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

Richard Bookstaber, Paul Glasserman, Garud Iyengar, Yu Luo, Venkat Venkatasubramanian, and Zhizun Zhang

Richard Bookstaber
is a research principal at the Office of Financial Research, U.S. Department of the Treasury, in Washington, DC.
richard.bookstaber@treasury.gov

Paul Glasserman
is Jack R. Anderson Professor of Business at Columbia Business School, Columbia University, in New York, NY.
pg20@columbia.edu

Garud Iyengar
is Chair of the Department of Industrial Engineering and Operations Research at Columbia University in New York, NY.
garud@ieor.columbia.edu

Yu Luo
is a PhD candidate in the Department of Chemical Engineering at Columbia University in New York, NY.
yl2750@columbia.edu

Venkat Venkatasubramanian
is Samuel Ruben-Peter G. Viele Professor of Engineering at Columbia University in New York, NY.
venkat@columbia.edu

Zhizun Zhang
is a PhD candidate in the Department of Chemical Engineering at Columbia University in New York, NY.
zz2216@columbia.edu

Modern financial systems are characterized by a very 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 recent financial crisis that was precipitated by counterparty exposures revealed by the Lehman bankruptcy, the near bankruptcy of AIG, and the European debt crisis that was caused by the exposure of European banks to sovereign default risk emphasizes the critical need for a fundamental understanding of the structure and dynamics of this system. 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.[1]

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

Furthermore, financial systems have the particular feature that the steps taken by a single agent to mitigate its risk, under extreme circumstances, can become the very source of destabilizing positive feedback through the interaction of multiple agents. We refer to these steps as *locally* stabilizing yet *globally* destabilizing. This phenomenon is illustrated by the phenomenon of the bank run. 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. Yet 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 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 (Shleifer and Vishny [2010]), liquidity spirals (Brunnermeier and Pedersen [2009]), leverage cycles (Adrian and Shin [2013], and Fostel and Geanakoplos [2008]), and panics (Gorton [2009]). 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. In order 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, whereas 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, the environmental state, and the causal relationships 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. Moreover, subjecting SDG to a process hazard analysis (PHA) (Venkatasubramanian et al. [2000] and Venkatasubramanian [2011]) pinpoints 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], and 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, the safety and stability of an assembled system is a design criterion. 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 article shows how the SDG framework makes this possible through a systemwide 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, a chemical process plant, that takes securities and funding as inputs and creates new financial products as outputs that are 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. Though researchers have suggested the Internet, electrical power grid, and transportation network as potential models for the financial system, none of these has the richness of phenomena seen in a large-scale chemical process plant. We demonstrate in this article that phenomena such as various physical or chemical transformations, feedback and recycle loops, and so on 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 production process.

In order to gain further insight into the underlying dynamics, one needs a richer, more detailed, modeling framework (Venkatasubramanian et al. [2000] and Venkatasubramanian [2009]). This is carried out in process systems engineering at three levels of increasing sophistication and effort: 1) qualitative causal models, such as SDGs, capture the underlying cause-and-effect relationships, 2) quantitative steady-state models, represented as a system of algebraic equations, capture the steady-state behavior of the process, and 3) quantitative dynamic models, generally represented as a system of ordinary or partial differential equations (ODEs or PDEs), predict the transient behavior of the process. The particular choice for the model depends on the need. For instance, for performing PHA, where one systematically identifies the potential hazards, their causes, and adverse consequences, it is often adequate to use the qualitative causal SDG models. On the other hand, for making process control decisions, one requires a detailed dynamic model that is derived from first principles (as ODEs or PDEs) or from a data-driven perspective as an input-output model. Generally speaking, in many industrial settings, given the complexity of the underlying process, it is often quite difficult or expensive to develop the quantitative dynamic models, particularly from first principles.

We now illustrate the SDG framework with the aid of a simple process engineering example, a continuous stirred-tank reactor (CSTR) process (see Exhibit 1 and Stephanopoulos [1984]) where an exothermic (that is, heat generating) reaction, $A \rightarrow B$, takes place. The heat generated by the reaction is removed by passing a coolant through the jacket of the reactor (shaded), thereby controlling the temperature T inside the reactor. If the temperature is not controlled, it could lead to a runaway reaction and explosion. The temperature is controlled by a feedback control loop that manipulates the coolant flow rate F_c to achieve the desired set point temperature.

We next build an SDG model for the CSTR process. A digraph is a graph with directed arcs between the

EXHIBIT 1

CSTR Example (adapted from Stephanopoulos [1984], Fig. 23.5c)

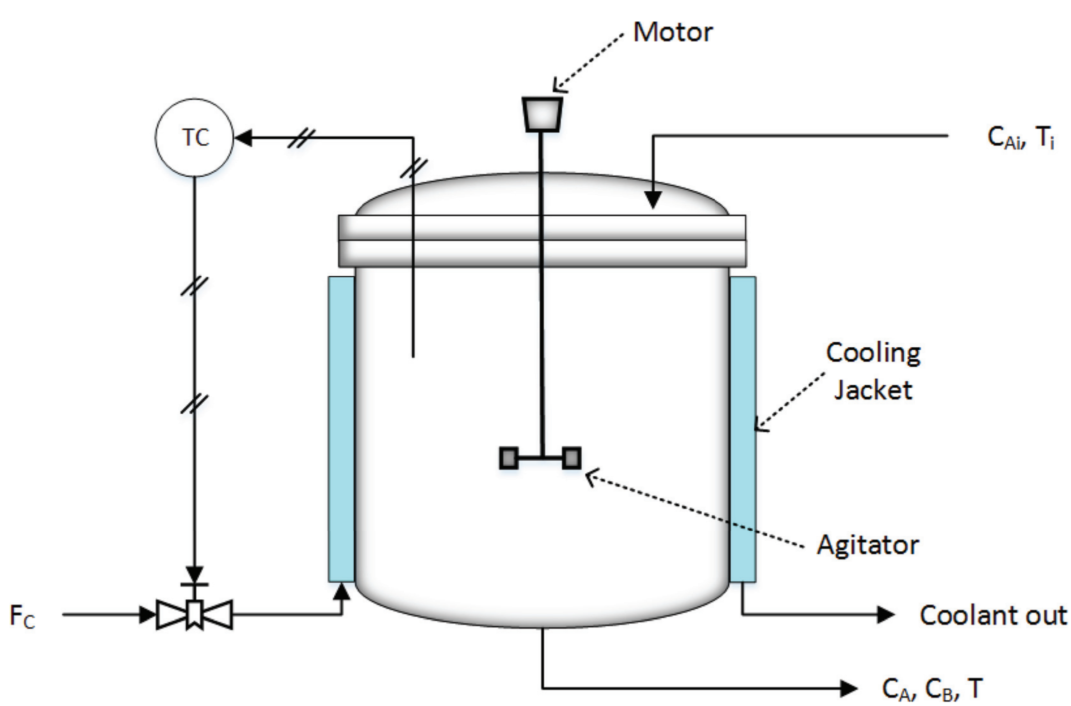

nodes, and a signed directed graph (SDG) is a graph in which the directed arcs have a positive (shown as solid lines) or negative sign (shown as dotted lines) attached to them. The nodes represent events or variables and edges relationship between the nodes. The directed arcs lead from the cause nodes to effect nodes, showing the direction of causality. In the typical use of SDG models, each node corresponds to a deviation from the steady-state value of a variable. SDG models are much more compact than truth tables, decision tables, or finite state models, and are, therefore, quite efficient in capturing the causes and effects represented in a process or equipment. The qualitative SDG models are easier to develop and analyze, in comparison to the dynamic models, and can yield quick and useful results in certain decision-making tasks such as process fault diagnosis and process hazards analysis (Venkatasubramanian et al. [2000], Venkatasubramanian and Vaidhyanathan [1994], Vaidhyanathan and Venkatasubramanian [1996], Srinivasan and Venkatasubramanian [1998a, 1998b], Maurya et al. [2003a, 2003b], Maurya et al. [2004], and Zhao et al. [2005a, 2005b]). Even when a dynamic model is available, it is generally faster and more efficient to use an SDG model to perform cause-and-effect reasoning for such applications. However, since SDG models are qualitative in nature, they can lead to ambiguities and hence are limited to certain kinds of tasks (Venkatasubramanian et al. [2003a, 2003b, 2003c]).

The SDG model for the CSTR example is shown in Exhibit 2. The figure is read as follows: a change in the inlet concentration of A, C_{Ai} positively affects the concentration of A inside the reactor, C_A; that is, if C_{Ai} increases, C_A will increase, and if C_{Ai} decreases, C_A will decrease. This is shown by the solid edge between these two nodes. And if C_A increases, then the reaction rate *r* will increase, which is shown by the solid edge between these two nodes. However, an increase in the reaction rate will increase the conversion of A → B, thereby reducing the concentration of A (a negative feedback here). This is captured by the negative edge in dotted line between *r* and C_A. An increase in the reaction rate *r* results an increase in *T*, which in turn causes an increase in *r*, potentially leading to a runaway reaction if the coolant flow fails to control this. The rest of the SDG is to be interpreted by following the direction of causality, as shown earlier. Maurya et al. [2003a, 2003b, 2004] discuss how the SDG model can be derived systematically from the underlying equations of the process or from a detailed causal understanding of the process.

Exhibit 2
SDG for the CSTR Example (exothermic reaction A → B)

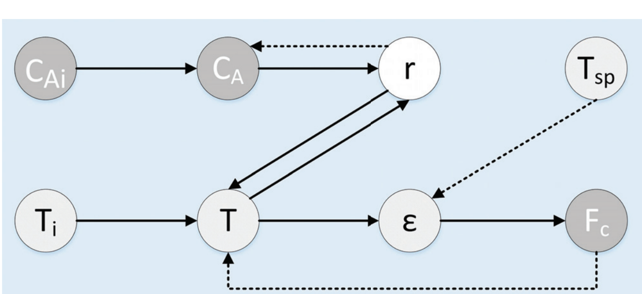

Although the SDG model of the entire process unit network (that is, flowsheet) for an industrial process is naturally more complicated, with hundreds of nodes and edges, it can be assembled from a library of unit-wise SDG models, as discussed by Maurya et al. [2003a, 2003b, 2004]. Venkatasubramanian and coworkers have also developed artificial intelligence-based systems that automate much of the cause-and-effect reasoning (both diagnostic and prognostic) using SDG models for entire flowsheets with recycle and control loops (Venkatasubramanian et al. [2000], Venkatasubramanian and Vaidhyanathan [1994], Vaidhyanathan and Venkatasubramanian [1996], Srinivasan and Venkatasubramanian [1998a, 1998b], and Maurya et al. [2003a, 2003b, 2004]) for process fault diagnosis and process hazards analysis applications. These methods can be adapted for developing a process systems engineering framework for modeling and analyzing risk in financial networks. We can 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

We now explain how SDG models can be used to analyze the dynamics of financial systems. 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 with other financial entities and among its business units, is complex. For the sake of simplicity, 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 the 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 finance 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 finance desk. The bank/dealer interacts with cash providers, such as money market funds, pension funds, and insurance companies; other banks/dealers through the over-the-counter market, which is the market for the bank/dealer to acquire or lay off inventory; and the hedge funds, which, as noted earlier, seek leverage and securities from prime brokers to support their long/short trading posi-

E X H I B I T 3
Simplified Bank/Dealer Network

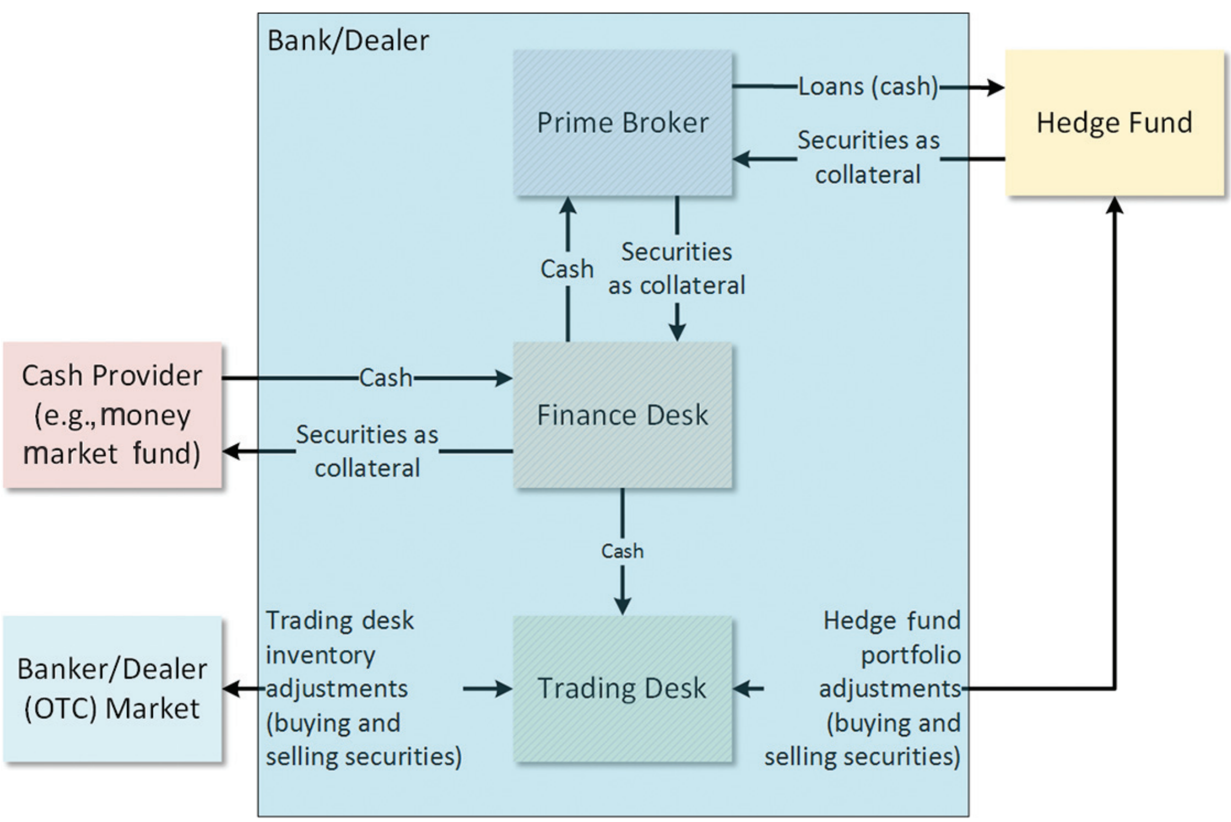

tions. The hedge funds also represent the wider swath of institutional customers that use the bank/dealer's market-making function, ranging from asset managers and hedge funds to pension funds, sovereign wealth funds, and insurance companies.

The interactions between the bank/dealer's functional areas create various financial transformations. The finance desk takes short-term loans from the cash providers and passes them through to clients that have 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 them 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, that is, nodes that receive the output from a bank/dealer, are also sources of inputs. A hedge fund that is borrowing in order to buy securities might also be lending other securities. A pension fund that is 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 in that they are both buyers and sellers of securities, and in that sense provide inputs as well as output in market making.

BANK/DEALER CASE STUDY

The network depicted in Exhibit 3, though illustrative of the layout of the components of the bank/dealer and its interactions, does not represent the effect of the various flows, and therefore 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 that is displayed in Exhibit 4.

For simplicity, we consider a system with a single market asset (for example, 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, respectively. These units need funding to finance their asset holdings; this 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. Thus, 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. More 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_{HF}. The prime broker's lending is determined by the bank/dealer's finance desk and by the prime broker's margin rate, χ_{PB}.

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 finance 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 that are 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. Thus 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 a financial crisis: funding runs and fire sales.

Exhibit 4

SDG Model for Bank/Dealer Example

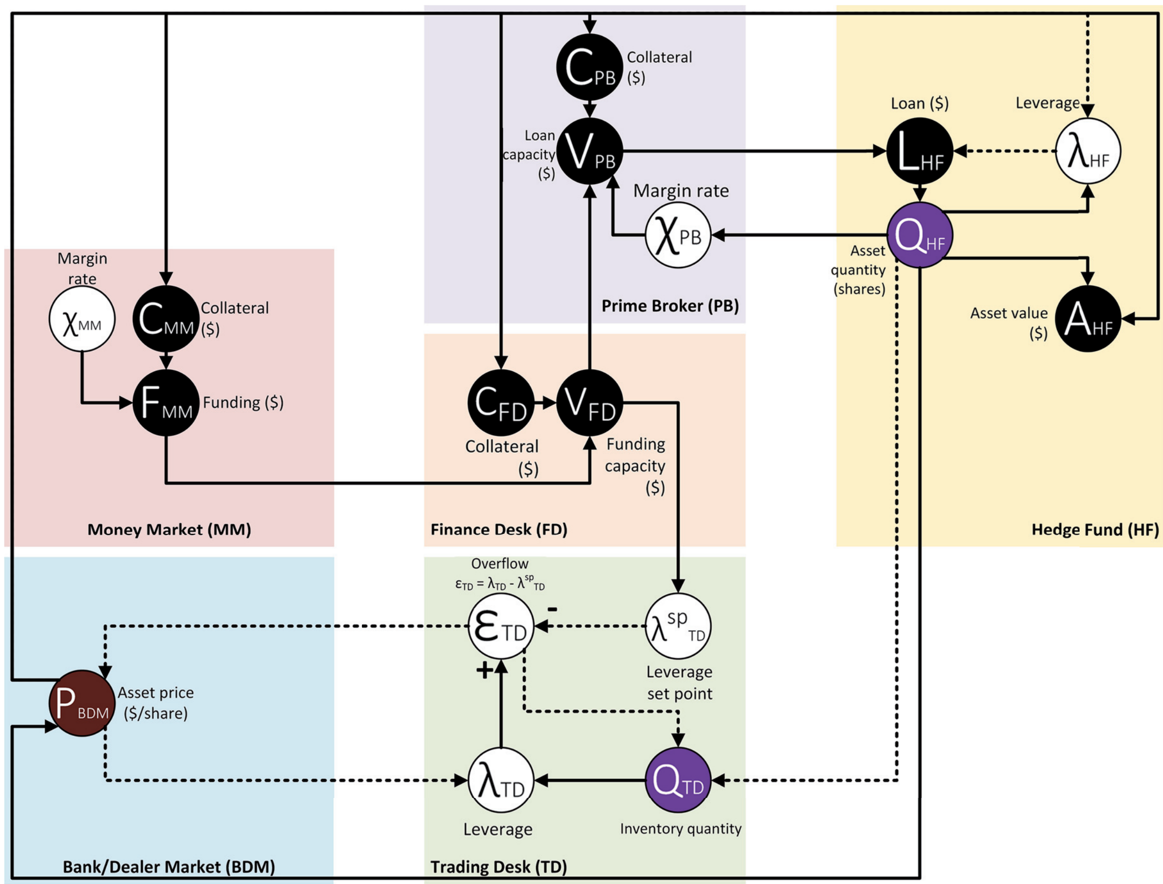

Process hazards analysis (Venkatasubramanian et al. [2000], Venkatasubramanian [2011], and Zhao et al. [2005a, 2005b]) is a methodology for systematically identifying abnormal causes and adverse consequences that can occur anywhere in the process system. In the context of an SDG model, PHA provides the 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. Using this framework, we can identify and examine the complete list of loops in an SDG model. This list can be computed via a depth-first search of the SDG (Russell and Norvig [1995]). 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 each feedback loop. Nevertheless, 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, respectively. Each of these units uses a leverage target as an internal control for the quantity held of the market asset. However, 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

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 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.

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 shown in Exhibit 6, 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 making the situation worse in every subsequent iteration.

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 in order to reduce its leverage, and this drops prices. The second loop has the same effect, a 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 λ_{TD}. This in turn leads the trading unit to sell into the market, with the end result again being a further drop in prices.

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 increases 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 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

Exhibit 6

SDG Model for Bank/Dealer Fire Sale Example

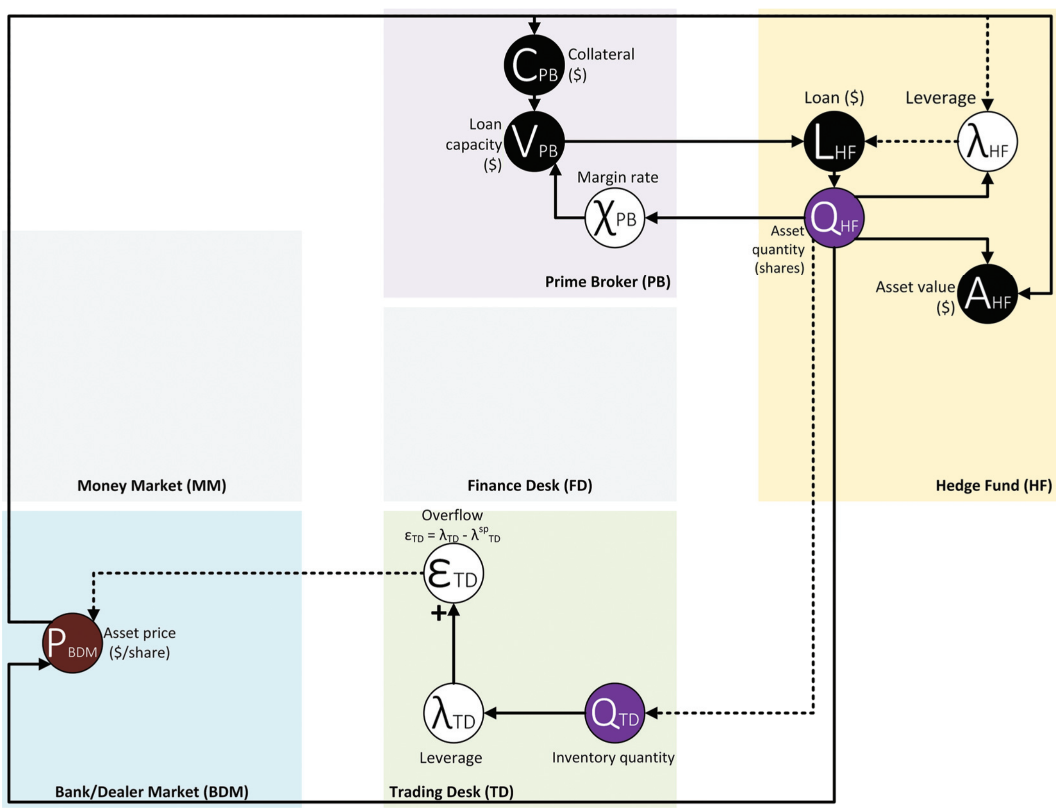

Exhibit 7

Fire Sale Loops

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

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 in 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 thereby 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 a drop in price-induced drop in funding leads to asset sales and subsequent price drops. For example, the drop in collateral value can affect the finance desk directly.

In both fire sales and funding runs, the SDG model identifies a critical dynamic 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 the hedge funds. But this leads to actions by the hedge funds that contribute to a positive feedback cycle that reduces funding for the prime broker even further. Similarly, a

EXHIBIT 8
SDG Model for Bank/Dealer Funding Run Example

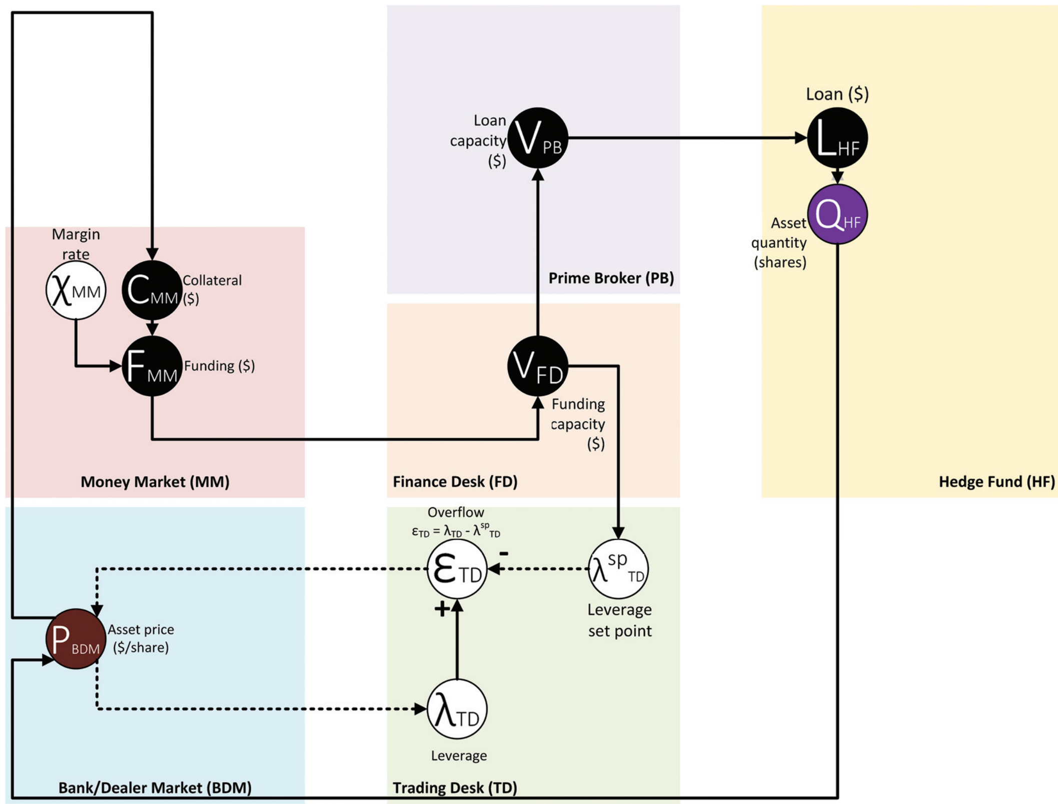

EXHIBIT 9
Funding Run Loops

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

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 thereby 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 providers, seeing erosion in their collateral, begin to reduce funding. The SDG model also shows that there is pathway in the opposite direction: a drop in funding to the trading desk leads to a reduction in inventory, causing a drop in prices that reduces the value of the hedge fund portfolio, leading the prime broker to increase its margin level, thereby 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 bank/dealer network in Exhibit 3, one can perhaps manually identify and analyze all the feedback loops listed in Exhibit 5. However, for a more realistic version of this network, as shown Exhibit 10, where there are multiple hedge funds, multiple banks/dealers, multiple clients, various derivatives and structured products, it is virtually impossible to identify and analyze all such loops manually. This, again, highlights the need for the SDG framework, which can be automated to handle larger systems.

Exhibit 10

More Realistic Bank/Dealer Configuration

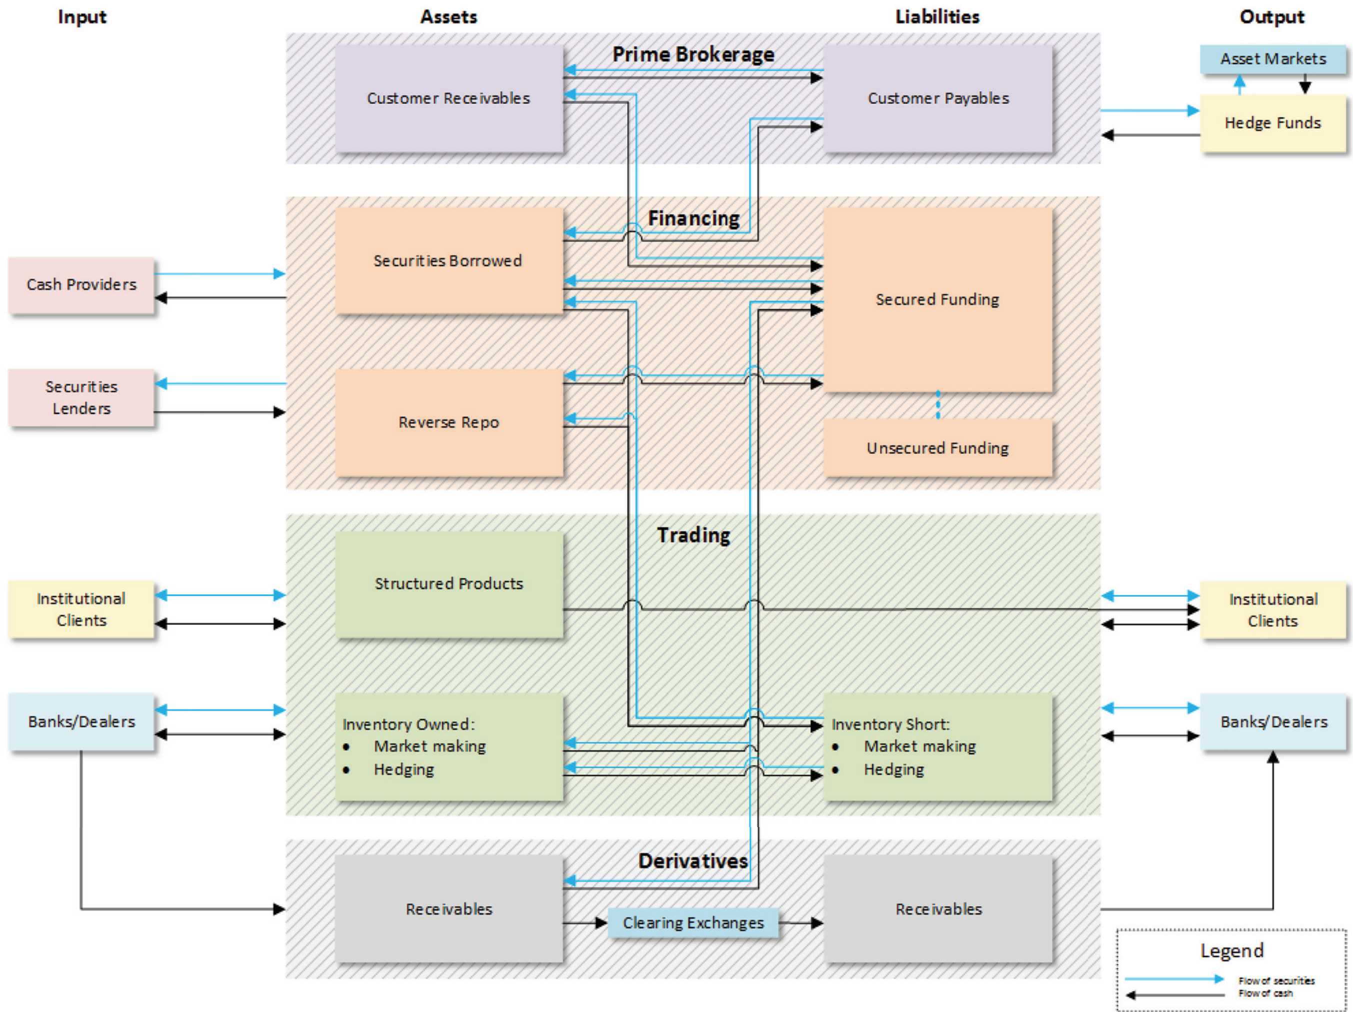

A further advantage is that the framework allows us to formulate more sophisticated models, as and when we need them, in a methodical manner. For instance, we now show how we can add numerical gains (Vaidhyanathan and Venkatasubramanian [1996]) on all the edges connecting various nodes and perform a 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 historical market data and the judgment of experienced market professionals.

SEMIQUANTITATIVE ANALYSIS

Consider a loop of the form $(v_1, v_2, \ldots, v_n, v_{n+1} = v_1)$, 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, …, 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

$$\delta v_i = \frac{\Delta v_i}{v_i} = \frac{f_{i-1}\left(v_{i-1}(1+\delta v_{i-1})\right) - f_{i-1}(v_{i-1})}{f_{i-1}(v_{i-1})}$$
$$= \frac{f_{i-1}\left(v_{i-1}(1+\delta v_{i-1})\right)}{f_{i-1}(v_{i-1})} - 1 \equiv F_{i-1}\left(\delta v_{i-1}; v_{i-1}\right) \quad (1)$$

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, that is, $f_{i-1}(v_{i-1}) = k_{i-1} v_{i-1}$, 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}, that is, the relative disturbance in the value of node v_1 after one iteration through the loop, by $\delta v_{1,f}$. From Equation (1), it follows that

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

For linear relationships, (that is, 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,

$$\delta v_{1,f} = k_n\, k_{n-1} \cdots k_1\, \delta v_1$$

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:

$P_{BDM} = \$10$
$C_{HF} = \$1$ billion
$C_{TD} = \$1$ billion

Exhibit 11
Loop 7 as an Example

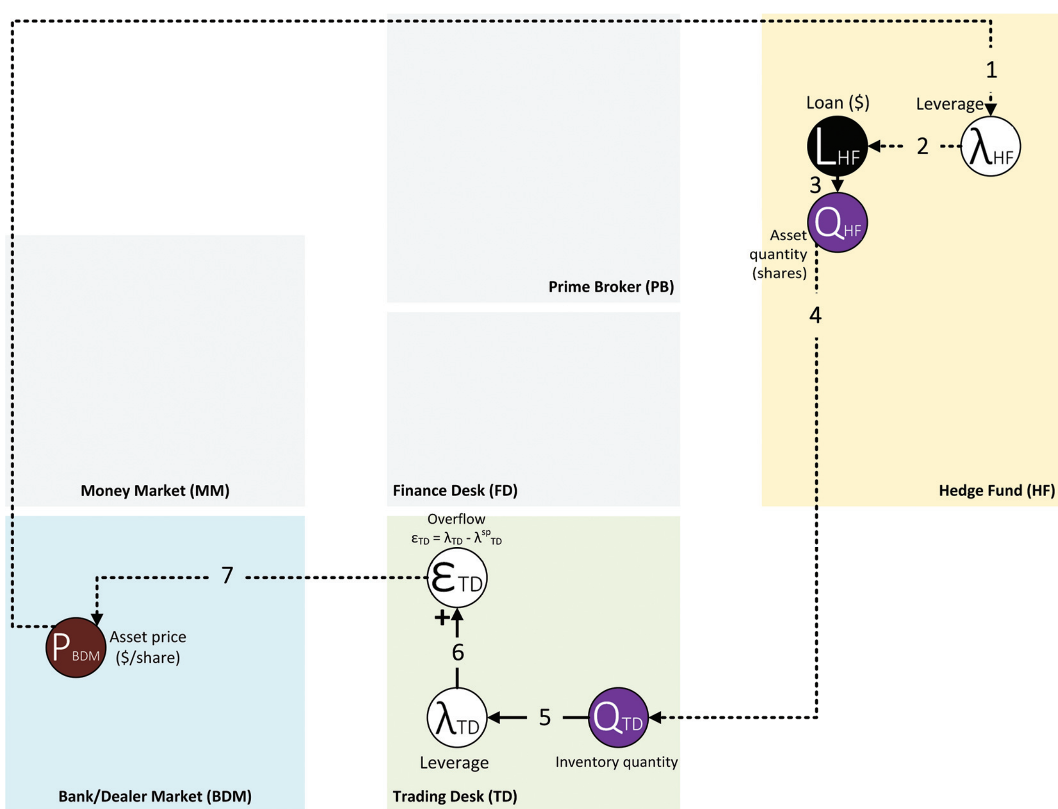

A_{PB} = \$5 billion
A_{HF} = \$5 billion
A_{TD} = \$15 billion
$A_{FD} = A_{PB} + A_{TD}$ = \$20 billion
$L_{HF} = A_{HF} - C_{HF}$ = \$4 billion
$L_{TD} = A_{TD} - C_{TD}$ = \$14 billion
Q_{HF} = 500 million shares
Q_{TD} = 1.5 billion shares
χ_{MM} = 25%
χ_{PB} = 25%

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

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

From Equation (1), it follows that

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

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 fund 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

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

that is,

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

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

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

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 (resp. sells) shares the capital C_{TD} of the trading desk decreases (resp. increases); moreover, the relationship is linear. Therefore, $\delta C_{TD} = -\delta Q_{TD}$. The relative change in leverage δL_{TD} is given by

$$\delta\lambda_{TD} = \frac{\dfrac{A_{TD}}{C_{TD}(1+\delta C_{TD})} - \dfrac{A_{TD}}{C_{TD}}}{A_{TD}/C_{TD}} = \frac{-\delta C_{TD}}{1+\delta C_{TD}}$$

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 $\epsilon_{TD} \le 0$, that is, 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 the SDG approach, we assume

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

Now we are in a position to compute the loop gain $\delta P_{BDM,f}/\delta P_{BDM}$ using Equation (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 14 loops for both normal and crisis conditions, and for a small (1%) and large (5%) initial decrease. Specifically, for Loop 7 under normal market conditions, a 1% initial decrease in P_{BDM} results in a 0.53% final decrease in P_{BDM}, that is, the feedback through the system stabilizes the price.

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

However, under crisis conditions, the same sale could trigger a 10.53% 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-) equation-based dynamic models, which provide a more detailed analysis of the dynamic behavior of the financial system.

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 the appropriate 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 that are 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 that this framework can reveal hidden instabilities and mechanisms of failure that may not be apparent in a network-based perspective for large financial systems. It can highlight and help us monitor dynamics such as fire sales and funding runs, where actions that are locally stabilizing—for example, where a financial institution takes risk-management actions without an understanding of the systemic implications—might cascade to globally destabilizing consequences.

ENDNOTES

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

[1]Authors are listed alphabetically. Corresponding author: Venkat Venkatasubramanian, Venkat@columbia.edu.

REFERENCES

Adrian, T., and H.S. Shin. "Procyclical Leverage and Value-at-Risk." *Review of Financial Studies*, 2013, hht068.

Allen, F., and A. Babus. "Networks in Finance." In *The Network Challenge: Strategy, Profit, and Risk in an Interlinked World, edited by* P.R. Kleindorfer, Y. Wind, and R.E. Gunther. Upper Saddle River, NJ: FT Press, 2009.

Battiston, S., G. Caldarelli, C. Georg, R. May, and J. Stiglitz. "Complex Derivatives." *Nature Physics*, Vol. 9, No. 3 (2013), pp. 123-125.

Brunnermeier, M.K., and L.H. Pedersen. "Market Liquidity and Funding Liquidity." *Review of Financial Studies*, Vol. 22, No. 6 (2009), pp. 2201-2238.

Diamond, D.W., and P.H. Dybvig. "Bank Runs, Deposit Insurance, and Liquidity." *Journal of Political Economy*, 1983, pp. 401-419.

Fostel, A., and J. Geanakoplos. "Leverage Cycles and the Anxious Economy." *American Economic Review*, Vol. 98, No. 4 (2008), pp. 1211-1244.

Gai, P., and S. Kapadia. "Contagion in Financial Networks." *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, (2010): rspa20090410.

Gorton, G.B. "Slapped in the Face by the Invisible Hand: Banking and the Panic of 2007." 2009, *SSRN* 1401882.

Maurya, M.R., R. Rengaswamy, and V. Venkatasubramanian. "A Systematic Framework for the Development and Analysis of Signed Digraphs for Chemical Processes. 1. Algorithms and Analysis." *Industrial & Engineering Chemistry Research*, Vol. 42, No. 20 (2003a), pp. 4789-4810.

——. "A Systematic Framework for the Development and Analysis of Signed Digraphs for Chemical Processes. 2. Control Loops and Flowsheet Analysis." *Industrial & Engineering Chemistry Research,* Vol. 42, No. 20 (2003b), pp. 4811-4827.

——. "Application of Signed Digraphs-Based Analysis for Fault Diagnosis of Chemical Process Flowsheets." *Engineering Applications of Artificial Intelligence*, Vol. 17, No. 5 (2004), pp. 501-518.

Russell, S., and P. Norvig. *Artificial Intelligence: A Modern Approach*. Englewood Cliffs, NJ: Prentice-Hall, 1995.

Shleifer, A., and R.W. Vishny. *Fire Sales in Finance and Macroeconomics*. No. w16642. National Bureau of Economic Research, 2010.

Srinivasan, R., and V. Venkatasubramanian. "Automating HAZOP Analysis of Batch Chemical Plants: Part I. The Knowledge Representation Framework." *Computers & Chemical Engineering*, Vol. 22, No. 9 (1998a), pp. 1345-1355.

——. "Automating HAZOP Analysis of Batch Chemical Plants: Part II. Algorithms and Application." *Computers & Chemical Engineering,* Vol. 22, No. 9 (1998b), pp. 1357-1370.

Stephanopoulos, G. *Chemical Process Control*. Upper Saddle River, NJ: Prentice Hall PTR, 1984.

Vaidhyanathan, R., and V. Venkatasubramanian. "A Semi-Quantitative Reasoning Methodology for Filtering and Ranking HAZOP Results in HAZOP Expert." *Reliability Engineering & System Safety*, Vol. 53, No. 2 (1996), pp. 185-203.

Venkatasubramanian, V. "Drowning in Data: Informatics and Modeling Challenges in a Data-Rich Networked World." *AIChE Journal*, Vol. 55, No. 1 (2009), pp. 2-8.

——. "Systemic Failures: Challenges and Opportunities in Risk Management in Complex Systems." *AIChE Journal*, Vol. 57, No. 1 (2011), pp. 2-9.

Venkatasubramanian, V., and R. Vaidhyanathan. "A Knowledge-Based Framework for Automating HAZOP Analysis." *AIChE Journal*, Vol. 40, No. 3 (1994), pp. 496-505.

Venkatasubramanian, V., J. Zhao, and S. Viswanathan. "Intelligent Systems for HAZOP Analysis of Complex Process Plants." *Computers & Chemical Engineering*, Vol. 24, No. 9 (2000), pp. 2291-2302.

Venkatasubramanian, V., R. Rengaswamy, K. Yin, and S.N. Kavuri. "A Review of Process Fault Detection and Diagnosis: Part I: Quantitative Model-Based Methods." *Computers & Chemical Engineering*, Vol. 27, No. 3 (2003a), pp. 293-311.

——. "A Review of Process Fault Detection and Diagnosis: Part II: Qualitative Models and Search Strategies." *Computers & Chemical Engineering*, Vol. 27, No. 3 (2003b), pp. 313-326.

——. "A Review of Process Fault Detection and Diagnosis: Part III: Process History-Based Methods." *Computers & Chemical Engineering,* Vol. 27, No. 3 (2003c), pp. 327-346.

Zhao, C., M. Bhushan, and V. Venkatasubramanian. "PHA Suite: An Automated HAZOP Analysis Tool for Chemical Processes: Part I: Knowledge Engineering Framework." *Process Safety and Environmental Protection*, Vol. 83, No. 6 (2005a), pp. 509-532.

——. "PHA Suite: An Automated HAZOP Analysis Tool for Chemical Processes: Part II: Implementation and Case Study." *Process Safety and Environmental Protection*, Vol. 83, No. 6 (2005b), pp. 533-548.