Smart Insurance:
How, Why and for Whom?

Leonid Kogan
Vinay Nair[1]
April 2018

[1] L. Kogan is the Chief Investment Scientist at 55ip and a Professor of Finance at the MIT Sloan School of Management. V. Nair is the founder of 55ip, and a visiting professor at MIT Sloan School of Management and The Wharton School at the University of Pennsylvania.

CONTENTS

Introduction

Insurance can be thought of as an action – usually with a cost – that insulates people from bad outcomes. When it comes to investing, people want protection from financial loss. This is done through risk management, which hedges against the risk of an uncertain loss. The simplest way to reduce downside risk is to trade the risky asset for a "safer" asset such as cash. The investor reduces the risk but has also limited the chance for higher returns. A more sophisticated approach is to use options, but they can be expensive and complex to manage.

There are also other approaches, which are often more flexible and cost effective. Institutional investors, for instance, typically protect themselves from stock market downturns by short selling stock index futures. But such insurance is, of course, not free. When the markets rally, the short position dampens any gains until fully unwound. Therefore, these direct and indirect costs of buying insurance may make the insurance undesirable.

Through dynamic risk management an investor can avail himself or herself of downside protection that we term "smart insurance" at a lower cost than the approaches mentioned above. In addition, dynamic risk management also provides the benefit of being flexible to adjust as market conditions change, thus further minimizing costs.

Much like the industry of smart beta has evolved around sources of returns other than the equity risk premium, we think that dynamic risk management can provide investors smart insurance that is better and cheaper than traditional portfolio insurance.

In this paper, we highlight the importance of downside protection in portfolios and help shed light on the promise of protection on the downside through dynamic risk management using volatility timing.

Benefits of Downside Protection

The Power of Compounding

To see why insurance can be a valuable aspect of an investment strategy, let's consider the importance of limiting drawdowns to wealth. If an investor loses 50% in the stock market, say through a period such as 2008, it will take her approximately 9 years to break even at a future annual return of 8%. A more protected lower volatility portfolio – e.g., one that is half as volatile by being half in cash, for instance – would by design have lost only 25%.

Obviously, this protection is enticing, but needs to be evaluated in the context of lower upside as well. The protected investor may become uncomfortable when the market starts rallying after the drop. Instead of the future annual return of 8%, the protected lower volatility investor only gets 4% (again, half of 8% to be consistent with the halved loss). Is such an insurance worth it? It can be.

For example, this lower volatility investor, who loses 25% instead of 50%, but also makes only 4% instead of 8%, will now break even more than a year and a half earlier than a higher volatility investor. But by the end of the 11th year after such a drawdown, the higher volatility investor would catch up and overtake the lower volatility investor (assuming she has sustained the heartburn for over a decade).

Nonetheless, this illustrates an important point. Portfolios that limit losses can be attractive even if accompanied by limited upside. This is because of the power of compounding – the same return increases your wealth more if you start with more. So, limiting losses allows your wealth to grow faster, even if the upside is equally limited. Such compounding characteristics don't exist in all phenomena that we seek insurance from, but when it does, it makes insurance even more valuable.

Avoiding Behavioral and Safety Triggers

There is another reason why such insurance is valuable. Consider the following. At the end of 2009, The Wall Street Journal reported that the best performing diversified U.S. stock mutual fund of the decade was CGM Focus Fund.2 The fund had generated annualized returns of 18%. Yet the article also pointed out that the typical investor in the fund lost 11% annually during the same period. The investor returns reflect their poor decisions on when they entered and exited the fund. Investor returns are often lower than fund total returns because they often buy a fund after it has had a strong run and sell after losses.

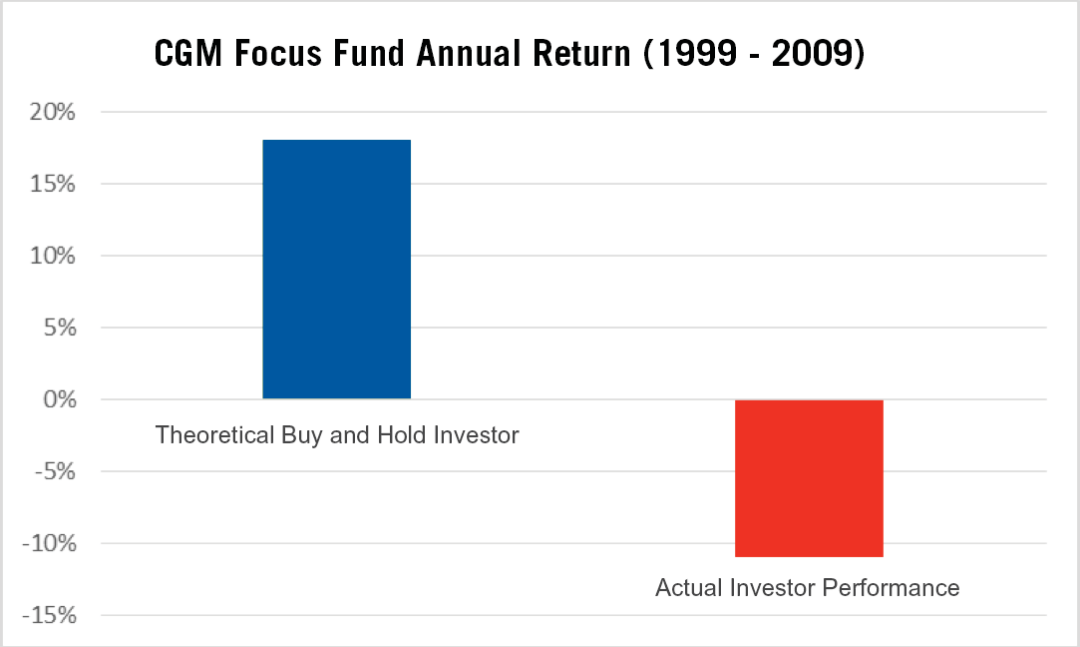

[2] *The Wall Street Journal*, "Best Stock Fund of the Decade: CGM Focus," December 31, 2009. Data sourced from Morningstar.

Higher volatility unmasks the behavioral demons in an investor's decision-making process such as overconfidence, optimism bias, hindsight bias, attribution bias, and confirmation bias, among others. CGM was a volatile fund that exacerbated the gap between a theoretical buy-and-hold return and what investors ended up getting for behavioral reasons. The fund surged 80% in 2007. Investors poured $2.6 billion into CGM Focus the following year, only to see the fund sink 48%. Investors then yanked more than $750 million from the fund in the first eleven months of 2009.

The case study of CGM shows that even well-performing strategies may not always pay off for investors if they are accompanied by high volatility.

Research shows that investors weigh gains and losses asymmetrically. Loss aversion was first convincingly demonstrated by Amos Tversky and Daniel Kahneman. The pains of losses swamp the utility derived from similarly sized gains. Some studies suggest that losses are as much as twice as psychologically powerful as gains.

It would be wrong to believe that irrational forces are the only factors at work. There could also be non-behavioral reasons for investors to tolerate losses up to a certain threshold, but then pull the safety trigger. For example, an investor may have some spending commitments to meet and cannot face possible further losses. As a result, higher volatility can create an asymmetric response between losses and gains.

In contrast to the CGM fund, consider the top performer during 2000-2009 among bond funds – the GMO Emerging Country Debt Fund, which rose 15% annually in the 10 years up to 2010. Here the investors realized a 14% return. Lower volatility of this fund, as well as the fact that it was held largely by institutional investors, were among the reasons for a better capture of positive returns.

Various biases in decision making or cash flow constraints lead to suboptimal choices and often generate a spiral of poor decisions. The decision to exit then leads to a poor decision on when to re-enter. These biases are also difficult to control even if one is aware of them. For example, even if an investor allocates 50% to the market and retains the remaining share of the portfolio in cash (to replicate the 50% downside portfolio), she may be tempted to evaluate the decision of whether to remain invested in the market in an exceedingly narrow context, based on the recent performance of only her stock investments. If, on the other hand, she allocated capital to a manager who does the same 50-50 cash-market allocation (but the investor only observes the combined returns), the behavioral implications may differ.

Asymmetric Returns – The More Powerful Goal

As the benefits of portfolio insurance become clearer, let's daydream a bit. What would the implications of free insurance be? Imagine that our investor only lost 50% of the losses, but remarkably made up 100% of the gains. How much wealthier would such an investor be relative to an uninsured investor? The insured investor's wealth is 1.5 times higher than the uninsured investor after 11 years. The insured investor would never regret having this insurance! The chart below shows how powerful the asymmetry of returns can be in the context of an investor who holds the SPY, the ETF tracking the S&P 500 Index. Starting on January 1, 1999, if an investor can achieve the full returns during an up year, but only suffer half the losses in a down year, they outperform an uninsured investor by 1.7 times by the end of 2015!

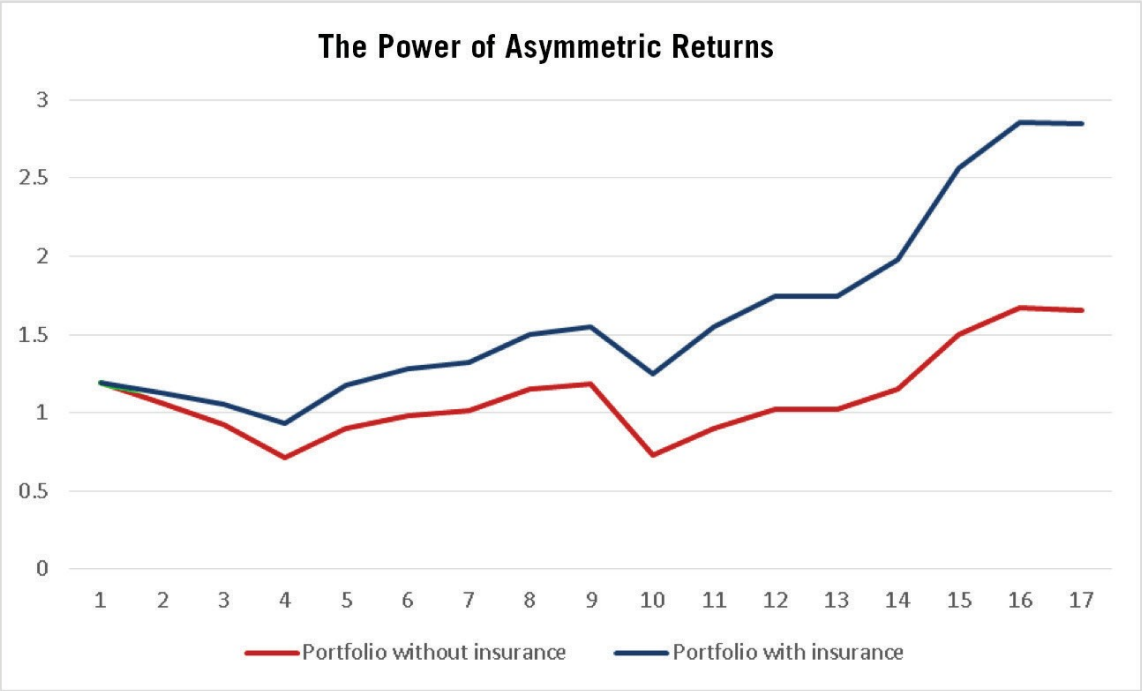

Source: 55ip

Smart Insurance

Clearly there are many ways to implement or buy insurance for a portfolio or investment position. But can we get closer to a "zero cost" insurance option? To achieve insurance against drawdowns without giving up an equivalent upside, one must have the ability to know when to buy insurance and when to discard it. This is dynamic purchasing of insurance. When the risk of a drawdown is high, you decrease your risk exposure. Conversely, you increase it back up when risk is low. This ability to adjust the amount of insurance is the core part of dynamic risk management and explains why we call it "smart insurance".

Some investors may look to market timing to achieve this goal as well. Even though, at short horizons, we are skeptical of market timing, this merits a comment. Insurance is related to protection from significant losses or events that are unusual. As a result, we are more interested in predicting not the average return but specifically the risk of extreme events.

Forecasting Risk is Possible, Practical and Robust

One form of dynamic risk management – volatility timing – allows an investor to use various measures of market volatility to identify periods with elevated risk of extreme events and significant drawdowns. This is good news since even as a significant amount of research documents the difficulty in forecasting returns – what many investors try to spend the most time and effort on, despite the odds – a large body of influential research documents the notable predictability of asset return volatility, which provides important implications for asset allocation. The evidence that volatility is forecastable also comes from different research traditions that include GARCH, stochastic volatility models, realized volatility models, and regime switching models.[3]

What is also helpful is that the math of estimation risk favors volatility over returns. Standard errors, a measure of the statistical accuracy of an estimate, are higher for means (returns) than for standard deviations (risk). For example, in a Gaussian setting, the standard error of returns is $\sigma/\sqrt{n}$ and the standard error of volatility is $\sigma/\sqrt{(2n)}$., where n is the number of observations. This improved precision in estimating volatility is further enhanced by the use of a high sampling frequency. For example, instead of n years of monthly data, one could use ~20n daily observations to reduce the standard error of the volatility estimate to $\sigma/\sqrt{(40n)}$. Interestingly, increasing the sample size of returns does not improve the returns estimate. The standard error remains the same as before.[4] As a result, the estimate for the volatility can be roughly 6 times better than the estimate for the larger sample of returns!

Volatility Translates into More Downside Than Upside

We know from the study of market level returns that the distribution of returns is not as well behaved as theoretically assumed. Returns do not exhibit a standard normal distribution with symmetric upside and downside. In fact, as you will see in the chart on the next page, aggregate stock market returns display negative skewness because markets have the propensity to have more months with very large negative returns than very large positive returns.

The exhibit specifically shows that over the 1966-2015 historical period, months with excess returns (defined as the difference between the return on the stock market index and the short-term treasury rate) of at least two standard deviations below the average are more than 1.4 times more likely than months with excess returns of at least two standard deviations above the average.

This feature is well documented and a large body of academic literature has aimed to explain this stylized fact about the distribution of aggregate stock returns (e.g., Fama 1965; Black 1976; Christie 1982; Blanchard and Watson 1982; Pindyck 1984; French, Schwert, and Stambaugh 1987; Hong and Stein 2003).

---

[3] Bollerslev, Chou and Kroner (1992) provide a fine review of evidence in the GARCH tradition, while Ghysels, Harvey and Renault (1996) survey results from stochastic volatility modeling, Andersen, Bollerslev and Diebold (2003) survey results from realized volatility modeling, and Franses and van Dijk (2000) survey results from regime-switching volatility models. The recent literature also contains intriguing theoretical work explaining the empirical phenomena, such as Brock and Hommes (1997).

[4] The standard error for the mean is 20 times standard error of the daily mean which is daily standard deviation/$\sqrt{(20n)}$)). The daily standard deviation is the $\sigma/\sqrt{20}$. As a result, the standard error for returns remains as $\sigma/\sqrt{n}$.

A Tale of Two Tails

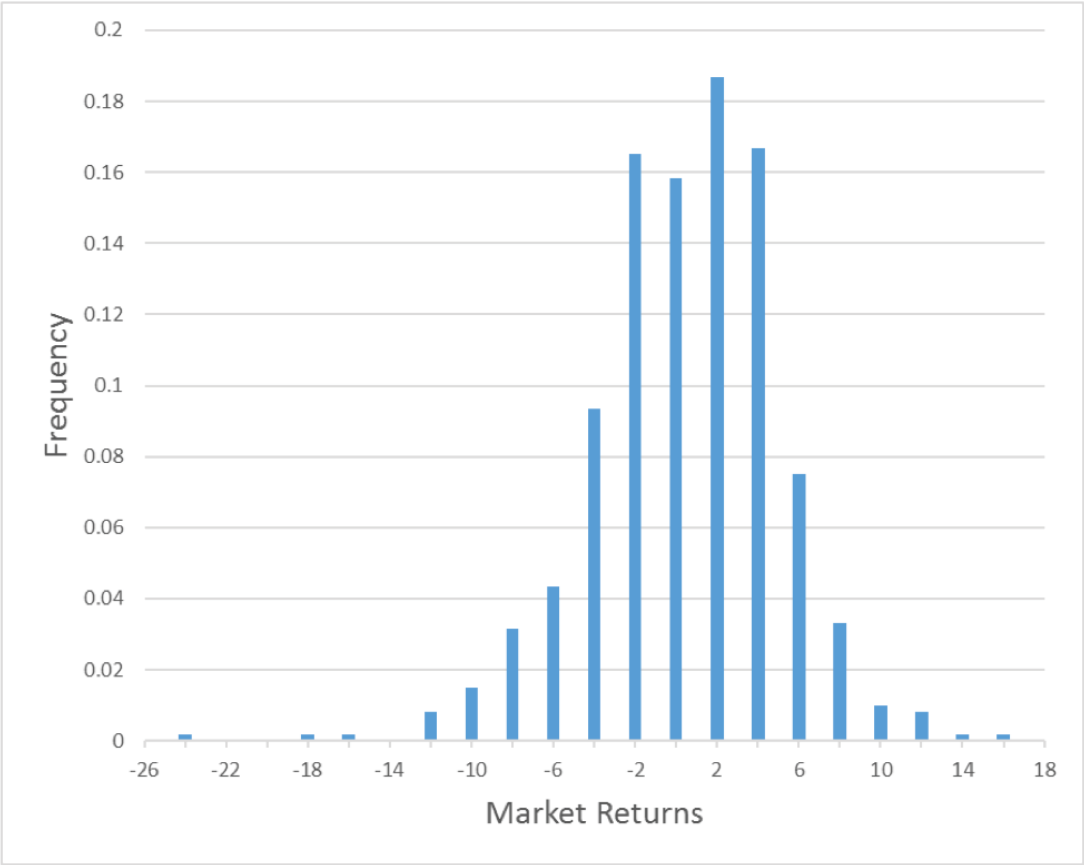

Source: 55ip

Higher Volatility and Higher Tail Risk

There are also economic reasons to believe that the benefits of buying insurance are higher than the costs, especially as volatility increases.

As volatility increases, the risk of investor liquidations increases. If these liquidations happen at the same time, as in a fire sale, asset returns drop significantly and may be a source of the observed skewness in returns. Worse yet, some large liquidations can move prices low enough to trigger other liquidations. This domino effect of sequential liquidations can have significant negative implications to returns.

These risks are further heightened in a world where there is increasing competition among asset managers and a significant overlap in the type of investment strategies and risk management tools they follow or use. For example, if all investors use the same VaR level on the market or a position, they all start liquidating at the same time. The increase in crowded positions and strategies increases the fire sale and contagion risks.

Traditional approaches to portfolio diversification assume stable relationships in prices across assets. But, in times of crisis, they may not be adequate for addressing tail risk created by crowded trades and fire sales. If everyone is diversified in a similar manner, the liquidations will create a situation where correlations across assets spike up and the historically low correlations that justified the diversifying mix would no longer be true.

This is an important lesson for investors who seek downside protection from unusual events. Low past correlations that can justify capital inflows may eventually increase liquidity-driven tail risk. In other words, diversification of normal risk can perversely lead to a concentration of tail risk. In higher volatility periods, investors are more vulnerable to such liquidity shocks and an increased probability of tail risks.

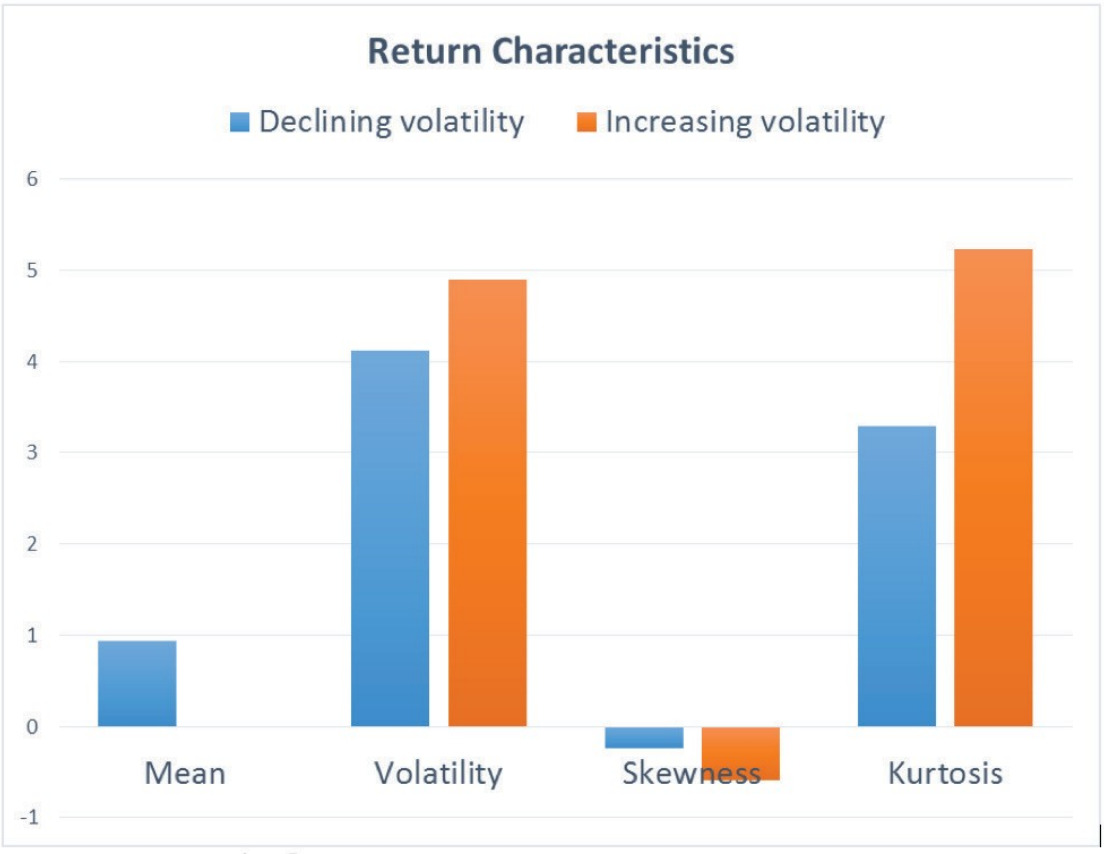

Source: 55ip

The exhibit above highlights the statistical evidence over the past 50 years (1966-2015) that supports the view that an increase in volatility is associated with an increase in tail risk. We compare the volatility of the past month over the prior month to characterize the current month into declining and increasing volatility periods.

During these two different types of regimes, we compute the average, volatility, skewness and kurtosis of the return distribution and find that the returns are indeed more negatively skewed (more downside than upside) and more likely to be in the extremes (high kurtosis or fat tails). The figure illustrates that return distribution exhibits higher tail risk after periods of higher volatility. This is manifested in more negative skewness and elevated kurtosis of return distribution.[5]

Dynamic Risk Management: A Solution to the Paradox

Since higher volatility is related to higher downside risk, if an investor can avoid such periods of higher volatility, the potential benefit is higher than the potential opportunity cost. Yes, there will be times when the buyer of such insurance regrets it. But over a long period of time, the negative skewness would make the decision of buying insurance a wise one.

Since we can predict periods of elevated risk, such dynamic risk management is implementable. During periods of higher predicted volatility, an investor can decrease their allocations or shift to safer asset classes, effectively buying some insurance. This opportunistic insurance purchase will remove the negative skew that market returns typically exhibit and may even generate a positive skew – higher upside than downside.[6]

Why doesn't everyone do this?

Markets have different types of participants, with different types of constraints and incentives governing investor behavior. For certain types of investors dynamic risk management may be undesirable or infeasible.

For example, many asset managers design their products relative to an asset specific benchmark, and are, in fact, compensated based on performance relative to a benchmark. Deviating from the benchmark would create unnecessary risk for which they are not sufficiently compensated.

For example, an equity manager might hold the individual securities in the S&P 500 Index even if she realizes that this is a period of increased fire sale risk because her performance is measured relative to the S&P 500. During these periods, such a benchmarked manager's version of insurance would be to mimic or hug the benchmark and take down active positions. It is a conundrum realized by few active fund pickers: no matter how bad an active manager who specializes in a single market does, so long as he or she stays close to the crashing benchmark, it can be blamed on the market, not the manager. On the other hand, multi-asset managers have more degrees of freedom to engage in dynamic risk management, to the benefit of their clients.

---

[5] The figure plots the first four moments of monthly returns on the value-weighted equity market portfolio in excess of the short-term risk-free rate. The first two moments are in percentages. Month t is said to be in a declining volatility regime if the intra-month daily return volatility is higher during month (t-1) than during month (t-2). The reverse holds for the increasing volatility regime. The sample spans the 1966-2015 period.

[6] Sellers of such insurance – those who deploy capital in such periods – will expect to get compensated. The level of compensation they need varies over time. An optimal insurance policy should calibrate the level of insurance based on this. We explore this in another note.

Another aspect that hinders dynamic risk management is position size. Very large asset managers simply cannot change positions efficiently. If they attempted to change their risk exposures drastically as the dynamic risk management strategy may prescribe – especially if they try to pick single securities – they would help create the very fire sale risk they seek to avoid. This could be one reason why smaller managers tend to find dynamic methods attractive, while large asset managers gravitate towards static approaches. Agility is a factor that affects the ability of managers to engage in dynamic risk management.

Who would value this the most?
Everyone likes some insurance, but some need it. Leland (1980) showed that two types of investors should buy insurance.

First, investors whose risk aversion increases with losses more rapidly than average could benefit from insurance. This group might include institutional investors such as pension funds, which are particularly averse to falling short of their performance target, but thereafter can accept reasonable risks. This group may also include individuals with inflexible spending requirements. Such "safety-first" investors would find portfolio insurance attractive.

Second, investors who have average risk tolerance, but whose expectations of returns are more optimistic than average, should buy insurance. This is particularly relevant in the current context. As monetary policy has unleashed a flood of capital into markets and as expected returns are falling dramatically, many investors are adjusting slowly to this new world. Even among the wary, return goals are slow to adjust because of inertia in policy return targets, forcing them to take on more risk than their risk appetites would otherwise dictate. In order to keep the risk of extreme losses within tolerable levels while reaching for higher levels of average returns, insured strategies are needed.

In addition to these types of investors, insurance is also valuable for financial intermediaries. Poor portfolio performance, particularly extreme losses, tend to trigger client redemptions as investors rightly or wrongly go looking for other solutions. Because of this amplification effect, intermediaries should be particularly sensitive to the left tail of the portfolio return distribution. As a result, financial intermediaries should benefit from insurance during times when risk of extreme events is high.

Times of significant losses in markets are also times when clients redeem, often for reasons unrelated to the performance of the intermediary. As a result, outflows in the business are significantly more likely precisely when risk of extreme events is high. As a result, any intermediary has an interest to maximize the capital that is not redeemed, especially at such times. One way to achieve this is to "trim the sails" during such high-risk regimes.

Conclusion

When investing, people often fear they cannot reduce their downside risk without giving up equal upside. This is not true. Dynamic risk management is effective insurance that can lower the risk of facing significant losses without an equivalent reduction in the upside. At 55ip, we use these insights to generate portfolios with reduced downside risk across equities, asset allocation and macro strategies.

55ip
617-960-9559
www.55-ip.com

BOSTON OFFICE
One Liberty Square, Suite 200,
Boston, MA 02109

NEW YORK OFFICE
12 East 49th Street, 11th Floor
New York, NY 10017

MUMBAI OFFICE
Aston Building, 10th Floor Shastri
Nagar, Andheri (W),
Mumbai 400 102, India