

Introduction

Finance is full of colourful stories and the most exciting ones tend to involve someone on the verge of collapse. We feel a mix of thrill and *schadenfreude* when we read about the traders who blew up or the elite hedge funds that had to liquidate after failing to meet their margin calls. In a moment of panic, investors can do the strangest things and this can make for great theatre. Arrogance and overconfidence are punished by the markets, which seem to have a life force of their own. Many shrewd investors have completely lost their way in a moment of crisis. There are numerous stories of portfolio managers who have patiently extracted profits from the markets for years, then had a large and unexpected loss. It might have been advisable for them to exit the position (“cutting their losses”) and try to claw back using their core strategy over time. Yet, the temptation is to put all the chips on black in an attempt to make the money back quickly. In principle, this is a wretched idea, as the profit from a long series of rational trades over time may be overwhelmed by a single irrational bet.

THE AIRPLANE TICKET TRADE

The legend of the airplane ticket trade is an extreme example of bad judgment under pressure, yet it is sometimes presented as rational decision-making. The story goes as follows. A trader has been losing money and is unlikely to collect much of a bonus this year. So the trader decides to dial up risk in an attempt to make it all back in one go. This backfires horribly, leading to further losses. The trader expects risk to be cut at any moment now, so he does two things. He makes a very large short-term trade that will either make or lose a large amount and he simultaneously buys a ticket to South America. It's a tactical play, with little edge but lots of risk. The trader then goes to the airport and repeatedly checks his price feed in the lounge. If the trade goes in his favour, he closes the position then goes back to the office. If it goes belly up, he buys a bottle of vodka from the duty free then takes the flight. The trader's behaviour might seem reasonable at 30,000 feet. In the best scenario, he gets a large bonus; in the worst, he takes a long tropical holiday. There doesn't seem to be much downside and one could argue that from the trader's standpoint, he is long an option. But would you want to be that trader at the moment of crisis? If the position is going slightly against you, are you willing to hang on for dear life, with no conviction that you are making the right trade? If it is your own money, do you want to risk everything on a roll of the dice? If you are a fund manager, how can you rationalise what you have done to clients if it all goes wrong?

THE BULL CYCLE

In reality, most institutional losses and disasters are not caused by trading reminiscent of the Wild West. Rather, they are caused by somewhat predictable behaviour through the market cycle. In bull markets, portfolio managers tend to increase exposure in an effort to chase the market and outperform competitors and benchmarks. Ten basis point differentials in performance seem important. By the “market”, we mean risky assets such as stocks and corporate bonds. Investors eagerly buy into every dip in the market, dampening volatility. As the value of collateral increases and volatility declines, banks lend more and the market eventually becomes overextended. This applies to equities, corporate bonds and other risky assets. When risky assets appear to be vectoring toward infinity, we would argue that it *is* a good time to hedge. Risk embedded in the system has increased, yet the market is practically giving away insurance. The painful memories of the last crash have been erased, making investors particularly vulnerable to a random shock.

Investors who chase returns after a large sustained move tend to have relatively low pain thresholds. They worry that they have missed the move, but are equally likely to bail out at the first sign of trouble. So long as the rally persists, the cost of insurance (i.e. options) tends to be low. The latecomers to the market do not want to erode their return by hedging and the longstanding bulls are complacent. You could sensibly argue that if the market continues to rally, hedging costs should be more than offset by profits in the rest of the portfolio. Yet there is a natural human reluctance to “waste” money on insurance when everything seems fine.

As the animal spirits take over, investors attempt to rationalise their behaviour in a variety of ways.

- “This time it’s different.” There is a central bank put on the market, as monetary conditions will be eased whenever there is a risk event. Regulators can prevent extreme intra-day moves by disqualifying trades that occur very far away from recent prices.
- Calm periods are persistent: they tend to last for a long time. Not very much happens from day to day, suggesting that there is plenty of time to prepare for the next correction.
- Over the long term, hedging is largely unnecessary. For example, some institutions don’t hedge their currency risk. Over the long term, they assume that currency moves will wash out. Buying insurance on risky assets such as equities is a losing strategy over the long term. According to academic theory, hedging must have a negative risk premium, as it reduces the non-diversifiable risks in your portfolio. Insurance companies are generally profitable because they sell individual policies that are statistically overpriced. So long as the policies are relatively uncorrelated, insurers are able to collect more than they pay out over the long term.

If you are not careful, you can convince yourself that *selling* insurance is an unbeatable strategy. Short volatility strategies tend to perform magnificently in back-tests, without much parameterisation. All you need to do is persistently sell downside protection on equity indices, risky currencies and corporate bonds, or so it would seem. When volatility is low, these options appear to be slightly but consistently overpriced. It is tempting to conclude that you can make small but very steady returns in this environment. As volatility rises, your profits become less reliable from day to day. However, this might be more than compensated for by an increase in the premium you collect when volatility is high. Most active management strategies are short volatility in one way or another. Whether you buy equities, take long positions in risky bonds or engage in spread trades, you will tend to perform better in flat to rising markets than highly volatile ones. The vast majority of hedge fund strategies are structurally short volatility. The incentive structures for many hedge funds and proprietary trading desks favour collecting pennies in front of the bulldozer. However, this does not imply that selling volatility universally has a positive expected return. Once you put a back test into action, you are vulnerable to large jumps that may not have appeared in the sample past. As soon as you introduce leverage, you are vulnerable to risk and margin constraints that can force you out of a trade at the worst possible time. Markets don’t usually collapse because investors want to sell, but because they *have* to. Liquidation is forced, in the presence of margin calls. We will examine the effect of margin constraints on short volatility strategies in Chapter 4.

THE RENEGADES

There is a small but dedicated group of defensive, bear market managers in the investment universe. The financial media trots them out every so often, typically after a market sell-off. However, in rising markets these managers are largely invisible or the subject of criticism. Profiting from panics, bankruptcies and liquidations requires patience and does not necessarily win you many friends. When equities are ramping up, bear-biased managers spend more time banging their heads against the wall than raising assets. The cost of insurance is steadily declining, yet there are no takers. The inveterate bears write long and engaging manifestos in an attempt to identify cracks in the financial system. In rising markets, the potential end users of these products generally can't or don't want to buy them. Some institutions take a crude "line item" approach, where they rank their funds according to recent performance and periodically redeem from underperforming managers. This approach seems oblivious to the idea of marginal risk, i.e. how much you can improve the risk-adjusted performance of an existing portfolio by adding a new asset or strategy. In reality, if you can find a strategy that performs strongly during crises yet doesn't lose too much over a market cycle, it can have a dramatic impact on portfolio performance over the long term.

Uncontaminated bear strategies have a hard time competing in a world where allocators believe that emerging markets, high yield bonds and carry trades are "diversifying" investments. While it is true that these asset classes can reduce realised volatility during normal market conditions, they typically amplify losses when conditions become extreme. Some strategies, such as the FX carry trade, seem innocuous during bull markets. They grind their way upward with low volatility. However, it is categorically *not* true that a strategy with relatively low volatility in a bull market will dampen risk during a crisis. If the strategy collects premium while taking extreme event risk, the opposite is in fact true. A manager who combines carry strategies with a modest number of equity index puts will often appear to be over-hedged most of the time and severely under-hedged when the protection is most needed.

In rising markets, dedicated bears have to overcome time decay as well as markets that are moving in the wrong direction. The portfolio manager who takes the opposite side of the trade by selling insurance has an optical advantage. Investors seem to prefer a sequence of returns of the form $\{+1\%, +1\%, +1\%, +1\%, +1\%, -5\%\}$ to $\{-1\%, -1\%, -1\%, -1\%, -1\%, +5\%\}$, even though the compounded return of the second strategy is a bit higher. In the first scenario, you can always say to your client that you are an alpha manager who had a few issues with risk control that have now been resolved. This cynical approach may well salvage the mandate. Even the most dedicated bears are incentivised to scale down their hedges when threatened with redemptions.

The best time to buy outright volatility is when it is low, in a counter-cyclical way. You want to swim against the tide of short-sighted overconfidence. Investors are more than happy to sell volatility when they are feeling confident. However, implied volatility is low precisely because there is virtually no demand for hedging or long volatility strategies in general. Hence, long volatility managers struggle to raise assets in situations when the best risk-adjusted returns are available. Our book acknowledges the perverse nature of hedging mandates. When assets are pouring in, outright volatility tends to be overpriced. We try to identify ways to minimise drag while still offering protection after markets have started to tumble.

CLAWS OF THE BEAR

[T]o borrow the term, your sense of time does change when you are running real money. Suppose you look at a cumulative return of a strategy with a Sharpe ration of 0.7 and see a three year period with poor performance. It does not phase you one drop. You go: "Oh, look, that happened in 1973, but it came back by 1976, and that's what a 0.7 Sharpe ratio does." But living through those periods takes – subjectively, and in wear and tear on your internal organs – many times the actual time it really lasts. If you have a three year period where

something doesn't work, it ages you a decade. You face an immense pressure to change your models, you have bosses and clients who lose faith, and I cannot explain the amount of discipline you need.

– Cliff Asnessⁱ

Once you put real money behind a short volatility strategy, the situation changes. Now you have some skin in the game and things aren't quite so comfortable. Your margin levels can change dramatically over time, requiring that you cut positions that look very attractive from a valuation standpoint. In Chapter 4, we show that wildly fluctuating margin requirements can force you out of a short volatility strategy at the worst possible moment. A historical series of daily NAVs is devoid of emotion and assumes that you have sufficient capital to keep playing indefinitely. It can't capture gut wrenching intraday moves or account for price action that is different from what has been observed in the past. If the worst 1 day historical loss is –10% and your strategy is down –9% at mid-day, there is no guarantee that losses will be bounded at roughly –1% thereafter. In rising markets, investors are quite happy to sweep latent risk under the carpet as risk and margin limits are never reached. Inevitably, at some point, risky assets take a significant leg down. The “stocks go up in the long term” bulls can no longer buy the dip as they approach their risk limits. Large institutions spend ages deciding whether “this is the one”, whether credit and equity markets will plunge further into the abyss. Their portfolios might already be down –5% or –10% on an unlevered basis and they really can't take much more. Do they hang on, cut exposure or hedge?

It has often been remarked that “hope is not an investment strategy”. Hanging on is a sign of desperation or delusion. Sometimes, an overconfident investor can become convinced that the market *has* to move a certain way and goes all in. It is almost as though the investor believes it is possible to move the market by force of mind. Solipsism doesn't seem to be a viable strategy, either. Some investors doggedly hold onto losing positions using “fair value” arguments. When combined with leverage, this approach can be toxic. The standard argument is that the expected return of a static portfolio goes up as its price drops, i.e. price and expected return move inversely. While this may be true over long horizons, there is a point at which every institutional manager has to cut risk. Most of us do not have an infinite investment horizon in which to capture a risk premium. There is a saying for the leveraged deep value investors who hang on during crises: “it looks good at 90, looks great at 80, looks absolutely fantastic at 70 and you're out of business at 60”. This is the classic value trap that needs to be avoided.

ZUGZWANG

In chess, zugzwang refers to a situation where a player has to move, but every move worsens the player's position. When a portfolio manager's risk limits are hit or losses are thought to be unacceptable, the situation is quite the same. There are two choices: cut risk or buy insurance. Neither seems appealing. If the manager slashes positions, the potential for further losses is reduced. This can be agonising for investors who believe that, given enough time, their portfolio is bound to bounce back. Some portfolios are large and complex, implying that they cannot be liquidated in one go. Finally, suppose an investor has been making small bets for years and now has to divest a large percentage of his or her portfolio. This one action can offset a large number of good decisions and successful trades. Some funds scale in and out of positions almost continuously as risk changes. They generally have sophisticated techniques for sampling volatility and correlation over time. However, even these funds are exposed to gap risk (i.e. when a currency peg is released) or situations where their alpha-generation systems have stopped working.

Faced with the choice of liquidating positions or hedging, institutions finally pick up the phone and contact managers who can protect capital during a crisis. Managed accounts that have not been used for months are reactivated, with a hedging overlay mandate. Assets begin to flow into bear-biased strategies. As the demand for hedging increases, its cost sky rockets. To a patient on the operating table

in a life-or-death situation, money is no object. Survival is all that matters. And so it goes for an individual or institution on financial life support, who hedges regardless of cost. The long volatility manager who gets the call is in two minds about it. On the one hand, the manager is more than happy to have a new allocation. It serves as vindication, as well as a new source of fees. On the other, hedging looks expensive now. If only the call had come a few weeks earlier, when there was a wide range of inexpensive hedges to choose from! Previously, an overlay could have been slowly and carefully constructed, with an emphasis on finding inexpensive hedges across a variety of asset classes. Now it's a case of making the best of a bad situation. You have to make sure that the patient survives (i.e. that there is a floor on further losses), while ensuring that you don't spend too much along the way. Once markets recover, your performance will be mercilessly scrutinised. Did you make enough on the way down? Did you monetise enough gains to avoid giving it all back during the recovery?

Whether you allocate to another manager or hedge yourself, the pressures are quite the same. Most of the time, you will be incentivised *not* to hedge, even when you can identify good short opportunities. Indirect hedges, such as buying calls on the VIX to hedge against long exposure to the S&P 500, will generally add to your exchange margin requirements. This reduces the degree to which you can lever the rest of your portfolio. Even if leverage is not an issue, hedging suffers from an optical standpoint. Unless you can bury your hedges in the rest of your portfolio, your supervisors and clients will see long strings of mildly negative returns punctuated by the occasional lumpy positive one. Once things get ugly, you will be asked whether you have hedged *enough*. Are you making money on every little drop in risk assets? Have you put a floor on how much can be lost in the overall portfolio?

THE SCEPTICS

Some investors, especially those with a “stocks for the long run” bias (e.g. Siegel, 1998), might argue that hedging is intrinsically wasteful. The hedging sceptics tend to intersect with the true believers in the equity risk premium. If you are prepared to wait long enough, there's no need to hedge, as equity market returns will exceed inflation. Over rolling 10-year horizons, the S&P 500 has nearly always outperformed CPI inflation on an annualised basis. It follows that, if equities deliver a positive real return over the long term, hedging *must* have a negative risk premium. After all, you are paying a premium to take a short position on the market, is that not so? Theory suggests that you earn a premium for bearing an undiversifiable risk. Conversely, an instrument that offsets market risk should have a negative expected return. Insurance companies are in business precisely because insurance is overpriced on average. Historical back-tests in the markets tend to support this idea. Insurance eats into your long-term expected return. Static options hedges tend to lose money at an alarming rate, with modestly positive spikes along the way. On paper, the appropriate strategy involves *buying into* market sell-offs, as risk premiums go up whenever the prices of risk assets go down. It would seem as though the last thing you want to do is buy options *after* volatility has gone up. If an option was expensive before, it must be egregious after a risk event. Our view is that listed options are somewhat different from insurance policies. While typical hedges are probably overpriced, as there is excess institutional demand for them, options are subject to the same cycles of greed and fear as equity markets.

A SAD TRUTH

Recently, a number of books and articles have appeared covering topics such as “tail risk protection”, “crisis alpha” or “extreme event hedging”. Many of these are thorough treatments of how institutions think about truncating the left tail of their return distribution. Bhansali (2014) is a thoughtful treatise on the nature of asset class distributions and institutional quality hedging strategies. However, they invariably ignore a sad truth. Almost no one wants to hedge much when the going is good. Institutional investors generally do not pay much attention to the independent economists and hedge fund managers

who warn that a new crisis is brewing. In bull markets, articles that focus on doomsday scenarios are viewed as nothing more than fearmongering. Indeed, it is notoriously difficult to predict where the next crisis will come from. Will it be credit derivatives, emerging markets or a change in Central Bank policy?

Several well-known hedge fund managers try to engage in crisis prediction by identifying potential cracks in the system. They typically screen for excessive leverage in some part of the economy and then direct their hedges to the places where danger seems to be lurking. This is a substantial improvement on not hedging at all, but it assumes that extreme events are predictable in place *and* time. If the manager places the doomsday bet too early, there may be a long string of losses before any material gain is realised. In the meantime, investors might redeem from the strategy. If the bet is placed too late, the risk of default may already be priced into the market, reducing potential returns. Most of us don't have the foggiest idea when the next crisis is coming and should be honest enough to admit it. Note that we *will* discuss crisis prediction in Chapter 8. It might seem contradictory that we are taking a stab at a problem as difficult as this. For the purposes of this discussion, however, it is best to assume that predicting financial crises is like predicting earthquakes. We can identify situations (geological fault lines) which are unstable, but can't with any certainty say *when* an event will occur.

Returning to the original problem, let us generalise and assume that investors only want to hedge after risk assets have taken a leg down. Hedging is not going to be cheap, as there is more demand for insurance. So what can you do to protect a portfolio against a systemic risk event, that isn't *too* bad? That is what this book is all about.

COMMON MISTAKES

In the chapters that follow, we identify strategies for protecting a portfolio of risky assets *after* a sell-off. Investors have suddenly become wary and are no longer just giving away protection at discount levels. It is not wise to just go in and buy index puts, as these are bound to be overpriced. Yet many institutions do exactly that. They react to an increase in perceived risk by identifying "plausible" downside scenarios and choosing options that target those scenarios. The risk committee might have a discussion about how bad things could get, before reaching a consensus on what constitutes a tolerable and plausible loss. We believe that this approach is flawed. While it is reasonable to average forecasted returns, taking an average of downside scenarios understates the risk of an extreme event.

If everyone is buying options to cover the risk of moderate losses, those options are likely to be overpriced. Our approach is to find other options to buy. We argue that an option does not have to wind up in the money to be profitable. All that is needed is a *repricing of risk*. Just as the price of hurricane insurance goes up when there is a thunderstorm, the price of extreme event insurance rises when there is a moderate sell-off in the equity market. You can always sell an option back to the market if it reprices substantially. In any case, implausible scenarios can appear plausible after a plausible scenario has occurred. This may sound thoroughly convoluted, but it is not meant to be. Our goal is to be as clear as possible. At first sight, a -30% one month collapse in the S&P 500 seems highly unlikely. Even in October 2008, the peak to trough drop was less than that. But suppose that the index drops -10% in the first week. Suddenly, that -30% drop does not seem so unlikely and investors are clamouring for insurance at levels (i.e. option strikes) far below what could be imagined. This is partly a function of perception. It is also based on the idea that in certain scenarios, markets exhibit positive feedback. A drop cannot be viewed in isolation, because that drop may force others to sell as they hit their risk limits. Sell-offs can occur in cascades.

Another common idea is to hedge extreme event risk using currency options. This is a play on the "Mrs. Watanabe trade", which will be analysed in greater depth in Chapter 2. Mr. Watanabe has a demanding job, so he delegates the family's personal investments to his wife. There is no point in depositing money at a local bank, since the bank rate is effectively 0%. The Nikkei 225 is still well below its peak in 1989 and there is no cult of equities in Japan, as there is in the US. The equity premium puzzle is irrelevant, as there is no premium to speak of. Investors are distrustful that Japanese equity indices will deliver a positive return over the long term. So why not sell the Yen to buy Australian dollars

(AUD) or another high yielding currency? When you buy Australian dollar forwards, you implicitly capture the Australian bank deposit rate. If a 1 year deposit yields 5% in Australia, you gain 5% carry in Australian dollars while borrowing Japanese Yen virtually for free. Theory suggests that the forward rate bias should be offset by an expected -5% annual decline in AUD. In practice, Kritzman (1999) and others have observed that spot exchange rates are not very correlated to yield differentials. If anything, high yielding currencies tend to outperform even in spot terms, as investors chase income. Assuming that AUD does hold its value relative to the Yen, you collect nearly 5% per year. This is a huge source of income in a deflationary environment. Once the carry trade gathers momentum, Mrs. Watanabe's investment club piles into the trade. Things sound rosy so far. Ultimately, the trouble with the trade is that it can become overleveraged and overcrowded. This poisonous combination can cause very steep declines when investors are exposed the most. Eventually, a random shock turns into a major reversal as the investment club heads screaming for the exits along with larger institutions.

Given that carry currencies go up the stairs and down the lift, it seems reasonable to hedge extreme event risk using puts on carry trades such as the Australian Dollar/Japanese Yen cross. More precisely, carry trades are negatively skewed, implying that the probability of downside surprises is higher than the probability of ones that work in your favour. The longer-dated the put, the more time you have to wait for a blow up. However, the shape of the currency forward curve can have a dramatic impact on the performance of a hedge. Once you buy the Aussie put, two forces are conspiring against you. Every day that passes, you lose money on time decay *and* drift as the forward rolls up to the spot. So the currency hedge usually winds up a loser.

IMPRECISE BUT EFFECTIVE

The hedges described in this book are *not* precise. We are not going to tell you to hedge your long position in Apple with Apple puts. While this may be the most accurate way to soften the impact of sell-offs, it is often egregiously expensive. If you really don't feel comfortable with Apple downside, your best strategy is probably to reduce the position. Yet institutional practice often suffers from a literal and somewhat narrow-minded approach to hedging. Many consultants in the pension fund industry seem overly focused on precise hedges. The solvency of a pension fund is often calculated relative to the *present value* of its liabilities. The liabilities are the expected payments the fund will have to make to its beneficiaries in the future. If interest rates drop, the value of those liabilities *today* will increase. This necessitates hedging against rate risk. A common solution is to use an actuarial number called the "average duration" of the liabilities as a crucial variable in the hedge. This is the time-weighted average of expected payments to beneficiaries. It is an imprecise number, based on projections of who will retire and when. Still, consultants often think of average duration as an exact number and offset rate risk with swaps or other instruments that precisely target it. While there may be regulatory reasons for transacting in this way, this approach seems wasteful. In a low interest rate environment, receiving fixed payments does not seem like the best idea. There may be other, slightly less precise-looking hedges that offer greater protection and are likely to cost less in the long term. If your shoe size is 8.5, do you want or need a shoe that is precisely calibrated to size 8?

We focus on overlay strategies as a mechanism for controlling portfolio risk. Our goal is to identify areas where insurance is relatively inexpensive, while recalling the idea that our hedges need to make money in a severe risk event. Most of the strategies involve exchange-traded options in deep and liquid markets, such as equity indices, the VIX, interest rates and currencies. We start with a non-technical overview of options theory. The goal is *not* to price options but rather to express the value of an option in terms of its implied volatility. Focusing on volatility, we can identify different combinations of options in different markets that are relatively cheap at a given point of time.

If an emerging markets index drops from 100 to 90 in a day, a 75 strike put may have a larger percentage gain than a 95 strike put, even though the stock hasn't come close to 75. The implausible scenario (a sudden -25% drop) has suddenly become plausible. As we will see, large moves sometimes beget even larger moves as the market enters a positive feedback loop. When investors start to worry

about a major loss, they bid up the prices of puts that are far out of the money. This causes those silly strike options to make multiples of what was initially paid.

HEDGING AGAINST IMPLAUSIBLE SCENARIOS

In Chapter 4, we explore hedges that take advantage of changing investor perceptions. One structure revolves around buying large quantities of options that only pay out for moves that appear ridiculously unlikely. We emphasise that we are not really concerned whether these options have value at maturity. When fear grips the market, investors re-price extreme event risk and the options can explode in price. It's a bit like betting on a long shot team in a football tournament. If the odds are 1000:1 at the beginning of the tournament and the team surprisingly wins the first game, the odds might drop to 500:1. The team is still unlikely to win the tournament, but if you exit now you have doubled your money. You don't need the team to actually win the rest of its games and can move on to other long shot bets over time. While sports betting odds are clearly skewed toward the bookies, financial markets offer a more level playing field. Certain options can become surprisingly cheap when investors are complacent. There also seem to be structural distortions in the options markets that reduce the cost of hedging. Our analysis suggests that investors like to hedge against "reasonable" looking downside scenarios over moderately long-time horizons. They wind up overpaying for options that pay out for moves that seem awful now, but won't account for much in a severe bear market. One strategy is to sell a few of those, while buying a large quantity of options that only pay out if there is a financial meltdown. This approach can work well for the VIX as well as for more conventional assets, as we will find out in Chapter 4.

A BLACK SWAN IN CORRELATION

Over the past 40 years, market crises have followed a familiar script. Stocks and risky bonds have sold off together, with a correlation close to +100%. US Treasuries have soared in the opposite direction, as investors have scrambled for safe havens. The market has operated in the "crash correlation" zone initially described by Hua (1997), with extreme diversification offered by a well-known collection of government bonds. So long as the familiar relationships hold under pressure, you can reduce aggregate risk by sprinkling some US or UK government bonds into your overall portfolio. Markowitz theory will work in your favour. But what if the next crisis challenges the core assumption that Treasuries offer protection against an equity collapse? This may sound like science fiction, but it may not be so far away as some might think. If money pours into products that rely upon asset class diversification, we could see a large-scale liquidation where stocks and bonds go down together. In this case, volatility hedging will be the last chance saloon. Options and volatility indices will become the only assets that can be relied upon for true protection in a crisis. This is one of the main reasons why we focus on options structures in Chapters 3, 4 and 5. Convex hedges are more reliable than diversifying assets. You are not at the mercy of historical relationships that may not hold in the future.

TAKING PROFITS

Suppose that gold prices are falling and you fear a further downside move. You can short gold futures or buy puts on the futures contract. The put strategy offers the advantage of reducing entry timing risk. If gold suddenly reverses, your put will dynamically de-lever, capping your maximum loss. You can only lose the premium you paid. However, losses on the short futures position are in principle unbounded until you cover the position. While the entry point is relatively unimportant, timing the exit from a winning long options trade is vital. If you buy an option and you find that its value has doubled, you wind up holding a much larger position than you started with. The gearing in your option has increased as the market moved in your favour.

In the pages that follow, we argue that the best way to manage a hedging overlay is to rotate across different strategies over time. You always want to have *some* kind of a defensive trade on, but you don't want all of your hedging profits to vanish if markets squeeze upward. Some profits have to be monetised. Otherwise, hedging will improve the mark-to-market performance of your portfolio during crisis periods, but you will never realise any profits from the hedge. If you unwind the entire hedge at some point, you might realise some gains. However, the rest of the portfolio will now be exposed to a further sell-off. So you need to walk the tightrope between maintaining enough protection whilst not holding on to hedges that are now overpriced. "Every dog has its day" is one of the many overused slogans from the finance community, but it succinctly characterises options hedging. No options strategy works all the time and it could equally well be said that, under the right conditions, nearly every hedge can work well. Behind each option is a trading strategy that gears dynamically over time. If the S&P 500 is trading at 2000 and you buy a put with strike 2000 (roughly at-the-money), your position grows as the index drops. You are long downside convexity: profits slingshot in tandem with the market. When you combine different options (e.g. buying the 2000 strike and selling the 1900 strike with a month to go), the return profile becomes more complex.

In Chapter 4, we analyse the back-tested performance of options with different deltas over time. This allows us to create hedges where we sell a small number of persistently "bad" options while overbuying the ones that don't look so bad. We explore the analogy between hedging and market making and demonstrate that buying extreme event insurance is worthwhile even if it never pays out at maturity. Hedging with the VIX is particularly intriguing, as VIX options do not burn off at the same rate as equity index options. They stay lively close to maturity. A VIX call is always in play, as any spike in volatility has a disproportionate impact on short-dated options. We discuss how to overcome time decay and futures roll down while building long options structures on the VIX.

THE GOOD, THE BAD AND THE UGLY

We recall that different types of hedges work in different market regimes: the "every dog has its day" argument. But how do we characterise the prevailing regime in a reasonably precise way? When can we say that global markets are calm and when can we say they are in a state of abject fear? Our approach is to use volatility indices, such as the VIX, as a guide. When the VIX is low, our analysis favours value-buying of volatility. As it rises, we transition to relative-value hedges in various markets. At the extreme, we recommend options combinations that provide significant payouts without too much exposure to volatility. We also explore the merits of trend following as a portfolio protection strategy during a crisis. Of course, this requires a definition of what is meant by low and high volatility. We delve into this question in Chapter 5, toggling back and forth between different types of hedges as conditions change.

This allows us to overcome the question of *when* to take profits in a hedging strategy. So long as the client wants a hedge in place, there is always something we can do. As conditions worsen, we simply rotate out of strategies whose cost is very sensitive to volatility, into other types of hedges. We emphasise that, for options that do not have a long time to maturity, stop losses are *not* an alternative to strategy rotation. Option prices can move radically from one day to the next, as volatility and price change for the underlying asset. They can quite easily crash through any internal profit-taking level you might have set or any price you may have flagged to the market.

THE GREAT ESCAPE

Suppose you are a portfolio manager who applies quite a bit of leverage, e.g. using derivatives. You have sold a large number of puts on an emerging markets currency, based on the idea that volatility is overpriced and the currency is likely to drift higher. Suddenly, there is an air pocket move of -5% down. Your portfolio has lost -10% and you need to meet a margin call. The situation is dire and you

feel sick. The natural thing to do is liquidate the position. This frees up margin and caps your loss. However, if you still like the trade, liquidating eliminates the possibility of making anything back. Cutting a bit allows you to stay in the game, but is sort of a halfway house between liquidating and doing nothing, while praying that the market comes back. If you cut the minimum amount, another down move could trigger a second margin call, forcing you to liquidate at an even worse level.

One way around this problem is to buy very short-dated puts on the currency. This allows you to hang on, in a responsible way. The short-dated puts act like a giant piece of duct tape on your portfolio. They contain your losses for a few days, while you regroup. As we will discuss later, you pay a small up-front cost that is not overly sensitive to the level of fear in the market. The gaping hole in your portfolio may be reopened when the options expire. Chapter 5 describes ways to use weekly options as a survival tactic. Weekly options are usually presented as either yield enhancers or lottery tickets. Since short-dated options decay rapidly as function of premium paid, you can capture a huge amount of yield by selling and reloading every Friday. As we will discuss later in the book, we believe that selling weekly options on an outright basis is an awful strategy. For now, we simply accept that many investors do precisely that, which may create distortions in the market. With the lottery ticket approach, you buy short-dated options in anticipation that something big is about to happen. That something may be a central bank announcement, an earnings report or anything else that could have a large impact on the price of an asset. The trouble with buying short-dated options in advance of announcements is that market makers will usually have adjusted the implied volatility of these options to protect against event-driven price discontinuities. You will have to pay up for the lottery ticket in this scenario.

But there is another, vastly more important, way to use weekly options. In the context of a hedging strategy, those weeklies let you live to fight another day. You can put a floor on potential losses in your portfolio at low fixed cost. Although the floor *will* disappear after a few days, it gives you time to reshuffle the rest of your portfolio at a measured pace. You can take that walk around the block before revisiting your positions with some clarity. Weekly option strategies are difficult to back-test, so we take a different approach to justifying their utility. In particular, we review some of the early econophysics research, which tabulated the distribution of stock index returns over various timescales. The beauty of this research is its naivety. It is not burdened by assumptions from traditional economics. We focus particularly on research from the Stanley group at Boston University in the 1990s. While the cutting edge has moved far beyond the early studies, the essential qualitative features of financial time series have not changed that much over time. The Stanley group simply tabulated large quantities of data in search of power law distributions, i.e. situations where the index returns had fat tails. They managed to show that returns over very short horizons (i.e. less than a week) had the greatest potential for outsized returns.

Over short horizons, stock index returns tended to have relatively high serial correlation, suggesting that there could be waves of buying or selling that were not interrupted by value investors. Assuming that the tails are not adequately accounted for in the market, weekly options can offer interesting value-buying opportunities as well as a blanket hedge on your portfolio. Chapter 5 concludes with a discussion of ways you can enhance the return stream of a long weekly options strategy, by trading the underlying as it approaches the strike.

HAVING A PLAN

Money managers typically focus on maximising risk-adjusted portfolio returns. From a psychological standpoint, though, avoiding regret is equally important. You don't want to fall prey to silly mistakes that could have been avoided. Having a plan, a strategy for dealing with any market eventuality, is vital. By definition, systematic funds have a plan for entering and exiting positions. Quoting Yeat's epitaph, models are able "to cast a cold eye on life and death" in the markets. Systematic trading strategies are typically based on data aggregated over a long time period. This implies that they cannot focus

acutely on current market conditions. However, they enable you to make decisions in an unemotional way, when asset prices are whipping around. In Chapter 4, we start with a digression on back-testing and model development. We review some research suggesting that fairly simple algorithms can outperform human experts when analysing certain types of data. We also discuss common pitfalls when developing trading strategies, including overfitting and the impact of selecting a strategy from a large number of alternatives.

TREND FOLLOWING AS A DEFENSIVE STRATEGY

One of the nice things about weekly options is that you can be *reactive*. You can wait until things get ugly before taking action. Hedging costs are not very sensitive to changes in volatility and the odds of an extreme event tend to increase after an initial spike in volatility. Most large-scale moves are preceded by choppy price action. You can just plug in a far out-of-the-money short-dated option whenever your position looks vulnerable. Let's take things a bit further. Perhaps the ultimate reactive strategy is trend following, where you are always chasing the market in the direction of an initial move. To clarify, trend following is a style of trading that relies upon buying assets that have already been going up and selling ones that are already moving down. Rather than following the value investing mantra "buy low and sell high", trend followers try to "buy high and sell higher". They make no attempt to identify turning points in the market. Instead, they prey upon situations where there are large-scale price moves, as institutions and leveraged investors reposition themselves. By definition, a pure trend follower wouldn't take a short position in S&P futures until they have already started to drop and will not cover the short until there is a reversal. Since trend followers typically focus on futures, they don't have to pay up for options after volatility has spiked. In high volatility regimes, trend following is a cheap alternative to a long volatility strategy *if* it can provide adequate protection. We examine this issue in Chapter 6.

A vaguely philosophical statement might be in order. It is hard to deduce *why* anything happens in financial markets. The newspaper articles about last week's move tend to be rationalisations, rather than accurate explanations. It's quite amusing to look at last month's, or last year's, research reports, where the recent move was extrapolated in an exaggerated way. The S&P 500 has dropped from 2000 to 1800 and the prophets of doom have come out in force, predicting a move to 1500 and below. In this book, we try to veer away from the financial entertainment industry (media headlines and so forth) and focus on ways that practitioners think about markets. Nevertheless, it is possible to speculate about the mechanisms behind market action. Credit and leverage play a larger role than is commonly recognised and it may be that trends are caused by predictable changes in gearing over time. Let's say you are a speculator who buys and sells commodity futures. You use a lot of leverage in an attempt to goose your returns. If the market is going your way, your credit situation automatically improves, because you can apply your profits to the margin account. This allows you to scale up your position. So if you had bought cotton futures, you can buy some more without damaging your margin situation. Conversely, if a long position shifts from a winner to a loser, at some point you have to sell. There is some threshold at which you will be wiped out and a nearer threshold which can turn you into a nervous wreck.

If the volume in a given asset class is dominated by leveraged speculators, trends are likely to emerge. The speculators have to manage their margin by following the policy of "cutting their losses and letting their profits run". This is tantamount to following trends. It is not always that the specs want to be trend followers. Rather, their style of trading demands it. So trends and ultimately bubbles form when the amount of leverage applied in a given direction increases. In Chapter 6, we take our analysis of trends one step further. Does trend following tend to generate positive returns during market crises? As we will see, the issue is quite complex. Historical studies show strong performance for trend followers during periods of market distress. This is called "crisis alpha" in the recent literature.

However, the relationship between the returns generated by a trend following algorithm and a rolling options structure is at best tenuous. We also take the direct approach to modelling trend following returns as a function of volatility. In particular, we will build our own basic trend signal and see how the signal performs in various markets as the VIX changes. It turns out that trend following probably has a mild correlation to volatility. While we can't rely on it as a hedge, trend following seems to offer real diversification benefits during adverse market conditions. Since the correlation to risk assets does *not* increase when volatility rises, it is a welcome addition to traditional portfolios.

TAKING THE OFFENSIVE

That's probably the best way of making money, to be a specialist in panics. Whenever there's panic hanging in the air, that's a great time to invest.

– Victor Niederhofferⁱⁱ

It might seem unusual to take investment advice from a manager who suffered damaging losses on two documented occasions. However, the basic idea is a sound one. Risk premia *do* go up after a severe sell-off. You can make more money from a piece of credit per unit of leverage than before. The prospective return for equities is inversely dependent on price. There is a reason why high yield bonds tend to outperform over time: only the most diehard investors are left holding distressed paper. The important thing to understand is that you can *only* invest aggressively after a panic if you haven't overreached beforehand or have a strategy for protecting against a further sell-off. In our view, the strategies developed in Chapters 4 and 5 allow you to snap up investment bargains without fear of blowing up. If you *are* going to specialise in panics, you not only need to know what to buy but how to hedge against disaster. In Chapter 6, we introduce some contrarian strategies for extracting alpha in a jittery market, while reminding the reader that it is necessary to size and hedge these strategies appropriately.

THE PRE-CONDITIONS FOR MARKET CRISES

There is of course, a larger issue. If we could solve it completely, the rest of this book would be largely unnecessary. Is it possible to *predict* the timing of market crises? Our view is that it is possible to identify conditions in the market that increase the odds of a crisis. Those are the limits of prediction. However, getting the timing right is nearly impossible. In the sciences, you can conduct experiments under controlled conditions. In the markets, you can't. A scientific idea can be zany, far from the mainstream, yet will be accepted if confirmed by experiment. In the markets, you need to get the aggregate of investors to agree with you in a reasonable amount of time. In 2008, it is probable that the portfolio managers who first predicted the mortgage-backed securities crisis made less money than other managers who jumped on the bandwagon at the last minute. The early buyers of default insurance were forced to pay a premium for many months before they were vindicated. The best one can do is identify situations that court disaster and hedging structures that have a large bang for the buck, i.e. that offer large payouts with low time decay.

At its core, the volatility cycle is intimately tied to the credit cycle. We will examine this idea further in Chapter 8. As banks lend more, investors become more vulnerable to a market shock. They have debts to pay off. So long as the random shock does not occur, volatility will tend to decline. Yet the system is becoming increasingly fragile. In Chapter 2, we argue that when investors scramble for the exits, seemingly "safe" assets can go down as much as risky ones. This was the case when quantitative equity funds were forced to unwind in August 2007. At the height of the 2008 crisis, investors generally sold whatever they had, including stocks that were usually considered to be defensive.

BANKS: THE GREAT MULTIPLIER

In the chapter on crisis prediction, we observe that the banking failures are responsible for most financial crises. If you think of the financial system as a network with agents at various nodes, the banking nodes are perhaps the most important. Once they are removed, the network collapses. Banks are the great multiplier in the economy. With a small quantity of deposits, they can lend a large amount of money. All of that lending goes into financial assets and the real economy, stimulating growth. The peak of the market cycle is generally characterised by a perverse relationship between volatility and leverage. Volatility tends to be low, as investors are complacent about the near future. The value of collateral (equities, real estate, etc.) has risen, so investors are able to apply more leverage to their overall portfolios. This implies that the risk in the system, the potential for a future collapse, is rising, while observed volatility is declining. Margin requirements at exchanges and prime brokers are low, enabling larger position sizes. Equity indices outperform equity hedge funds, leading some managers to chase the market and cut their hedge.

In order to keep lending, however, banks need access to a revolving credit line from investors and ultimately, from Central Banks. If a bank is overleveraged, the market may react by demanding a higher yield on loans. At some point, the bank has to repair its balance sheet by slashing some of its assets. Every loan that is not renewed has a knock-on effect, as corporations are unable to lend. When you take out a loan, you don't let the money just sit there, accruing negative interest. You put it to use, buying a house or investing in financial assets. Borrowing increases consumption, stimulating growth and increasing the value of risky assets. The amount of credit in the system is more important than the amount of money and deposits, as it is put to use in the economy. In Chapter 8, we develop this narrative more fully. In a non-rigorous way, we describe how bubbles form and burst over the market cycle. We take a multi-disciplinary approach, analysing the fundamentals and price dynamics that lead to market crises.

A CHANGE IN RISK REGIME

Some investors try to anticipate what could go wrong, what could cause a risk event. While many crises seem obvious in retrospect, the triggers take a long time to surface, are thought to be unimportant at the time or come as a complete surprise to investors. Many things *could* go wrong at a given time and one could argue that the most obvious ones are already priced into the market. Trying to understand “why” risk assets dropped on a given day is a game best left to market commentators. While text-based analysis of financial articles can shed some light on investor sentiment, we have deeper and darker problems to solve in the chapters that follow.

ENDNOTES

- 1 Pedersen, Lasse Heje. *Efficiently Inefficient: How Smart Money Invests and Market Prices Are Determined*. New Jersey: Princeton University Press, 2015.
- 2 Schulz, Kathryn. “Hoodoo, Hedge Funds, and Alibis: Victor Niederhoffer on Being Wrong.” *Slate.com*, June 21, 2010. http://www.slate.com/blogs/thewrongstuff/2010/06/21/hoodoo_hedge_funds_and_alibis_victor_niederhoffer_on_being_wrong.html