The Black-Scholes PDE from Scratch

chris bemis

November 27, 2006

Goal: Derive the Black-Scholes PDE

To do this, we will need to:

- ⋆ Come up with some dynamics for the stock returns
- ⋆ Discuss Brownian motion
- ⋆ Look at Ito's lemma
- ⋆ Discuss replicating and self-financing portfolios
- ⋆ Cleverly put some pieces together

In the (additive) binomial tree model, we are led to model the returns from a stock as

$$\frac{\delta S_t}{S_t} = \mu\delta t + \sigma\sqrt{\delta t}. \tag{1}$$

We may like to find the continuous version of (1). To do this, we need to use Brownian motion.

What is Brownian motion? A Brownian motion is a stochastic process; i.e. a family of random variables indexed by t: $\{W_t\}_{t \geq 0}$ such that

- $\star$ The function $t \rightarrow W_t$ is almost surely continuous
- $\star$ The process has stationary, independent increments
- $\star$ The increment $W_{t+s} - W_s$ is normally distributed with variance t.

How does this apply to the trees we have already seen?

For $n \geq 1$, consider the stochastic process $\{W_t^n\}_{t\geq 0}$ given by

$$W_t^n = \frac{1}{\sqrt{n}} \sum_{1\leq j\leq \lfloor nt \rfloor} \varepsilon_j$$

with each ε_1, ε_2, ... a sequence of independent standard normal random variables ($\varepsilon_j \sim N(0,1)$).

W_t^n is a random walk that takes a new step every $1/n$ units of time. For n large, we can see the connection to trees.

By the Central Limit Theorem

$$\frac{1}{\sqrt{\lfloor nt \rfloor}} \sum_{1 \le j \le \lfloor nt \rfloor} \varepsilon_j$$

converges (in distribution) to a standard normal random variable, Z. Now

$$W_t^n = \frac{\sqrt{\lfloor nt \rfloor}}{\sqrt{n}} \frac{1}{\sqrt{\lfloor nt \rfloor}} \sum_{1 \le j \le \lfloor nt \rfloor} \varepsilon_j$$

And since $\lim_n \to \infty \frac{\sqrt{\lfloor nt \rfloor}}{\sqrt{n}} = \sqrt{t}$, in the limit we have

$$W_t = \sqrt{t} Z$$

One may rigorously define the infintesimal increment of a Brownian motion. We won't. But we will use it. Before doing so, we notice that for s, $t \in \{0,\, 1/n,\, 2/n,\, \ldots\}$,

$$\begin{aligned} W^n_{t+s} - W^n_t &= \frac{1}{\sqrt{n}} \sum_{1 \le j \le n(t+s)} \varepsilon_j - \frac{1}{\sqrt{n}} \sum_{1 \le j \le nt} \varepsilon_j \\ &= \frac{1}{\sqrt{n}} \sum_{nt+1 \le j \le n(t+s)} \varepsilon_j \end{aligned}$$

Again we have that $\frac{1}{\sqrt{ns}} \sum_{nt+1 \le j \le n(t+s)} \varepsilon_j \to N(0,1)$ in distribution.

So that $W^n_{t+s} - W^n_t \to N(0,s)$. Or, $W_{t+s} - W_t = \sqrt{s} Z$.

It therefore seems plausible that dW_t is like $\sqrt{dt}$

From the binomial tree with drift equation (1), we could guess that

$$\frac{dS_t}{S_t} = \mu dt + \sigma dW \tag{2}$$

is a reasonably similar model. In fact, this model is the continuous time analogue of the binomial tree.

To derive the Black-Scholes PDE, we will need the dynamics of (2) we just stated.

We will also find that we need to take differentials of functions, $f(S_t, t)$, where S_t has the dynamics of (2). This is handled using Ito's lemma.

Before looking at this lemma, though, we will see why we need to take differentials of such functions.

We'll first talk about arbitrage, and then see how arbitrage can determines prices.

We have already seen how to determine the price of a contingent claim using risk-neutral probability (martingales, change of measure, etc.).

Just to be clear, examples of contingent claims are call options and put options.

A call option gives the holder the right (but not the obligation) to buy a specified item for an agreed upon price at an agreed upon time.

A put option gives the holder the right (but not the obligation) to sell a specified item for an agreed upon price at an agreed upon time.

We may also use *arbitrage* arguments. *Arbitrage* is simply (risk-free) free money. And an *arbitrage argument* says that there should be no (risk-free) free money.

How do we 'use arbitrage' to price a claim? We try to replicate the claim with stocks and bonds. We call stocks and bonds securities.

A contingent claim, f, is *replicable* if we can construct a portfolio Π such that

- The values of Π and f are the same under every circumstance.
- Π is self financing. As time goes on, we only shift money around within the portfolio, we don't put anymore in (or take any out).

We will call Π the replicating portfolio (of f).

Why does arbitrage work? Let's do an example with gold.

Suppose the price of gold today is \$200 and the risk-free interest rate is 3%.

You don't want gold today (because it's out of fashion), but you do want gold in 6 months (when, of course, it will be all the rage).

You therefore buy a *forward contract*. This says that you will receive gold in 6 months. You are locking in a price today for something you'll buy in half a year.

How much should you pay for this wonderful opportunity?

Suppose the forward contract costs \$250. You should then go to the bank, and borrow \$200. Use this money to buy some gold right now. Then *short* (sell) the forward (to a sucker).

In six months, what happens?

- You sell your gold for \$250
- You pay back your loan with your newly received funds
- You are left with \$250-\$200$e^{.5(.03)}$=\$46.97

Which is a lot of free money.

What if the forward contract, F_0, is selling for less than $\$200e^{.5(.03)}$? Well, you have to be able to sell an ounce of gold today.

Assuming you have gold lying around, you'll (because you know the trick) sell your gold today and get \$200. Next, you put this \$200 in the bank. Finally, you go *long* (buy) the forward contract.

So what happens at the end of 6 months?

- Take your money, $\$200e^{.5(.03)}$ out of the bank.
- Use it to buy your gold back for $\$F_0$.

You have your gold back, and $\$(200e^{.5(.03)}\text{-}F_0)$. Since this number is positive, you are very happy.

Arbitrage therefore sets the price of the forward contract to be \$$200e^{.5(.03)}$. If the price is anything else, there is risk-free free money to be made.

This is true of any forward contract on an asset with no storage costs and which does not pay dividends and if we assume interest rates are constant.

Even more generally, we have that any replicable claim will have the same price as its self-financing replicating portfolio.

Forward contracts are simple(!) to price. This is due in large part to the linearity of the payoffs at maturity.

Options are not so easy. The payoff at maturity has a kink.

However, we may construct a self-financing portfolio.

Now we will need Ito's Lemma.

If $dS_t = S_t \mu dt + S_t \sigma dW$, and $f : (S_t,\, t) \to \mathbb{R}$, we would like to determine df.

In Newtonian calculus, if $dx = (dS_t, dt)'$, we would simply have

$$
\begin{aligned}
df &= (\nabla f, dx) = \frac{\partial f}{\partial S} dS + \frac{\partial f}{\partial t} dt \\
&= \left(\frac{\partial f}{\partial S} S_t \mu + \frac{\partial f}{\partial t} \right) dt + \frac{\partial f}{\partial S} S_t \sigma dW
\end{aligned}
$$

But we observed that dW is like $\sqrt{dt}$. So our first order expansion should include one second order term.

If we believe that $(dW)^2 = dt$, we need to look at

$$\frac{1}{2}(dx, \nabla^2 f dx)$$

If we do, we see that:

$$\frac{1}{2}\frac{\partial^2 f}{\partial S^2}(dS)^2 \quad = \quad \frac{1}{2}\frac{\partial^2 f}{\partial S^2} S_t^2 \sigma^2 dt$$

up to first order.

We therefore have Ito's Lemma

$$df = \left(\frac{\partial f}{\partial S} S_t \mu + \frac{\partial f}{\partial t} + \frac{1}{2} \frac{\partial^2 f}{\partial S^2} S_t^2 \sigma^2 \right) dt + \frac{\partial f}{\partial S} S_t \sigma dW \tag{3}$$

with the same dW from (2).

How will we use this?

The only randomness in df is the dW term. So if we can construct a portfolio that eliminates the random part, we know exactly how the portfolio should behave.

For the first showing of this derivation, we will rely on the discrete versions of (2) and (3). We can prove this with much more rigor, but it is not much more enlightening.

Our goal is to price a contingent claim, or derivative.

We set Π to have

$$\begin{array}{rcl} -1 & : & \text{derivative} \\ \Delta & : & \text{shares} \end{array}$$

where $\Delta = \frac{\partial f}{\partial S}$.

We get that for a small change in time, δt, the corresponding change in Π is given

$$\delta\Pi = -\delta f + \Delta \delta S$$

From the discrete versions of (2) and (3), we get

$$\delta\Pi = \left(-\frac{\partial f}{\partial t} - \frac{1}{2}\frac{\partial^2 f}{\partial S^2}\sigma^2 S_t^2\right)\delta t. \tag{4}$$

But this implies the change in the portfolio is riskless (no uncertainty), and so arbitrage arguments, we must have

$$\begin{aligned}
\delta\Pi &= r\Pi\delta t \\
\left(-\frac{\partial f}{\partial t} - \frac{1}{2}\frac{\partial^2 f}{\partial S^2}\sigma^2 S_t^2\right)\delta t &= r(-f + \Delta S)\delta t \\
\left(\frac{\partial f}{\partial t} + \frac{1}{2}\frac{\partial^2 f}{\partial S^2}\sigma^2 S_t^2 + r\Delta S\right)\delta t &= rf\delta t \\
\frac{\partial f}{\partial t} + \frac{1}{2}\frac{\partial^2 f}{\partial S^2}\sigma^2 S_t^2 + r\Delta S &= rf \qquad (5)
\end{aligned}$$

The pde in (5) is the Black-Scholes-Merton differential equation:

$$\frac{\partial f}{\partial t} + \frac{1}{2}\frac{\partial^2 f}{\partial S^2}\sigma^2 S_t^2 + r\frac{\partial f}{\partial S}S - rf = 0$$

with Cauchy data $f(S_T, T)$ known.

By using only (2) and arbitrage, we must have that

- Any function f that satisfies (5) is the price of some theoretical contingent claim.
- Every contingent claim must satisfy (5).

When

$$\frac{\partial f}{\partial t} + \frac{1}{2}\frac{\partial^2 f}{\partial S^2}\sigma^2 + r\frac{\partial f}{\partial S}S = rf$$

is solved with boundary conditions depicting a European call option with strike K,

$$f(S,T) = \max(S-K,0),$$

we get the Black-Scholes price of the option.

The BS price of a European call, c, (on a stock with no dividend) is

$$c = c(K, r, S_t, t, T, \sigma) = S_t \Phi(d_1) - K e^{-r(T-t)} \Phi(d_2) \tag{6}$$

$$\begin{aligned} d_1 &= \frac{\ln(S_t/K) + (r + \sigma^2/2)(T-t)}{\sigma\sqrt{T-t}} \\ d_2 &= d_1 - \sigma\sqrt{T-t} \end{aligned} \tag{7}$$

Φ is the cumulative distribution function of standard normal random variable ($N(0,1)$)

Here are a few properties of the BS price of c (a benchmark test, really)

- We would expect that if S_t is very large, c should be priced like a forward contract (why?). We see that if S_t is large,

$$c \approx S_t - Ke^{-r(T-t)}$$

  which is, in fact, the price of a forward contract (why?).

- When σ is extremely small, we would expect that the payoff would be

$$c \approx \max(S_t e^{r(T-t)} - K, 0) \tag{8}$$

  (why?).

We also have

- c is an increasing function of σ.
- $\frac{\partial c}{\partial S} = N(d_1)$.

From the last point, we can estimate the Δ to use in the replicating portfolio of c.

So we see that the price determined by risk-neutral expectation is the same as the price determined by solving the Black-Scholes pde.

Everything seems to be going swimmingly.

Next up....

Implied Volatility, and Where Black-Scholes is Going Wrong

Prices are not set by the BS options price. Rather, markets set prices (and if you believe some economists, they set prices near perfectly).

We may therefore go to the market to see what a call option on a certain underlying is selling for right now at $t = 0$.

We observe K, r, S_t, T. We can't observe σ, though.

We solve for σ using (6). This is relatively easy since the BS call option price is monotonic in σ. The number we get is called the implied volatility.

If we check market data for different strike prices, K, with all else being equal, we get different implied volatilities.

In fact we get what is called a volatility smile, or a volatility skew depending on the shape.

Why is this a problem? We have assumed that σ is some intrinsic property of the underlying. It shouldn't vary with K.

Below are the prices for (European) call and put options on the QQQ (a NASDAQ 100 composite) for January 9, 2004. Expiration dates are January 16, and February 20.

	Calls		*Puts*	
Strike	*January*	*February*	*January*	*February*
34	3.9	4.1	0.05	0.25
35	2.8	3.2	0.05	0.35
36	1.85	2.35	0.1	0.55
37	1	1.65	0.25	0.85
38	0.35	1.05	0.6	1.25
39	0.1	0.6	1.4	1.9
40	0.05	0.35	2.35	2.6

As we have seen, BS depends on $(K, r, S_t, t, T, q, \sigma)$, and the only unobservable quantity is σ. In the present case, for the February options, the data give

$$
\begin{aligned}
S_0 &= 37.73 \text{ (the price at closing Jan. 9, 2004)} \\
T - t &= 42/365 = .1151 \\
r &= .83 \\
q &= .18
\end{aligned}
$$

This gives

Implied Volatility

Strike	February Call	February Put
35	0.323	0.29
36	0.2592	0.2493
37	0.2455	0.2369
38	0.2279	0.2198
39	0.2156	0.2279
40	0.2181	0.2206

Graphically, plotting strike prices on the x-axis and implied volatility on the y-axis, we have:

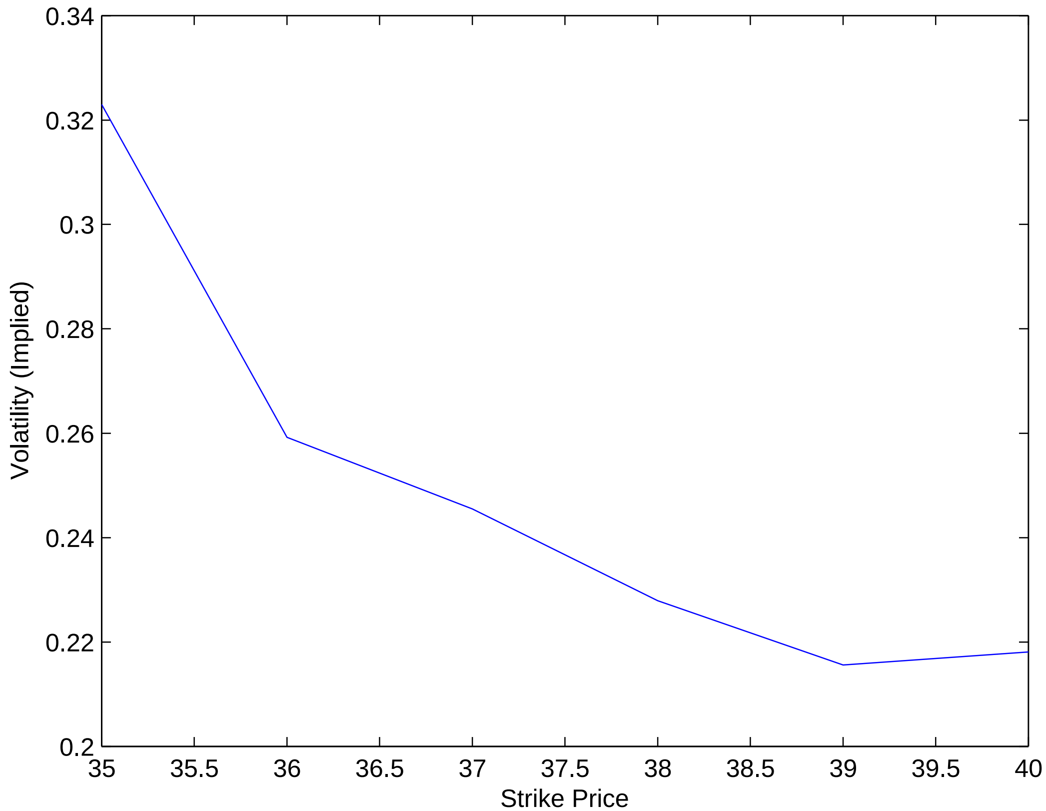

Sometimes things are not so perfect. Suppose the volatility smile we observe looked more like:

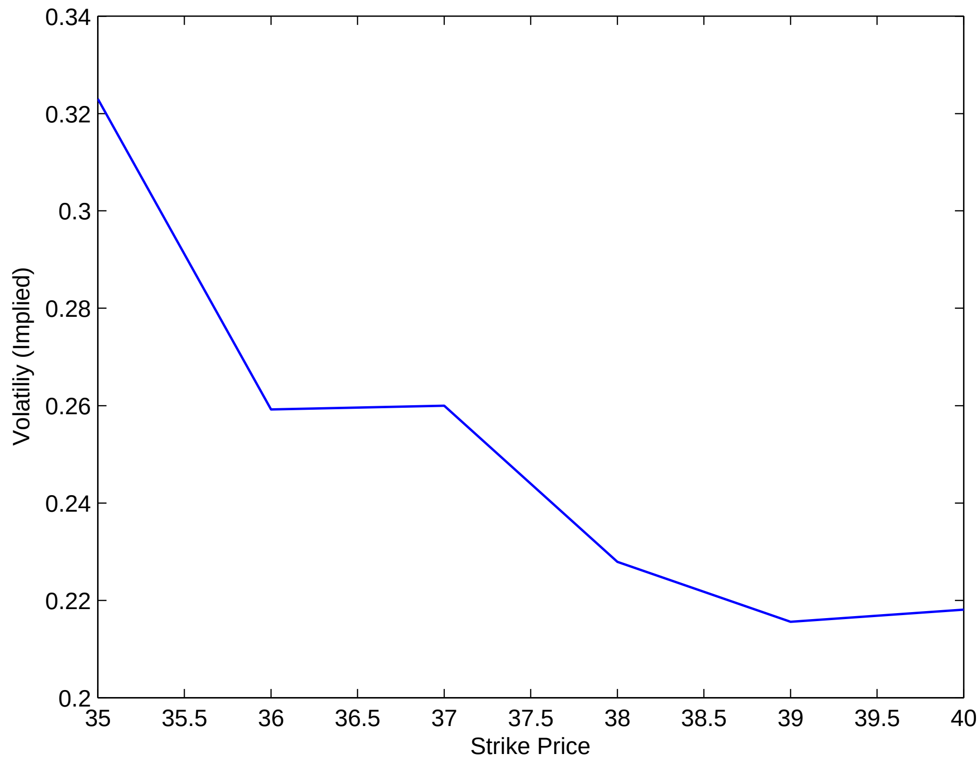

We would likely think that the market was overpricing the call for one of the strike prices (which one?), and take a position.

Volatility smiles also occur with commodities. Below are examples of smiles for both calls and puts for crude oil.

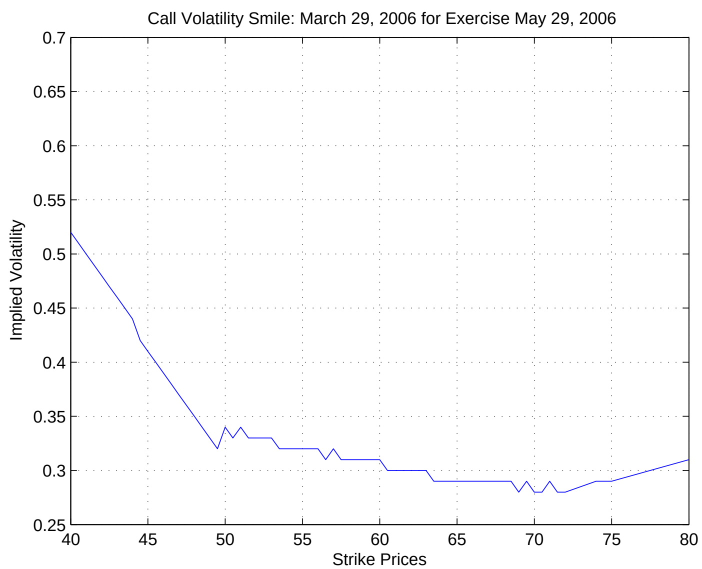

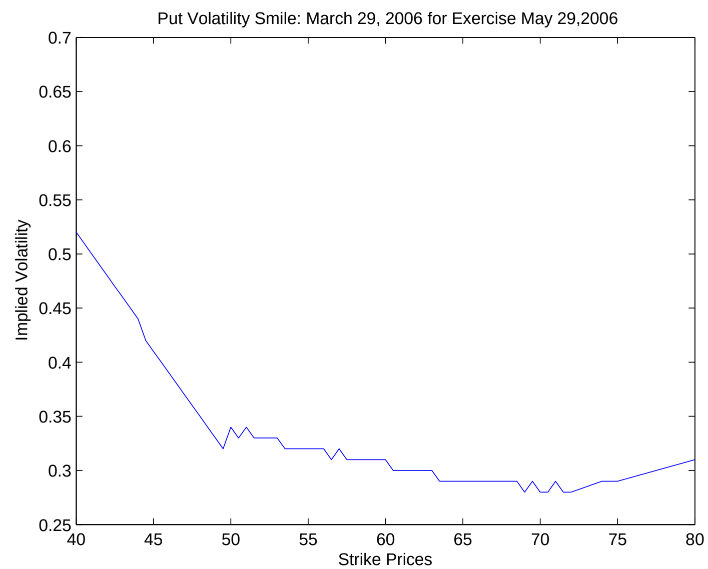

So σ not only varies with the strike price, but also depends on whether we are pricing a call or a put. Below are the volatility smiles of the call and put above in one plot.

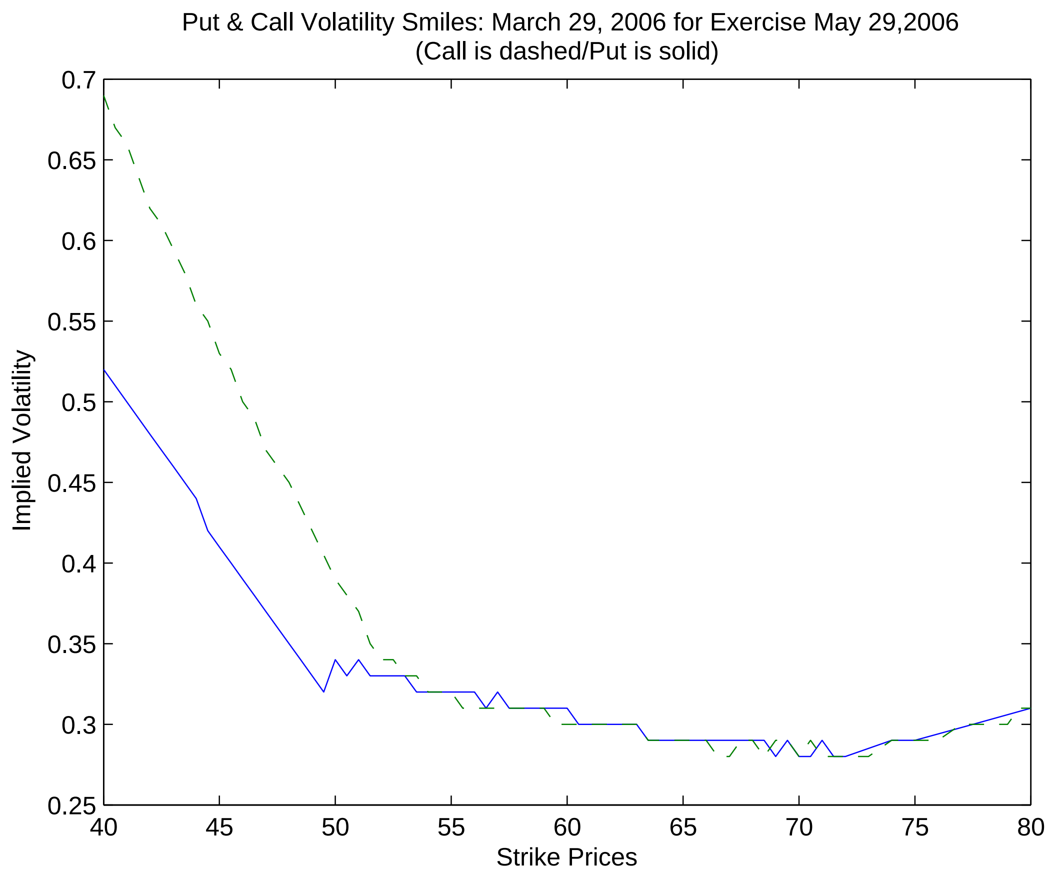

As a final kicker, implied volatility varies with the expiration of the option. We may therefore plot a volatility *surface*.

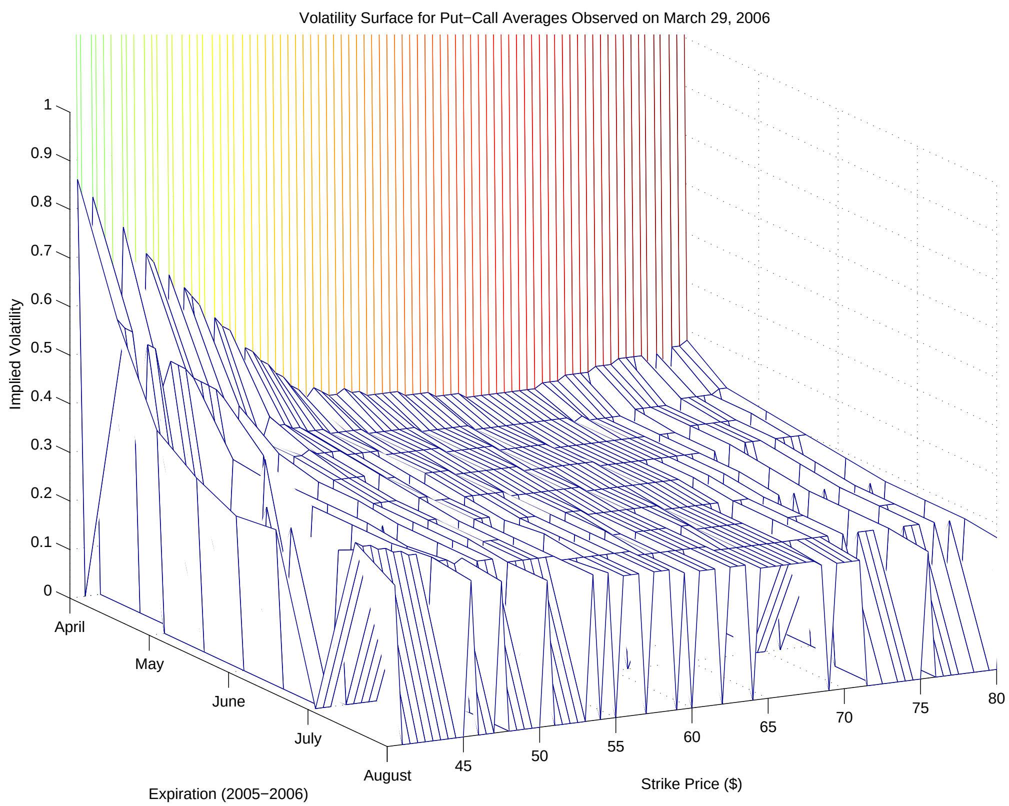

In the end, Black-Scholes is used to show that Black-Scholes is lacking. We could enrich the model. Some prime suggestions are

- Assume volatility is stochastic. That is, let $\sigma = \mu_\sigma dt + \hat{\sigma} dW$.
- Assume volatility is local. That is, $\sigma = \sigma(S, t)$.
- Assume the process that the underlying follows is a jump-diffusion process.
- Assume interest rates are, at the very least, nonconstant.

Everything that is tweaked, however, leads to more issues. Today, there is no clear successor to the BS model.