

NOTES FOR THE FINAL PANEL DISCUSSION: OBSERVATIONS ON NEW DIRECTIONS FOR UNDERSTANDING SYSTEMIC RISK IN THE FINANCIAL SECTOR. FEDERAL RESERVE BANK OF NEW YORK AND THE NATIONAL ACADEMY OF SCIENCES

(to be pruned according to time allocated)

Stated Objectives for the Final Panel Discussion:

- 1) Provide summary impressions of the topic and of the conference as a whole**
- 2) Amplify topics covered**
- 3) Give direction for further work**

Panel Participants: Timothy Geithner (FRBNY, Pres), Robert Litzenger (Azimuth), Donald Kohn (Gov, US FED DC), Simon Levin (Princeton), Morten Blech (FRBNY), George Sugihara (Moderator, Rapporteur)

SUGIHARA COMMENTS:

I. Definition of Systemic Risk

As the medical term implies, “systemic” risk is risk that affects the whole financial organism, not just isolated parts. It corresponds to widespread coordinated (correlated) system failure characterized by a catastrophic change in the overall state (valuation) of the system. Such changes may be hysteretic, in the sense that recovery may be much slower to achieve than collapse, and they may be irreversible in that the original state may not be fully recoverable.

A systemic event can be triggered either *exogenously* by a common event such as war (akin to a trauma) or *endogenously* by an imbalance in key parts of the system or a structural instability (akin to disease). The triggering event may be sufficiently global to itself be systemic or it can cause a chain reaction.

In the latter case, the domino effect may simply represent physical disruption of downstream physical flows, such as eliminating a node in the payments system that causes the downstream contingent payments network to fragment. Here the propagation of systemic effects can be determined directly by the network structure of the system and the sensitive bottleneck nodes can be identified (see the discussion of dominator nodes in section VIII below).

Alternatively, the mechanism of contagion may arise endogenously from destabilizing positive feedbacks. These self-reinforcing positive feedbacks can arise in various contexts such as the “rumor mill network” of social contacts, the network of balance sheets, the web of claims and obligations or in the payments system. The topology of the relevant networks will affect propagation and the robustness of the system to such instabilities (see section VII and VIII below).

Box 1 gives a list of specific mechanisms operating in the financial sector that could trigger and propagate a systemic failure through the relevant network.

II. The Potential for Large-Scale Catastrophic Failures in Complex Systems is a Cross Cutting Theme of Large Social Importance.

Social importance

Complex systems abound, and worries about instability, catastrophic change and regime shifts, known common features of many natural systems, are a dominant social concern. Therefore, it is not surprising that these worries are at the leading edge of many disciplines particularly in the environmental and engineering sciences.

- atmospheric scientists interested in global climate change (eg. Keeling)
- ecologists interested in regime shifts (Hsieh, Sugihara et. al. NATURE 2005).
- fisheries managers trying to avoid the sudden collapse of certain economically important fish stocks (Worm and Meyer, Sugihara).
- engineering risk experts interested in catastrophic machine failures and failures in the electrical grid and the internet (Heimes and Amin).

The ubiquity of such problems across so many fields suggests the possibility of finding common principles at work. For example, engineers and medics may both be interested in the effects of short-term solutions that address high frequency low amplitude events (small floods, small disease outbreaks) but that predispose the system to low frequency very high amplitude disturbances (eg. the levees designed for intermediate-strength storm surges but that amplified the effects of Katrina, and the risk of virulent pandemics arising from excessive use of antibiotics). In the financial sector it is important to understand whether implementation of any local or short-term policy (eg. managing risk by VAR) makes the system more vulnerable to the rare but truly catastrophic systemic event that could destroy that system as we know it. When does local or short-term risk management increase risk from a more global system-wide perspective? Can the systemic risk problem benefit from other perspectives? Are there generic principles that can inform the systemic risk problem?

Potential for leverage

The cross-cutting nature of the problem of large-scale system collapse suggests an unusual opportunity to leverage effort from other fields and apply it to investigating systemic risk. This effort can be coordinated by the National Academies and by the Federal Reserve Bank. Much of the discussion at this meeting is aimed at the feasibility and value of exploiting this leverage.

Leveraging systemic risk as a generic problem

As Simon Levin suggests, ecological and economic systems share common elements as complex adaptive systems. To the extent that the analogy holds, these two disciplines have potential for mutual leverage.

Recent studies in nonlinear complex systems show rapid and large transitions in state to be common features of many “generic” interconnected dynamic (and cybernetic) systems. Beyond the specific analogy between ecology and economics, certain dynamical behaviors and structural (topological network) constraints are common to broad classes of systems. Behaviors and network topologies that are truly generic (not necessarily system-specific) can inform many disciplines. For example, it is useful to the systemic risk problem to know the general properties of complex systems (particularly the structural ones) that promote stability or collapse.

Leverage can also occur by sharing methods across disciplines. As Yakov Haimes, Masaoud Amin, and Simon Levin have shown, diverse fields such as engineering risk analysis, epidemiology and ecology employ methods and research styles that are not in common use in the financial sector but that may be useful there. The study presented here of network topology by Soramaki et. al. is a specific demonstration of this point. I reckon most participants at this conference share the view that a larger tool chest of methods is valuable.

Obviously, this does not discount traditional discipline-specific methods. System-specific details and more classical discipline-specific modeling methods and historical event narratives are valuable as we saw in our first sessions yesterday, and in Rob Litzenberger’s detailed comments on LTCM. Engineering risk methods fit well with such analyses and can provide a quantitative implementation framework for evaluating risk associated with the various event chains. As Massaoud Amin describes, these risk methods also allow for the exploration of events that go beyond the historical record (through the use of genetic algorithms and such); a feature that could greatly increase their utility to the financial sector.

As Hendricks et.al. point out, systemic risk is a well-established topic in the finance literature and useful insights have been gained. The importance of historic event-specific analyses is obvious. Nonetheless, they have the same potential shortcoming that Rob Litzenberger describes in the use of historic VAR for setting firm risk limits. Normal

volatility does not apply during crisis times, which is when risk management earns its keep. Each crisis has its own characteristic volatility and contingent event cascade. At that level, systemic events are historic and each is unique, so that a normal scenario-event analysis can be difficult to generalize or apply accurately to future events. Focusing in a more discipline-specific way on the interesting but possibly distracting natural historical epiphenomena that defines financial reality is necessary. However, we should also recognize that it has an implied cost in terms of loss of generality and possible loss of future applicability. For this and other equally direct reasons caution needs to be taken when extrapolating details of history to shape future policy.

Multiple approaches should be taken. In this vein it is important to include some analysis of the problem as a member of a broader generic class that as Simon Levin mentions “is free of distracting epiphenomena” where one can extrapolate beyond any specific historical scenario. More specifically, it seems prudent to know whether any specific policy potentially makes things worse from the generic systems perspective. Again, the aim is not to replace conventional approaches, but to explore complementary approaches. Exploiting commonalities is one way that leverage is achieved.

A cost-benefit view

Two questions tell all: *(I muffed the first part of this one ... but here is what was written in my notes, based on Charles Taylor’s comment at dinner)*

How much money is spent on studying systemic risk, versus funds spent on conventional risk management in individual firms? (Answer: very little ... a ratio close to 0).

How expensive is a systemic risk event to the economy and nation as a whole? (Answer: probably very costly).

These are admittedly vague and difficult questions to answer with any precision. However, they make a compelling case for the sensibility of focusing more attention and resources on the problem. Indeed the scenario in financial markets is very similar to what has happened in fisheries science.

III. A Meeting of Research Cultures

An analogy with fisheries management

The above resource allocation scenario in financial markets is very similar to what happened in fisheries science. For the past half century, classical fisheries management was done on a single species-by-species basis (analogous to single firm risk analysis). The general failure of the classical single-species models to fit the harvest data (eg. the classical Ricker model that relates stock size to future recruitment) and the consequent collapse of many of the world’s fisheries (Worm and Meyer 2002) has made people

aware of the inadequacy of the old paradigm. At least 4 generations of fisheries managers have been trained to see ideal single-species curves emerging from amorphous clouds of points. Their jobs require them to make stock assessments based on these poorly fitting models of reality. For lack of alternatives, this along with other similar assessment models remain in wide use today, although they are also widely criticized by the fisheries community who are nonetheless mandated to use them to produce annual stock estimates of some kind. Early efforts to improve the model fits were largely aimed at better statistical fitting techniques. However, this single-species view is now being replaced by the idea that the models may be fundamentally inadequate, and that the wider ecosystem context and environmental context (by analogy, the full banking system and market system context) are required to make responsible decisions. A new initiative by federal and state agencies toward “ecosystem-based management” though daunting, is beginning to take shape.

Empirically hardened theory

The broader systems view that is now being acknowledged in fisheries science is being echoed in these proceedings. The systems view is explicit in the theoretical papers on current macroeconomic research directions by Roberto Rigobon, Hyun Song Shin, Markus Brunnermeier and Lasse Pedersen. Shin’s model of a “web of claims and obligations” speaks directly to this idea. Though one could pick nits about the simplifying assumptions, of practical consequence is the fact that all three papers, each invoking different simplifying assumptions, converge on the same point. Namely, that it is generally bad to increase margin requirements during a market crisis. This sensible recommendation is convincing (though slightly counter-intuitive for risk managers) because of independent theoretical support, not because it was validated empirically.

The question arises whether the aesthetic benefits that come from invoking simplifying assumptions that produce neat closed-form solutions (such as constructing a model whose dynamics have a fixed point solution and occur on a regular lattice etc.) may be made at the sacrifice of informing reality. There is a general point to be made. Namely, mathematical elegance, which has clear iconographic importance in many fields, may sometimes dominate a discipline both in terms of its culture and in terms of its research vision. I reckon this may be more the case in theoretical settings than is typical in applied science. In either case few would argue against the benefits of bringing macroeconomic models closer to data and testing them in a way similar to the way fisheries and other models in the physical sciences are ideally validated. That is, by seeing how well nature (observation) is captured, and better yet, predicted. (It took fisheries science over 40 years to finally recognize that the emperor had no clothes). The models may remain simple, but by being more tightly constrained by data and observation, they would be strengthened and their utility to policy makers maximized. This suggestion applies equally to research on ecological catastrophes as they would to modeling the systemic risk problem.

Models and empirical constraints

Two different empirical studies of the Fed payments network were presented at this meeting: the Duffy-Ashcroft study of temporal payment flows and the study of the topology of these flows by Soramaki et. al. Both studies demonstrate the value of good empirical work to inform this problem. (Also witness recent papers by Kaminsky and Reinhart, Dungevy et. al., Cohen and Shin, and Danielson and Payne among many others).

Empirical studies serve two important functions. First, they represent real constraints that should either be built into models or should at least be reproducible in models. Secondly, and perhaps more interestingly, they inform (or confirm) intuition and may point to specific classes of generating mechanisms (and lead to better model structures). Describing the unique topology and dynamics of payment flows should lead to a better understanding of how robust these systems might be and how they operate as a whole. It is a credit to the vision of the Federal Reserve Bank to allow the redacted data on payment flows to be analyzed in this manner.

IV. A Recommendation to Make a Hierarchy of Models: From Simple to Complex (optional)

A recommendation made several times during our deliberations and earlier by Robert Oliver, Charles Taylor, Chuck Lucas, David Levermore, and myself is the idea of approaching the systemic risk problem with a hierarchy of models from simple to complex. I summarize some of these remarks below.

(i) Make minimal (simple) models (such as those presented by Shin and others) or agent-based models with simple sets of rules to see how much real variation in the data can be explained. The latter would include models such as the one proposed by Lux and Marchesi (Nature 1999) which have two-simple kinds of behavioral agents, but which can reproduce certain statistical properties of aggregate price series. The work in progress by the FRBNY-Skandia team on building an agent-based model for the Fed Payments Network, is a step in this direction that needs to be constrained more closely by the empirical work they have done. The importance of empirical validation should not be overlooked, and the meaning of the topological patterns uncovered by Soramaki et. al. needs to be understood (see below). There is much to be learned from simple models that can inform the systemic risk problem at the most general level.

(ii) Make more complex mechanistic models to complement the simpler ones. This is an ideal to aim for, but it needs to be done carefully and in tandem with the simple models. Nonlinearities in functional relationships fix the scale of the model mechanisms (aggregation problem) and can hinder the applicability (generality) of this across different scales of markets (firm-industry-regional-national-global). Again this needs to be done carefully. Another cautionary experience is from early ecosystem models that appear at first glance to be very complex (they incorporated everyone's favorite variable, so people believed them). Despite the apparent complexity, the overall model dynamic was

actually quite simple. Thus the apparent realistic detailed mechanisms were actually a kind of architectural decoration on a basic A-frame that was essentially simple logistic growth (e.g., the FORET model for forests in East Tennessee). Again, this is not to say this cannot be done well, but it needs to be done carefully and is very difficult (as Brian Arthur's experiments attest).

V. Lessons From Ecology and Complexity Theory

The analogy that Simon Levin advocates between ecosystems and financial systems and the generalization of complex systems thinking to this problem is compelling. In my view, the important function of this analogy is that it generates productive hypotheses that in turn can help to guide thinking and policy formation. Obviously, empirical corroboration of the analogy is one way to strengthen its utility.

Ecosystems are survivors of extreme stress testing.

Ecosystems are robust by virtue of their existence. They are the selected survivors of billions of years of upheaval and perturbation (continental drift, meteor extinctions etc), and show some remarkable constancy in structure that persists for hundreds of millions of years (eg. the constancy of predator/prey ratios, Baumgartner et. al.). As such, enumerating the common attributes of these diverse naturally selected surviving systems should be of interest to the systemic risk problem. Because experimental stress testing is not feasible in the financial sector, such common structural properties of ecosystems should be of interest, and may help to guide policy. They represent robust solutions for systems that have undergone extreme stress testing. They are templates of efficient design.

Are there empirical properties thought to affect stability in ecological systems that are shared by the Fed Inter-Bank Payments Network?

I enumerate a few features discovered here in the empirical pilot studies of the Fed Payments Network by Duffy-Ashcroft and by Soramaki et. al. that intriguingly appear to resonate with special patterns found in ecological systems.

VI. Empirical Attributes Of The Payments Network That Can Influence Robustness

- 1) Ignoring the Fed hub itself, the connectance of the Fed Payments Network is very low (0.003) and this connectance is characterized by a relatively small number of strong flows between nodes, with the vast majority of linkages being weak to nil. On a daily basis 75% of the payment flows are conducted over less than 0.1% of the nodes and only 0.3% of the observed linkages (which are already extremely sparse). Of interest here, is that these particular nonrandom characteristics of the network (very low connectance by **strong linkages** and ubiquitous weak interactions) are shared by most

ecological systems. This is in contrast to other man-made networks such as the electrical grid or the internet, which have much higher connectance, with linkages having more equal strength. Moreover, in the sense described by May-Wigner (May, Sinha & Sinha) the sparseness of **strong linkages** is thought to confer greater stability in systems whose components (nodes or banks) have some self-regulation or self-damping effect in operation (for example, where flows are potentially limited internally by the size of each institution). Here the stabilizing and self-damping behavior of individual nodes is in constant battle with the potentially destabilizing effects of the other nodes in the network. In the systemic risk context of debt obligations and assets, stability may be achieved, for example, where market pricing of assets is determined locally and spill-over pricing is subordinate. Regardless of specific mechanism, viewed generically, policies that encourage stronger-self limitation among banks and weaker inter-bank effects should be stabilizing.

- 2) To further strengthen Levin's analogy is the finding by both Duffy-Ashcroft and Soramaki et. al. showing that the frequency distribution of payment flows (linkage strengths) on a daily basis has a strong positive skew that is described as being approximately lognormal. This matches the observed frequency distributions of linkage strengths found in ecological systems, which are also positively skewed (approximately lognormal) (Goldwasser; Sala; Bascompte). The positively skewed distribution of interaction strengths implies a predominance of weak interactions, a feature of ecological systems that is thought to be stabilizing.
- 3) During the 9/11 crisis, the linkage pattern changed so that path lengths increased (higher modularity), connectivity decreased markedly, flows or interaction strengths decreased, and the effective size of the system fell (fewer nodes). The network behaved much like a percolation network with reduced flow. Nonetheless, it is interesting to note that these modifications of the inter-bank payment network during crisis are changes to the system that are thought to confer greater overall dynamic stability to a network. That is to say, during the 9/11 crisis the Fed inter-bank payments system acquired a configuration that to a first-approximation is consistent with enhanced stability.
- 4) The linkage pattern in the Fed Wire System was highly nonrandom in being strongly disassortative. That is, large banks were selectively connected to small banks and visa versa. I suggest a possible reason for this might derive from a simple business motivation. Small banks prefer the implied safety and stability of doing business with larger institutions, who, in turn find they can charge a larger premium from smaller banks. It seems like common sense. However, the dynamic implications of the resulting disassortative linkage pattern are interesting. Bank payment networks can be thought of as mutualistic (characterized by positive feedbacks) in their payment flows. However from an abstract mathematical viewpoint, such systems are difficult to stabilize, and this is a reason given to the relative rarity of mutualistic networks in the ecological world (May). Quite interestingly this curious disassortative pattern is also found uniquely in mutualistic plant-pollinator webs (Bascompte) where generalist pollinators are found to associate with plant specialists, and plants visited by many species of pollinator, do so with pollinators who specialize in that plant. Such disassortative or asymmetrical interaction patterns are characterized by the presence of a few hubs with many spokes. Again, in ecology they appear to be unique

to mutualistic networks and it has been suggested that this peculiar feature confers dynamic stability to simple mutualistic (positive feed back) systems. I think this is the most interesting parallel implied in the descriptive results obtained in these pilot exercises. To the extent that the analogy applies, it suggests that policies that disrupt this asymmetry such as ones that reduce the size inequity between banks might be destabilizing.

- 5) Soramaki et. al. claim that the average daily network of inter-bank payments in Fed Wire System is a scale-free network. This property is shared by other systems such as the internet, but it is not a general characteristic of ecological food webs. Scale-free networks can arise as an expression of positive feedback where the rich (highly connected nodes) get richer (attract more connections). Networks having this property are generally robust to random elimination of nodes but are susceptible to targeted attack (in the sense of network fragmentation). Although scale-free networks can be problematic to document for unique statistical reasons (Stumpf & May), this does not appear to be a problem in their analysis.

VII. More on the Analogy Between Food Webs the Payments Network.

- 1) The intersection graphs (where two species are connected if they use the same resource) of simple food webs tend to be dominated by rigid or choral circuits (circuits effectively paved with triangles). This special property of niche overlap graphs confers dynamic stability to the web. It would be interesting to see if the intersection graphs for the payments network (where two banks are connected if they receive payments from the same bank) have a tendency to contain an excess of rigid circuits. As in all analyses above, the Fed source links would of course be ignored.
- 2) Food chains tend to be short, typically having 4 to 7 trophic levels. It is believed that these chains are shorter in more heavily disturbed environments. It would be interesting to see if the payments network also contains short chains, and if these became shorter during the 9/11 crisis.
- 3) Weaker links in food webs are believed to be more ephemeral and have more temporal variability than the strong linkages. It would be interesting to see if this applies to the payments network.
- 4) The payment flows network is similar to webs in ecology that describe energy or nutrient flow. These ecological webs are essentially accounting diagrams and in various simple models it is assumed that there is conservation of energy or biomass. As an ecologist it was interesting to see how weakly this assumption is upheld in the payments network, in terms of the daily variability of total payment throughputs.

VIII. Future Work:

The initial results of the pilot study to describe the topology of payment flows in the Fed Wire Network suggests to a first approximation that it is likely a robust dynamic system. Moreover, in many curious ways it shares robust properties with ecological systems (network topology and the distribution of its interactions). The scale-free nature of the

network is consistent with other man-made networks but this property does not apply to ecological systems with any generality. The attack sensitivity associated with scale-free structure does not necessarily imply systemic risk, beyond the immediately contingent nodes (see 1 below).

Suggestions for future work are as follows:

- 1) Given the apparent attack sensitivity of the Payments network, it seems worthwhile to identify the major nodes that are most sensitive to targeted removal, that is, the nodes resulting in the largest collapse of the downstream network. These nodes can be found by constructing so-called dominator trees (sensu Tarjan) for the Payments graphs. A dominator tree is an operation on directed graphs that allows one to locate the most sensitive nodes in a network. It involves the construction of a tree graph whose nodes are linked to each other by obligate paths. That is, it identifies the bottlenecks and all of the downstream nodes that will effectively perish if that bottleneck is targeted. Here again sensitivity is in terms of network fragmentation.
- 2) The study of payment flows is of immediate interest to central bankers but it misses an essential aspect of systemic collapse. Namely the contagion dynamics of asset valuation mediated by the interaction of balance sheets. In the spirit of producing multiple models it would be interesting to look at the dynamic network of balance sheets and if possible quantify the mutual effects of valuations, credit policies, hedging etc. between financial institutions, and especially investment banks. These balance sheet networks would be useful for studying the effects asset pricing bubbles, credit crises, and the potential effects of the current widespread use of derivatives and dynamic hedging by investment banks.

Footnotes:

- *NOTE: A “regime shift” is defined as a change in government or a change in governing rules. Here the interest is in whether the local character of the system dynamics is the same globally or differs regionally. As such nonlinear systems will contain different regimes each corresponding to the different regions of an embedded attractor in a state-space representation of the system. Therefore, separate “regimes” refers to the operation of different “local linear” (i.e. nonlinear) rules as opposed to global linear rules (a single global regime).
- **Note: Ricker’s son is quoted as describing his conservative father as “disliking change or any kind of innovation” to the extent that he never owned a dishwasher or other common appliances.