Arbitrage-Free Parametric Volatility Surfaces and Real-Time Fitting

Timothy R. Klassen

CEO & founder, Vola Dynamics LLC

timothy.klassen@VolaDynamics.com

Global Derivatives Chicago Nov. 2, 2017

Outline

- Dividend modeling (very briefly, details in SSRN paper)
- Implied volatility curve design
- No-Arbitrage constraints in price- and vol-space
  - Including some non-trivial, exact results on vol parameter constraints
- Arbitrage-free vol surface fitting in practice

For background see (available at SSRN):

Pricing Vanilla Options with Cash Dividends

Necessary and Sufficient No-Arbitrage Conditions for the SSVI/S3 Volatility Curve

J. Gatheral, A. Jacquier: *Arbitrage-free SVI volatility surfaces*

J. Gatheral: *The Volatility Surface*, Wiley 2006

Implied Vols and Surfaces

- Implied volatility surfaces (and borrow cost curves) are the standard approach to summarizing 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 flow derivatives and exotics.
- Options market makers use a Black-Scholes framework with some added bells and whistles for valuation and risk management. (BS as language...)
- For more exotic applications: (i) An arbitrage-free vol surface is equivalent to the existence of a local volatility surface. (ii) Having an arbitrage-free, sensible surface for all (T, K) makes calibrating fancy "SLVJ" models much easier.
- Many of the most liquid options allow American exercise: "De-Americanization" is necessary and in principle introduces model-dependence, since American options are (very light) "exotics".
- So, from now on all implied vols will be the "European-equivalent" ones.
- But what do we mean by implied vols in the presence of cash dividends... ?

Dividend Modeling

- Forty years after Black-Scholes there is no consensus on how to model cash dividends (even for vanillas)!
- Cash dividends mean that the observed stock price can not follow geometric Brownian motion (GBM).
- In a vanilla context the question is how to combine the stochastic part of underlier evolution (e.g. *who* follows GBM?) with...
- Three types of dividends:
  - A dividend yield – used to model borrow cost
  - Cash dividends – how most dividends are actually paid
  - Discrete proportional dividends
- Most firms use a *blending scheme* to transition from cash dividends on short end to proportional dividends in long term.
- Proportional divs are also useful in times of extreme uncertainty (market-wise or name-specific). E.g. during 2008 crisis.

Dividend Models

- Main two classes of dividend models are:
  - *Spot model*: The dividends come out of the observed stock price. (Need to modify cash dividends at low stock price.)
  - *Hybrid models*: The dividends come out of a "cash buffer", related to the PV of future dividends: $S_t = \tilde{S}_t + D_t$
- Spot model might seem naively more reasonable, but in practice leads to a lot of complications, since not GBM.
- Hybrid models are much simpler to handle for both vanillas and exotics, since *pure stock* $\tilde{S}_t$ still follows GBM. Can also easily handle credit risk, extension to (light) exotics, local vols, etc.
- We will assume a hybrid model from now on.
- For a detailed review, synthesis, and new results see SSRN paper.

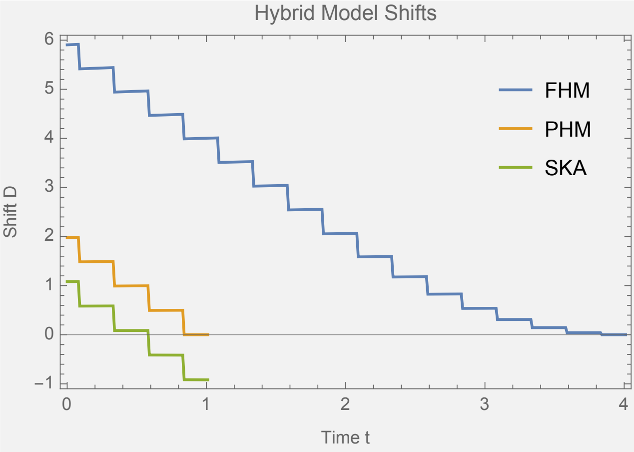

The shifts D_t for various hybrid models with $r = 3\%$, $q = 1\%$, and a quarterly cash dividend of 0.5 first paid at $t_1 = 0.085$, using blending scheme $(2, 4)$.

For PHM, SKA: using $T = 1.01$. FHM = HM2, PHM = HM1, SKA = HM3

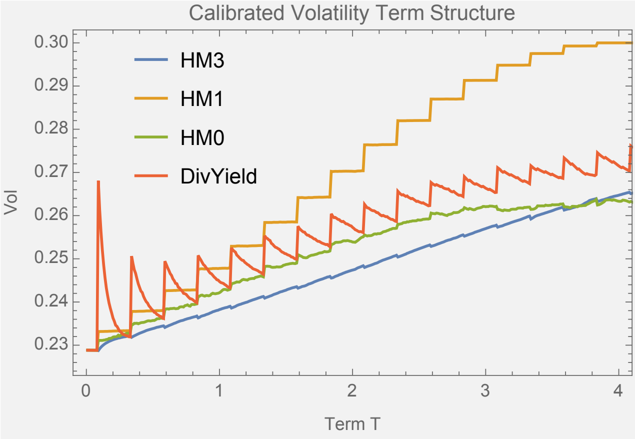

Implied ATF volatility term-structure of various models calibrated to reference market HM2 with $r = 3\%$, $q = 1\%$, $\sigma = 30\%$, and a quarterly cash dividend of 2 first paid at $t_1 = 0.085$. Exercise-style is American.

Volatility Curve Parametrization Wish List

Parametric vol curves are better than non-parametric ones!

- Parameters should have simple, intuitive meaning, esp. first three.
- Parameters should be as "independent" as possible, and stable from day to day (parsimonious).
- Little term-structure, if possible.
- No-arbitrage constraints should be "easy" to incorporate.
- Parametric vols should be easy/fast to compute.
- No hacks! (in wings, etc)
- Vol curves arising from standard "SLVJ"-type model should be fittable within a few bps (at worst).

Our parametrization approach

- Work one term at a time, impose smoothness across terms.
- Factor out overall vol level (ATF) as: $\sigma_0 := \sigma(T, K = F)$.
- Define the "shape" curve $f(z) = f(z|\mathbf{p})$ as a function of normalized strike

$$\text{NS} \;=\; z \;:=\; \frac{y}{\hat{\sigma}_0} \;=\; \frac{\log(K/F)}{\sigma_0 \sqrt{T}}$$

  such that

$$\sigma(z)^2 \;=\; \sigma_0^2 \; f(z|\mathbf{p})$$

- There are no standard definitions – we define the dimensionless "skew" and "smile/convexity" as slope and curvature of the shape curve:

$$f(z) \;=:\; 1 + s_2\, z + \tfrac{1}{2} c_2\, z^2 + \ldots$$

Our parametrization approach (cont'd)

- s_2 and c_2 tend to have mild term-structure (except maybe as $T \to 0$). They are even comparable across names. Have been range-bound for decades.
- Sometimes it is useful to work with s_1, c_1 defined via

$$\sigma(z) \;=:\; \sigma_0\,(1 + s_1\,z + \tfrac{1}{2}c_1\,z^2 + \ldots)$$

- Trivially: $s_2 = 2s_1$, $c_2 = 2(c_1 + s_1^2)$.
- Note that

$$\sigma(z) \;=\; \sigma_0 + \frac{s_1}{\sqrt{T}}\,log(K/F) + \ldots,$$

so that an alternative definition of skew

$$\tilde{s}_1 \;:=\; K\frac{\partial\sigma}{\partial K}|_{K=F} \;=\; \frac{s_1}{\sqrt{T}}$$

- No simple relationships between alternative definitions of curvature/convexity/smile.

No-Arbitrage: Basic issues

- In price space arbitrage conditions are well-known.
- It's easy to detect when there is arbitrage.
- But it is not at all clear, a priori, how to remove it.
- There are an infinity of ways of doing so – almost all of them are "bad", especially when working in price space.
- We find it easier to remove arbitrage in vol-space.
- Due to de-Americanization issue, one has to be a bit careful when moving from price- to vol-space for American options.

No-Arbitrage Constraints in Vol-Space

- **No butterfly arbitrage**: Implied density ρ should be positive:

$$\hat{C}(T,K) \;=\; \int_0^\infty dS_T\;(S_T - K)_+\;\rho_T(S_0 \to S_T)$$

$$\Rightarrow \quad \partial_K^2 \hat{C}(T,K) \;=\; \rho_T(S_0 \to S)|_{S=K}$$

- **No calendar arbitrage**: Total BS variance $w(y) := T\sigma(y)^2$ has to be increasing in T at any fixed y.
- Necessary (but generally not sufficient) constraint on the asymptotic wing behavior of implied vols (TRK 2001, R. Lee, 2004):

$$w(y) \leq 2|y| \qquad\qquad \text{as} \;\; |y| \to \infty$$

Simple consequences: Implied density 1

- Local vols and implied densities can be calculated most neatly in terms of the total variance $w(y) = T\sigma(z)^2$. Eg the implied density:

$$\rho(y) \;=\; \frac{g(y)}{\hat{\sigma}(y)}\; n(d_-(y))\;,$$

where $n(x) = N'(x)$ is the normal density, and

$$g(y) \;=\; \left(1 - \frac{y\,w'(y)}{2w(y)}\right)^2 \;-\; \frac{1}{4}\left(\frac{1}{w(y)} + \frac{1}{4}\right) w'(y)^2 \;+\; \frac{1}{2}\, w''(y)$$

- Absence of butterfly arbitrage: $g(y) \geq 0$ for all y.
- In Black-Scholes case: $g(y) = 1$ for all y.

Simple consequences: Implied density 2

- For our vol curve parametrizations $w(y) = \hat{\sigma}_0^2 \ f(z)$.
- Then $w'(y) = \hat{\sigma}_0 f'(z)$ and $w''(y) = f''(z)$, so that

$$g(y) = \left(1 - \frac{z\,f'(z)}{2f(z)}\right)^2 - \frac{1}{4}\,\frac{f'(z)^2}{f(z)} - \frac{\hat{\sigma}_0^2}{16}\,f'(z)^2 + \frac{1}{2}\,f''(z) \ =: \ g(z)\,.$$

- The vol level appears in only one place! All else only depends on shape parameters.
- Makes the analysis of butterfly arbitrage significantly simpler (but is still very hard in general).
- But for the S3/SSVI curve the butterfly arb question can be completely answered! See below, but first...

ATF No-Arbitrage Constraints

- If $w(z) \;=\; \hat\sigma_0^2\,(1 + s_2 z + \frac{1}{2} c_2 z^2 + \ldots)$, then

$$g(z\!=\!0) \;=\; 1 \,+\, \tfrac{1}{2} c_2 \,-\, \tfrac{1}{4} s_2^2\,(1 + \tfrac{1}{4}\hat\sigma_0^2)$$

- $g(0) \geq 0$ implies upper bound on slope

$$s_2^2 \;\leq\; \frac{4 + 2c_2}{1 + \frac{1}{4}\hat\sigma_0^2}$$

or lower bound on curvature ($c_1 = \frac{1}{2} c_2 - \frac{1}{4} s_2^2$)

$$c_1 \;\geq\; -1 \,+\, \tfrac{1}{16}\, s_2^2\, \hat\sigma_0^2 \;\approx\; -1$$

- Very relevant around FOMC and earnings where not just $c_1 < 0$ but even $c_2 < 0$ can happen!

Specific Curves: Parabolas

- What are simplest possible implied vol curves? Need at least 3 parameters for ATF behavior.
- Vendors often use

$$\sigma(y)^n \;=\; \sigma_0^n \,(1 + s\; y + \tfrac{1}{2}c\; y^2) \qquad \text{(or in terms of } z\text{)}$$

- Obviously has arbitrage in wings for $n = 1, 2$.
- Slight hope for $n = 4$, but would imply symmetric wings, which is intuitively and empirically wrong.
- Positivity has to be enforced too.
- Must do better...

Specific Curves: S3/SSVI

- Simplest sensible curve with 3 parameters ($c_2 \geq 0$):

$$\sigma^2(z) = \sigma_0^2 \left(\frac{1}{2}(1 + s_2 z) + \sqrt{\frac{1}{4}(1 + s_2 z)^2 + \frac{1}{2} c_2 z^2} \right)$$

- Was independently discovered by TRK (2003, "S3") and Gatheral/Jacquier (2013, "SSVI" = Simple SVI).
- Allows surprisingly varied skew shapes, including "takeover-for-cash" curves as $c_2 \to 0$.
- Allows fitting of vast majority of US equity names.
- Very easy to avoid (butterfly) arbitrage.
- In fact, in terms of the dimensionless variables $\hat{\sigma}_0, s_2, c_2$ can completely answer the butterfly-arbitrage question...

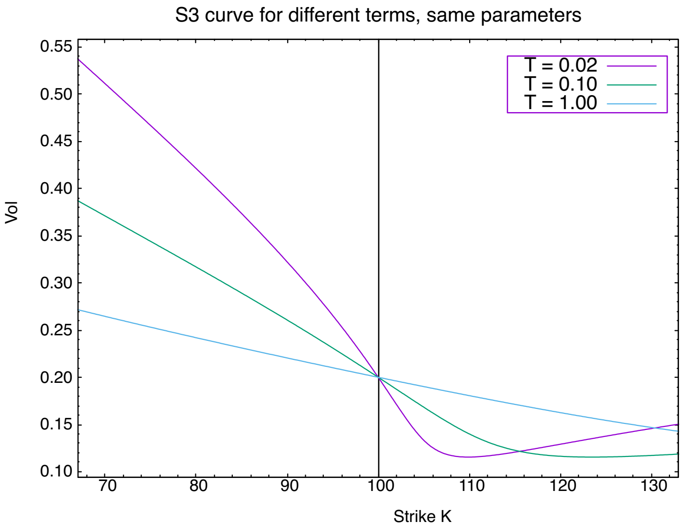
S3 curve for different terms, same parameters
T = 0.02
T = 0.10
T = 1.00
Vol
0.55
0.50
0.45
0.40
0.35
0.30
0.25
0.20
0.15
0.10
70
80
90
100
110
120
130
Strike K

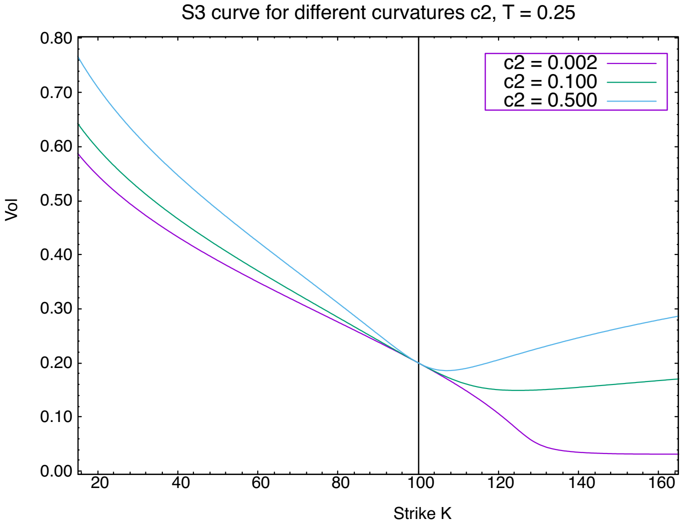
S3 curve for different curvatures c2, T = 0.25
c2 = 0.002
c2 = 0.100
c2 = 0.500
Vol
0.80
0.70
0.60
0.50
0.40
0.30
0.20
0.10
0.00
20
40
60
80
100
120
140
160
Strike K

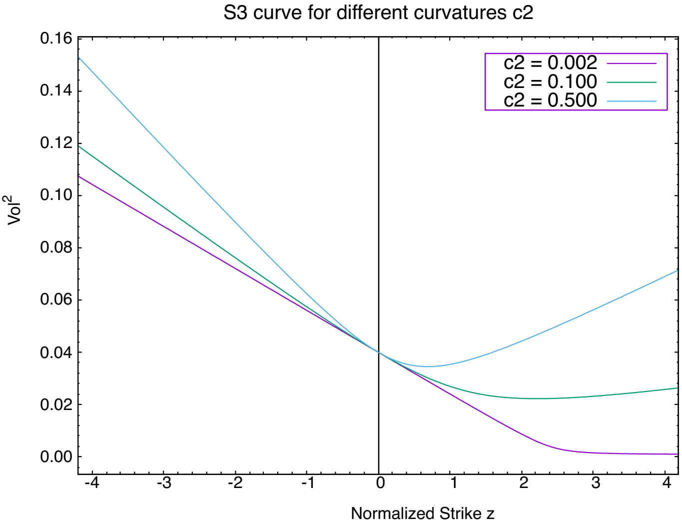
S3 curve for different curvatures c2
c2 = 0.002
c2 = 0.100
c2 = 0.500
Vol²
0.16
0.14
0.12
0.10
0.08
0.06
0.04
0.02
0.00
-4
-3
-2
-1
0
1
2
3
4
Normalized Strike z

Necessary and Sufficient No-Arb Conditions for S3/SSVI

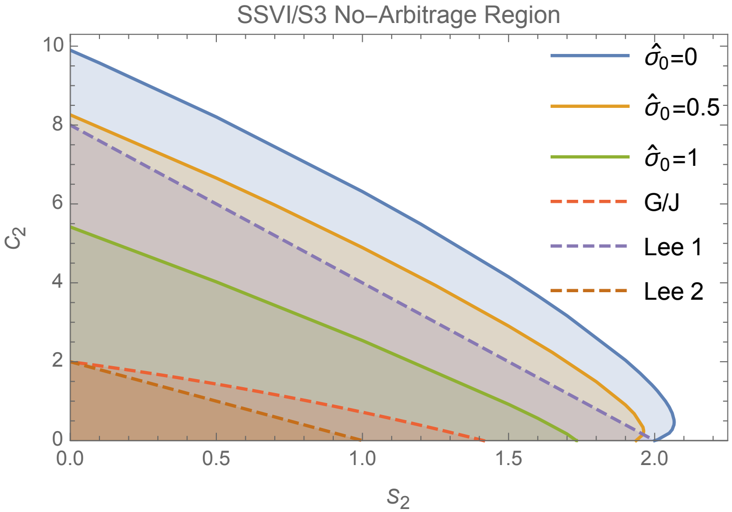

Necessary and Sufficient No-Arb Conditions for S3/SSVI

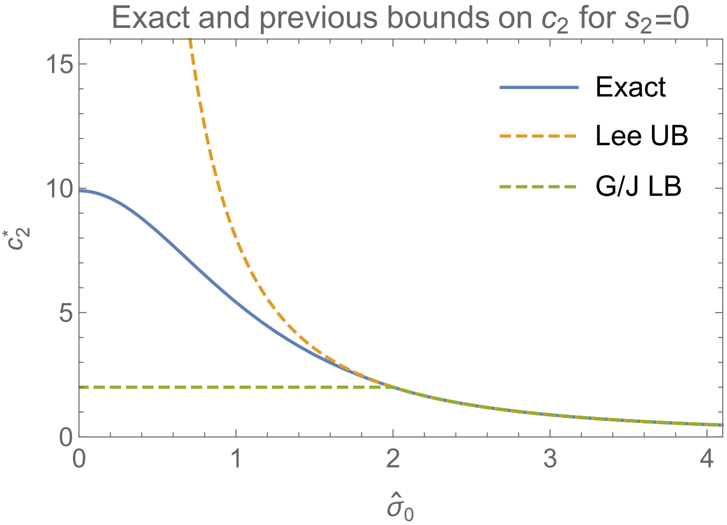

Beyond the Simplest Curves: 5 Parameters (SVI, etc)

- Besides 3 parameters for ATF would be nice to have independent parameters $C_\pm$ for wings:

$$\sigma(z)^2 \;\to\; \sigma_0^2\; C_\pm\; |z| \qquad \text{as} \quad z \to \pm\infty \qquad (\hat{\sigma}_0 C_\pm \leq 2)$$

- For S3/SSVI: $\quad C_\pm \;=\; \sqrt{\tfrac{1}{4}s_2^2 + \tfrac{1}{2}c_2} \;\pm\; \tfrac{1}{2}s_2$
- For Gatheral's SVI and others (JW/L5, TRK) the $C_\pm$ are independent parameters (constrained by $-C_- \leq s_2 \leq C_+$).
- Just some algebra to re-express their "raw" parametrization in terms of natural parameters $\sigma_0, s_2, c_2, C_-, C_+$.
- Can fit some names better than with S3/SSVI.... but not much better in many cases!?
- Certainly can not fit W-shaped curves around events (still $c_2 \geq 0$).

What curves to use for most liquid names?

- For very liquid names (ES, SPX, SPY, other ETFs, AAPL, KOSPI, etc) none of the standard curves (SVI, L5 or amendments) work well, even in the absence of events.
- There is a fundamental problem with the shapes allowed by these curves: Curvature has unique maximum around ATF, but that's not what the market wants!
- Need more flexible shapes that can handle more generic curvature structures, incl. negative curvature around ATF: C5, C6, C7, C8, C10, ...

Volatility fitting framework

- Input to fitter are implied vols with error bars (after proper div modeling, borrow implication, etc).
- All our vol curves have sensible dimensionless parameters (first three are universal), which allows the use of curve-independent heuristics from decades of vol fitting experience across many names, geographies and asset classes.
- Fit one term at a time, for speed, transfer information between terms, for smoothness and stability.
- Minimize chi-square + soft penalties, for robustness and to allow the fitting of terms with less (effective) data than parameters.
- Good microprices help, but even then various heuristics are needed to deal with data issues in real-time.
- Keeping track of quality-of-fit metrics and error bars for final outputs is crucial for real-time trading applications.

Volatility fitting examples

- Examples of fitting American-style stock, ETF and futures options, as well as European-style index options.
- Starting with options and underlier prices for stocks/ETFs, we:
  1. pick interest rate (we like "SPIBOR")
  2. pick cash divs (if appropriate)
  3. imply borrow cost for each term to get "American PCP"
  4. imply vol-by-strike
  5. fit all terms to various vol curves
- Advanced topics:
  - Smart temporal filtering, to fill in missing data or discount "bad" snapshots or fits.
  - No-arbitrage fitting (see later).
- Simple metrics shown below: chi = reduced chi-square, avE5.

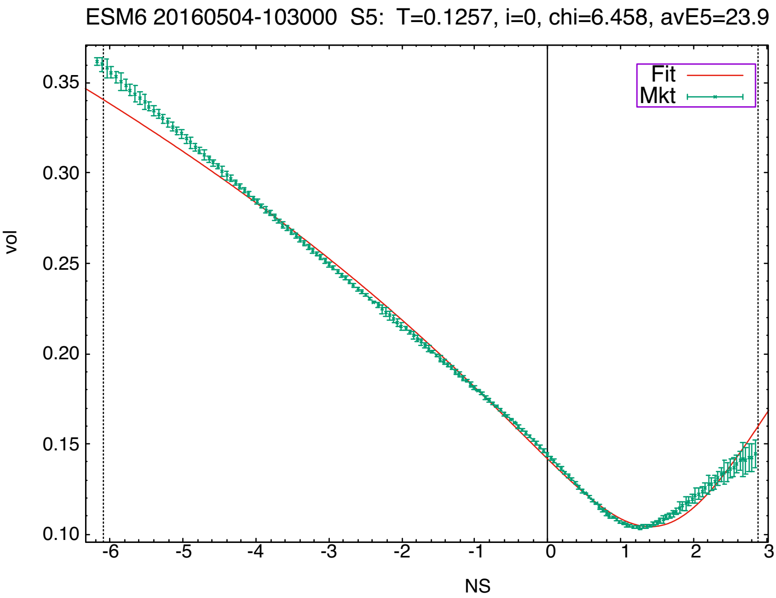
ESM6 20160504-103000 S5: T=0.1257, i=0, chi=6.458, avE5=23.9
Fit
Mkt
vol
0.35
0.30
0.25
0.20
0.15
0.10
-6
-5
-4
-3
-2
-1
0
1
2
3
NS

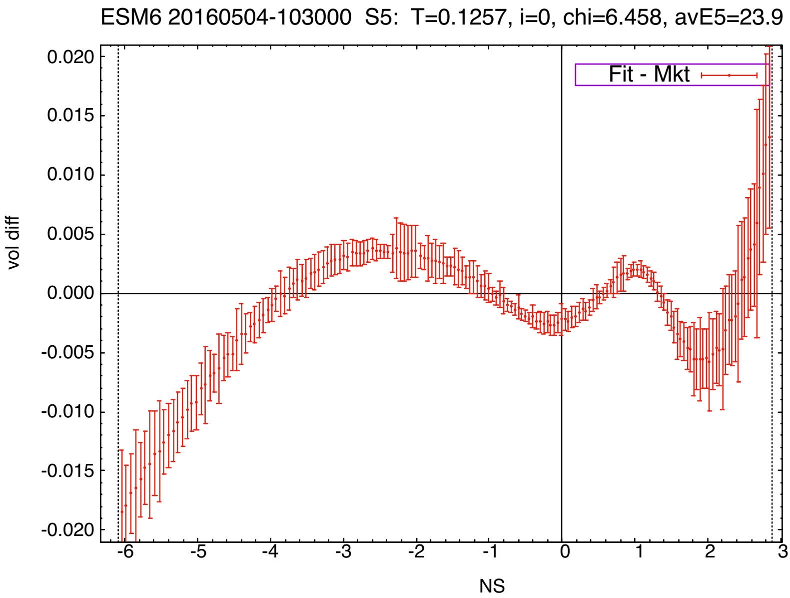
ESM6 20160504-103000 S5: T=0.1257, i=0, chi=6.458, avE5=23.9
Fit - Mkt
vol diff
NS
0.020
0.015
0.010
0.005
0.000
-0.005
-0.010
-0.015
-0.020
-6
-5
-4
-3
-2
-1
0
1
2
3

ESM6 20160504-103000 C8: T=0.1257, i=0, chi=0.599, avE5=3.2

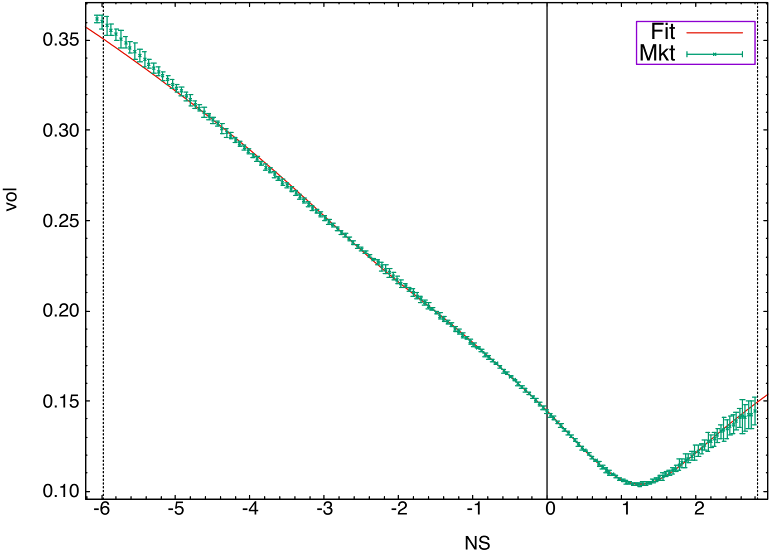

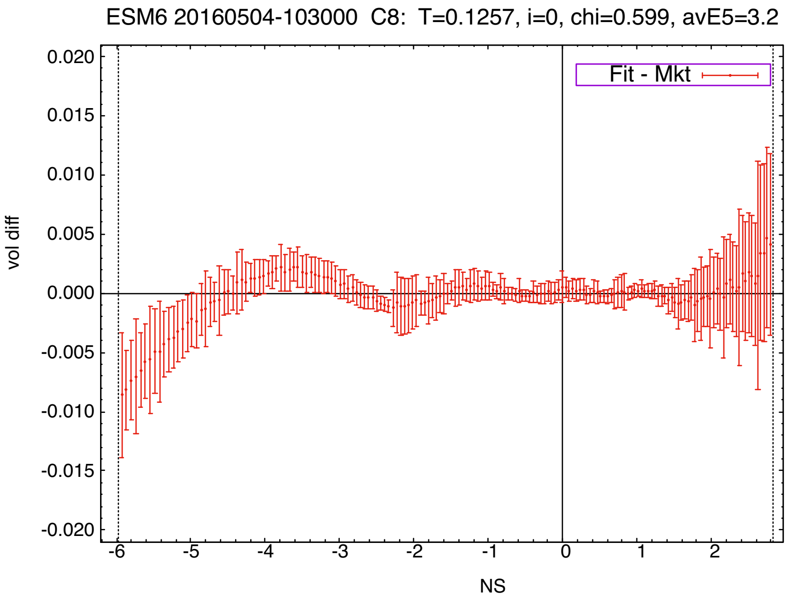
ESM6 20160504-103000 C8: T=0.1257, i=0, chi=0.599, avE5=3.2
Fit - Mkt
vol diff
NS
0.020
0.015
0.010
0.005
0.000
-0.005
-0.010
-0.015
-0.020
-6
-5
-4
-3
-2
-1
0
1
2

ESM6 20160504-103000 C12m: T=0.1257, i=0, chi=0.021, avE5=1.4

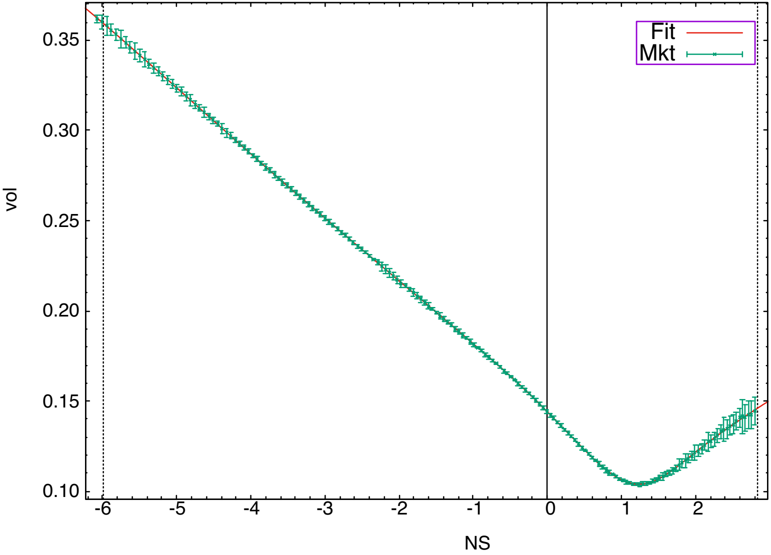

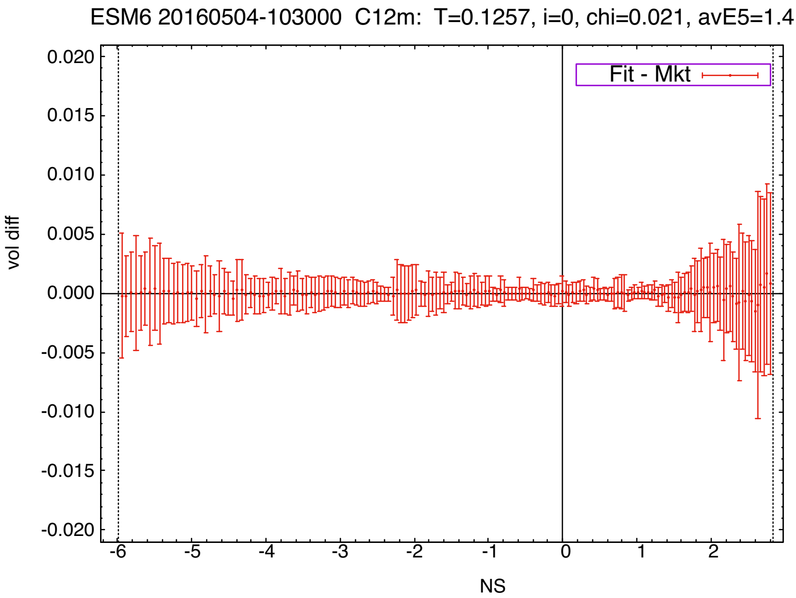
ESM6 20160504-103000 C12m: T=0.1257, i=0, chi=0.021, avE5=1.4
Fit - Mkt
vol diff
0.020
0.015
0.010
0.005
0.000
-0.005
-0.010
-0.015
-0.020
-6
-5
-4
-3
-2
-1
0
1
2
NS

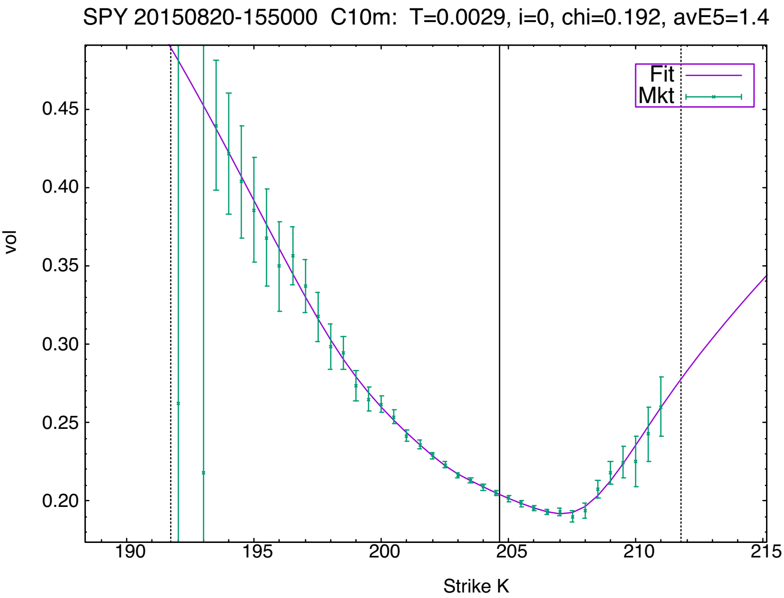
SPY 20150820-155000 C10m: T=0.0029, i=0, chi=0.192, avE5=1.4
Fit
Mkt
vol
0.45
0.40
0.35
0.30
0.25
0.20
190
195
200
205
210
215
Strike K

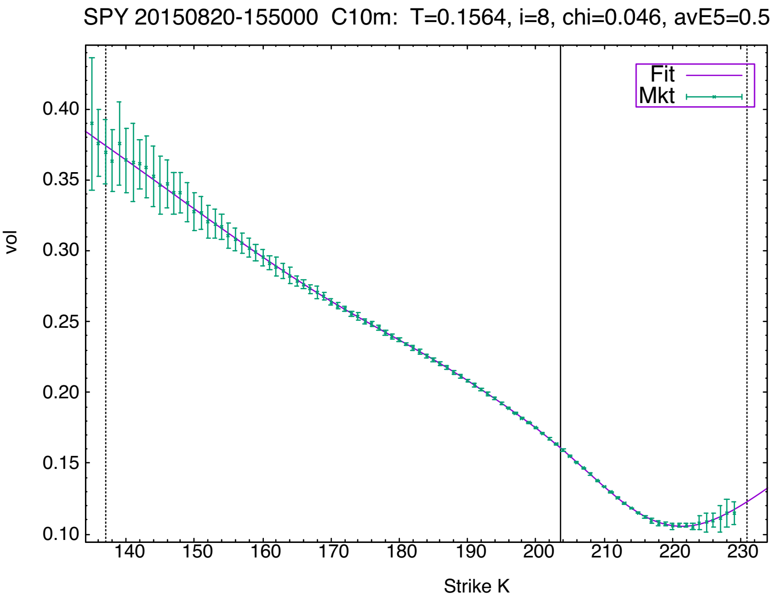
SPY 20150820-155000 C10m: T=0.1564, i=8, chi=0.046, avE5=0.5
Fit
Mkt
vol
0.40
0.35
0.30
0.25
0.20
0.15
0.10
140
150
160
170
180
190
200
210
220
230
Strike K

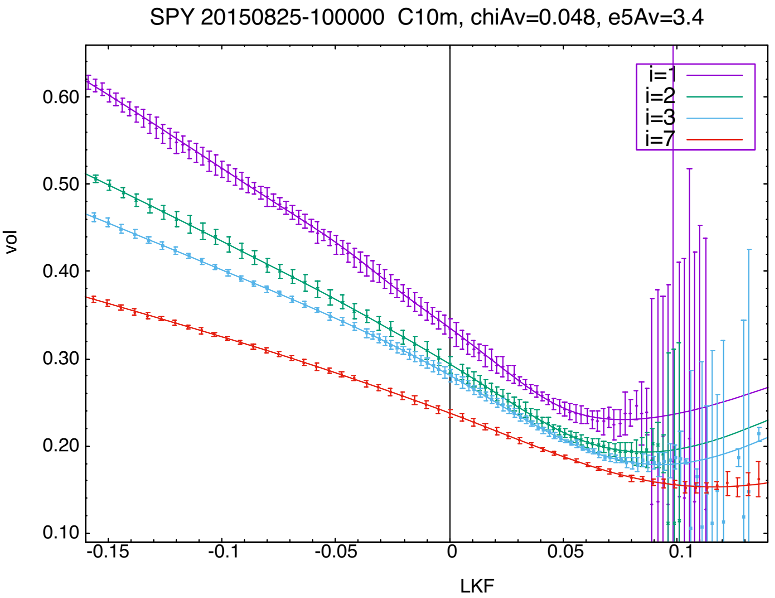
SPY 20150825-100000 C10m, chiAv=0.048, e5Av=3.4
i=1
i=2
i=3
i=7
vol
0.60
0.50
0.40
0.30
0.20
0.10
-0.15
-0.1
-0.05
0
0.05
0.1
LKF

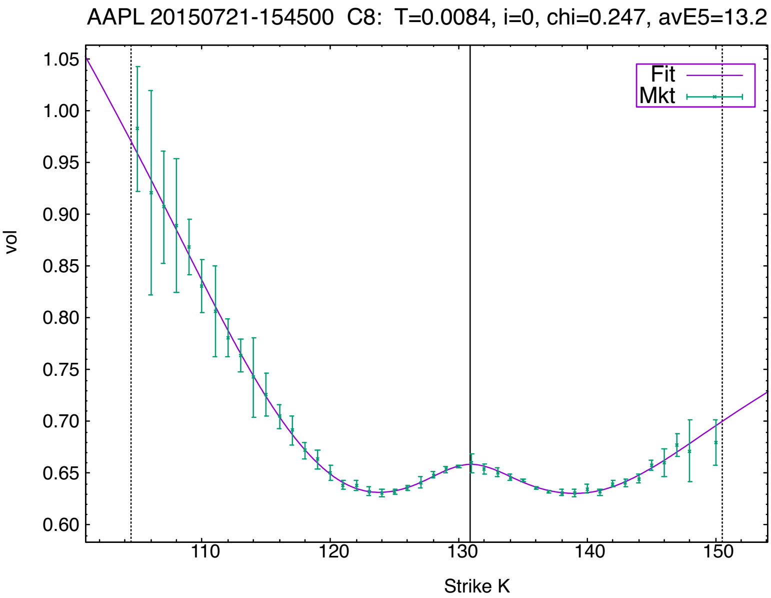
AAPL 20150721-154500 C8: T=0.0084, i=0, chi=0.247, avE5=13.2
Fit
Mkt
vol
1.05
1.00
0.95
0.90
0.85
0.80
0.75
0.70
0.65
0.60
110
120
130
140
150
Strike K

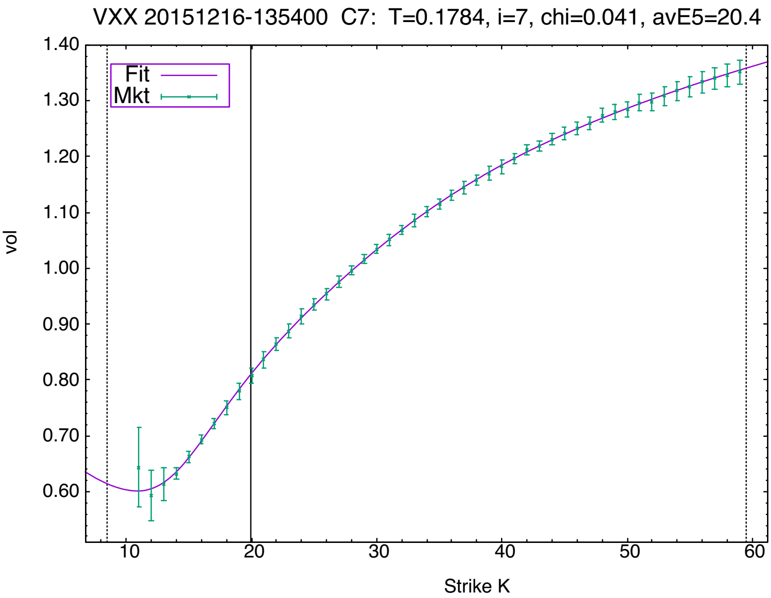
VXX 20151216-135400 C7: T=0.1784, i=7, chi=0.041, avE5=20.4
Fit
Mkt
vol
1.40
1.30
1.20
1.10
1.00
0.90
0.80
0.70
0.60
10
20
30
40
50
60
Strike K

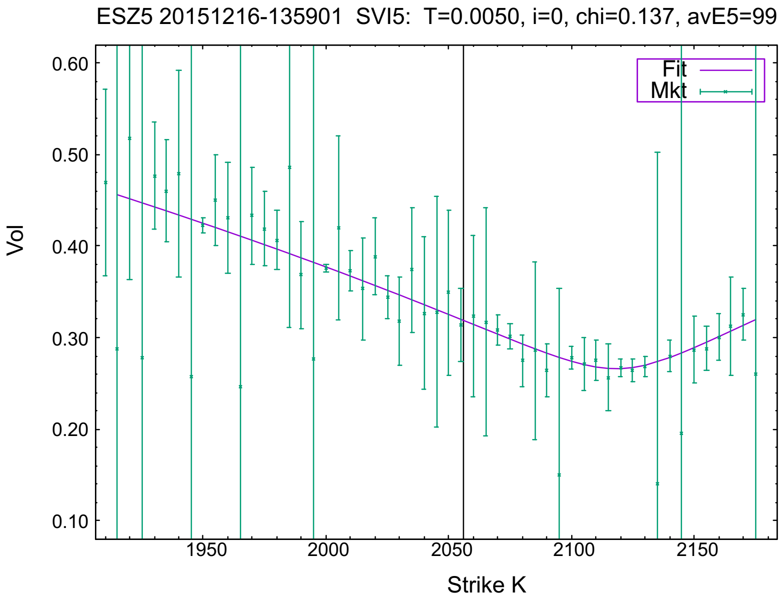
ESZ5 20151216-135901 SVI5: T=0.0050, i=0, chi=0.137, avE5=99
Fit
Mkt
Vol
0.60
0.50
0.40
0.30
0.20
0.10
1950
2000
2050
2100
2150
Strike K

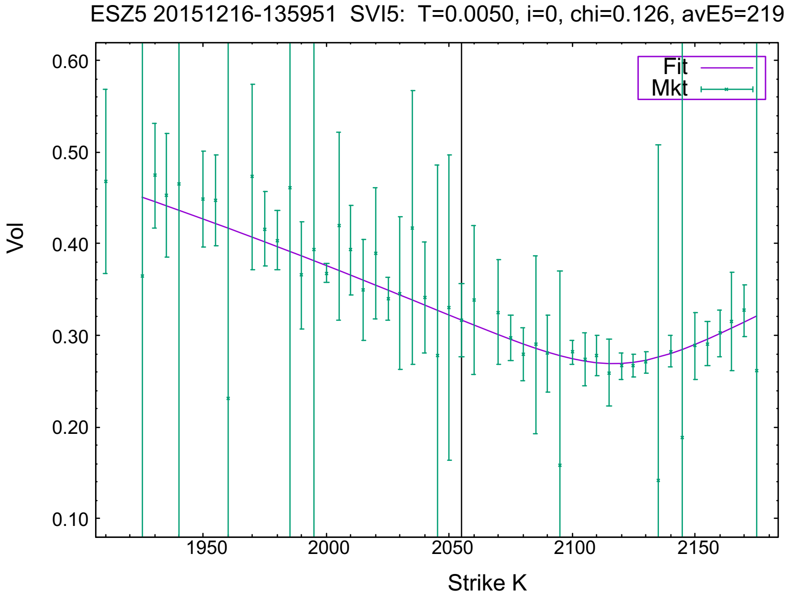
ESZ5 20151216-135951 SVI5: T=0.0050, i=0, chi=0.126, avE5=219
Fit
Mkt
Vol
0.60
0.50
0.40
0.30
0.20
0.10
1950
2000
2050
2100
2150
Strike K

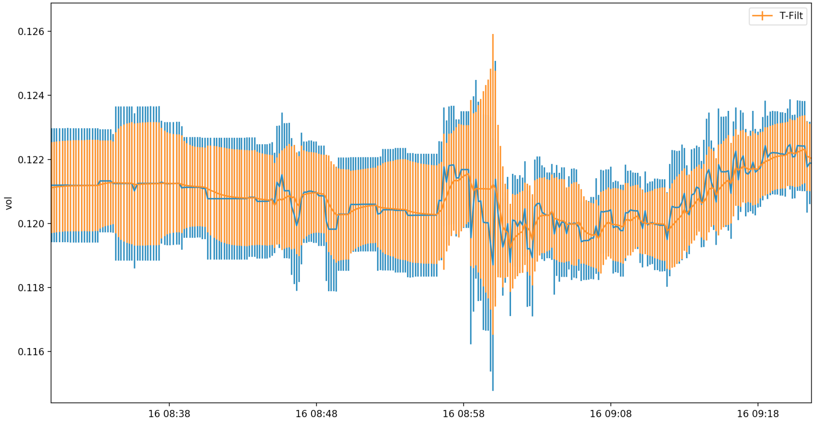

Time series of the ATF vol of the front month, with error bars, from snapshot fits (blue) and after temporal filtering (orange) for the NKY/N225 around the open (9:00am) on 20170516.

Arbitrage Elimination in Practice

- Our volatility curve framework makes it (almost) impossible to have butterfly arbitrage in the fitted curves.
- Complete calendar arbitrage elimination is accomplished by algorithmic means ("no-arb mode") considering:
  - With a proper shape parametrization, the term-structure of the (shape) parameters should be "pretty smooth" (except maybe due to events).
  - Enforce total variance constraint taking error bars into account to spread shape information across terms.
- Many potential arbitrages in a vol surface fit of a snapshot can already be eliminated via smart temporal filtering.
- Final result: arbitrage-free vol surface "closest" to the vol surface that fits market prices best, in parametric form.
- Some examples, with "regular" and "no-arb" fit mode.

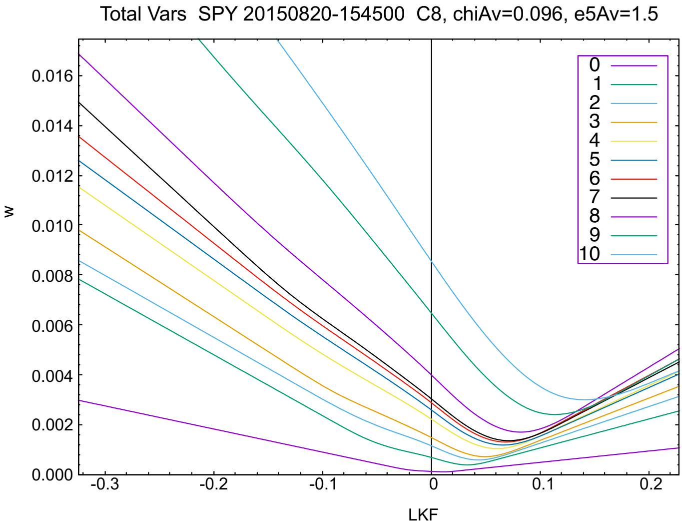

Regular fit mode

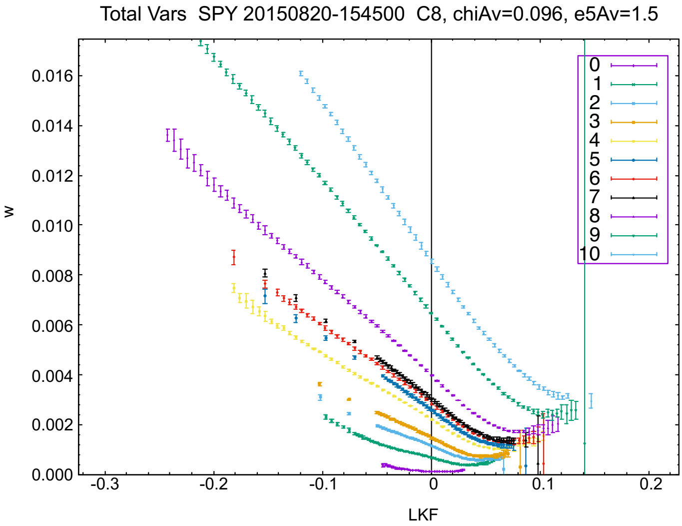

Regular fit mode

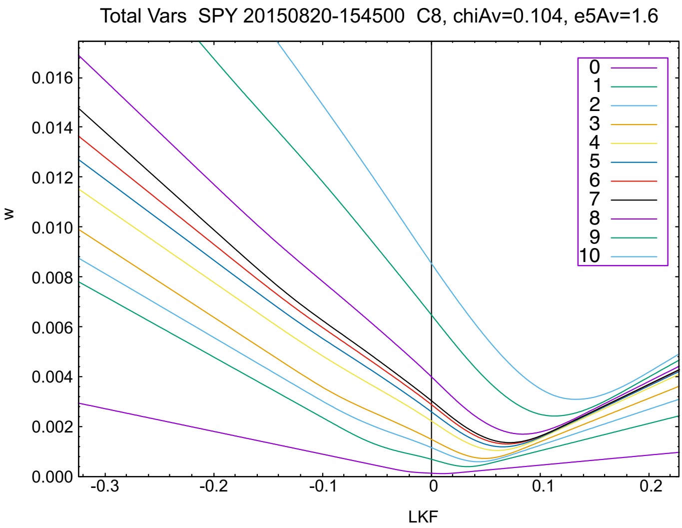

No-arb fit mode

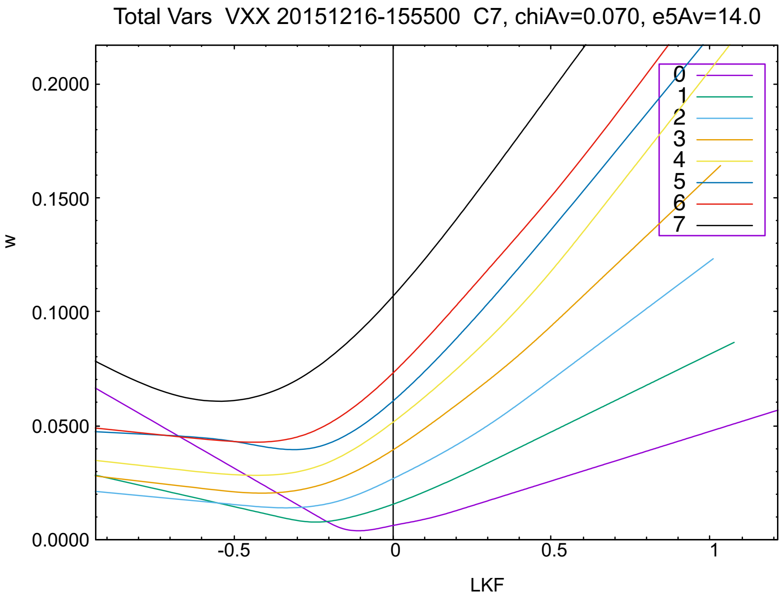
Total Vars VXX 20151216-155500 C7, chiAv=0.070, e5Av=14.0
w
LKF
0.2000
0.1500
0.1000
0.0500
0.0000
-0.5
0
0.5
1
0
1
2
3
4
5
6
7

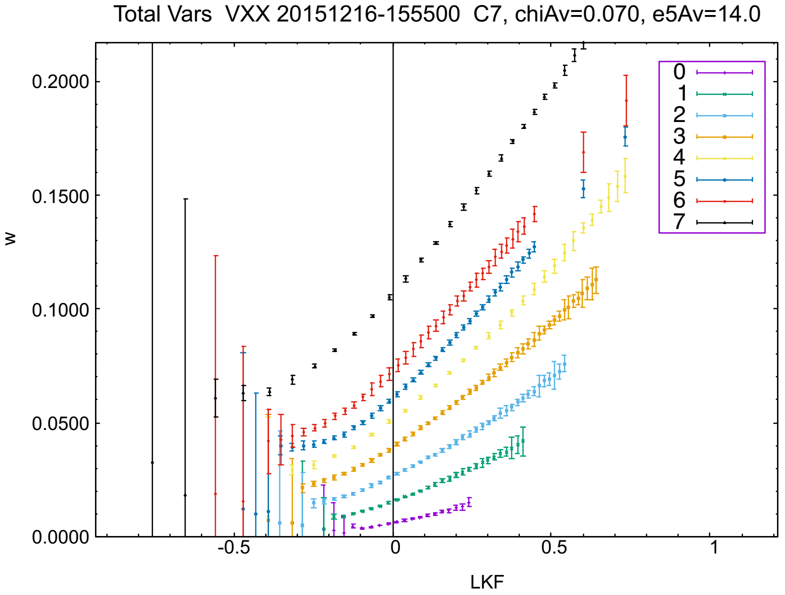
Total Vars VXX 20151216-155500 C7, chiAv=0.070, e5Av=14.0
w
0.2000
0.1500
0.1000
0.0500
0.0000
-0.5
0
0.5
1
LKF
0
1
2
3
4
5
6
7

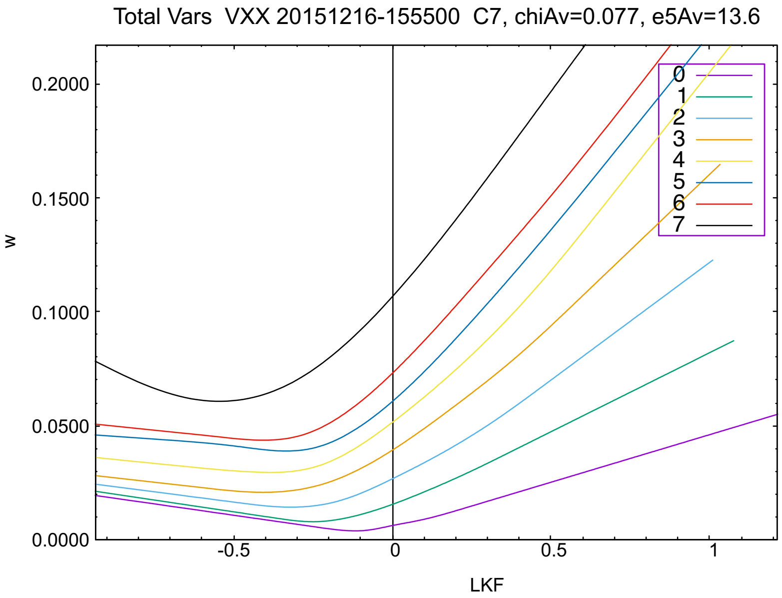
Total Vars VXX 20151216-155500 C7, chiAv=0.077, e5Av=13.6
w
0.2000
0.1500
0.1000
0.0500
0.0000
-0.5
0
0.5
1
LKF
0
1
2
3
4
5
6
7

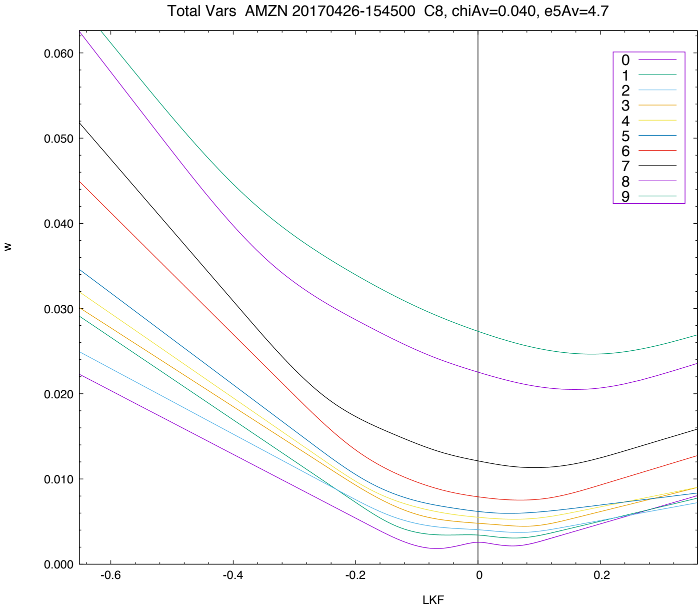
Total Vars AMZN 20170426-154500 C8, chiAv=0.040, e5Av=4.7
w
0.060
0.050
0.040
0.030
0.020
0.010
0.000
-0.6
-0.4
-0.2
0
0.2
LKF
0
1
2
3
4
5
6
7
8
9

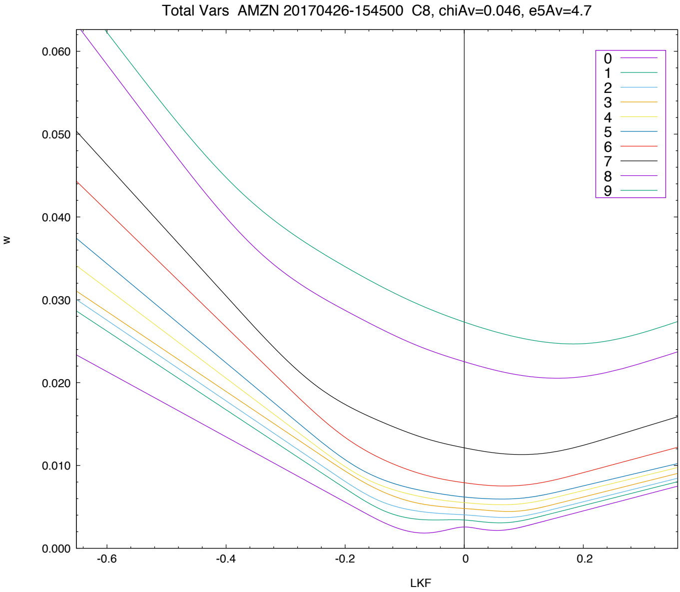
Total Vars AMZN 20170426-154500 C8, chiAv=0.046, e5Av=4.7
0.060
0.050
0.040
0.030
0.020
0.010
0.000
w
-0.6
-0.4
-0.2
0
0.2
LKF
0
1
2
3
4
5
6
7
8
9

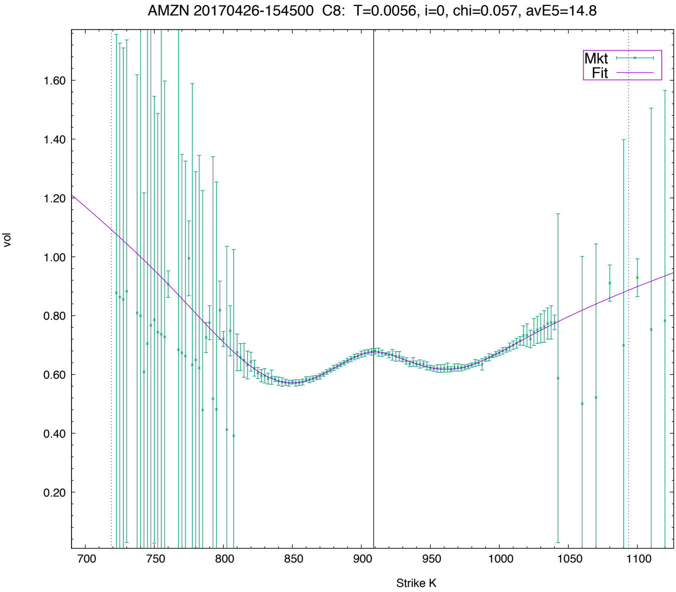
AMZN 20170426-154500 C8: T=0.0056, i=0, chi=0.057, avE5=14.8
Mkt
Fit
vol
Strike K
1.60
1.40
1.20
1.00
0.80
0.60
0.40
0.20
700
750
800
850
900
950
1000
1050
1100

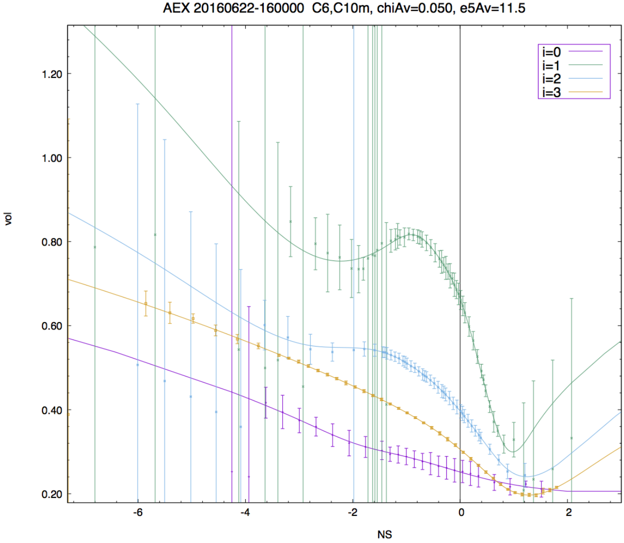

AEX vols on the day of the Brexit vote!

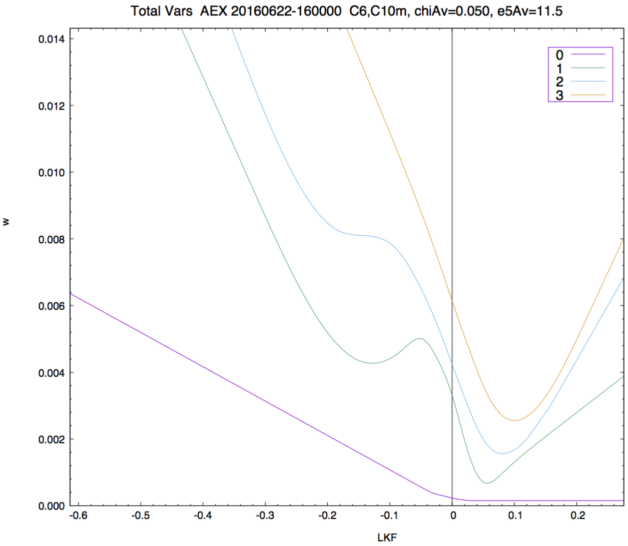

AEX total variances

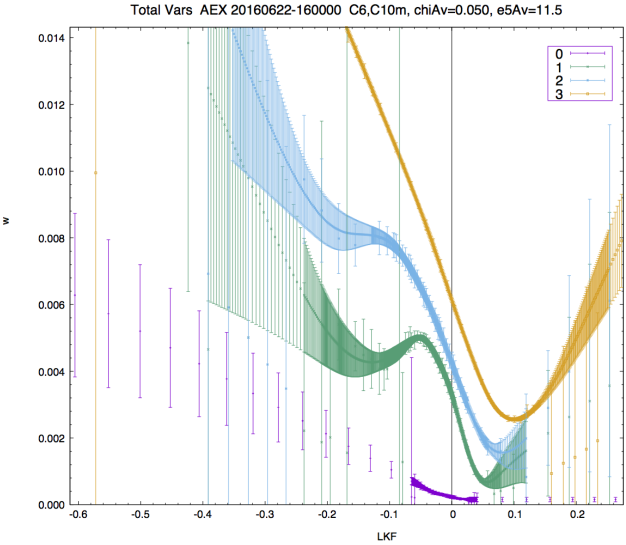

AEX total variances with input and output error bars

Fitting Options on Illiquid Names

- The vast majority (95%+) of the 4000+ underliers with options in the US are relatively illiquid.
- Mids and even "naive" theos (from some curve fit...) will often have arbitrage (at least when spreads are ignored).
- As far as vol curve fitting is concerned, they can fortunately usually be fitted with the simple 3-parameter curve "S3" (aka "SSVI").
- No-arbitrage constraints and careful consideration of error bars are extremely helpful in spreading information in the limited set of more liquid options across all options and to a complete surface.
- Let's consider the example of Sprint, ticker S, with closing data. We will show fits in "regular" and "no-arb" mode again.

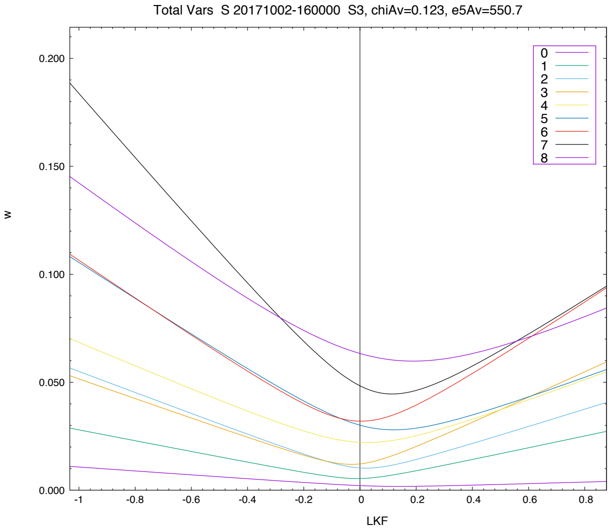

Total variances in regular fit mode

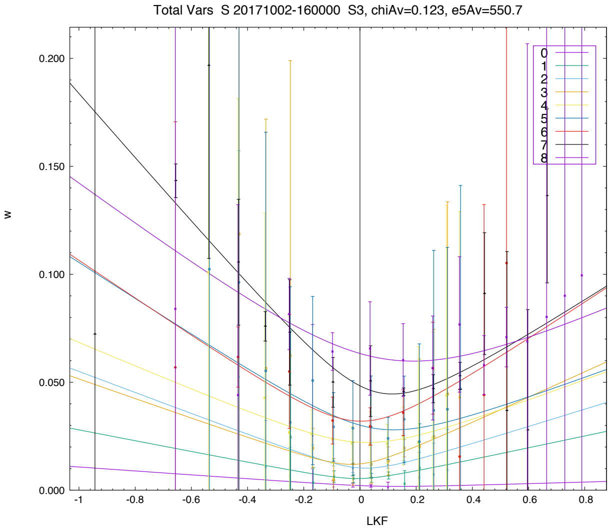

Total variances with error bars in regular fit mode

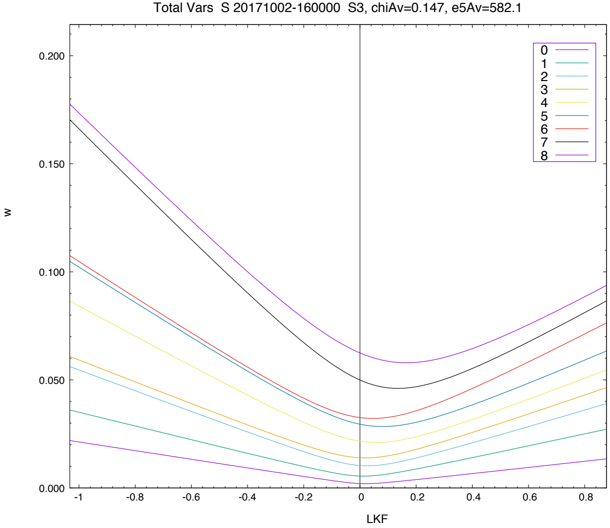

Total variances in no-arb fit mode

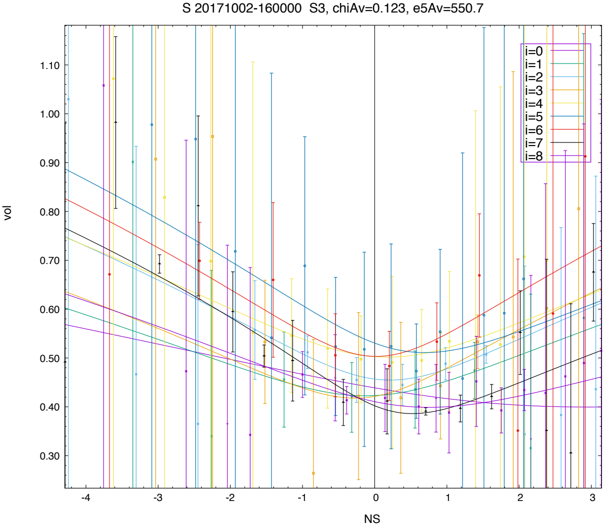

Vols in regular fit mode

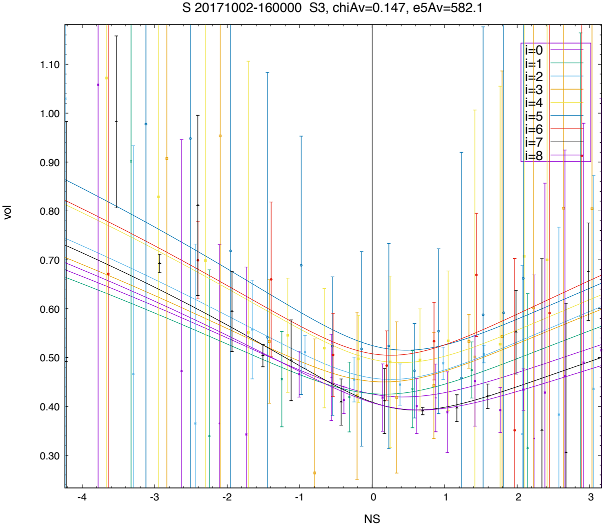

Vols in no-arb fit mode

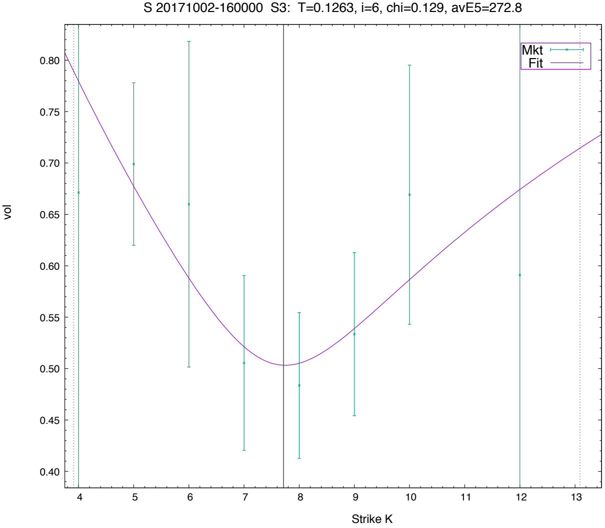

S3 fit of S (Sprint) options, $T = 0.1263$y, $i = 6$, regular fit mode

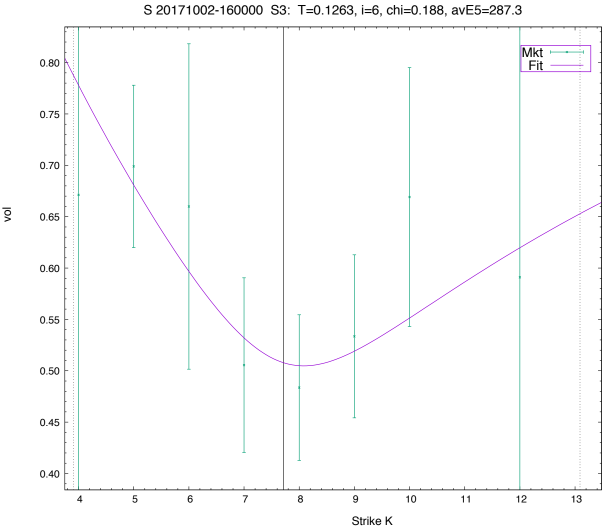

S3 fit of S (Sprint) options, $T = 0.1263$y, $i = 6$, no-arb fit mode

Other important topics we do not have time for

- Speed: Price and fit the whole US options universe on one box.
- Spot vol dynamics: How do vol surfaces move when the underlier moves? (Simple rules like "sticky-strike" or "sticky-delta" vols are not realistic.) This is relevant for:
  - Realistic parametric scenario generation; overnight moves
  - Smart delta
  - Smart temporal filtering
- How do vol surfaces behave under shifts of "as-of-time" (events...)?
- How does new vol level and skew information in one term spread across the surface?
- How "local" are our parametrizations? Can our parametric volatility surfaces be simply and realistically deformed?

Summary and Conclusion

- There is no standardization in the equity options markets around dividend modeling, borrow costs, or vol curves and their calibration.
- No vol curves in the public domain can fit liquid names like ES, SPX, SPY, AAPL, AMZN, etc.
- Guaranteed arbitrage elimination while still fitting the market "as good as possible" is a hard problem even for simple curves/surfaces.
- Our fitting framework allows to efficiently eliminate arbitrage for any underlier, liquid or illiquid, with any curve types, in real time.
- Standards for div modeling, borrows & vol fitting are at hand:
  - Improved price transparency facilitates wider use of options and the efficient transfer of vol information across related products.
  - Hope to achieve same as in transition from old to new VIX: A healthier market, larger volumes, esp. from smaller players.