Buffett and Black-Scholes: What Does Volatility Mean?[1]

Pablo Triana[2]

The Berkshire Hathaway derivatives selling case has taught us many valuable things. From the synthetic funding benefits that the strategy can generate, to the valuation and risk factors of options, the characteristics of more exotic products, as well as the impact of credit considerations in the presence of lax collateral agreements. I have explored all those in some detail in previous pieces. Here I attempt to tackle what is possibly the last main theme that had yet to be uncovered, namely the volatility parameters used by the firm when calculating its put option liabilities and what such numbers tell us both about Berkshire´s approach to valuation as well as about the valuation tool itself.

As is widely known, Berkshire Hathaway uses the famed Black-Scholes option pricing model to calculate its liabilities on the massive long-term equity index puts it sold between 2004 and 2008. The fair value of the options, as churned out by the model, equals Berkshire´s discounted theoretical expected cash obligations on the trade or, in other words, the liquidation cost of the portfolio (should Berkshire be able to find someone willing to take the risk off its hands for precisely that amount). Berkshire´s boss Warren Buffett has long been critical of Black-Scholes but nevertheless chose to employ it for the accounting representation of this particular exposure.

We, in fact, should be glad that Berkshire is using Black-Scholes, as its use of the model can teach us lots about the nature of the model. In particular, about the meaning of volatility in the Black-Scholes context. We already analyzed in

[1] Forthcoming, *Corporate Finance Review*

[2] ESADE Business School

previous work the impact of adding a credit risk premium to the model´s typical assumption of a risk-free rate of interest. We also wondered about liability numbers that intuitively looked smaller than may have been expected. Now we deal with what has perhaps traditionally been the most discussed and controversial aspect of the model: what number should we insert under the formula´s volatility parameter and what should that number truly stand for. The answers to those questions can lead to vastly different option prices and to vastly different interpretations of those prices. The implications in terms of fair value accounting (and thus in terms of net income) can be very big: a small liability instead of a big one, a large profit instead of a loss.

Berkshire has opted for a peculiar way of dealing with volatility. The firm is using the Black-Scholes volatility parameter as a static forecast and has stoically stuck by such prediction even in the face of some of the most wildly swinging markets ever contemplated. In this piece, I argue that there may be a more efficient, let alone realistic, way of using that parameter. Not as a predictor but more humbly as a price adjustment mechanism, reflecting not some precise futuristic view but rather more modestly trying to dynamically adapt option prices to unavoidable worldly developments. In that sense, Berkshire´s numbers would be a less than perfect representation of actual liquidation costs, its mark-to-model numbers perhaps far off true mark-to-market figures, on occasion at least. By opting for a fixed prediction, Berkshire may have obtained liability metrics that would have been too low some times and too high other times.

Volatility is an important part of the Berkshire option selling story for several reasons. For one, very long term contracts (the puts expire in 2019-2028) can be particularly exposed to the volatility number. Also, these options have danced furiously throughout their lives, moving from at-the-money to deep in-the-money to deep out-of-the-money. A key characteristic of modern option markets is that traders would use a different volatility number depending on the moneyness of the contract at any given time, in contrast to pure Black-Scholes that assumes constant volatility independent of moneyness intensity. By sticking with a constant figure throughout, Berkshire appears to be ignoring or neglecting this reality, the same option portfolio being endowed with essentially an identical volatility input whether at-the-money or deep in-the-money or deep out-of-the-money.

What is volatility? Or, better yet, what should volatility be? The Berkshire episode, already rich in many other lessons, can help illustrate this critical conundrum.

Price, don´t predict

Why should an underlying asset´s volatility matter when pricing and valuing an option? Unless the underlying is itself directly referenced to volatility, the final payout will not be directly determined by volatility. So shouldn´t we only care about the actual spot and forward prices of the asset and not about how much those happen to move around? Because options provide an upside (a potentially very large one) while limiting the downside (to a possibly very small upfront premium sum), option buyers would enjoy the sight of the underlying asset dancing vertiginously, as movement can only deliver benefits on a net basis: the more movement, the higher the potential for a large gain all the while keeping the loss perennially constant. Even deep in-the-money options can profit from extra volatility, even though it would seem that they don´t need any extra “help” in terms of additional dancing. In fact, some of the biggest mark-to-market profits that can be obtained from a long option position (and thus biggest losses for the shorts) derive from changes in volatility.

So volatility is important because it tells us at any point how swingy the underlying is, and thus whether we should gauge the future payout potential as modest or as mouthwateringly sizable. Volatility aids us in option pricing by incorporating a reality-informed view as to the underlying asset´s possibilities, beyond the irrepressibly isolated picture provided by the asset´s spot price at any particular point. It is good that we can incorporate something called "volatility" into the pricing equation (whether mathematical or mental) because we need to make presence for a variability-representing parameter when trying to ascertain the proper value of a variability-enjoying instrument. The asymmetry of options payouts determines that an asset with the capacity to swing is a more attractive underlying than one without such capacity.

One of the beauties of the Black-Scholes formula is that it contains a place for the volatility parameter. It allows you to put a number for volatility. The key question is, how should we take advantage of that? We could try to predict turbulence from here till expiration date. But that´s going to be hard, maybe well nigh implausible. And we may all have different predictions, making option pricing quite subjective (here we are talking mostly about less liquid longer term options rather than the more liquid short-term contracts typically listed on official exchanges with prices coming from sizable supply and demand streams). By trying to forecast and nothing else, we may be entirely wasting the benefit of being able to put a number on volatility. What if we gave up on the prediction stuff and instead used the volatility parameter to gauge the capacity of and potential for the underlying asset to fluctuate? Volatility would now serve the purpose of increasing and decreasing option prices as underlying markets show more or less fluctuation.

The benefit of the volatility parameter would thus be directional, rather than precisely numerical: revise it up or down based on recent market events, but don´t presume to get the right future figure at three decimal points. Forecasting precisely is hard and maybe naïve, but adjusting in the right direction should be easier and more grounded. That would be the real value of volatility, as prompter of option value correction, not so much as alibi to make turbulence predictions.

In this light, the volatility parameter in the formula should not be seen as platform for end-users to express their forecasts, but as a way to add (or take away) premium to the option´s value as the underlying asset proves its potential (or lack of) to swing wildly. Can´t predict, but can tweak a parameter to incorporate recent market developments and what they say about the potential and capacity for the asset to be swingy and thus worth more.

An option on an underlying that can (demonstrably) move around a lot should be worth more, given its asymmetric payoff and its convexity. The way to make that upward adjustment is through the formula´s volatility parameter. Starting from some reasonable benchmark (perhaps the asset´s average historical volatility, whichever way you want to measure that), the parameter should be increased or decreased following obviously significant market behavior. The option´s value should reflect its immediate liquidation cost, thus making it unavoidable to present a realistic assessment of the underlying´s current volatility. Otherwise, the price may reflect neither true market value (what people would pay for it today) nor fundamental value (the, updated, nature of the underlying asset).

Many times we are instructed to forecast future volatility based on what volatility did in the past. But you can´t rely solely on History when the market is making History. If the vol number does not reflect the latest developments, you would be effectively treating the option as if written on a different asset (yesterday´s asset, not updated).

There are high-profile cases of firms that got in trouble by assuming that volatility would abide by historical tenets and by disregarding the market-driven liquidation costs. Sticking by a forecast did not work well here, as others thought it more prudent to incorporate real-life events into the prices of even long-dated options.

Take the notorious case of UBS´ Ramy Goldstein, who had built an apparently successful business selling long-dated (five years) equity index volatility via structured products back in the early and mid 1990s. When the Asian crisis erupted in 1997, short-term implied equity index vol shot up, driving five-year implied vol up. Goldstein´s desk experienced a mark-to-market blow up, involving huge liquidation costs materialized when they were ordered to close down the positions by buying them back from other institutions. It was considered at the time quite reckless to be selling such long dated volatility (they apparently did not hedge their vega exposure). The market did build the short-term tremors into long-dated volatility, with the five-year tenor reportedly jumping from 17% to 25%. For UBS the mark-to-market realities obviously mattered: the bank was forced into a shotgun wedding with Swiss rival SBC.

Or take the related case of long gone hedge fund LTCM. With Goldstein out of the game, LTCM took on the role of Central Bank of Volatility, or unique seller of long-dated index call options to banks that had sold long-dated structured products. LTCM firmly believed that the options were mispriced ("free money", implied vol being sold at 23% with historical vol at 15%; the fund believed in convergence to the "normal"). But when the Russian crisis exploded in summer of 1998, LTCM´s counterparts too reacted to those short-termish yet impossible to ignore developments and marked long dated vol at 30-40%. LTCM, bound by collateral agreements, faced gargantuan mark-to-market-driven margin calls and had to be rescued and then orderly liquidated.

These classic episodes show that long term volatility can be difficult to buy back. It can be a highly illiquid, unique asset that enjoys no fresh supply ("like a Picasso", in the words of a market participant). In other words, a very expensive luxury. Liquidation costs do matter, as you may be forced (by yourself or your counterparts) to get rid of the exposure when things are tough. Even though Berkshire got good collateral treatment that doesn´t mean that it would be absolved from paying a market price if it had to liquidate. Thus, its liabilities, many may conclude, should properly reflect such cost.

For Berkshire to not modify its volatility input in the face of such changing equity markets as the put option selling strategy witnessed (incredible chaos, supreme calm), with the excuse that the firm is interested only in at-expiration long-term volatility, would be akin to saying that what happens in an option´s underlying asset´s market throughout the option´s life should not matter too much if at all, and that an option´s replacement cost is not an important variable. That you shouldn´t judge a seller´s accounting (and, if it comes to that, economic) exposure according to the price one would have to pay to cancel the exposure. By religiously abiding by a forecast, Berkshire would have both deflated (during chaotic times) and exaggerated (during placid days) its liability.

Does it really make sense to build the same turbulence premium into a put option as Bear Stearns or Lehman Brothers are sinking and (option cost-driven) volatility indices are hitting all-time highs, as when no major crisis is taking place and those same indices are close to all-time lows?

It is true that Berkshire´s settlement exposure was and is to very long dated maturities, while the volatility indices that reached historical highs in 2008-09 and then historical lows in 2012-13 stand for short-term volatility. But that should be no reason to disregard market dynamics, if only for the purposes of more representative accounting and mark-to-market metrics. Also, traders can hedge long dated options with short term contracts, rolling them over as they expire, so if prices of short term options go up (as indicated for instance by a rising VIX index) that would drive up the hedging costs and thus prices of long dated options.

Too little vol for so much vol?

The tables and graphs below show the volatility numbers that Berkshire has been inserting into its Black-Scholes model, and the levels of the (shorter-term) implied volatility indices for the equity markets that concern us here (S&P, FTSE, Nikkei, Eurostoxx). While those indices are often described as representing traders´ volatility expectations, they may best be seen as indicators of option costliness, drawing on a large range of strike prices.

We can see that while Berskhire stuck with volatility estimates of 20-22%, equity markets witnessed quite the chute-the-chute during those periods. As vol indices reached extraordinary peaks in late 2008, in particular, Berkshire seems to have made no correction for such fact, its Q4 2008 number identical for instance to its input for Q4 2009 (much calmer markets). Conversely, the much deflated short-term fluctuations experienced in 2013 seem to have made little dent in Berkshire´s vol forecasts, with just seemingly modest downward revisions to the inputs. It is obvious that Berskhire has seen no reason to subject its volatility numbers to an accounting rollercoaster, notwithstanding the very real rollercoaster playing out in the markets.

	Berkshire´s Vol Inputs
q2 2008	23,0%
q3 2008	22,0%
q4 2008	22,0%
q4 2009	22,0%
q4 2010	21,5%
q4 2011	21,4%
q1 2012	22,0%
q2 2012	22,0%
q3 2012	22,0%
q4 2012	20,9%
q1 2013	20,0%
q2 2013	21,0%
q3 2013	21,0%
q4 2013	20,7%

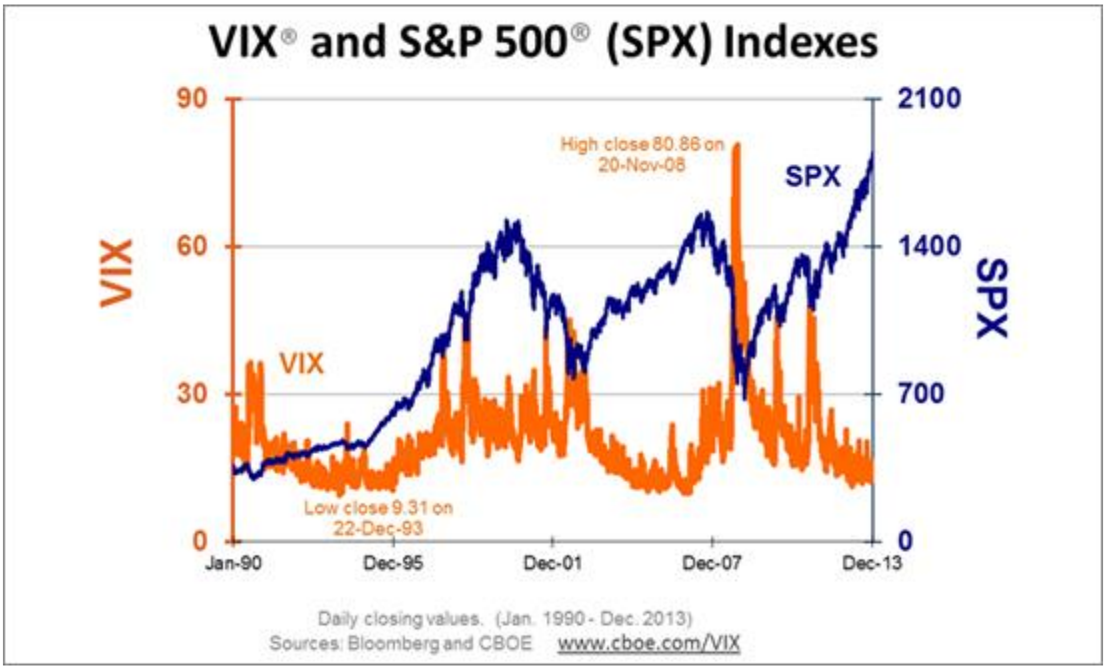

	30 Day Volatility			60 Day Volatility			90 Day Volatility			180 Day Volatility			360 Day Volatility		
Year	avg	max	min	avg	max	min	avg	max	min	avg	max	min	avg	max	min
YTD	17.0	27.4	11.2	17.9	27.7	12.0	18.6	28.7	12.1	21.2	30.3	14.3	22.9	29.5	16.7
2012	17.9	27.4	12.2	18.8	27.7	12.8	19.5	28.7	14.3	22.1	30.3	16.5	23.8	29.5	18.6
2011	23.6	45.5	12.4	23.6	41.4	16.1	23.6	38.3	15.1	25.4	38.8	19.1	25.6	34.3	21.0
2010	21.8	41.1	13.8	22.6	34.2	16.8	23.0	32.5	15.9	26.1	36.2	18.0	27.2	37.2	16.1
2009	29.3	49.5	19.0	29.2	49.9	20.5	28.4	40.0	20.6	30.8	43.5	24.6	31.0	41.2	25.8
Since 2000	21.5	84.8	9.0	21.6	64.2	9.1	21.5	55.1	9.1	22.6	49.6	10.2	22.6	43.8	11.2

FTSE 100 IVI – Implied Volatility

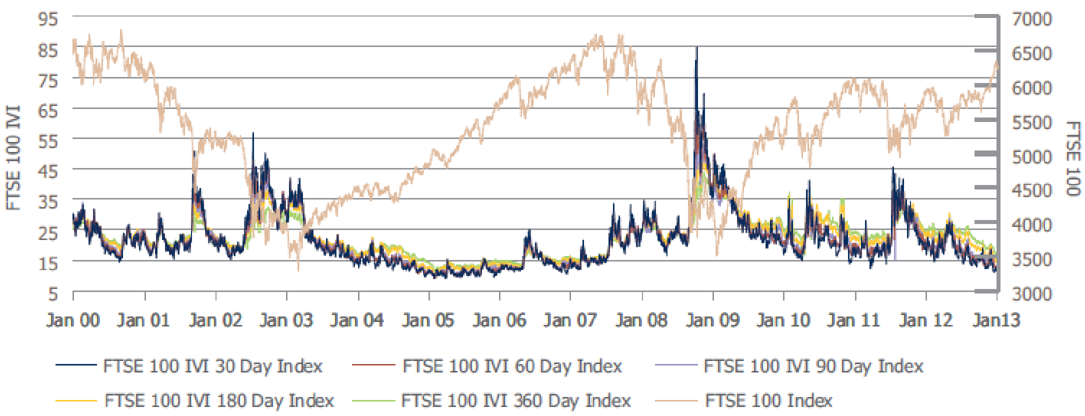

■Graph (Dec/30/2013)

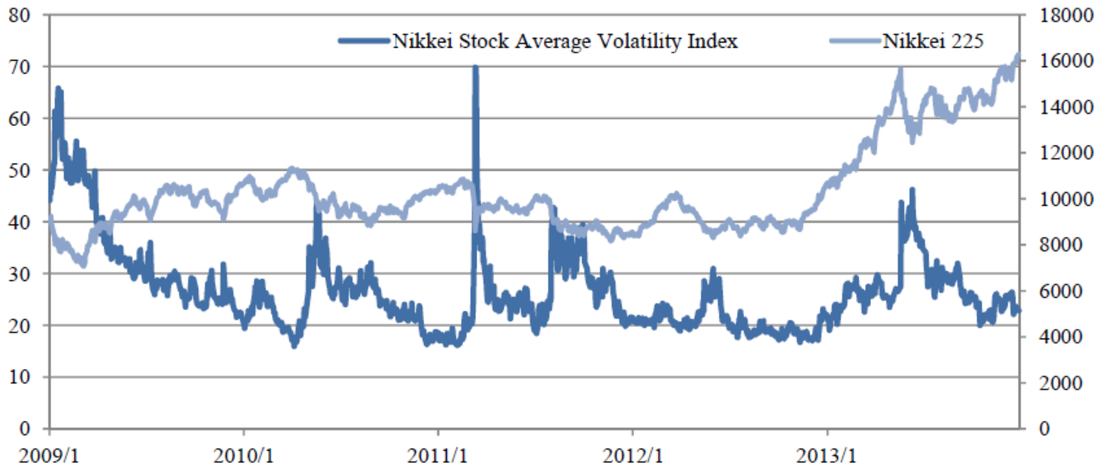

High, Low & Mean

Index	YTD			2011			2010			2009			Since Incept.*		
VSTOXX	High	Low	Mean	High	Low	Mean	High	Low	Mean	High	Low	Mean	High	Low	Mean
30 days	25.90	14.86	18.40	53.55	18.48	30.06	49.87	18.36	26.45	56.33	22.85	33.59	87.51	11.60	26.12
60 days	23.85	15.78	18.94	50.48	19.86	30.14	49.96	21.12	27.31	54.41	24.34	33.95	74.02	11.66	25.87
90 days	23.21	15.08	19.55	49.08	20.90	30.45	47.65	21.76	28.01	53.39	24.34	34.43	65.30	12.00	26.16
120 days	23.64	17.26	20.20	47.49	21.79	30.74	46.46	23.16	28.85	52.83	26.28	35.02	64.32	12.58	26.58
150 days	23.86	17.81	20.60	46.84	22.45	30.92	45.18	23.61	29.32	53.28	27.18	35.25	60.93	12.88	26.64
180 days	24.16	18.51	20.96	46.25	23.10	31.07	44.57	24.02	29.67	52.40	27.60	35.32	58.56	13.08	26.63
210 days	24.37	18.99	21.27	45.73	23.56	31.13	44.14	24.43	29.97	51.77	27.07	35.34	56.81	13.22	26.60
240 days	24.52	19.39	21.56	45.33	23.84	31.13	43.65	24.73	30.20	51.29	25.62	35.32	54.00	13.50	26.55
270 days	24.67	19.83	21.83	44.77	24.02	31.10	43.26	24.91	30.36	50.52	24.43	35.15	51.70	13.57	26.45
300 days	24.97	20.18	21.88	44.09	24.15	31.00	42.95	23.25	30.42	49.90	23.43	34.88	49.90	13.33	26.31
330 days	25.21	19.21	21.94	43.53	22.14	30.92	42.44	21.19	30.47	49.39	22.58	34.64	50.95	13.62	26.25
360 days	25.42	17.67	22.03	43.01	20.17	30.86	42.00	19.32	30.55	48.68	21.88	34.54	51.99	13.87	26.29

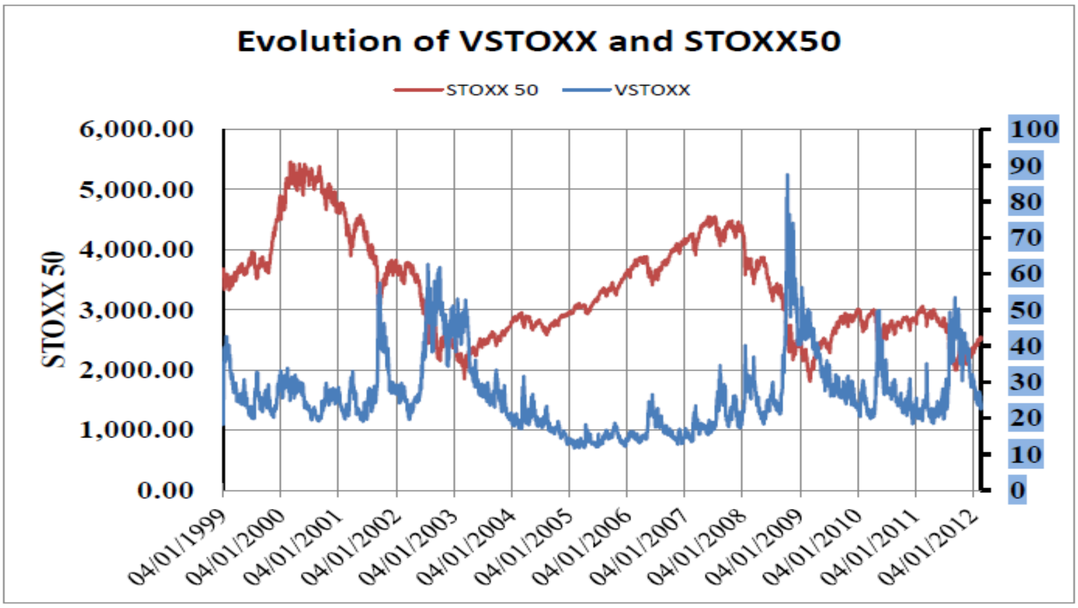

Where do Berkshire´s volatility numbers come from? Luckily the firm tells us, for instance as of year-end 2009:

“The Black-Scholes model incorporates volatility estimates that measure potential price changes over time. The weighted average volatility used as of December 31, 2009 was approximately 22%, which was relatively unchanged from year end 2008. The weighted average volatilities are based on the volatility input for each equity index put option contract weighted by the notional value of each equity index put option contract as compared to the aggregate notional value of all equity index put option contracts. The volatility input for each equity index put option contract is based upon the implied volatility at the inception of each equity index put option contract”

That is, the numbers correspond to the implied volatility of each contract at inception, thus at-the-money volatility as all the options were struck at the same level as the index spot level at the time. The different volatility numbers are weighted by the notional size of each contract so that larger contracts matter more and also presumably there would be a foreign exchange adjustment here to reflect the dollar-denominated size of the contracts linked to international non-US indices. In later regulatory filings, Berkshire clarified that those volatility numbers reflected the firm´s long term volatility expectations.

We should note that Berkshire´s volatility marking tactic may fly in the face not only of inescapably noteworthy real-life developments, but also of some key hard core features of modern option markets. For almost thirty years, options have traded under the so-called “volatility smile”, whereby options on the same underlying asset and with the same maturity but with different strike prices register implied volatility numbers that are different based on each strike price. The smile (or the less symmetrical skew, more typical in the case of equities) reflects traders’ desires to bump up the values of options that are further from at-the-money strikes, in particular those with deep out-of-the-money or in-the-money strikes. Given the real-world chance that those extremes will be reached and given Black-Scholes' underpricing of such contingencies given its assumption of a normal probability distribution, traders feel that such options need to be worth an extra more if they are to be sold at all. Thus, when working the model backwards from actual market price to the implied volatility parameter that would correspond to that price we obtain numbers that are higher for extreme strikes (the ones that need the additional bump; as volatility is the only number that can be freely manipulated in the formula, this is where the adjustment is made). In the case of equities, deep out-of-the-money puts have the highest implied volatility number, a reflection of traders' crash-o-phobia as well as the potential illiquidity of those contracts.

Berkshire, we are told, is marking its put options at the at-the-money implied volatility at which the contracts were originally sold. The firm seems to believe those numbers to be accurate predictions of the relevant long term volatility. But in marking in such fashion, it is neglecting the smile. Many of those options have swung violently from at-the-money into in-the-money, deep in-the-money, out-of-the-money, and deep out-of-the-money. And once the underlying asset moves so far below or so far above the strike, the implied volatility figure should change significantly so as to adjust for the desired proper price. Keeping your volatility input constant under those circumstances would seem rather odd, a subliminal disregard of one of the most basic and elemental tenets of option markets. Even if you believed that your original at-the-money volatility forecast was solidly reliable, the fact that those options are no longer at-the-money would, as per the smile, yield different volatility numbers for those same contracts. Berkshire did not just stick by a forecast in the face of untold turbulence, it is sticking to the number even as the options are far away from the original moneyness. Some may conclude that Berkshire is disrespecting the smile.

Finally, not only is Berskshire apparently disregarding short-term market volatilities, but possibly also changes in long term vol. Given the over-the-counter and illiquid nature of these options it may be not be easy to come by hard data, but some evidence points towards long-dated implied volatility figures that spiked up during the 2008-09 equity markets dark days. For instance, some data indicates that S&P at-the-money ten-year implied volatility jumped from around 25% in May 2008 to around 35% during December 2008-April 2009. Since implied volatility is a measure of option costliness, the bottom line is, not surprisingly, that long term equity protection became quite more expensive in the face of tanking equity markets and explosive equity turbulence. Another piece of data suggests that five-year S&P options saw their market prices increase three-fold between the start of 2007 (9% of notional) and the end of 2008 (27% of notional). It seems only logical that purchasing optionality, even if way into the future, would cost a tad more than prior to an epic market meltdown.

Nebraskan Vega

What if Berkshire had marked its volatility in a dynamic, market-revering manner? Would it have made much of a difference with regards to the constant-volatility path eventually taken? Would Berkshire´s put options liabilities have changed much? In other words, would giving another meaning to volatility have been relevant in this case or not? If the answer is a resounding “no” then perhaps we shouldn´t make much of a fuss about Berkshire´s attitude towards volatility, even if we may find it less than perfectly appropriate. Who cares if Berkshire chose 20-22% vol if nothing would have changed had the number been, say, 30-32%?

In order to conduct this analysis we need a critical piece of information: Berkshire´s vega. By how much would the mark-to-model value of the puts portfolio vary as a result of inserting a higher volatility number into the model? Thankfully, Berkshire is amply generous here, having disclosed the numbers for at least some quarters.

The table below displays the puts’ fair value (in millions of $) as a result of increasing implied volatility by two percentage points and by four percentage points, with everything else constant. We measure the increase in fair value in percentual terms, so as to gauge the relative importance of volatility at any given moment.

	puts value	vol up 2%	change	vol up 4%	change
q2 2008	$5.845	$6.408	9,63%	$6.969	19,23%
q3 2008	$6.725	$7.231	7,52%	$7.733	14,99%
q4 2008	$10.000	$10.451	4,51%	$10.882	8,82%
q4 2009	$7.300	$7.885	8,01%	$8.459	15,88%
q4 2010	$6.700	$7.221	7,78%	$7.732	15,40%
q4 2011	$8.500	$8.950	5,29%	$9.407	10,67%
q4 2012	$7.500	$7.955	6,07%	$8.414	12,19%
q4 2013	$4.700	$5.067	7,81%	$5.479	16,57%

We can see that had Berkshire in the critical Q4 2008 chosen 24% (26%) instead of 22% for the volatility parameter, its liabilities would have been $10.45 billion ($10.88 billion) rather than the reported $10 billion, or a 4.51% (8.82%) increase. Falling into the temptation of interpolation, one may venture to conclude that marking a 32% (42%) volatility would have led to around $12 bn ($14 bn) in liabilities, a decidedly larger accounting setback for the firm. Big deal? Well, for starters Q4 2008 would have registered a large quarterly loss rather than the small gain that was reported (at 22% vol), and annual profits could have decreased by 50-25%.

Other periods saw a much more significant vega, in particular Q2-Q3 2008, and Q4 2009-2010-2013, with sensitivities to 4% volatility increases in excess of 15% of option value. In those periods, not having selected a higher volatility figure would have paid especially handsomely for Berkshire, preventing a sizable relative rise in liabilities. In many of those dates short-term market implied volatility (and thus option prices) was much higher than Berkshire´s volatility estimates.

It makes sense that exposure to vol would be lowest in Q4 2008 as all the puts were very deep in-the-money and vega is highest at-the-money and then shifts downwards for strikes further away, in a bell shape fashion. As equity markets recovered (the S&P and the FTSE much more than the Nikkei and the Eurostoxx) and returned to levels closer to at-the-money, sensitivity to volatility became greater, only to diminish again as the S&P and the FTSE kept rising upwards and those options now became deep out-of-the-money and thus endowed with much lower vegas. Following a very strong 2013 for the Nikkei, which saw those options get away from deep in-the-moneyness and towards at-the-moneyness, the overall portfolio vega went up again. The table below shows equity index levels as well as simulated moneyness levels, assuming hypothetical average strike prices for each of the underlyings (of 1300 for S&P puts, 6000 for FTSE puts, 16000 for Nikkei puts, and 3800 for Eurostoxx puts; those levels attempt to represent average index levels for the 2004-2008 period when the options were being sold).

	S&P	Moneyness	FTSE	Moneyness	Nikkei	Moneyness	Eurostoxx	Moneyness
q2 2008	1262	2,92%	5410	9,83%	13370	16,44%	3360	11,58%
q3 2008	1100	15,38%	4370	27,17%	8570	46,44%	2590	31,84%
q4 2008	932	28,31%	4140	31,00%	7990	50,06%	2230	41,32%
q4 2009	1120	13,85%	5180	13,67%	10190	36,31%	2770	27,11%
q4 2010	1260	3,08%	5860	2,33%	10230	36,06%	2950	22,37%
q4 2011	1260	3,08%	5680	5,33%	8800	45,00%	2410	36,58%
q4 2012	1425	-9,62%	6270	-4,50%	11130	30,44%	2700	28,95%
q4 2013	1848	-42,15%	6750	-12,50%	16290	-1,81%	3110	18,16%

Vega is also influenced by time to maturity and by volatility itself. Longer-dated options would have a higher vega, while more volatile times would typically yield a lower vega. Thus, the further away from at-the-money and the higher the vol parameter and the less time to maturity the lower the exposure to vol; conversely, the closer to at-the-money and the lower the vol and the more time to maturity the higher the exposure to vol. Because Berkshire´s options were all long-dated, this guaranteed a decent exposure to volatility, that would be tamed if spot traded away from the strikes and if volatility increased. This all goes to indicate the timely relevance of fudging with the model´s vol parameter, when it would matter the most to not insert a higher number, when it would have the biggest effect per one percentage point in extra vol. Given the very long lives of these options, and thus the higher vegas, the benefits in terms of lower mark-to-model liabilities from choosing relatively modest implied volatility figures were significant, especially when at least some of the options were trading close to their strikes. Also, and given the inverse relation between vega and volatility with vega greatest for low vol, once you have selected a low volatility number (perhaps because it would lead to a lowish option fair value) it is particularly beneficial not to increase that number as the percentual increase in liabilities would be quite large.

The table below shows how a put´s vega varies with time to maturity, spot prices, and volatility. We see how vega has the greatest influence at-the-money, and then decreases as the option gets in-the-money. Why vega highest at-the-money? At those levels, a small change in vol can make a world of difference for the option: between having intrinsic value and not having intrinsic value. When the option is deep in-the-money extra vol doesn´t change things much, similarly deep out-of-the-money.

We also see that, everything else constant, shorter maturity means a less relevant vega. For very long maturities, even deep in-the-money puts can have high and relevant vega (not much lower than at-the-money vega); the option must be extremely deep in-the-money for vega to be very small. So for Berkshire in the critical late 2008-early 2009 days, vega would still have been highly significant even though the puts were getting quite below the strike. In other words, it paid off accounting-wise to keep the volatility number lowish. For much shorter maturities, in sharp contrast, even slightly in-the-money options can have much smaller and much less relevant vega than at-the-money.

Finally, we see how vega´s impact decreases as volatility increases. This reduction in influence is especially true for long term puts both at-the-money (especially) and (less so) in-the-money. Late 2008-early 2009 was a propitious time for Berkshire not to increase its volatility estimates, given the combination of very long maturities and a low base volatility number, though the deep in-the-moneyness at the time would have tempered the accounting damage from a higher vol number. The impact is smaller for short dated at-the-money, with a negligible impact for short dated in-the-money. Why vega lower for higher vol? When vol is very high, a small increase in vol changes little; when vol is low, a small change in vol counts for a lot.

strike	spot	vol	time	interest	put price	vega	vega/price
100	100	0,22	15	0,04	8,51	0,82	9,59%
100	80	0,22	15	0,04	11,67	0,85	7,26%
100	60	0,22	15	0,04	16,49	0,81	4,88%
100	100	0,22	10	0,04	9,69	0,82	8,51%
100	80	0,22	10	0,04	14,13	0,84	5,96%
100	60	0,22	10	0,04	21,00	0,74	3,54%
100	100	0,22	1	0,04	6,77	0,38	5,64%
100	80	0,22	1	0,04	18,26	0,25	1,35%
100	60	0,22	1	0,04	36,18	0,03	0,08%
100	100	0,25	15	0,04	11,00	0,84	7,64%
100	100	0,35	15	0,04	19,41	0,83	4,25%
100	100	0,45	15	0,04	27,26	0,74	2,71%
100	80	0,25	15	0,04	14,22	0,84	5,94%
100	80	0,35	15	0,04	22,40	0,78	3,50%
100	80	0,45	15	0,04	29,75	0,68	2,30%
100	60	0,25	15	0,04	18,87	0,79	4,16%
100	60	0,35	15	0,04	26,33	0,70	2,67%
100	60	0,45	15	0,04	32,89	0,61	1,84%
100	100	0,25	10	0,04	12,19	0,84	6,89%
100	100	0,35	10	0,04	20,60	0,83	4,03%
100	100	0,45	10	0,04	28,63	0,77	2,69%
100	80	0,25	10	0,04	16,64	0,83	5,00%
100	80	0,35	10	0,04	27,75	0,78	2,82%
100	80	0,45	10	0,04	32,23	0,71	2,21%
100	60	0,25	10	0,04	23,22	0,73	3,16%
100	60	0,35	10	0,04	30,31	0,68	2,25%
100	60	0,45	10	0,04	36,83	0,62	1,68%
100	100	0,25	1	0,04	7,92	0,38	4,84%
100	100	0,35	1	0,04	11,75	0,38	3,26%
100	100	0,45	1	0,04	15,56	0,38	2,44%
100	80	0,25	1	0,04	19,03	0,27	1,39%
100	80	0,35	1	0,04	21,89	0,30	1,37%
100	80	0,45	1	0,04	24,97	0,31	1,26%
100	60	0,25	1	0,04	36,30	0,05	0,14%
100	60	0,35	1	0,04	37,17	0,12	0,33%
100	60	0,45	1	0,04	38,64	0,17	0,44%

Conclusions

- Volatility matters for pricing options due to the asymmetry of the payout function; the more potential for swinging the underlying asset has, the more valuable the option should generally be

- The Black-Scholes volatility parameter may be best employed if used to dynamically gauge the potential of the underlying asset for moving, adapting the number as markets change; using it rather as a platform for volatility forecasting could be an ineffective waste

- Berkshire Hathaway´s view of the volatility parameter as a static volatility forecast stands in sharp contrast to the extremely bouncy path experienced by the equity indexes on which it sold puts; by not adapting (in fact, making essentially no modifications at all) its volatility input to real-world realities, the firm´s liabilities numbers may be categorized as too small at times and as too large at other times

- Drawing on Berkshire´s own estimates, mark-to-model losses on the puts could have increased by several billion dollars during several different quarters if the firm had employed a more market-sensitive volatility figure; the very long-term nature of the options made sure that sensitivity to the volatility number remained quite high, in spite of other factors often pulling in an opposite direction, and thus there were significant accounting benefits to be reaped from inserting a lowish figure in the model´s volatility parameter