AARON BROWN

Financial Risk Management... for Dummies

Respected probabilist David Aldous reviews Aaron Brown's demystification of risk management

Readers will know Aaron Brown as a columnist for this magazine and as the author of *Red-Blooded Risk: The Secret History of Wall Street* and *The Poker Face of Wall Street*, and the co-author of *A World of Chance: Betting on Religion, Games, Wall Street*. Unlike those more traditional style books, the book under review is in the familiar ... *for Dummies* style of short, brisk sections replete with stories and asides. Brown writes that he does not seek to teach basic financial theory and the basic related mathematics. What he achieves is a very comprehensive, and engagingly opinionated, account of both the explicit responsibilities and the implicit world-view of a good financial risk manager (FRM). Here are a few of his many nuggets of wisdom.

"*To a portfolio manager, risk is something bad to be minimized, just like a cost ... To a risk manager, risk is something to be set at the correct level.*"

"*If you're going to hedge, select the best hedge, and explain the likely outcomes clearly to everyone. One of those likely outcomes is that the hedge will turn out to be an overhedge, and will lose more money than the exposure gain. If that outcome is unacceptable, don't hedge in the first place.*"

"*I'd like to tell you how to use shareholder communications to help shareholders understand the forward-looking risks of the company and the actions taken to manage that risk. Unfortunately, I don't know how – I've never done it successfully.*"

And, finally, Brown's comment on extreme value theory – "*When you're looking at a few of the largest past events and trying to think about plausible extreme future events, there's just enough data to support enthusiastic mathematical modelling without having enough data to rein in wild ideas*" – is one I will gleefully relay to some mathematics colleagues.

At this point, your reviewer should admit to having no personal knowledge of FRM, and so I cannot assess the main substance of the book beyond emphasizing that Brown reveals it to be a fascinating and amazingly multifaceted job. My own academic career at Berkeley has centered on obscure mathematical probability theory research, although as a background activity over the last 15 years I have intermittently taught a course with the grandiose title "Probability in the Real World." The course seeks to explore the breadth of contexts where we perceive chance, via 20 lectures anchored (ideally) by real data and on topics not covered in other courses. (To digress to an example, a lecture on "Everyday Perception of Chance" includes data from the search engine Bing on 100,000 queries containing the strings "chance of" or "probability of"; typical examples are "chance of getting pregnant while breastfeeding"; "chance of getting a brain tumor"; "chance of getting shot if you run from an attacker"; "chance of becoming a USAF airlift pilot"; "probability of winning a traffic ticket court case".) A broad theme of my course – and therefore what I personally look out for when reading books like Brown's – is: "when is it useful to try to assess probabilities (and impacts) of future events quantitatively?" That broad theme recurs throughout Brown's book in rather intricate ways. He emphasizes the point that you have front-line risk takers seeking to exploit uncertainty for profit in some specialized context. To the extent that there are predictable aspects or reliable models in a given context, the FRM can and should use the conclusions without engaging the statistical details – "... *stay scrupulously away from kibitzing*

Any risk manager who's been making risk decisions longer than five years has an unshakeable faith in VaR ... The discipline of estimating and checking VaR every day forces you to learn about an awful lot of stuff, and a surprising amount of that stuff is useful

Musings on Random Musings

Aaron Brown returns the compliment …

As a rule, I do not respond to reviews of my books because I feel that after having 384 pages to express myself in any manner I chose, readers should be spared additional clarification. The book is what it is, and the reviewer's reactions are what they are; any further discussion should be taken up in a new book or article. In this case, however, Professor Aldous raises some fascinating issues that go beyond the book and are worth embellishing.

The only point of disagreement I have with the review is Aldous' modesty. His "*obscure mathematical probability research*" is actually a stellar 40-year career as one of the leading theoreticians in the field, decorated with most of the honors on offer. And the "*background*" "*intermittent*" activity he refers to is his *Probability in the Real World* project, the best attempt anyone has ever made to link mathematical probability to generalized notions of chance and practical applications. I can say that, because it's not about my book.

(1) I'm also a big fan of Philip Tetlock (I reviewed *Superforecasting* in the July 2015 issue of *Wilmott*) and agree that we can assess the relative accuracy of individuals in estimating probabilities without knowing true probabilities, or even without assuming that true probabilities exist. The disadvantage of that is that we can assess only after outcomes are known and, in fact, only after many predictions are made and many outcomes known. A related point made by Tetlock is that we can assess the calibration of individual forecasters without knowing anything about true probabilities. I am personally more impressed by the accurate calibration of his superforecasters than by their track records based on outcomes. If we supplement asking for the probability that *China will declare an Air-Defense Identification Zone over part of the South China Sea before October 1, 2016*, with *How about by July 1, 2016*? and *How about China making any declaration of increased sovereignty over the South China Sea before October 1, 2016*? we can provoke deeper thinking by the forecaster and get more material to evaluate the predictions against outcomes. Predictions with consistent calibrations do not have to be accurate but inconsistent calibration is a sign that we have not captured the predictor's knowledge completely.

(2) In a January 2015 column in *Wilmott*, "The Invisible Hand Rolling Darwin's Dice," I reviewed Curtis Johnson's book *Darwin's Dice*. That book carefully parses Darwin's public and private writings over his lifetime, to uncover his belief about the source of what we now call random variation. Everyone knew that offspring were not identical to parents, and it was obvious that population frequencies of traits changed based on whether they favored successful reproduction or not. Darwin's controversial key belief was that *individual* frequencies of traits did not change based on whether or not they were favorable in the current environment; the population frequencies changed only because individuals with favored traits were more likely to survive and reproduce, and because offspring have a tendency to resemble their parents. Survival of the fittest, not birth of the fittest. This is random in the sense of *uncaused*, or at least, *uncaused by design considerations or concern for individuals*.

Professor Aldous raises some fascinating issues that go beyond the book and are worth embellishing

(3) It is true that refusal to put a dollar value on human life, or anything else, will generally decrease the quantity of it. It's also true that people saying, "*You can't put a dollar value on human life*," are usually unaware of that. But most people care more about how lives are lost than how many lives are lost. A proposal that promised to prevent 1,000 gun homicides per year would get more support than one that promised to prevent 1,000 gun suicides but a lot less than one that promised to eliminate the indiscriminate mass shootings in public places that kill an average of about 50 people per year. A bureaucrat who shot a retiring teacher to save the state $1 million in pension benefits would go to jail, but one who cut healthcare funding by $1 million, with an expected consequence of one additional 65-year-old person dying, might get promoted. You might argue that people shouldn't feel this way, but they do. You might argue that human life is a complicated asset, like an acre of real estate, with varying value depending on which life or acre we're talking about and what you want to do to it. I argue that, because how the decision involving human life is made is often more important than the result of the decision, it's a mistake to put it into a standard expected utility decision theory framework.

(4) I wouldn't make voluntary participation a defining characteristic of games. I think they are contests governed by rules (not ethics, as Aldous points out), played in a court or arena marked off from everyday life, in which external considerations are not supposed to matter, often involving randomization and where people usually wear funny clothes and engage in rituals, including funny words, chants, alcohol, and music. You pass to your teammate in basketball, not to your friend, or to the guy you owe money to, or to the player who needs the ball most, or to the highest bidder. A lawyer in court tries to introduce or suppress evidence based on whether it helps their client or not, not based on whether it is accurate or relevant, or on her personal judgment about admissibility. Ethics are better than rules when they're good ethics and everyone agrees on them, but rules are necessary when clear decisions accepted as fair are more important than either morally or consequentially good decisions. Sports have both types of regulation; hitting a player's hand to prevent a layup is part of the game because your team benefits from giving them two one-point foul shots instead of a sure two-point basket. Grabbing their ankles with the intention of injuring them is outside the game. I'd say that the danger in financial

Continued on page 44

regulation is that principles-based regulations are interpreted by market participants as a rules-based game. The opposite problem occurs as well, when arguable technical violation of arbitrary rules which should, at most, carry fines to offset the damage is treated as moral failings to be punished with long prison sentences. It's possible to have good principles-based financial markets and good rules-based financial markets, but mixing the two is dangerous, especially without clear distinction.

(5) This is a useful point. One of the first hurdles you face in making money is that if your trading is successful, people don't want to trade with you. This creates a more or less deterministic struggle as you try to put your trades on efficiently, and others seek to siphon off your profits. Another hurdle is that a successful strategy attracts imitators, and the imitation makes the strategy appear both more profitable and less risky than it really is; this, in turn, creates the negative tail event when the strategy overshoots its economics owing to traders piling in, so it falls, causing traders to pile out, exacerbating the decline. Avoiding this is a probabilistic struggle. Other traders aren't siphoning off your profits, they're adding to them – but in the process they're adding a hidden toxic outcome to your probability distribution of returns, and obscuring even the normal risks of the strategy. The strategy returns are still a random walk, but the die has changed to one that can ruin you if you don't notice the shift, say, from 1–2–3–4–56 to 5–6–6–6–6–negative 8. It's a jungle out there.

(6) Of course, I agree that the three points from a last-second field goal are no more important than three points scored at any other time in the game. It's somewhat more meaningful to say that a few big plays won an American football game. Quantitative analysis of football suggests that about 60 percent of game outcomes are determined by the 10 percent or so of plays that can be called 'big,' turnovers, long gains, big sacks, fourth down-plays, and so on. In an otherwise even football game, if a team intercepts a pass thrown from near its own endzone and returns it 100 yards or so for a touchdown, it has about a 90 percent chance of winning the game. So, it's not silly to single out that play; it's likely the biggest single factor contributing to the win (if the team does win, of course). With the S&P500 over a five-year holding period, the 99.7 percent of days with moves of less than 5 percent have a standard deviation of total return over the five years of 12 percent, while the 0.3 percent of days that move more than 5 percent have a 62 percent standard deviation. That is, if you want to know the total return of the stock market over a five-year period, knowing the total return on big days is much more informative than knowing the total return on all the other days; the big days explain 97 percent of the total variance of five-year total return. You get an average of 4.5 big days in five years, but the standard deviation of that number is 6.3; moreover, of course, the big days are much more volatile than the other days. Drawing a histogram of daily returns is fundamentally misleading because the important source of variation is invisible.

(7) This is, of course, a subject of active research, and opinions differ. The big arguments for overreaction are: (a) short-term volatilities of prices are greater than long-term volatilities scaled by the square root of time; (b) shrinking prices toward estimates of fundamental value reduces volatility and leads to more accurate predictions of future prices; and (c) most large price movements cannot be plausibly traced to news events. None of these are overwhelming logical points, or uncontested empirical ones. The massive uncertainties in predicting the economic value of Apple stock certainly could justify the volatility in the stock price, the 30 percent or so annual volatility, or even much more. But it's hard to credit that, on an average day with no significant overall stock market move or company announcement, news comes out that changes the best estimates by 2 percent. I have no idea what happens to people after death, but my opinion doesn't change much from day to day. Hard-to-credit does not mean false, but the weight of academic opinion is on overreaction, and people who make their livings by trading the stock spend nearly all their energy thinking about short-term technical factors rather than long-term economic value.

on any actuarial decisions." But the total level of risk to the entire business is not just some simple combination of separate risks; there are complex interactions between different markets and the real economy, and possible events that no front-line risk takers would consider. This big picture is the job of the FRM, and there is no formula for doing it.

In one chapter, Brown gives a very clear explanation of what VaR does and does not mean, and approves its use while acknowledging its limitations:

> "*Any risk manager who's been making risk decisions longer than five years has an unshakeable faith in VaR ... The discipline of estimating and checking VaR every day forces you to learn about an awful lot of stuff, and a surprising amount of that stuff is useful ... Time after time in the past, little disturbances in VaR have been the only warning that the markets give of crises to come ...*"

So, here, Brown is concerned with actual numerical values of VaR calculations. In the different context of stress tests, he emphasizes "*preparing for anything that might happen*," rather than worrying about precise likelihoods.

Random musings

Reiterating that I have no expertise in FRM, my remaining comments are seven (out of several dozen) thoughts arising from Brown's early chapters and their relationships to matters outside finance.

(1) Brown's short chapter on frequentism, Bayesianism, game theory, etcetera, will provide food for thought for anyone interested in the philosophical foundations of probability. My favorite actual data relevant to these philosophical questions is from Tetlock's Good Judgment Project on geopolitical forecasting, where contestants were required to assess probabilities of specified events occurring before specified deadlines. By calculating mean-square errors, one can reliably score the relative abilities of different individuals to assess the true probabilities, even though no one knows what the true probabilities were. My point is that we can actually do something which dogmatic frequentists and dogmatic Bayesians claim (for different reasons) is impossible, in that neither believes there is a *true* probability for events like *China will declare an Air-Defense Identification Zone over part of the South China Sea before October 1,*

2016. Incidentally, academics who actually do Bayesian statistics (instead of philosophizing about it) are much more sensible about its limitations; check the article *What are the Open Problems in Bayesian Statistics* by my colleague Michael Jordan.

(2) "*Darwinian evolution is defined as random variation and natural selection. It was the random part that was revolutionary when Darwin published ... in 1859.*" Not quite: Darwin talks about variation but admits he doesn't know how it arises, and (I believe) he never refers to randomness; it took another half-century for genetic theory and the notion of random mutations to be formulated and accepted.

(3) Regarding the Ford Pinto story – "*The outcome (burn deaths) cannot be measured in dollars ...*" – there is the (awkwardly named) *value of a statistical life* used mainly for government planning purposes (and typically around US$9 million in the US). If you have US$200 million to spend on road improvements, then, among various possible projects, some are more likely to save lives in accidents, so you should take that into account in any cost–benefit analysis. People who reflexively say "*You can't put a value on life*" believe that they are assessing the value as infinite. But, in fact, they are assessing it as zero if they ignore it in such decisions. On a somewhat related matter, the magnitudes of different physical dangers to individuals are best communicated in *micromorts* (a one in a million chance of death) and *microlives* (a 30-minute decrease in life expectancy, being about one millionth of an adult lifetime). In communicating assessments of very unlikely but possible causes of financial ruin of your institution, would some similar word be useful?

(4) In a brief digression on games, Brown cites legal disputes and elections as analogous to games. It is interesting to debate what differentiates a *game* from other human activities; to me, one criterion is that participation is voluntary. But a more interesting point is illustrated by boxing and chess. In an obvious sense, they are almost complete opposites, but they share a *shake hands and come out fighting* aspect. My point is that, in a game, rules substitute for ethics.

It is interesting to debate what differentiates a game from other human activities; to me, one criterion is that participation is voluntary

Brown writes about the increasing legal regulation of the finance industry; I worry that an increasing focus on regulation leads to a decreasing focus on ethical behavior – to an ethos of *if it's not actually illegal, then we can do it*.

(5) The usual phrase *random walks* used for basic probability models of stock prices is somewhat misleading because it suggests strong statistical regularity, like repeated throwing of a single die. A better mental image is that there is a different die each day; we don't know the numbers on the dice but they are constrained to have a slightly positive expectation. Mathematics tells you that the belief *it's very hard in practice to beat the market by some clever investment strategy* (being long/short in stocks; I'm not talking about derivatives) is logically equivalent to the belief *it's very hard in practice to distinguish actual daily stock price changes from those which would occur in the model above*. I am saying this as a complement to Brown's discussion of the market as a game which you should worry is somehow out to get you. Brown treats this as antithetical to the random walk image. But one can partly incorporate it into the random walk framework by imagining that the die is chosen each day by an adversary who knows your strategy but cannot control the outcome of the throw.

(6) Brown repeats a common type of assertion: "*…all the gain in the S&P500 from 1927 to 2014 comes from the best 90 days.*" This strikes me as a red herring – literally true but meaningless. Saying that a certain football game was won by a last-second field goal is literally true. But the inference that this field goal was the only thing that mattered "*because the other points net out*" is absurd; the result is a combination of all the scoring plays that succeeded or failed, and there is no logical justification for arbitrarily picking out the final such play, rather than (say) the first one – that's just an artifact of human time-oriented perception. Similarly, the S&P gain depended on the sum of contributions from all 'up' days; there is no logical justification for arbitrarily picking out 90 of them as of special significance. For a reinsurance company with a handful of large claims each year, it may well be true and meaningful that "*long-run outcome depends mainly on a small subset of extreme events*," but there is no reason to believe that this is meaningful for the stock market.

(7) Brown repeats another common type of assertion: "*Financial market prices are far more volatile than can be reasonably explained by changes in fundamental economic information.*" Really? More precisely, are they more volatile than in a hypothetical world, where buy and sell decisions were based only on estimated discounted future profits (crudely, if Apple's stock price today were proportional to the expectation of its profits in 2025)? I have never seen a convincing argument. Whatever economic information is observed during 2016 (are those iWatches selling well or badly?), our hypothetical people must somehow estimate the effect on profits in 2025. This surely involves so much extrapolation that it is impossible for you or I to know how volatile their estimates would be.

But my comments above represent very minor quibbles. If you have any interest in FRM, buy the book. Even if you know it all, you can usefully give it to the next young person you encounter who expresses an interest in a career in finance.