Kent Osband

FINFORMATICS

# Pictures of Learning

Time to calculate "the weight of the evidence" and apply the concept to solve particularly tricky market mysteries

My previous article showed that smart learners ought to adjust the log odds of beliefs – that is, the logarithm of the ratio of the probabilities of two competing hypotheses – by a statistical quantity known as "the weight of the evidence". To calculate the weight of the evidence:

- compute the probability that the evidence would have occurred if the first hypothesis were true
- compute the probability that the evidence would have occurred if the second hypothesis were true
- form the ratio of the two probabilities
- take the logarithm of the ratio.

Since new evidence continually spews up, smart learners are obliged to keep repeating this process, endlessly updating their beliefs. That makes the log odds for any smart learner the cumulative sum (or integral, in continuous time) of all the weights of evidence observed in the past, plus the log odds of whatever beliefs you started with.

In this article we're going to practice applying these concepts. In the process we'll find remarkably simple explanations for some market behaviors that baffle classical finance.


*I didn't see that one coming*

## Tides of disagreement

The first problem we'll tackle is the influence of the initial or "prior" beliefs. Suppose two smart learners start with different priors, so that the difference in the initial log odds equals some nonzero $\Delta$. If each of them observes the same evidence and incorporates it properly, they will adjust their log odds by identical amounts. So what happens to the difference in their log odds? That's right: nothing. It will remain $\Delta$, regardless of the evidence.

At first glance, this seems doubly strange. It's

## Opinions can ultimately converge if regimes change; that is, if the past occasionally ceases to be relevant

**LOGISTIC RELATIONSHIP BETWEEN LOG ODDS AND PROBABILITY**

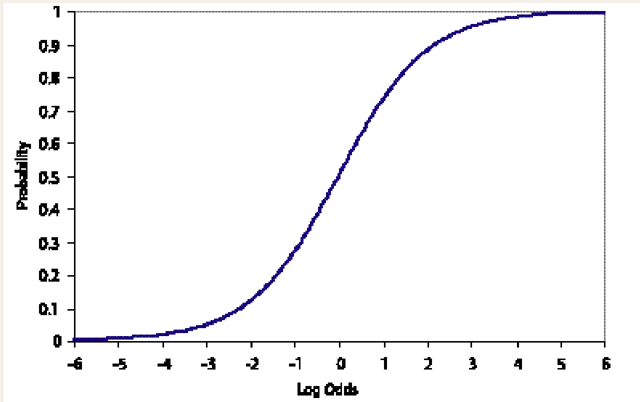


strange theoretically, because you would think that two rational people carefully analyzing the same evidence ought to eventually converge in beliefs. But it's also strange practically, because even among rational people disagreements ebb and flow like tides, with few signs of constant spacing. Or perhaps I should say " especially among rational people", because irrational people tend to get well and truly stuck in the rut of wishful thinking.

In fact, all three perspectives are easily reconciled. The "constant $\Delta$" result refers not to beliefs directly but to the log odds of beliefs. Denoting beliefs by $p$, the log odds $\eta$ are defined as

$$\eta = \ln\left(\frac{p}{1-p}\right).$$

Hence, given $\eta$, $p$ can be calculated as:

$$p = \frac{1}{1+e^{-\eta}}.$$

This makes $p$ what is known as a logistic function in $\eta$, more precisely a standard logistic function. A standard logistic function veers toward zero or one outside of a transitional range centered on and symmetric around the origin. That shouldn't be too surprising, given that probability is bounded by 0 and 1 while log odds stretch from $-\infty$ to $+\infty$, and given that switching a core hypothesis with its complement shouldn't make any fundamental difference.

The relationship is graphed above. The horizontal distance between two points represents the log odds difference between two beliefs, while the vertical distance represents the corresponding probability difference. Clearly, a constant horizontal distance can make for a wildly varying vertical distance.

To illustrate, let's suppose Punch's prior about something is $\eta_0$=+3 while Judy's is $\eta_0$=-3. That is, initially Punch is 95 per cent sure that some hypothesis is true, whereas Judy is 95 per cent sure it's false. If new evidence comes along that raises $\eta$ by 10, even Judy will be 99.91 per cent confident.

But disagreements can wax as well as wane. Suppose Punch starts out 99.9998 per cent confident in something, whereas Judy is slightly less confident, by a measly 0.0909 per cent. Neither Punch nor Judy can bother fighting over such microscopic slivers of disagreement. But suppose the hypothesis is false. Eventually, if they keep observing long enough, both will swing toward the correct view. But here Judy will be much quicker to the punch than Punch. Indeed, there is bound to be one day when Judy has lost 95 per cent of her confidence while Punch has lost only 5 per cent. How do I know that? Because all I've done is run the first example in reverse, with a stable difference of 6 in the log odds.

Do differences in priors always persist without change? No. Opinions can ultimately converge if regimes change; that is, if the past occasionally ceases to be relevant. On the other hand, uncertainty about the nature of new regimes can breed more divergence. So I would caution readers against dismissing what we've just seen. Remember that standard finance imposes an even stricter assumption, where the regime is fixed for all time and everyone knows what it is. As a result, standard finance can't even explain why reasonable people disagree, much less why disagreements ebb and flow. The model here can. Which do you think is a sounder foundation for finance?

### The volatility of volatility

Suppose we know a coin is biased by some fixed amount $\epsilon$ but we're unsure about the direction. Denoting heads by 1, tails by 0, and the unknown mean by $\mu$, and the weight of the evidence if heads appears as $\delta$, we can calculate:

$$\delta = \ln\left(\frac{P(1|\mu = 1/2+\epsilon)}{P(1|\mu = 1/2-\epsilon)}\right)$$
$$= \ln\left(\frac{1/2+\epsilon}{1/2-\epsilon}\right) = \ln\left(\frac{1+2\epsilon}{1-2\epsilon}\right).$$

Similarly, the weight of the evidence if tails appears is:

$$\delta = \ln\left(\frac{P(0|\mu = 1/2+\epsilon)}{P(0|\mu = 1/2-\epsilon)}\right)$$
$$= \ln\left(\frac{1-(1/2+\epsilon)}{1-(1/2-\epsilon)}\right) = \ln\left(\frac{1-2\epsilon}{1+2\epsilon}\right) = -\delta.$$

Hence the log odds $\eta$ will be a random walk with increments $\pm\delta$. If the bias is positive, the drift will be

$$\left(\frac{1}{2}+\epsilon\right)\delta + \left(\frac{1}{2}-\epsilon\right)(-\delta) = 2\epsilon\delta.$$

If the drift is negative, the drift will be

$$\left(\frac{1}{2}-\epsilon\right)\delta + \left(\frac{1}{2}+\epsilon\right)(-\delta) = -2\epsilon\delta.$$

Last but not least, suppose initially we think either regime is equally likely, so that the prior log odds $\eta_0$ is zero and the initially expected drift is zero. Hence, everything is neatly symmetrical around the origin and as simple as this kind of problem can be.

Now the price $S$ of a claim on a sequence of coin flips will rise with the expectation of heads,

## MONTE CARLO SIMULATIONS OF VOLATILITY

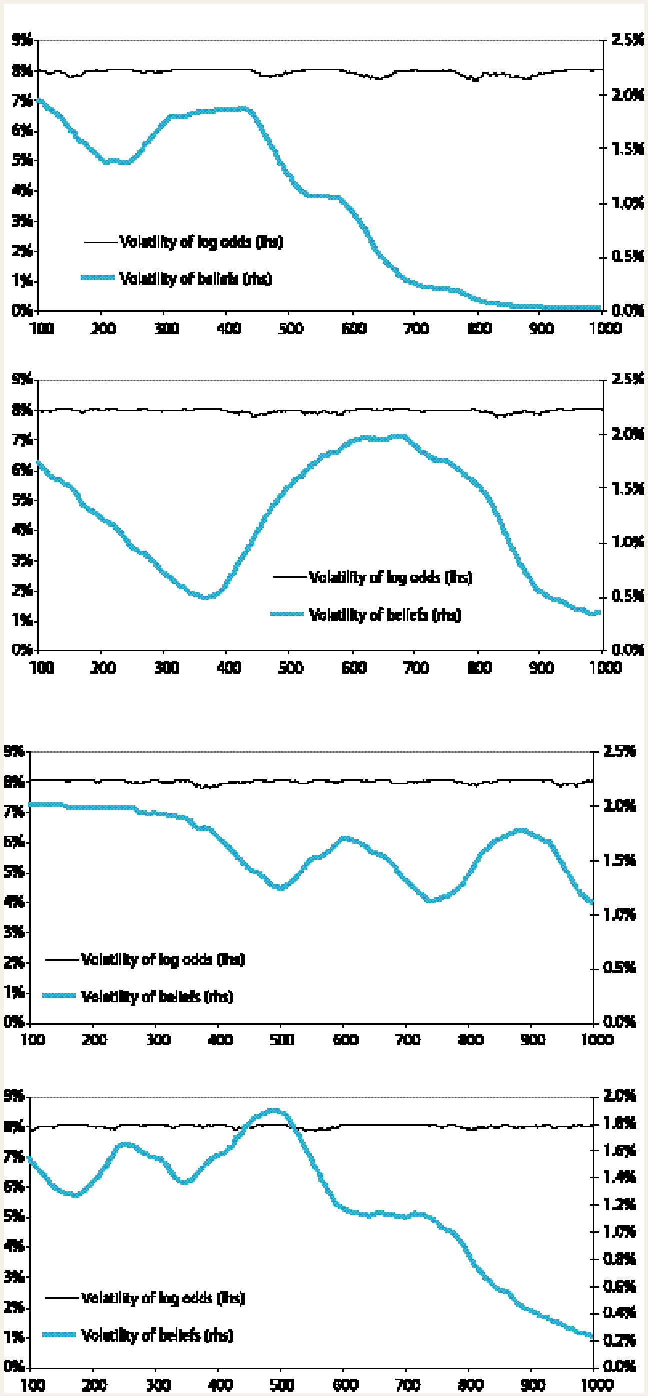


## The more nearly convinced you are about something to begin with, the less noticeable impact any new information will make

all else being equal. In other words, holding the coin flipping schedule and payoffs for heads and tails are fixed, $S$ will be a monotonically increasing function of $p$. To simplify some more, let's assume buyers are risk neutral. In that case, they must be indifferent between a claim with perceived probability $p$ and a claim consisting of fractional shares $p$ in the flip of a good coin and $1$-$p$ in the flip of a bad coin. That implies $S$ is linear in $p$, with

$$S(p)(pS(1) + (1 - p)S(0)$$

$$= S(0) + p(S(1) - S(0)).$$

It follows that the volatility of the price of a claim will be $S(1) - S(0)$ times the volatility of $p$, where $p$ is the standard logistic of $\eta$ and $\eta$ is an ordinary random walk. To give you a feel for what this looks like, volatilities for four simulations are charted below. In each of these simulations, $\epsilon$=0.02. Volatilities are calculated on a rolling 100 flip basis and reported for flips 100 to 1000. For comparison I chart volatilities both for log odds and for beliefs and 1000 coins flips.

In all the simulations, the sample volatility of logodds is very stable, with just a tiny flutter. This is true for all ordinary random walks, including the continuous variant known as Brownian motion. In contrast, the sample volatility of beliefs and/or prices shows much more variety, with pronounced and irregular fluctuations, although volatility does tend to decline over time.

The explanation is simple. Volatility of beliefs varies with current beliefs for the same reason that the impact of news varies with current beliefs. Mathematically,

$$\frac{dp}{d\eta} = \frac{e^{-\eta}}{(1 - e^{-\eta})^2}$$

$$= p(1 - p)$$

which is maximized at $p$=1/2 and shrinks toward zero at the edges. In other words, the more nearly convinced you are about something to begin with, the less noticeable impact any new information will make – at least insofar as you process new information sensibly.

So the current model, simple as it is, is far more realistic than standard finance theory, except for the secular decline in volatility as uncertainty eventually dissipates. The obvious fix for that is to allow sources of new uncertainty through regime switching. That's where we're heading in subsequent models. For those of you who can't wait, do a Google search for "regime switching finance" and you'll find a host of interesting stuff. I particularly recommend some papers by Pietro Veronesi of the University of Chicago, W