Government Office for Science

Foresight

Systemic risk arising from computer based trading and connections to the empirical literature on systemic risk

Driver Review DR29

Foresight, Government Office for Science

Contents

Systemic risk arising from computer based trading and connections to the empirical literature on systemic risk

Jon Danielsson and Ilknur Zer

London School of Economics

May 2012

This review has been commissioned as part of the UK Government's Foresight Project, The Future of Computer Trading in Financial Markets. The views expressed do not represent the policy of any Government or organisation.

1. Introduction

No clear consensus exists on what constitutes systemic risk, neither empirically nor theoretically. Although systemic risk might be defined as the risk of the entire financial system collapsing, translating that conceptual view into practical measurements of systemic risk and policy is not straightforward. The extant literature has developed in isolated silos for most parts, with research and proposals focusing on one particular aspect of systemic risk rather than the concept in its entirety.

Concerns about what is now termed systemic risk arising from trading are not new in the history of the financial system. Just a few examples are the crises of 1907, 1914, 1929 and 1987. Moreover, market participants have always employed algorithmic trading (AT) rules and endogenous risk type adverse feedback loops have historically been common. However, the ever–increasing prominence of computer based trading (CBT), that is trading where computer algorithms directly interface with trading platforms and placing orders without immediate human intervention, makes the study of systemic risk arising from CBT an especially pertinent issue, especially after the flash crash of 2010.

The objective of this paper is to provide explicit links between the literatures on systemic risk and CBT and make proposals as to how one might develop a more unified systemic risk research agenda. The first literature we consider is concerned with empirical systemic risk measures (SRMs). This relates to a set of tools developed over the past few years aiming at translating some observed data on the financial system into a measure of systemic risk. Our main interest here is SRMs that use market data, typically daily, as the main input into the systemic risk measure. We term those MSRMs (market data based systemic risk measures).

The second literature is on CBT and its subset high–frequency trading (HFT) which have come to play a prominent role in the exchanges all over the world. While CBT has been a key factor in anomalous market outcomes for several decades, and its main ingredient AT even longer, extant formal studies of CBT currently only provide a tenuous link to commonly understood concepts of systemic risk and almost no direct connection to empirical studies of systemic risk, nor MSRMs.

The connection between the CBT and systemic risk depends on whether the aggregation of individual trading outcomes, however extreme at high frequencies, leads to a random noise at immediate sampling frequencies or not. We suspect both are true.

The speed of trading may make it easier for anomalous outcomes and asynchronicities between different markets to self–correct, compared to human only trading. Therefore, events like the 1987 crash might have been less severe if trading had been faster. The fact that the flash crash played out within a single day, and high frequency extreme outcomes do not seem to result in extreme daily outcomes, provides strong evidence for the self–healing nature of high–frequency trading. In other words, HFT might reduce the incidence of market crashes.

On the other hand, a sudden coordination between high frequency traders in selling or buying the same assets may create feedback loops which in turn may lead to a structural break in the prices. Hence, CBT is more likely to affect systemic risk via endogenous risk and resulting adverse feedback loops. While none of the extant MSRMs capture this phenomenon, we expect that combining extreme value theory (EVT, i.e. power laws) type approaches to MSRMs with an explicit focus of endogenous feedback loops might provide a roadmap for future research on the systemic implications of CBT.

2. Systemic risk measures

Systemic risk is a term used frequently, both in popular media and by specialists. Unfortunately, most usage is imprecise and contradictory. One common definition is from Trichet (2009):

> "systemic risk, in a broader sense, relates to the risk that widespread instabilities in the financial system translate into adverse effects on growth and welfare in the economy at large."

whilst the BIS in 1994 finds

> "risk that the failure of a participant to meet its contractual obligations may in turn cause other participants to default, with the chain reaction leading to broader financial difficulties."

These are representative for two different types of views, either focused on the financial system adversely affecting the real economy (Trichet) or particular problems in the financial system leading to financial difficulties (BIS).

These high level definitions of systemic risk provide multiple different paths for translation into more operational constructs, either theoretically or empirically, many of which have little relation with each other or are even mutually inconsistent.

2.1. Operationalizing systemic risk analysis

When it comes to translating the high level definitions of systemic risk into concepts that can be used operationally, many authors have opted to ignore the interconnectedness between financial institutions, perhaps because it is harder to model empirically. More recent studies of systemic risk are more likely to explicitly focus on the inter–linkages in the financial system, with cascading failures amplified by the inherent pro–cyclicality of banking and regulations, see for example Kritzman et al. (2010) or Drehmann and Tarashev (2011).

However, identifying inter–linkages and pro–cyclicality as the main ingredient of systemic risk still leaves multiple different ways of defining it on a conceptual level, with yet harder problem of how it can be measured. Perhaps, a single measure of systemic risk is neither feasible nor desirable, with policymakers considering different definitions of threats to financial stability in their analysis.

2.2. Probability levels

Besides identifying the main ingredients of systemic risk, it is also necessary to identify the probability threshold required for an event to be systemic, along with levels of granularity and accuracy. No consensus exists on such probability thresholds. The definitions at the beginning of this section would suggest an event that happens very rarely, perhaps a couple of times in a century. That frequency of (near) systemic events is consistent with the view that the last three (near) systemic crises happened in 2008, 1929 (the Great Depression) and 1914.

Such high level views of the probability of systemic risk do not translate well into operational tools, and especially not into empirical measures of systemic risk. After all, this suggests that 2008 is the only (near) systemic event in recent history, making statistical analysis rather difficult.

As a consequence, most studies aiming to operationalize systemic risk measures focus on less extreme probabilities. Some authors and organizations even view a typical financial crisis as a systemic event, giving rise to difficulties in identifying and analyzing more serious crisis events.

In practice, most MSRMs focus on probabilities such as daily 1% (2.5 times a year) outcomes, and the issue of whether that is extreme enough to be considered systemic remains an open and controversial question. While such non–extreme probabilities facilitate empirical modelling, our view is that a daily 1% event is far from systemic, and does not provide by itself any useful information to policymakers concerned with systemic risk.

Furthermore, the definition, theoretical methodology, empirical methodology and data need to be directly consistent with the probabilities one considers systemic. Trying to somehow extrapolate from non–extreme to more extreme systematically related probabilities by means of probability shifting, that interpret interference at one probability level for outcomes at another level, is not advisable.

2.3. Measurement challenges

The empirical modeling is further made difficult by the fact that systemic events seem to happen quite suddenly, even if the underlying risk builds up slowly over time, and hidden away, as noted by Crockett (2000):

> "The received wisdom is that risk increases in recessions and falls in booms. In contrast, it may be more helpful to think of risk as increasing during upswings, as financial imbalances build up, and materialising in recessions."

This suggests that most indications of increasing systemic risk only flash up after a crisis is already underway, at which time every single indicator points in the same direction.

Trying to propose a forward–looking method/indicator for identifying systemic risk is challenging since we only have one recent event to test such method/indicator on. Therefore, there is a significant probability of spurious results. This becomes especially problematic if the probability of the underlying crisis events is relatively non–extreme and model risk high.

2.4. The literature on market based SRMs

A large number of authors have proposed various different empirical systemic risk measures (SRMs). For a comprehensive survey see Bisias et al. (2012). The connection between most of these and CBT is rather remote, but some are more directly related, in particular those who make use of the same type of data (securities prices), albeit at different frequencies.

Within the broad universe of SRMs, a common approach is to use medium frequency market data, such as daily stock returns. For this reason we term these market data based SRMs (MSRMs).

Often, the authors start with some theoretical model of interconnectedness between financial institutions and the transmission of adverse outcomes. Since estimating such structural models would be challenging, generally the authors then propose a reduced form statistical model. Many of these reduced form approaches consider both the time series and cross–sectional dimensions, with medium frequency securities data as the main input.

Of these, one example is the banking stability measures of Zhou (2010), a set of quantitative measures of the financial stability of a portfolio of banks. A different approach is taken by Adrian and Brunnermeier (2010) (conditional value at risk, CoVaR) and Acharya et al. (2010); Brownlees and Engle (2011) (marginal expected shortfall, MES). These measures first identify the asset or equity returns of publicly traded financial institutions and define a systemic event as simultaneous losses among those financial institutions.

Focusing on aggregate imbalances and considering the financial interconnectedness, Kritzman et al. (2010) propose the 'absorption ratio' whereas Billio et al. (2010) propose several systemic risk measures, derived by using Granger–causality test and principal components analysis, where the latter enables them to gauge the degree of commonality among a vector of asset returns.

Finally, although related, but focusing more on the credit risk as an underlying source of the systemic risk, Tarashev et al. (2010) proposed the Shapley value (SV) approach and Huang et al. (2010) construct a systemic risk indicator, the distress insurance premium, using the structural model of Vasicek (1991).

All of these MSRMs make use of (conditional) volatilities and/or VaR, and we hence refer to them as V–MSRM.

An alternative to the V–MSRMs is to use methods focused on more extreme probabilities, typically EVT. Several proposals have been made for EVT based MSRMs, for example Hartmann et al. (2004) and Gravelle and Li (2011). The main problem with these approaches is that they are based on unconditional probabilities for very large outcomes and therefore provide an imprecise signal for systemic risk, unless one is willing to assume commonality in the distribution of extreme outcomes across crisis. This could be problematic since the statistical properties of different crisis events, and the underlying direct causes, are generally quite different.

2.4.1 The unifying factors

Of the MSRMs, perhaps CoVaR and SES/MES are the most prominent. Danielsson et al. (2011) study their theoretical and empirical properties and we adopt that analysis below.

Let R_i indicate the risky outcomes of a financial institution *i* on which the risk measures are calculated. This could be for example, daily returns of such an institution. Similarly, we denote the risky outcomes of the entire financial system by R_s. We can then define the joint density of an institution and the system by

$$f(R_i , R_S).$$

The marginal density of the institution is then $f(R_i)$, and the two conditional densities are $f(R_i \mid R_S)$ and $f(R_S \mid R_i)$. If we then consider the marginal density of the system as a normalizing constant, we get by Bayes theorem:

$$f(R_i \mid R_S) \propto f(R_S \mid R_i) f(R_i) .$$

Suppose we use as a risk measure, we arrive at CoVaR. Define Q as the event such that:

$$[R \leq Q] = p$$

where Q is some extreme negative quantile and p is the probability. Then, CoVaR_i is the value at risk of the financial system given that the institution i is under financial distress. i.e.,

$$\text{CoVaR}_i = \text{pr}[R_S \leq Q_S \,|R_i \leq Q_i] = p$$

If alternatively we use ES:

$$[R|R \leq Q]$$

we get MES; institutions expected equity loss in tail market outcomes:

$$\text{MES}_i = \text{E}[R_i|R_S \leq Q_S] \quad (1)$$

We could just as easily have defined MVaR as

$$\text{MVaR}_i = \text{pr}[R_i \leq Q_i|R_S \leq Q_S] = p$$

and CoES as

$$\text{CoES}_i = \text{E}[R_S|R_i \leq Q_i]$$

Table 1 provides a summary, which suggests that ultimately, regardless of the risk measure, or conditioning, the main determinant of the empirical performance of each measure is VaR.

Table 1: Summary.

Marginal risk measure	Condition on system	Condition on institution
	MVaR	CoVaR
VaR	$[R_i \leq Q_i\|R_S \leq Q_S] = p$	$[R_S \leq Q_S\|R_i \leq Q_i] = p$
	MES	CoES
ES	$[R_i\,\|R_S \leq Q_S]$	$[R_S\,\|R_i \leq Q_i]$

2.4.2 Empirical performance

The empirical performance of the MSRMs are important. To be reliable, they need to be to be of a high–quality, not just weakly better than a noise. Taking this into account, (Danielsson et al., 2011, 2012) study the model risk of MES and CoVaR and along with VaR by employing a sample of daily returns from the 92 largest financial institutions in the US from January 1997 to December 2010. Their analysis finds that the signals provided by these V–MSRMs are highly unreliable.

One of the reasons for the poor performance of the V–MSRMs is because the underlying VaR forecast is subject to a high degree of model risk at the best of times, which is even worse during heightened market uncertainty. The risk forecasts lag behind market outcomes and the statistical models perform poorly at the structural break when the markets transit from stable to a crisis period.

3. Computer based trading

The ongoing developments in CBT have affected the operations of financial markets in a profound way. In this survey we limit the analysis to how CBT may affect systemic risk, especially through extant MSRMs. For a comprehensive survey of CBT see Gomber et al. (2011).

3.1. Implications for market dynamics

Most of the existing literature focuses on the effects of HFT, or in general CBT, on market quality; whether CBT makes prices more informative, its effect on market liquidity or intraday volatility.

One of the earliest work of this kind is by Brogaard (2010) where he analyses the impact of HFT on the US equity market by using a sample of 120 Nasdaq stocks. In his sample, the group of HFT is identifiable. The results show that much of HFT activity follows a price reversal strategy to make higher profits. This in turn suggests a role of HFT in diminishing intraday volatility. Moreover, he shows that HFTs add significantly to the price discovery process. Although HFTs consume almost the same amount of liquidity they supply, they provide best bid and ask quotes for most of the trading day. Overall, those findings advocate HFT over non–HFT as an improved market quality.

Hasbrouck and Saar (2010) use order–level NASDAQ data and therefore are not able to identify HTF activity explicitly. Instead, they use messaging bursts as a proxy for strategic runs, which in turn is used as a measure of low–latency activity. Statistical analysis of 500 stocks for the period of October 2007–June 2008 suggests that increased high–frequency, at least, low–latency activity improves market quality measures such as short–term volatility, spreads, and depth in the limit order book.

By developing a theoretical model that allows high–frequency machine trading, Cvitanic and Kirilenko (2010) examine the effects of HFT on transaction prices. They find that in a market with a high–frequency trader, the distribution of transaction prices has more mass around the center and thinner extreme tails.

Hendershott et al. (2011) test whether AT improves liquidity by analyzing the time series evolution of AT in a sample of NYSE stocks spanning 2001–2005, finding that AT increases liquidity by narrowing the bid–ask spread, especially for large cap stocks by reducing adverse selection, or equivalently by increasing the informativeness of the quotes.

Similar question is considered by Hendershott and Riordan (2012). In a statistical analysis of 30 DAX stocks on the Deutsche Boerse for a total of 13 trading days in January 2008, the authors examine the role of algorithmic traders in liquidity supply and demand. They conclude that AT consumes liquidity when it is plentiful, i.e. when the inside spread is narrow and supply liquidity when the spreads are wide. Hence, CBT is beneficial for market liquidity since ATs react more quickly and actively compare to human traders.

Although many of the academic literature provide empirical evidence suggesting that CBT is beneficial for market quality, some studies reach a different conclusion. Zhang (2010) shows that HFT is positively correlated with stock price volatility and this positive correlation is even stronger during periods of high market uncertainty. Similarly, contrary to the findings mentioned above, he argues that HFT is harmful for price discovery.

Haldane (2011) notes in a speech that HFT is increasing abnormalities in the market prices, where the latter is proxied by the volatility and correlation. He concludes that instead of solving the liquidity problem, HFT even amplifies market stress and illiquidity since it adds liquidity to the market in the good times, but absorbs it during turmoil.

As an investigation into the HFT and algorithmic trading on Swedish equity market, Finansinspektionen, Financial Supervisory Authority of Sweden, conducted a survey with 24 companies; 10 Swedish banks and 14 large institutional investors. Results reveal that although the negative effects related to HFT is limited (for example increased volatility in the market), there is still considerable concern about market abuse among survey participants.

3.1.1 Relevance for systemic risk

While the question of the impact of CBT on securities markets generally is of interest in many applications, it is less clear how relevant it is to the question of systemic risk. For example, establishing that CBT increases volatility does not mean systemic risk is increased because the connection between volatility and systemic risk is tangental. The reason is that unconditional volatility as a measure of risk is generally only relevant if and only if returns on financial assets are Gaussian, which is not the case. At the daily frequency it has been recognized at least since the pioneering work of Mandelbrot (1963) and his student Fama (1963, 1965) that financial returns exhibit fat tails. Consequently, it is the tail thickness that provides information about the risk, including the systemic risk, not the volatility.

CBT also has implications for liquidity. Liquidity is a multifaceted concept and CBT only imparts on a small part of the concept, liquidity in high–speed trading environments. There, liquidity is measured either by spreads of various types or price impacts of relatively small trades. While such market liquidity is of fundamental importance to many market participants, the connection to the more general concept of liquidity and systemic risk is more tangental. There are however direct connections in special cases discussed below.

Most extant studies on the impact of CBT on market dynamics, such as those discussed above, are made with pre flash crash data. As a consequence, the focus of these studies is more on the general operation of markets rather than anomalous in crisis situations. While HFT might normally make markets, this can be quickly reversed during extreme turmoil. Clear mechanisms exist for CBT leading to extreme outcomes in financial markets. Rapid liquidation of positions to avoid further losses may accelerate the downward pressure on prices, aggravating a crisis situation. Based on several arguments such as feedback loops, large trading volumes and lack of regulations, several authors consider the post–flash crash period and underline the HFT as a potential source to systemic instabilities and crashes.

Zigrand et al. (2011) argue that feedback loops are the underlying force behind most of the financial crises and those loops are more likely to arise, or at least harder to supervise in AT environments. Feedback loops arise because in a stress situation many algorithms quickly coordinate and act simultaneously and feed each other. Under CBT, this coordination will be even quicker.

Similarly, Leland (2011) underlines forced selling as a possible source of positive feedback loops, hence the crashes. He notes that forced selling can lead to price declines which in turn force further selling and further price declines, a positive feedback situation that can lead to extreme market volatility and crashes. The 1929 crash, the 1987 crash, 1998 Long Term Capital Management collapse, the 2007 Slaughter of the Quants, and elements of the 2008 financial crisis are all discussed as examples of forced selling triggered market crashes. Finally, he concludes that the 1987 crash and 2010 flash crash are similar in nature, but not in speed; the latter is accelerated because of the HF traders who turn to be liquidity consumers rather than providers hence created rapid and further price drops.

Sornette and von der Becke (2011) question the value of liquidity provided by HFT by arguing that beyond a certain cutoff, excess liquidity may not be beneficial for the real economy, and can increase the risk of herding and possible systemic instabilities, even crashes. They conclude that HFT is likely to create a future crash.

Kirilenko et al. (2011) ask whether HF traders played any role in triggering the flash crash. In contrast to the findings of Sornette and von der Becke (2011), they clear HFT as a trigger, however, the study still shows that due to the large selling pressure on that day, presence of HFT exacerbated market volatility and drive withdrawals of liquidity in times of stress.

Farmer and Skouras (2011) study the impact of CBT on stability of the markets by employing an ecological microstructure perspective of financial markets. Since there is significant homogeneity and interdependence in AT, only a few algorithms and their interactions may dominate the market dynamics and can lead to systemic risk. The authors argue that this may be the underlying cause of the 1987 crash, the 2007 quant crisis and the 2010 flash crash.

A different approach is taken by Johnson et al. (2012) who analyze a set of 18,520 ultrafast large price movements (‘black swan events’) for the period 2006–2011. They conclude that a system–wide transition from a mixed human–machine phase (mixed phase of humans and machines, in which humans have time to assess information and act) to a all–machine phase is characterized by lower durations but more frequent black swan events. This does however not provide a channel from these high frequency price movements to medium frequency market outcomes more likely to be associated with systemic risk.

3.2. Relevance for systemic risk

These results suggest that a necessary condition for CBT to increase systemic risk is that it creates channels for adverse feedback loops whereby the relatively small events like micro crisis or sudden disappearance of liquidity, can be amplified into a systemwide market crisis. However, this is not a sufficiency condition. After all, one of the two major event caused by AT, the flash crash was not systemic. Furthermore, the fact that none of the high–frequency extreme events documented by Johnson et al. (2012) is associated with anything resembling a systemic event, further argues against CBT contributing to systemic risk.

The preponderance of available evidence argues against CBT being of systemic concern. This however does not lead us to the conclusion that one should disregard the systemic consequences of CBT. We discuss this in the next section.

4. Impact of CBT on MSRMs

The literature on SRMs and CBT has generally developed independently with authors not providing linkages between the two research areas in most cases, except perhaps in general statements. This is problematic because one would expect any systemic risk arising from CBT to manifest itself in the SRMs, in particular the MSRMs. This leads us to the following questions:

1. Does the presence of CBT affect price dynamics (typically at the daily frequencies) and in turn do these changes constitute a contribution to systemic risk?
2. Does CBT contribute to systemic risk in a way that is not captured by extant SRMs?
3. Can we use empirical and theoretical CBT modelling to improve MSRMs?

4.1. Effects of CBT on price dynamics

In order to address this issue, it is necessary to consider how high–frequency trading outcomes get aggregated into lower frequency price movements, such as daily returns. To help in the analysis, we identify three levels of aggregation of financial data.

The first level of aggregation is individual trading outcomes and data aggregated to very high frequencies, perhaps seconds.

The second level of aggregation is to intermediate frequencies, such as daily, the aggregation level often considered by financial institutions, supervisors and MSRMs. Without loss of generality, it could also include weekly or monthly, even annual aggregation of securities prices.

The final level of aggregation constitutes the impact of securities prices on other economic variables, such as macroeconomic data, investment decisions, funding liquidity, and the like.

We consider two cases below, both the situation where high frequency trading outcomes aggregate into noise, and also where they lead to feedback loops, culminating in extreme outcomes.

4.1.1 CBT aggregates to noise

Most of the time, aggregated outcomes from CBT resemble random noise. Under these conditions, no systemic risk would arise from CBT, since lower frequencies are more dangerous than higher frequencies, and mini crashes, microfractures and the like are self–healing, aggregating into essentially noise. Even large events like the flash crash aggregated into noise. This suggests that the systemic risk from CBT does not arise from the presence of very large short duration price changes, in the absence of acceleration mechanisms. Therefore, we can view risk as usually being *exogenous* in the terminology of Daníelsson and

Shin (2003), as compared to *endogenous*, even if high frequency data exhibits large jumps or other anomalies.

In some historical cases we have seen CBT aggregate into extreme market movements, most prominently in 1987. Regardless, 1987 being the largest price shock in modern history, it was not systemic, according to either the Trichet or BIS definitions discussed in Section 2 as it had neither adverse effects on growth and welfare and the economy nor caused brought financial difficulties. This illustrates a key fact that an extreme market movement does not by itself constitute an increase in systemic risk.

Furthermore, is plausible that if the speed of trading had been higher, the price recovery would have been faster – HFT increasing the speed of self healing.

4.1.2 CBT aggregates to extreme outcomes

CBT might create systemic risk when the aggregation of very high–frequency outcomes does not result in noise, but instead leads to a chain reaction whereby market participants suddenly coordinate in selling and buying the same assets, causing a structural break in price dynamics, because risk suddenly stops being approximately exogenous and becomes highly endogenous.

This might happen because of the interaction between unthinking algorithms following mechanical rules, heterogeneous trading venues and the speed of trading making real–time human supervision impossible. This creates informational and efficiency problems, since market participants not only focus on the basic worth of a security, but also on the beliefs about other traders' beliefs about yet other traders' beliefs and so forth.

The immediate impact of this structural break is the failure of models assuming the risk is exogenous, such as most extant MSRSs, at the precise time when they are needed the most.

4.2. Does CBT contribute to systemic risk in a way that is not captured by extant SRMs?

At the second level of aggregation, the focus is not on individual trading strategies but rather how prices across assets change from one day or week to the next. The problem of modelling risk at this frequency has been under intense study since the pioneering work of Engle (1982) and his ARCH model. A decade later this got translated into formal statistical risk measures, first Value–at–Risk (VaR), which remains the most widely used risk measure at the daily frequency. Following the crisis that started in 2007, we have seen several applications to systemic risk with the various MSRMs discussed in Section 2.4 above.

Even though these risk measures are derived from daily data, which is mostly just aggregated outcomes from HFT, the link between CBT and the SRMs is somewhat tenuous. As noted above, most of the time CBT generates noise that may not be all that relevant for systemic risk.

In this, there is a significant difference between the V–MSRMs and the EVT–MSRMs. The former derive from the daily price dynamics whilst the latter focus on unconditional extreme movements.

Given the analysis above, we find that a more likely chain of causality would be for CBT to result in an extreme market movement which in turn would affect price dynamics. Therefore, a more direct link of causality between CBT and the MSRMs is with EVT–MSRMs rather than the V–MSRMs.

However, the evidence seems to suggest that if CBT creates systemic risk it would not be captured by the MSRMs. The reason is that the MSRMs are based on using historical observed data to forecast the possibility of an extreme future market outcome. However, there is no historical data of CBT and HFT causing extreme outcomes (1987 is hardly fast). For this reason, if one is concerned with systemic risk arising from CBT, it is necessary to focus on high–frequency data, along with direct modeling of the possible feedback loops, rather than the medium frequency data, without any mechanism modeling, central to the MSRMs.

4.3. Implication for CBT–SRMs and possible improvements

Although not discussed in the extant literature, beyond perhaps general statements, there are several directions one could take in to develop a CBT–SRM, a SRM using CBT as a fundamental building block. An SRM incorporating elements from the EVT–MSRMs and considering endogenous feedback loops from CBT seems to us a sensible way forward.

On the CBT side, this involves explicit modeling of endogenous risk, arising for example from the interaction between algorithms. This might be done in a hybrid real world—simulated trading environment, whereby the researcher can explicitly identify the interaction between algorithms and the potential for endogenous adverse feedback loops. Johnson et al. (2012) is a useful first attempt at such modelling.

We suggest that a combination of such an approach with EVT could be fruitful. EVT is focused on the formal modelling of unconditional probabilities of extreme outcomes and extreme tail dependence (power law analysis is a subset of EVT). This might then help research in approaching the problem from both ends: aggregating from individual trading outcomes in the presence of endogenous adverse feedback loops from algorithms, and the probabilities of extreme aggregated outcomes from the other side.

5. Conclusion

In this paper we critically survey the recent empirical literature on the systemic risk and CBT along with its subset HFT. In particular, our focus is to explicitly connect the literature on systemic risk arise from CBT to the broader empirical literature on systemic risk.

We find that most MSRMs focus on modeling daily price dynamics and are not reliable as a tool to get at systemic risk. Similarly, most of the time HFT is not a systemic concern as high frequency extremes tend to aggregate into medium frequency innocuous noise.

There are however good reasons to believe that under special circumstances high frequency extremes might interact with trading algorithms to generate endogenous feedback loops, resulting in extreme medium frequency outcomes and even on to a systemic event. One way to study these phenomena is to combine the analysis of HFT with EVT to formalize the probabilities of extreme outcomes. This in turn might then be useful as an input into non-MSRMs.

References

Acharya, V. V., L. H. Pedersen, T. Philippon and M. Richardson, "Measuring Systemic Risk," (May 2010).

Adrian, T. and M. K. Brunnermeier, "CoVaR," Technical Report Staff Report no. 348, Federal Reserve Bank of New York, 2010.

Billio, M., M. Getmansky, A. Lo and L. Pelizzon, "Econometric measures of systemic risk in the finance and insurance sectors," (2010), NBER Working Paper, 16223.

BIS, "Annual Report, 1993-1994," Technical Report, Bank for International Settlements, 1994.

Bisias, D., M. Flood, A. Lo and S. Valavanis, "A Survey of Systemic Risk Analytics," Technical Report Working Paper, 0001, U.S. Department of Treasury, January 2012.

Brogaard, J. A., "High Frequency Trading and its Impact on Market Quality," (July 2010).

Brownlees, C. T. and R. Engle, "Volatility, Correlation and Tails for Systemic Risk Measurement," (June 2011).

Crockett, A., "Marrying the micro- and macro-prudential dimensions of financial stability," (2000), the General Manager of the Bank for International Settlements; http://www.bis.org/review/rr000921b.pdf.

Cvitanic, J. and A. Kirilenko, "High Frequency Traders and Asset Prices," (March 2010).

Danielsson, J., K. R. James, M. Valenzuela and I. Zer, "Model Risk of Systemic Risk Models," (December 2011), http://www.RiskResearch.org.

Danielsson, J., Danielsson, J., Kevin R., Valenzuela, M and, Ilknur, K. "Dealing With Systemic Risk when We Measure It Badly," (February 2012), http://www.RiskResearch.org.

Daníelsson, J. and H. S. Shin, "Endogenous Risk," in *Modern Risk Management – A History* (Risk Books, 2003), http://www.RiskResearch.org.

Drehmann, M. and N. Tarashev, "Measuring the systemic importance of interconnected banks," Technical Report Working Papers No 342, BIS, March 2011.

Engle, R. F., "Autoregressive conditional heteroskedasticity with estimates of the variance of the United Kingdom Inflation," *Econometrica* 50 (1982), 987–1007.

Fama, E., "Mandelbrot and the Stable Paretian Hypothesis," *Journal of Business* 36 (1963), 420–429.

Fama, E., "The Behavior of Stock-Market Prices," *Journal of Business* 38 (1965), 34–105.

Farmer, J. D. and S. Skouras, "An Ecological Perspective on the Future of Computer Trading," Technical Report Foresight Driver Review, DR 6, Foresight, 2011.

Gomber, P., B. Arndt, M. Lutat and T. Uhle, “High Frequency Trading,” (2011).

Gravelle, T. and F. Li, “Measuring Systemic Importance of Financial Institutions: An Extreme Value Theory Approach,” Technical Report 2011-19, Bank of Canada, 2011.

Haldane, A. G., “The race to zero,” Technical Report, Bank of England, July 2011.

Hartmann, P., S. Straetmans and C. G. de Vries, “Asset Market Linkages in Crisis Periods,” *The Review of Economics and Statistics* 86 (2004), 313–326.

Hasbrouck, J. and G. Saar, “Low–Latency Trading,” (November 2010).

Hendershott, T., C. M. Jones and A. J. Menkveld, “Does Algorithmic Trading Improve Liquidity? ,” *Journal of Finance* 66 (2011), 1–33.

Hendershott, T. and R. Riordan, “Algorithmic Trading and the Market for Liquidity,” (January 2012).

Huang, X., H. Zhou and H. Zhu, “Measuring Systemic Risk Contributions of Large and Complex Financial Institutions,” (2010).

Johnson, N., G. Zhao, E. Hunsader, J. Meng, A. Ravindar, S. Carran and B. Tivnan, “Financial black swans driven by ultrafast machine ecology,” (February 2012).

Kirilenko, A., A. S. Kyle, M. Samadi and T. Tuzun, “The Flash Crash: The Impact of High Frequency Trading on an Electronic Market,” (January 2011).

Kritzman, M., Y. Li, S. Page and R. Rigobon, “Principal Components as a Measure of Systemic Risk,” (2010).

Leland, H. E., “Leverage, Forced Asset Sales, and Market Stability: Lessons from Past Market Crises and the Flash Crash,” Technical Report Foresight Driver Review, DR 9, Foresight, 2011.

Mandelbrot, B. B., “The variation of certain speculative prices,” *Journal of Business* 36 (1963), 392 – 417.

Sornette, D. and S. von der Becke, “Crashes and High Frequency Trading,” Technical Report Foresight Driver Review, DR 7, Foresight, 2011.

Tarashev, N., C. Borio and K. Tsatsaronis, “Attributing systemic risk to individual institutions,” Technical Report Working Papers No 308, BIS, May 2010.

Trichet, J. C., “Systemic Risk,” (2009), http://www.ecb.int/press/key/date/2009/html/sp091210_1.en.html.

Vasicek, O. A., “The Limiting Loan Loss Probability Distribution,” KMV Working Paper (1991).

Zhang, X. F., “High–Frequency Trading, Stock Volatility, and Price Discovery,” (December 2010).

Zhou, C., “Are banks too big to fail? Measuring systemic importance of financial institutions,” *International Journal of Central Banking* 6 (2010), 205–250.

Zigrand, J.-P., H. S. Shin and D. Beunza, “Feedback Effects and Changes in the Diversity of Trading Strategies,” Technical Report Foresight Driver Review, DR 2, Foresight, 2011.

Glossary

AT algorithmic trading.

CBT computer-based trading.

HFT high–frequency trading.

SRM systemic risk measure. A statistical device to quantify systemic risk.

MSRM market data based systemic risk measure. An SRM using market prices, typically daily, to quantify systemic risk.

VaR Value–at–Risk. A quantile of return distributions, typically 1% or 5% (99% or 95%) for daily data. The main market risk measure, used in regulations and internal risk control.

ES expected shortfall, also known as tail VaR. Expected losses conditional on losses exceeding VaR.

EVT extreme value theory. The probability theory of power laws.

V–MSRM MSRMs that use volatilities and/or VaR (or ES which is a function of VaR) as a fundamental building block.

EVT–MSRM MSRMs that use EVT as a fundamental building block.

CBT–SRMs MSRMs that use CBT as a fundamental building block.

Foresight
1 Victoria Street
London SW1H 0ET
www.foresight.gov.uk

URN 12/1062