

Process Systems Engineering as a Modeling Paradigm for Analyzing Systemic Risk in Financial Networks

Richard Bookstaber

Office of Financial Research

richard.bookstaber@treasury.gov

Paul Glasserman

Office of Financial Research and
Columbia University

pg20@columbia.edu

Garud Iyengar

Columbia University

garud@ieor.columbia.edu

Yu Luo

Columbia University

yl2750@columbia.edu

Venkat Venkatasubramanian

Columbia University

venkat@columbia.edu

Zhizun Zhang

Columbia University

zz2216@columbia.edu

The Office of Financial Research (OFR) Working Paper Series allows members of the OFR staff and their coauthors to disseminate preliminary research findings in a format intended to generate discussion and critical comments. Papers in the OFR Working Paper Series are works in progress and subject to revision.

Views and opinions expressed are those of the authors and do not necessarily represent official positions or policy of the OFR or Treasury. Comments and suggestions for improvements are welcome and should be directed to the authors. OFR working papers may be quoted without additional permission.

ABSTRACT

Financial instability often results from positive feedback loops intrinsic to the operation of the financial system. The challenging task of identifying, modeling, and analyzing the causes and effects of such feedback loops requires a proper systems engineering perspective lacking in the remedies proposed in recent literature. We propose that *signed directed graphs* (SDG), a modeling methodology extensively used in process systems engineering, is a useful framework to address this challenge. The SDG framework is able to represent and reveal information missed by more traditional network models of financial system. This framework adds crucial information to a network model about the direction of influence and control between nodes, providing a tool for analyzing the potential hazards and instabilities in the system. This paper also discusses how the SDG framework can facilitate the automation of the identification and monitoring of potential vulnerabilities, illustrated with an example of a bank/dealer case study.

Modern financial systems are characterized by a complex set of interdependencies among a large number of institutions. Stress to one part of the system can spread to others, often threatening the stability of the entire financial system. The critical need for a fundamental understanding of the structure and dynamics of this system has been emphasized by the recent financial crisis precipitated by counterparty exposures revealed by the Lehman bankruptcy and the near-bankruptcy of AIG, as well as the European debt crisis caused by the exposure of European banks to sovereign default risk. In the aftermath of the 2008 crisis, regulators have come to recognize that interconnectedness can pose substantial threats to the stability of the financial system.

Financial instability typically results from positive feedback loops intrinsic to the operation of the financial system; the instability results from responses to shocks that reinforce and amplify the initial shock. The structures and mechanisms that create the positive feedback must, therefore, be a focus of analysis of financial stability, and new tools are needed to identify and model these structures and mechanisms.

In addition, under extreme circumstances the steps taken by individual agents to mitigate the risk of financial systems can become the very source of destabilizing positive feedback through their interaction with other agents. We refer to these steps as *locally* stabilizing but *globally* destabilizing. This phenomenon is illustrated by bank runs. Suppose a bank is weakened by losses. The prudent action for each individual depositor is to withdraw funds, yet this very response will drive the bank to failure if followed by every depositor (Diamond and Dybvig [1983]). The longer the line of customers outside grows, the greater the incentive for more customers to join the line, and the stronger the amplifying feedback.

The problem of traditional bank runs was largely solved through deposit insurance, which effectively eliminates any reason for depositors to react to news about a bank. However, similar dynamics operate throughout the financial system. For example, a bank/dealer facing a shortfall in funding might reduce the lending it provides to hedge funds, and to control their risk, the hedge funds might respond by liquidating positions. But this circuit of actions, reasonable and prudent for each of the two sectors, can lead to global instability: The resulting decline in prices reduces the value of collateral, reducing the cash provided to the bank/dealer on the one hand, and leading to further margin calls and demand for forced liquidation by the hedge funds on the other.

Examples of these patterns have been identified as fire sale dynamics.¹ But to understand these critical aspects of the financial system comprehensively, we need a systematic way to identify the paths of feedback globally wherever they may arise. To do so, one must understand the conduits for the transmission of information and the control mechanisms applied by the various financial entities based on their observations of flows and the financial environment. A further complicating fact is that the nature of this feedback is scale dependent. For example, a small change in prices, funding, or a bank's financial condition might be absorbed by the system, but a large shock might trigger a destabilizing cascade.

We introduce signed directed graphs (SDGs) as a tool for understanding the feedback effects in financial systems. SDGs are extensively used in process systems engineering. An SDG representation captures the information transmission, environmental state, and causal rela-

¹ See Shleifer and Vishny [2010], and Brunnermeier and Pedersen [2009] for liquidity spirals, Adrian and Shin [2013] and Fostel and Geanakoplos [2008] for leverage cycles, and Gorton [2009] for panics.

tionships that underlie feedback. It encodes the control rules and responses, followed by individual units within a financial system, and provides a framework for systematically investigating the resulting interactions between these units. In particular, the SDG representation can be used to identify cycles of positive feedback that may not be immediately apparent, and to pinpoint areas of potential stress and instability in a systematic manner.

The SDG framework is able to represent and reveal information missed by more traditional network models of financial interconnections. Network models typically describe payment obligations and flows, and they can be effective in quantifying the degree and complexity of the connections among the financial entities. Standard network models represent financial entities as nodes and the flows between them as edges. Research questions in this area focus on which types of networks provide robust structures for the financial system (Allen and Babus [2009]; Battiston et al. [2013]; Gai and Kapadia [2010]). But these models lack a representation for the flow of information and responses to information; they do not provide a vehicle for understanding how responses and controls of multiple agents interact or the inner workings of an institution summarized by a single node.


In engineering systems, safety and stability are design criteria. In contrast, the financial system is self-organized. Individual financial entities generally have risk management procedures and controls to preserve their own stability, but the system as a whole was never engineered for safety and stability. Because of this, it is all the more critical to understand the paths of positive and negative feedback, alternative routes for funding and securities flows in the event of a shock to one node or edge of the network, and more generally, how the interactions of the system can create vulnerabilities and instability.

This paper shows how the SDG framework makes this possible through a system-wide view of transformations and dynamical interactions in the financial system. With an SDG representation, it becomes possible to automate the systematic identification and monitoring of vulnerabilities. In particular, this approach contributes to the critical task of systemic financial risk management: It can highlight, and help us monitor, dynamics such as fire sales and funding runs where actions that are locally stabilizing might cascade to be globally destabilizing.

Financial Network as a Process Plant: A Systems Engineering Framework

An appropriate process systems engineering analogy is to view each financial entity as a production or manufacturing plant, for example, as a chemical process plant that takes securities and funding as input and creates new financial products as outputs delivered to other processing units. This analogy opens the possibility of using tools that are applied in engineering for network analysis to gain a better understanding of the dynamic process underlying the financial system. Although researchers have suggested the Internet, electrical power grid, and transportation network as potential models for the financial system, none of these have the richness of phenomena seen in a large-scale chemical process plant. Various physical/chemical transformations, feedback and recycle loops, etc., can serve as relevant and useful analogies for modeling the financial system. In the existing network-based models risk travels along edges. However, these models ignore the financial transformations executed within the nodes that generate and compound risk. Although flows and connections are important, the picture of risk creation and contagion is incomplete without understanding the control of the production process.

Feedback control is a fundamental concept in process engineering, and it provides a useful setting to illustrate the SDG framework. Consider the temperature in a reactor that is modulated through coolant flow. This mechanism can be illustrated through an SDG as follows:




The circle on the left represents the temperature inside the reactor; the circle on the right represents the coolant flow rate. A solid arrow means that a change in one variable causes a change in the other variable in the same direction. A dashed line means that the effect is in the opposite direction. The figure illustrates stabilizing feedback: an increase in temperature causes an increase in coolant flow rate, which causes a *decrease* in temperature, which then lowers the flow rate. The control of an entire chemical plant can be described by assembling these types of building blocks, and the result is useful for fault diagnosis and process hazards analysis.² Because SDG models are qualitative in nature, they can lead to ambiguities and are limited to certain kinds of tasks. The key point for our purpose is that they illustrate controls or influences, and not material flow.

² The use of SDG in process hazard analysis in chemical engineering applications is presented in Venkatasubramanian [2000] and Zhao et al. [2005a, 2005b]. Some of the limitations of SDG due to their qualitative nature are discussed in Venkatasubramanian et al. [2003a, 2003b, 2003c].

Contrast this example with what would happen if both lines were dashed. In that case, the temperature increase would cause a *decrease* in the coolant flow rate, causing a further temperature increasing and, indeed, causing the temperature to rise out of control.

Now consider the following SDG, representing a basic economic relationship:



Here, P denotes the price of some good, and Q denotes the quantity of the good supplied by the market. A price increase pushes supply up, and a supply increase pushes price down, so the two influences are mutually stabilizing. If both arrows were solid, both price and quantity would spiral out of control, as in a bubble or a crash.

As these examples indicate, we can get valuable information about the stability of a system by examining loops of positive and negative influence: a loop that is net negative is stabilizing, and one that is net positive is destabilizing. We will apply this idea to study a financial network in the next section.

The SDG model of an entire industrial process is naturally very complicated, with hundreds of nodes and edges. It can be assembled from a library of unit-wise SDG models and applied using artificial intelligence-based systems that automate much of the cause-and-effect reasoning. These methods can be adapted for developing a process systems engineering framework for modeling and analyzing risk in financial networks. The goal is to develop automated systems that can identify the potential hazards lurking in a complex financial network by systematically examining various “what if” failure scenarios.

SDG Modeling Framework for Financial Networks




Exhibit 3: Simplified Bank/Dealer Network

Source: Aguiar, Bookstaber, and Wipf [2014]

We now explain how SDG models can be used to analyze the dynamics of financial systems. The financial system can be represented in a manner that is analogous to a processing plant by mapping the flows of funding, assets, and collateral through the various financial agents and delineating the transformations the agents perform on those flows (Aguiar, Bookstaber, and Wipf [2014]). A bank/dealer acts as an intermediary between buyers and sellers of

securities and between lenders and borrowers of funding. Its clients are investors, such as asset management firms, hedge funds, and pension funds, as well as other bank/dealers. There are specific business units within the bank/dealer that process funding and securities to create products for these clients. The bank/dealer's network, with its connections to other financial entities and between its business units, is complex. To demonstrate the process systems engineering inspired modeling framework, we now consider a simplified version of the reality and focus only on two types of bank/dealer activities shown in Exhibit 3:

1. Funding and securities lending: The Bank/Dealer goes to sources of funding such as money market funds through the repo market, and to security lenders such as pension funds and asset management firms through their custodian banks.
2. Providing liquidity as a market maker: The Bank/Dealer goes to the asset markets, to institutions that hold assets, and to other market makers to acquire positions in the securities that clients demand. This function also includes securitization taking securities and restructuring them. This involves liquidity and risk transformations.

The functions we show within the Bank/Dealer include the Prime Broker, which lends cash to hedge funds in order for the hedge funds to buy securities on margin; the Financing Desk, which borrows cash with high-quality securities used as collateral; and the Trading Desk, which manages inventory in its market making activities that it finances through the Financing Desk. The Bank/Dealer interacts with Cash Providers, such as money market funds, pension funds, and insurance companies; other bank/dealers through the over-the-counter market, which is the market for the Bank/Dealer to acquire or lay off inventory; and hedge funds, which seek leverage and securities from Prime Brokers to support their long and short trading posi-

tions. Hedge funds also represent the wider swath of institutional customers that use the Bank/Dealer's market making function, ranging from asset managers and hedge funds to pension funds, sovereign wealth funds, and insurance companies.

The interactions between the Bank/Dealer's functional areas create various transformations, like parts of a processing plant. The Financing Desk takes short-term loans from the Cash Providers and passes them through to clients with lower credit standing, often as longer-term loans. In doing this, the Bank/Dealer is engaging in both a maturity and a credit transformation. The Trading Desk inventories securities until it can either lay it off based on the demand of another client or to the over-the-counter market. In doing this, it provides a liquidity transformation.

The network for the Bank/Dealer is more interconnected than that of a chemical plant, because some clients, which are the nodes that receive the output from a bank/dealer, are also sources of inputs. A Hedge Fund borrowing in order to buy securities might also be lending other securities. A pension fund providing funding might also be using the Bank/Dealer for market making. Hedge funds and related institutional investors are on both sides of the production because they are both buyers and sellers of securities, and in that sense they provide inputs as well as output in market making.




Exhibit 4: SDG Model for Bank/Dealer Example

Bank/Dealer Case Study

The network depicted in Exhibit 3, although illustrative of the layout of the components of the Bank/Dealer and its interactions, does not represent the effect of the various flows and cannot by itself suggest conditions and areas where a disruption will create instability through positive feedback cycles. To achieve this, we need a cause-and-effect representation of this network, as we did in the chemical processing example of the previous section. We accomplish this by creating the SDG model for this network displayed in Exhibit 4.

For simplicity, we consider a system with a single market asset (such as a stock or a bond). Its price is represented by the node P_{BDM} , and this price level influences, and is influenced by, the rest of the system. Quantities of the asset Q_{HF} and Q_{TD} are held by the Hedge

Fund and Trading Desk. These units need funding to finance their asset holdings. The funding is provided by the Money Market, the Prime Broker, and the Finance Desk. In each case, funding availability depends on the unit's collateral level, and collateral is held in the form of the market asset. Changes in the market price change the value of the collateral, which in turn changes the level of funding available. A margin rate controls the ratio of funding capacity to collateral at the Money Market and the Prime Broker; a leverage target controls the level of borrowing relative to asset holdings at the Hedge Fund and the Trading Desk. Specifically, the Hedge Fund determines its dollar borrowing based on the availability of loans that are provided through the Prime Broker and a comparison of its assets to its target leverage ratio, λ . The Prime Broker's lending is determined by the Bank/Dealer's Financing Desk and by the Prime Broker's margin rate, χ .

The Trading Desk provides a market making function; it stands ready to take on any quantity sent its way by the hedge fund. This increases its inventory of shares, and when this inventory becomes too large relative to a set point, it opens the overflow control to pass shares through to the market, dropping the price as a result. The Trading Desk's market making function distinguishes its control mechanism from that of the Hedge Fund. As with the hedge fund, the Trading Desk depends on the Financing Desk to fund its inventory, and a drop in funding might force the Trading Desk to release more shares into the Bank/Dealer Market.

The Money Market provides funding for both the Hedge Fund and the Trading Desk through the Finance Desk; and it is changes in the funding of the Funding Desk that lead to changes in the quantity held by the Hedge Fund and the Trading Unit, ultimately changing the price.

The entire system is driven by, and feeds back into, the prices set in the bank/dealer market. These prices are determined by the actions of the Trading Desk and the Hedge Fund, and determine the collateral value that helps drive the willingness of the various agents along the path to provide funding.

The SDG model clearly illustrates why the financial system becomes embroiled in one crisis after another: Nearly all of the pathways extending from the Money Market through the bank/dealers to the hedge funds are positive, so a shock to one node may create a positive feedback, exacerbating the shock. This can be seen by applying the SDG framework and its associated process hazard analysis methodology to the two most common sources of financial crisis — funding runs and fire sales.

Process hazards analysis (Venkatasubramanian et al. [2000]; Venkatasubramanian [2011]; Zhao et al. [2005a, 2005b]) is a methodology for systematically identifying abnormal causes and adverse consequences that can occur anywhere in the process system. In the context of an SDG model, process hazards analysis provides the framework that can guide us in identifying methodically what can go wrong at each node and edge and how that failure would propagate through the rest of the system. Given the self-organized nature of financial networks, here we focus on identifying and examining feedback loops in an SDG model. The complete list can be computed via a depth-first search of the SDG (Russell and Norvig [2003]). Not all positive loops are necessarily significant sources of vulnerability because the edges of the SDG record the direction of influence but not its magnitude. An individual node is typically subject to multiple competing effects, so the net effect ultimately depends on the gain associated

with reach feedback loop. However, the list of loops provides a valuable tool for identifying vulnerabilities; indeed, we know of no other systematic approach to this problem.

Exhibit 5 gives a complete list of loops for the SDG model of the bank/dealer network, with each row describing a loop. A positive (negative) loop is one in which the product of the signs along the edges defining the loop is positive (negative). Only the last two loops in the table are negative, and these have a simple interpretation: They are the internal risk management processes of the Hedge Fund and the Trading Desk. Each of these units uses a leverage target as an internal control for the quantity held of the market asset. But when we combine these stabilizing negative feedback loops with the rest of financial system, we get a range of potentially destabilizing positive feedback loops through the interactions across units. We will examine two types of positive loops in greater detail, because these represent fire sales and funding runs, two key examples of crisis dynamics. We emphasize that these dynamics are discovered automatically by the SDG analysis, which highlights the value of this approach.

Exhibit 5: List of loops

Index	Sign	Loop
1	+	$[P_{BDM}, C_{MM}, F_{MM}, V_{FD}, V_{PB}, L_{HF}, Q_{HF}, Q_{TD}, \lambda_{TD}, \epsilon_{TD}, P_{BDM}]$
2	+	$[P_{BDM}, C_{MM}, F_{MM}, V_{FD}, V_{PB}, L_{HF}, Q_{HF}, P_{BDM}]$
3	+	$[P_{BDM}, C_{FD}, V_{FD}, V_{PB}, L_{HF}, Q_{HF}, Q_{TD}, \lambda_{TD}, \epsilon_{TD}, P_{BDM}]$
4	+	$[P_{BDM}, C_{FD}, V_{FD}, V_{PB}, L_{HF}, Q_{HF}, P_{BDM}]$
5	+	$[P_{BDM}, C_{PB}, V_{PB}, L_{HF}, Q_{HF}, Q_{TD}, \lambda_{TD}, \epsilon_{TD}, P_{BDM}]$
6	+	$[P_{BDM}, C_{PB}, V_{PB}, L_{HF}, Q_{HF}, P_{BDM}]$
7	+	$[P_{BDM}, \lambda_{HF}, L_{HF}, Q_{HF}, Q_{TD}, \lambda_{TD}, \epsilon_{TD}, P_{BDM}]$
8	+	$[P_{BDM}, \lambda_{HF}, L_{HF}, Q_{HF}, P_{BDM}]$
9	+	$[P_{BDM}, C_{MM}, F_{MM}, V_{FD}, \lambda_{TD}^{SP}, \epsilon_{TD}, P_{BDM}]$
10	+	$[P_{BDM}, C_{FD}, V_{FD}, \lambda_{TD}^{SP}, \epsilon_{TD}, P_{BDM}]$
11	+	$[\chi_{PB}, V_{PB}, L_{HF}, Q_{HF}, \chi_{PB}]$
12	+	$[P_{BDM}, \lambda_{TD}, \epsilon_{TD}, P_{BDM}]$
13	-	$[\lambda_{HF}, L_{HF}, Q_{HF}, \lambda_{HF}]$
14	-	$[\epsilon_{TD}, Q_{TD}, \lambda_{TD}, \epsilon_{TD}]$

Fire Sales

Exhibit 6 shows a segment of the SDG model of Exhibit 4 that focuses on the interaction of the Hedge Fund with the Bank/Dealer's Prime Broker. The fire sale occurs when there is a disruption to the system that forces a hedge fund to sell positions. As Exhibit 6 shows, this disruption can occur through three channels: a price drop and resulting drop in asset value, an increase in the margin rate that leads to a margin call from the Prime Broker, or a drop in the loan capacity of the Prime Broker. As the Hedge Fund reduces its assets, prices drop again, leading to a second (and subsequent) round of feedback, which makes the situation worse in each iteration.




Exhibit 6: SDG Model for Bank/Dealer Fire Sale Example

The fire sale is best depicted by the two loops listed in Exhibit 7. The first of these loops shows a price shock increasing the leverage of the Hedge Fund. The Hedge Fund then reduces

its holdings to decrease its leverage, which drops prices. The second loop has the same effect, the drop in prices increases leverage, which in turn leads to a drop in the quantity held by the Hedge Fund, but the effect in this case works its way through the Trading Desk. The quantity sold by the Hedge Fund raises the quantity held by the Trading Desk, increasing its lambda. This in turn leads the Trading Unit to sell into the market, with the end result being a further drop in prices.

Exhibit 7: Fire sale loops

Index	Sign	Loop
7	+	$[P_{BDM}, \lambda_{HF}, L_{HF}, Q_{HF}, Q_{TD}, \lambda_{TD}, \epsilon_{TD}, P_{BDM}]$
8	+	$[P_{BDM}, \lambda_{HF}, L_{HF}, Q_{HF}, P_{BDM}]$

Note that each of the units is acting to maintain stability: the Prime Broker is keeping its loans within bounds given its collateral; the Hedge Fund is maintaining a target level of leverage to control its risk, and the Trading Desk is governing its inventory level through an outflow if its market-making activities increase its inventory above a target level. Yet the stabilizing activities at the local level still lead to instability at the global level. This underscores a central point in the functioning of the financial system, namely that it can exhibit global instability even in the face of each unit acting to control its risk.

Funding Runs




Exhibit 8: SDG Model for Bank/Dealer Funding Run Example

Exhibit 8 shows another segment of Exhibit 4, focusing on the interaction of the Bank/Dealer with the Money Market. A funding run can be triggered by a disruption in funding flows from the Money Market. This may happen if there is an increased uncertainty about the quality of the collateral, or a drop in the market value of collateral, or by a change in the Money Market's margin rate, which might occur due to an erosion of confidence. The drop in funding negatively affects the amount of inventory the Trading Desk can carry, and as a result, it sells into the market. As is the case with dynamics associated with fire sales, selling drops prices, which feeds back to the value of collateral, and can precipitate a further reduction in funding from the Money Market.

The funding run is demonstrated by the two loops in Exhibit 9 that focus on the effect of a price drop on the collateral held by the Money Market. The price shock drops the value of the collateral being held by the Money Market, which reduces the funding available to the Bank/Dealer's Finance Desk. This has two effects: In Loop 2 it feeds through to ultimately reduce the funding available to the Hedge Fund through the Prime Broker, forcing a reduction in quantity held, and further reducing price. In Loop 9 the reduction in funding from the Money Market reduces the funding available to the Trading Desk, and its reduction in inventory again leads to a further price drop. These are only two of the possible loops where price-induced drop in funding leads to asset sales and more price drops. For example, the drop in collateral value can affect the Finance Desk directly.

Exhibit 9: Funding run loops

Index	Sign	Loop
2	+	$[P_{BDM}, C_{MM}, F_{MM}, V_{FD}, V_{PB}, L_{HF}, Q_{HF}, P_{BDM}]$
9	+	$[P_{BDM}, C_{MM}, F_{MM}, V_{FD}, \lambda_{TD}^{SP}, \epsilon_{TD}, P_{BDM}]$

In both fire sales and funding runs, the SDG model identifies a critical dynamic of that leads to market crises: Actions that dampen risk on a local level can contribute positive feedback and cascades on the global level. The proper response for the Prime Broker when faced with a reduction in funding is to reduce funding to hedge funds. But this leads to actions by hedge funds that contribute to a positive feedback cycle that reduces funding for the prime broker even more. Similarly, a locally proper response for the Trading Desk in the face of lower funding is to reduce inventories, but this leads to a drop in prices that feeds back to affect the value of collateral and reduces funding even further.

The unintended consequences are even more widespread than this. There are links between the segments representing fire sales and funding runs, so a funding run might precipitate a fire sale, and vice versa. From the SDG model, it is clear that a fire sale can lead to funding run, if the fire sale by the Hedge Fund drops prices to the point that the Cash Provider, seeing erosion in their collateral, begin to reduce funding. SDG model also shows that there is pathway in the opposite direction: drop in funding to the Trading Desk to lead to a reduction in inventory, causing a drop in prices which reduces the value of the Hedge Fund portfolio, leading the Prime Broker to increase its margin level, inducing a forced sale. The forced sale will add yet another positive feedback loop to the initial price impact that came from the Trading Desk. So actions that are reasonable locally can contribute to adverse global consequences.

For the simplified map of the Bank/Dealer network in Exhibit 3, one can perhaps manually identify and analyze all the feedback loops listed in Exhibit 5. However, for a network based on a more realistic map, such as shown in Exhibit 10, with multiple hedge funds, banks-dealers, and clients, various derivatives, as well as structured products, it is virtually impossible to identify and analyze all such loops manually, which again highlights the need for the SDG framework that can be automated to handle larger systems.




Exhibit 10: More Realistic Bank/Dealer Configuration

Source: Aguiar, Bookstaber, and Wipf [2014]

A further advantage is that the framework allows us to formulate more sophisticated models as necessary in a methodical manner. For instance, we now show how we can add numerical gains (Vaidhyanathan and Venkatasubramanian [1996]) on all the edges connecting various nodes and perform a quantitative analysis of how shocks of different magnitudes might propagate through the system. The gains used in this example are for illustrative purposes only and are not meant to reflect actual market conditions. In practice, these gains can be estimated using a combination of historic market data and the judgment of experienced market professionals.

Semiquantitative Analysis

Consider a loop of the form $(v_1, v_2, \dots, v_n, v_{n+1} = v_1)$ where each pair of nodes (v_i, v_{i+1}) is connected by a directed edge. Suppose the value of node v_{i+1} as a function of the value of node v_i is given by the functional relationship $v_{i+1} = f_i(v_i)$. The semiquantitative analysis proceeds in two steps:

1. Initiate a disturbance at node v_1
2. Propagate the deviation through the nodes v_2, v_3, \dots, v_n back to $v_{n+1} = v_1$.

We are interested in quantifying whether the loop amplifies or diminishes the initial disturbance.

Let $\delta v_i = \Delta v_i / v_i$ denote the relative change in the value of node i . Then

$$\begin{aligned}
 \delta v_i &= \frac{\Delta v_i}{v_i} \\
 &= \frac{f_{i-1}(v_{i-1}(1 + \delta v_{i-1})) - f_{i-1}(v_{i-1})}{f_{i-1}(v_{i-1})} \\
 &= \frac{f_{i-1}(v_{i-1}(1 + \delta v_{i-1}))}{f_{i-1}(v_{i-1})} - 1 \\
 &\equiv F_{i-1}(\delta v_{i-1}; v_i)
 \end{aligned} \tag{1}$$

Thus, the relative change in the value δv_i is a function of both the relative change δv_{i-1} and the current value v_{i-1} . Note that when $f_{i-1}(v_{i-1})$ is linear, i.e., $f_{i-1}(v_{i-1}) = k_{i-1}v_{i-1}$, the function $F_{i-1}(\delta v_{i-1}) = \delta v_{i-1}$. In the sequel, we will suppress the dependence on the current value v_{i-1} .

We will denote δv_{n+1} , i.e., the relative disturbance in the value of node v_1 after one iteration through the loop, by $\delta v_{1,f}$. From (1) it follows that

$$\delta v_{1,f} = F_n \left(F_{n-1} \left(\dots F_1(\delta v_1) \right) \right) \tag{2}$$

For linear relationships, (i.e., F_i is replaced by a constant gain k_i)

$$\delta v_{i+1} = F_i(\delta v_i) = k_i \delta v_i$$

Thus, when a loop contains only linear edges,

$$\delta v_{1,f} = k_n k_{n-1} \dots k_1 \delta v_{1,f}$$




Exhibit 11: Loop 7 as an example

We now illustrate this approach on Loop 7 displayed in Exhibit 11. Suppose the starting node $v_1 = P_{BDM}$. Our goal is to determine the relative change in the value of $v_1 = P_{BDM}$ after one iteration. We assume that the market conditions are described as follows:

$$P_{BDM} = \$10$$

$$C_{HF} = \$1 \text{ billion}$$

$$C_{TD} = \$1 \text{ billion}$$

$$A_{PB} = \$5 \text{ billion}$$

$$A_{HF} = \$5 \text{ billion}$$

$$A_{TD} = \$15 \text{ billion}$$

$$A_{FD} = A_{PB} + A_{TD} = \$20 \text{ billion}$$

$$L_{HF} = A_{HF} - C_{HF} = \$4 \text{ billion}$$

$$L_{TD} = A_{TD} - C_{TD} = \$14 \text{ billion}$$

$$Q_{HF} = 500 \text{ million shares}$$

$$Q_{TD} = 1.5 \text{ billion shares}$$

$$\chi_{MM} = 25\%$$

$$\chi_{PB} = 25\%$$

These values are chosen simply to illustrate the methodology; we do not claim that the values chosen are representative of true market conditions. We will first compute the functions $F_i(\delta v_i)$ for each of the nodes, and then compute the feedback effect.

1. $\delta \lambda_{HF} = F_1(\delta P_{BDM})$. The leverage

$$\begin{aligned} \lambda_{HF} &= \frac{1}{1 - L_{HF}/A_{HF}} \\ &= \frac{1}{1 - L_{HF}/(P_{BDM}Q_{HF})} \\ &\equiv f_1(P_{BDM}) \end{aligned}$$

From (1), it follows that

$$F_1(\delta P_{BDM}) = \frac{-L_{HF}\delta P}{P_{BDM}Q_{HF}(1 + \delta P) - L_{HF}}$$

2. $\delta L_{HF} = F_2(\delta \lambda_{HF})$. The relationship between L_{HF} and λ_{HF} is as follows. The price change δP_{BDM} results in a change in the leverage λ_{HF} ; this change triggers a trade since

the hedge fund is targeting a fixed leverage λ_{HF} . Thus, the hedge either takes on more loan or pays down some of the loan in order to reset the leverage back to λ_{HF} . Thus, the relative change δL_{HF} can be computed from the relation

$$\lambda_{HF} = \frac{A_{HF}(1 + \delta P_{BDM}) + \delta L_{HF}L_{HF}}{A_{HF}(1 + \delta P_{BDM}) - L_{HF}}$$

i.e.,

$$\delta L_{HF} = \frac{A_{HF}(\lambda_{HF} - 1)}{L_{HF}}(1 + \delta P_{BDM}) - \lambda_{HF}$$

Using the relationship that $\delta \lambda_{HF} = F_1(\delta P_{BDM})$ it follows that

$$F_2(\delta \lambda_{HF}) = \frac{A_{HF}(\lambda_{HF} - 1)}{L_{HF}}(1 + F_1^{-1}(\delta \lambda_{HF})) - \lambda_{HF}$$

3. $\delta Q_{HF} = F_3(\delta L_{HF})$, $\delta Q_{TD} = F_4(\delta Q_{HF})$, and $\delta \epsilon_{TD} = F_6(\delta \lambda_{TD})$. The functions f_3 , f_4 , and f_6 are all linear; therefore, it follows that $F_3(\delta L_{HF}) = \delta L_{HF}$, $F_4(\delta Q_{HF}) = \delta Q_{HF}$, and $F_6(\delta \lambda_{TD}) = \delta \lambda_{TD}$.
4. $\delta \lambda_{TD} = F_5(\delta Q_{TD})$. When the trading desk purchases (or sells) shares the capital C_{TD} of the trading desk decreases (or increases), and; the relationship is linear. Therefore, $\delta C_{TD} = -\delta Q_{TD}$. The relative change in leverage δL_{TD} is given by

$$\begin{aligned} \delta \lambda_{TD} &= \frac{\frac{A_{TD}}{C_{TD}(1 + \delta C_{TD})} - \frac{A_{TD}}{C_{TD}}}{A_{TD}/C_{TD}} \\ &= \frac{-\delta C_{TD}}{1 + \delta C_{TD}} \end{aligned}$$

Therefore, it follows that

$$F_5(\delta Q_{TD}) = \frac{\delta Q_{TD}}{1 - \delta Q_{TD}}$$

5. $\delta P_{BDM} = F_7(\delta \epsilon_{TD})$. The relationship between P_{BDM} and ϵ_{TD} is as follows. So long as $\epsilon_{TD} \leq 0$, i.e., the trading desk leverage λ_{TD} is less than or equal to the leverage set point λ_{TD}^{SP} , no action is taken. However, when the $\epsilon_{TD} > 0$, the trading desk sells assets to reset the error $\epsilon_{TD} = 0$. This trading impacts the price P_{BDM} . Thus, there is a complex nonlinear relationship between $\delta \epsilon_{TD}$ and δP_{BDM} that needs to be calibrated from data. For the purpose of illustrating SDG approach, we assume

$$F_7(\delta \epsilon_{TD}) = \begin{cases} 0.1 \delta \epsilon_{TD} & \text{normal market condition} \\ 2 \delta \epsilon_{TD} & \text{crisis conditions} \end{cases}$$

Now we are in a position to compute the loop gain $\delta P_{BDM,f} / \delta P_{BDM}$ using (2) and the nominal market condition described above. $\delta P_{BDM,f}$ can be determined for a given $\delta P_{BDM,i}$.

Exhibit 12 reports the loop gains for all the 14 loops for both normal and crisis conditions, and for small (1 percent) and large (5 percent) initial decrease. Specifically, for Loop 7 under normal market conditions, a 1 percent initial decrease in P_{BDM} results in a 0.53 percent final decrease in P_{BDM} , i.e., the feedback through the system stabilizes the price. However, under crisis conditions, the same sale could trigger a 10.53 percent decrease in price. Thus, iterating over the loop several times leads to a fire sale situation.

Since the SDG approach allows one to model how the system might behave to price shocks under normal and abnormal conditions, this approach can serve as a framework for methodical stress testing and monitoring the critical nodes and edges. The next level of sophistication would be to develop differential (or difference) equations based dynamic models, which provide a more detailed analysis of the dynamic behavior of the financial system.

Exhibit 12: Results for all loops

ID	Sign	Loop	Deviation	Situation	Final value	Threshold	Remarks
1	+	[$P_{BDM}, C_{MM}, F_{MM}, V_{FD}, V_{PB}, L_{HF}, Q_{HF}, Q_{TD}, \lambda_{TD}, \epsilon_{TD}, P_{BDM}$]	Low	Normal	-0.10%	-10%	safe
			Low	Abnormal	-2.02%	-10%	safe
			High	Normal	-0.53%	-10%	safe
			High	Abnormal	-10.53%	-10%	not safe
2	+	[$P_{BDM}, C_{MM}, F_{MM}, V_{FD}, V_{PB}, L_{HF}, Q_{HF}, P_{BDM}$]	Low	Normal	-0.10%	-10%	safe
			Low	Abnormal	-2.00%	-10%	safe
			High	Normal	-0.50%	-10%	safe
			High	Abnormal	-10.00%	-10%	not safe
3	+	[$P_{BDM}, C_{FD}, V_{FD}, V_{PB}, L_{HF}, Q_{HF}, Q_{TD}, \lambda_{TD}, \epsilon_{TD}, P_{BDM}$]	Low	Normal	-0.10%	-10%	safe
			Low	Abnormal	-2.02%	-10%	safe
			High	Normal	-0.53%	-10%	safe
			High	Abnormal	-10.53%	-10%	not safe
4	+	[$P_{BDM}, C_{FD}, V_{FD}, V_{PB}, L_{HF}, Q_{HF}, P_{BDM}$]	Low	Normal	-0.10%	-10%	safe
			Low	Abnormal	-2.00%	-10%	safe
			High	Normal	-0.50%	-10%	safe
			High	Abnormal	-10.00%	-10%	not safe
5	+	[$P_{BDM}, C_{PB}, V_{PB}, L_{HF}, Q_{HF}, Q_{TD}, \lambda_{TD}, \epsilon_{TD}, P_{BDM}$]	Low	Normal	-0.10%	-10%	safe
			Low	Abnormal	-2.02%	-10%	safe
			High	Normal	-0.53%	-10%	safe
			High	Abnormal	-10.53%	-10%	not safe
6	+	[$P_{BDM}, C_{PB}, V_{PB}, L_{HF}, Q_{HF}, P_{BDM}$]	Low	Normal	-0.10%	-10%	safe
			Low	Abnormal	-2.00%	-10%	safe
			High	Normal	-0.50%	-10%	safe
			High	Abnormal	-10.00%	-10%	not safe
7	+	[$P_{BDM}, \lambda_{HF}, L_{HF}, Q_{HF}, Q_{TD}, \lambda_{TD}, \epsilon_{TD}, P_{BDM}$]	Low	Normal	-0.53%	-10%	safe
			Low	Abnormal	-10.53%	-10%	not safe
			High	Normal	-3.33%	-10%	safe
			High	Abnormal	-66.67%	-10%	not safe
8	+	[$P_{BDM}, \lambda_{HF}, L_{HF}, Q_{HF}, P_{BDM}$]	Low	Normal	-0.50%	-10%	safe
			Low	Abnormal	-10.00%	-10%	not safe
			High	Normal	-2.50%	-10%	safe
			High	Abnormal	-50.00%	-10%	not safe
9	+	[$P_{BDM}, C_{MM}, F_{MM}, V_{FD}, \lambda_{TD}^{SP}, \epsilon_{TD}, P_{BDM}$]	Low	Normal	-0.10%	-10%	safe
			Low	Abnormal	-2.00%	-10%	safe
			High	Normal	-0.50%	-10%	safe
			High	Abnormal	-10.00%	-10%	not safe
10	+	[$P_{BDM}, C_{FD}, V_{FD}, \lambda_{TD}^{SP}, \epsilon_{TD}, P_{BDM}$]	Low	Normal	-0.10%	-10%	safe
			Low	Abnormal	-2.00%	-10%	safe
			High	Normal	-0.50%	-10%	safe
			High	Abnormal	-10.00%	-10%	not safe
11	+	[$\chi_{PB}, V_{PB}, L_{HF}, Q_{HF}, \chi_{PB}$]	Low	Normal	-1.00%	-10%	safe
			Low	Abnormal	-1.00%	-10%	safe
			High	Normal	-5.00%	-10%	safe
			High	Abnormal	-5.00%	-10%	safe
12	+	[$P_{BDM}, \lambda_{TD}, \epsilon_{TD}, P_{BDM}$]	Low	Normal	-1.65%	-10%	safe
			Low	Abnormal	-32.94%	-10%	not safe
			High	Normal	-28.00%	-10%	not safe
			High	Abnormal	-560.00%	-10%	not safe
13	-	[$\lambda_{HF}, L_{HF}, Q_{HF}, \lambda_{HF}$]	Low	Normal	-1.23%	-10%	safe
			Low	Abnormal	-1.23%	-10%	safe
			High	Normal	-5.88%	-10%	safe
			High	Abnormal	-5.88%	-10%	safe
14	-	[$\epsilon_{TD}, Q_{TD}, \lambda_{TD}, \epsilon_{TD}$]	Low	Normal	-0.10%	-10%	safe
			Low	Abnormal	-1.96%	-10%	safe
			High	Normal	-0.50%	-10%	safe
			High	Abnormal	-9.09%	-10%	safe

Conclusion

The financial system is self-organized; it did not develop as a carefully engineered system with proper consideration given to the stability and the management of its complex interactions. Because of this, it is all the more critical to understand the paths of positive and negative feedback, alternative routes for funding and securities flows in the event of a shock to one node or edge of the network, and more generally, how the dynamic interactions in the system can create vulnerabilities and instabilities.

We suggest that a process systems engineering framework is a useful modeling paradigm for this challenge. In particular, causal models represented as SDGs and the associated process hazards analysis framework can add the critical capabilities missing in the current network-based approaches emerging as the leading modeling framework for the financial system. The SDG framework adds crucial information to the context of linkages in a network in terms of the direction of various flows and whether they contribute positive or negative feedback, thereby providing a systematic framework for analyzing the potential hazards and instabilities in the system. We show this framework can reveal instabilities and mechanisms of failure that may not be apparent in a network-based perspective for large financial systems. This framework can highlight and help us monitor dynamics such as fire sales and funding runs, where actions that are locally stabilizing, such as a financial institution taking risk management actions without an understanding of the systemic implications, might cascade to have globally destabilizing consequences.

Acknowledgments

This work is supported in part by Columbia University and the Center for the Management of Systemic Risk.

References

Adrian, Tobias, and Hyun Song Shin. "Procyclical Leverage and Value-at-Risk." *Review of Financial Studies* 27, no. 2 (2014): 373-403.

Aguiar, Andrea, Richard Bookstaber, and Thomas Wipf. "A Map of Funding Durability and Risk." Office of Financial Research Working Paper No. 14-03 (2014).

Allen, Franklin, and Ana Babus. "Networks in Finance." In *The Network Challenge: Strategy, Profit, and Risk in an Interlinked World*, Paul R. Kleindorfer, Yoram Wind, and Robert E. Gunther, eds. (2009).

Battiston, Stefano, Guido Caldarelli, Co-Pierre Georg, Robert May, and Joseph Stiglitz. "Complex derivatives." *Nature Physics* 9, no. 3 (2013): 123-125.

Brunnermeier, Markus K., and Lasse Heje Pedersen. "Market Liquidity and Funding Liquidity." *Review of Financial Studies* 22, no. 6 (2009): 2201-2238.

Diamond, Douglas W., and Philip H. Dybvig. "Bank Runs, Deposit Insurance, and Liquidity." *Journal of Political Economy* 91, no. 3 (1983): 401-419.

Fostel, Ana, and John Geanakoplos. "Leverage cycles and the anxious economy." *American Economic Review* 98, no. 4 (2008): 1211-1244.

Gai, Prasanna, and Sujit Kapadia. "Contagion in financial networks." *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science* 466 (2010): 2401-2423.

Gorton, Gary B. "Slapped in the Face by the Invisible Hand: Banking and the Panic of 2007." Available at SSRN, no. 1401882 (2009).

Russell, Stuart, and Peter Norvig. *Artificial Intelligence: A Modern Approach*. Prentice-Hall, Englewood Cliffs (1995).

Shleifer, Andrei, and Robert W. Vishny. "Fire Sales in Finance and Macroeconomics." *Journal of Economic Perspectives* 25, no. 1 (2011): 29-48.

Venkatasubramanian, Venkat, Jinsong Zhao, and Shankar Viswanathan. "Intelligent systems for HAZOP analysis of complex process plants." *Computers and Chemical Engineering* 24, no. 9 (2000): 2291-2302.

Venkatasubramanian, Venkat, Raghunathan Rengaswamy, Kewen Yin, and Surya N. Kavuri. "A review of process fault detection and diagnosis: Part I: Quantitative model-based methods." *Computers and chemical engineering* 27, no. 3 (2003a): 293-311.

Venkatasubramanian, Venkat, Raghunathan Rengaswamy, and Surya N. Kavuri. "A review of process fault detection and diagnosis: Part II: Qualitative models and search strategies." *Computers and Chemical Engineering* 27, no. 3 (2003b): 313-326.

Venkatasubramanian, Venkat, Raghunathan Rengaswamy, Surya N. Kavuri, and Kewen Yin. "A review of process fault detection and diagnosis: Part III: Process history based methods." *Computers and chemical engineering* 27, no. 3 (2003c): 327-346.

Venkatasubramanian, Venkat. "DROWNING IN DATA: Informatics and modeling challenges in a data-rich networked world." *AIChE journal* 55, no. 1 (2009): 2-8.

Venkatasubramanian, Venkat. "Systemic failures: Challenges and opportunities in risk management in complex systems." *AIChE Journal* 57, no. 1 (2011): 2-9.

Zhao, C., M. Bhushan, and V. Venkatasubramanian. "PHASuite: An automated HAZOP analysis tool for chemical processes: Part I: Knowledge engineering framework." *Process Safety and Environmental Protection* 83, no. 6 (2005a): 509-532.

Zhao, C., M. Bhushan, and V. Venkatasubramanian. "PHASuite: An automated HAZOP analysis tool for chemical processes: Part II: Implementation and Case Study." *Process Safety and Environmental Protection* 83, no. 6 (2005b): 533-548.