NORTHFIELD | modeling financial markets worldwide

The Liquidity Risk Time Bomb

Chris Kantos

Northfield Research Seminar London

April 2019

Motivation

- During the financial crisis years of 2007-2009, much of the declines and volatility experienced by global markets purportedly had to do with liquidity related concepts.
- Regulators such as the US SEC and the various aspects of MIFID II in Europe have begun to require that asset managers of open-end funds and ETFs carry out analyses of their liquidity risk.
- In addition, regulators desire trading practices that do not unfairly shift the cost burden of large liquidations to remaining investors from those investors withdrawing.
- Various approaches industry are being taken to analyze liquidity risk. Unfortunately, we find that in all but a few cases the analytical approaches being undertaken are unsound.
- These flawed analyses give the impression that market liquidity to transact securities is far greater than actually it is.

The Time Line Since the GFC

- Liquidity effects of interest began with the "hedge fund meltdown" of August 2007 and the destabilization of money markets triggered by the failure of Lehman Brothers.
- The near-failure of numerous other financial institutions contributed further to the misery, to which central banks responded with unprecedented injections of massive liquidity into financial markets
- Since then, the equities world has been subjected to lots of discussions on "crowding" of strategies and factors. Interest rates have gone to zero or even negative in many countries.
- The rapid growth of ETFs makes the current problem worse, as ETFs are traded with high liquidity but without regard to the fact that many of the underlying securities may not be equally liquid.

A Typical Current Analysis

- Encouraged by regulators, many asset managers now analyze liquidity by considering how long it would take them to liquidate their positions, assuming a given rate of trade participation and current (i.e. usually normal) market conditions.
- For example, how long would it take an asset manager to liquidate the positions in a $10 billion mutual fund assuming their trades could be 20% of the trading volume in the relevant assets.
  - Obviously such a measure is calculable and typically the result is a matter of a relatively small number of days, *which seems ok*
- However, why would all investors in a given fund want to liquidate but investors in the firm's other funds would not?
  - What if our $10 billion fund was part of a $1 Trillion fund family
  - If a crisis or scandal hit a firm, it is likely all funds would suffer

Now Let's Dig a Little Deeper

- If our subject fund was an equity fund, why would we think it realistic to assume that this one fund could be 20% of trading volume for an indefinite period.
  - The global equity market is worth in excess of $70 Trillion, so how would a $10B fund manage to be one fifth of the total volume?
  - Put another way, a maximum five funds can possibly be 20% participants in trading volume, when many thousands exist.
  - If we repeat the analysis at a participation rate comparable to the fund's participation in global AUM, the results would show that most funds would require months or years to liquidate.
- A much more likely scenario is that either the entire firm is impacted by a "run on the bank" or that a macro event (e.g. the August 2007 hedge fund mess) impacts entire markets.
  - Now the simple analysis is even more implausible.
  - *If there is a fire in a crowded theatre, and everyone stampedes to the exists, it is not possible for everyone to be the fastest runner.*

Our Earlier Proposal

- Several years ago we proposed that fund's should calculate their risk metrics such as volatility, tracking error, or Value at Risk in a liquidity adjusted fashion.
- As a portfolio gets larger and larger in value, it becomes more cumbersome to run in the same way that a large ship is not as agile as a small boat.
- As the market impact of larger trades increases the cost of a crisis liquidation, the effective risk level of a portfolio is a function of portfolio size
  - All traditional portfolio theory is based on the assumption of infinite liquidity so risk is a function only of portfolio weights
  - *In the real world, trading costs are an increasing function of size so both portfolio size and weights matter in risk calculations*

Formulating a Liquidity Policy

- Theoretical discussion is in Acerbi and Scandolo (2008)
- We can formulate a liquidity policy as:
  - We have to be able to liquidate X% of the portfolio in N trading days
- Given our models of cost and the size of our positions, we can estimate the cost of liquidation during crisis conditions
  - Consider the impact of "crisis volatility" on trading costs
- To adjust our portfolio risk estimate for liquidity
  - Convert our portfolio volatility estimate to parametric Value at Risk for the length of time specified in our liquidity policy
  - Add the expected cost of fulfilling liquidation to VaR
  - Convert the new VaR value back to the equivalent volatility

A Simplified Example

- Our liquidity policy:
  - We must be able to liquidate 30% of the portfolio in 10 trading days
- Our estimated portfolio volatility is 25% per year
  - Assume 3 standard deviation VaR (covers 99.8% of normal distribution)
  - % Parametric VaR = 14.94 [25 * 3 * (10/252)^.5]
- Assume our forecast liquidation cost is 4% under normal conditions
  - We can get this from model we will discuss later
  - % Parametric VaR with Cost = 18.94 [14.94 + 4]
  - Revised portfolio volatility = 31.70 [18.94 / 3 * (252/10)^.5]
- *Volatility estimate increased by more than 23% under normal conditions*
  - *For typical equity portfolios the effect is about a 60% increase in risk during crisis periods just from the greater liquidity costs*

Playing “Show and Tell”

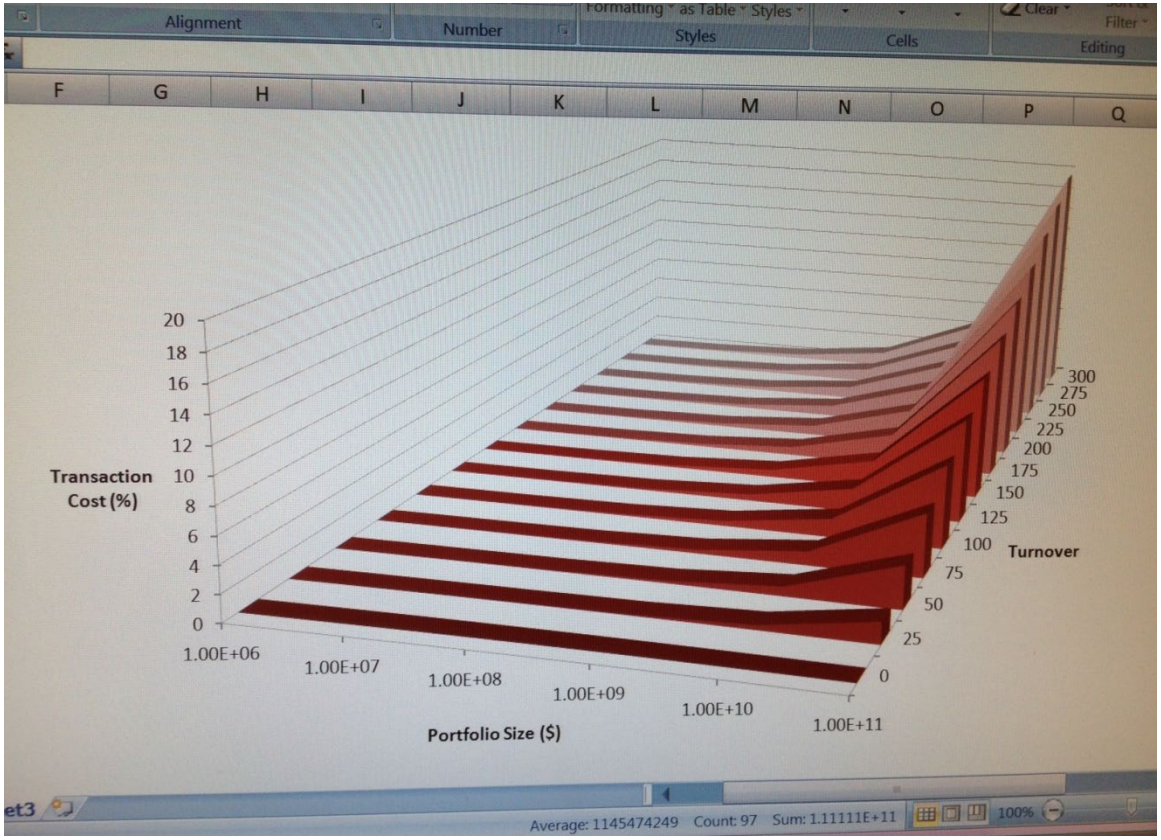

Sample Calculation Inputs and Assumptions

- The portfolio is the fifty smallest market cap stocks in the S&P 500
  - Highly liquid by global standards but less liquid than many US stocks
  - Equal weighted
- The annual turnover dimension runs from 25% to 300%
- The AUM dimension runs from $1 Million to $100 Billion
- We assume that conditions are “normal” for risk and trading costs
  - We estimate that the costs would increase by a factor of 2.5 during a “crisis” period as seen in 2007
- We assume that traders get one trading day to execute transactions
  - If we give traders a week to execute each trade we expect per share costs to drop by about 65%
- The 3-D graph illustrates the expected annual trading costs of a hypothetical portfolio at various levels of AUM and annual turnover
  - You can see how much alpha is lost to trading for any combination
  - The maximum annual cost is over 20% (300% Turnover, $100 B)

A Recent Contributor to Illusions of Liquidity

- ETFs and related ETNs have grown exponentially in recent years.
- While the routine operations of these securities imply liquidity as both an exchange traded asset, and the creation/destruction of units. There is also the structural possibility of *in-kind contributions and withdrawals of securities*
- Most ETFs and ETNs are based on passive indices. Some more exotic ETFs (levered, inverse, VIX related) have daily rebalancing rules required by the prospectus.
- The rebalancing of the underlying portfolios is predictable to a material extent, putting a bound on "normal" trading. Some funds also use prearranged "in kind" flows of securities around rebalances to reduce realization of taxable capital gains.
- Debt funded market makers and hedge funds have been providing liquidity for rebalances in return for transaction spreads, *particularly those securities with low trading volume.*

Basics of Security Market Liquidity

- You can think of market liquidity as a simple problem of supply and demand. Basic economics tells us that price will change whenever there is an imbalance. We call this market impact.
- Understanding the extent of imbalances can be observed from limit order book data, the actual flow of orders (if available) or inferred by methods such as Lee and Ready (1991).
- Large investment banks and broker dealers have access to the order flow information from the clientele, and have created models to predict market impact such as that of Kissell (Northfield conference, 2015).
  - Such models usually limit prediction to transactions involving participation of not more than 25% of expected trading volume as empirical data is very sparse for larger trades
  - The data is sparse because traders know that very large trades are likely to cause large adverse price movements.

A Framework That Works

- For many years, Northfield has used a functional form for predicting trading costs that captures the empirical feature that costs rise rapidly when trades get too big for the market to absorb.

$$M = A + BS + CS^{1/2} + D * MAX[S-L,0]^2$$

M = percentage transaction cost

S = proposed trade size in shares

L = the number of shares where routine liquidity is exhausted

A = agency costs (broker fees, normal bid-asked spread)

B, C, D are estimated constants per security

B, C, D and L assume a known time horizon for trade completion. All inputs can be restated as percentages of expected trading volume

Defusing the Last Bomb

- When trades get larger than S > L, then the price impact of trading imbalances gets very large, very fast.
- Traders know this so they avoid doing such large trades
- In the GFC period, costs got large because both S values went up (lots of people trying to sell) and L went down (market makers withdrew from activity due to margin calls).
- Central bank first response was to flood markets with money to bring L values back up. The second response was to start buying securities so that larger values of S (selling trades) could be accommodated without creating imbalances.

A Blast From the Past

By Dan DiBartolomeo

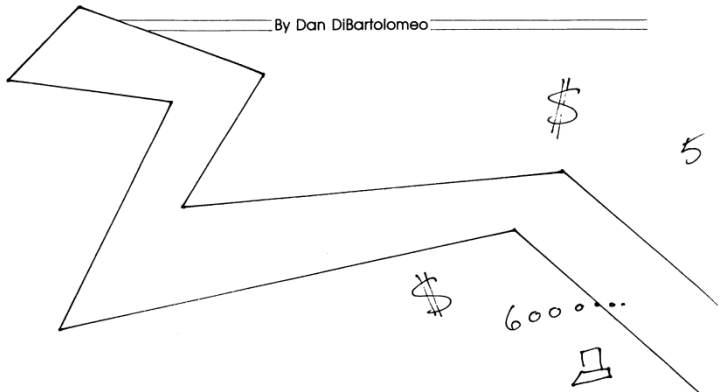

COMPUTERS AND THE CRASH: DID THE MACHINES REALLY DO IT?

A lot of people blamed computers for last year's stock market crash. We look at the role they really played.

The stock market crash on October 19, 1987, reduced the wealth of every man, woman, and child in the United States by an average of $5,000. In total, more than *$1 trillion* in the value of traded stocks was wiped out in the span of just seven hours. The morning newspapers of Tuesday, October 20, began the inevitable fingerpointing. One of the most popular rationalizations was that computer-aided trading practices, which had come into vogue in recent years, had gone haywire. In short, it was all "the computer's fault."

To begin to understand the role computers played in what has come to be realized as a near-total collapse of the world's capital markets, we first must come to grips with the enormity of the financial events themselves. In addition to a one-day loss of nearly a third of the value of all stocks traded in the United States, every major stock exchange in the world dealt with near-panic conditions. Total losses around the world exceeded $2 trillion in a matter of a few days. This sum is greater than the accumulated national debt of the United States, which has taken over 200 years to accrue. In Hong Kong, things got so bad they simply closed the stock exchange for a week, bringing finance and commerce to a near standstill.

The Dow Jones Industrial Average lost 508 points, over 23 percent of its value, on October 19. This percentage is more than double the 11.2-percent loss recorded on October 29, 1929, the event normally cited as having triggered the Great Depression. The likelihood of such an event happening randomly, given the history of stock market ups and downs, is so small as to be nearly incalculable.

Dan DiBartolomeo is president of Northfield Information Services, a consulting firm specializing in computer-aided investment analysis.

The October 1987 Crash

- The preceding page is an illustration of an article Dan wrote on the October 19th, 1987 stock market crash.
- The factual data was taken from federal study known as the Brady Commission Report. Lead investigator was Robert Glauber (now at Harvard). The report is named for Nicholas Brady who went on to be Treasury Secretary. I know both and have discussed the events.
- The seminal feature is that the value of the US stock market dropped by $1 Trillion or 23%. Globally the effect was roughly $2 Trillion or 20%.
- Total volume traded in the US markets was only around $15 Billion.
- Even if we assume that every transaction was initiated by a seller, the market impact was 70 times the total volume of trading, or a market impact of 7000%

Now We're Equipped

- If we take L as 25% of expected volume, and S as 100% in expected volume terms, we can back out a value for D of 1.244.
- Empirical studies of trading costs during crises periods suggest that such costs were around 2.5 the corresponding routine costs.
- Dividing 1.244 by 3, we get a D value of around .5.
- Given values for A, B and C from traditional transaction cost models we can estimate the cost of any size one day trade from a single fund, fund family or the market as a whole.
- We can also estimate the temporary and permanent components of such trades when spread over multiple days. Our internal research suggest costs decline at a less than linear rate when trading is slowed down. Our default coefficient estimate is .71 so if we spread the same trade over 3 days

  Cost(3 days) = Cost (1 day) / 3^.71 = Cost (1 day) * .459

Lighting The Fuse to the Time Bomb

- Many liquidity analyses being carried out for regulatory reporting assume participation rates that would routinely be sufficient to put markets into imbalances such that S > L, triggering potentially very large price impacts.

- In the event of some kind of market wide imbalance, there is little doubt that S > L and large price impacts would arise. The question is whether asset managers would make matters worse by trying to sell out at the participation rates suggested in previous liquidity analyses.

The Fairness Problem

- If we assume that large price impacts would arise from a liquidity event, there is a fairness issue between investors seeking to liquidate and investors choosing to hold their investment.
- Once the explicit cost of liquidation are considered, asset managers might be incented to sell off more liquid assets to minimize costs, but leaving the remaining investors with less liquid securities that could be even worse in the event of a second liquidity event.
- The incentive comes from wishing to preserve AUM fees and minimize the perceived impact on investment performance.
- One sensible approach is to force the post liquidation portfolio to have a maximum amount of tracking error relative to whatever the portfolio was before the sell-off. As long as this pre-set limit is not breached, managers would have the flexibility to select which assets to sell to minimize costs, and would have clear reasoning for spreading trades over multiple days.

Conclusions

- The concept of analyzing liquidity as the number of days necessary to liquidate the positions of a fund under typical conditions is overly simplistic.
- Viewing liquidity in this way may actually contribute to future crises by encouraging asset managers to trade at participation rates likely to trigger crisis conditions in which market values of assets would be substantially impacted.
- We have presented a tractable and realistic functional form that describes how overly aggressive trading by even a few entities could cause massive, if temporary, dislocations of asset prices.
- To the extent that large price movements could be anticipated, a fair balance must be struck between investors seeking to liquidate and those investors seeking to remain. Existing optimization techniques can be sensibly employed to form the solution.