American Option Pricing in a Tick - Calibration in a Click

Quant Insights Conference
Online
October 2021

Jesper Andreasen
Saxo Bank, Copenhagen
kwant.daddy@saxobank.com

Outline

- Introduction: equity option markets and the task at hand.
- American option pricing: PDEs, early exercise boundary, integral representation and a few asymptotic results.
- Fix point method for American options.
- Numerical performance and calibration.
- Arbitrage considerations and implementation.
- Conclusion.

References

- Andersen, L, M Lake, and D Offengenden (2016): "High Performance American Option Pricing." *Journal of Computational Finance 20, 1, 39-87*.

- Andreasen, J (2006): "The Fastest Bermuda Pricer in the West." *Bank of America WP*.

- Andreasen, J (2020): "American Option Pricing in the Normal Model." *Saxo WP*.

- Andreasen, J and B Huge (2011): "Volatility Interpolation." *Risk March*.

- Kim, J (1990): "The Analytic Valuation of American Options." *Review of Financial Studies 3, 4, 547-572*.

- McKean, H P (1965): "A Free Boundary Problem for the Heat Equation Arising from a Problem in Mathematical Economics." *Industrial Management Review 6, 32-39.*

Intro: Task at Hand

- You can trade approximately 1,000,000 listed options on 1,500 underlying stocks on the Saxo Bank trading platform. On average 700 options per underlying.
- We wish to construct a risk system where you are able to price and risk [first order and scenarios] these options.
- To do so we need to back out the parameters of the option prices [dividends, volatilities, etc] from the quoted prices.
- The challenge here is that 90% of the equity options are American style.
- American options require pricing by a numerical method such as finite difference or binomial tree.

Two Challenges: Computer and Man Power

- The best we can do FD pricing is probably O(1ms) CPU time per pricing.
- Hence, calibration to, say 100 American options, would take O(1s) to do.
- However, I am a little nervous about the stability and robustness of such an approach on the industrial scale we're looking at here.
- Particularly, as we also need to back out dividend information.
- At Danske we maintained approx. 100 volatility surfaces with approx. 10 FTEs (traders + quants).
- Here we can't afford any manual hand holding.

- So I need some abracadabra here: speed, stability, automation.

American Option Pricing

- … is an optimal stopping problem

$$C(t) = \sup_{\tau} E_t[e^{-r(\tau-t)}(S(\tau)-K)^+] \tag{1}$$

- Under the assumption of one factor Markov diffusion

$$\frac{dS}{S} = (r-q)dt + \sigma dW \tag{2}$$

- … American call options can be priced as the solution to the PDE

$$rC = C_t + (r-q)SC_S + \frac{1}{2}\sigma^2 S^2 C_{SS} \quad , C \geq (S-K)^+ \tag{3}$$

- The free boundary condition can also be written as

$$C(t,S) = \max(\underbrace{C(t+,S)}_{rollback\,value}, S-K) \tag{4}$$

- Technically, this includes models with state dependent (local) interest rates, dividend yields and volatility:

$$r = r(t,S)\ ,q = q(t,S)\ ,\sigma = \sigma(t,S) \tag{5}$$

- Conventional attack is to use finite difference methods to directly solve (3).

- “Market convention” is to use a binomial tree with constant interest rate, dividend and volatility.

- … and then use different volatilities for different strikes and different dividend yields and interest rates for different expiries.

- Yes, I know, doesn't win any beauty or consistency awards but we shouldn't neglect the information content in this being considered "enough" for battle.

- Note, I will only be considering calls here as we can use the American put/call duality

$$C(t) = \underbrace{\sup_{\tau} E_t[e^{-r(\tau-t)}(S(\tau)-K)^+]}_{American\ call\ in\ domestic\ ccy} = KS(t)\underbrace{\sup_{\tau} \bar{E}_t[e^{-q(\tau-t)}(\frac{1}{K}-\frac{1}{S(\tau)})^+]}_{American\ put\ in\ foreign\ ccy}$$

- … to price American puts.

Early Exercise Boundary: Some Standard Results

- Let the early exercise boundary be given by

$$S^*(t) = \inf_S \{S : C(t,S) = S - K\} \qquad (6)$$

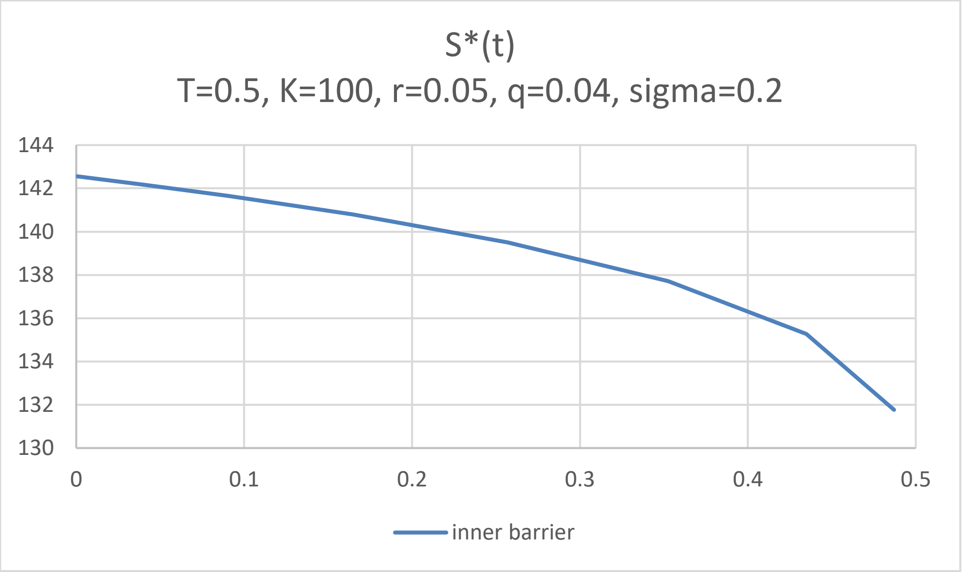

- Smooth pasting:

$$C_S(t,S^*(t))=1\quad ,C_t(t,S^*(t))=0\quad ,C_{SS}(t,S^*(t)-)=\frac{qS^*(t)-rK}{\frac{1}{2}\sigma^2 S^*(t)^2} \tag{7}$$

- Discontinuity at expiry:

$$\lim_{t\to T} S^*(t)=\max(1,\frac{r}{q})K\quad ,q>0 \tag{8}$$

- Interesting two boundary case for $r<q<0$:

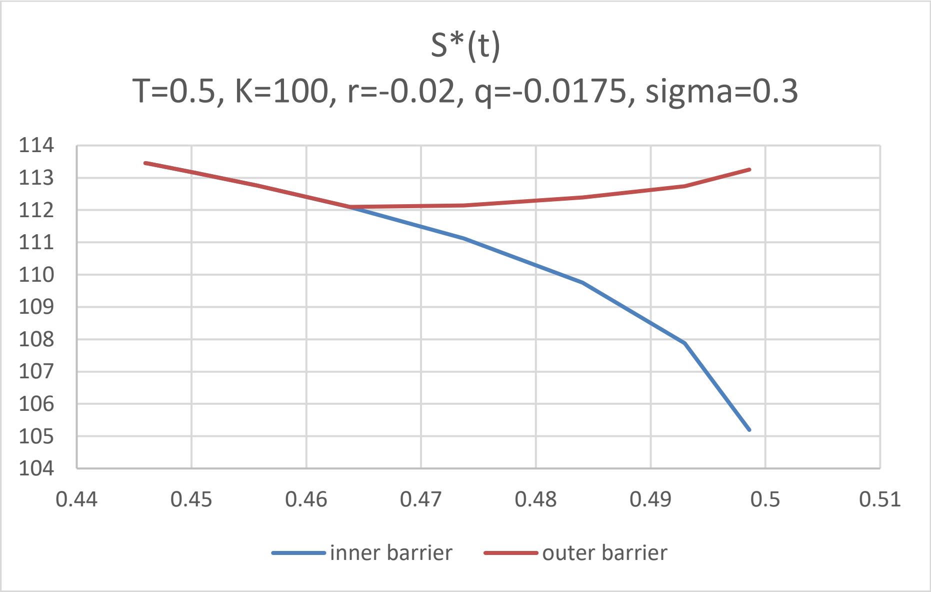

- Square root behavior of the early exercise boundary towards expiry

$$\ln S^*(t) = O(1) + O((T-t)^{1/2}) \qquad (9)$$

- Surprisingly, some asymptotic results can be proven by use of the put call parity and geometric arguments.

- For example, for $r > q > 0$:

$$S^*(t) - K = C(t, S^*(t)) \geq c(t, S^*(t)) \geq S^*(t)e^{-q(T-t)} - Ke^{-r(T-t)}$$
$$\Rightarrow \lim_{t \to T} S^*(t) \geq \frac{r}{q} K \qquad (10)$$

Integral Representation

- The American call can be written as

$$C(t,S(t)) = \underbrace{c(t,S(t))}_{\substack{european \\ call}} + \underbrace{\int_t^T e^{-r(u-t)} E_t[(qS(u) - rK)1_{S(u) \geq S^*(u)}]du}_{early\ exercise\ premium} \qquad (11)$$

- If the early exercise boundary $\{S^*(t)\}_{t \leq T}$ is known and parameters are at most time dependent then the integrand in (11) can be written in closed form.

- Hence, if S^* is known then we can use numerical integration to price the American option.

- However, (10) can be solved to identify the boundary S^* and then subsequently price the American option.

Recursive Integral Equation

- We can discretize (11) on a discrete time grid $0 = t_0 < t_1 < \ldots < t_{n-1} < t_n = T$ and then solve

$$S^*(t_i) - K = C(t_i, S^*(t_i); \{S^*(t_j)\}_{j>i}) \tag{12}$$

- … recursively backwards starting at t_{n-1}. Bootstrap type of solution.
- This has been known, though not widely spread, since Kim (1990).
- Personally, I have implemented American option pricing along these lines as a student back in the early 90s.
- And later used it for fast pricing of Bermuda swaptions at BofA in 2005-08.

- Using this we can price O(10,000) American options per second.
- That’s an order better than finite difference and binomial tree pricing of O(1,000) options per second.
- Can we do better?

Fix Point Integral Equation

- Equation (12) can also be seen as a curve or vector equation -- in time:

$$S^{*}=K+C(S^{*}) \quad ,S^{*}=\{S^{*}(t)\}_{0\le t<T} \tag{13}$$

- This suggest to use (13) as a fix point generator

$$S^{(n+1)}=K+C(S^{(n)}) \tag{14}$$

- Alternatively, we can derive a fix point generator from the smooth pasting condition $C_{S}(S^{*})=1$.

- Andersen, Lake and Offengenden (2016) combine this idea with careful interpolation and numerical integration techniques to achieve a very high performance algorithm.

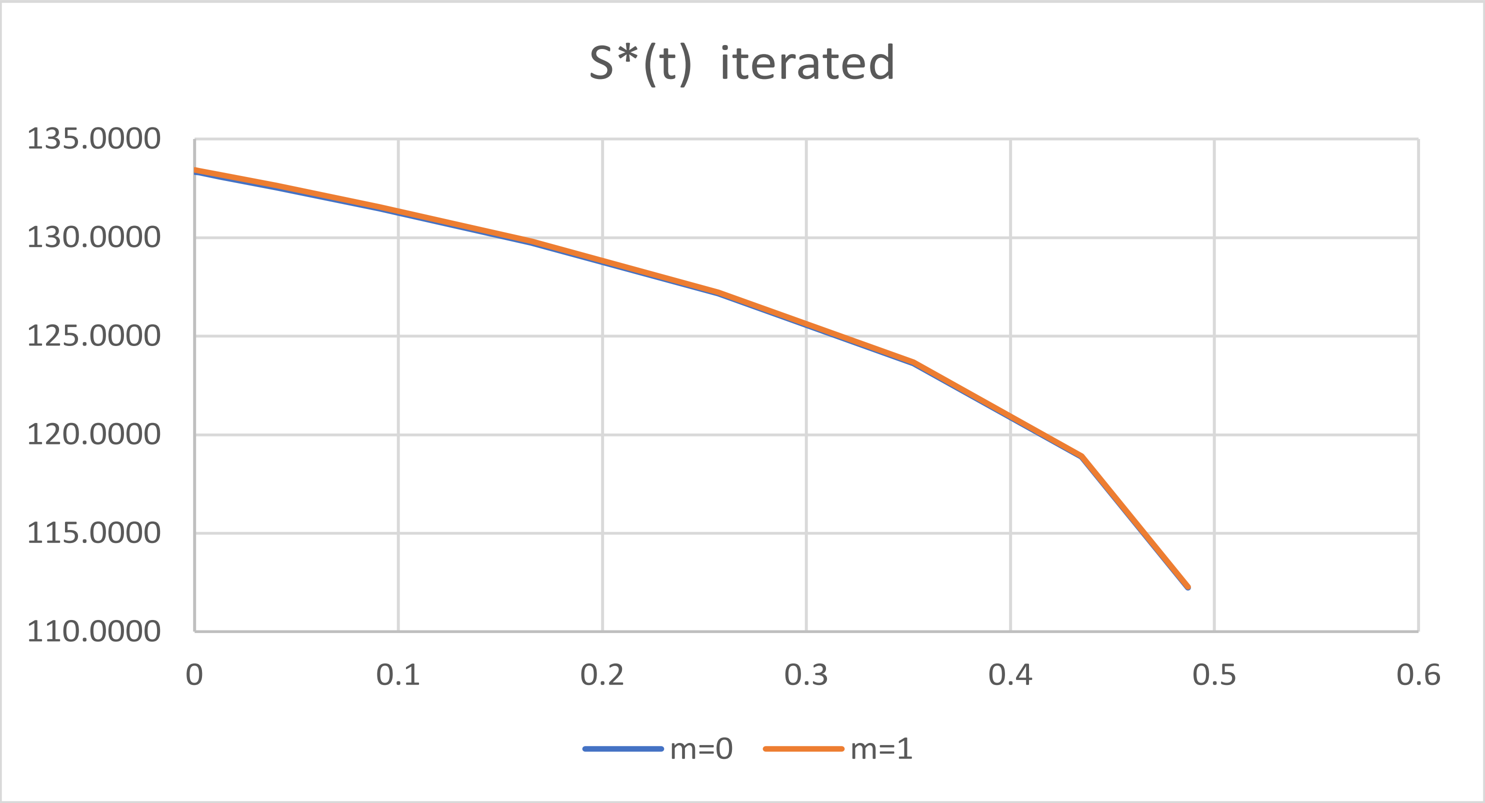

Andersen & Co American Option Pricing Algorithm

- … combines a number of tricks:
  - Fix point integral equation on $C_S = 1$ or $C = S - K$ iterated m times.
  - QD+ approximation (see Li (2009)) as starting point.
  - Interpolation of the early exercise boundary using Chebyshev polynomial with n spanning points, respecting known limits.
  - Numerical integration with l quadrature points. Different points are used for each integral.
- Normally, $l = 5, m = 1, n = 4$ gives an accuracy that is better than good enough.

- At this resolution you will be able to price O(100,000) American options per second.

- Roughly 25-50 Black-Scholes type calls per American option.

- Which is approx. 100 times faster than conventional methods.

- In fact, we can simultaneously calibrate to dividends and implied volatilities of approx. 200 American calls and puts in O(0.01s).

Why Does it Work so Well?

- If parameters are constant then the American option price is infinitely smooth in all dimensions.

- In this case, the higher order methods: Chebyshev interpolation and quadrature integration really come to their right.

- At the same time they have been very careful with selecting a fix point algorithm that has good stability properties.

- Likewise with the particular choice of interpolation.

Remarks

- The method's accuracy does not deteriorate for strikes far OTM, because the iterative bit is done around the early exercise boundary.
- The use of high order methods, somewhat, goes against traditional kwant instincts which suggest to be concerned about their lack of robustness.
- For piecewise constant parameters (say), we would need to split the integral equation in steps.
- This could potentially reduce performance.
- The algorithm requires a closed-form for the value of the early exercise premium.

- It is possible that the algorithm could work with local volatility if approximations were used. But this is untested.
- The algorithm is not particularly easy to code.
- The two-boundary special case is not completely trivial and does require a bit of TLC.
- … and this can’t generally be ignored because interest rates are negative in large parts of the world.
- The algorithm actually allows vectorising (AVX/GPU style) and this can be used to further speed it.
- My implementation is AD’ed and multi-threaded but I haven’t looked at vectorising.

Absence of Arbitrage

- In this framework there is an implicit assumption that European options can be priced at the same implied volatilities as the American options.

- If this is the case, or assumption, then we want to preclude arbitrage in the European option prices.

- I.e. for undiscounted option prices with strike given in forward space

$$g(T,X)=E[(\frac{S(T)}{F(T)}-X)^{+}] \quad ,F(T)=E[S(T)] \tag{15}$$

- … we need two conditions to be satisfied:

  - Positive maturity spreads: $g_T(T,X)>0$

- Positive butterfly spreads: $g_{XX}(T,X)>0$

- The mother of all arbitrage free option prices is the implicit finite difference scheme

$$[1-\Delta T\frac{1}{2}\vartheta(T,X)^2 X^2\delta_{XX}]g(T+\Delta T)=g(T) \quad ,g(0)=(1-X)^+ \tag{16}$$

- ... as described in Andreasen and Huge (2011).

Implementation

- Use Andersen & co (2015) algorithm to jointly estimate the dividend yield (and interest rate) for each expiry and implied volatility for each strike from American put and call mid prices.

- Weed out strikes with negative butterfly spreads.

- Fit the A&H volatility interpolation (implicit finite difference grid) to European option prices on remaining strikes.

Example Tesla 1m Smile

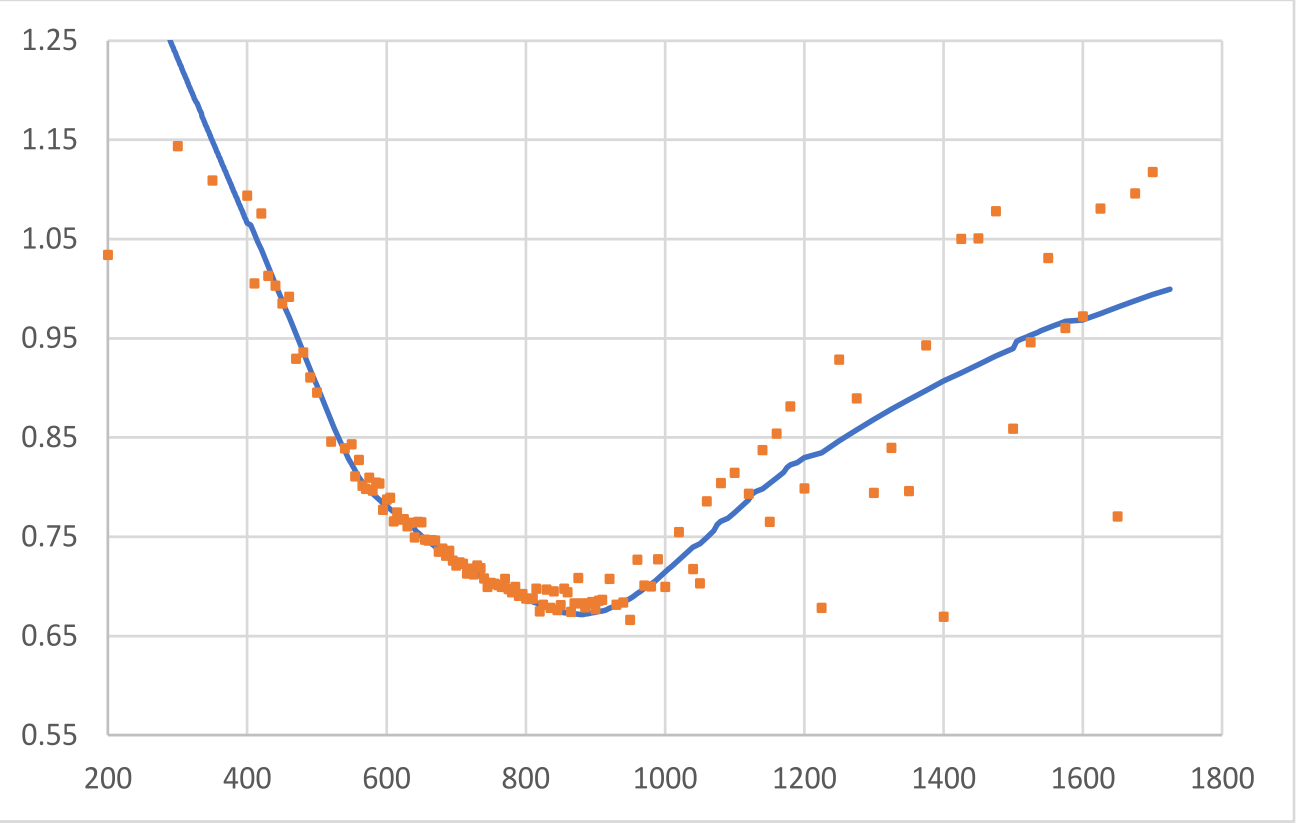

Example Tesla 6m Smile

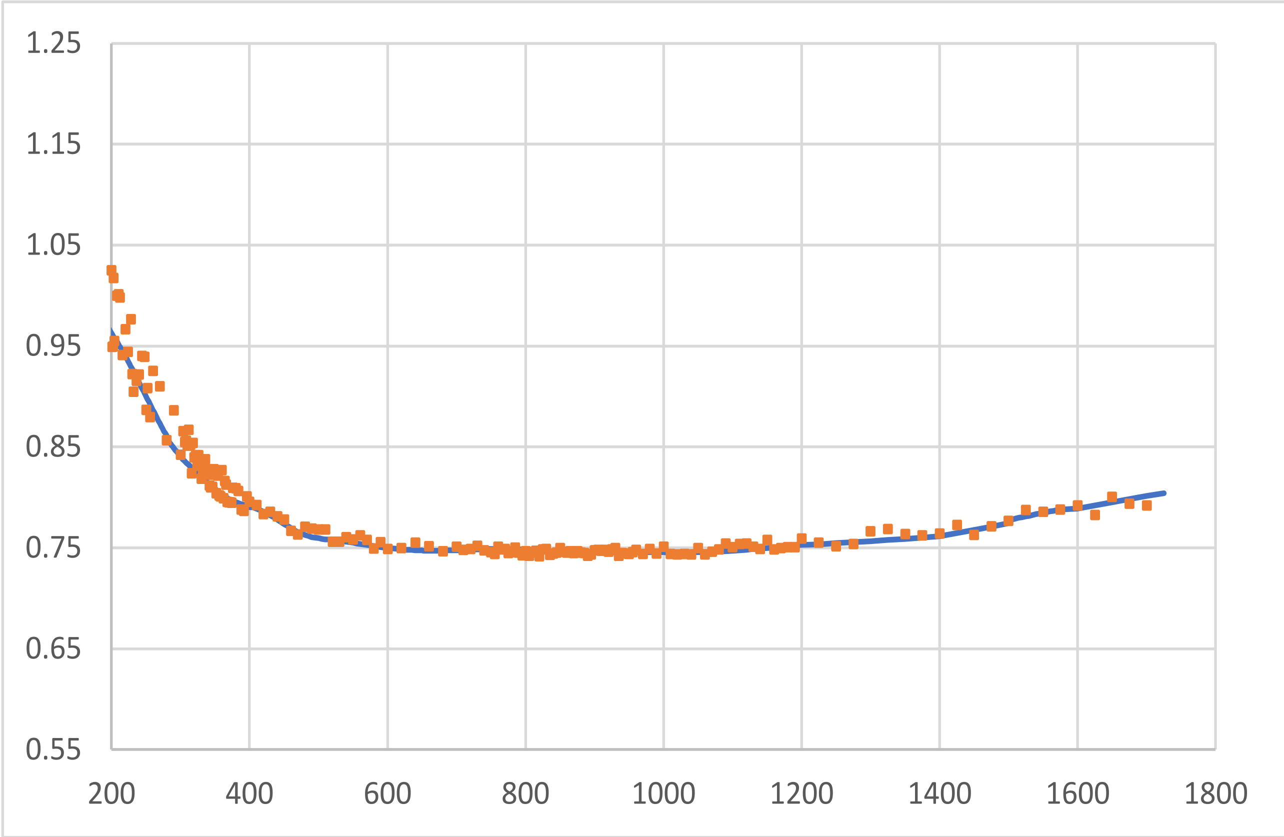

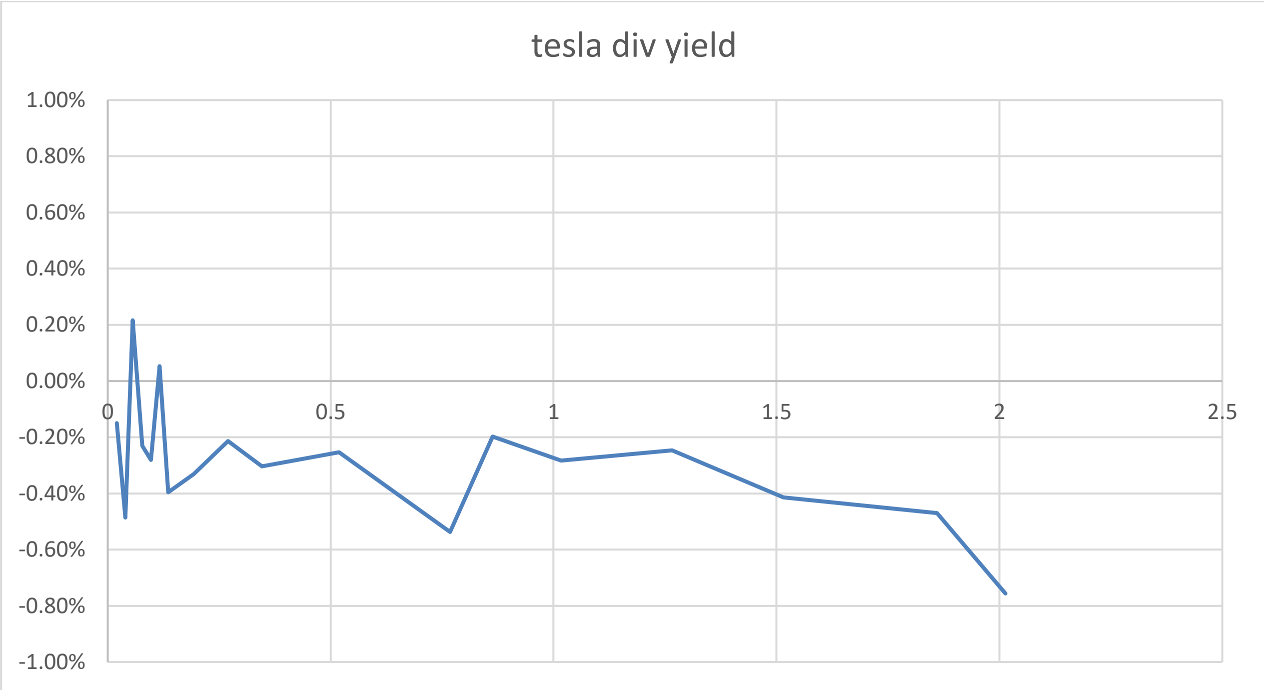
tesla div yield
1.00%
0.80%
0.60%
0.40%
0.20%
0.00%
-0.20%
-0.40%
-0.60%
-0.80%
-1.00%
0
0.5
1
1.5
2
2.5

Example: 3M 1m Smile

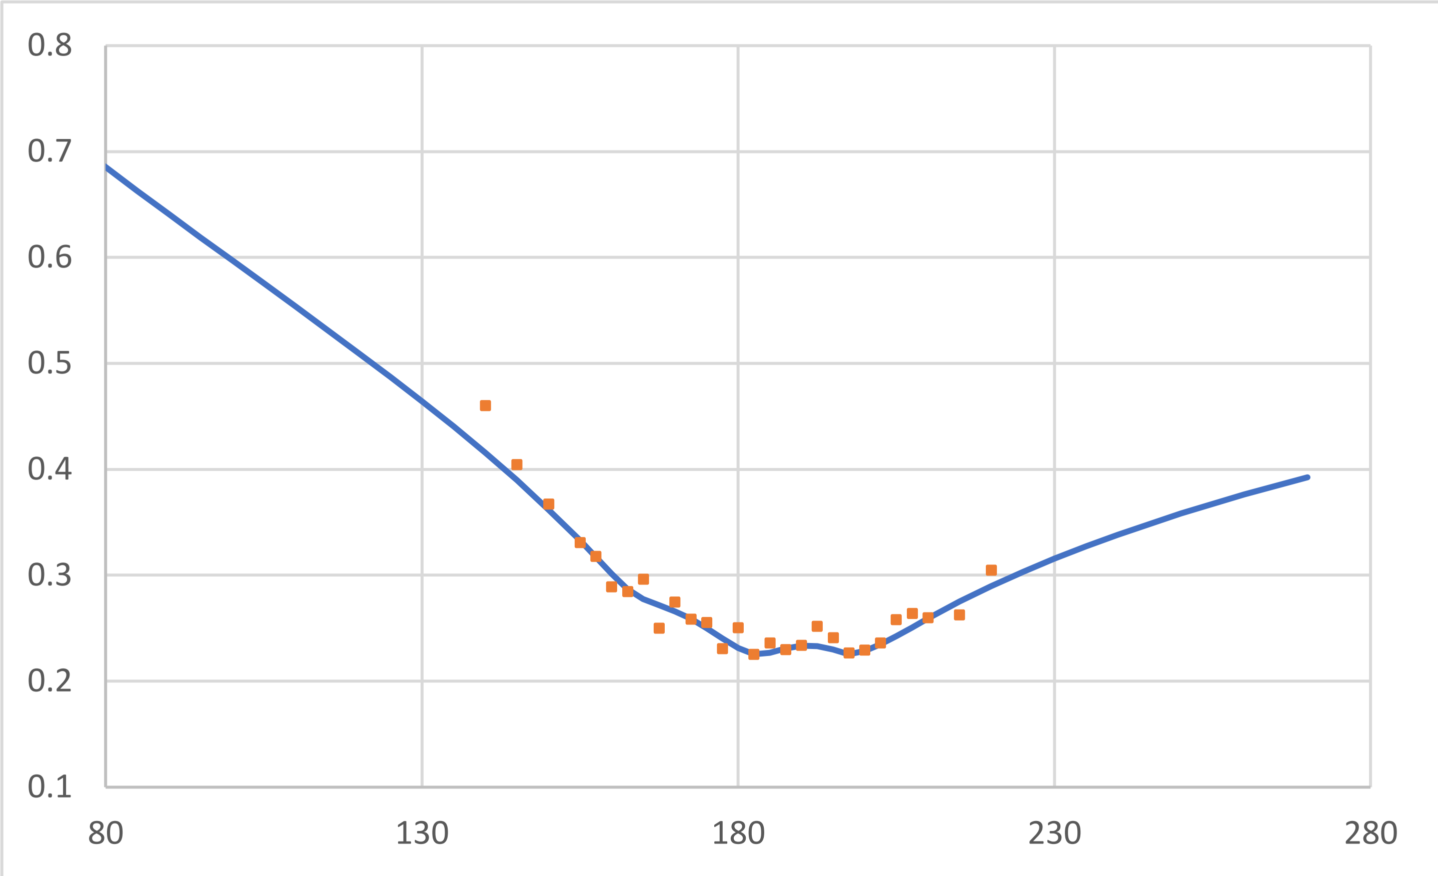

Example: 3M 6m Smile

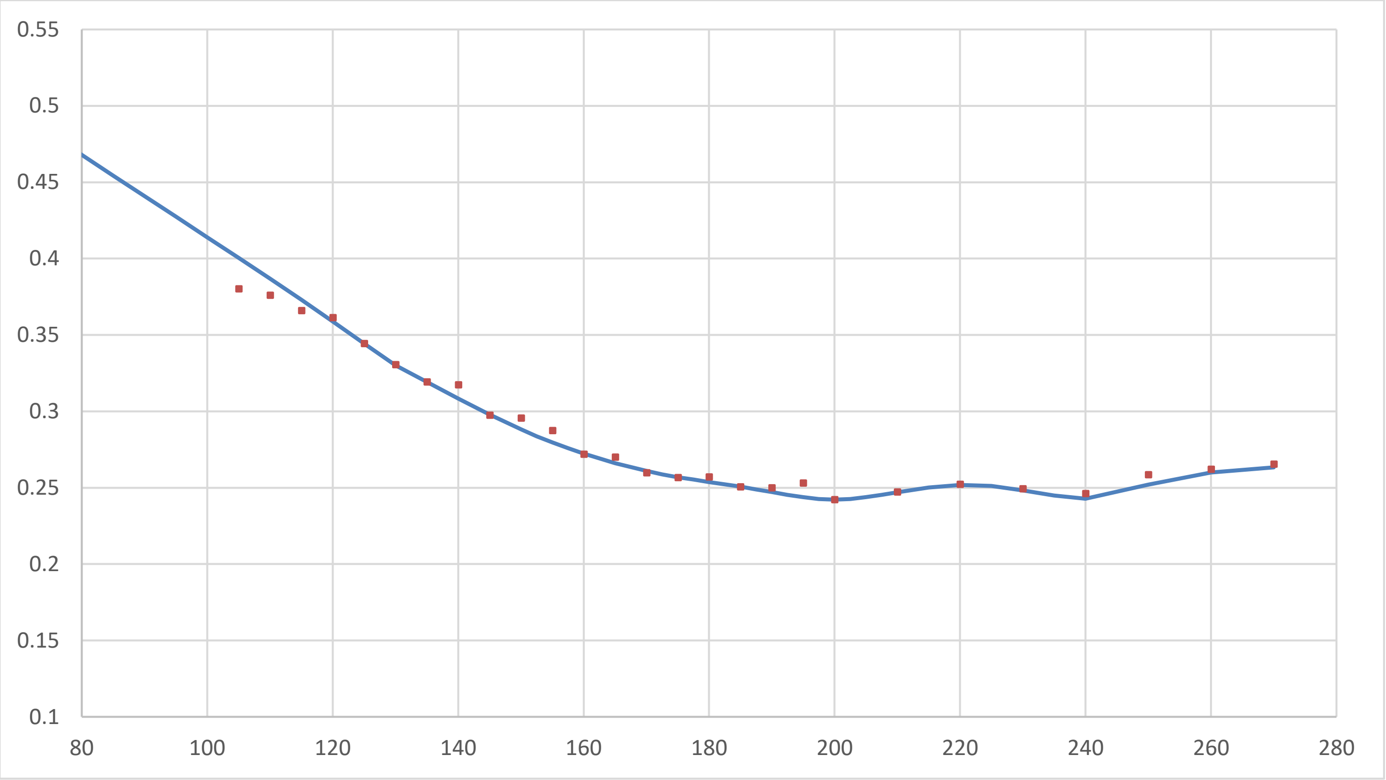

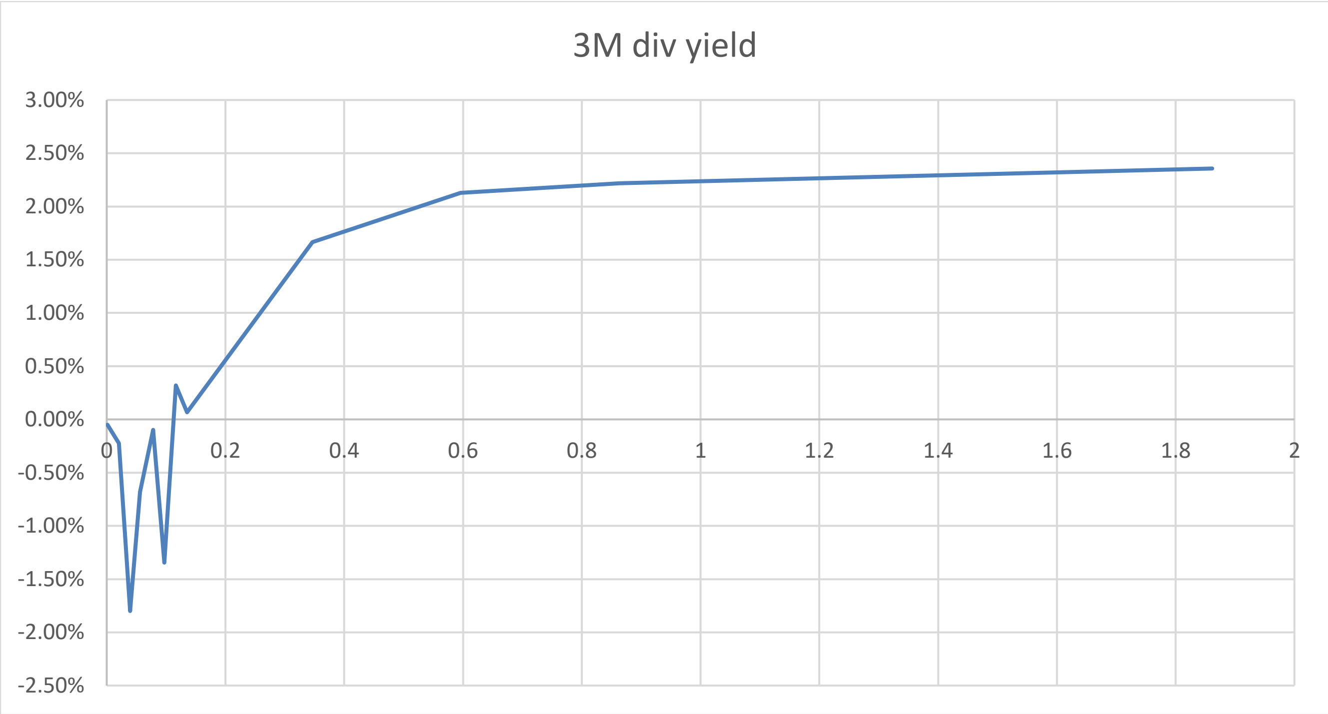
3M div yield
3.00%
2.50%
2.00%
1.50%
1.00%
0.50%
0.00%
-0.50%
-1.00%
-1.50%
-2.00%
-2.50%
0
0.2
0.4
0.6
0.8
1
1.2
1.4
1.6
1.8
2

Practical Experiences

- At the moment the machinery runs live at Saxo on 1,500 underlying stocks and a total of 1,000,000 options prices.

- All these option volatility surfaces and dividend curves are updated approximately every 5-10 minutes.

- For this we use only one server.

- Very ESG.

- The methodology is robust and we have very few errors.

Applications: Risk Modeling

- The arbitrage free volatility surfaces are used for stochastic local volatility models.

- For all options where we have positions we solve for these option prices in 3D ($t \times S \times z$) backward finite difference grids.

- The finite difference grids are kept in memory and we can now do über fast margin calculations – by simply looking up values in stored matrices of prices.

- This applies to all sorts of market and counterparty credit risk calculations.

- Value-at-risk on 10,000 options in 1ms.

Other Applications

- Automated option market making.
- Data and risk models for our clients.

Conclusion

- This is not the last word on American option pricing.
- There are significant and complicated issues around dividends, local and stochastic volatility, that are, somewhat, swept under the carpet here.
- However, under the ruling "market convention" it does allow you to calibrate quicker than the others can price.
- It would be interesting to see if there are other problems that we should/could revisit with this high order fix point machinery.