Aaron Brown

The Unbearable Lightness of Cross-Market Risk

Common assumptions could mean disaster for effective risk management writes **Aaron Brown**

Do you ever wonder if the Earth has an orbit? It seems to go around the sun pretty regularly, and its minor deviations from a perfect ellipse can be computed thousands of years into the future. But can we be confident it will stay within its comfortable 150 ± 3 million kilometers from the sun during our lifetimes, so we do not, like Icarus, perish from being too near the sun or too far?

We might derive some comfort from Isaac Newton's famous result that two incompressible spheres of uniform density form a stable gravitational system. But chaos theory teaches that even simple dynamic systems can behave regularly for a while, then fly off in another direction without measurable cause. Could the cumulative effects of tidal forces, gravitation from the Moon and Jupiter, relativity and other forces neglected in the classical two-body problem send the Earth to regions of solar inferno or icy death?

Does the Earth have a climate? Weather is a complicated system and a glance at the historical temperature record looks more like a random walk than a series with mean reversion. We like to think about a normal climate, with periodic ice ages and warm spells, but is there really any reason to believe that? How do we know the temperature, or the oxygen content of the air or the level of rainfall, will not randomly walk up or down to a value inconsistent with human survival?

Can we assume it tastes like chicken?

Choking on PIP

Most people spend little time worrying about such things. It is not a deep knowledge of astrodynamics or meteorology that reassures them. It is the Personal Insignificance Principle. When people use words like "orbit" and "climate," or for that matter "normal" and "usual," they are typically invoking some form of PIP. Risk managers have learned to be very suspicious about such words as they often reveal unexamined (and dangerous) assumptions.

PIP says there is nothing special about me. Therefore, there is no reason for me to be born at any particular point of an interval. If the earth's orbit and climate has been stable for over a hundred million years, as the fossil record seems to indicate, the chance of that changing during my lifetime seems insignificant.

PIP can be made into a precise statistical principle. If I assume that I am born at a random point of all intervals, then I expect everything to last as long into the future as it has into the past. When I was born, the American President, British Monarch and the male half of the Wimbledon Mixed Doubles Champions had each held their respec-

tive positions for four years. So, had I been precocious, I would have predicted that each would continue to reign for four years. Queen Elizabeth II surpassed her target by 46 years and counting, the tennis player did not last out the year and Dwight Eisenhower hit four years exactly.

In 1956, most people would have predicted correctly that Queen Elizabeth would outlast Dwight Eisenhower as head of state, and both would be in office when Vic Seixas lost a Mixed Doubles match at Wimbledon. Based on information available in 1956, I would estimate a 95% probability on the Queen reigning more than four years, using a mortality table and factoring in small probabilities of abdication, usurpation and abolition of the monarchy. I would give the President about a 90% shot of going exactly four years (he was not eligible for reelection in 1960). Thirteen of thirty-four Wimbledon Mixed Doubles winners from 1914 to 1956 had at least one player in common with the prior year's winners, implying about a 1% probability that the 1956 male champion would continue to win another four years.

These estimates could be further refined with better models, however they do not contradict PIP, which assigns all three events the same 50% probability. In a subtle sense, both probability estimates can be correct. Although there is controversy about the fundamental meaning of probability, all major factions agree that probability is meaningful only against some reference set of information, model or subjective prior belief.

Well, what do you know?

Consider, for example, a casino trying to compute the VaR of its blackjack tables. The house edge is determined mainly by the probability of the customer being dealt a blackjack, an Ace and a Ten or face card. Playing from a single deck (casinos use multiple decks which reduces the chance of blackjacks slightly) the probability of a blackjack on the first deal is 32/663. This implies a house edge of about 2%; the house has an expected profit of $2 for every $100 bet.

In blackjack, the deck is not shuffled after each deal, instead the remaining cards are used for the next deal. Suppose after a few deals 50 cards have been used, what is the probability that the last two cards form a blackjack? To someone who has seen the first 50 cards, the probability is either 1 or 0, the last two cards either form a blackjack or they do not. 32 times out of 663 the deck will come down to a blackjack on the last two cards, 631 times it will not. So before any cards are dealt, the probability that the last two cards form a blackjack is 32/663. Both probability estimates, 1 or 0 and 32/663, are correct relative to their information sets.

As the cards in the blackjack deck are played the house edge varies with a standard deviation of about 2.5%. A significant fraction of the time, the odds are actually in the customer's favor. "Card counters" take advantage of this by making the minimum bet when the odds are in the casino's favor, and the maximum bet otherwise. If a table has a minimum bet of $100 and a maximum bet of $1,000, a card counter will bet an average of about $300 and make a profit of about $3, 1%, per bet.

One version of PIP that is central to formal and informal risk management is the idea that there is nothing special about tomorrow

It is very important to the casino whether customers are card counting or not. If they are, the casino expects to lose $3 per bet, if they are not, the casino stands to make $6 per bet.

However, for VaR purposes, it does not matter much. If customers are card counters, the standard deviation of the casino's take on each bet is $484, if the customers do not count the standard deviation is $488. Over a million bets the casino expects to make $6 million if play is random or lose $3 million if customers count cards. But the 1% million-bet VaR is $1.1 million in either case. One percent of the time the casino will make $1.1 million less (or lose $1.1 million more) than the expected amount.

Normal is as normal does

This principle holds for financial markets. Sophisticated financial models that take into account lots of information are necessary for pricing, but risk management can usually be done using simple general principles and ignoring most market information. We tend to prefer simple models, like PIP, because they do not depend on complicated assumptions. They can give accurate results even if the pricing models are flawed (and, of course, flawed pricing models are one of the major financial risks). Moreover, using simple models allows us to focus on important systemic risks, which cannot be incorporated into specific models accurate enough for pricing.

One version of PIP that is central to formal and informal risk management is the idea that there is nothing special about tomorrow. Therefore, I can set an exact 99% one-day VaR on $1,000 worth of Microsoft stock by looking at the largest loss such a position experienced over the last 99 trading days. At this writing, that happens to be a $95 loss on July 23, when the stock went from $47.51 to $43.01. Tomorrow is as likely as any other day to be the worst of the hundred days up to and including tomorrow. Therefore, there is a 1% chance that the portfolio will lose more that $95 tomorrow.

If we instead assume Microsoft stock returns follow a Normal distribution with negligible daily mean, we get a 99% one-day VaR of $76, based on a standard deviation of $33, using the same 99 days of data. Both VaRs can be correct. Over long periods of time, both methods could give VaRs that were breached exactly 1% of the time, even if the two methods gave different VaRs every day. But the PIP method is guaranteed to work unless we have relevant information about tomorrow's return on Microsoft stock that makes it different from the past. The Normal method will only work if the PIP method works and if the

Common sense observation of cycles do not provide reliable risk estimates once we get beyond three or four independent factors

actual distribution of Microsoft stock returns is independently and identically Normal.

Long before VaR was invented, in fact long before any sort of quantitative risk management, financial professionals used a qualitative version of PIP to assess risk. The common sense way to set limits and reserves is to look at the distribution of hypothetical or actual losses in the past. This is a reliable method unless something unusual is occurring.

Armageddon, etcetera

Consider the end of the universe. Based on a 10 billion year history, PIP tells us the universe has a 99.999% chance of lasting 100,000 years. That is enough for most people to ignore the possibility, which undercuts the appeal of apocalyptic religion. The long-haired, long-bearded guy walking around with a sandwich board saying "Repent! The End is near!" is a cartoon cliché, a figure of fun, because his prediction violates PIP.

The Book of Revelation gets around this. It lists about three hundred signs that presage the End, things like the sun turning black, the moon turning red, the stars falling, a river of blood (about one hundred million people's worth) and a giant angel sitting on a cloud with a sharp sickle, while another angel yells "thrust in thy sickle and reap for the harvest of the earth is ripe." This kind of thing is clearly intended to trump PIP. If you see events like this, you have to throw out the assumption that current times are no different from any other. It is not that any of these things have an obvious causal relation to the end of time, it is that any unprecedented observations undercut the key assumption of PIP.

The power of this logic is illustrated by the 1999 movie *End of Days*. Rod Steiger (the mysterious priest) tells Arnold Schwarzenegger (the hero) that Gabriel Byrne (the devil) must father a child between 11 pm and midnight on December 31, 1999 to kick off the Apocalypse. Skeptical Arnie replies with an assertion of PIP. He asks, "Eastern Standard Time?" In other words, if New Year's Day 2000 in New York City is determined by arbitrary convention, it is unlikely to be a time of any cosmic importance. This is why your more cerebral end-of-the-world movies, such as *The Fifth Element*, rely on things like planetary alignments.

"Unusual" in this context depends on the time interval over which we are applying PIP. Over the history of the universe things like exploding stars, colliding galaxies and mass extinctions are commonplace. It would take something unexplainable by physical law to get most modern people to take seriously the idea that the universe will end soon. In financial risk management, we are typically assessing risk based on a few months or years of data. Over that period of time, historically commonplace events like recessions, elections, major bankruptcies and legal changes can make the application of PIP doubtful.

The common sense way to set limits and reserves is to look at the distribution of hypothetical or actual losses in the past

Not a spoke in the wheel?

The art of risk management is to integrate long-term information, which is necessarily qualitative and market specific, to general, precise, quantitative statistical principles like PIP based on short-term data. The Czech novelist Milan Kundera dealt with this dilemma in *The Unbearable Lightness of Being*. If time is cyclic, as Australian aborigines and German philosophers are supposed to think, then even the most inconsequential events acquire great significance, heaviness, from infinite repetition. If, however, time is linear, with each day unique and never to be repeated, then it seems as if nothing matters, being is unbearably light.

Individual markets are clearly cyclic. Stocks go up, stocks go down. Interest rates go up, interest rates go down. Oil prices go up, oil prices go down. Although the timing of the cycles is too erratic for useful prediction, at least according to mainstream theorists, risk managers can rely on the accumulated wisdom revealed by past cycles. Market specialists know the signs that a shift in statistical behavior is likely. If those signs are absent, PIP is a reliable principle. If the signs are present, longer term data, such as the distribution of peak to trough prices, can be used instead.

I think this is a deeper truth than might appear at first. I think markets are defined in such a way that the price behavior of individual securities is cyclic. In other words, securities are grouped into markets in such a way that their statistical behavior changes at the same times. By "market" I do not mean all securities that trade on the same exchange, although exchanges are often grouped the same way as markets, but that clearing and settlement procedures are designed such that combinations of securities can be traded almost as easily as individual securities. Thus I consider all global common equities make up one market, despite trading on many different exchanges, but preferred stock and bonds are different markets, although both trade on the New York Stock Exchange. When markets grow to include too many disparate factors, risk management becomes too difficult and they split.

Forget yesterday. You'd better worry about tomorrow

For a single security, PIP gives an exact VaR. The chance is $1/(n+1)$ that tomorrow will be worse than the worst of the previous n days. The worst day out of the last 99 is an exact one-day 99% VaR. You need n data points to set a $1 - 1/(n+1)$ VaR, each point must represent an interval equal to the interval of your VaR. So you need 200 days of data to set a ten-day 95% VaR.

But what if your position depends on the prices of two securities? What is the chance that tomorrow will be worse than the worst of the previous n days?

The answer is simple if the two prices move independently. For all i from 0 to n, there is $1/(n+1)$ chance that the movement of security A will be worse than exactly i of the previous n values for it. There is $1/(i+1)$ chance that the movement of security B will be worse than all i of its movements on those i days. Therefore the chance is the sum from $i = 0$ to n of $1/(i+1)^*(n+1)$. This is approximately $\log(n)/n$ so it takes about 647 trading days of data, about 2 years' worth, to set a PIP VaR on two independent market factors.

For k independent factors, the chance that tomorrow is worse than any of the n previous days is approximately $\log(n)k-1$ divided by $n \times (k-1)!$ (k must be small relative to n for this approximation to be valid). For one-day 99% VaR we need 13 years of data for three factors, 57 years for four, 237 for five, 944 for six and 3,616 for seven. Eight factors takes us beyond the limit of human history.

This does not just mean PIP does not work. It means that common sense observation of cycles do not provide reliable risk estimates once we get beyond three or four independent factors.
On the scale of human history, these multidimensional cycles never repeat. The oldest market veterans have not seen it all before, each day is truly new.

Things are much better if the factors are not independent. With dependent factors, the $1/(i+1)$ chance in the argument above is too high. Moreover, the computation assumes that all factors have equal economic risk and that participants hold completely arbitrary combinations of factors. In practice, if most of the profits and losses of most participants in a market can be reduced to a few principal components, then PIP and other common sense measures will work. There will be cycles and predictability. Being will be heavy.
However, the movement in finance has been toward connecting markets. Derivatives are traded with underlyings in different markets. New markets spring up that cut across traditional lines. Integrated financial institutions use one pool of capital to support risk of many different types of activities in many different markets. Hedge funds pursue highly levered cross market arbitrages.

The idea that markets can only go up or down is not reliable in multidimensional space

This is a fundamentally different world. The idea that markets can only go up or down is not reliable in multidimensional space. Data patterns that can be traced back before the Bronze Age cannot be extrapolated into the future. PIP is not a sound principle; the markets can change quickly, and without warning signs. Every day is new, there are no reruns. Being is light.

This is a massive experiment. Giving up the discipline of segregated markets of low dimensionality will have unpredictable effects. Certainly the character of financial crises seems to have changed in the last fifteen years, it is not yet clear whether for better or worse. We have thrown off the heaviness of repetition; we have no choice but to bear the lightness of being.