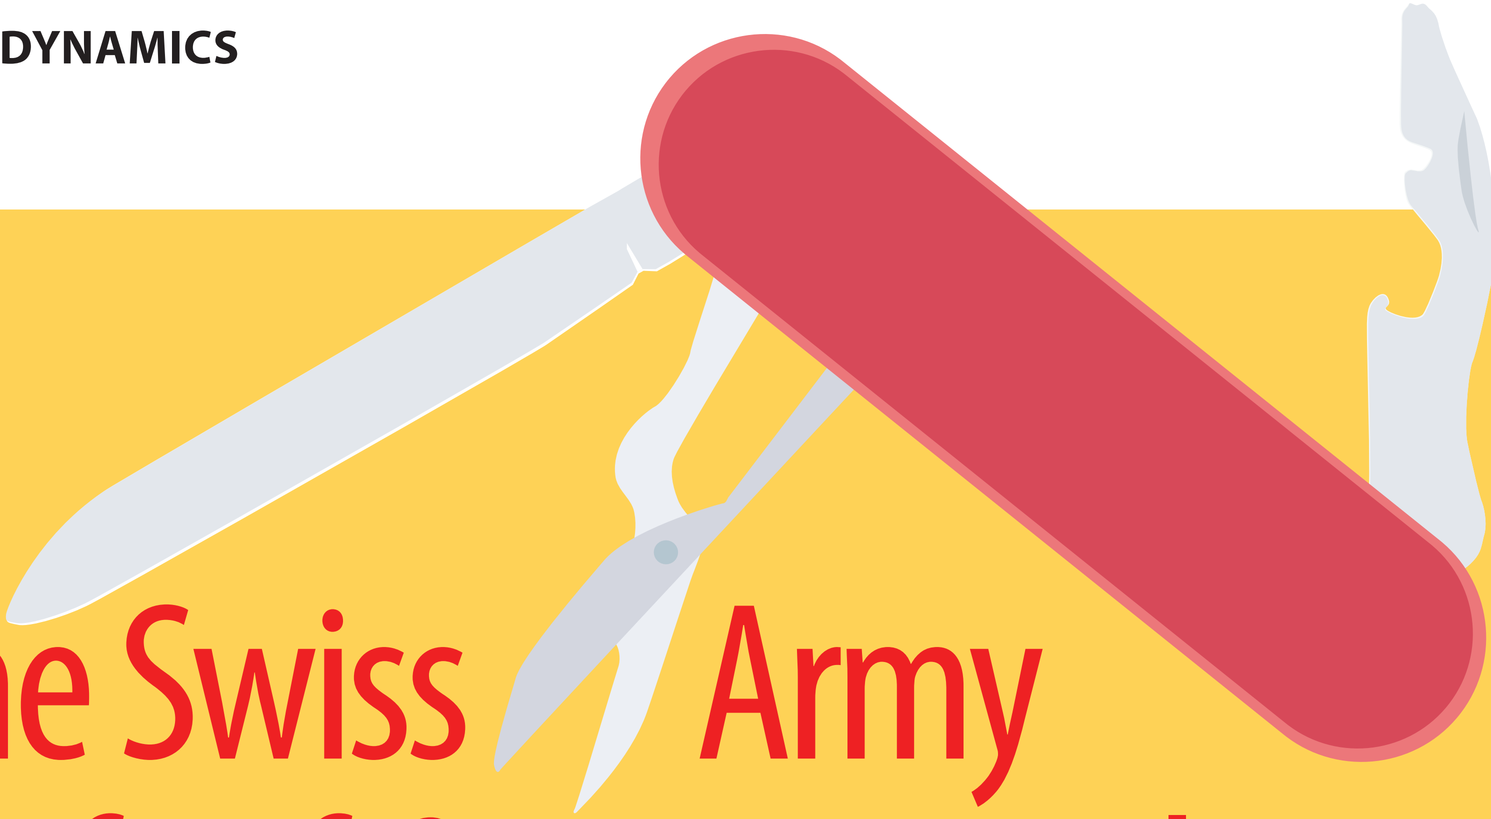

The Swiss Army Knife of Options Analytics

Finally, a solution to the most vexing problems of options pricing and fitting.

Vola Dynamics provides analytics for options trading and risk management, as well as portfolio, PnL, and scenario analysis. Since its founding in New York in 2016, it has quickly established itself as the only third-party vendor to provide top-tier market-maker-quality valuations for vanillas and vol derivatives. Its vol fitter is generally acknowledged to be the best in the industry, and its American option pricer (with proper handling of cash dividends, etc.) is probably the fastest available anywhere.

Vola's analytics library produces easy-to-use, tradable, real-time volatility curves and surfaces, that can match the most challenging markets in a sensible, arbitrage-free, and robust manner. Intuitive, parametric vol curves are the crucial ingredient in allowing efficient signal research. The accurate spot-vol dynamics available as part of the library gives better Greeks and more realistic scenarios, among other benefits.

As a battle-tested easy-to-integrate solution in C++11 (with wrappers for Python, Java, and C#) for any platform, it is available at a fraction of the cost it would take to build and maintain a comparable system in-house (if possible at all), thereby eliminating one of the big barriers to entry in both the listed options market and derivatives markets more generally.

The founders, Timothy Klassen, Jiri Hoogland, and Misha Fomytskyi, have a combined 50+ years of experience in trading and modeling listed vanillas, as well as flow and exotic derivatives in all asset classes, at firms like Goldman Sachs, Morgan Stanley, and Getco. Timothy Klassen designed the 'new VIX' for the CBOE in 2003 (when at Goldman).

The pricing and fitting problem

A fundamental problem that Vola Dynamics solves for its clients is what is often called the 'volatility surface-fitting problem.' It is one of the holy-grail problems of quantitative finance, especially, but not exclusively, in the equity markets. It is nowadays really a complex of multiple, interconnected problems, all of which contribute to the sizable barriers to entry that exist for derivatives markets these days, and arguably make the options market significantly less efficient and useful than it could be.

Recall that implied volatility surfaces (and borrow cost curves) are the standard concepts used to summarize the vanilla options market in an intuitive and compact manner. They provide the fundamental building blocks for trading and risk-managing vanillas (listed and OTC), as well as the foundation for flow and exotic products modeling and trading.

Here is a list of the main subproblems:

Cash dividend modeling [1]: Surprising as this may sound, even to some professionals, half a century after Black–Scholes there is no agreement on how to handle cash dividends, even in the context of vanilla options! Different dividend models will give rise to (potentially very) different implied volatilities – one of the main reasons why every (equity) options trading team speaks a somewhat different dialect of the Black–Scholes language.

Borrow costs: To have proper put-call-parity, stock and ETF options have to be priced with a term-dependent borrow cost. There is no way to get up-to-date and useful borrow cost information directly, so borrow costs have to be implied in real time (or at least daily, using intraday data) from the options market itself, just like implied vols.

Vol curves: None of the curves discussed in the public domain (SABR, SVI, SSVI, two parabolic or cubic polynomials joined at-the-money (ATM), etc.) can actually fit vol skews of liquid underliers/terms, like SPX, SPY, E-mini futures, QQQ, AAPL, AMZN, etc., in a bias-free manner. One of the reasons is that vol curve shapes these days include many with negative curvature around ATM. This started around 2005, with the now (in-)famous W-shaped vol curves of big tech names around earnings, then spread to US indices, and is now also present in most global index vol curves (KOSPI, NKY, HSI, etc.) for shorter maturities. Most recently (2019), vol curve shapes have become even stranger, with super-liquid names like SPX having features variously described as being due to the 'Trump put' and such. Such shapes are simply an expression of the bimodal (or higher-modal) expected underlier distributions around future events with potentially very different outcomes, of which in the current global environment there is no end in sight. The 'funky' shapes seen in equity vol curves for a while should therefore become a common feature also for other asset classes (and create serious friction with simple curves/models like SABR, the Vanna–Volga method, etc., that are often used there). Many firms have given up on the use of parametric curves, since they could not design any that match the market for liquid options. This approach is possible, but parametric curves have a number of significant advantages in practice.

No-arbitrage: Finding a vol surface (theos) without butterfly and calendar arbitrage that is 'closest' to the observed implied volatilities (market options prices) is, to repeat, one of the holy-grail problems for practitioners and academics alike. [Note that the existence of an arbitrage-free vol surface is equivalent to the existence of a local vol surface, which is still an important starting point for many problems in the pricing and hedging of light and heavy exotics.] Even the big banks struggle to produce arbitrage-free vol surfaces that match the market for their flow and structured products desk in an efficient, robust, and timely manner.

The speed, scale, and spread race; complexity; and changes in market structure: The last decade has seen important changes in listed options markets around the world, following the US lead. The penny pilot, the adoption of high-frequency technology, the listing of many more individual instruments (weeklies, mid-weeklies, smaller strike increments, options on new underliers like leveraged and inverse ETFs, VIX-related products), the proliferation of exchanges, order types, fee schedules, and trading protocols (e.g., auctions) have led to a reduction in bid-ask spreads for liquid instruments. At the same time, more recently, market-making requirements have largely disappeared, and spreads for some of the less liquid options have widened. The market for these options has become less efficient because market makers do not have to show a two-sided market anymore to trade – they can, for example, effectively internalize their options flow via auctions, at least in the US. (The prevalence of, for example, RFQs has prevented many other listed markets from becoming transparent and efficient in the first place.) Even the largest market participants have struggled to build and maintain the infrastructure necessary to deal with the increased complexity in modeling and engineering. All this has led to a dramatic increase in the barriers to entry for (listed) options market makers, and derivative desks more generally.

Another general problem facing the financial industry these days is the competition for quantitative talent. Even the most famous banks and hedge funds, let alone smaller trading firms, have trouble hiring good PhDs in the face of competition from big tech. And hiring is not enough – the quants have to stay for the long term to be able to design and maintain the increasingly complex algorithms and analytics libraries used in finance today (even for trading shops that can rely on decades of accumulated institutional knowledge, there are typically teams of 15–30 people maintaining the required infrastructure). Related to this, budgets are not as lavish as they used to be in finance, so just being able to pay a quant team up to the task is increasingly difficult. Hence, it is not surprising that many market participants are now interested in outsourcing these kinds of problems.

The solution

The founders of Vola Dynamics have been thinking about the problems outlined above for almost two decades, and have had the good fortune to be able to refine solutions to these problems multiple times.

For the hardest part of the volatility surface-fitting problem, the steps involved are:

1. The design of a family of intuitive, nested volatility curves that can fit all the vol skew shapes seen in the market while making butterfly arbitrage (practically) impossible.
2. The design and tuning of super-fast and robust fitting algorithms for these curves.
3. As part of these algorithms, it is crucial to transfer information from liquid to less liquid options, to eliminate calendar and butterfly arbitrage, all while staying consistent with the market bid-asks, if possible.
4. The algorithms involved take a Bayesian approach (as in probabilistic robotics) of always updating 'error bars' around each of the quantities being calibrated, be they borrow costs or volatility surface parameters. This is important for robustness.
5. Crucially, the algorithms involved at all stages have a notion of what makes financial sense, whether it be put-call-parity, no-arbitrage, or spot-vol dynamics. The notion of error bars allows the algorithms to find the proper compromise, when needed, between enforcing these constraints versus matching the market data as literally as possible.

As alluded to above, another important building block that permeates the Vola Dynamics library is a simple and accurate spot-vol dynamics – that is, a simple parametrization of how a vol surface moves when the underlier moves. In most cases, it requires just one number (the 'vol sensitivity', a.k.a. Bergomi's 'SSR') to realistically describe the behavior of the full parametric vol surface for any underlier move.

The highlights of the current offering of the Vola Dynamics library can be summarized as:

- Super-fast, robust, bias- and arbitrage-free calibration of parametric volatility surfaces for all vanilla options under any market conditions, thereby providing top-tier

Figure 1: Volatility curve fit for the first monthly expiry on the 15:30:00 ET snapshot of 2019-08-26, plotted in normalized strike (NS) space, with the recommended curve type. The input market-implied vols are shown in green

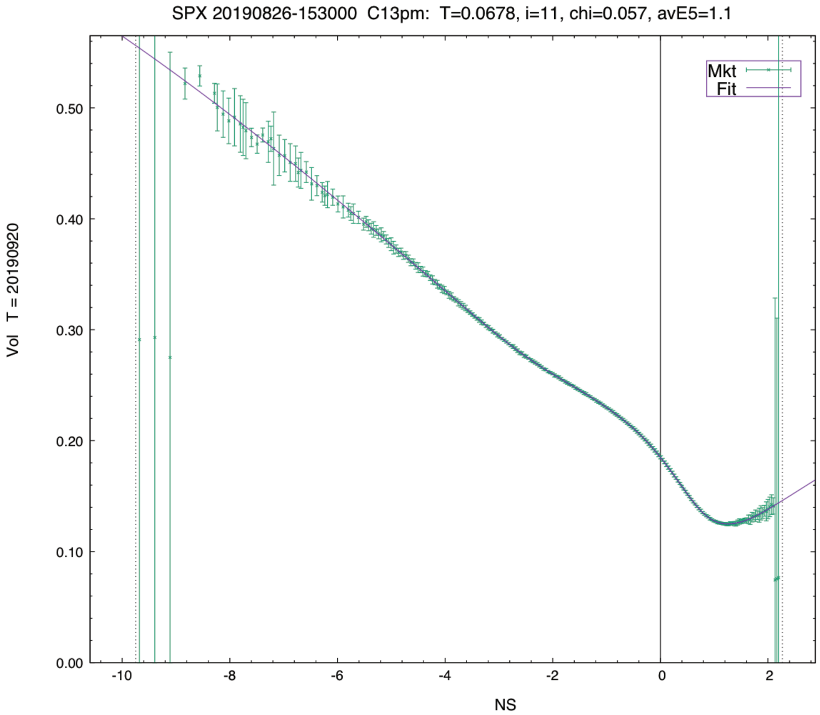

market-maker-quality valuation and Greeks. Even on the level of fitting one snapshot of the market are the fits extremely stable (e.g., no 'floppy wings'). There is also a 'temporal filtering' mode to use information from past fits (without introducing a lag).

- Super-fast, consistent pricing, and hedging of volatility derivatives with vanillas under the same spot-vol dynamics assumptions.
- A simple, universal, and empirically accurate implementation of spot-vol dynamics. This leads to extra benefits, including:
  - Smart Greeks like 'skew delta' and 'skew gamma' (a big deal these days – for example, for SPX calls, as will be discussed elsewhere).
  - Easy and realistic scenario generation for vanillas and volatility derivatives.
- Well-designed APIs (in C++11, Python, and Java; soon also C#) to allow easy integration into any valuation and risk infrastructure.

The library's functionality is constantly growing. The latest addition, for example, is a 'VIX module' that provides, among other features, an analytic beta of VIX futures prices with respect to SPX moves.

There are many concrete examples at VolaDynamics.com of how the Vola fitter produces sensible volatility surfaces in difficult situations. Here, we just show two examples – one involving SPX, the other AMZN.

Let's start with SPX on 2019-08-26, a market up-day (Monday), after the big down-day of 2019-08-23 (Friday). A typical fit is shown in Figure 1.

There are days with even 'funkier'

Figure 2: As in Figure 1, but fitting the same input data to the SSVI/S3 curve

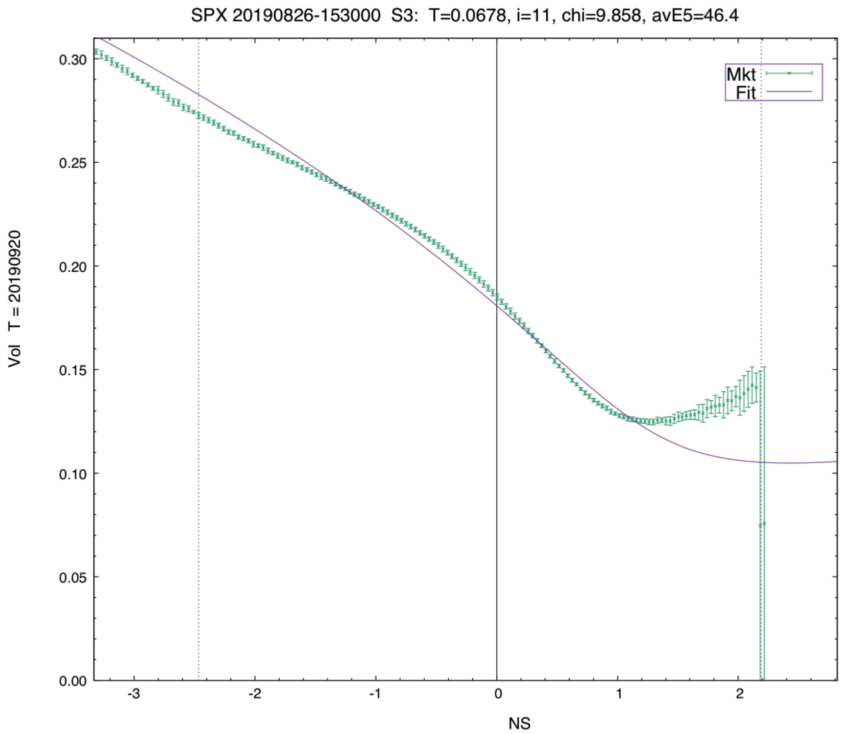

Figure 3: The implied 'total variances' (w) of the fit in Figure 1, as a function of log-strike-over-forward (LKF), for all SPX terms on this snapshot

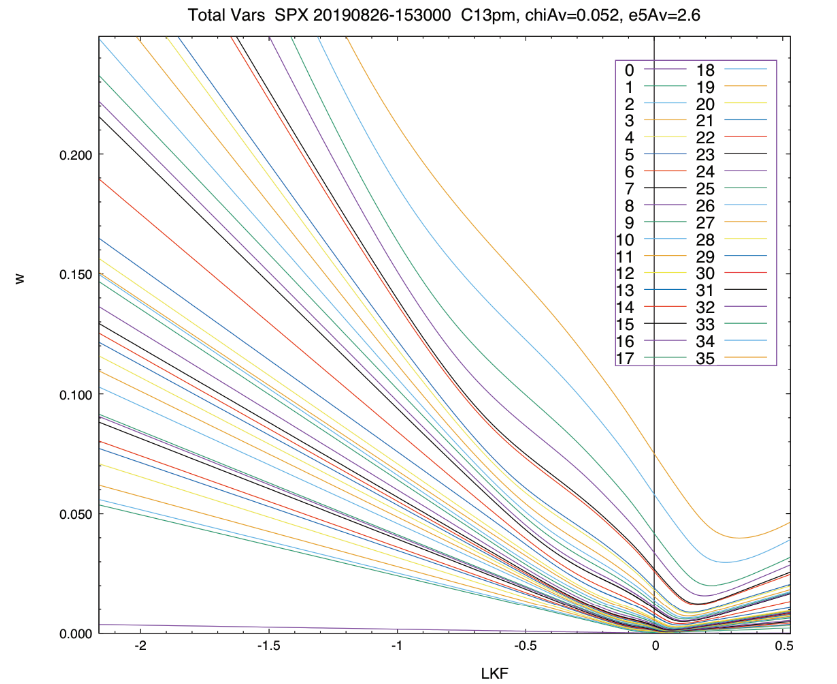

Figure 4: Volatility curves for the first ten expiries of AMZN on the 2018-04-26, 15:45:00 ET snapshot, just before earnings are released, as a function of normalized strike (NS)

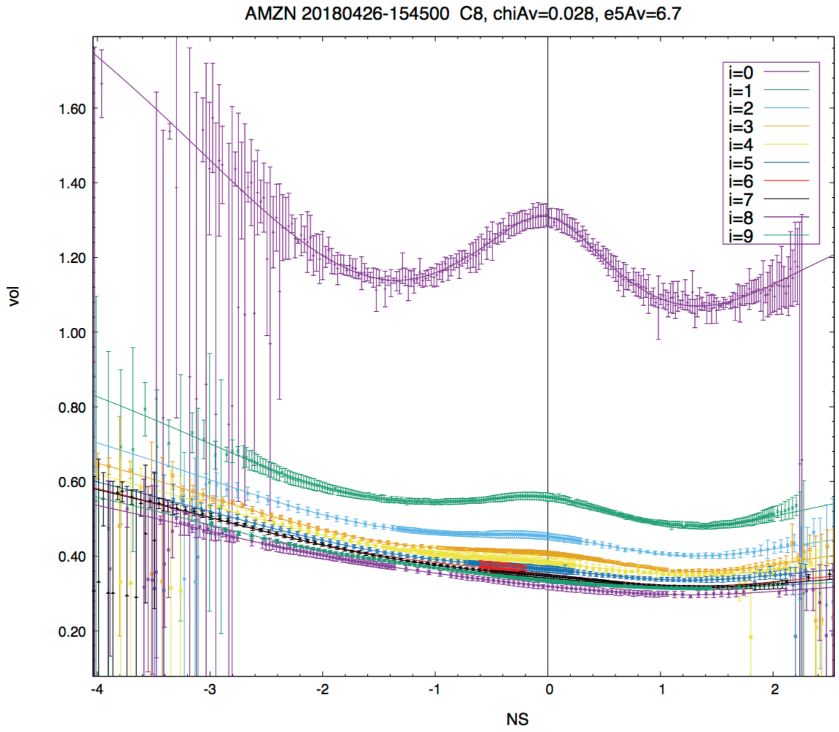

curves for SPX (e.g., 2019-08-23), but note that even here the first monthly expiry has negative curvature around ATM, and quite a steep (positive curvature) 'knee' around NS = 1 in the call wing. Just how strange these shapes are, perhaps only becomes clear when we fit the data with simpler curves. In Figure 2, for example, we show the fit with the well-known SSVI (a.k.a. S3) curve, a three-parameter curve.

Even though we're trying to fit only a much smaller range with the S3 curve – about NS = −2.5 to +2.3, rather than −10 to +2.4 (meaning all data) — the quality of the fit is of the order of 100× worse, by standard statistical metrics, e.g., the reduced chi-square. SABR, another three-parameter curve, is even worse, and SVI (a.k.a. 'S5'), a five-parameter curve, is only slightly better (in the call wing). All of these curves suffer from the fundamental problem that they do not allow one to match, even in principle, the qualitative shapes often seen in the market these days, like negative ATM curvature. In practice, this is the difference between a tradeable fitted curve and a useless one.

To conclude this example, Figure 3 shows the so-called total variance plot, useful, for example, to assess the presence or absence of calendar arbitrage.

The lines for the almost 40 terms do not intersect, showing that there is no calendar arbitrage in the fits (although not proven by this plot, there is also no butterfly arbitrage). This is quite difficult to achieve, as any practitioner who has worked on these problems knows.

In Figure 4, we show some volatility curves for Amazon (AMZN) shortly before the earnings announcement after the close on 2018-04-26.

Serving a wide variety of market participants

Vola Dynamics' clients include a large variety of firms, both small and large, from sophisticated prop shops, market makers, and banks to the biggest hedge funds. They trade equity, futures, and index options across all geographies, as well as vol derivatives. (And some clients are not trading firms, but entities that have to value a large universe of options reliably and accurately every day.) Clients have traded off Vola valuations, even during challenging market conditions around Brexit, the US and French elections, and the February 2018 'Volmageddon'.

The response to the Vola library has been enthusiastic, from day one, as expressed in some client quotes:

"Your vol curves work like magic."

"Impressive library, vastly superior to any other vendor product."

"Modern options market making is incredibly competitive. Automated volatility curves are critical and Vola provides the most reliable and flexible curves in the business. Vola support is also top-notch; you have access to a level of quant knowledge that would be hard to hire."

"We've reviewed a few vendors and none came close to the performance and fitting quality we've seen with your library. In some instances [of other vendors], there were serious gaps between what was advertised and what we've seen, in terms of readiness and fitting quality."

The massively negative curvature in the first term is close to, but does not cross, the lower bound on the curvature coming from the no-ATM-butterfly-arbitrage condition (see section 3 in [2] for details on this bound). Again, it is not trivial to produce an arbitrage-free volatility surface that matches the market extremely well in a situation such as this.

Conclusion

Virtually all market participants need some combination of speed, scale, accuracy, and robustness in their options valuations, but had nowhere to turn for a solution. Every serious market participant had to develop their own in-house solution, with the encompassing opportunity cost, large actual upfront and maintenance costs, as well as a significant risk of outright failure. Numerous vendors have previously tried to solve this problem. The Vola analytics library finally provides all the basic building blocks that serious options and derivatives market participants – whether new or established – need to build and grow their business.

More information is available at VolaDynamics.com.

REFERENCES

[1] Klassen, T.R. 2015. Pricing vanilla options with cash dividends. Vola Dynamics White Paper.
[2] Klassen, T.R. 2016. Necessary and sufficient no-arbitrage conditions for the SSVI/S3 volatility curve. Vola Dynamics White Paper.

W