Risk Topography

Author(s): Markus K. Brunnermeier, Gary Gorton, and Arvind Krishnamurthy

Source: *NBER Macroeconomics Annual*, Vol. 26, No. 1 (2011), pp. 149-176
Published by: The University of Chicago Press
Stable URL: http://www.jstor.org/stable/10.1086/663991
Accessed: 24/02/2015 10:32

3

Risk Topography

Markus K. Brunnermeier, *Princeton University and NBER*
Gary Gorton, *Yale University and NBER*
Arvind Krishnamurthy, *Northwestern University and NBER*

I belong to those theoreticians who know by direct observation what it means to make a measurement. Methinks it were better if there were more of them.

—Erwin Schrödinger (quoted in Walter Moore, *Schrödinger: Life and Thought*, 1989, 58–59)

I. Introduction

The financial crisis of 2007–2008 dramatically revealed that it is time to rethink the measurement of economic activity. In particular, because of derivative securities, off-balance sheet vehicles, and other financial innovations, traditional measures of aggregate risk, such as leverage, are inadequate. It is imperative that we build an economy-wide risk topography, and submaps of different financial sectors of the economy. Measuring only cash instruments and income and balance sheet items is not sufficient for understanding the economy; instead we should measure risks, and think in terms of risks, in addition to quantities.

The situation today, and during the crisis, is not so different from the 1930s when Simon Kuznets, Arthur Burns, Wesley Mitchell, and their colleagues developed the first official measures of economic activity for the overall US economy, the National Income and Product Accounts (NIPA), and business cycle chronology. This occurred in the midst of and just after the Great Depression. Referring to the Great Depression, Richard Froyen (2009) put it this way:

> One reads with dismay of Presidents Hoover and then Roosevelt designing policies to combat the Great Depression of the 1930s on the basis of such sketchy data as stock prices indices, freight car loadings, and incomplete indices of industrial production. The fact was that comprehensive measures of national

978-0-226-00214-9/2012/2011-0301$10.00

income and output did not exist at the time. The Depression, and with it the growing role of government in the economy, emphasized the need for such measures and led to the development of a comprehensive set of national income accounts (Froyen 2009, 13).[1]

During the financial crisis of 2007–2008 policymakers faced a similar problem. Relevant information about the financial sector and its linkages to the real economy was missing. Very basic measures were inadequate. For example, a measure such as "leverage" has little meaning in a world with derivatives and off-balance sheet vehicles. "Liquidity" was not clearly defined, let alone appropriately measured. Existing measures did not account for the shadow banking system, the size of the repo market, or the extent of different financial institutions' exposure to residential mortgages and credit derivatives.

Measurement is the root of science, and is also the basis of macroprudential regulation and of firms' risk management systems. Recognizing these measurement problems, the Dodd-Frank Wall Street Reform and Consumer Protection Act (Pub. L. 111-203, H.R. 4173) includes a provision for the establishment of the Office of Financial Research (OFR), a new division within the Treasury. The OFR is tasked with providing research and information to the newly created Financial Stability Oversight Council. The OFR has subpoena power to require financial institutions to produce data that the OFR requests. One possible role for the OFR would be to implement new measurement systems. Similarly, in Europe, the European Systemic Risk Board (ESRB) was established to oversee the build-up of systemic risk.

In this paper we outline a system of measuring risks and liquidity in the financial sector and producing a risk topography for the economy. We see two tangible benefits to implementing these ideas.

First, such a measurement system would improve significantly on the standard accounting paradigms in capturing the risks that are most relevant for systemic risk assessment by regulators and financial market participants. The basic idea behind the measurement metrics is to elicit from financial firms their sensitivity to a number of prespecified factors and scenarios on a regular basis. Essentially, firms report their "deltas" with respect to the specified factors; that is, the dollar gain or loss that occurs when the specified factor changes by a specified amount. In addition, they report their liquidity deltas: the increase or decrease in their liquidity as defined by a liquidity index, the Liquidity Mismatch Index (LMI).[2] For example, we ask what the capital gain or capital

loss is to your firm if house prices fall by 5%, 10%, 15%, and 20%, and what if they rise by the same increments. By deviating from standard accounting paradigms for measurement and moving closer to risk-management scenarios, these metrics reflect derivatives, liquidity, and other important features of a modern financial system. For example, an important point we develop in Section III is that the liquidity/capital delta measures are more informative than accounting measures of leverage. Standard measures of leverage may be meaningless in a world of derivatives, while the liquidity/capital deltas will better measure the "fragility" of the financial sector.

The data can reveal risk and liquidity pockets in the economy. Currently, the absence of information about the risk exposures of the financial system mean that firms can be in "crowded trades" without knowing it. That is, their risk exposures may be viewed as small for their firm, but may be large if all other firms have a similar exposure. Data on risk pockets can trigger better private risk management as well as enhance regulatory risk assessment. The data can also detect trends in liquidity or risk imbalances in the economy. For example, the data may show the financial sector's reliance on the repo market grew over the 2000 to 2007 period and resulted in a significant liquidity imbalance for dealer banks. We discuss these types of uses of collected data in Section V.

Second, current macroeconomic models, which for the most part do not incorporate a financial sector, would have the essential data to guide such an endeavor. Theorists do not need data, but their thinking and their models are strongly influenced by what is measurable.

Solow (1970) is explicit that the stylized facts that were the outcome of the work of Kuznets and others were at the root of conceptualizing the neoclassical growth model, which is the current workhorse macro model. Burns and Mitchell's (1946) measurement of business cycles and Kuznets' work allowed Kaldor (1961) to state six "stylized facts" about the macroeconomy, which were instrumental in subsequent business cycle and growth research. The intellectual history is recounted by Lucas (1977) and Kydland and Prescott (1990). It seems clear that systematically collecting the relevant financial sector data will have an impact on the set of macro models that are developed.

Such macro modeling is essential for understanding systemic risk and financial crises. While the triggers for crises are varied, the amplification mechanisms that play out in crises exhibit common pat-

terns. These patterns may be direct due to contractual links or indirect through equilibrium feedbacks on asset prices and liquidity. The important question in assessing systemic risk is to ask, for example, if the commercial banking sector takes a $500bn loss next year, how will this spill over to other asset markets and players, and what will be the resulting system-wide or aggregate general equilibrium dislocations? Developing a data set on the actions and exposures of different parts of the financial sector in varying economic conditions can allow a researcher to develop quantitative models of common amplification mechanisms. For example, a key response indicator of the spillover effects is the liquidity mismatch index (LMI). We expect that firms with a very negative LMI will be forced to fire-sell assets and hence amplify the crisis and lead to excessive spillover effects. On the other hand, firms with positive (or moderately negative) LMI can ride out adverse effects and not cause any externalities. We discuss the use of data in modeling of systemic risk in Section VI.

Like the construction of the National Income and Product Accounts, it will take a significant effort and time to build this risk topography, although financial firms already currently produce much of the data that we suggest gathering.[3] We take advantage of the data and knowledge of the private sector internal risk models. Truth-telling can be ensured by cross-checking the various internal models across all market participants.

There would be substantial benefits to making such measurements publicly available, just as with other government-collected data (e.g., National Accounts, Bank Call Reports, Federal Reserve Flow of Funds, etc.). The responses can be aggregated, suitably anonymized, and then made public. An important principle is that the data be made publicly available to all (in a form that protects some proprietary responses).

Related Literature. Three strands of literature are related to the ideas in this paper: the first is on measurement, the second concerns stress testing. On measurement, in the United States the current systems include the bank Call Reports of Condition and Income and the Federal Reserve Flow of Funds data. Both of these data sets were explicitly developed to aid regulators to monitor banks. The Call Reports were mandated by the National Bank Act (1863) (Sec. 5211) and have continued (and been expanded) to this day. In essence, these reports contain fairly detailed balance sheet and income statement information of regulated banks, but fail to capture other financial institutions and risk sensitivity measures. Similar to the Call Reports, we emphasize eliciting the same sce-

narios repeatedly and regularly to develop a risk map. Over time, such data will become a large library of information that can be used to build and fine-tune models. Secondly, we emphasize that the data elicited (suitably anonymized) be made public so that academics, regulators, and industry participants will be in a position to build their own models of systemic risk.

The Flow of Funds data was designed by Morris Copeland (1947, 1952) to characterize money flows in the economy. Notably, at first, economists did not see how to use the Flow of Funds; see, for example, Dawson (1958) and Taylor (1958).

Central banks currently recognize that existing measurement systems are not up to the task and have begun to think about revisions and additions. See, for example, Eichner, Kohn, and Palumbo (2010) and European Central Bank (2010). Compared with these proposals and ideas, we suggest to fundamentally change the nature of the information that is collected by deviating from the accounting paradigm and operating under a measurement paradigm closer to risk management scenarios. We want to collect data that will, over time, be useful for developing macroeconomic models of crises. We argue that this requires data on risk. In addition, we emphasize that measuring "liquidity" is central to understanding crisis.

The second related literature concerns bank stress testing. Bank stress testing is an evaluation of the impact of a particular scenario or event on a firm, the scenario usually being a movement in financial variables. Stress testing is an adjunct to statistical models, such as value-at-risk models. There are many papers that provide a general introduction to stress testing. Examples include Blaschke et al. (2001); Jones, Hilbers, and Slack (2004); Cihák (2007); and Drehmann (2008). Collections of articles that discuss stress testing include Quagliariello (2009). International organizations have developed stress testing procedures: the Bank for International Settlements (BIS 2009), the Committee on the Global Financial System (2005), and the International Monetary Fund, which started the Financial Sector Assessment Program in May 1999. Other articles include Haldane, Hall, and Pezzini (2007) and Hoggarth and Whitley (2003). Hirtle, Schuermann, and Stiroh (2009) discuss the US Supervisory Capital Assessment Program (SCAP)—these were the stress tests applied to the largest US bank holding companies from February to May 2009 (see Board of Governors of the Federal Reserve System 2009a, 2009b). The data that we would like to collect are akin to that collected in the stress tests.

The third strand of literature is macroeconomic and banking theory, which guides our thinking concerning what data to collect. This strand is discussed in Section II.

Following Section II the paper proceeds as follows. In Section III we present simple examples to illustrate the data issues that arise in practice, and to motivate our approach to measurement. In Section IV we more formally present our approach of eliciting risk and liquidity deltas. In Section V we first discuss certain simple risk indicators for fundamental risks and liquidity risk. Section VI discusses the use of the data for macro modeling of amplification effects within the financial sector and the economy as a whole. Section VII concludes.

II. Guidance from Existing Theoretical Research

What data should be collected in order to better understand the vulnerability of the economy to systemic risk? Existing research in macroeconomics and finance can guide us in answering this question. Macro models with financial frictions focus on leverage and the dynamics of net worth/capital, limiting the leverage ratio, while models in finance highlight in addition the important role of liquidity.

The most influential macroeconomic models of financial market frictions are the works of Bernanke, Gertler, and Gilchrist (BGG) (1996) and Kiyotaki and Moore (KM) (1997). Technically, these models only feature a corporate sector that is subject to financial frictions rather than a financial sector subject to such frictions, but as Brunnermeier and Sannikov (2010) show, it is possible to rework these models so that the results are driven by frictions in the financial sector. We henceforth discuss these models in such terms and dispense with this qualification.

The BGG model emphasizes that the "net worth" of the financial sector is an important state variable in driving macroeconomic phenomena. Net worth is commonly thought of as the equity capital of the financial sector. Thus, in this model, when banks take losses that deplete their equity, they increase the rates charged on loans and/or cut back on lending, thus causing a credit crunch. The Kiyotaki-Moore model adds an important ingredient to this analysis. Agents in the model have collateral that they pledge to raise funds from lenders. Since the market value of agents' collateral is partly dependent on their financial health, it affects leverage in the system, which in turn affect the value of capital. With high leverage, losses deplete capital more dramatically and feedback to further reducing the market value of collateral, and so on.

The roles of net worth and leverage in these models are suggestive, but the challenge is to determine what these correspond to in reality. Most notably, as we will show with some simple examples in the next section, reliance on cash measures to capture net worth or leverage misses the effects of derivatives.

Work in the finance tradition emphasizes in addition the importance of "liquidity" for understanding financial crisis. Diamond and Dybvig (1983) is the canonical model in this literature. In this model, it is not just borrowing or leverage of the financial sector that is salient, but rather the proportion of debt that is comprised of short-term demandable deposits. More broadly, the literature describes that when the financial sector holds illiquid assets financed by short-term debt, the possibility of "counterparty run" behavior emerges that can precipitate a crisis. This literature also describes a feedback mechanism between capital problems and liquidity problems. See, for example, Allen and Gale (2004). When the financial sector runs into liquidity problems, triggered by runs by lenders, the sector sells assets whose prices then reflect an illiquidity discount. The lower asset prices lead to losses that deplete capital, further compromising liquidity. Brunnermeier and Pedersen (2009) model the interaction between funding liquidity and market liquidity for modern collateralized (wholesale) funding markets. Importantly, they model liquidity spirals and "collateral runs." An adverse shock heightens volatility, leading to higher margins/haircuts. This lowers funding liquidity and forces institutions to fire-sell their assets, thus depressing market liquidity of assets and increasing volatility further.

In sum, the existing micro-founded literature points to net worth/leverage of the financial sector, and liquidity exposure, often expressed as maturity mismatch, as key state variables that drive systemic crises.

III. Measurement Challenges—Four Examples

In this section we present some extremely simple examples to illustrate the measurement issues and to emphasize the weaknesses of traditional measures of leverage and maturity mismatch.

Even though leverage is well-defined in simple stylized models, it is an ill-defined measure in practice in current financial markets. Given derivatives and off-balance sheet vehicles, the standard leverage measure (on-balance sheet debt/equity) is at best noisy, and more likely useless, as a measure of the fragility of the financial sector.

Liquidity refers to many related concepts. Following the banking literature, liquidity mismatch in banks emerges when the market liquidity of assets is less than the funding liquidity on the liability side of banks' balance sheets. However, insurance of demandable deposits since 1934 make the textbook Diamond-Dybvig bank runs unlikely. On the other hand, it is widely understood that "collateral run" phenomena have been important in the asset-backed markets and the shadow banking sector in the 2007–2009 crisis (see Gorton and Metrick 2010). As another example, when a major financial institution—AIG is a good example here—is downgraded, its derivative counterparties will require that the institution post a large amount of collateral. This is a liquidity drain for the institution that is conceptually similar to the run by a number of short-term lenders.

The measurement issues that arise in practice are best presented in a series of very simple examples. The examples are simplified in the extreme and so they are clearly not realistic, nor are they intended to be. All values should be thought of as market values.

Benchmark: *Consider a firm with $20 of equity and $80 of five-year debt with a coupon rate of 4.5%. The firm makes loans to two different firms, each for $50 for one year at an interest rate of 5%.*

This example is a benchmark; it is a plain vanilla firm that resembles a traditional bank, though it does not take deposits. Call Report-type data would record the income and balance sheet items from this bank, and in this example that might suffice. The debt-to-assets ratio for this firm is 80%.

There are, however, some measurement issues even in this case. For example, the loans are one-year loans, but the debt is five-year debt. This bank is potentially facing a large loss if at the end of the year the term structure of interest rates were to change, resulting in a lower competitive rate for loans. For example, if the loans can only be made at 3% in one year's time, then this bank is facing a loss. Simple concepts like duration would capture this, but nothing that is currently reported would measure this interest rate sensitivity. Our measurement ideas involve asking what happens to firm value if, for example, the one-year loan rate in one year's time moves up by 100 basis points (bps), by 500 bps, down by 100 bps, or down by 500 bps, and so on?

Liquidity Mismatch: *Consider a firm with $20 of equity and $80 of debt as above, but now half the debt is overnight repo financing at 1% and the other half is five-year debt*

at 4.5%. The firm buys one Agency mortgage-backed security (MBS) for $50 (which is financed via repo at a zero haircut) and loans $50 to a firm for one year at an interest rate of 5%.

This example complicates the benchmark case by making the bank sensitive to funding risk, in addition to interest rate risk. What if the firm cannot renew the repo financing, and is forced to liquidate some of its assets? Standard measures, such as leverage, will not pick up this funding risk. That is, they will treat the overnight debt and the five-year debt symmetrically. One could construct a leverage measure that focused on the maturity mismatch in this example—such as a short-term leverage measure—but this too may prove inadequate. For example, suppose that instead of the Agency MBS, the bank owned $50 of private-label MBS, which is less liquid than the Agency MBS. Now this bank has more of a liquidity mismatch, stemming from the asset side. Thus it is clear that a liquidity measure needs to incorporate information from both the asset side of the balance sheet and the liability side (market liquidity and funding liquidity).

For this firm the Liquidity Mismatch Index (LMI), which roughly reflects the market liquidity on the asset side (price impact if sold at fire-sale prices) minus the funding liquidity on the liability side (effective maturity structure), we construct would be negative. Because the repo is overnight the firm is exposed to funding risk. In the next section, we discuss a liquidity index that measures funding and market liquidity risk. In this specific example, there is an MBS worth $50, which has a liquidity weight of, say, $\lambda_{ABS} = 0.9$, so the asset liquidity is $45. (Cash or Treasuries have a liquidity weight of one.) On the liability side, the MBS bond is funded by repo of $50 with $\lambda_{Repo} = 1$, so the Liability Index is –$50, which gives a net liquidity index of $–5. What happens if repo haircuts suddenly increase to 20%? Then in renewing the repo financing, the firm can raise less money against the MBS. The asset is less liquid in that borrowing against it raises less cash—say, now $\lambda_{ABS} = 0.8$. Then the Liability Index goes to $–10.

There is currently no measuring system (accounting or regulatory) that detects the sensitivity of a firm to change in market and funding liquidity conditions. The Liquidity Mismatch Index is designed to understand such potential stresses. In this example, one could further ask what would happen if the securitization secondary market were to become less liquid. That is, we could ask the firm to report its LMI if the liquidity weight on MBS was $\lambda_{ABS} = 0.5$, for example.

Rehypothecation: *The bank lends $100 to a hedge fund for three days and receives a bond with a market value of $100 as collateral (a reverse repo). The bank then uses the bond as collateral to borrow $100 in the overnight repo market. (Whatever else the bank is doing we ignore for purposes of the example.)*

The bank has a liquidity mismatch since the repo is overnight, but the reverse repo is for three days. If the repo does not roll over, then the bank must sell the bond or find some other funding. This sensitivity would be captured by our Liquidity Mismatch Index, discussed following. The liquidity weight on the three-day reverse repo loan is lower than on the overnight repo, entering negatively in the firm's liquidity index.

The Liquidity Mismatch Index is designed to capture the sensitivities to these kinds of issues, which were particularly important in the recent crisis, but which, again, are not captured by any current reporting system.

Synthetic Leverage: *Consider a firm with $20 of equity and $80 of debt; half the debt is overnight repo financing at 1% and the other half is five-year debt at 4.5%. The firm buys $100 of US Treasury securities and writes protection (using credit default swaps [CDS]) on a diversified portfolio of 100 investment-grade US corporates, each with a notional amount of $10; so there is a total notional of $1,000. The weighted-average premium received on the CDS is 5%.*

This firm is sensitive to movements in the term structure of interest rates, and also to funding risk, as in the previous examples. But now it is also quite significantly exposed to a macro risk that could cause failures of investment-grade firms. The risk is not idiosyncratic because the CDS portfolio is diversified, but if there were a recession in which three or four firms failed there would be losses on this portfolio. If we ask what would happen if four investment-grade US firms in their portfolio failed, with 50% recovery, the answer would be that the firm would be bankrupt because a loss of (50%)(4)(10) = 20, which is the amount of equity in the firm. Thus, the CDS creates "leverage" in this firm, which any standard measure will miss.

Note that the CDS position would be marked-to-market for accounting purposes. Thus, the marks would contain expectations about future defaults of the firms in the portfolio (and risk premia). However, we want to detect what would happen in specific events (e.g., four firms fail), rather than the probability weighted market price.

Also, note that this firm has another complication. Derivatives trade under the International Swaps and Derivatives Association master

agreement. This agreement usually has a Credit Support Annex (CSA), a legal document, which sets forth the conditions under which each party must post collateral. Suppose that in this example the CSA has collateral-posting requirements based on the market value of the CDS position. If the marks widen—that is, when it is more likely that a firm or firms in the portfolio will default—this firm will have to post collateral to the counterparty. It has a Treasury bond, which could be posted. The LMI calculation takes into account the CSA provisions. To see the issue note that if half the Treasury holdings are posted, then this falls out of the LMI. In the extreme, imagine that the entire Treasury holding is posted. Then the only remaining asset the firm has is the CDS portfolio. Measuring the liquidity index of this firm in the event that four firms fail will capture the liquidity risk of this firm.

As another example of a liquidity event triggered by derivatives, consider the effect of a ratings downgrade. The CSA typically prescribes that if the bank is downgraded during the term of the derivative contract, it will have to post more collateral, which again uses liquidity. Moreover, if the firm had written many derivative contracts—that is, the CDS as in the example, plus interest rate derivatives—the need for liquidity will apply to all derivative contracts. Thus, the downgrade is potentially a significant liquidity risk that arises when firms use derivatives.

Cross Scenarios: *Consider a firm with $20 of equity and $80 of debt; half the debt is overnight repo financing at 1% and the other half is five-year debt at 4.5%. The firm buys a Spanish residential mortgage-backed security (MBS), denominated in Euros, for the equivalent of $50 and lends the other $50 to a US firm. The firm does not hedge its Euro exposure.*

This firm is sensitive to house prices in Spain and to the Dollar/Euro exchange rate (as well as other risks). The bank may be fine if (1) Spanish house prices go down, but the exchange rate stays the same; or (2) the Euro weakens against the dollar, but house prices do not decline. But, if Spanish house prices decline and the Euro weakens against the dollar, then the bank may be in trouble.

This possible stress scenario would not be revealed by anything the firm would report to regulators or in Securities and Exchange Commission (SEC) filings, under the current system. The example would be a bit more complicated if the firm did hedge the exchange rate risk. Since the firm receives Euros on the MBS, but pays dollars on its debt (and equity dividends), it enters into a swap to receive dollars in exchange for

Euros. But, this transaction is with a counterparty, which might weaken in some states of the world, introducing counterparty risk. Also, as seen earlier, the firm might have to post collateral.

These examples are intended to illustrate the difficulties of measuring risks based on accounting measures. Many more such examples, increasingly complicated, can be produced. The point is that measurement systems based solely on accounting-type measures are inadequate.

IV. Measurement Metrics

In this section we explain our ideas using simple notation. We then introduce the Liquidity Mismatch Index, and discuss reporting. Finally, we say a bit more about what the exact scenarios could be.

A. Basic Set-up

There are two dates. Date 0 is the ex ante date at which each firm makes risk and liquidity decisions by choosing cash assets and cash liabilities, as well as derivative positions and off-balance sheet positions. Derivative positions may have a market value of 0 at date 0, but are sensitive to the risk factors. At date 1 a state $\omega \in \Omega$ is realized, some of which may be a systemic crisis, depending on what decisions firms have made. Firm i chooses assets A^i and liabilities L^i. The assets are a mix of cash, repo lending to other firms, derivative exposure, and outright asset purchases. Liabilities include short-term debt, long-term debt, secured debt, equity, and so forth.

The equity value of a firm i is given by $E^i_\omega = A^i_\omega - L^i_\omega$, where A^i_ω is the asset value in state ω and L^i_ω is the value of the total liabilities in state ω. The equity value E^i measures how close firm i is to insolvency and can feed into how the firm is likely to behave given considerations such as capital constraints, the risk of bankruptcy, managerial compensation contracts, and so on. In addition to the total value of assets and liabilities as well as the equity value, we are interested in the liquidity position of each firm.

For reporting purposes, the regulator specifies ω and elicits ΔA^i_ω, that is, the change in asset value as a result of state ω. This will be a dollar amount gained or lost in that scenario. The firm simply calculates the effect of the state on their current position; that is, the dollar gain or loss due to the current position being affected by changing to the specified

state. The firm *does not* take into account what their response would be to a change to that state. The response of the system as a whole can, in principle, be calculated using a macro model, discussed in the following.

B. The Liquidity Mismatch Index (LMI)

Each asset and liability is assigned a liquidity weight λ_ω^j for each state of the world. We index assets with positive j, while liabilities j takes on a negative value. Super-liquid monetary assets such as bank reserves and Treasuries of "flight to quality sovereigns" to have a λ_ω^{money} of one across all states. For something like an MBS, we can imagine measuring λ_ω^{MBS} as one minus the repo haircut on that MBS in state ω. Alternatively, λ_ω^{MBS} could measure the price discount that firm i has to accept if it immediately wanted to convert the asset into cash. The key point is that λ_ω^j measures the *immediate cash-equivalent value* of asset j across states. Aggregating liquidity across the asset side, one obtains firm i's asset liquidity $\Lambda_\omega^{A,i}$ for the different states in the economy. We also measure the liquidity of the liabilities, funding liquidity, as $\lambda_\omega^{j<0} < 0$. Overnight debt has liquidity of –1 in all states, while longer-term debt has $-1 < \lambda_\omega^{LT} < 0$. Common equity is $\lambda_\omega^{equity} = 0$ for all states ω. Note that assets that are held on margin contain a short-term debt component. If margins can be reset from 10% to 50% on a daily basis, 40% is essentially overnight debt. Aggregating all liability positions gives the total funding liquidity of firm i, $\Lambda_\omega^{L,i}$. Overall, firm i's liquidity position is: $\Lambda_\omega^i \equiv \Lambda_\omega^{A,i} - \Lambda_\omega^{L,i}$. Further, this can be aggregated across all firms (or across particular sectors) to determine the economy's liquidity state in state ω at time t.

As shown earlier, the regulator specifies the ω and the firm reports $\Delta\Lambda_\omega^i$, the change in liquidity due to being in that state or scenario. Also as previously shown, the firm simply calculates this delta. It does not try to take into account its response to a change to that state of the world. The firm also reports its current liquidity index, which we denote as Λ_0^i.

The liquidity weights are not zero or one. Determination of the liquidity weights is an empirical question. One way to determine these weights would be to set the base case weights based on spreads to LIBOR (the London Interbank Offered Rate). However the base case is determined, different liquidity scenarios correspond to different specifications of weights, shocking one or more at a time. Brunnermeier, Gorton, and Krishnamurthy (2011) discuss further details of the construction of the liquidity weights and the LMI.

C. *Reporting*

The dimensions of the Ω state space that describes a firms' asset, liability, and liquidity positions can be huge. We focus on states s within an S-dimensional factor space, a subspace of Ω. Factors consist of certain prices (risk factors) or liquidity/funding conditions (liquidity factors). For example, a risk factor might be a change in real estate prices, while a liquidity factor could be a change in haircuts or margins, or the shutdown of a given market. Individual market participants take these as a given, but they are endogenously determined in the financial system. The selection of factors is discussed in more detail in Section IV.

For the specified factors, firms report a "value-liquidity" vector that consists of their calculated equity delta and liquidity mismatch delta for each specified factor/state. For example, if there is only one risk factor (e.g., with N real estate price levels) and one liquidity factor (e.g., with M overall haircut levels), then the state space can be characterized by an $N \times M$-matrix. Firms have to report their estimated value-liquidity indices in the first row and first column. From this one can derive the partial sensitivities of each firm along each single factor. In addition, firms will be asked to report their value and liquidity indices for some prespecified cross scenarios, in the nth column, mth row.

In addition to repeating the elicitation of data for the same set of scenarios every reporting date, we also look for special, one-off stress scenarios, as described in the following.

D. *Factor Scenarios*

The "s-states" just described are stress scenarios. The choice of scenarios is critical to the assessment of systemic risk. There are two considerations driving the choice. First, the propagation and patterns of a crisis are similar across events. Crises invariably involve capital and liquidity problems in important parts of the financial sector. Shocks interact with these capital and liquidity problems and lead to adverse general equilibrium feedback loops. By collecting data on a core set of factors that are held constant over time, the data can shed light on the common propagation patterns that underlie all financial crises.

Second, history suggests that the trigger for crises varies from event to event. Thus, at any time the regulator needs to choose factors that are informed by prevailing economic conditions. For example, the regulator may choose to focus on the effects of an Internet stock price shock

in the late 1990s, but such a shock may not have been relevant in 2007, where subprime mortgages were a more significant concern.

Third, in most cases, particular cross-scenarios are of special interest; for example, a scenario involving a simultaneous change in house prices, unemployment, and liquidity. This is the thinking behind the successful bank stress tests (SCAP). One approach to institutionalizing the cross-scenario stresses is as follows. In each quarter, we ask each firm to submit a suggested cross-scenario that the firm deems to be the "worst-case" for itself. The regulator then examines these suggested cross-scenarios across firms to see if many firms have a similar worst case. If so, then that cross-scenario will be part of next quarters' survey.

Notably, scenarios can include events that have never happened before, that is, events that are not in recorded experience. Broadly, stress scenarios fall into four groups. The first three are specified changes in market risks, idiosyncratic risks, and in liquidity factors. The scenarios in the first three categories are orthogonal stress scenarios. These correspond to partial derivatives of value and liquidity indices with respect to the factor. The last group asks for more complicated cross scenarios; for example, what if house prices fall 20% and repo haircuts rise to 10%. We provide examples of scenarios.

Examples of market risk scenarios include specifications of changes in:

- Interest rates (yields) on major government bonds (e.g., United States, United Kingdom, Germany, Japan, China), bond rates for different maturities; also swap rates in LIBOR, FIBOR (Frankfurt), PIBOR (Paris), HIBOR (Hong Kong), and so forth, at different maturities.
- Credit spreads: Changes in major credit derivative indices (CDX, CMBX, LCDX) at different maturities.
- Exchange rates of major currencies.
- Stock prices, measured by major indices.
- Commodity prices, measured by sector and aggregate indices.
- Commercial real estate prices, for example, the NCREIF (National Council of Real Estate Investment Fiduciaries) property index.
- Residential house prices, for example, Case-Shiller index.

A liquidity risk scenario corresponds to a shock to liquidity as follows:

- Firms are unable to access the market to raise new cash for one month, three months, and six months.
- Repo haircuts on some asset classes rise.
- The syndicated loan market, or the securitization market, shuts down for some period.

Idiosyncratic risk refers to scenarios specific to the reporting firm. Such scenarios include:

- Default by the largest (second largest, third largest) of the firm's counterparties.
- Default by largest supplier of bank lines.
- Default by a major clearing bank, or clearing system.
- Reputational event, which prevents new security issuance for six months (one year).
- Inability to issue new securities for one month, three months, six months.
- Inability to clear for three days, ten days.
- Rating downgrade of the senior unsecured debt of the company by three notches, six notches.

These idiosyncratic scenarios can shed some light on the network-linkage effects that play an important role in crises. For example, we envision that the first item on counterparty exposures will be measured for the largest financial firms. However, we should stress that our measures are not designed to inform a regulator on the dynamics during a crisis. That is, it seems clear from recent experience that it would be useful for regulators to know in real-time the counterparty exposures to the default of AIG or the default on Lehman bonds. With such information, the regulator can make decisions on how or when to intervene during a crisis. For this, regulators would need data that is much more detailed (e.g., individual position data) than what we are suggesting here (e.g., see Duffie 2010). Our coarser measures, on the other hand, shed light on the risk build-up and possibility of systemic risk in advance of a crisis. From this standpoint, what is important is to measure the extent of, say, real estate risk held by major financial institutions, including, for example, Lehman and AIG. These firms are in a market equilibrium with other financial firms, so that real-estate losses can be expected to affect other firms; whether the loss transmission is through

a direct default, or through a firm unwinding a large risk position, lowering prices (and thus inflicting losses on other firms) is more detail than we think is necessary to understand systemic risk.

E. Discussion

If the goal is to provide information on a broad class of risk exposures, then elicited information should include events or scenarios that have never happened before; that is, events that are not in recorded experience. The bulk of the elicited information should remain constant over time to develop a panel data set, as with the Call Reports.

Elicited information will involve the use of models by firms. We view this as a desirable feature. Furthermore, they constantly evolve as innovation occurs. But we recognize that there will be problems. First, these models are not homogeneous across firms, so the elicited information will have different degrees of accuracy. Also, these models will evolve over time so that the accuracy of the responses will change over time (presumably it will improve). Second, some verification of models and outputs will have to be performed by bank regulators (who, in any case, will have to become more sophisticated). Cross-checks will need to be developed in order to verify the integrity of reported data. This could be done by horizontal comparison across similar firms. Also, adding-up constraints and zero-sum conditions should be applied wherever possible in order to make it difficult for market participants to misreport their exposures. See the next section on risk and liquidity aggregates for adding-up conditions.

One of the important aspects of the current crisis was that the amount of information supplied to regulators varied widely across different types of firms. Our view is that all firms significantly engaged in financial activities should report scenario responses. However, the amount supplied should vary by size. Larger firms should supply more data, smaller firms less data.

V. Firm and Regulatory Risk Management

This section first discusses model-free indicators of systemic risk that can be constructed from the value and liquidity deltas. Secondly, such data can be useful for regulatory risk assessments as well as to provide data that may enable firms to improve their own risk management. We discuss several applications of the risk measures. In the next section, we

turn to a more significant and more challenging macroeconomic step: in order to fully understand systemic risk one has to model the *endogenous response* of various market participants to adverse shocks and analyze the general equilibrium.

Risk Aggregates. The "delta" of asset value, A^i, with respect to a particular risk factor $\Delta^{i,A}$, is a measure of risk exposure that naturally aggregates over firms. That is, the sum

$$\sum_i^I \Delta^{i,A}.$$

is the total exposure of all measured firms to, say, real estate risk. We would expect that some firms are long exposure and others are short exposure. If we measured all important parts of the economy, the sum should equal the physical supply of risk. For risks in positive net supply such as real estate, we can arrive at what the sum should be considering how much real estate exists in the economy. For risks in zero-net supply such as pure derivatives, the sum should be zero. In both cases, the risk measures have the feature that they can be aggregated into something meaningful. Even at less aggregated levels—say, sectoral—the risk aggregates are likely to be informative. They will reveal pockets of risk concentration and can serve to diagnose systemic risk.

Liquidity Aggregates. The LMI measures can also be aggregated. An interbank loan that is a liquid asset for firm i is a drain on liquidity for the borrower, firm j (i.e., negative liquidity weight). Aggregating across firm i and firm j the interbank loan will net out. Consider the net LMI for firm i,

$$\Lambda^i = \Lambda^{A,i} - \Lambda^{L,i}.$$

Again consider the sum,

$$\sum_i^I \Lambda^i.$$

Summed across all sectors, the liquidity aggregate equals the supply of liquid assets: the Λ-weighted sum across all relevant liquid assets. The aggregate measures are analogous to Barnett's (1980) divisia indices for monetary aggregates. Barnett devised indices to weight different components of the money supply based on their usefulness as a transaction medium. The LMI index is similar but is based on both assets and liabilities, and has weights that reflect the financial liquidity of the asset and liability.

The aggregates are most interesting in describing the liquidity position of particular sectors. We may expect to find, for example, that the banking sector always carries a negative liquidity position, while the corporate sector or household sector carries a long liquidity position. The extent of liquidity transformation done by the banking sector may also be informative for diagnosing systemic risk. For example, in the period from 2000 to 2008, it is likely that the aggregate LMI grew substantially. However, for systemic risk purposes, what would have been most interesting is a diagnosis that the aggregate growth reflected a growing mismatch between the banking sector and the other sectors in the economy.

Intermediation chains. The LMI index can be constructed with either symmetric weights or asymmetric weights. We have discussed the index in the case where the weight of a loan from firm i (asset for that firm) is equal to the negative of the weight of that loan to firm j (liability for that firm). However, it may be interesting to construct asymmetric weights, so that the asset weight is set equal to, for example, 90% of the liability weight on the same transaction. With asymmetric weights, the LMI aggregate will decrease as liquidity chains grow. For example, it is widely thought that financial fragility is created by the long chains of assets and liabilities that underlie the securitization model (i.e., household mortgage, packaged into MBS, further packaged into collateralized debt obligation [CDO], and then serving as collateral for a repo, which may be rehypothecated many times). The aggregate LMI can measure this fragility.

Liquidity risk. A second dimension of liquidity that our measures shed light on is liquidity risk. The LMI can be aggregated in the different macro stress-events. For example, the LMI for a given bank in the event that housing prices fall by 10% tells us how much the liquidity of that bank will suffer with a housing shock. This liquidity risk measure can be aggregated, at sectoral levels and across the economy, to diagnose the liquidity risk of the economy.

Mutually inconsistent plans. The data on risk and liquidity exposures can reduce systemic risk by allowing the private sector to improve its own risk management. Consider the problem of mutually inconsistent plans, which is a recurring theme in many financial crises. As Grossman (1988) argues, in the 1987 market crash, firms were following dynamic trading strategies to insure them against market downturns. These strategies involved firms selling stocks as prices fell, replicating a synthetic put option. In the crash, it became apparent that such strategies were mutually inconsistent: if everyone follows a portfolio insurance

strategy, markets will not clear. Perhaps more relevant, in a bank panic, withdrawals—failure to roll over repo—can cause the very dangerous deleveraging that was the fear motivating the run in the first place.

To go back to the model, suppose that firm i (and all other firms), carrying out its risk management, planned that in the event certain prices fell, so that the firm's liquidity and solvency were reduced, it would sell some of its assets to reduce risk and at the same time cut back on interbank repo lending to preserve liquidity. The regulator elicits information on all firms and makes such information public. The firms then recognize that their plans are mutually inconsistent. As a result, all firms adjust their positions to reduce risk exposure and enhance liquidity.

In addition to simply making public the exposure data, we can imagine that a regulator will additionally provide guidance for private risk management. It could reveal the results of its own analysis that private plans are mutually inconsistent. Or, it could survey firms on how they would behave in the low price scenario and then make public the result of this survey. Such a survey may be similar to the senior loan officer's survey of the Fed or the new Fed survey of dealer credit terms (see Eichner and Natalucci 2010).

Systemically important institutions. New banking regulations require greater oversight and higher capital requirements for systemically important institutions. One cut at judging who is systemically important is to rank institutions by size of assets. However, this type of ranking suffers from all of the shortcomings of relying on balance sheet entries for asset holdings, which we have discussed earlier. Economically, it is more meaningful to judge firms in terms of their magnitude of their (aggregate) risk exposures and liquidity exposures. Thus, our data at the firm level can provide guidance on which institutions should be judged systemically.

Regulatory Capital Requirements. Notions of "risk-based" capital need some metric for measuring "risk." The approach outlined here naturally leads to the idea that regulatory capital be based on measures of risk that correspond to the deltas and LMI that we discussed earlier.

VI. Macroeconomic Modeling of Systemic Risk

The measures we have discussed in the previous section allow us to answer the following type of question: If real estate prices decline by 20% over the next year, how much capital and liquidity will the commercial banking sector lose? On one hand, it is possible that the banking sector

is well capitalized to handle such losses and the event is not systemic. On the other hand, it is possible that the losses trigger a credit crunch, fire-sales of assets, and so forth, so that the losses are amplified into a systemic crisis.

We define *systemic risk* as the risk that shocks affect the financial sector and trigger an endogenous adverse feedback significantly amplifying these shocks, causing further deterioration in the financial sector, and leading to significant output losses.

Systemic risk is inherently endogenous. There are endogenous feedbacks that can lead to a small shock having a large effect. This key point is often obscured (or absent) in policy discussions. To diagnose systemic risk, the data on shock-deltas must be viewed through the lens of a model. That is, it is not enough to measure the losses that may arise from a 20% fall in real estate prices. Measurement does not show how the impact of such a response can endogenously lead to a crisis. The important step for systemic risk assessment is to compute the general equilibrium response of the economy to such a shock. What behavioral response can we expect of different parts of the financial sector to the losses? What is the resulting general equilibrium? And when does that equilibrium feature a significant amplification mechanism?

To provide another example, the measures of the previous section describe how a firm's LMI will fall in some state, described by the vector of asset/liability λ's for that state. However, the λ's themselves are endogenous. That is, the liquidity of assets depends on the behavior of agents. If many key liquidity providers are insolvent in some state, then the asset market λ's will fall in that state, and such a fall in λ's may further compromise the financial sector, and so on. To understand systemic risk, one needs to compute a general equilibrium in which liquidity is treated endogenously.

As should be clear, this is a macroeconomic model building exercise and is a fertile area for both academic and policy research. This section describes how the data we have suggested collecting may be useful to further this research. To be clear, we do not propose here a specific model of systemic risk – such models do exist in the literature, and there are probably more that will be developed. Rather, we wish to illustrate how the data can be useful for macro modeling.

A. Modeling Behavioral Responses

Any macroeconomic model of systemic risk needs to posit a relation between a firm's decisions (lending, trading, etc.) and the firm-specific

variables such as the deltas we have discussed as well as macroeconomic state variables that may affect the investment opportunities faced by the firm. For example, moral hazard considerations may imply that as the capital of a bank falls, it reduces its lending and investments and moreover increases the premium it requires to take on risky investments. The LMI is another key response indicator. Models may predict that firms with very negative LMI will be forced to fire-sell their asset and cause negative spillover effects to others—even to market participants they have no direct contractual relationship with. Firms with positive (or moderately negative) LMI might hold on to their positions and hence limit the negative impact on others.

The data as outlined can help model agents' behavioral responses to a shock. Let us suppose that the deltas are collected over time to form a long panel data set.

Consider, for example, the measured Δ^i for real estate risk for firm i. The Δ^i comes from firm i's choices over the amount of MBS to own, both through derivatives and through direct holdings, as well as the amount of real estate lending the firm has undertaken. The time series of these Δ^i's is a time series of portfolio choices of the firm in different economic environments. These choices are influenced by the expected return and risk profile (both fundamental and liquidity-driven) of MBS. They are influenced by the overall macroeconomic environment. They are also influenced by the current capital and liquidity of the firm.

Likewise, consider the measured liquidity delta of the firm. This liquidity delta is influenced by whether MBS is owned directly or through derivatives, and whether it is funded using short-term debt, long-term debt, or equity. It is influenced by the collateral arrangements in the derivatives taken on by the firm. Again, the liquidity choices can be thought of as portfolio choices of the firm in varying economic conditions.

A macroeconomic model makes choices that result in a behavioral function $f^i(\cdot)$:

$$\Delta_{i,t} = f^i(\textit{own firm characteristics}_{t-1}, \textit{aggregate } \Delta's_{t-1}, \textit{macroeconomic state}_{t-1}).$$

Here, firm characteristics may include capital and liquidity, in addition to typical measures of organizational function (size, industry, etc.). The macroeconomic state may measure expected returns, volatility, and so forth. We have suggestively included a separate role for the aggregate Δs to highlight interdependences within the financial sector.

Since, as we have just discussed, the time series data reflects these choices in varying economy conditions, they will be useful in disciplin-

ing models of systemic risk. We view this as an important part of our approach, since a panel data set allows us to verify various future models using the whole history of data from the beginning of the data collection onwards.

B. Financial Aggregates

In much of the existing literature, the only financial measure that a macro-finance model aims to match is the credit spread (e.g., the Commercial Paper-Treasury Bill spread). The large gap between theory and measurement is striking and needs to be closed.

The data we would like to collect can form the "financial aggregates" that macro-finance models should aim to match. For example, suppose we used the data to construct aggregated risk-deltas across the financial sector, and formed a time series of these variables. Likewise, we aggregate the LMIs and liquidity-deltas across the financial sector to measure the liquidity and liquidity risk of the sector. The behavior of these variables would be akin to that of "leverage" in stylized models, in that high deltas may tell us something about the susceptibility of the financial sector to a meltdown.

In the spirit of Kaldor (1961), potential stylized facts concerning the interaction between the financial sector and macro could be:

1. The risk-deltas tend to display a high coherence with more traditional measures of economic activity, such as output or hours worked. That is, the risk-deltas of different sectors (e.g., housing, credit, commodities) tend to be positively correlated with output.

2. The risk-deltas in financial firms rises from trough to peak, and falls from peak to trough.

3. Risk becomes more concentrated over the cycle. This does not necessarily mean that, on average, risk in individual firms becomes concentrated in certain risk sectors.

4. Real estate-related risk is the main type of risk for 1 and 2.

5. The liquidity aggregate is countercyclical, declining as output rises.

6. The liquidity aggregate is positively related to commercial and industrial loans.

7. Liquidity risk is procyclical.

8. Risk-deltas and liquidity are negatively related to the Commercial Paper-Treasury Bill spread.

These conjectured stylized facts are suggestive of the types of moments that would need to be matched, in addition to moments relating to the real economy.

C. *Medium and Large Crises*

To model systemic risk, ideally we would like data that includes periods of extreme financial crises with large real economic fallouts, such as the current Great Recession. Fortunately or unfortunately, these extreme events are rare. However, there are numerous medium-size crises that occur more frequently. For example, recent experience has included the 1987 stock market crash and the 1994 mortgage market crisis, as well as the 1998 hedge fund crisis. These crises all reflected significant shocks to financial intermediaries. The shocks were amplified and spilled over across asset markets, but either because the shocks were small or because of significant government intervention, the crises involved negligible real effects.

If one assumes that the amplification mechanisms present in the medium-sized crises are also present in larger crises, collecting data on these crises can be extremely useful. Models can be built to match behavior in medium-sized crises, and these same models can be used in counterfactual exercises to gauge the effects of large shocks. The obvious caveat here is that such an exercise has to mechanically take nonlinear effects into account. We would like to stress that any model of the 1987 crash, for example, will likely already include significant nonlinearities.

Finally, it is worth highlighting the commonality with and differences to extreme event analysis in general. Extreme value theory and other methods covering rare events rely critically on certain statistical assumptions. The probability distribution of outcomes deep in the tails is typically assumed. In comparison, macroeconomic modeling involves assumptions about structural parameters that govern behavior both in medium events and in tail events. Such modeling is less subject to the Lucas critique. In addition, models of financial market frictions often describe behavior in terms of constraints, rather than beliefs or preferences. If constraints are tighter in extreme events, then it seems plausible that the models may better approximate behavior during such events so that a modeling exercise may perform better than a statistical exercise. However, like statistical models, modeling assumptions for

extreme events or behavior in extreme events are unavoidable due to limited data.

VII. Final Comments

The financial crisis is a strong reminder that measurement is at the root of science. The measurement systems that we currently have are outmoded, leaving regulators, academics, and risk managers in a dangerous position. Assessing systemic risk requires viewing data on the financial sector through the lens of a macroeconomic model. However, macroeconomics in particular frames questions and builds models based on available data, and we have so far lacked the data to construct macro-finance models.

In order to track systemic risk in the economy, there needs to be an overhaul of the reporting done by financial firms. While not perfect, our suggestions can be implemented, as the SCAP stress tests demonstrated. Based on our discussions with risk managers at major banks, most of the data can be produced by these firms already. Nevertheless, we recognize that there will have to be a lot of work to determine reporting standards (especially with respect to foreign subsidiaries, branches, etc., and with respect to different currencies). The reported numbers will change over time as the underlying models change even if the underlying risks remain the same. Presumably, over time models will get better, so reported numbers will become more accurate. This is true of any measurement system. There is no alternative. The OFR has already been mandated, so the first step with regard to infrastructure for measurement has already occurred.

Ours is not the only set of ideas for what new data collections systems should look like. In terms of what the Office of Financial Research (OFR) and the Federal Reserve System should collect, some have proposed that financial firms should submit *all* position data and *all* transaction data. There are several troubling problems with this proposal. First, the regulators would then have to develop summary measures themselves, independently of the firms from which they collected the data. Meaningful measures would be some kind of summary statistics of the data. We have suggested such summary statistics. But moreover, our ideas rely on the firms' own models. Without this the regulators would have to invent their own models, which not only seems like reinventing the wheel, but reasonably seems like they would not be as

good as firms' models. There is no reason for everyone to have to calculate the National Income and Product Accounts themselves. Another problem would be that the data that would be made available publicly would be the regulators' summary statistics, rather than the underlying data.

Throughout the earlier discussion we have mentioned several principles that seem important for any data project. We summarize them here.

- There is no substitute for better data. Without the proper measurement, science, macroeconomics, and risk management cannot make progress.
- Science is based on replication, reusing, and criticizing published work. This requires that scientific data be publicly available.
- Not all data should be collected, but a subset. Measurement must be intelligent and based on theory in designing summary statistics.

Endnotes

We thank Robert Berry, Slim Bentami, Darrell Duffie, Richard Evans, Mike Gibson, William Gonska, Vinit Joshi, Anil Kashyap, Sasa Pilipovic, Cheryl Ratheun, Martin Schmalz, Hyun Shin, Paul Walker, Marc Weinberg, and participants at the NBER Summer Institute (Capital Markets and the Economy), NBER Systemic Risk Initiative, IMF, CMU Accounting Conference, Banco de Chile, 2011 AEA meetings, and the 16th International Financial Risk Institute Annual Roundtable for their comments. For acknowledgments, sources of research support, and disclosure of the authors' material financial relationships, if any, please see http://www.nber.org/chapters/c12412.ack.

1. This quotation is also cited by Landefeld, Seskin, and Fraumeni (2008).
2. We focus on measuring these two dimensions—capital gain/loss and change in liquidity—because the theoretical literature on financial crises has centered on capital and liquidity as the most significant factors underlying the behavior of financial firms during crises.
3. The bank stress tests that took place during the financial crisis (the Supervisory Capital Assessment Program; see Board of Governors of the Federal Reserve 2009a, 2009b) also show that (large) financial firms are (at least embryonically) in a position to produce the numbers that we propose.

References

Allen, Franklin, and Douglas Gale. 2004. "Financial Intermediaries and Markets." *Econometrica* 72:1023–61.

Bank for International Settlements (BIS). 2009. "Principles for Sound Stress Testing Practices and Supervision." Basel Committee on Banking Supervision.

Barnett, William A. 1980. "Economic Monetary Aggregates: An Application of Aggregation and Index Number Theory." *Journal of Econometrics* 14:11–48.

Bernanke, Ben, Mark Gertler, and Simon Gilchrist. 1996. "The Financial Accelerator and the Flight to Quality." *Review of Economics and Statistics* 78:1–15.
Blaschke, Winfrid, Matthew T. Jones, Giovanni Majnoni, and Soledad Martinez Peria. 2001. "Stress Testing of Financial Systems: An Overview of Issues, Methodologies, and FSAP Experiences." IMF Working Paper No. 01/88. Washington, DC: International Monetary Fund.
Board of Governors of the Federal Reserve System. 2009a. "The Supervisory Capital Assessment Program: Design and Implementation." White paper.
———. 2009b. "The Supervisory Capital Assessment Program: Overview of Results." White paper.
Brunnermeier, Markus, Gary Gorton, and Arvind Krishnamurthy. 2011. "Liquidity Mismatch Measurement." Working Paper, Northwestern University.
Brunnermeier, Markus, and Lasse Pedersen. 2009. "Market Liquidity and Funding Liquidity." *Review of Financial Studies* 22 (60): 2201–38.
Brunnermeier, Markus, and Yuliy Sannikov. 2010. "A Macroeconomic Model with a Financial Sector." Working Paper, Princeton University.
Burns, Arthur F., and Wesley C. Mitchell. 1946. *Measuring Business Cycles.* New York: National Bureau of Economic Research.
Cihák, Martin. 2007. "Introduction to Applied Stress Testing." IMF Working Paper WP/07/59. Washington, DC: International Monetary Fund.
Committee on the Global Financial System. 2005. *Stress Testing at Major Financial Institutions: Survey Results and Practice*. Report issued by the Bank for International Settlements.
Copeland, Morris. 1947. "Tracing Money Flows through the United States Economy." *American Economic Review*, *Papers and Proceedings* 37:31–49.
———. 1952. *A Study of Money Flows in the United States.* New York: National Bureau of Economic Research.
Dawson, John. 1958. "A Cyclical Model for Postwar U.S. Financial Markets." *American Economic Review* 48:145–57.
Diamond, Douglas, and Phillip Dybvig. 1983. "Bank Runs, Deposit Insurance, and Liquidity." *Journal of Political Economy* 91:401–19.
Drehmann, Mathias. 2008. "Stress Tests: Objectives, Challenges and Modeling Choices." Riksbank *Economic Review* (June):60–92.
Duffie, Darrell. 2010. "Systemic Risk Exposures: A 10-by-10-by-10 Approach." Working Paper, Stanford University, Graduate School of Business.
Eichner, Matthew, Donald Kohn, and Michael Palumbo. 2010. "Financial Statistics for the United States and the Crisis: What Did They Get Right, What Did They Miss, and How Should They Change?" Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series, No. 2010-20.
Eichner, Matthew, and Fabio Natalucci. 2010. "Capturing the Evolution of Dealer Credit Terms Related to Securities Financing and OTC Derivatives: Some Initial Results from the New Senior Credit Officer Opinion Survey on Dealer Financing Terms." Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series, No. 2010-47.
European Central Bank. 2010. *Central Bank Statistics: What Did the Financial Crisis Change?* Frankfurt: European Central Bank.
Froyen, Richard. 2009. *Macroeconomics: Theories and Policies.* Englewood Cliffs, NJ: Prentice Hall.
Gorton, Gary, and Andrew Metrick. 2010. "Securitized Banking and the Run on Repo." *Journal of Financial Economics*, forthcoming.
Grossman, S. J. 1988. "An Analysis of the Implications for Stock and Futures Price

Volatility of Program Trading and Dynamic Hedging Strategies." *Journal of Business* 61 (3): 275–98.

Haldane, Andrew, Simon Hall, and Silvia Pezzini. 2007. "A New Approach to Assessing Risks to Financial Stability." Bank of England, Financial Stability Paper No. 2.

Hirtle, Beverly, Til Schuermann, and Kevin Stiroh. 2009. "Macroprudential Supervision of Financial Institutions: Lessons from the SCAP." Federal Reserve Bank of New York, Staff Report no. 409.

Hoggarth, G., and J. Whitley. 2003. "Assessing the Strength of UK Banks through Macroeconomic Stress Tests." Bank of England *Financial Stability Review* (June).

Kaldor, Nicholas. 1961. "Capital Accumulation and Economic Growth." In *The Theory of Capital*, edited by F. A. Lutz and D. C. Hague, 177–222. New York: St. Martins Press.

Kiyotaki, Nobuhiro, and John Moore. 1997. "Credit Cycles." *Journal of Political Economy* 105:211–48.

Kydland, Finn, and Edward Prescott. 1990. "Business Cycles: Real Facts and a Monetary Myth." *Quarterly Review of the Federal Reserve Bank of Minneapolis* (Spring):3–18.

Jones, Matthew, Paul Hilbers, and Graham Slack. 2004. "Stress Testing Financial Systems: What to Do When the Governor Calls." IMF Working Paper No. 04/127. Washington, DC: International Monetary Fund.

Landefeld, J. Steven, Eugene P. Seskin, and Barbara M. Fraumeni. 2008. "Taking the Pulse of the Economy: Measuring GDP." *Journal of Economic Perspectives* 22 (2): 193–216.

Lucas, Robert. 1977. "Understanding Business Cycles." In *Stabilization of the Domestic and International Economy*, edited by Karl Brunner and Allan Meltzer, 7–29. Carnegie-Rochester Conference Series on Public Policy 5. Amsterdam: North-Holland.

Quagliariello, Mario. 2009. *Stress Testing the Banking System: Methodologies and Applications.* New York: Cambridge University Press.

Solow, Robert. 1970. *Growth Theory: An Exposition*. New York: Oxford University Press.

Taylor, Stephen. 1958. "An Analytic Summary of the Flow-of-Funds Accounts." *American Economic Review, Papers and Proceedings* 48:158–70.