Black Scholes hedging a diffusion with true volatility σ_R over a short period dt at a vol σ_A is $\frac{1}{2}(\sigma_A^2 - \sigma_R)\gamma_R dt$. For the first half of the period, then, hedging at a vol of around 15% – close to the realised for these two weeks – will be better than using a hedge volatility of 32%.

In the final two weeks the realised volatility goes up, and the index falls, then recovers. However the option has relatively little gamma in this period, so the gamma-weighted delivered volatility is not large and the P/L is relatively insensitive to the hedge volatility. Thus while the high realised volatility in the second half of the option's life increases the average realised volatility significantly, it does not increase the prophetic volatility much, as it occurs when we do not have much gamma to collect this increased volatility[5]. Figure 7 shows the hedge performance at both realised and prophetic volatility.

The only time, then, when hedging at realised volatility does not produce close to optimal hedge ratios is when realised volatility changes, and the option's gamma in one volatility regime is significantly different from that in the other.

[5]The term 'volatility collector' for gamma, due to Dupire, is particularly apposite here.

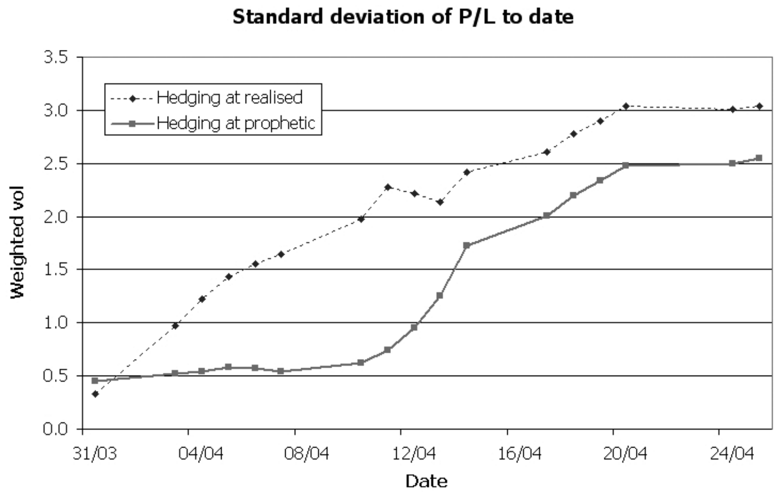

Figure 7: Hedge performance for the month starting 27th March 2000

3.3 How much do you need to charge to make a profit?

We have seen that hedging a prophetic volatility produces a stable daily P/L. However, a constant loss is stable. Therefore it is worth investigating the profitability of using prophetic volatility. In particular suppose that we not only hedge at prophetic – we price our option at this level too.

This would be foolhardy for several reasons: firstly, it ignores the market; and secondly it provides no compensation for other costs, such as bid/offer spreads, excluded from this analysis. Nevertheless, examining this profitability will give some useful insight into the profitability of strategies that use implied volatility alone. The table in figure 8 summarises the results of this analysis.

	Whole Period 1999-2009		Quiet Period mid 2003-2005		Crunch Period 2008-2009	
	Average	SD	Average	SD	Average	SD
Realised P/L at σ_P (%)	0.013%	7.0%	0.2%	6.2%	-0.42%	6.0%
Realised P/L at σ_I (%)	7.3%	26%	16.5%	16.3%	-10.8%	43.3%

Figure 8: Realised P/L from delta hedging a short call position at various volatilities as a percentage of the option price

Option pricing using prophetic volatility produces a fairly flat P/L with low variance. On average this has been centered on fair value (ignoring bid/offer spreads). Implied volatility, in contrast, typically produces a profit, but at the cost of significant earnings variance. Moreover, using implied during the Crunch period was loss-making.

A dealer might well think of the option pricing problem as determining how many vol points over the mid-market implied volatility to charge for selling an option. The bigger a vol spread we charge, the more likely we are to make a profit, but the fewer deals we win. Figure 9 illustrates this for pricing at spreads to both prophetic and realised volatility.

These results suggest that selling an option at a small spread to mid-market implied volatility and then hedging at implied is far from certain to make money: wide spreads are needed to give a high likelihood of making money. For instance, in order to be 90% certain of breaking even from selling a one month S&P 500 index option and delta hedging at implied, we would need to charge a 4.5 vol point spread over mid market implied volatility. That is a big spread for an index option, one that the market will not often bear. In contrast, if we knew what prophetic

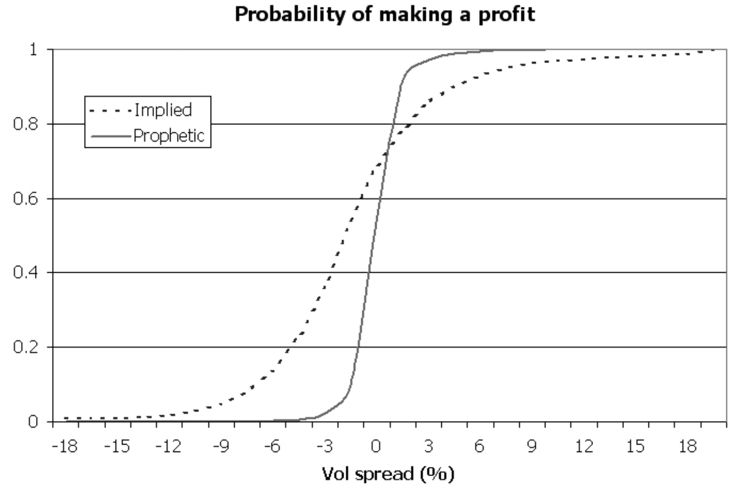

Figure 9: The cumulative probability of no loss from selling and delta hedging an option as a function of the spread it is sold at

vol was, a mere 1.5 vol point spread over it would give the same confidence that money will not be lost.

Evidently, the ability to prophesy future volatility would be a useful skill for an options trader to have (and being able to guess what gamma-weighted volatility is going to be would be even more useful).

4 Conclusions and Further Work

It would obviously be dangerous to generalise from one underlying, one strike and one maturity to hedging in general. Nevertheless it is suggestive that Black Scholes hedging at realised volatility works as well as it does despite the fat tails and autocorrelation of the S&P 500. Moreover, we would expect hedging longer term options to work better, at least if vega P/L on volatility remarks is ignored or minimised by vega hedging.

An extension of this work to different underlyings is difficult due to the paucity of accurate public domain implied volatility information of a constant moneyness. Therefore, while it would be interesting to see if the robustness of hedging at realised volatility on the S&P 500 applies more generally, this would require the use of private data sets. This would be worthwhile, though, as it may be that implied volatility on equity indices is (or at least in the past has been) biased by the demand for long-dated calls embedded in retail products: the spread of implieds over realised may in part be due to the demand for vega hedges to dealers' short positions in OTC calls.

The results reported here indicate that the ability to calculate hedge ratios at volatilities other than implied can be useful. If the trader's rôle is to hedge the book accurately, then implicit in that job description is the requirement that they predict gamma weighted realised volatility, and hedge accordingly. Since implied volatility is not the optimal volatility for generating hedge ratios, the feature of allowing both 'mark' and 'hedge' vols may be considered more necessary than it currently is in options trading systems.

References

[1] C. Alexander, L. Nogueira, *Hedging with Stochastic and Local Volatility* ISMA Centre Discussion Papers in Finance Number 10, 2004.

[2] L. Canina, S. Figlewski *The Informational Content of Implied Volatility*, The Review of Financial Studies, Volume 6, Number 3, pages 659-681, 1993.

[3] P. Carr *A survey of preference free option valuation with stochastic volatility*, Risk 5th annual European derivatives and risk management congress, Paris, 2000.

[4] T. Coleman, Y. Kim, Y. Li, A. Verma *Dynamic Hedging With a Deterministic Local Volatility Function Model*, Cornell University Technical Report 2003, available at `techreports.library.cornell.edu:8081/Dienst/UI/1.0/Display/cul.tc/2003-285`

[5] E. Derman *Fear and Greed in Volatility Markets*, Presentation, Goldman Sachs Quantitative Strategies Group, 1999.

[6] E. Derman, I. Kani, J. Zou *The Local Volatility Surface*, Goldman Sachs Quantitative Strategies Research Notes, 1995.

[7] B. Dumas, F. Fleming, R. Whaley *Implied Volatility Function: Empirical Tests*, Journal of Finance, Volume 53, Number 6, pages 2059-2106, 1998.

[8] Duan Jin-Chuan *The GARCH option pricing model*, Mathematical Finance, Volume 5, pages 13-32, 1995.

[9] B. Dupire *Pricing with a smile*, Risk, Volume 7, pages1820, 1994.

[10] M. Forde *The real P&L in Black-Scholes and Dupire Delta hedging*, University of Bristol Working paper, 2003, available at `www.maths.bris.ac.uk/research/stats/pub/ResRept/2003/12.pdf`.

[11] J. Gatheral *The Volatility Surface: A Practitioner's Guide*, Wiley Finance, 2006

[12] S. Heston *A closed-form solution for options with stochastic volatility, with application to bond and currency options*, Review of Financial Studies, Volume 6, pages 327343, 1993.

[13] R. Jarrow (Ed.) *Volatility - New Estimation Techniques for Pricing Derivatives*, Risk Books, 1998.

[14] International Accounting Standards Board *Financial Instruments: Classification and Measurement*, IFRS 9 Exposure Draft, 2009.

[15] M. McIntyre *Performance of Dupire's Implied Diffusion Approach Under Sparse and Incomplete Data*, Journal of Computational Finance, Volume 4, Number 4, pages 33-84, 2001.

[16] A. McNeil, R. Frey *Estimation of Tail-Related Risk Measures for Hetroscedastic Financial Time Series: An Extreme Value Approach*, Journal of Empirical Finance, Volume 7, pages 271 300, 2000.

[17] D. Murphy *The Empirically Optimal Volatility for Black/Scholes Hedging in Various Market Conditions*, preprint, 2010.

[18] P. Ritchken, R. Trevor *Pricing Options under Generalized GARCH and Stochastic Volatility Processes*, Journal of Finance, Volume 54, Number 1, pages 377-402, 1999.