Spot-vol Dynamics and Deltas for SPX Options[1]

OR: Does anyone know how to calculate the correct delta of a vanilla option?

The casual observer might think that hedging European-style vanilla options on SPX, some of the most liquid options on the planet, might be a well understood and nearly exact science. The truth is far from it, especially in this day and age.

Ignoring another large topic, that of optimal portfolio hedging, we will here concentrate on the question of optimal delta hedging with the underlier, which arises already for one option. Even this topic raises many questions, and has attracted a lot of study over the years.

The reason is that the greeks of vanillas, even delta, are model-dependent. Recall that the Black-Scholes formula is in practice mostly used as a dictionary, to translate back and forth between (tradable) prices and (more intuitive) implied volatilities. Hence the well-known deficiencies of the Black-Scholes model do not matter (much) for pricing. However, this is not the case for the greeks. One will, for example, get the wrong delta if one ignores the fact that implied volatilities systematically move, generically, when the underlier moves.

The correct instantaneous delta is given by

$$\text{delta} = \text{deltaBS} + \text{vega} \cdot \text{dVol/dF} \quad (1)$$

where deltaBS denotes the standard Black-Scholes delta, and vega is the usual Black-Scholes vega. For readability we suppress the dependence on contract details (time to expiry T, strike K, etc) and market data (rates, vols, etc). In the above we consider the delta with respect to the forward of the term in question (i.e. technically we're using the Black rather than the Black-Scholes model), which of course is easily adjusted, if needed.

The above is exact and well-defined for the case of European exercise-style options, i.e. covers the SPX case we mostly care about here, and has the usual (relatively minor) subtleties for the American vanilla case (i.e. most stock and ETF options).

The quantity dVol/dF, that we will abbreviate as **dVdF** from here on, is the derivative of the implied volatility of the option with respect to the underlier value (here the forward).

---

[1] This article originally appeared on LinkedIn:
https://www.linkedin.com/pulse/spot-vol-dynamics-deltas-spx-options-timothy-klassen

So, knowing the optimal deltas is equivalent to knowing dVdF for any option. If we know the "**spot-vol dynamics**" (aka **vol dynamics** for short) -- i.e. how the implied vol surface moves when the underlier moves -- then of course we also know dVdF for any option.

At least some of the sophisticated equity options market participants have known, at least in principle, how to properly think about vol dynamics for a while: If one factors the vol curve at a given term into an overall vol level (say the at-the-forward aka ATF vol), and a "shape" curve, suitably defined in dimensionless terms in e.g. normalized strike (NS) space, then vol dynamics is quite simple. In a good approximation the shape does not change at all, and the dynamics of the underlier can be expressed in terms of one dimensionless number, the **ratio** of slopes of the curve on which the ATF vol actually moves over the ATF slope of the vol curve of the term in question. (The former curve is a function of the underlier value, also known as the "vol path", the latter a function of strike, the "vol skew".)

This dimensionless number has many names: Bergomi calls it the "skew stickiness ratio" (SSR), some firms call it "super-skew"; we've also seen the term "skew multiplier" used. We will refer to it as the **vol sensitivity** or **SSR**, interchangeably. Given the empirical stability of suitably defined shape curves -- with a good fitter! -- this one number is the secret to an efficient understanding of vol dynamics. [One can also define sensitivities for the shape parameters, and use them if evidence emerges that they are needed. There is some for the curvature on longer horizons these days; ask us for details.]

Its value tends to be around 1.3, somewhat lower in quiet/rising markets, somewhat higher in volatile markets. On big down days it tends to be close to 2.0, at least for shorter terms. There are some theoretical reasons to expect it to "usually" be in the interval [1,2]. (The only prolonged exception to this rule of thumb we are aware of are KOSPI options, where SSR is often above 2.)

Actually, we are ignoring one wrinkle in the story here, namely term-structure. In the "olden days", till perhaps 2014, there where very few days on which the vol sensitivity showed evidence of term-structure. Slowly, more and more days started to show such evidence. Nowadays, it is quite common to see clear, usually decreasing, term-structure for SPX vol sensitivities. (Perhaps due to the effect of auto-callable structured products.)

That being said, there are probably few market participants for which the prediction and use of an accurate SSR term-structure is crucial. For most traders, just knowing whether SSR is closer to 1.3 or 1.0 or 1.7 is sufficient. At that level, even term-structure is not that important. At this level of precision, the vol sensitivities also show significant stickiness from day to day, so using a value determined from the previous day's data is a good start (though one can do better).

In any case, the (realized) vol sensitivities can be determined reasonably well using, say, one day's worth of 1-minute or 5-minute snapshot data using a linear regression, with some outlier treatment for robustness. If needed one could even run an on-line regression to update it intraday.

There is clear statistical evidence that on some days the vol sensitivities are larger for larger horizons, especially for shorter terms. In other words, using 5-minute rather than 1-minute snapshots one might find larger values for the vol sensitivities. We leave a more detailed discussion of this issue and its repercussions for trading to another day. But independent of the details here, we would in all cases recommend to smooth the term-structure of the vol sensitivities before using them in trading. An example of the final result of our approach using 1-minute snapshots is shown below (in green) for SPX on 2019-08-05 (contact us for details):

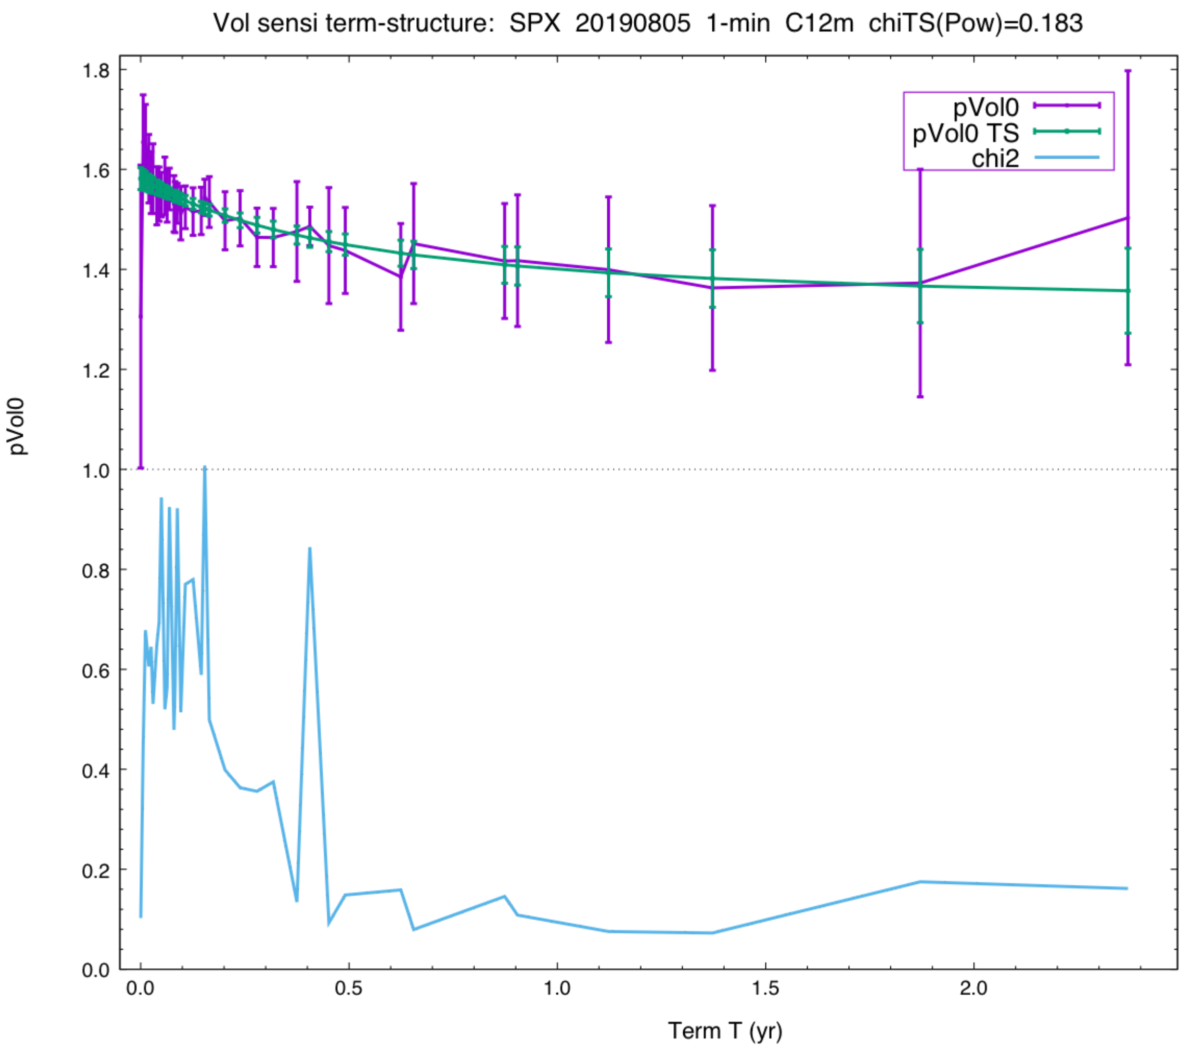

With this out of the way, what does the vol dynamics described above predict for the correct SPX deltas? An example is shown here, for the second monthly expiry (T = 2019-09-20) on 2019-08-05:

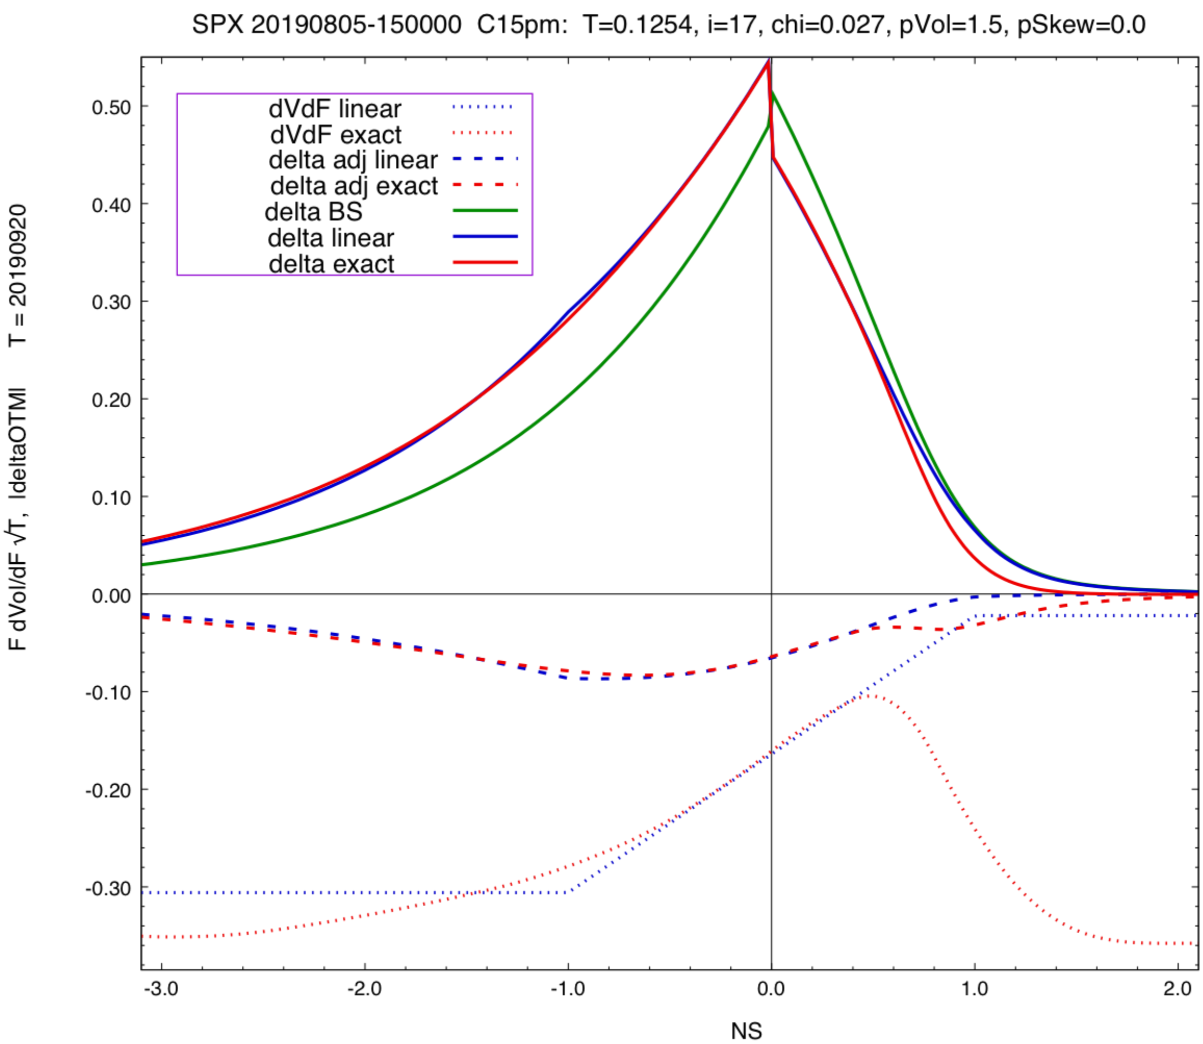

There is a lot of information in this plot:

- The dotted lines represent the calculated dVdF according to the estimated vol dynamics (i.e. with measured vol sensitivity 1.5). For ease of intuition, we use normalized-strike rather than strike on the x-axis; the dVdF (and other) shapes are comparable across terms then. We also normalized dVdF by multiplying by F and sqrt(T); that makes it dimensionless and now not just the shape but even the overall levels are comparable across terms and underliers.

- The red dotted line is the “exact” normalized dVdF, the blue dotted line a linear approximation (flattened at +- 1 NS). On the put side the linear approximation works quite well, but on the call side it breaks down quite suddenly around +0.5 NS. We’ll discuss below why this happens. Let us mention why we show the linear approximation at all: Namely, for various reasons some firms use the value of dVdF that can be calculated analytically for the ATF strike -- there is a simple formula for it in terms of the SSR and the dimensionless skew -- for all strikes. The linear approximation is an improvement on using just the ATF constant term, and there are many names and/or market regimes where it works reasonably well. However, this is mostly not true for SPX nowadays, at least in the call wing.
- The dashed lines are the predicted delta adjustment, i.e. the second term in eqn. (1). Again, we show the “exact” answer, and the result if one uses the linear approximation for dVdF (with flattening in the wings).
- Finally, the solid lines are the various predictions for delta. To be precise, we show the absolute value of the OTM delta, i.e. -deltaPut for K < F and deltaCall for K > F. The green line is the Black value. The red and blue lines represent the final delta according to the exact vol dynamics and the linear approximation for dVdF, respectively.
- Here the delta adjustment is always negative. (This is not always the case, even for SPX, as we will see; for big tech names it tends to be positive for most of the put wing). The biggest absolute adjustment is on the put-side in this case, about -0.08 for NS = -1. But even the ATM value is a substantial -0.06.
- What is particularly interesting is the call side: Even though the *absolute* value of the adjustment is significantly smaller for the (OTM) call deltas than for the puts (this is not always the case, as we will see), in *relative* terms it is huge. In fact, for NS > 1.2 or so is the exact call delta essentially zero here, when the naive delta can still be 0.03. Even at NS=1 is the correct delta about half of the naive one.

There is nothing special about the as-of-date 2019-08-05 or the expiry shown in regards to the delta adjustment. Yes, this was a big down-day for the market, but the vol sensitivities were only about 0.1 higher than on the surrounding days of Aug 2 and 6. Around ATF the delta-adjustment is directly proportional to the difference SSR - 1 (and the dimensionless skew parameter), so the bigger the SSR, the bigger the ATF delta adjustment. To some extent this also tends to hold in the put wing, but in the call wing around and above NS=1 the dependence on SSR is quite weak.

Instead, the behavior in the call wing is essentially all due to the vol shape there. In fact, we can now explain why the linear approximation for dVdF breaks down so dramatically for NS > 0.5 on this and other days. The reason is that the linear approximation for dVdF works well for strikes where the vol curve is well-described by the same parabola that describes the behavior in some region around ATF. On this and many other days that is not the case; the “knee” on the call side where the vols start increasing again is not described by that parabola at all. The proof is easy: The ATF curvature is negative on this day and for this expiry, so it can’t possibly describe the region around the “knee”, which has positive curvature. For reference, here the vol curve from which the above dVdF prediction has been derived:

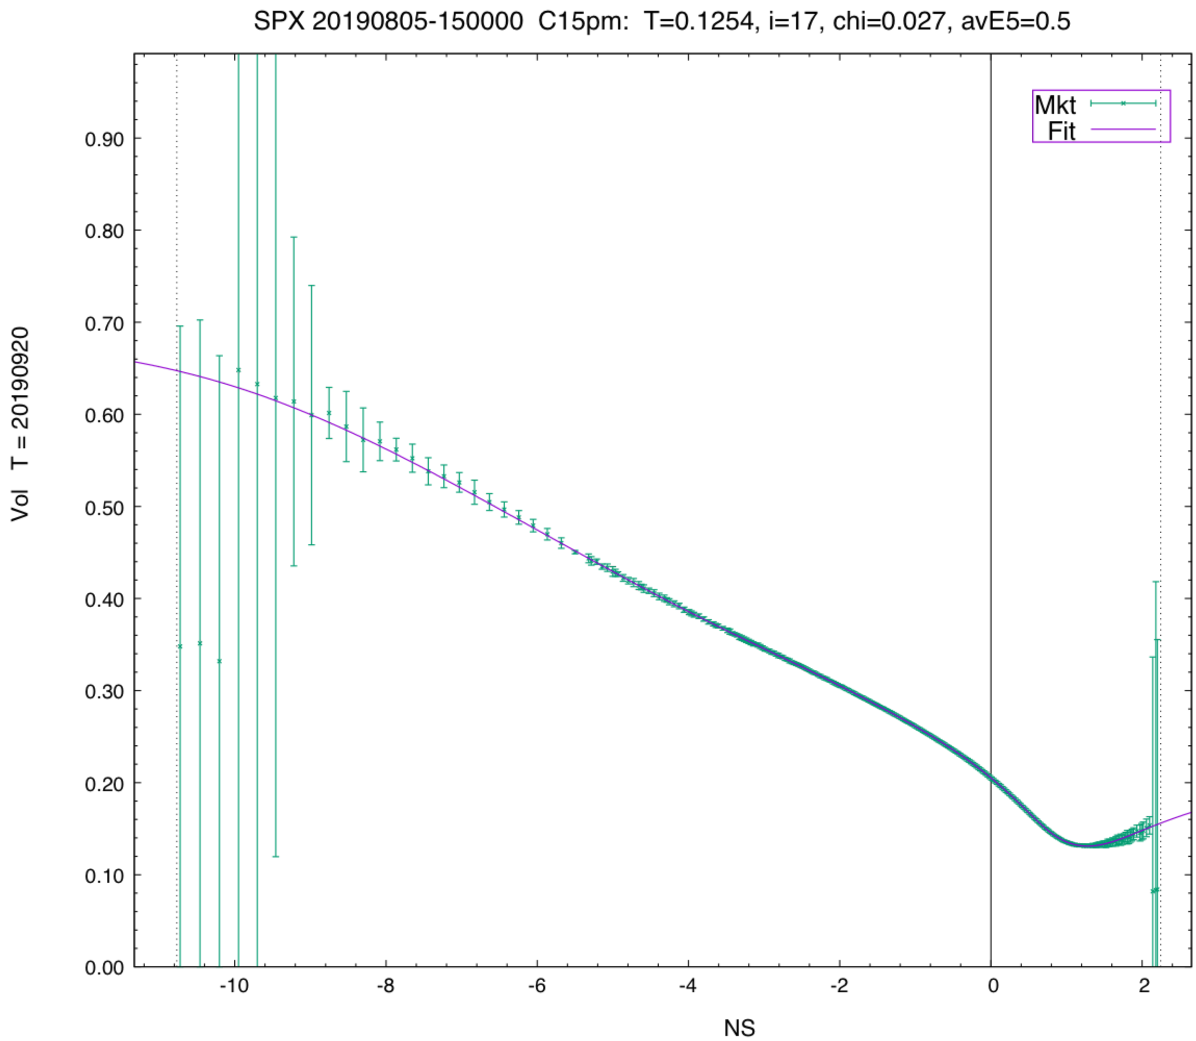

Even when the ATF curvature is not negative will the ATF parabola often very much fail to describe the behavior around the “knee”. Strange -- and occasionally stranger -- vol shapes in SPX have been with us for many years now. (Negative ATF curvature around FOMC and other big events has been around since about 2011 in SPX; for much longer around earnings for big tech names.) Various phrases, like the “Trump put”, have recently been used as the strange SPX vol curve shapes have gained wider attention; but the general phenomena of “rich shape curves” has been around for longer than Trump as president.

Getting back to dVdF, it is somewhat surprising how the very non-trivial shape we saw above derives from just two inputs: (i) a very simple specification of the spot-vol dynamics in terms of just one number (perhaps with a little bit of term-structure these days), and (ii) the precise shape of the vol curve at a given term.

The reader might ask for a more direct proof of this claim. Here it is. Working with the same 1-minute snapshots we used to estimate the vol sensitivities, we can for each strike regress vol changes versus forward changes for a direct empirical estimate of dVdF for each strike. Of course, there is a little bit of “hair” on this calculation, since, for example, the moneyness of a strike changes during the day as the underlier moves around. One can estimate how much of an effect this is -- e.g. by doing this regression only over one hour at a time, etc -- and it is not that large. In the following plot for the (normalized) empirical dVdF obtained this way we use NS for the x-axis (using averages of forward, ATF vol and time-to-expiry to convert from strike K to NS), and indicate the range of the underlier move over the day as dotted lines.

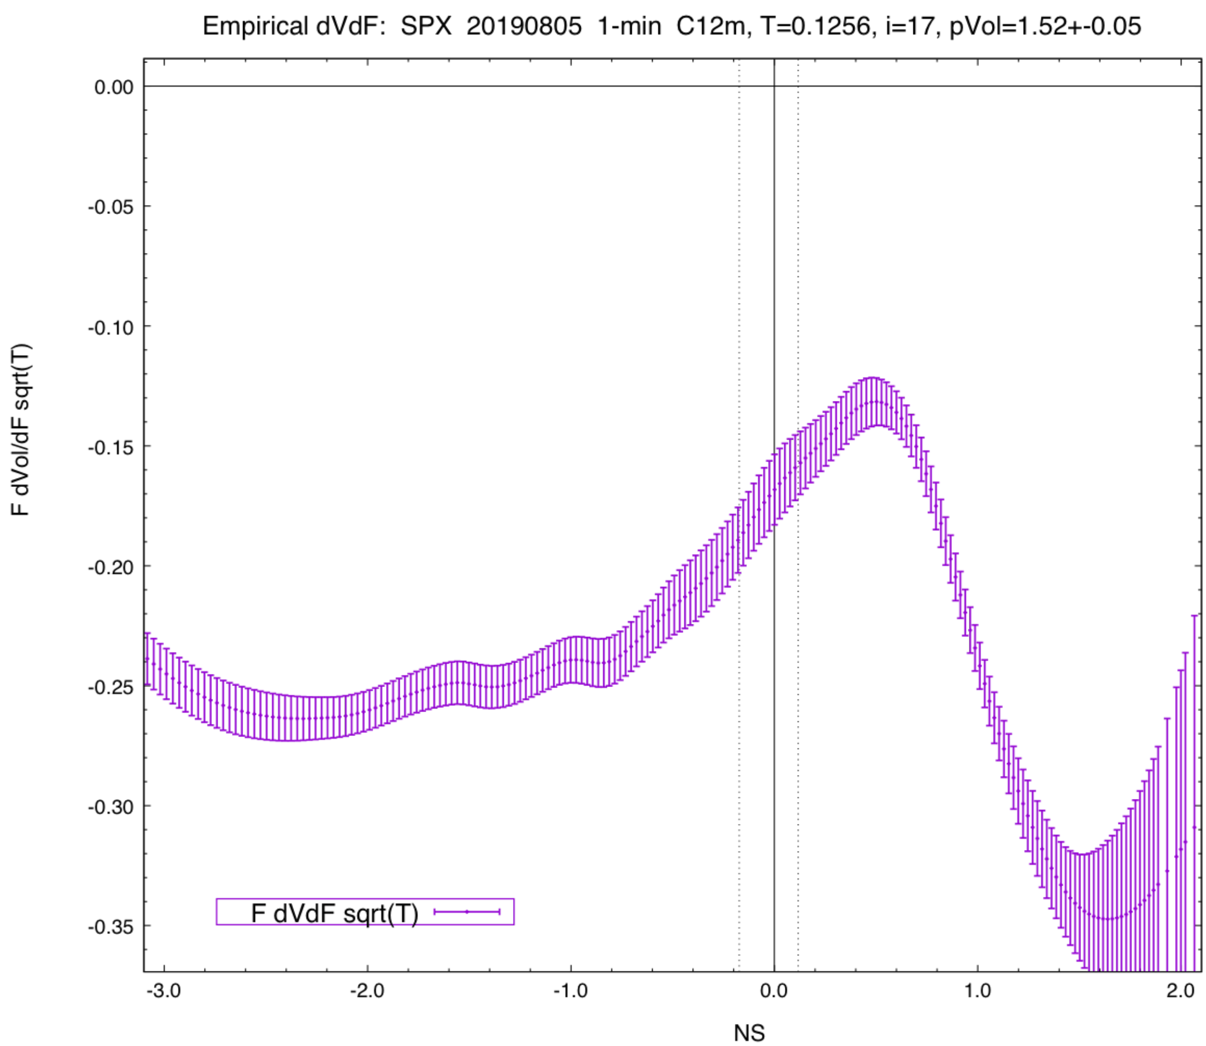

Note that the dVdF estimates are quite smooth as a function of normalized strike, since we used the implied vols from our fitted curves to estimate dVdF (in all cases are our fitted vol curves excellent representations of the market implied vols). Using market implied vols directly gives results consistent within errors, with more noise. The “true” error bars on the above estimates of dVdF should perhaps be 2x larger in the wings than the ones shown above (which are essentially just the “statistical” ones, ignoring several “systematic” ones). Considering this, the agreement with the previous “exact” dVdF estimate (the red dotted line two figures up) is quite amazing, especially given the simplicity of the assumptions used for the latter. Note that the latter is also estimated from just one snapshot of the market (15:00:00 in this case). Using different snapshots, or, for that matter, different (sensible) algorithmic parameters for the fitter or curves, gives very similar results.

Note that the “null-hypothesis” for spot-vol dynamics that one might use in the equity domain, namely “sticky-strike”, would mean that dVdF is flat at 0. The above results reject that idea on both the absolute and standard deviation level conclusively (ditto for “sticky-delta” dynamics, or any simple “combination” of the two).

We should also mention that at least parts of the above story are well-known to serious SPX traders these days; for example that the correct call delta is *much* smaller than the naive Black(-Scholes) one for NS above 1.0 or so.

Before we summarize our conclusion, let us look at some periods in the past where vol shapes or vol sensitivities or both were quite different than in the above example.

Let us look at two days in January 2018. On 2018-01-22 vols had just come off the historical lows of 2017 but the market was still setting new all-time highs. That the market was still complacent can be seen in the vol sensitivities that were close to 1.0 (or even below) on this day. On 2018-01-29 the market was beginning its first down move, and vol sensitivities had moved up to more typical values around 1.3.

Here the dVdF plot for the term with 11 days to expiry on 2018-01-22:

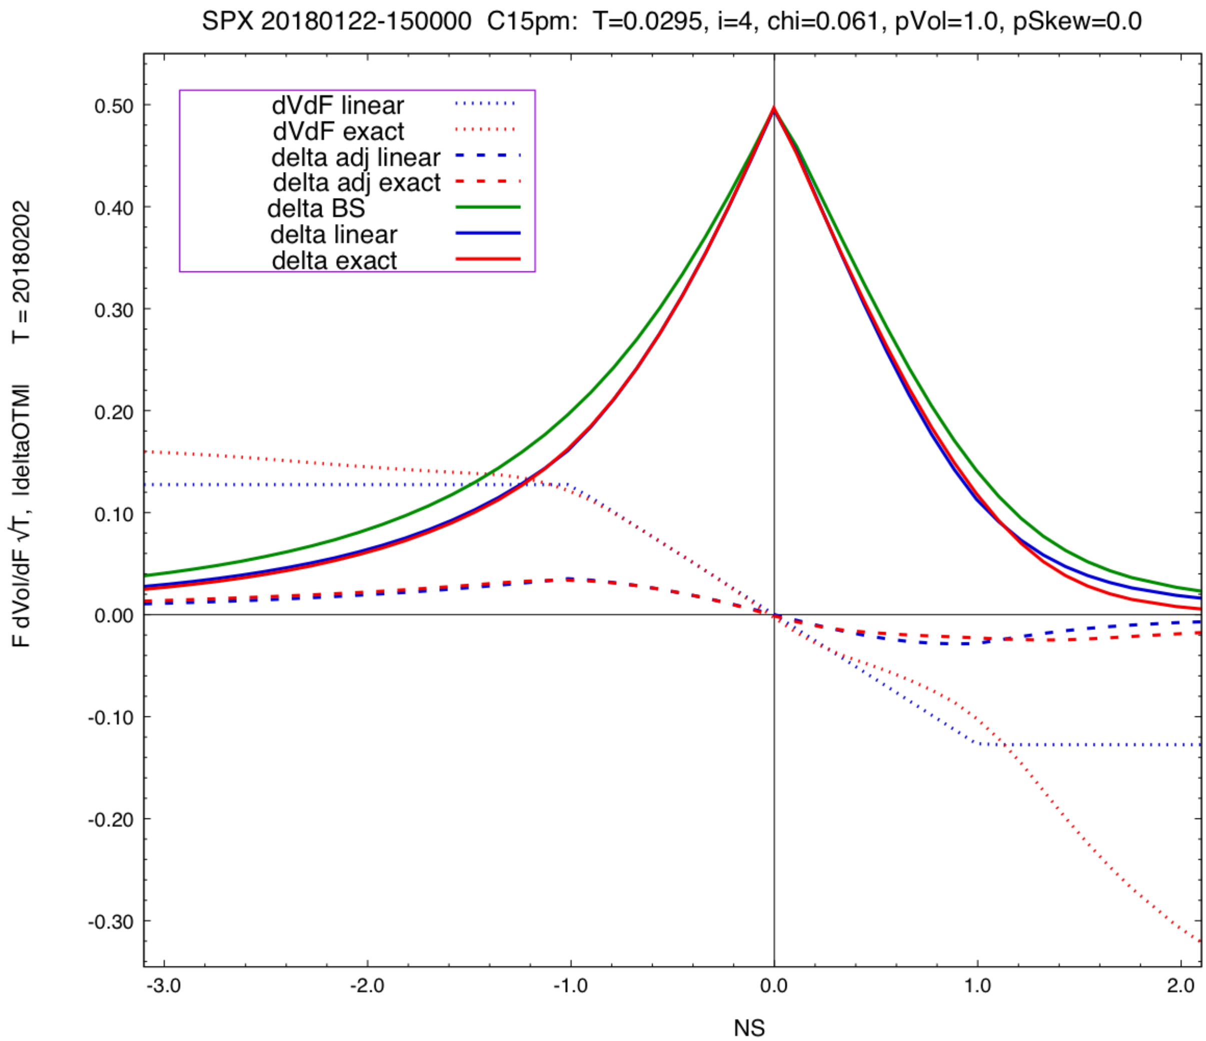

Note the qualitatively very different shape of dVdF from the previous example. It is positive on the put side, for example (so OTM put deltas are *smaller* in magnitude than the Black-Scholes ones). Does this agree with the empirical dVdF estimate? Yes, as can be seen here:

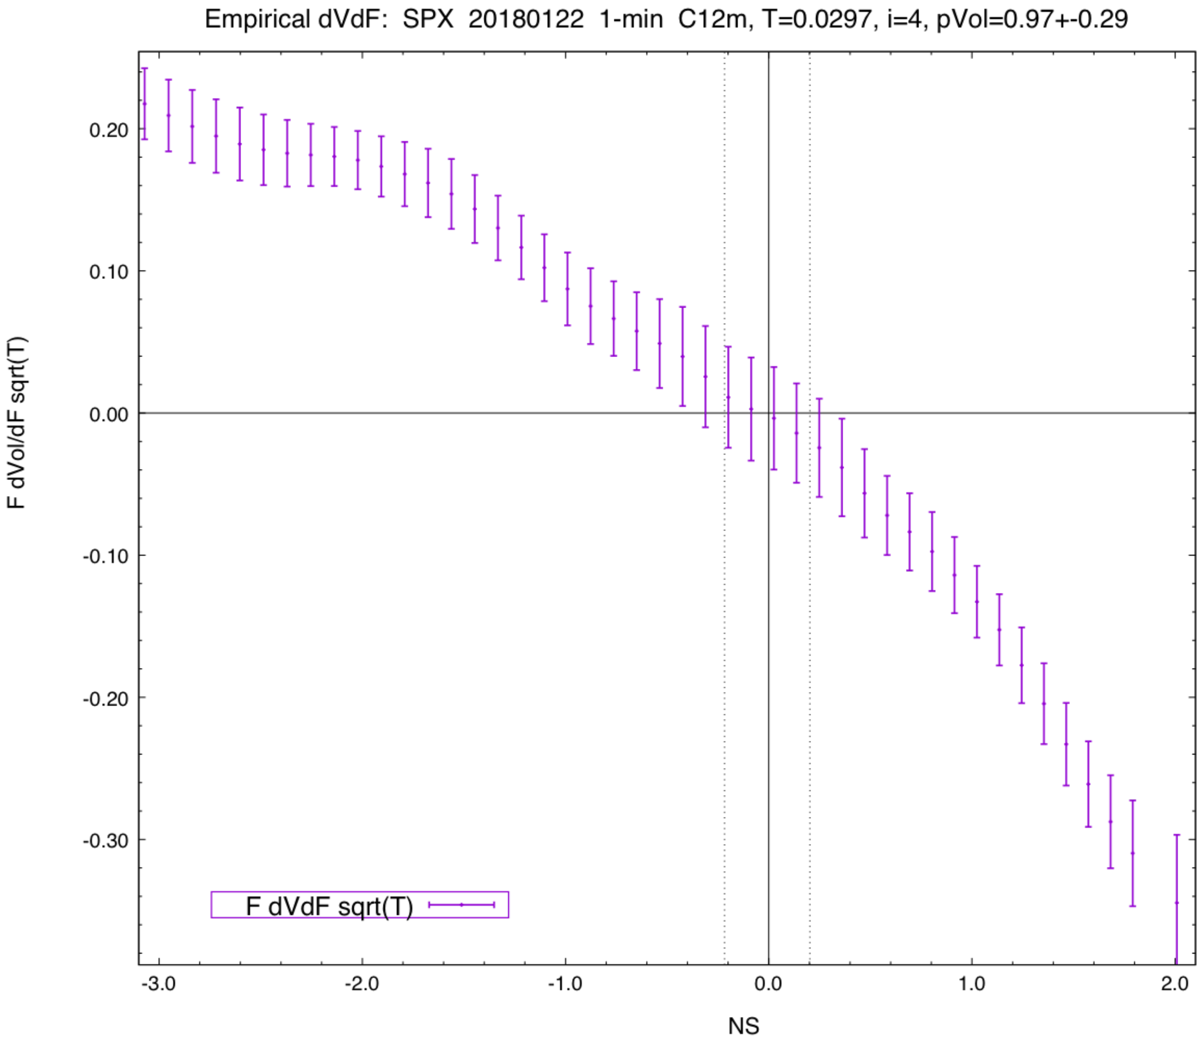

Note that SSR = 1, like here, does not mean that vols are sticky-by-strike; it only implies that for ATF vols, as illustrated above.

Finally, here the dVdF plot for the term with 14 days to expiry, on 2018-01-29:

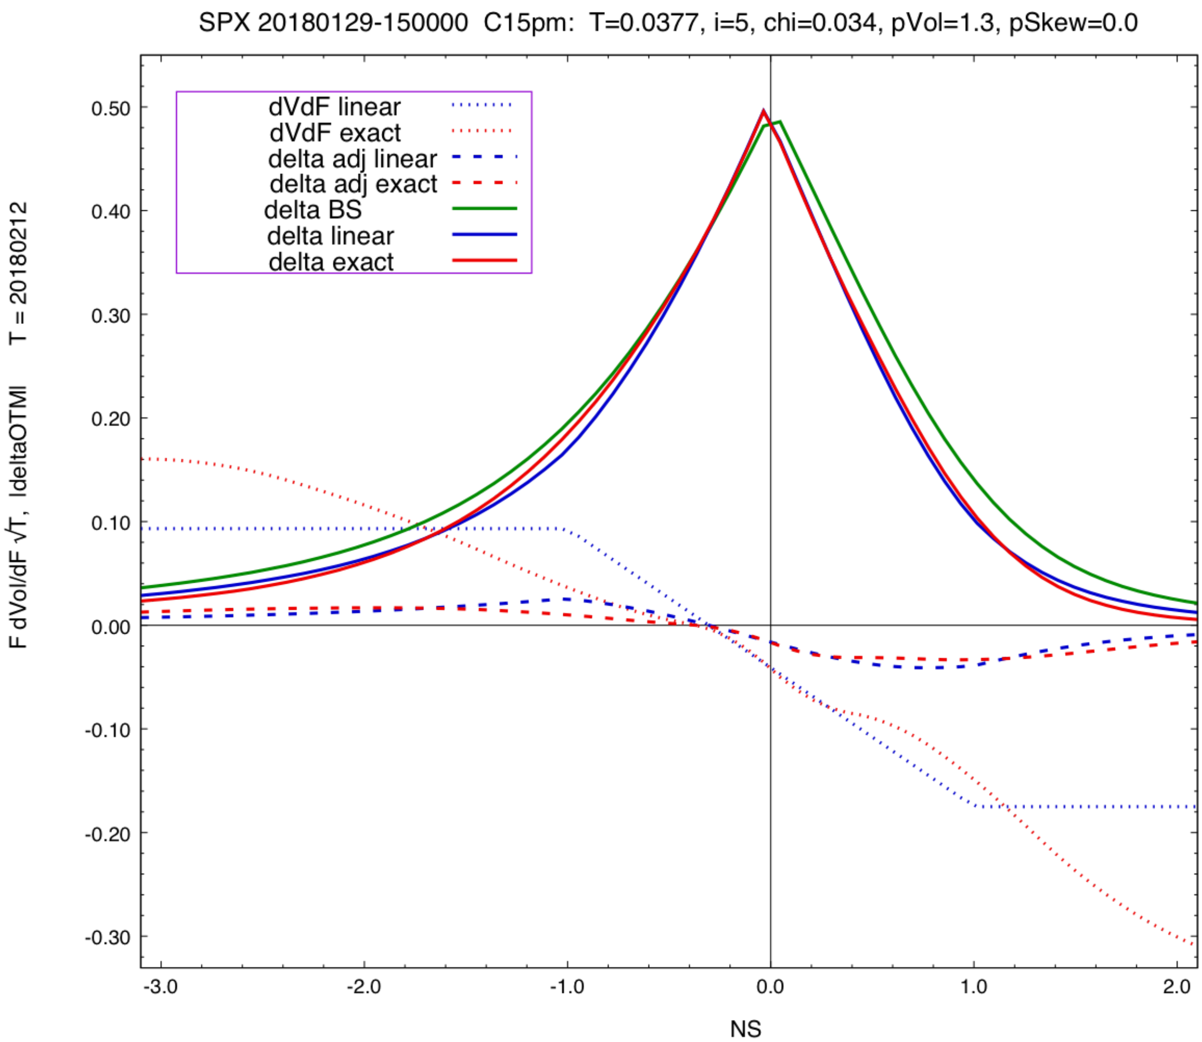

On this day the empirically regressed SSR is close to the typical value of 1.3. Again, the empirical dVdF from the whole day is in excellent agreement with the one-snapshot prediction:

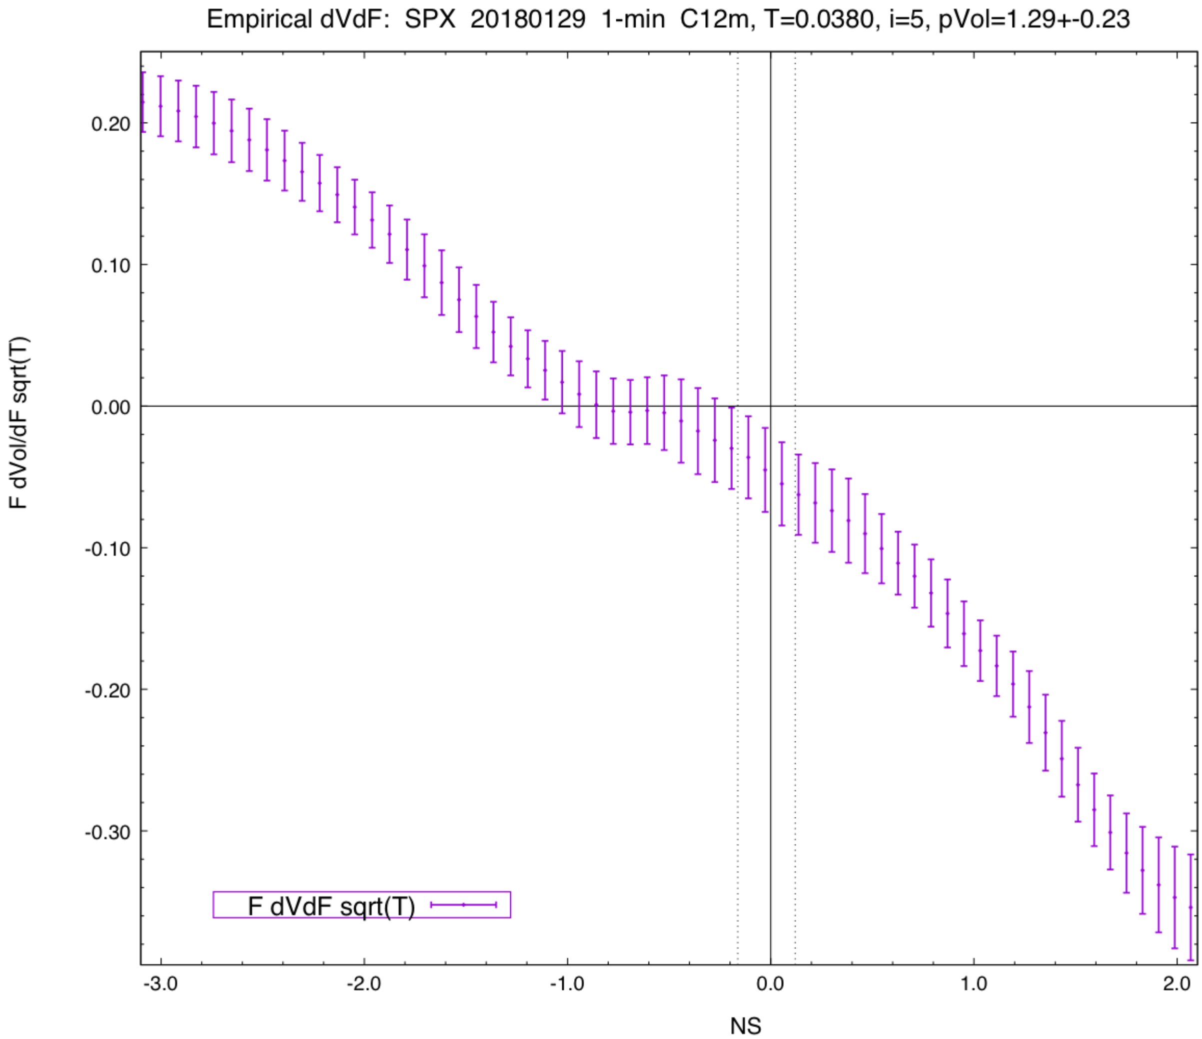

Let us summarize. Starting from first principles we have seen that calculating the optimal delta of a vanilla option is equivalent to knowing spot-vol dynamics. By factoring vol curves into an overall vol level and a suitably defined dimensionless shape curve, vol dynamics becomes very simple. One essentially just has to estimate one number (per term, perhaps), the vol sensitivity, aka skew stickiness ratio SSR. This can be done from one day -- or even one hour -- worth of data, using a (robust) linear regression, at least for liquid names.

The adjustment to delta can be very large, for both calls and puts. The absolute value of these adjustments typically reaches a maximum somewhere between NS -1 and +1. The *relative* amount of the adjustment tends to be largest on the call side (for equity-like names), where it can reach almost 100% for tradable OTM calls; in other words, the optimal delta is very close to 0 while the naive delta can still be significant. The detailed pattern of dVdF and the adjustment

-- even the sign -- varies in complicated and sometimes perhaps surprising patterns. But it can all be explained in quantitative detail using just the vol sensitivities and the precise shapes of the vol curves.

We have illustrated the power of this approach for SPX, but the same approach can of course be used for other liquid indices, ETFs and stocks. For less liquid names, it is statistically harder to estimate the vol sensitivity directly, due to wider spreads and more noise. (In many cases a lower skew also makes it harder). Fortunately, for the same reasons, a precise value of SSR does not matter as much; delta adjustments in the wings can still be very large. One can use some simple models and intuition backed up by historical analysis to tie stock and ETF vol sensitivities to index estimates, for more robust daily estimates.

The important inputs in out approach are:

- Good vol curves that can actually match the market implied vols for liquid names (none in the public domain do).
- A robust and fast fitter to estimate those vol curves and tie them together into sensible (arbitrage-free, tradable, etc) vol surfaces.
- A proper understanding of vol dynamics, and the data and analytics to estimate the vol sensitivities reliably every day.
- For real-time trading use all of this should be packaged into an efficient and easy-to-use C++ library (with higher-level language wrappers on top).

Finally, let us note that a proper understanding of vol dynamics does not just lead to more efficient and/or cheaper hedging of an options portfolio, but also matters significantly for a proper understanding of PnL, risk management, and realistic scenario analysis.

For more information please contact us at info@VolaDynamics.com