

Blind Ambitions

Some of the risk management tools that investors have embraced to try to avoid a repeat of 2008 could end up causing more harm than good.

By Max Stone, Daniel Michalow and Trey Beck

ILLUSTRATION BY MAXWELL HOLYOKE-HIRSCH



The idea goes by different names —

unintended consequences, negative externalities, perverse incentives, moral hazard — but regardless of what it is called, the result is the same: Measures taken to solve one problem can exacerbate that very problem or cause others. Examples abound, from the ability of infections to resist drugs to new roads that result in more congestion to forms of insurance that encourage dangerous behavior. An unprecedented amount of household wealth was destroyed in 2008, with the Federal Reserve estimating losses of more than \$11 trillion in the U.S. alone that year. Given the potential stakes, it's useful to consider whether some of the key steps that investors have taken in the past two years to protect themselves against the next financial crisis may do more harm than good. This naturally puts the spotlight on risk management.

Have market participants learned the right risk management lessons from 2008? In many cases, it appears they have learned some very good ones. Investors are paying more attention to tail risks; they are more conscious of the need to carefully manage liabilities as well as assets; and they are appropriately more skeptical of reliance on third parties like rating agencies. However, some important exceptions are worthy of discussion. This article examines four investor behaviors that, to our mind, reflect a fundamental misreading of the last financial crisis: relying on stop losses, buying options, stress testing and crowding into certain trades. More important, these behaviors are at best inefficient from a trading perspective and at worst could cause substantial trading losses regardless of whether markets themselves are behaving normally.

Stop-loss constraints have received an unusual amount of attention in the past couple of years as investors rattled by 2008's severe losses have asked themselves how they could avoid a repeat performance. What is a stop loss, exactly? It's a regime that mandates, in advance of entering into an investment position, the later reduction or elimination of that position if a predetermined loss threshold is breached. (This concept can be applied at a variety of levels, from individual positions to specific traders or business units.) On the surface, this sounds like a tempting solution for stopping losses, as the name itself implies. Recent performance appears to lend credence to the idea. For example, macro funds and commodity trading advisers (CTAs), which tend to use stop-loss mechanisms more than other hedge fund strategies, outperformed in 2008.

We believe that stop losses are generally ineffective as a risk management tool and thus almost never use them, relying on other solutions instead. In our view, "stop losses are good" is a poor lesson to take away from the last crisis.

Stop-loss protocols are simple, and simplicity is a virtue. But this simplicity comes at a price because stop losses require trading on the basis of stale or incomplete information. That hurdle arises

for two reasons: Stop losses are set in advance, which accounts for the staleness, and they are a function of only one form of information — price. This means a lot of other crucial market dynamics are ignored. Normally, when investors think about a trade and look at recent market data, they try to understand what's driving a given asset price. Stop losses, however, don't allow for such reasoning. To illustrate the potential pitfalls, let's consider a real-world example.

On the afternoon of May 6, 2010, the Dow Jones industrial average plummeted 573 points in a matter of minutes, only to recover most of that loss just as quickly. Although the "flash crash" started in equities, the panic immediately spread to the currency markets, which interpreted the equity sell-off as a "risk off!" signal. Riskier positions in emerging-markets currencies and popular currency carry trades began hemorrhaging money. The Brazilian real weakened by 5 percent against the U.S. dollar while the Australian dollar dropped 9 percent relative to the yen. It's noteworthy that a large fraction of the currency trades executed during the flash crash were stop outs. And most of those currency liquidations, it turns out, were awful outcomes for those affected. The stops hit when people didn't really want to get out, both because of their original trade thesis — in some cases, the positions were becoming fleetingly more attractive — and because there was virtually no liquidity (meaning that there were unusually high costs of execution). The price movements that triggered the stops reverted almost immediately.

In a sense, these stop-loss currency trades were similar to equity market sell orders filled at the lows that day. Most investors that sold Procter & Gamble Co. shares when they were down 50 percent during the flash crash probably didn't want that outcome, but their execution tactics left them vulnerable. There are many other examples. Of note more recently is trading in the Japanese yen a few days after the Tohoku earthquake. On March 17 the yen plunged 3.5 percent against the U.S. dollar in a matter of minutes because of stop outs, then rebounded within a few hours.

By the very nature of their simplicity, stop losses don't take into account other risks in a portfolio. That's a big deal. Consider an admittedly extreme example. Suppose that a macro fund's currency portfolio has a \$100 million long position in the South Korean won, is short the Taiwan dollar by the same amount and has 5 percent stop losses on both positions. Let's also assume that the fund manager doesn't have a view on Asian foreign exchange overall and therefore has put on these positions at the same size in an effort to hedge regional risk. Now imagine that Asian currencies sell off and that the won and Taiwan dollar perform similarly: down 2 percent on Monday, another 2 percent on Tuesday and 1 percent on Wednesday. The fund really hasn't made or lost any money — losses on the South Korean won will be matched by gains on the Taiwan dollar. Nonetheless, the won position shows a large loss, and the 5 percent stop loss consequently will be triggered, forcing a sale to cut risk. The problem is that the book is now short Asian currencies, leaving the portfolio riskier than it was before the stop loss kicked in.

The above example illustrates that stop losses may help limit losses on a particular position but they don't necessarily limit losses on the portfolio. A trader can try to mitigate this problem by setting stops on groups of positions, but he or she cannot escape it entirely. This property of stop losses is an important limitation because fund managers ultimately care about portfolio performance, not position performance.

ASSET MANAGEMENT

Stop losses are often promoted as sound risk management, but it's important to remember that managing risk is not just about putting guardrails around trades or even portfolios. It's also about creating structures that encourage responsible behavior. Stop losses can promote a number of unhelpful behaviors, often through complacency borne of a sense that these safeguards will do much of the risk management heavy lifting. In short, they foster a binary mind-set: "If I haven't hit my stop loss, the trade's still okay; if I've hit my stop loss, I must reduce or eliminate the trade."

Traders may rely on their stop losses to knock them out of a losing position, to avoid continually reassessing the risk-reward merits of the position. Similarly, risk overseers will tend to allow positions to run if they haven't hit their stop losses, even if that's not the right thing to do given changing market conditions. This isn't to say traders using stop losses won't necessarily conduct such analyses. It's just that they are less encouraged to do so by the apparent safety afforded by a stop-loss program.

There are other stop-loss implementation problems as well, especially with their most extreme form (used more widely by smaller funds and individual investors), whereby a trader leaves the stop order with a trading counterparty. As an execution tactic, this is unlikely to produce good results. After all, giving a stop-loss order to a dealer not only reveals the price at which one is willing to exit the position, but also cedes control of the order. And there's typically no guarantee that an investor will actually get out at the stop loss, only that the order will kick in at that point.

This reservation highlights an important illusion about stop losses: They are built on the assumption that at the stop point it's possible to trade out of a losing position at reasonable cost and thereby avert a lot more damage. This assumption often doesn't hold, however, particularly when a position is large relative to the market or when a trade is crowded and multiple investors get stopped out simultaneously.

Even in the most-liquid markets, stop losses can inflict substantial damage. For example, it's widely believed that in March 2010 the

For starters, we believe that much of the critical risk management function must be carried out when one is initially researching a trade. This is vital. Rather than apply an inflexible rule to all positions if they suffer some prespecified drawdown, we try to anticipate an array of possible bad outcomes and the specific circumstances that might cause them, and we then size positions at the outset by assuming that some of the worst of those outcomes may occur and that downsizing may not be possible except at prohibitive costs. The practical upshot of this is often that we just put on a smaller position than we likely would if we used stop losses. Simple, but effective.

In addition, it's easier to be nimble when responding to bad outcomes if you've been reasonable about the number of positions that your portfolio holds. For example, we tend to run our fundamental (as opposed to systematic) portfolios with a manageable number of positions so that we can make informed decisions when interesting news hits or prices move.

What do we do when a position goes against us? There are a few scenarios to consider. If the fundamentals have not changed but the price moves against us nonetheless, we might look to increase the size of the trade. If the price moves for a reason unrelated to our forecast, so that our go-forward forecast is the same, we would generally maintain the position size. If we're wrong and the trade loses money because our fundamental forecast turned out to be incorrect, then we would liquidate the position. In addition to those forecast considerations, a large adverse price move may signal increased risk, in which case we would tend to downsize the position. Of course, in the real world, prices will tend to go against us for several of the above reasons. The main point here is that, in contrast to the robotic nature of a stop-loss regime, we evaluate the utility of the position dynamically and in the context of the current set of available facts.

Now, a stop-loss proponent might argue that the approach we're advocating is undisciplined and affords much wiggle room for a crafty trader to rationalize a bad investment all the way to zero. It's true that the human tendency to indulge in "thesis creep" is real. But rather than give up hope that we can be rational and disciplined in the face of new

Managing risk is not just about creating guardrails around trades, it's also about encouraging responsible behavior.

triggering of stops greatly exacerbated a huge move in ten-year U.S. swap spreads. These stop losses were unfortunate. A substantial portion of the realized loss apparently resulted from transaction costs associated with liquidating such a large position. We believe the trade would have made a modest profit if these stops had not been tripped. What's especially remarkable is that the specific instruments at issue here — ten-year U.S. Treasury bonds and interest rate swaps — are two of the most liquid in the world, yet they moved significantly (and, as a side note, very visibly in public policy circles) as a result of a concentrated burst of stop-loss unwinding. It is simply untrue, therefore, that concerns about realistic liquidation assumptions are limited to smaller, less-liquid markets.

We have now discussed in detail why we rarely employ stop losses and why we don't give credence to some of the arguments used to defend them. So what do we propose instead?

information, we think it's our job to try to be just that. We think it's an overreaction to outsource this decision to a simple price-based rule.

When viewed purely on the economic merits, stop-loss strategies seem to have little to recommend them. Yet we know that some very successful firms use stop losses. How did such an inefficient trading practice gain such a wide following? Are we missing something?

In some contexts there's an internal logic to stop losses. Certain quantitative trend-following models widely employed by CTAs effectively incorporate stop losses into the investment thesis itself. A two-day momentum forecast, for example, has an implicit stop loss after a given instrument loses money for two consecutive days. But such cases are a minority. To get our heads around the wide popularity of stop losses, we need to move outside the bounds of portfolio management logic and consider principal-agent relationships within investment firms.

Some investment funds or proprietary trading desks employ a strictly siloed structure in which traders or business units operate relatively independently after receiving an allocation of capital. Often those individual traders are also monitored by a centralized risk management group. Stop losses tend to be more popular among siloed firms. Why? We believe two factors relating to principal-agent dynamics may be at play.

First, a siloed firm may use stop losses to impose order in a centralized or top-down fashion. Siloed firms may have dozens of individual traders running separate books that in the aggregate contain a multitude of positions across numerous asset classes. Under these conditions it may be virtually impossible for a centralized risk manager to know exactly what a given trader is doing all of the time. If the trader's profit-loss fluctuations are unexpectedly large and it would take a lot of time to understand why, it might be reasonable for the firm to implement a "shut it down and ask questions later" approach using stop losses. Here the risk manager uses stop losses to fulfill his or her duties in the most reasonable and effective manner possible given the firm's decentralized structure.

Second, given the potential penalties for exceeding risk limits, it may be rational for traders to impose stop losses themselves. Principals at siloed firms seek simultaneously to motivate and discipline their traders. That entails balancing opportunity and profit incentives against limitations designed to prevent an individual from taking so much risk that it threatens the firm's overall profitability or solvency. This balancing act is typically implemented through a strict "eat what you kill" policy whereby traders are compensated solely based on a percentage of the profits from their trades and by imposing stop losses that, if breached, will result in traders having their risk allocations pared back or, worse, their employment terminated. Under these conditions it may be rational for traders to impose (or willingly have imposed on them) stop losses on their individual positions to prevent a reduction in earnings potential or the career risk that can result if too many of their investments suffer significant drawdowns.

Let us briefly recap why we think stop losses are not the answer to a year like 2008:

- They are set in advance and based solely on price movements, so they demand trading activity that may prove irrational or destruc-

tive after subsequent or nonprice market information becomes available.

- They operate in isolation (on, for instance, a single trade) and thus may fail to capture important portfolio dynamics.
- They implicitly (and often very wrongly) assume that future trading activity can be accomplished, and at a reasonable cost.
- They may create a false sense of security about trading positions.

ANOTHER STRATEGY THAT SOME managers and investors have offered as a way to avoid large, 2008-like losses is to express trading views using option contracts. The argument here usually goes something like, "Buying options allows us to structure asymmetric bets with limited downside and unlimited upside and, by doing so, avoid the potential large losses we might suffer in an outright position." As with stop losses, we've noticed considerably greater interest in this risk management technique since 2008, so it's worth scrutinizing this claim in more detail.

To begin with, any argument that one should implement a view on instrument X by trading a different security, instrument Y, needs some serious explanatory support. It's true that option payoffs are asymmetric, which is a significant benefit for risk-averse investors. One important factor mitigating that benefit, however, is that the payoff profile only matters to the extent that it applies to the entire portfolio. For all but the largest individual trades, the practical benefits of asymmetry begin to wash out as the number of positions increases. So when it comes to implementing individual trades by buying options, the benefit to the end investor may not be all that significant.

Options are also not the most efficient way to express a directional trade view. For one thing, although the price of an option is obviously directly related to the performance of the underlying instrument, it's also affected by other factors, such as volatility, interest rates and, in the case of stocks, dividends. This is an important point; consider the following example.

Suppose a trader has a view that the exchange rate for the euro will go from \$1.40 to \$1.50. This directional view could be implemented simply by buying the euro or euro forwards. Or the trader could instead buy unhedged euro call options, thereby structuring an asymmetric payoff. Perhaps

\$1.46 three-month calls would do the trick. Unfortunately, this option implementation introduces two side wagers that are not part of the original forecast.

First, there's a bet on the volatility of the underlying security because the price of the option, at least close to inception, is a function of that volatility. Another way to think about this is that the ultimate payoff will be very sensitive to which option strike price is chosen. Suppose the trader's bullish euro view is correct and the currency goes to \$1.45 (in the right direction but not all the way there). Because the trader chose a \$1.46 strike price to generate a more asymmetric payoff, the trade as implemented will lose (because the options will expire worthless out of the money), even though the basic directional forecast was correct.

Second, the options implementation introduces an added and critically important timing component. Although timing always matters (all else equal, we'd rather see our investments work out quickly), timing is vital with options given their expiry feature. If the above positive forecast on the euro was realized over six months but the trade was implemented using three-month options, the fundamental insight would have been directionally right, but the trade could well have lost money, depending on the path the euro took to get to that price. To summarize, although the asymmetry offered by options is a nice risk property, they also introduce volatility- and timing-related uncertainty, which is a meaningful risk negative.

The most significant problem with expressing directional views using options is that their exposure to the underlying instrument varies over time. The typical argument for buying options focuses only on the terminal up-and-down scenarios and ignores the path to those final outcomes. Consider the previous currency example. Rather than buying \$100 million in cash euro, suppose the trader instead buys \$500 million of notional exposure to euro call options that have a 20 percent delta (meaning the option gains in value by 20 cents for each dollar increase in the underlying), so that the exposure to the euro at the trade's inception is also \$100 million. Now imagine that the euro goes quickly from \$1.40 to \$1.46. Suddenly, the effective exposure to the euro, as computed using standard option models, is about \$250 million.

It's great that the trade is profitable so far,

ASSET MANAGEMENT

but from a risk management standpoint, the position is now more than double its initial size. And note that the size increase has occurred as the price of the underlying security has moved closer to fair value. This is a general property of implementing a directional view by using options: The effective economic exposure to the underlying instrument will tend to grow when it would be better for it to shrink, and vice versa. As a trade moves favorably and approaches the target price, it's often prudent for a trader to reduce the position size given that there's less juice left in it. Similarly, as a position moves away, it may become a more attractive investment and thus deserving of increasing.

It's worth noting that the above property on its own is sufficient to prove mathematically that using options to implement investment views yields a lower Sharpe ratio than relying directly on the underlying instrument. This is true even if we ignore the previously mentioned timing and volatility uncertainties that options introduce. The intuition behind this math fact is that under the standard assumptions of the Black-Scholes pricing model, the payoff structure of an option can be replicated by dynamically trading the underlying instrument. So we can compare holding the instrument to maturity over one year against dynamically trading it over the same period (which is equivalent to buying an unhedged option). In the first case, our forecast over a certain time period is expressed by holding the instrument over that time period; in the latter dynamic sizing case (equivalent to owning a call option), additional noise is introduced because of the way the size of the trade fluctuates. This means that the volatility of this implementation is higher than the volatility of the direct implementation and that therefore the options implementation has a lower Sharpe ratio.

Additionally, using options is expensive. They're costly to buy, relative to underlying instruments, because of their volatility risk premium. (Implied volatility is typically higher than expected future realized volatility.) Moreover, options markets typically are far less liquid and have higher transaction costs than markets for the underlying securities.

To be clear, while we have argued against the implementation of most directional views using options, we think options can be wonderful instruments for certain investment ends. We frequently use them when

seeking to exploit apparent inefficiencies in the volatility markets, some of which we believe occur because of trading activity much like that described above. (In other words, we do think it's appropriate to use options for implementing views on options.) We also use out-of-the-money options as portfolio hedges. The skewness of options makes them particularly useful for hedging against crash events, which are discontinuous and produce portfolio tails that, as risk-averse investors, we want to hedge.

To recap our concerns with the post-2008 infatuation with options-based directional investments:

- Incorporating options into an investment position can muddy the waters by introducing additional forms of risk (volatility and timing, for example), about which the trader may have no opinion or insight.
- Options are often less liquid and involve higher transaction costs than the underlying securities.

Investors typically underappreciate risky scenarios that haven't previously occurred and worry too much about bad outcomes that have recently occurred.

- Option strategies have the perverse feature of increasing exposure to a trade as it approaches a trader's theoretical target price and decreasing that exposure as the trade moves away from the trader.

- In practice, the asymmetries that are obvious at the level of an individual option position often largely disappear in a broader portfolio of such positions.

STRESS TESTS HAVE BEEN ALL THE rage since 2008. They've been conducted on U.S. and European banks, and recently the concept has even been extended to nuclear power plants. We believe in the value of stress tests and conduct them as part of our risk management process. But we also think that leaning too heavily on narrow stress testing entails certain risks that should be borne in mind, particularly as this risk management practice gains more adherents.

A standard approach to risk management typically includes some form of value-at-risk analysis, or VaR, that places limits around portfolio volatility. One of the risk manage-

ment truths reinforced by the events of 2008 is that VaR methodologies by themselves are not enough. This shouldn't have been news, but markets hadn't experienced a major crisis for quite some time.

Crises reinforce the idea that VaR is nice but insufficient for two reasons. First, volatility increases and correlations change in a crisis; this means that precrisis VaR fails as a predictor of realized crisis performance. Second, in every crisis some set of factors that had not been known to carry unusual risk takes center stage and makes that crisis unique relative to previous events. Consider as an example 2008's funding risk — that is, the risk of losing financing on an asset purchased with borrowed money.

The main complement to VaR modeling is more-extensive and more-sophisticated stress testing. That people are doing more of this now is in many ways a good thing. Investment managers should worry about how their portfolios will perform in highly risky but plausible

scenarios — stock market crashes, liquidation crises, oil spikes and so forth — that are unlike the most recent past that VaR captures.

The problem, though, is that investors often overfit to the recent past. Investors typically underappreciate risky scenarios that haven't previously occurred and worry too much about specific bad outcomes that have occurred most recently. As a former director of the Fukushima Daiichi nuclear power plant said after Japan's recent earthquake, "We can only work on precedent, and there was no precedent." The engineers involved discounted the possibility of such a large tsunami because they hadn't seen one before.

Because the recent financial crisis is so fresh and was so painful, we worry that investors are stress testing based on scenarios that are overly specific. The simplest example of this is a tendency to ask, "What would have happened to this current portfolio in 2008?" (We're routinely asked exactly this by investors and consultants.) The answer is certainly worth knowing, but we don't believe portfo-

lios should be managed to parameters defined by that analysis, for at least a few reasons.

It's a general risk management principle that one should worry about a range of risks and try to use a diverse set of broad hedges when attempting to mitigate them. A risk model that is too narrow can easily be tricked or sidestepped by risk-takers, either intentionally or unintentionally. To take an extreme example, suppose an investor is allowed to own any portfolio so long as he or she would not lose more than \$10 million if, say, July 21, 2010 (pulled from a hat), happened all over again. This model is both specific and random, and it's clear what incentives it produces. With such a model in place, the portfolio manager can gain exposure to many risks without much penalty as long as the risks of this one historical scenario are avoided. The model also distorts the portfolio by skewing it away from a scenario that is unlikely to repeat itself.

More specifically, we're confident in predicting that the next liquidation crisis will always differ from the last one in important respects, because the world changes. New products appear, rules and regulations evolve, macroeconomic forces shift, and new economic actors arrive while old ones depart.

Even more worrisome for the "2008" stress test, the world will tend to change in ways that make recent bad experiences less likely to repeat themselves. For the present

The next financial crisis will necessarily differ from the last one because the world will change to avoid the major problems of the recent past.

situation, that means that the next liquidation crisis may have less of a "financing is going away forever" feel to it than did 2008, precisely because of how much pain investors and financing providers alike experienced when credit evaporated. Regulatory constraints, risk management practices, investor preferences and other forces continue to push market participants of all stripes to reduce their asset-liability mismatches, obtain more-secure funding, capitalize themselves more robustly and reduce counterparty risk. It would be very surprising to see moves in the cash-CDS basis (the difference in price between the cash price of a bond and the price

of its credit default swap) of the magnitudes seen in 2008 during a crash in, say, 2018. So managing to a highly specific past scenario can be bad because it permits "arbiting the risk model" and because it has characteristics that are less likely to recur.

An example of overfocusing on 2008 is the Basel Committee on Banking Supervision's proposal that banks include a "stressed value-at-risk" measure in certain calculations of market risk in bank trading books. The plan envisions calculating this risk based on a recent period of market stress. For many applications, that will be March 2008 through March 2009, and it will likely deter banks from taking the risks that hurt most in the recent crisis rather than the risks that will hurt the most in the average or next crisis.

Overall, although it's good for regulators to make changes to prevent 2008 from happening again, it's likely they will spend too much time and resources fighting the last war. We appreciate that the task of the Basel Committee and other regulators is not an easy one, yet building a regulatory regime and corresponding hedges around the most recent crisis could blind banks to non-2008 risks or create incentives to take on disproportionate exposure to new or emerging risks that don't fit the 2008 model. To be clear, we think investors should incorporate 2008 risk in their risk model but not overfit to it.

We can speak from considerable experience in this area, having learned some painful lessons from fixed-income trading losses we sustained in 1998 following the collapse of hedge fund firm Long-Term Capital Management. We have investigated past specific scenarios such as "If LTCM blew up today, how would our portfolios do?" This is a difficult exercise because it's not easy to know precisely how to apply that specific set of facts to today's markets. Some instruments weren't around in 1998; some of those that were are at vastly different levels today (two-year interest rates have gone from 6 percent then to about 1 percent); and some relative relationships have

inverted signs today (the 30-year swap spread turned negative in 2008).

What does "if 1998 happens today" even mean for these instruments? Such stress tests are also not realistic, in that portfolios do not remain static over specific periods, such as 1998–2008. In our experience, the results of such exercises are not all that meaningful and are reasonably different from what we would predict would happen in the "average" crash of that type. If 13 years from now we're still conducting stress tests based on 2008, what are the consequences? Will the results look as inapplicable as a "1998" stress test seems today?

To paraphrase Peter Bernard, our chief risk officer, if the 2008 risk management failure is now generally agreed to have been the result of an overreliance on VaR modeling, the failure in the next crisis may be a reliance on overly specific stress testing. VaR and standardized stress tests are popular because they are objective. And that's a useful property, particularly in the presence of agency issues (regulator versus bank or risk manager versus trader). These risk models are also relatively easy to implement, which is important as well.

But we believe that a broader and more subjective approach is called for. We try to test a portfolio based on a wide and changing set of scenarios in which the performance estimates reflect what might happen if a certain scenario occurred in today's market. We also try to put error bars on these calculations. A statement that a portfolio is expected to lose \$100 million in a given event will generate a very different risk management outcome than a statement that a portfolio is expected to lose \$100 million plus or minus \$1 billion. Subjective and broad-based risk management is difficult, and we do not expect regulators to be able to implement such a regime (at least, not without enormous cost). But we should not delude ourselves into thinking that stress scenarios based on what happened in 2008 will alone protect us in the next crisis.

To summarize our concerns about stress testing:

- The next crisis is unlikely to resemble the last crisis even in broad terms.
- Stress testing based on the last crisis may create perverse incentives by allowing risk-takers to pile on a host of risks—just not those characteristic of the most recent precedent.
- Stress tests may create a false sense of security unless they are sensitive to the dynamic

relationship between individual portfolios and a broad range of constraints and evolving market conditions.

THE NEXT FINANCIAL CRISIS WILL necessarily differ from the last one because the world will change to avoid the major problems of the recent past. The precipitating events will likely differ, and the financial damage will be concentrated in other areas. And importantly, positioning will differ. Despite these deviations, however, a common thread does tend to run through most financial crises: liquidity.

One inherent property of liquidity crises in particular is that crowded trades will be punished. Every cash-funded asset is owned by someone, and swaps, futures and options will always have buyers and sellers. A trade is crowded when a group of similar investors becomes concentrated on one side. The consequences of a trade being crowded are greater when this group of similar investors is unsteady, perhaps because they have overextended themselves from a risk perspective or relied on unstable leverage (or because they have used stop losses).

In a liquidity crisis, markets become thin, and it's much easier to move asset prices. For crowded trades in which investors might need to trade in a similar direction, there will be an imbalance between buyers and sellers, and the market will move against them. The imbalance between buyers and sellers will reduce liquidity, creating a vicious cycle.

Investors tend to understand all of this, and 2008 reinforced the lesson, but there is a tension between two of the mantras that investors have often repeated since then: "We need to commit ourselves now to actions that would have helped during the crisis" — for example, run stress tests based on 2008 — and "We need to avoid crowded trades." We see well-intentioned investors overemphasizing the first lesson at the expense of the second.

By way of example, consider the following investor behaviors:

Collectively bidding up tail-risk hedging products. This has made such products much more expensive and reduced their expected benefits in a liquidity crisis.

Favoring strategies that did well in 2008 on the assumption that they will also do well in the next crisis. This has resulted in more capital being allocated to strategies like managed futures, making them more crowded and likely to perform worse in the next liquidity crisis because of their popularity as a hedge.

Allocating predominantly to liquid assets and strategies. Financial products that offer cosmetically investor-friendly liquidity terms have attracted disproportionate capital since 2008 and thus have become more crowded. Consequently, the liquidity they end up providing may be disappointing.

These behaviors may stem from a collective action problem. That is, it's hard to account for the behavior of other investors in one's own decision making. Instead, investors study recent events, draw similar conclusions and then decide upon similar actions, all the while believing they're independent actors. Only later do they realize that they all made the same trade.

It's genuinely hard to understand in advance which trades or strategies are crowded. And we are not the only ones who worry about these developments. But worry about them we do, and for at least two reasons. First, at a minimum, this "fishing in the same pond" phenomenon necessarily limits the opportunity set for investors. With everyone nibbling away at the same limited set of market inefficiencies, returns will tend to erode and disappoint investors who have not factored this into their investment decision making.

Second, and more ominously, it's possible that this search for liquidity has perversely increased the likelihood of another 2008-style blowup, the very thing the quest for liquidity is seeking to avoid. The ultraliquid investment strategies and products so popular these days are likely to be the first place

investors will go to access liquidity in a crisis, just as many of the most-liquid products were tapped first in 2008. We worry that this not only increases the risk exposure of investors in those products but also threatens the health of the financial system more generally.

In summary, we're concerned that investors are so heavily focused on avoiding the bad outcomes associated with 2008 that they have piled into crowded trades that at best could dampen returns and at worst could provide a fulcrum for the next crisis. We're also worried that investors have become excessively partial to products that are viewed mostly as antidotes to the last crisis but may result in unexpected bad outcomes given collective action dynamics that are overlooked or not well understood.

Each financial crisis is in some sense unique and therefore unprecedented. And yet investors and traders must learn from previous crises even if the exact lessons will not perfectly apply to the next crash, which history tells us is certain to come. Of course, there are no easy solutions, and we are not the only ones worried about the next crisis. But ironically, if we all learn the same specific lessons from 2008, we probably have learned nothing at all. ●●

Max Stone is a managing director and a member of the executive committee at the D. E. Shaw group, a global investment and technology development firm headquartered in New York. He jointly supervises the firm's worldwide asset management businesses, with a particular focus on its macro trading activities. Daniel Michalow is a senior vice president at the D. E. Shaw group and a member of its macro trading group, where he contributes to the generation of investment ideas across a number of asset classes and to overall portfolio construction. Trey Beck is a managing director at the D. E. Shaw group. He has oversight responsibility for the investor relations and external affairs groups and participates in the development of the firm's overall business strategy. The authors thank their colleague Erik Yesson for his meaningful assistance in the preparation of this article.

Disclaimer: The attached document is provided for your information only and does not constitute investment advice or convey an offer to sell, or the solicitation of an offer to buy, any securities or other financial products. Furthermore, the D. E. Shaw group does not endorse any information, beliefs, and/or opinions discussed in the document and makes no representation as to their accuracy or adequacy. Please note also the date of the document as the information contained in it has not been updated for any information that may have changed.

Past performance should not be considered indicative of future performance. No assurances can be given that any aims, assumptions, expectations, and/or goals described in the document will be realized or that the activities described in the document did or will continue at all or in the same manner as they were conducted during the period covered by the document.