Vladimir Uhmylenko

Integrated Management of Risk and Capital in Insurance Organizations

Article outline:

Summary

This commentary article presents an overview of the capital modeling frameworks adopted by large and complex (re)insurance organizations. The article focuses on risk metrics/measures used, analytical/modeling tools, lessons learned from the recent financial turmoil, and the emerging trends in capital-adequacy analyses. The main conclusions we draw are: 1) The approaches to capital-adequacy analyses vary widely, and consensus across the industry on what may constitute best practices has been slow to develop. 2) Stochastic capital models are establishing themselves as predominant capital-adequacy analysis tools. 3) Although concerns remain about the robustness of some of the adopted methodologies, consistency of implementation and integration into decision-making, we also recognize the great progress the insurance industry has made over the last three years.

1. Background

Capital needs of an insurer are driven by its risk profile, that is the risks it has assumed and is likely to assume in the future (for example, over the next year), as they are mitigated by the risk controls, and seen through the lens of potential future adversities that may affect the company, the industry, or the whole economy. The ability to continue to operate under extreme circumstances is a key factor in determining whether the insurer's capital is adequate. Capital management decisions are taken in the context of the firm's liquidity and financial flexibility, as well as regulatory requirements and rating-agency expectations.

The recent trend in regulatory and rating-agency analyses is to view an insurer's risk profile through the lens of it potential impact on the volatility of earnings and capital –

factoring in the firm's risk management and its ability to contain the risks within the predefined tolerances and limits.

This commentary article presents an overview of the most common analytical practices that insurers use to evaluate their capital needs relative to their risk profiles. Specifically:

1. Stress testing: Extreme scenarios used to evaluate risk-bearing capacity and identify capital and funding needs; as well as to establish risk appetite, test risk controls, explore risk contagions, test liquidity, etc.
2. Stochastically calibrated stress scenarios and factor-based models: A structured scenario approach to evaluate economic capital (EC) from "silo" capital calculations. It is a bridge between stress testing and fully stochastic capital-at-risk modeling.
3. Integrated stochastic capital-at-risk models: Dynamic financial models that are used to simulate the development of an insurer's risk profile as an interplay of the legacy, in-force, new-business, and macro-economic risks, the size and structure of the firm's capital, financial flexibility, existing risk controls, and the firm's prospective risk appetite.

This commentary article is based on our discussions with many of the leading Enterprise Risk Management (ERM) practitioners. The article will also highlight the developing trends and the lessons learned from the recent financial turmoil. It will also provide an overview of tactical and strategic applications based on internal capital models.

2. A Comment on Capital-Adequacy Metrics

How do insurers measure their capital needs? Perhaps the most widely used approaches involve VaR-type (Value-at-Risk) metrics that are often interpreted as representing worst-case outcomes in a particular financial indicator at a given level of confidence. More precisely, as an indicator of the "downside", VaR is the best of the worst, say 0.5%, outcomes. For example, if a firm ties its capitalization to the 0.5% Capital-at-Risk metric, then it may imply that they target a level of capital that 99.5% of the time would be sufficient to absorb next year's potential adversities.

Based on our discussions with many leading insurers, they see VaR as a useful indicator of risk, but are aware of the limitations of decision-making that is driven solely by VaR. For example, the 0.5% VaR is "blind" to the extreme events and event contagions in the worst 0.5% of outcomes, and therefore may lead to a bias against using risk-transfer mechanisms such as reinsurance, catastrophe bonds, and hedging.

In the context of capital-adequacy analyses, VaR may be indicative of the likelihood of insolvency, which probably explains its popularity and use in regulatory regimes (e.g. Basel 2-3 and Solvency 2) and rating-agency capital-adequacy analytics. At the same time, practitioners are aware that VaR tells only part of the story. For example, it tells little about how damaging these insolvencies may be to the policyholders.

Another popular risk indicator is TVaR, or Tail Value at Risk (a.k.a. Conditional Tail Expectation, or CTE). Unlike VaR, this metric focuses on the "tail", e.g. the worst 1%, of the probability distribution, and thus may better represent extreme outcomes.

However, the 1% TVaR may lose sight of what happens before the "tail", i.e. in the 99% of the probability distribution. These metrics are also very sensitive to tail-behavior assumptions, tend to be computationally more demanding, and more difficult to interpret and communicate. In the capital-adequacy context, TVaR may better reflect the loss to policyholders under the most extreme outcomes, but it does not explicitly capture the likelihood of insolvency.

Transforms and offshoots of these metrics are also used in decision-making applications of capital models. For example, co-TVaR and co-xTVaR are used for the purpose of cost and capital allocation. In fact, at the organizations that we view as leaders in ERM the measurement of risk and its capital implications is quite sophisticated and tends to combine multiple metrics. These measurements are designed to facilitate an integrated view of risk and capital, and to reflect the diminishing appetite for progressively larger losses. These practices try to weigh conflicting objectives, constraints, time horizons, etc., but do not rely on any single metric to fully reflect the capital implications of the risk profile of a complex insurance enterprise.

3. Stress-Testing

Although stress scenarios have long been part of insurers' decision-making, over the last two years they have become truly indispensible in the risk-management toolbox of many insurers. They are used in capital-adequacy and liquidity analyses, to test the effectiveness of risk mitigation (e.g. reinsurance and hedging), to evaluate potential risk contagions, and even to validate complex stochastic models. Moreover, stress-testing is being increasingly incorporated into regulatory and rating-agency reviews to evaluate insurers' ability to withstand hypothetical shocks over a given time horizon.

With that, financially strong firms are expected to withstand more severe stresses. For instance, Standard & Poor's has articulated a set of economic stress scenarios associated with particular rating levels (as discussed in Appendix IV of: "Understanding S&P's Rating Definitions", Standard & Poor's, June 3, 2009). However, even more stressful would be scenarios that combine multiple unfavorable factors and risks, such as macro-economic shocks coupled with industry- and company-specific events. For example, larger-than-usual hurricane damage experienced in the midst of a major financial downturn can exacerbate the economic stress for a Property-Casualty insurer. During the recent financial crisis, many life insurers experience a real-life example of stress-event concurrencies, particularly affecting variable annuities with guaranteed living benefits. When both the equity market and interest rates declined, the adverse impact on the insurers' earnings and capital was magnified further by industry-wide lower-than-expected lapsation as policyholders took advantage of the in-the-money guarantees.

Many insurers also use stress scenarios to describe their risk tolerance. For instance, an insurer may conclude that in the business environment that they anticipate for the next year they would be able to absorb $200 million of unbudgeted losses from their insurance operations, while in a severely depressed business environment they would only be able to withstand a $100-million loss. Armed with this analysis, they may decide that they should try to prudently confine such losses to $80 million. In our experience, leading

ERM practices can readily demonstrate that their risk tolerances are set quite prudently – keeping a comfortable distance away from the brink of financial distress.

In addition to testing for solvency and liquidity, stress scenarios may consider various other implications. For example, a stress event may undermine the firm's competitiveness, erode the market share, affect the availability and the cost of borrowing, trigger contingent obligations, prompt exodus of key personnel, and force the firm to forego strategically important plans. With their focus on solvency, economic capital models tend to ignore these effects, but some insurers try to reflect them through scenario analyses. When these effects are factored in, the analysis will more adequately assess the true financial impact of stress scenarios, and identify the maximum adversity the insurer can withstand without falling into financial distress.

4. Stochastically Calibrated Stress Scenarios

If credible quantification is possible, stress scenarios – such as for market, credit, insurance, or operational risks – can be generated stochastically and used to analyze adverse deviations from the expected cash flows. For example, stochastic projections – such as the 0.5% and 99.5% percentiles for interest rates – may help analyze potential asset-liability mismatches in the in-force business of a life insurer. The most common practice is to calibrate each of these "shocks" as VaR to a level consistent with the rating or regulatory solvency (e.g. 99.5%) targets, and then compare the present-value impact of these shocks against the insurer's existing capital. With that, some insurers go further and interpret these VaR-calibrated present-value impacts as indicative of the capital needed to offset the risk.

Moreover, by aggregating the per-risk shocks ("capitals") – usually via a correlation matrix (variance-covariance approach) – many insurers hope to estimate their economic capital (EC), i.e. the capital sufficient to avoid potential future insolvencies in all but the most extreme events (e.g. in all but 0.5% of the outcomes). Mathematically, such aggregation of VaR-calibrated shocks may be valid under the assumption that these risks are symmetrically ("elliptically") distributed – as in the case of normal probability distributions, – and that linear (Pearson) correlations are known and fully capture the interdependencies between the risks.

Unfortunately, generally these assumptions are overly simplified. Indeed, except for, perhaps, small time horizons, the risks affecting an insurance enterprise – individually and in the aggregate – tend to be "fat-tailed", i.e. highly asymmetric. Moreover, the interdependencies between an insurer's risks tend to be non-linear and much more pronounced in extreme (tail) events, which was evident in the risk contagions of the recent financial turmoil. Thus, despite the seeming precision implied by the stochastic calibration of the scenarios, this VaR-based methodology may produce only a rough EC estimate. This is not to say that statistical methods cannot be used to develop risk scenarios, but rather that a more nuanced risk-aggregation approach is needed to credibly assess economic capital.

Recognition of these methodologies' limitations is growing, and over the past several years most of the large P&C (re)insurers have adopted more advanced internal capital models (as described below).

5. Integrated Stochastic Capital Models

For the lack of a better term, these more advanced models are often referred to as "integrated models" – to emphasize their attempt to re-create the insurers' risk profiles as a dynamic interplay of all material risks and risk controls – in the context of the available (existing and accessible) capital and in the environment of projected balance-sheet valuations. Rather than developing just a handful (e.g. "shock up" and "shock down") of scenarios, these models use random-sampling simulation techniques to generate large numbers, often in hundreds of thousands, of future scenarios, including extreme events and contagion effects that may afflict an insurer's financial health – its profitability, liquidity, and even solvency. Rather than relying on diversification or correlation assumptions, whenever possible these models try to re-create the relationships between risk drivers and the risks, and recognize ripple effects and the heightened interconnectedness between the risks in extreme circumstances. These models do not try to calculate the "capitals" for each risk category, but instead focus on the overall capital needs caused by the complex risk interdependencies, and then allocate the capital to specific risk sources based on their contribution to the aggregate risk.

These models' starting point is to define their capital-adequacy targets – group-wide and for individual operating companies. Most of the models we have seen focus on rating-specific targets – for example, maintaining capital at a level consistent with the "A" rating, which some interpret as the level that assures ongoing solvency in 99.92% of the modeled one-year outcomes. The most nuanced of these models perform operating-company solvency checks in each simulated scenario; and in the most extreme scenarios the capital is drawn from the group's resources (e.g. the holding company's capital) to try to offset the simulated breaches of operating-company regulatory capital requirements. Then, if the model indicates capital deficiency or redundancy (e.g. relative to the 99.92% ongoing-solvency target), a "what if" analysis may help determine the targeted capital by varying the amount of capital to raise or release (taking into account the capital-structure and cost-of-capital implications).

5.1 Dependence Structure

Of the models we have seen, most use outputs from Economic Scenario Generators (ESG). In our opinion, an ESG is a critical component in a capital model's dependence structure, as each simulated macroeconomic environment naturally brings together ("correlates") many of the risks of an insurance enterprise. The idea behind an ESG is to create realistic and consistent (non-contradictory) sets of economic scenarios, e.g. for future interest rates, bond spreads, inflation, stock-market indices, GDP growth, and unemployment; and then link them to the insurer's projected asset and liability valuations. In some – especially multi-year – models, the macroeconomic projections may also be linked to future business volume, claim-cost trends, asset sales and new investments, as well as management actions such as raising capital, scaling down unprofitable business, etc. With the explicit linkage to the underlying cause (an economic environment), ESG-based models are likely to better re-create risk contagions at an insurance enterprise, such as those between a drop in business volume, escalation of credit/counterparty risk, drop in asset valuations and increase in liability valuations.

Unfortunately, with respect to ESGs the practices vary widely. For example, a debate is still going on among the proponents of "market-consistent" and "real-world" ESG parameterizations. The argument in favor of market consistency emphasizes the alignment with current valuation of the balance sheet; while "real world" ESG proponents point out that the economic amount of capital should be sufficient to absorb the impact of potential market dislocations that may not be fully[1] reflected in the valuations currently observed in the markets.

Another debate is between the proponents of in-house-built ESGs and those provided by 3rd-party vendors. While some of the in-house-built ESGs we have observed are transparent and easy to follow and replicate, they also tend to be somewhat simplistic, and may not fully capture the potential contagions between the economic variables. On the other hand, while vendor-developed ESGs may be built on sophisticated macroeconomic theories, they also appear as "black boxes" to many users. In fact, a frequent complaint from the users is that they do not feel they have full "ownership" of the vendor model outputs. The inability to fully understand and trust the models may go against what is known as the "use test" that is a prerequisite for regulatory (e.g. Solvency II) and rating-agency (e.g. S&P ERM Level III) model approval. This may raise doubts about the capital model's ability to satisfy other tests, such as the quality, validation, and documentation tests.

In addition to ESGs, the dependence structures of integrated stochastic capital models often contain other – insurance-specific – layers of interconnectedness. Whenever the interdependencies lack clear causal relationships, they tend to be modeled using (rank-order) correlations or copulas. On the other hand, in some cases a risk driver can be identified. For example, if a catastrophe or a pandemic were to cause a large insurance loss, they would also likely escalate the reinsurance counterparty risk, while also pushing up reinsurance pricing in subsequent renewals. Recreating such cause-and-effect relationships is, in our view, a superior (to correlations and copulas) way of capturing the contagion effect and the resulting capital-adequacy impact.

With its tangled layers of interconnectedness, the dependence structure of an integrated stochastic capital model can be very difficult to parameterize and validate. For example, its calibration may involve an iterative approach with multiple model runs to assure that the resulting cross-risk correlations agree with the empirical data and intuitive perception of risk correlations. This is just one of many complexities often left unaddressed, but potentially aggravating the model risk.

5.2 Model Risk

Evaluation of capital needs tends to be driven by the modeled extreme events and risk contagions. Yet, quantification of extreme events and their interdependencies with other variables is inherently imprecise. In fact, even when based on substantial data, the particular theories, methodologies, and assumptions chosen to model the extreme-event behavior may heavily influence the outcome. Moreover, complex stochastic models are prone to various parameter-estimation, approximation, extrapolation, and sampling errors. Human error and inadequate execution also contribute to model risk.

[1] Unless a market-consistent ESG is parameterized right in the midst of financial turmoil, it is likely to project lower volatilities than a real-world ESG.

Model risk is unavoidable. However, it can be reduced through data and methodology improvements, through model validation, and through controls over the modeling process and execution. Deliberate conservatism in assumptions / parameterizations may help avoid underestimation of capital needs. Increased computer-processing power and a more efficient model design may help minimize the random-sampling risk.

A model's robustness is always a function of its granularity and completeness, i.e. its ability to capture the vast majority of risks that may potentially cause the insurers' distress. Most insurers note that, unfortunately, some risks and their interdependencies may defy reliable quantification (e.g. emerging risks, man-made catastrophes, pandemics, operational risks, new-business volume and pricing, etc.). As a result, the impact of these risks tends to be modeled rather simplistically, or even entirely omitted.

In our experience, few insurers attempt to robustly evaluate the capital impact and materiality of such simplifications and omissions. Nevertheless, most insurers try to, at least partially, compensate for the "imperfect knowledge". For example, some of them scale down their risk appetite and tolerances, when in doubt about the robustness of the models. Others pursue a deliberate conservative bias in parameterizing extreme-event models. Yet others address these uncertainties by requiring a capital "buffer" in excess of the model-indicated economic capital.

Reflecting the size and complexity of an insurer, a capital model can be quite large and complex, which often poses big computational challenges. This is especially true for higher-rated insurers, as they tend to measure their capital adequacy against more-demanding targets (e.g. 99.92% one-year solvency rate), and thus may need very large simulations – with hundreds of thousands (and more) of iterations – to (credibly) evaluate potential insolvency events. For many models this may involve long hours of distributed computing, while for some models this may be computationally impracticable.

A model's granularity tends to increase its complexity, and may also exacerbate the model risk caused by the compounding of the inevitable parameter-estimation errors. Moreover, a model's complexity tends to slow down the calculations, forcing the users to reduce the simulation sample and thus increase the sampling error. The complexity and size may also make it more difficult to maintain, govern, and interpret the model. These risks are especially heightened in multi-year risk-and-capital models.

In our discussions with large multi-line insurers, they report that it may take years (at least two) to settle on an efficient model design that adequately captures the aggregate risk without resorting to overly simplistic methods. Yet, other insurers, recounting the analytical and computational challenges, give up on building more sophisticated models and continue to rely on simplified methodologies. Some insurers, instead of looking for a more efficient model design, settle for inadequate (too small) simulation samples. Others, trying to speed up the calculations, use replicating portfolios to approximate their liabilities. Other shortcuts involve extrapolation from smaller percentiles (e.g. 99%, 99.5%) to evaluate the capital needs relative to higher percentiles – say, 99.92%. In our opinion, such simplifications and shortcuts are likely to aggravate the model risk (e.g. the sampling, approximation, and extrapolation errors), but may be necessary for practical reasons.

Particularly worrisome are those shortcuts and parameterizations that may understate the capital needs. To minimize these concerns, regulators and rating agencies are starting to look for evidence of robust model testing and validation, thorough documentation, and management buy-in, and integration into decision-making (the "use test"). Although, as of this moment few insurers would be able to readily provide such evidence, we have observed definite improvements and focused efforts to address these concerns.

5.3 Time Horizon

The insurance business is accustomed to stochastic multi-year projections of cash flows. Take, for example, stochastic asset-liability-management (ALM) models. For a given and predictable liability profile, ALM models test the assets' ability to fully offset all in-force contractual obligations and related business expenses projected dozens of years into the future. Accordingly, some insurers define Economic Capital (EC) as the amount of assets needed today over and above the current value of reserves and liabilities to policyholders in order to fully run off the in-force obligations with only a minimum – to a pre-determined level of confidence – risk of insolvency.

The runoff approach carves out the capital needed to offset the existing, that is in-force, liabilities as if they were isolated from the rest of the business, and disregards the evolving capital structure, the cost of capital, new business, or a strategic change of direction. However, unless it truly is a runoff operation, over time a firm's risk profile tends to undergo significant changes. While some of these changes are gradual and predictable, others can be explosive and dramatic. Consider, for example, the life-insurance industry transformation brought about by the growth in guaranteed living benefits in variable-annuity products. Can a runoff-focused capital model capture the impact of such dramatic risk-profile changes? The limitations of the runoff (read: in-force only) approach are particularly evident in short-tailed insurance businesses. For example, at the end of a benign hurricane season a Property insurer in Florida may not need much capital to run off its in-force liabilities. However, this tells little about the capital needed to support the ongoing operations, or about the capital the insurer may release to the shareholders.

If left unmanaged, future uncertainties and risks would escalate over time, and more capital would be needed today to offset them. However, insurers do manage their risks. For example, if a wide duration or cash-flow mismatch develops between the assets and liabilities, some corrective actions will likely be taken. Conceptually, this should result in a lower EC number. Accordingly, a model may try to "foresee" such future actions through a set of built-in rules emulating the management's responses. These rules may also address capital management, such as raising or releasing capital when the built-in solvency and capital-adequacy checks indicate a deficiency or a redundancy. However, the EC concept typically does not take into account future capital-management actions, and focuses instead on the capital needed today to offset potential future adversities. The limitations caused by the present-value focus of the EC concept are especially evident in multi-year models.

To support their decision-making, a small number of practitioners have been using multi-year models that try to mimic the real-world risk and capital dynamics, stochastically simulating the future balance sheets with the rolling view of adding new business and period-to-period runoff of the existing obligations. One can think of these models as financial planning under hundreds of thousands of stochastically generated scenarios. In each scenario, such models assume that under the simulated circumstances the management will take a particular course of action. For example, if six quarters from now the capital falls below a pre-defined threshold, the rules built into the model may dictate that the capital be replenished by raising debt or equity, or that new business writings be reduced, or more risk be transferred, etc. Unfortunately, with each modeled time increment (a quarter or a year), the credibility of the balance-sheet projections diminishes, as the model risk balloons fueled by the layered assumptions. The built-in feedback loops and the nested stochastic projections (e.g. to determine the risk margins in liability valuations) place severe computational demands. As a result, beyond just several years such granular multi-year dynamic financial models may become impracticable.

This is probably why capital models that focus on one-year solvency dominate the insurance industry. We should also note that most of the implementations of one-year models are not granular enough to perform intermediate (e.g. quarterly) solvency checks, and therefore are likely to understate the modeled rate of insolvencies.

With its focus on one-year solvency, a one-year model may only partially reflect longer-term capital implications of most strategic decisions. Exclusion of new business from the analyses further reduces a capital model's ability to serve as strategic-guidance tool. This was one of the lessons some Life insurers learned in 2008. Not only did they fail to adequately parameterize their models, but they also failed to capture the nascent threat of the risk-profile transformation caused by rapid variable-annuity expansion (VA) in their risk portfolios. By the time these large VA writers started to scale back new sales or de-risk new products, their balance sheets had already been strained by the accumulation of VAs with guaranteed living benefits. To make things worse, these companies could not resort to traditional risk-transfer mechanisms because of diminishing availability of reinsurance and prohibitive hedging costs.

6. Risk Vs. Capital Analytics

In our discussions with insurers, they frequently refer to using internal capital models to support tactical and strategic decisions. Below we provide an overview of the decision-making applications we commonly see at insurance organizations.

Integrated risk-vs.-capital models have proven to be instrumental in analyses of the capital's sensitivity to risk drivers, and as such, they are routinely used to set and monitor risk tolerances and limits. For example, pressured by the recent financial turmoil many insurers have recognized the increased sensitivity to catastrophe, market and credit risk exposures, which has prompted many of them to scale down the tolerance for these risks and impose additional and stricter limits on the size and concentration of these exposures. While in many cases this has affected their earnings, some of these insurers argue that these new practices have reduced excess earnings volatility, and on a risk-adjusted basis may actually improve profitability.

A recent trend for many insurers has been to publicly state their targeted returns (e.g. on equity) – sometimes along with the probability distributions from their capital models. If the targeted return is likely to be achieved, while the risk to capital is well contained, then it may imply that the capital is utilized efficiently. To segment the capital efficiency analysis, the most common practice is to transform the enterprise return target into business-unit or line-of-business targets for returns on allocated (risk-adjusted) capital. Where practices diverge is in how *exactly* the capital is allocated. While all capital-allocation methodologies try to reflect a segment's contribution to the potential depletion of capital, they may produce dramatically different results, and thus cause divergent views on the profitability of a business segment. Moreover, many insurers have discovered that a decision to change one business segment may change other segments' allocated capital, and thus, the evaluation of their profitability. Furthermore, sometimes there is also a significant degree of subjectivity, arbitrariness and judgment, which may raise questions about the robustness of the capital-allocation process. These issues can become particularly sensitive when they affect the view of a segment's performance, or influence the incentive compensation.

In our opinion, business-segment RoE / RoRAC analyses are among the most useful and meaningful applications of internal capital models. They help assess each segment's contribution to the overall strategic objective (e.g. RoE), and have become a standard feature of leading ERM frameworks. At the same time, the choices made in the capital allocation process are often not transparent to internal stakeholders and not fully understood; with little consensus in the industry on what constitutes best practices.

Another popular application of capital models is in analyses of business decisions relative to their impact on Economic Capital (EC) or other capital-at-risk metrics. For example, many insurers use these models to compare risk-transfer (e.g. reinsurance or insurance-linked securities) alternatives relative to a capital-impact metric, such as the EC or the Solvency-2-inspired 99.5% capital-at-risk. Apparently, the premise there is that risk transfer creates value by minimizing the EC or the solvency capital requirement. Another application is common among property-catastrophe reinsurers who have been using capital models to price individual transactions relative to their incremental capital consumption in the context of the overall (rolling or projected) portfolio. In life insurance, a few leading companies are testing product features and pricing alternatives relative to their impact on EC.

Most of these applications are essentially "what if" analyses: How would the EC (or another capital-consumption indicator) change if a particular alternative were implemented? However, capital consumption does not necessarily reflect *value creation or destruction*. For example, a hurricane may erode a reinsurer's capital, but it may also increase the demand (and price) for its capacity, helping the firm to re-capitalize and position itself for better profitability going forward. On the other hand, capital depletion may not tell the full story of value destruction. For example, in addition to eroding capital, an underwriting loss may get interpreted as evidence of chronic underperformance – undermining the insurer's credibility and financial-strength rating, impairing growth opportunities, and thus destroying value far in excess of the lost capital. This illustrates why capital depletion is not always equivalent to value destruction, and why EC may not be a good value-creation measure. If this is true, what role does EC play in decision-making? To answer this question, we explain and generalize the value-creation analytical framework we typically see at insurance organizations.

Suppose that an A-rated insurer's capital model generates 0.05% of insolvencies out of all simulated one-year outcomes. Also, suppose that this insurer defines its EC as the capital sufficient to assure a one-year solvency rate of 99.92% (which is consistent with the 0.08% historical average default rate for S&P A-rated bonds). Since the modeled 0.05% rate is lower than the allowable 0.08%, the insurer may interpret this as capital redundancy relative to EC. This redundancy may open a range of capital-management and risk-taking alternatives for the insurer to consider – such as buying back shares or increasing the allocation to riskier instruments in the investment portfolio, or expanding sales of riskier products, or purchasing less reinsurance – as long as they do not push the modeled insolvency rate above 0.08%. Of course, various regulatory and rating-agency expectations, as well as the insurer's own risk-appetite framework, may further constrain the range or scale of available alternatives. The alternatives are then compared against a value-creation metric, such as a one-year RoE or three-year growth in book value per share. To use Decision Science terminology, this is an optimization exercise under a set of constraints[2].

Observe that under this framework the EC is relegated to a secondary role of a constraint (one of many), while the target (objective function) is to maximize value creation. Moreover, even as a constraint VaR-based economic capital may not be sufficient, as it only reflects the probability of insolvency – without addressing the loss given insolvency. Conceptually, focusing solely on EC will likely create a bias against managing (e.g. via reinsurance or hedging) tail risks. Therefore, additional constraints may be needed to help contain the loss given insolvency (the risk to policyholders). For example, in addition to imposing a constraint on the probability of insolvency, a model may also seek to constrain the average loss given insolvency.

Mindful of VaR's limitations, some insurers define their economic capital needs relative to tail value at risk (TVaR). Although these TVaR-based EC definitions may better reflect the risk to policyholders (i.e. the risk that the assets supporting the reserves and the capital are insufficient to meet all policyholder obligations), some of these insurers could not clearly explain the choice of the particular metric (e.g. 0.5% TVaR, but why not 0.62% TVaR?). Furthermore, it was unclear how those TVaR-based economic-capital frameworks could be linked to rating-specific capital adequacy at the targeted confidence level.

Overall, VaR-based capital modeling approaches continue to dominate, and only a small number of insurers have taken deliberate steps to compensate for the potential bias that VaR-focused decision-making may cause. For example, some property-catastrophe reinsurers have been using composite metrics (derived from VaRs or TVaRs) to evaluate the marginal capital impact when pricing reinsurance deals. A few others have been using probability transforms or expected-utility scoring to value-rank large reinsurance transactions. These approaches are more likely to discourage excessive risk-taking or accumulation of risk, as they try to recognize the "disutility" of the alternatives that aggravate risk to the capital and to the policyholders. Moreover, they attempt to better

[2] Such as $C \geq EC$ and $C \geq RC$, i.e. making sure that economic capital (EC) and regulatory capital (RC) do not exceed the actual capital (C) held.

reflect the whole risk profile – and not just one point (VaR) on the probability distribution.

Despite the claims by many insurers that their capital models are used to support strategic decisions, the decision-making we have observed is mostly tactical (non-strategic) in nature, such as supporting pricing and reinsurance purchasing decisions. This is because EC models tend to view the future through the lens of today's capital – usually only one year ahead, and typically ignoring future capital-management and other actions that are likely to be taken to adjust the balance of risk and capital. Be it a strategic business line expansion or an acquisition, modeling the impact of a strategic decision is inherently uncertain and complex, as it tries to project the future paths conditionally on the simulated business and economic environments. Despite this complexity, a small number of insurers have been using models that attempt to go beyond the one-year EC construct and reflect long-term value creation and risks from strategic decisions. The few cases we have observed address very specific objectives, such as strategic asset allocation or an acquisition, and are built on multiyear (3-5 years) dynamic financial models. Unlike a typical ECM, these models de-emphasize the present-value capital impact of future uncertainties, focusing instead on the need to assure ongoing solvency over a multiyear horizon (e.g. 99.32% solvency over five years – for an "A"-rated company), and set out to maximize a long-term value-creation target, such as book value per share. These models combine assumptions about the future business with a set of built-in rules allowing for discretional management actions, such as future share buy-backs or replenishing capital when it falls below an acceptable level.

We applaud these attempts to build more nuanced decision-making frameworks that address long-term risk-vs.-reward implications. However, we also note that the complexity and the computational challenges in such models can make them impracticable – especially for frequent use or where quick results are required, while the model risk may raise doubts about the credibility of the recommended decisions. These challenges are particularly evident in the life insurance business, mainly due to the long-term nature of the risk profile, as well as the uncertainty of the policyholder optionality. As a result, to facilitate strategic decision-making most insurers continue to rely on more traditional Financial Economics methodologies, such as RAROC, that compare strategic alternatives against risk- and duration-adjusted hurdles rates.

In our opinion, the science of strategic decision-making under *long-term* uncertainties of an insurance enterprise is still in early stages of development – both theoretically and computationally. At the same time, we have observed many examples of frameworks that appear to successfully facilitate decision-making at insurance businesses, especially those with *shorter-tailed* risk profiles.

7. Analytical Lessons Learned

As we explained above, economic capital models are highly prone to model risk[3], because of their complexity and focus on extremely rare events. Given the high model

[3] Generally, the higher the rating, the greater the focus on the tail scenarios, and thus the higher the model risk.

risk, is it even possible to assign much credibility to the assessment of capital needs such models produce? This leads to another interesting question with respect to the firm's management and corporate strategy: If the risk profile is too complex to credibly quantify, is the business too complex to manage?

On the one hand, there are insurers that operate under very strict risk limits and controls, and whose size and complexity, in our opinion, do not preclude them from successfully containing their risks. This may explain why some of the highest-rated (many of whom happen to be redundantly capitalized) insurers, when assessing their capital needs, opt for a simpler approach and "stack" together a set of extremely adverse scenarios – on the premise that in extreme circumstances most risks tend to correlate – and still are able to demonstrate that the capital is sufficient to absorb even the most adverse outcomes.

On the other hand, as was evident in the financial crisis of 2008, many large financial institutions failed to adequately measure and manage their risks. Especially challenged have been the business models of those insurers whose risk profiles are dominated by complex financial products. Some of these products are long-tailed in nature and often involve complex long-term guarantees; for example, universal life with secondary guarantees, long term care, and variable annuities with guaranteed living benefits. Many of the affected insurers continue to struggle with finding appropriate methodologies and overcoming technological constraints when trying to quantify the aggregate risk.

In our opinion, one of the biggest lessons learned from the recent financial crisis is that sophisticated models may create a false sense of confidence and that overreliance on complex models further magnifies the effect of model risk. As a result, many companies are moving towards simpler, more transparent risk profiles. Moreover, many insurers are now paying heightened attention to extreme-event and event-contagion stress tests as a means to verify whether the risk is contained. The recognition of this issue seems to be growing, as many insurers have been scaling down ("de-risking") their exposures, imposing stricter limits, and embracing stress-testing as a critical risk-management tool.

In fact, we view favorably the use of stress tests as a validation method for capital-adequacy analyses – especially those that rely on complex capital models. For example, some insurers compare the simulated model outputs with the stress tests. If the severe stresses (including concurrences and contagions) cannot be identified within the simulated outcomes, then most likely the model has not adequately assessed the capital needs.

With the shift toward mark-to-market and market-consistent valuations, the future uncertainties caused by products with "long-tail" exposures become visible on the balance sheet very quickly, as the volatility and uncertainty tend to erode the asset valuations and amplify the volatility in valuation of liabilities (e.g. by means of a risk margin). Moreover, at least theoretically, the simulated future balance sheets will be affected even more, as the uncertainty of future projections tends to increase with every modeled time increment. With that, the greater the uncertainty of cash flows – in terms of the amounts and timing – the greater the volatility of projected asset and liability valuations, and the greater the economic capital that is needed to assure ongoing solvency. Accordingly, viewed through the lens of EC, volatile longer-tailed risks will tend to be more capital consumptive.

The silver lining to this is that it imposes a more economic view on profitability, and hence pricing, of long-tailed risks. Consider, for instance, long-tailed liability (casualty) lines, such as Excess Workers' Compensation in the US. Insurers are getting increasingly uneasy about these lines' future uncertainties and, in particular, about the high sensitivity to future claim-cost inflation. For example, having quantified the capital implications from carrying such volatile and long-term liabilities, especially in light of the uncertainties around future medical costs, one large diversified insurer has questioned the viability of its current pricing, and is drastically reducing new writings. This is an example of risk selection that better reflects the economic view of capital.

By itself, economic capital (EC) is just a capital-adequacy indicator, and is not immune to theoretical limitations or significant model risks. The benefit of economic capital modeling is not necessarily in precise quantification of the EC, but rather in determining the drivers of a firm's capital needs, and in promoting the economic principles in decision-making.

8. Conclusion: Generally Positive Trend, but Uneven Progress

Leading practitioners of Enterprise Risk Management tend to use advanced approaches to modeling and aggregation of risk in the context of the available and targeted capital. At the same time, they have a good understanding of risk-modeling limitations, offsetting them with experience-driven judgment. They also emphasize forward-looking assessment of risk, with continual monitoring of trends, prompt re-parameterizations, and frequent portfolio roll-ups within the capital models. Complex stochastic models are supplemented with deterministic stress testing.

In our opinion, the level of sophistication and the degree of adoption of these advanced practices has greatly increased since the financial turmoil of 2008-09. In particular, most large insurers and reinsurers have embraced advanced stochastic capital models. According to many of these firms, this trend may be partially explained by the increased attention from regulators and rating agencies. On the other hand, the industry as a whole is still in early stages of adoption and development of such models. Although there are some promising trends at medium-size companies, smaller insurers tend to lag far behind in the depth of expertise, modeling tools, quality of data, and management's buy-in.

We should also note that the capital-modeling practices and methodologies used in the insurance industry vary widely. While there is a clear trend toward greater analytical sophistication, some of the approaches we have observed remain simplistic or poorly implemented.