

The Known, the Unknown, and the Unknowable in Financial Risk Management

Francis X. Diebold, Neil A. Doherty, and Richard J. Herring

University of Pennsylvania

June 2008

I. Knowledge as Measurement and Knowledge as Theory

II. **KuU** Lessons for Financial Markets and Institutions

A. Invest in Knowledge

Better Measurement

Better Theory

B. Share Risks

*Simple insurance for **K***

*Mutual insurance for **u***

*Ex Post Wealth Redistribution for **U***

C. Create Flexible and Adaptive Structures

D. Use Incentives to Promote Desired Outcomes

*Organizations and Relationships: Principal/Agent Concerns in
Corporate Governance*

Contracts: Intentional Incompleteness and “Hold-up”

Investment Vehicles: Riding “Sidecar” with Better-Informed Agents

E. Use Policy to Limit Vulnerability to Shocks *Ex Ante* and Mitigate Consequences *Ex Post*

Regulating with Imperfect Information

Crisis Prevention

Crisis Management

III. Onward

I. Knowledge as Measurement and Knowledge as Theory

Knowledge is both measurement and theory. Observed or measured facts about our world have no meaning for us outside our ability to observe them into a conceptual model. For example, the numerous stones we find with what appear to be reverse images of animals and plants would be unremarkable if it were not for their place in our intellectual model of the world we live in. Without the evolutionary theories associated with Darwin, the fossil record would be no more than a collection of pretty stones. And indeed without the pretty stones, Darwin may not have conceived his theory.

When we speak of knowledge, there is no bright line which separates measured observations, from our theories. Though we may see the deficit at one, or the other, end of the spectrum, knowledge joins together observations of phenomena with conceptual structures that organize those facts in a manner that is meaningful to our wider human experience. Consider the following.

“When you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meager and unsatisfactory kind: it may be the beginning of knowledge, but you have scarcely, in your thoughts, advanced to the stage of science” Lord Kelvin (*Popular Lectures and Addresses*, 1891-1894)

“The whole machinery of our intelligence, our general ideas and laws, fixed and external object, principles, persons and gods, are so many symbolic, algebraic expressions. They stand for experience, experience which we are incapable of retaining and surveying in its multitudinous immediacy. We should flounder hopelessly, like the animals, did we not keep ourselves afloat and direct our course by these intellectual devices. Theory helps us to bear our ignorance of fact.” George Santayana (*The Sense of Beauty*, 1896).

Thus, if we talk of what is known and what is unknown, we may be referring to the presence or absence of data to collaborate our theories, or to the inability of our theories to provide meaning to the curious phenomena we observe and measure.

For this volume, we have adopted the taxonomy of knowledge used in a famous article by Ralph Gomory (1995). Gomory classifies knowledge into the Known, the Unknown and the Unknowable, for which we adopt the acronym **KuU** (the **K**nown, the **u**nknown and the **U**nknowable). As applied to both knowledge-as-measurement and to knowledge-as-theory, this classification can be described as:

1. **Knowledge as Measurement.** If we define a probability distribution as a complete and exhaustive set of outcomes, together with associated probabilities, that may befall some variable of interest, then:
 - a. **K** refers to a situation where the distribution is completely specified. For example; the distribution of automobile or life insurance claims for an

insurance company is more or less known. This is equivalent to Frank Knight's (1921) definition of *risk* – both outcomes and probabilities are specified.

- b. ***u*** refers to a situation where probabilities cannot be assigned to at least part of the event space. The systemic risk to financial systems and terrorism risk might fall into this category. This is equivalent to Knight's definition of *uncertainty* – events are specified but probabilities are not.
 - c. ***U*** refers to a situation where even the events defining the space cannot be identified in advance. Once they occur, they enter the domain of ***u***. Retrospectively, the surge of asbestos claims for long standing (real or imagined) injury is an example as indeed are many legal actions caused by innovative legal theories.
2. ***Knowledge as Theory***. The second approach focuses on the conceptual model that helps us to understand the underlying structure of the phenomenon of interest.
- a. ***K*** refers to a situation where the underlying model is well understood. We may refer to this a paradigm. This is not to say that the model is correct, only that experts are in broad agreement. For example, scientific models of evolution based on Darwin refer to a situation of scientific knowledge. We may not agree on all the details, but there is almost universal agreement amongst scientists on the broad structure. We might say there is “knowledge” on the broad principles of corporate governance, multi factor asset pricing and on the general principles of credit and interest rate risk management. Thus, in short ***K*** refers to *theory*
 - b. ***u*** refers to a situation where there are competing models, none of which has ascended to the status of a paradigm. Operational risk management probably falls into this category. Other examples might include the performance of markets and financial institutions in emerging economies. If ***K*** refers to theory, then ***u*** refers to *hypothesis* or, more weakly, *conjecture*.
 - c. ***U*** is where there is no underlying model (or no model with scientific credibility). This does not mean that we cannot conceivably form hypotheses and even theory in the future. But until some conceptual model is constructed, we cannot understand certain phenomena that we observe, or we may not even be able to identify phenomena because it never occurs to us to look. For example we would never look for black holes unless we had a theory about how matter behaves under extreme gravitational forces.

The two taxonomies are complementary. For example, the inability to specify the tail of a distribution might be due to both the absence of data and the deficiencies of the statistical theory. Thus innovations such as extreme value theory can lead to upgrading of knowledge (from ***U*** to ***u*** or from ***u*** to ***K***) under both taxonomies. As another illustration, the absence of a theory for a yet-to-be-identified phenomenon is hardly surprising and the emergence of such events will generate an interest in both measurement and theory.

For the most part, the various authors in this volume adopt the **KuU** framework (perhaps this is not too surprising since we did bully them gently toward a common terminology), though most use it to address knowledge-as-measurement issues and some papers modify the framework. For example, Richard Zeckhauser notes that, as regards measurement, we could otherwise describe **KuU** as *risk*, *uncertainty* and *ignorance*. Similarly, Howard Kunreuther and Mark Pauly use the alternative *ambiguity* in a similar manner to our **u** and Knight's *uncertainty*. However, the most common chomping at the **KuU** bit was in insisting that we look at informational asymmetries. For example, Ken Scott looks at corporate governance in **KuU** terms, driven partly by informational (and skill) differences between managers and owners. Similarly, Zeckhauser observes that some investors have informational and skill advantages over others and then examines how uninformed investors, who recognize their inferior status, form strategies to benefit from the higher returns that can be earned from the knowledge and skills they lack.

A related theme that arises in some of the papers is that the language used by different stakeholders depends on what is important to them. Clive Granger in particular notes that risk means different things to different people. Most particularly, many people think mostly of the downside of risk because that is what worries them. Thus he emphasizes downside measures of risk many of which (such as the various value at risk measures) have become increasingly important in risk management. Similarly, Scott notes that the conflict of interest that lies behind corporate governance is partly due to the fact that different stakeholders emphasize different parts of the distribution; undiversified managers may be more focused on downside risk than diversified shareholders.

II. **KuU** Lessons for Financial Markets and Institutions

What to do in the nitty-gritty of life in the financial trenches, in which people grapple daily with risk measurement and management in situations of **K** and **u** and **U**? We think of **KuU** as more than simply an acronym for “the known, the unknown and the unknowable;” indeed, we think of it as a *conceptual framework*. We believe that “**KuU** thinking” can promote improved decision making – recognizing situations of **K** and **u** and **U** and their differences, using different tools in different situations, while maintaining an awareness that the boundaries are fuzzy and subject to change.

Perhaps the broadest lesson is recognition of the wide applicability of **KuU** thinking, and the importance of *each* of **K** and **u** and **U**. **KuU** thinking spans all types of financial risk, with the proportion of **uU** increasing steadily as one moves through market, credit, and operational risks. In addition, **KuU** thinking spans risk measurement and management in all segments of the financial services industry, including investing, asset management, banking, insurance, and real estate. Finally, **KuU** thinking spans both the regulated and the regulators: regulators' concerns largely match those of the regulated (risk measurement and management), but with an extra layer of concern for systemic risk.

In the following sub-sections we highlight several practical prescriptions that emerge from **KuU** thinking, distilling themes that run through subsequent chapters.

That we will treat \mathbf{K} first is hardly surprising. \mathbf{K} is comparatively easy to navigate, and it is natural to take the low-hanging fruit first. (We hasten to add, of course, that \mathbf{K} is also very important.) Indeed the existing risk management literature focuses almost exclusively on \mathbf{K} , as summarized for example in the well known texts of Jorion (1997), Doherty (2000) and Christoffersen (2003) and emphasized in the Basel II capital adequacy framework, which employs probabilistic value-at-risk methods to set minimum capital requirements.

Perhaps surprisingly in light of the literature's focus on \mathbf{K} , however, we ultimately focus more on situations of \mathbf{u} and \mathbf{U} here and throughout. The reason is simple enough: reflection (and much of this volume) makes clear that a large fraction of real-world risk management challenges fall largely in the domain of \mathbf{u} and \mathbf{U} . Indeed a cynic might assert that, by focusing on \mathbf{K} , the existing literature has made us expert at the least-relevant aspects of financial risk management. We would largely disagree, as \mathbf{K} situations are often of relevance, but we *would* assert that \mathbf{u} and \mathbf{U} are at least equally relevant, particularly insofar as many of the "killer risks" that can bring firms down lurk there.

A. Invest in Knowledge

Although life is not easy in the world of \mathbf{K} , it is easier in \mathbf{K} than in \mathbf{u} , and easier in \mathbf{u} than in \mathbf{U} . Hence one gains by moving leftward through \mathbf{KuU} toward knowledge, that is, from \mathbf{U} to \mathbf{u} to \mathbf{K} . The question, then, is how to do it; how to invest in knowledge? Not surprisingly given our taxonomy of knowledge as measurement and knowledge as theory, two routes emerge: better measurement and better theory. The two are mutually reinforcing, moreover, as better measurement provides grist for the theory mill, and better theory stimulates improved measurement.

Better Measurement

Better measurement in part means better data, and data can get better in several ways. One way is more precise and timely measurement of previously-measured phenomena, as for example with increased survey coverage when moving from a preliminary GDP release through to the "final" revised value.

Second, better data can correspond to intrinsically *new* data about phenomena that previously did not exist. For example, exchange-traded house price futures contracts have recently begun trading. Many now collect and examine those futures prices, which contain valuable clues regarding the market's view on the likely evolution of house prices. But such data could not have been collected before – they simply did not exist. Papers like Bardhan and Edelstein's sweeping chronicle of \mathbf{KuU} in real estate markets call to mind many similar such scenarios. Who, for example, could collect and analyze mortgage prepayment data before the development of mortgage markets and associated prepayment options?

Third, better data can arise via technological advances in data capture, transmission and organization. A prime example is the emergence and increasingly-widespread availability of ultra-high-frequency (trade-by-trade) data on financial asset prices, as emphasized in Andersen, Bollerslev, Christoffersen and Diebold (2006). In principle, such data existed whenever trades occurred and could have been collected, but it was the advent and growth of electronic financial markets – which themselves require huge computing and database resources – that made these data available. These technological advances, however, may threaten to leave us data rich, but information poor. How to summarize risks in an efficient way is a continuing challenge.

Finally, and perhaps most importantly, better financial data can result from new insights regarding the determinants of risks and returns. It may have always been possible to collect such data, but until the conceptual breakthrough, it seemed pointless to do so. For example, traditional Markowitz risk-return thinking emphasizes only return mean and variance. But that approach (and its extension, Sharpe's celebrated CAPM) assumes that returns are *Gaussian* with *constant variances*. Subsequent generations of theory naturally began to explore asset pricing under more general conditions, which stimulated new measurement that could have been done earlier, but wasn't. The resulting explosion of new measurement makes clear that asset returns – particularly at high frequencies – are highly non-Gaussian and have non-constant variances, and that both important pitfalls and important opportunities are embedded in the new worldview. Mandelbrot and Taleb, for example, stress the pitfalls of assuming normality when return distributions are in fact highly fat-tailed (badly mis-calibrated risk assessments), while Colacito and Engle stress the opportunities associated with exploiting forecastable volatility (enhanced portfolio performance fuelled by volatility timing).

Thus far, we have treated better measurement as better data, but what of better tools with which to summarize and ultimately understand that data? That is, what of better statistical / econometric models? We hasten to add that if better measurement sometimes means better data, it also sometimes means better models; the two are obviously not mutually exclusive. Volatility measurement, for example, requires not only data but also models. Crude early proxies for volatility, such as squared returns, have been replaced with much more precise estimates such as those based on ARCH models. This allows much more nuanced modeling, as for example in the paper by Colacito and Engle, who measure the entire *term structure* of volatility. They construct a powerful new model of time-varying volatility that incorporates nonstationarity and hence changing distributions, nudging statistical volatility modeling closer to addressing uU . Similarly, Litzenberger and Modest develop a new model that allows for regime switching in the data, with different types of trading strategies exposed to different crisis regimes, and with regime transition probabilities varying *endogenously* and sharply with trading, hence allowing for “trade-driven crises.”

In closing this subsection, we hasten to add that although better data and models may help transform u into K , the role of better data in dealing with U is necessarily much more speculative. To the extent that U represents a failure of imagination, however, the collection and analysis of data regarding near misses – disasters that were narrowly

averted – may provide a window into the domain of U and alternative outcomes. The challenge is how to learn from near misses.

Better Theory

As we indicated earlier, the literature on the behavior of markets and institutions, and the decision making that drives them, is almost exclusively based in the world of K .

Accordingly, risk prices can be set, investors can choose strategies that balance risk and reward, managers can operate to a known level of safety, regulators can determine a standard of safety, and so forth. Similarly, a variety of information providers, from rating agencies to hazard modeling companies, can assess risk for you, if you want to verify or supplement your own efforts.

Not only does the literature rely on the assumption of K , but proceeds also with the assumption that actors are sophisticated and rational. For example, our theory of individual decision making is based largely on expected utility maximization. A similar level of rationality is required in sophisticated enterprise risk management models that are now available and increasingly in use.

Even in situations of K , however, the assumption of sophistication and rationality is questionable. As Granger emphasizes in his paper, people don't always behave according to the seemingly innocuous axioms of expected utility, as many experiments have shown. If the rationality battle is joined, however, its outcome is not yet determined. The behavioral economists are providing evidence of anomalies and the neo-classical economists are showing that many of these anomalies can be explained by properly formulated models.

If behavioral economics has had some success in the K world, one might suppose that it will become even more important in the uU world of scant knowledge. This point is addressed in several of the papers. For example, Kunreuther and Pauly address unknown but catastrophic losses such as major acts of terrorism and point to framing anomalies. For example, there may be an “it can't happen to me” mentality that forestalls action.¹

Construction and application of such “better theories” – theories geared toward the worlds of u and U – appear throughout the volume. For example, forecasting terrorism from past data for the purpose of countering future attacks would play into the hands of terrorists who would therefore be able to predict and avoid counter-terrorism strategies. Thus, forecasting would be self defeating. Instead, such situations can be modeled by game theory. Similarly, as Zeckhauser notes, investing in a K world may play into the hands of the math jocks; but not so when probabilities are unknown and the potential scenarios that can drive them unknowable. Zeckhauser outlines some innovative strategies to cope in this world, always supposing one has the tolerance for the unknown and a relatively low aversion to blame when things turn out badly. Other papers ask whether, given that we cannot anticipate what might happen in the future, we can

¹ However, it is somewhat difficult to entertain the usual alternatives to expected utility, such as prospect theory where the derivation of a weighting function for unknown probabilities seems an empty exercise.

nevertheless, arrange our affairs such that the appropriate decisions will be made in response to the new circumstance? For example, Scott stresses the importance of long-term incentives to focus the CEO on the long-term survival and value of the firm in an unknown future. Doherty and Muermann stress that some unknowns might be insurable by means of incomplete contracts in which insurers deliberately offer the reputation as a bargaining chip under which policyholders can bargain for unanticipated and thus uninsured losses. And finally, despite normal reservations about the Samaritan's dilemma, government might happily accept their inability to pre-commit not to bail out the victim's of surprise catastrophic losses – and indeed, such *ex post* allocations might represent efficient response to unknowable calamities.

B. Share Risks

The desirability of risk sharing is hardly novel. An emergent and novel theme, however, is the desirability – indeed the necessity – of tailoring risk-sharing mechanisms to risk types.

Simple Insurance for K

Operations of financial institutions are greatly simplified when dealing with “known” assets and obligations. If a bank or insurance company has assets and liabilities with known distributions, appropriate risk adjusted prices (interest rates and insurance premiums) can be set. The main challenge comes from the correlation structure of the assets and liabilities and from incentive problems, such as adverse selection and moral hazard. Enterprise risk management models help institutions to estimate the economic capital required to operate with an appropriate level of survival probability.

The task of the regulator is likewise simplified when distributions are known. Regulators typically impose minimum capital requirements. They monitor institutions and may intervene in the event of distress. Moreover, as described by Charles Goodhart, the sophistication of bank solvency regulation may have been ratcheted upwards in the form of Basel II. In the case of insurance, Europe's Solvency II parallels the Basel II initiative, though in many cases insurance regulators (particularly in the United States) seem to be conflicted between their desire for solvency and their desire for cheap and available insurance, especially in high-risk and low-income locations. This regulatory consumerism has been particularly active in several areas subject to high levels of catastrophe risk, notably Florida where prices in the highest risk locations (e.g. barrier islands) are subsidized either from insurer capital or from lower risk policyholders (Kunreuther and Pauly). The risk in question here is storm risk, which though it can be catastrophic in nature, is essentially a **K** risk in our taxonomy. Storm risk is indeed intensively modeled, and insurers routinely construct their portfolios and set prices using this model information. This provides an example of market disruption arising for **K** risks, though one can debate whether this is really a market failure or a regulatory failure.

The role of the regulators is supplemented by that of rating agencies. By supplying the market with information on default risk, they support the process of risk-adjusted pricing.

Furthermore, agencies such as Standard and Poors are now not only building their own ERM models, but also rating companies on the basis of their internal models

The general picture we see here is that, in the case of **K** risks, financial institutions can pool idiosyncratic risk and reserve against expected losses, such that the risks to bank depositors and insurance policyholders are small. The remaining systematic risk is controlled by setting economic/regulatory capital. Any defaults of institutions, in this stylized model, are not informational failures, but rather the result of inadequate provision of costly economic capital. Of course, there are in practice some frictions, such as adverse selection and moral hazard, which can be addressed through contract design. However, the implication of pure **K** would be that all parties know the distribution, and the information asymmetries would disappear.

Mutual Insurance for u

For risks that are unknown, the potential events can be identified, but the problem is in attributing probabilities. One could choose how to refine this definition to address correlation. One possible definition would suggest that, for **u** risks, we neither know the probabilities nor the correlations. This definition is appropriate if we define events in terms of their consequences to individuals. Thus you and I are each exposed to storm risk, bank default, melt-down on our investment portfolio, etc. Another tack would be to define the events to include both the individual and the collective impact. Thus the following four events define an exhaustive and mutually exclusive set: (1) I may suffer a storm loss and you do not, (2) you may suffer a storm loss and I do not, (3) neither of us suffers a storm loss, and (4) both of us suffer a storm loss.

Which definition is appropriate depends on context. However, the point is that we can imagine uncertainties that affect individuals in different ways and those that impact people collectively. Consider climate change. It can, by its nature, impact the entire world, and we are certainly unsure about the particular form and magnitude of many of its effects. Thus it is clearly in the realm of **u**, and the nature of the uncertainty surrounding global warming spans both the global and local impacts. The extent of any rise in sea level is unknown, but it will have a common impact on all coastal areas, exposing them all to increased flood and tidal surge risk. However, the impact of rising temperatures on drought risk may vary considerably across regions in ways that we do not fully understand and cannot predict. In the former case of rising sea level, the uncertainty is “correlated,” and in the latter case of drought risk, the uncertainty is of “lower correlation.” This distinction is important in determining how effectively **u** risk can be pooled in insurance-like structures.

In the case of uncorrelated uncertainty, there is no real obstacle to the pooling of risk; for some locations, the probabilities (and therefore the randomly realized outcomes) turn out to be higher and, for others, the probabilities (and thus the realized outcomes) turn out to be lower. This is simply a two stage lottery; in stage 1, the distribution is randomly chosen and, in stage 2, the outcome is realized. Insurance mechanisms can cover both the stage 1 distribution risk and the stage 2 outcome risk as long as both are random and

uncorrelated. Insurance on stage 1 is essentially hedging against future premium risk, and insurance in stage 2 is hedging future outcome risk.

For correlated uncertainties, pooling of risk is more challenging. At any realization in stage 1 of the lottery will have a common impact on both the overall level of prices and on the level of economic capital required to undertake stage 2. In principle, the optimal sharing of risk is not too difficult to envision. Indeed, it is similar to that proposed by Borch (1962) for the case of known, but correlated risks. A mutual-like structure can still achieve some degree of risk pooling for the idiosyncratic risk, and the systematic risk (whether from stage 1 or stage 2 of the lottery) can be shared across the population at risk by devices such as *ex post* dividends, assessments or taxes.²

Kunreuther and Pauly present a case study of insurance of catastrophic risk, which may be known or unknown. Losses from natural hazards such as hurricanes and earthquakes, although correlated, fit more comfortably into the **K** category. Intensive modeling of these losses is undertaken both by insurers and by boutique modeling firms. Nevertheless, the correlation structure does pose problems which calls for some mutualization of risk. In contrast to natural hazard risk, terrorism risk falls more squarely in the **u** category. Here past data do not provide a good basis for estimating future risk. Indeed, as noted by Kunreuther and Pauly, terrorists will deliberately try to be as unpredictable as possible, thus frustrating any counter-terrorism measures.

Kunreuther and Pauly propose a policy response that applies to catastrophic risks ranging from **K**-type natural catastrophes to **u**-type terrorist attacks. They argue that the same type of layered public-private program applies to both risks types (though presumably the details such as attachment points could differ). The idea is to provide potent incentives for mitigation of losses, while still using the risk-bearing capacity of the private insurance and capital markets to provide the greatest possible diversification benefit. The lowest layer would be self-insured, followed by a layer of private insurance contracts with risk-based premiums. The third layer would be reinsured or otherwise spread through insurance linked securities. The final layer would allocate highest level of losses on an economy-wide basis; though the details are not specified, this could comprise multi-state risk pools or federal government intervention as the reinsurer of last resort.

Ex Post Wealth Redistribution for U

The Borch-like arguments for mutualizing risk gather force as we move from **K** to **u**. As we move to **U** it becomes impossible to specify, let alone price, risks that could be transferred by standard contractual devices and it is correspondingly difficult to provide incentives to mitigate risks that cannot be identified. However, we do of course know that surprises, of a yet unimaginable nature, could arise in the future and we can anticipate that we might react in predictable ways.

² One can even imagine two-stage *ex post* assessment. After more information on probabilities is revealed, assessments are made to determine the premiums for stage 2. Then stage 2 assessments are made to cover the difference between the overall realized losses and funds raised from stage 1 assessments.

When major catastrophes occur, it is not uncommon for governments to provide resources to the victims. For example, the two most notable catastrophes to hit the United States in recent years were 9/11 and Hurricane Katrina. In both cases, the federal government was generous in providing compensation. However, it does seem that there are “inconsistencies”. For example, the Victim’s compensation allocated a total of \$7 billion, which amounted to an average payment of \$1.8 million and compensation was paid to 93% of the families. For the earlier terrorism event in Oklahoma City, there was no such fund. Although some aid may be available through FEMA for lesser disasters, there appears to be a political imperative to be especially generous when the scale of the disaster exceeds some threshold of saliency. It would appear that bailouts of failed financial institutions also must meet a scale criterion, although this appears to be highly dependent on the perceived fragility of markets when the failure occurs. It is noteworthy that neither Northern Rock nor Bear Stearns, both of which were bailed out, were counted among the large, complex financial institutions that the International Monetary Fund has identified as critical to the functioning of the international financial system.

The efficiency and equity arguments for *ex post* government compensation for catastrophic losses are several. For example, Borch’s principle of mutuality would have the optimal apportionment of risk to be such that all idiosyncratic risk of individuals would be fully insured, but all social risk would be shared. This indeed is the same theory as the capital asset pricing model, in which all people have scaled shares in the social wealth (i.e., the market portfolio). There may be several ways to achieve such a redistribution, including an *ex post* allocation that redistributes taxpayers’ funds to victims, thus equalizing the impact across the population (Kunreuther and Pauly). On the other hand, the incentives from such bailouts can be perverse. Banks may be less inclined to practice financial discipline and their customers less inclined to monitor them. Individuals may be less inclined to buy insurance and unwilling to invest in mitigation measures. On this later argument, one might think it would behoove governments to pre-commit not to bail-out in the future. However, such pre-commitments may not be enforceable given the political pressures brought to bear when the event strikes. This is the well-known “Samaritan’s dilemma.”

This tension over *ex post* allocations for risky events is real. Since risk, in our terms, relates to events that are “known” then appropriate risk management strategies can be adopted such as hedging and mitigation. *Ex post* allocations will undermine incentives for implementing these risk management measures, thus creating an efficiency loss. However, when we switch the setting from K through u to U , the issues re-align. Insofar as it is difficult to mitigate unknown or unknowable risks, then there can be little efficiency loss. We cannot really plan for complete surprises. However, as some of the writers recognize, e.g. Clive Granger, we can nevertheless spread the impact of these complete surprises by compensating the victims from public funds.

The equity case for bailout from hitherto unknown or unknowable events can be interpreted in a Rawlsian (1971) framework in which the principles of justice are chosen behind a “veil of ignorance.” In voting on a policy for wealth distribution, people will be influenced by their current endowment. However, if one were to conceal endowment

behind a Rawlsian veil of ignorance, then policy preferences might change. This allows us to separate distributional preferences from wealth endowments. Now consider people's *ex ante* preferences for the desired *ex post* wealth endowment after some major risky event (i.e., whether people are likely to vote for future bailouts). Each person can anticipate that their *ex post* wealth W_1 will be their starting wealth W_0 plus a random shock μ ; that is, $W_1 = W_0 + \mu$, where the distributional properties of μ are known. In other words, we know in a stochastic sense what our *ex post* wealth endowment will be, and this will be influential in any vote about future bailouts.

The more we venture from the known towards the unknowable, (i.e., as μ in the above equation becomes a u or U risk) it becomes more difficult to specify the relationship between W_1 and W_0 . Indeed, the *ex post* distribution for the unknowable is naturally concealed behind a veil of ignorance. If it were true that we have an innate preference for an egalitarian distribution (when the contaminating influences of our endowment are removed), we are left with a strong equity arguments for *ex post* allocations following the occurrence of unknowable events.

C. Create Flexible and Adaptive Structures

In a comment that spans both parts of our taxonomy, which grounds KuU in knowledge-as-measurement and knowledge-as-theory, Paul Kleindorfer notes that the balance between aleatory risk (dependent on an uncertain outcome) and epistemic risk (related to an imperfect understanding or perception) changes as we move from K through u to U . In particular, in a K setting, risk management will tend to focus on risk mitigation or transfer, but as we move towards U , risk management will stress adaptation and flexibility. Where risk is K nown, it is likely that a market for trading that risk will emerge, as with commodity and financial derivatives and insurance contracts. Alternatively, if the process that generates the risk is understood, then risk can be mitigated *ex ante*. For example, as Kunreuther and Pauly point out, much exposure to known catastrophe risk can be mitigated by choosing location and by constructing wind- or earthquake-robust structures. K conditions imply that risk can be priced and that we will be able to make informed capital budgeting decisions.

Transfer and *ex ante* mitigation become more difficult in the case of u and U , in which risk management emphasis shifts to adaptation and flexibility, and to robustness and crisis response. These strategies are both *ex ante* and *ex post*. The knowledge that unknown and unknowable losses might strike suggests caution with regard to capital and investment decisions; for example, extra capital provides some buffer against the unknown, and required rates of return on investment might be more conservatively chosen in a uU environment.

Heightened awareness of the existence of uU risks is valuable, despite the fact that, by definition, such risks cannot be identified *ex ante*. We have mentioned the value of holding more capital as a buffer against such shocks, and we can think of this as simply a change in financial leverage (a change in the ratio of variable to fixed financing costs). Similarly, a change in operating leverage, the ratio of variable to fixed factor costs, can

make any firm more robust to shocks, whether of the K or u or U variety. For example, Microsoft has operated with a high ratio of contract as opposed to payroll labor, which it can more easily reduce in unforeseen bad times.

Other examples of strategies that create organizational flexibility and adaptability are given by several contributors. For example, Scott notes that if the compensation of CEO's and other top managers is based on long-term wealth maximization, then it will motivate managers to manage crises in a way that protects shareholders. But this need not be purely after-the-fact improvisation. Indeed, properly structured compensation may motivate managers to anticipate how they should respond to crises and invest in crisis response capability. This may prove useful even if the anticipated crisis does not occur. Precautionary measures for dealing with Y2K, which proved uneventful, are nevertheless widely credited with enhancing the resilience of the financial system after the 9/11 terrorist attack on New York's financial center.

Kleindorfer also stresses crisis management, emphasizing the role of a crisis management team in creating response capability. Although crises may have unique and perhaps unanticipated causes (either u or U), the responses called for are often similar, and a well prepared crisis management team can often limit the damage. For example, large shocks create uncertainty, and clearly articulated responses can reassure customers and investors, ensure that supply chains are secured, etc. But well-designed responses can even snatch victory from the jaws of defeat. For example, after the cyanide scare with its Tylenol product, Johnson and Johnson withdrew and then redesigned its product and packaging, setting a standard that led its competition and secured competitive advantage.

Sound management of crises can not only mitigate their impact, but also generate new knowledge. Indeed, and perhaps ironically, crises are sometimes portals that take us from U and u toward K . For example, Hurricane Andrew in 1992, the Asian currency crisis in 1997, and 9/11 in 2001, all spurred new approaches to modeling extreme event risk. Hurricane Andrew led to an explosion of interest in catastrophe modeling and a consequent refining of modeling methodology. The Asian crisis led to a new interest in the properties of tail risk (fat tails and tail correlations), that have been incorporated in new generations of models. Finally, 9/11 led to the development of game theoretic catastrophe modeling techniques.

D. Use Incentives to Promote Desired Outcomes

Risk management strategies must confront the issue of *incentives*, thwarting moral hazard by coaxing rational and purposeful economic agents to act desirably (whatever that might mean in a particular situation). We take a broad interpretation of "designing strategies," and here we discuss three: Designing organizational/governance arrangements, designing contracts, and designing investment vehicles.

Organizations and Relationships: Principal/Agent Concerns in Corporate Governance

K risks can be identified and probabilities assigned to them. If knowledge is symmetric and actions are commonly observable by all parties, then simple contracts can be written in which actions are specified contingent on states of nature. In this simple world, there are no moral hazard and adverse selection problems. For example, insurance or loan contracts can be written, and banks and insurers would know the quality of each customer and price accordingly. Private wealth transfer caused by inefficient actions (the insured under-investing in loss mitigation or borrowers taking on excessive risk) would be avoided because they were excluded by contractual conditions, and institutions would be able to monitor accordingly.

In principle, special purpose vehicles (SPV) work that way. They may be organized as trusts or limited liability companies and are set-up for a specific, limited purpose, often to facilitate securitizations. In the securitization process, the SPV buys pools of assets and issues debt to be repaid by cash flows from that pool of assets in a carefully specified way. The SPV is tightly bound by a set of contractual obligations that ensure its activities are perfectly transparent and essentially predetermined at its inception. SPVs tend to be thinly capitalized, lack independent management or employees and have all administrative functions performed by a trustee who receives and distributes cash according to detailed contracts. SPVs are designed to be anchored firmly in the domain of **K** in order to fund assets more cheaply than they could be funded on the balance sheets of more opaque, actively managed institutions. The turmoil in the subprime market during the summer of 2007, however, revealed that the claims on some of these institutions were much less transparent than assumed and that investors (and the ratings agencies they relied upon) were in fact operating in a world of **u** rather than **K**. This led to a repricing of risk and disruption of markets that spread well beyond the market for subprime-related securitizations.

Governance of most firms is not as straightforward as with SPVs. Partly the problem is one of complexity and division of labor. There are numerous potential states of nature facing firms, and contracts anticipating the appropriate response of managers to each possible state would be impossibly cumbersome. Moreover, envisioning shareholders (or their agents) writing such contracts presupposes that they already have the managerial skills they are seeking to employ. Indeed the reason for employing managers is that they alone know the appropriate responses.

The division of labor issue can be cast as an information problem. Managers have much better knowledge than shareholders of how to deal with managerial opportunities and crises. In this light, the issue is not whether knowledge and skills can be acquired, but how they are distributed across stakeholders. Ken Scott digs much deeper into the informational aspects of corporate governance and explores how governance mechanisms may be designed when risks are unknown, **u**, to any party, or indeed unknowable, **U**.

Scott's particular focus is the risk management aspect of governance, and he starts by contrasting the risk preferences of the major stakeholders. While diversified shareholders seek to maximize the value reflecting both the upside and downside of the distribution; relatively undiversified managers are probably more risk averse and compensation is

often designed to enhance their preference for risk. Government and regulators, presumably protecting the interests of consumers, also are more interested in downside risk, particularly the prospect of contagion, or systemic risk. Their attention is focused on how to avoid the prospect that firms will incur unsustainable losses.

For **K** risks, shareholder and societal interests are promoted by risk-neutral decision making and the governance problem is in large part one of furthering this objective through appropriate compensation design. But the fine tuning of compensation to provide risk-neutral incentives and correct reporting is not an easy task. Thus Scott stresses the importance of establishing a “risk culture” within the firm. This can start with board members who demand that top managers articulate their risk management strategy, and it can flow down to division and project managers who conduct (marginal)³ risk analysis. Coordination can be addressed by the appointment of a chief risk officer.

For **u**, part of the governance problem is mainly to encourage the acquisition of more information; to convert **u** into **K** by investing in information. Another part of the issue is to design internal controls and board oversight of management actions. Scott makes the important point that these efforts might be more effective if management were unable to keep ill-gotten gains (resulting from manipulated earnings) in their bonuses.

The externalities caused by bank failure are classified as **u**. Control of this risk can be influenced by contract design, and Scott points to the perverse case that arises when derivative counterparties are given a favored position when banks go into receivership, which diminishes their incentives to monitor banks and price risk appropriately. For their part, regulators have addressed the bank failure risk in detail through the Basel I and II requirements, which are designed to provide more information on risk and establish appropriate levels of regulatory capital. But more interesting is the shift in decision-making authority that occurs as a bank’s capital declines. The “prompt corrective action” measures in the US law permit the downside risk preferences of regulators to trump the risk neutral perspective shareholders as the bank’s capital declines.

For **U** risks, Scott proposes financial flexibility and managerial incentives linked to the long term survival of the bank. Thus, as unknowns appear, managers will be rewarded for finding responses that protect the interests of the other stakeholders. The very nature of this risk implies that one cannot estimate with accuracy the additional capital needed to cushion against unforeseeable failures. Nevertheless, additional capital margins will reduce this prospect and an ordinal ranking of institutions is feasible.

The process by which knowledge is acquired (either facts or understanding) creates not only opportunities to use the new knowledge but also institutional stresses. Stresses occur because institutions are designed around increasingly outdated knowledge, with changing knowledge shared asymmetrically by the various stakeholders. In the past two decades or so, capital markets have undergone considerable changes. These include changes in our conceptual model of how markets work, as well as changes in the volume of data. The evolution in asset pricing, from the one-factor capital asset pricing model

³ Marginal risk analysis ascertains the incremental contribution of each activity to the total risk of the firm.

though to more recent multi-factor models, as well as the revolution in derivative pricing together with advances in corporate financial theory, have changed the way investor and management decisions are made. These theoretical innovations have created a demand for new data to verify and calibrate these new models. This push for data has been exacerbated by a phenomenal growth in computing power. Enhanced understanding of the underlying economic mechanisms and better data potentially allow all stakeholders to make decisions that better create value.

Accompanying the revolution in financial theory and explosion in data has been a market enhancement that Bravler and Borge label *capital market intensity*. More information and better understanding allow investors to monitor changes in a firm's fortune quickly and to act accordingly. Passive investors may simply buy or sell. But an increasing tendency to shareholder activism, especially in hedge funds, has led to investor involvement in corporate decision making. This is exercised by applying direct pressure to management, influencing board composition, removing management, etc. In this way, investors can exert direct influence over management to seek preferred risk-reward profiles. At the same time, of course, the innovations bestow better tools on management to attend to investor needs. In particular, the sophisticated tools of financial engineering and the bewildering array of financial instruments permit almost unlimited flexibility to management to redesign its risk-reward profile.

In Bravler and Borge's view, increased capital market intensity challenges the traditional principal-agent based model of corporate governance. The traditional model assumes that managers have a comparative advantage in both information and decision skills over investors, but investors induce managers to create value by means of incentive compatible employment contracts. In the new world of capital intensity, the comparative advantage between (at least some) investors and managers is largely removed. Braver and Borge see a new model which is analogous to a two-sided market structure. The CFO acts as an intermediary between the management and investors: "The CFO is the agent of the company in the capital markets and the agent of capital market discipline inside the company". In fact, we would probably suggest that the CFO is still properly regarded as the agent of the company. However, this does not diminish the power of the Bravler-Borge observation that the CFO's role needs to be redefined to refocus corporate attention on value creation for investors and to use the potent strategies and instruments now available to achieve this end. If the CFO falls down on this task, increasingly activist investors may simply do it for themselves.

Contracts: Intentional Incompleteness and "Hold-Up"

Things that are unknown now may become known as events unfold. 9/11 informed us of a different form and magnitude of terrorism. Recent financial crises, notably the Asian crisis and the subprime crisis, informed us of hitherto unsuspected correlations in tail risk that have now deepened our understanding of systemic risk. However, not only new events, but also new theory, can shift us from U toward K . For example, assets that might appear mis-priced under a simple single-factor pricing model may appear well-priced under a multi-factor model. Unfortunately, retrogression also occurs. Statistical

relationships which have proven stable over many years may suddenly break down. Institutional structures which were well-understood may prove to have hidden flaws. Policies that seem reliable in normal times may fail to work in a crisis.

If we cannot anticipate events, or do not understand their consequences, it becomes difficult to write effective contracts. For example, the insurance industry was recently surprised by several new classes of claims that it had not suspected and therefore had neither written them into coverage or explicitly excluded them. These included toxic mold damage to buildings and to the health of their residents, as well as the new forms of terrorism that blur the distinction between traditional terrorism and actual warfare. Other examples include innovative legal rulings that have substantially changed coverage from what seems to have been written into policies; examples include the (sometime) removal of the distinction between flood and wind coverage in post-Katrina claims and the earlier reinterpretations coverage for “sudden and unexpected” liabilities to include “gradual and expected.”

Insurance contracts are usually written for named perils, or written to include a broad class of perils insofar as they are not specifically excluded by contract language. Either way, the contract defines what is covered and what is not. If the covered perils are “known” in our terminology, a price can be set relative to the (known) expected loss and other distributional parameters that indicate the cost of capital. Even if events are unknown in the sense that they can be identified but cannot be assigned probabilities, contracts can still be written, although the setting of premiums becomes a challenge. But when events cannot even be specified, contracts cannot easily be imagined.

Doherty and Muermann ask whether risks that are indeed unknowable can be effectively transferred to insurers. Using incomplete contract theory, they argue that such risks can be, and are, allocated to insurers. When writing through independent agents and brokers, insurers vest the intermediary with considerable “hold-up” power. Agents and brokers can move their books of business and may do so if they believe this serves the interests of their policyholders. Moreover, they argue that this hold-up power is used to extend insurance coverage to include some non-specified events. If a hitherto unknown event arises, they broker can decide whether it is one that can, and should be, insurable going forward (i.e., now that it has graduated from U to u or K). If so, the broker might use its leverage to bargain with the insurer for a settlement for its client. Indeed, such *ex post* bargaining may even be anticipated when contracts are written, and premiums adjusted upwards accordingly. In this way, brokered markets can provide an orderly market in which unknown events can be insured despite the fact that the coverage is not formally specified in the contract.

Incomplete contracts may indeed be a common device for coping with the unknown. For example, employment contracts for CEO’s and other top executives are incomplete insofar as they do not anticipate detailed scenarios and prescribe specific managerial responses. Instead, they rely on alignment of the interest of the CEO and shareholders through compensation design and allocate considerable discretion to the CEO to respond to events drawn from the whole KuU space. In this way the CEO’s skill is given

considerable scope to respond quickly to new information, as emphasized for example in Ken Scott's paper.

Investment Vehicles: Riding "Sidecar" with Better-Informed Agents

Richard Zeckhauser looks at investment in, uU events, for which there is scant knowledge. The usual training in financial economics avails us little in dealing with uU events. Consequently, markets are thin and potentially enormous excess returns are available to those with resources and talents to venture into these little explored places. It helps to have billions to invest, complementary skills, steady nerve and freedom from blame when things don't work out (as often happens). Warren Buffett and his ilk can prosper in this realm, but what about the rest of us? Can we also make sensible and profitable forays into this compelling and intimidating territory? Richard Zeckhauser raises this question in an unorthodox essay that draws both on his own experience (as a "sidecar investor") and on his deep insights in areas of economics not usually considered relevant to investors.

Consider the investor with money, steady nerve and complementary skills. He or she may well be willing to make a speculative investment, accepting large risk for extraordinary expected returns. Can investors with lesser skills and resources attach themselves to this powerful motor-bike and go along for the ride as a "sidecar"? It is certainly dangerous territory. The risks for the biker with the wherewithal to handle it can be enormous, and risks may be relatively greater for those riding sidecar. Moreover, dealing with those who are better informed exposes us to adverse selection risk, and this must be balanced against the absolute advantage from their superior resources and skills. Yet if we understand this trade off, then there are opportunities for sidecar investments. Robert Edelstein notes that real estate syndicates are often structured that way. Investors have the money and the developer has the complementary skill and experience. In the end, however, the developer often has the money and the investors have the experience.

A little game theory and behavioral economics helps balance the adverse selection against the absolute advantage, and Zeckhauser gives some instructive examples ranging from Russian oil investments to Warren Buffet's reinsurance ventures. In writing the first draft of his paper – around the time that Katrina was devastating New Orleans – he presaged the development of an active sidecar market in insurance, as described by Kunreuther and Pauly.

Disasters such as 9/11 and Katrina both diminish insurance capacity and usually enhance insurance demand, thus leading to excess demand, which is felt in a hardening of the insurance market. This hard market is felt most acutely in reinsurance, where post-loss supply is especially scarce and prices soar. Often excess demand is fueled by a shift along the continuum from K to u to U ; for example, 9/11 created major uncertainty about future terror risk, and Katrina fed our fears on the unknowns of global warming. The hardened reinsurance market, together with enhanced uncertainty of the future risk, create just those conditions which Zeckhauser considers ripe for very high returns. Reinsurers possess the complementary skills (if indeed anyone does), but hedge funds have the funds

and tolerance of ambiguity to partake. As a result sidecar structures have bloomed – usually with hedge fund (and some other) investors taking quota share investment on the same terms as reinsurance contracts. These differ from equity investment in the reinsurer in that they cover only specified risks and usually for a short time frame.

The example of reinsurance-based sidecars represents an example where absolute advantage probably trumps adverse selection. While there may be some adverse selection in the original contractual relationship whereby the reinsurer “insures the primary insurer,” there is unlikely to be much additional adverse selection in the derived relationship between the reinsurer and the sidecar investors. Thus, the value creation is driven by absolute advantage, and sidecar investors can share in this added value.

E. Use Financial Policy to Limit Vulnerability to Shocks *Ex Ante* or Mitigate the Consequences *Ex Post*

Financial policy becomes most relevant when a shock that was unknown or unknowable shifts the financial system from the domain of the known into the unknown. Financial policy-makers are charged with limiting the vulnerability of the financial system to such shocks and mitigating the consequences of these shocks once they occur. Financial policy-makers aim to promote monetary and financial stability. In practice, virtually every aspect of financial policy is subject to uncertainty. For example, how precisely should these objectives be defined? With regard to monetary policy, what amount of inflation is consistent with achieving stable, sustainable growth? What measure of inflation is appropriate? Is it feasible, both technically and politically, for the monetary authorities to prevent asset bubbles during periods of low and stable inflation? Jacob Frenkel (Thornhill and Michaels, 2008, p. 4), former Governor of the Bank of Israel, has expressed doubt about whether the monetary authorities know enough to deflate bubbles before they become dangerous. He asserts that the real choice is “Which system do you want: one in which the [monetary authority] pricks three bubbles out of five or five out of three bubbles? Because we know for sure that it will not be able to solve four out of four.”

With regard to prudential policy, the primary goal of financial stability must be to protect the functioning of the financial system in providing payments services and facilitating the efficient allocation of resources over time and across space. This may be threatened by a loss of confidence in key financial markets or institutions. But how safe should financial institutions be? Should all failures be prevented? Would the required restrictions on risk-taking by financial institutions reduce the efficiency of financial intermediation and reduce investment? Would this deprive the economy of the dynamic benefits of creative destruction? But if financial institutions should not be required to be perfectly safe, what degree of safety should the prudential authorities try to achieve?

What tools should be used to achieve these objectives? And what governance structure is most likely to motivate policy-makers to act in the public interest? Public-sector compensation contracts are much more highly constrained than compensation contracts for senior executives in financial services firms. More fundamentally, when

objectives are not crisply defined, it is difficult to establish and enforce accountability. Blame avoidance is, by default, the primary objective of most bureaucrats.

Although the prudential supervisory authorities have enormous, if ill-defined, responsibility, they have relatively little power to constrain risk-taking by profitable institutions that they believe to have excessive exposures to uncertain shocks. In order to guard against the arbitrary use of regulatory and supervisory power, most countries subject disciplinary decisions by officials to some sort of judicial or administrative review. To discipline a bank, a supervisor must not only know that a bank is taking excessive risk, but also be able to prove it to the satisfaction of the reviewing body – perhaps beyond a reasonable doubt. This leads to a natural tendency to delay disciplinary measures until much of the damage from excessive risk-taking has already been done. (As Kenneth Scott notes, the Prompt Corrective Action measures adopted in the US are intended to constrain this tendency to forbear in the enforcement of capital regulations by removing a degree of supervisory discretion.) It also leads officials to react mainly to what has already happened (and is, therefore, objectively verifiable) rather than to act on the basis of expectations about what may happen (which are inherently disputable). In Charles Goodhart’s refinement of the *KuU* framework in which *K* is partitioned into actual past data and expected values, supervisors generally react to past actual losses rather than mean expected losses, much less the unexpected losses, even when the governing probability distribution is believed to be known. Alan Greenspan (2008, p. 9), former Chairman of the Board of Governors of the Federal Reserve System, has expressed doubt about whether regulators know enough to act preemptively: “Regulators, to be effective, have to be forward-looking to anticipate the next financial malfunction. This has not proved feasible. Regulators confronting real time uncertainty have rarely, if ever, been able to achieve the level of future clarity required to act preemptively.”

Regulating with Imperfect Information

Information issues present a fundamental challenge to supervisory authorities who must oversee the solvency of regulated financial institutions. Neither actual past data, nor expected values can be relied upon in times of crisis when difficult supervisory decisions must be made. Bank accounting has traditionally been a mix of historical cost accounting, accrual accounting and mark to market accounting. This has sometimes undermined incentives for hedging risks by valuing a risky position and the offsetting hedge differently, thereby increasing the volatility of earnings, even though risk has been reduced. Many doubt this mix of standards conveys a true and fair account of the current position of a financial institution. New financial accounting standards require firms to classify assets in three different categories: (1) assets that can be marked to market based on quoted prices in active markets for identical instruments; (2) assets that are marked to matrix, based on observable market data; and (3) assets that are marked to model, based on judgment regarding how the market would price such assets if they were traded in active markets. This third category presents significant difficulties for regulators, who face a severe asymmetric information problem vis-à-vis the regulated institution. How can the regulatory authorities comfortably rely on the estimated values of category three

assets? Opinions of auditors and ratings agencies may help the authorities avoid blame, but the key question, as Goodhart notes, is “Who has legal liability if the values are wrong?”

Part of the problem, as noted by Stewart Myers in a workshop that preceded this volume, is that financial theory offers only two kinds of tools for valuing assets that are not traded in active markets: (1) the present value of discounted cash flows, which works well in a world of K , where cash flows can be predicted and risks estimated; and (2) real option theory, which works well only if you can write a decision tree that captures most of the key uncertainties and decision points in the future. Fundamental values rest on relatively shaky foundations, and a shock may shift a price from the realm of K to that of u .

Even category 1 assets may present problems in a crisis. Setting aside the issue of asset price bubbles, market values can be relied upon so long as assets are traded in broad, deep resilient markets. In such markets, however, assets tend to be priced on the basis of comparisons to their own past prices or to the prices of comparable assets. When a shock undermines confidence in these relative values and causes losses, traders tend to withdraw from markets until they regain confidence in their valuation models. Such shocks move prices from K to u . Concerns may arise about counterparties who may have had excessive exposures to the shock and markets become thin. A flight to quality may occur and liquidity will be restored only when confidence in valuation models and counterparties is restored.

Crisis Prevention

Most policymakers would agree with Don Kohn that it is better to prevent such crises than to try to manage and mitigate them once they have occurred. However, crisis prevention is an enormous burden, which falls mainly on the shoulders of the prudential authorities. Prudential regulation attempts to establish rules for the sound operation of financial institutions and critical elements of the financial infrastructure such as clearing and settlement arrangements. Ideally, prudential policymakers should be looking beyond K to anticipate emerging sources of systemic vulnerability in order to calibrate appropriate prudential policies. In the dynamic world of modern finance this requires trying to understand how changing institutions, products, markets and trading strategies create vulnerabilities to new kinds of shocks and new channels of contagion. But K cannot be neglected. Institutions still fail in familiar ways by taking excessive concentrations of credit risk or by imprudently borrowing short and lending long as recent experience with Northern Rock and dozens of Structured Investment Vehicles have shown.

Prudential supervisory authorities confront a number of trade-offs that must be made on uncertain terms. How safe should banks be? Goodhart notes that it is relatively easy to establish a set of penalties that would make the banking system perfectly safe, but largely irrelevant in intermediating between savers and investors. Scott argues that a central feature of corporate governance is aligning the risk neutral preferences of well-

diversified shareholders with risk-averse managers. This calculus is unlikely to take account of the systemic costs of an institution's failure and so the prudential authorities will presumably prefer a higher degree of safety, but how much higher?

How much competition is desirable? Competition is generally viewed as a positive feature of the financial system. It stimulates innovation and lowers the cost of financial services. But, it also reduces the charter values of incumbent banks and may lead to increased risk taking. Goodhart notes that, over time, the official view regarding competition has swung from one extreme to another. During the Depression, the authorities tended to regard competition as a source of instability and implemented a number of reforms to constrain competition. More recently, the dominant trend has been liberalization of competition, although the current crisis in credit markets may cause a reversal.

Should financial innovation be encouraged? Securitization has facilitated diversification of risk, reduced costs and liberated borrowers from dependence on particular lenders, but the subprime crisis has shown that it can also undermine credit standards and enable banks to achieve higher leverage by evading capital requirements. Derivatives have enabled financial institutions to partition and manage risks much more efficiently, but they can also be used to take enormous, highly leveraged risks. The growing sophistication of risk management techniques has enabled institutions to push out the boundaries of the known, but the very complexity of these techniques presents a challenge in the event of a crisis because it is very difficult for the authorities to comprehend the full range of positions and how they are managed. As Ralph Gomory warned in his essay on *KuU*, “[A]s the artifacts of science and engineering grow ever larger and more complex, they may themselves become unpredictable.”

The supervisory authorities have a number of tools. These include licensing requirements, restrictions on certain kinds of activity believed to be excessively risky, liquidity requirements, capital requirements and disclosure requirements. The authorities may also try to identify and encourage the widespread adoption of best practices in risk management, in effect urging the private sector to convert *u* into *K*.

By far the most ambitious effort at prudential regulation has been the development and of the Basel II standards for capital adequacy. Andrew Kuriztkes and Til Schuermann provide a framework for analyzing *KuU* in bank risk taking and show how the Basel II capital requirements correspond to this framework. They argue that a risk can be classified as *K* to the extent that it can be identified and quantified *ex ante*. They observe that the ability to estimate downside tail risks at a high level of confidence has enabled financial institutions to develop the concept of economic capital, the amount of capital needed to protect against earnings volatility at a prescribed level of confidence, usually set equal to the default rate associated with the financial institution's target debt rating. Economic capital has become the common denominator for measuring and aggregating risks in the financial services industry. But it is firmly rooted in the known and does not transplant readily to the unknown.

Kuritzkes and Schuermann classify a risk as *u* to the extent it can be identified *ex ante*, but not meaningfully quantified. A risk is classified as *U* if the existence of the risk is not predictable, much less quantifiable *ex ante*. Since these risks can't be quantified, they can't be managed. They can, however, sometimes be transferred. Kuritzkes and Schuermann employ this framework to analyze how *KuU* varies by risk type based on the richness and granularity of the data available to estimate each kind of risk and conclude that *K* decreases and *u* and *U* increase moving along a spectrum from market risk to credit risk, to asset/liability management risk to operational risk to business risk. In addition they analyze bank holding company data on earnings volatility to estimate the total amount of risk in the US banking system and to allocate this total risk across risk types. Financial risks – market risk, credit risk and asset/liability management risk account for 70% of the volatility of earnings. Non-financial risks – operational risk and business risk – account for the remaining 30%. Within financial risks, market risk is 6%, credit risk is 46% and asset-liability management risk is 18% of total volatility. Unfortunately, we know the most about the least important type of risk (although as the authors note this may be due in part to the success banks have achieved in measuring and managing it). Within non-financial risks, operational risk accounts for 12% and business risk accounts for 18% of the volatility of earnings. Regrettably, we know the least about business risk which has been three times larger than market risk.

Bank regulators began to take note of the evolving concept of economic capital when they expanded the original Basel Accord on Capital Adequacy to take account of market risk. The 1996 Market Risk Amendment provided an entirely new approach to setting capital requirements that relied on the way that leading banks were measuring and managing this risk. The original Accord set capital requirements roughly in line with expected losses. The concept of economic capital made clear that the role of capital should be to absorb unexpected losses, with reserves established to absorb expected losses. And so, instead of requiring banks to allocate their positions to crude risk buckets, or applying mechanical asset price haircuts to positions in an attempt to approximate risks, the regulatory authorities provided the opportunity for qualifying banks to rely on the supervised use of their internal models to determine their capital charges for exposure to market risk.

The internal models approach was expected to deliver several benefits. First, it would reduce or eliminate incentives for regulatory capital arbitrage because the capital charge would reflect the bank's own estimate of risk. Second, it would reward diversification to the extent that a bank's internal models captured correlations across risk positions. Third, it would deal more flexibly with financial innovations, incorporating them in the regulatory framework as soon as they were incorporated in the bank's own risk management models. Fourth, it would provide banks with an incentive to improve their risk management processes and procedures in order to qualify for the internal models approach. And fifth, compliance costs would be reduced to the extent that the business was regulated in the same way that it was managed. By and large, the internal models approach for market risk has proven to be highly successful, even when it was severely tested by the extreme market disruption in 1997, 1998, and 2001, which is consistent with the view of Kuritzkes and Schuermann that market risk is largely *K*. This

success, in combination with the progress made in modeling credit risk, led to calls from industry to revise the original Basel Accord to incorporate an internal models approach to capital regulation of credit risk.

Basel II attempts to extend this new approach to setting capital requirements to credit risk and operational risk. Although the supervisory authorities were convinced that credit scoring models had significantly expanded the amount of credit risk that could be regarded as falling in the domain of the known, they were skeptical that internal models of credit risk were as reliable and verifiable as models of market risk. While some kinds of credit risk like retail lending have rich and granular data sets comparable to market risk, other kinds of credit risk are less amenable to empirical analysis because data are sparse relative to past credit cycles and distinctly non-granular. In the end, the regulators rejected the supervised use of internal models, but permitted qualifying banks to use their internal model inputs – estimates of probability of default, loss given default, exposure at default and duration of exposure -- as inputs in the regulatory model that would determine capital requirements. These Pillar 1 capital requirements recognized the analytical and empirical advances banks had made in expanding the extent to which credit risk can be regarded as known.

Moving further to the right in the *KuU* spectrum, the decision to establish a Pillar 1 capital charge for operational risk has been much more controversial. In this instance the regulators were not simply adopting industry best practice as in the case of market risk and credit risk. They were attempting to advance best practice by requiring greater investment in measuring and managing operational risk. Moving operational risk into the domain of the known presents major challenges. Until quite recently, the industry lacked even a common definition of operational risk. Moreover, it is difficult to quantify and disaggregate, data are sparse and theory is weak.

Because Basel II is an agreement negotiated among the members of the Basel Committee on Banking Supervision it reflects a number of political compromises that undermine its aspirations for technical precision. This is most evident in the definition of regulatory capital which is based on accounting values and includes a number of items that do not reflect an institution's capacity to bear unexpected loss. This undercuts the link to best practices in risk management the logic of the approach.

Pillar 1 capital charges are intended to deal with known risks. Pillar 2, the supervisory review process, is intended to deal with unknown risks that can be identified, but are not sufficiently well-quantified to establish Pillar 1 capital charges. Presumably, as theoretical and empirical advances succeed in moving some of these risks into the domain of the known, Pillar 1 capital charges will be established for them as well. In view of the analysis by Kuritzkes and Schuermann, it is surprising that asset-liability management risk is treated under Pillar 2, while operational risk is treated under Pillar 1. Although liquidity is inherently difficult to measure because it has at least three dimensions – price, time and size – interest rate risk, another important aspect of asset-liability management risk, is much more easily quantified than operational risk and it has been a much more important source of volatility in bank earnings than operational risk.

Kuritzkes and Schuermann raise the question of whether regulatory and industry resources might have been more usefully directed to standardizing the approach for characterizing and measuring asset/liability risk.

Benoit Mandelbrot and Nassim Taleb warn that there is much more u in K than is commonly acknowledged. The past is never a perfect predictor of the future. New factors may become important and relationships estimated in times of normal market functioning tend to break down at times of market stress. What we thought was mild randomness often proves to be wild randomness – or at least more often than it should if it were governed by a Gaussian distribution. In Will Roger’s phrase, one the key risks may be what we think we know “that just ain’t so.”

The principal tools of supervisory analysis in the domain of the unknown are stress testing and scenario analysis. Stress-testing requires economic judgment to formulate and calibrate scenarios that expose potential vulnerabilities. It requires a careful consideration of which relationships will continue to hold and which relationships will break down in time of stress. Mandelbrot and Taleb caution that traditional stress testing, which relies on selecting a number of worst-case scenarios from past data, may be seriously misleading because it implicitly assumes that a fluctuation of this magnitude would be the worst that should be expected. They note that crashes happen without antecedents. Before the crash of 1987, for example, stress testing would not have included a 22% drop in share prices within a single day. They note that just ten trading days account for 63% of the returns on the stock market over the past 50 years. In their view fractal methods should be used to extrapolate multiple projected scenarios that would enable risk managers and prudential supervisors to evaluate the robustness of a portfolio over an entire spectrum of extreme risks.

Goodhart emphasizes a different concern regarding stress testing and scenario analysis. What may matter most in crises are interactive effects that occur when many institutions attempt to adjust their portfolios in the same way at the same time. These are critical to understanding an institution’s vulnerability in a crisis, but are omitted from most scenarios.

Stress-testing and the simulation of crises may be of value even if such crises never occur. The data necessary to simulate a crisis may prove useful in monitoring vulnerability and a careful consideration of the consequences of such a crisis may lead to changes in strategy and/or risk management. Crises seldom unfold according to the anticipated scenario, but strategies for responding to one kind of shock may prove useful when a different kind of shock occurs. For example, evacuation procedures that Morgan Stanley established after the bombing of the World Trade Center in 1993 enabled the firm to safeguard all of their employees in the much more severe terrorist attack on September 11, 2001.

The key element of regulatory discipline under Pillar 2, however, is the ability of the prudential supervisor to impose an additional capital charge on an institution if they are uncomfortable with the results of its stress tests. This places supervisors in the role of

imposing discipline on an institution thought to be vulnerable to a shock of unknown probability even though they are less-well-paid and less-well-informed than bank managers. The history of bank supervision does not provide much basis for optimism that they will succeed. Northern Rock provides a recent example of a notable failure to do so. In June 2007, just before the near collapse of the bank, the British Financial Services Authority (FSA) authorized Northern Rock to apply the Pillar I internal-ratings-based risk weights, which reduced its required regulatory capital and Northern Rock increased its dividends by 30% the next month. The FSA made no attempt to offset the reduction in Pillar 1 capital charges with a Pillar 2 capital nor did it require Northern Rock to conduct a stress scenario that would have shown that it was fatally exposed to a liquidity shock.

How should prudential supervisors deal with *U*? As Scott notes, firms can limit their leverage and maintain enough capital and liquidity to absorb unknowable losses if they should occur. But how much slack is sufficient? By assumption that is unknowable, but almost all of the things that banks could do to cope with the unknowable are very costly and competitive pressures may make it very difficult to sustain such precautions. Should regulators therefore require that banks hold capital substantially in excess of the regulatory minimum as a safeguard against unknown and unknowable shocks? Increasing capital charges for risks that cannot be identified becomes a deadweight cost and may lead to the circumvention of regulation and riskier outcomes. As Andrew Crockett observed at the opening conference, which began this project, it is inherently difficult for policy-makers to strike the proper balance between the efficiency losses associated with excessively onerous preventative policies and the cost effectiveness of responding *ex post* to adverse events. For regulators as well as firms, the appropriate amount of financial slack is an unknown.

Pillar 3 of the Basel II approach is intended to enhance market discipline by improving disclosure. The authorities may collect and publish data that helps market participants understand the current state of the economy and financial markets and the condition of regulated financial institutions. But growing reliance on dynamic trading strategies to manage risk has made it increasingly difficult to provide a meaningful picture of risk exposures. Positions may change so rapidly that information is out-of-date before it can be published. Moreover, the chief motive for market discipline – the fear of loss – is often undermined by the reluctance of the authorities to permit the creditors and counterparties of systemically important financial institutions to suffer loss. The extraordinary bailout of Bear Stearns by the Fed is a recent case in point.

The ambitious new Basel II approach attempts to incorporate in capital regulation what is known about risk management, but it may generate several unintended consequences that could shift the financial system into the domain of the unknown. The attempt to force all major firms to adopt one version of “best practice” and especially the imposition of a regulatory model of credit risk may increase the likelihood of herding. Banks may be much more likely to attempt to move the same direction at the same time which can undermine the liquidity of markets and increase volatility. To the extent that Basel II succeeds in making capital requirements more risk sensitive, it will make bank

lending more pro-cyclical. In a boom, risks are likely to decline and, consequently, so will required capital thus facilitating additional lending. In a recession, internal ratings will migrate downward, thus increasing required capital just as credit losses erode the bank's capital position thus constricting the supply of loans. More risk-sensitive capital requirements may inadvertently accentuate booms and busts. It may make individual banks safer, but weaken the banking system. More fundamentally, Basel II fails to deal with systemic risk.

Crisis Mitigation

Because it is so difficult for prudential supervisors to fulfill their responsibilities *ex ante*, policy-makers must often shift into crisis management mode to mitigate, *ex post*, the consequences of a shock. Kohn observes that in a financial crisis, the ratio of u and U will be especially large relative to K . Policy-makers must deal with unknowns such as the size of the disruption. How large will it be? How many firms will be involved? How long will it last? How likely is it to have serious spillover consequences for real economic activity?

Part of the problem is in anticipating the channels of contagion. Which firms have direct exposure to the shock? Which firms have indirect exposure because they are counterparties or creditors of the firms that sustain a direct impact or because they have similar exposures and could lose access to external financing? Which other firms might be placed in jeopardy because of the forced liquidation of assets in illiquid markets? Risk preferences and perceptions of risk are dynamic and so a flight to quality often occurs. Market participants may sell assets whose prices are already declining and avoid any counterparty that might be impaired.

Another part of the problem is that policy-makers must operate with incomplete knowledge about the current state of the economy and how their actions (or inaction) may affect economic activity. Moreover, monetary policy operates with long and variable lags and it is difficult to anticipate market responses to shocks. Yet the monetary authorities, Kohn argues, must immediately determine whether there is adequate liquidity in the financial system and whether monetary policy needs to be adjusted to counter the effects on the economy of a crisis-induced tightening of credit.

In a crisis, policy-makers must try to convert u into K as quickly as possible. This requires close cooperation across regulatory authorities with a country and, increasingly, across borders. Inevitably, the primary source of information is major market participants. But conflicts of interest may corrupt flows of information. Information may be selectively communicated to serve the self interest of market participants who might be the beneficiaries of crisis management policies. Does this argue for a direct role of the crisis manager in supervising systemically important institutions? The Fed insists that it does, but central banks lack such authority in many other countries (Herring and Carmassi, 2008) and the new Treasury proposal for reforming the US financial system removes supervisory authority from the Fed while increasing its responsibility for crisis

management. How best to organize prudential supervision and crisis management remains a significant unknown.

Policy-makers must also convey information in a crisis. Kohn raises the question of what is the appropriate response. They may urge firms to do what the policymakers believe they should do in their own self-interest, as happened in the LTCM crisis in 1998. But when is it appropriate to be reassuring? When might reassurance prove counterproductive?

Crisis management may inadvertently lead to larger future crises. If risk-takers are protected from the full negative consequences of their decisions, they may be likely to take greater risks in the future. This presents a difficult dilemma for crisis management. The costs of inaction are immediate and obvious. It's easy to imagine damaging outcomes and self-interested market participants will press for official support and can easily muster political support. Inaction in a crisis is likely to be subject to blame even when it is appropriate which may contribute to an inherent tendency to oversupply public support. Once it has been provided, entrenched interests will lobby to keep it and new additional activity may depend on it. Moreover, moral hazard manifests itself slowly and may be difficult to relate to any one particular policy choice.

Kohn argues that moral hazard is less likely if policy is directed at the broad market rather than individual firms. From this perspective open market operations are a better means of adjusting aggregate liquidity to meet the demands that arise from a flight to safety. Although this kind of response may encourage risk-taking, it may also genuinely lower risk. Direct lending and bailouts are much more likely to distort incentives. Ultimately, efficient resolution policy may be the best safeguard against moral hazard. But in most countries policy-makers lack the appropriate tools to resolve a large, complex financial institution without jeopardizing the rest of the financial system (Herring, 2004).

III. Onward

The chapters that follow heighten our awareness of the existence of and distinctions among K , u and U risks, pushing in a variety of contexts toward improved risk measurement and management strategies. Because K risks are often amenable to statistical treatment, whereas uU risks are usually not (despite their potentially large consequences), substantial resources will continue to be deployed in academia, industry and government to expand the domain of K when possible. Surely Ralph Gomory (1995) was correct in noting that “In time many things now unknown will become known,” so it is appropriate for a significant part of this book to focus on K .

But as we also emphasize, there are sharp limits to expanding the domain of K , so it is also appropriate for a significant part – indeed the larger part – of this book to focus on uU . The important issues in the world of uU are more economic (strategic) than statistical, and crucially linked to *incentives*: How to write contracts (design

organizations, formulate fiscal or monetary policies, draft regulations, make investments ...) in ways that create incentives for best-practice proactive and reactive risk management for *all types* of risks, including (and especially) *uU* risks. As Gomory also notes, often “*We do not even know* if we are dealing ... with the partly known, the mainly unknown or the unknowable” (our emphasis).

We look forward to the evolution of financial risk management toward confronting **K** and **u** and **U** equally, and we are pleased that essays collected here take us in that direction.

References

- Andersen, T.G., T. Bollerslev, P.F. Christoffersen and F.X. Diebold (2006), "Practical Volatility and Correlation Modeling for Financial Market Risk Management," in M. Carey and R. Stulz (eds.), *Risks of Financial Institutions*. Chicago: University of Chicago Press for NBER, 513-548.
- Borch, K. (1962), "Equilibrium in a Reinsurance Market," *Econometrica*, 30, 424-444.
- Christoffersen, P.F. (2003), *Elements of Financial Risk Management*. San Diego: Academic Press.
- Doherty, N.A. (2000), *Integrated Risk Management: Techniques and Strategies for Reducing Risk*. New York: McGraw-Hill.
- Gomory, R. (1995), "The Known, the Unknown and the Unknowable," *Scientific American*, June.
- Greenspan, A. (2008), "The Fed is Blameless on the Property Bubble," *Financial Times*, April 7, 2008, p. 9.
- Herring, R.J. (2004), "International Financial Conglomerates: Implications for National Insolvency Regimes," in G. Kaufman (ed.), *Market Discipline and Banking: Theory and Evidence*. Amsterdam: Elsevier, 99-129.
- Herring, R. and J. Carmassi (2008), "The Structure of Cross-Sector Financial Supervision," *Financial Markets, Institutions and Instruments*, 17, 51-76.
- Jorion, P. (1997), *Value at Risk: The New Benchmark for Managing Financial Risk*, (Third Edition, 2006). New York: McGraw-Hill.
- Kelvin, W.T. (1891-1894), *Popular Lectures and Addresses*. Three volumes. London: MacMillan and Company.
- Knight, F.H. (1921), *Risk, Uncertainty and Profit*. Boston: Houghton Mifflin Company.
- Rawls, J. (1971), *A Theory of Justice*. Cambridge, Mass.: Belnap.
- Santayana, G. (1896), *The Sense of Beauty*. Dover Edition, 1955.
- Thornhill, J. and A. Michaels (2008), "Bear Stearns Rescue a 'Turning Point'," *Financial Times*, April 7, p. 4.