

NMS EXCHANGE



By Mark A. Schmid
Vice President and
Chief Investment Officer
The University of Chicago



By Mike Edleson
Chief Risk Officer
The University of Chicago



By Patrick O'Hara
Chief Operating Officer
The University of Chicago

A Framework for Endowment Risk Management

Introduction

In the waning years of the previous decade, before the onset of the financial crisis, The University of Chicago's Investment Committee decided to take a 'risk view' of the endowment and its governance process. The goal was to have a better understanding of the investment risks faced by the endowment, within the broader context of the strategic mission and overall risks faced by the university. Instead of taking a purely traditional asset class view (e.g., hedge funds, private equity), the investment portfolio and decisions at the governance level would be framed in terms of the key risk drivers affecting risk and return. With a new CIO in place in 2009, the Investment Office transitioned to a risk-based investment model over the next several years, enhancing rather than replacing the traditional endowment model, with a more customized solution.

A risk-based approach is not the subjugation of the pursuit of investment return to the minimization of risk. Instead of chasing the highest return at the exclusion of any consideration of risk, the focus shifts to a conscious choice of the type and size of risks best suited to our University's long-term strategies. Investors may pursue higher returns, but we don't have direct control over the investment outcomes. We can, however, directly choose the risks that we take in pursuit of those outcomes. Risk is not only an integral part of the investment process; the amount and type of risk chosen by managers are the most critical factors in determining their investment success. Risk and return are the yin and yang of the financial universe. And, risk management is not just defensive, but is a key offensive weapon in the investment manager's arsenal.

The financial crisis provided a litmus test of the importance of risk management in the long-term outlook of endowed universities. In the financial crisis, the largest endowments reported FY2009 returns ranging from -15% to -28%. From peak-to-trough, the typical large endowment losses were in the mid-30s range, as reported. Given the return lags present in reported returns, though, these reported returns suppressed the true economic losses. Including lagged or suppressed losses,

the typical endowment's economic loss estimate for the crisis exceeds -40%, with a couple of large endowments estimated at over -50%. These were real losses, with real impact on universities for a number of years. This highlights the need for a top-down proscription of adverse investment impact, which helps inform the overall risk framework. At the University of Chicago, our endowment risk management and investment strategy are driven by an enterprise view of the entire university, including the Medical Center.

Endowment Risk Management for the Enterprise

"A Total Enterprise Approach to Endowment Management," by Mark Schmid and Que Nguyen in the January 2012 *NMS Exchange*, introduced our integrated approach to investment strategy. A couple of examples will serve to illustrate this enterprise view of the endowment within the context of the university's overall business and risk profile. First, suppose the present value of the university's future gifts is directly affected by the market return environment (i.e., gifts have 'equity beta'). When considering the amount of overall exposure of the university's wealth to the vagaries of the market, it would be wrong to focus on a partial analysis of the endowment's equity risk alone—one must consider the market exposure inherent in gifts as well. A degree of equity risk that might seem acceptable to a stand-alone portfolio, may in fact be overly aggressive in an enterprise context, when considering the additional equity exposure from the gifts' value, which is sizable and illiquid. Second, suppose the university's debt issuance is sizable relative to its financial base. This financial leverage could be compounded, or offset, by the cash, fixed income, or leverage within the university endowment. Making a standalone decision on the endowment's asset allocation could leave the university far too levered, or even not levered enough. An integrated look at the risks across the enterprise should result in the right amount and right types of risk to support a university's long-term needs and risk appetite.

Our own top-down analysis of key risk sensitivities showed potential tangible adverse impact in the event of:

(1) too much equity risk, (2) too much illiquidity, (3) too little equity risk, (4) excessive leverage, and (5) nonlinear pain response to the degree of market decline. Thus, even if a stand-alone endowment manager were somehow unconcerned with these five parameters, these sensitivities would have to be considered in the endowment's asset allocation and governance, to have any hope of leaving the university in its comfort zone. These issues should be considered in the endowment's risk framework to give the university it serves the best chance of achieving its strategic goals.

Risk Framework (Risk Management Process)

Consistent with the TEAM strategy and vision of our administration and governance body (investment committee), we have developed a comprehensive, formal risk management framework at The University of Chicago. Risk management is not a "one size fits all" solution, but is customized and essential to the enterprise's needs and goals. But every risk management framework should address the following key factors:

- ◆ Risk governance
- ◆ Risk identification and measurement
- ◆ Risk infrastructure
- ◆ Defined policies and procedures
- ◆ Risk monitoring, mitigation and management
- ◆ Communications (more than just reports)
- ◆ Strategic analysis and integration

To start building this, initially the critical elements are governance, strategy, and resources. Knowing how risk management will fit into the governance structure is a key starting point. Is there strong sponsorship for risk management at the committee or CIO level? What questions is the committee trying to answer, and what are their key risk concerns? Is there a desire to truly integrate risk into the investment and decision process, or is it more of a lip service exercise, or 'checking the box'? Is the intended outcome just a few reports, or a 'risk culture'? Risk management is deeply embedded in the governance here—our committee was driving this top-down, and hired a CIO of like mind (Mark Schmid had previously built comprehensive risk programs at Boeing and Chrysler).

But while our commitment to risk management is full immersion, that won't necessarily be the right answer at other endowments. You need a strong sense, going in, of what you want out of your risk program. You should be open and honest about this, even if the answer is not politically correct.

Related to governance is the element of strategy. What are the investment strategy and philosophy, and how will risk management inform and integrate with them? At the outset you need to develop some idea of this, even if the mechanics aren't worked out. At The University of Chicago, even before the first risk hire, it was decided that there would be an enterprise view of endowment risk, that we would take a 'risk view' of asset allocation, focus our attention on key risk drivers, and integrate risk management fully into investment management. After the fact this all seems fairly obvious

to us, but in the initial months much of this seemed both controversial and unreasonably ambitious.

The other foundation element is resources, and that can't be well addressed until the elements of governance and strategy are worked out. Consistent with the goals and vision of risk management that you develop, you will need to commit the people, dollars, systems, and time to implement that vision. Risk management is hard work, requiring both deep and broad expertise. It is not easily done on a shoestring by having 'everyone' do a little bit of it, or by expecting the boss to do it in his or her spare time. The resources can be internal, outsourced, or a mix of both (an internal risk manager using a vendor system is quite common)—but a keen focus and clear responsibilities are needed. Consistent with our deep dive as determined in the governance and strategic version of risk, we have invested heavily with an experienced risk team representing about 10% of our manpower, and have pursued a fully internal approach, built mostly from scratch since the new risk team's 2010 arrival.

Risk Governance

Governance (the top-level system of structures, rights, duties, etc.) is a very broad topic, and that is true as it relates to risk as well. At an informal level, as discussed above, the key issues surround what the governing body hopes to get out of risk management, and their key questions and expectations. This is dynamic and subtle, and best developed by frequent and open discussion with the board or committee. At a formal level, risk governance is about the expectations and concerns of the board, as formalized in authorities, limits, and other constraints on risk. Best practice would result in a clear statement of risk appetite or tolerance from the board, along with high-level controls consistent with that risk appetite.

What does this mean in concrete terms? You must tease out, by a variety of methods, what are the key risk drivers of your investment performance, and what are the key concerns of your governance body, especially as relates to risks facing the overall enterprise. Then you have to figure out what to do about those concerns and drivers. At The University of Chicago, this meant deep analysis and frequent discussions with the investment committee—on top of our four 'informal' strategic guidelines mentioned earlier (enterprise approach, risk view, focus on drivers, and integrated approach) which were developed early on. The first step, identifying our primary risk driver, also led us to how to allocate with a risk view, formalize a risk appetite, develop primary risk controls to effect that, and integrated the risk discussion with the enterprise strategy (TEAM approach).

A variety of statistical analysis showed us that the clearly dominant risk factor in our returns was global equity market returns—this isn't surprising, and is probably the case for most endowments. The investment committee and staff agreed quite early to use the amount of overall global equity risk exposure as the basis of our risk view of the endowment—the 'language' we would use to allocate assets. This board committee then took a detailed look at the endowment's contribution to the university's risk

(and wealth) profile, and developed a clear risk appetite couched in terms of global equity exposure, in an enterprise risk context. This complex governance process was described in detail in the TEAM article referenced above, and resulted in a stated risk tolerance with a global equity factor beta range between 0.70 and 0.80. This range then became our *primary* governance control, with a central target of a 0.75 beta to the global equity factor. So, we know what our main risk is, and we have determined how it affects the university as a whole, how much of it our governing body is willing to take, how to translate this into numbers, and how to use it to both control and communicate our investment allocation and risk, and how to integrate this into our investment decisions.

While this all was a huge leap forward, the risk governance, risk appetite, and key driver risk control work did not stop there. We continued to analyze and discuss a variety of risk drivers with our investment committee, and discovered a few other dimensions in which to narrow the stated risk appetite, and develop secondary governance controls on our investment authority. The following table outlines the four most important risk drivers from an enterprise perspective, and the risk controls that became part of our formal risk governance process. (*Figure 1*)

Although each of these risk drivers could merit its own article, we'll just briefly mention each one here. Our governing body and investment team agree that by far, the most important investment risk to the university's future success is what happens globally in equity markets. How we measure, monitor and manage this will be addressed later; but we calculate this exposure weekly and communicate it to our board committee and university leadership on a weekly basis, and discuss it with them at each meeting, frequently revisiting the risk appetite decision (the target has always been between 0.75-0.77, and the range has been expanded slightly as shown in the table).

The three secondary risk drivers shown rose to the level of a formal governance control not because of a stand-alone endowment risk concern, but more from their interplay with non-investment risks within an enterprise context. Liquidity is not a major concern for the endowment in isolation, but the endowment's need to maintain or even increase substantial cash flows to support university operations in a future crisis, means that liquidity must merit close analysis and monitoring. The investment committee examined this in detail with the staff, and decided to set a long-run target of 35% in private fund structures. In terms of convexity, endowments tend to have strongly accelerating betas in market declines (nonlinear beta

exposure), for a variety of reasons that we are not immune from. Whether or not negative convexity makes sense as a risk profile for the endowment alone would be an interesting argument, but that is moot. The university feels tangible pain from a market drawdown in a quite nonlinear way, especially in view of enhanced operating and financial leverage coming out of the last crisis, when a very successful strategic program of 'enhanced eminence' was implemented here. While we cannot linearize the university's inherent utility function, we can at least take steps to linearize the endowment's return profile, and have been mandated by our governing body to do so. Finally, from a university-wide perspective, we are fairly levered relative to the size of the overall balance sheet, so to mitigate enterprise risk, we prohibit leverage at the endowment level, and additionally hold cash to 'fund' implicit derivative leverage, plus a minimum level of Treasuries as a partial offset, as well. The fact that these three additional risk drivers only merited formal controls due to an enterprise-wide risk assessment is a testament to the sincerity of our governance body's commitment to a true enterprise risk framework. Helping the board develop a clear statement of its risk appetite across the most important dimensions is extremely challenging—but this is among the most important things we will do as investment managers.

One of the many side benefits of all this risk governance effort and controls is the extent to which it reinforces taking a risk view in our investment thinking and allocation decisions—which has been one of the key governance precepts from the outset. When we are considering an investment, say, a hedge fund, we no longer primarily think of it as a hedge fund. Instead, we think of it in terms of its characteristics across those key risk dimensions defined above (and to a lesser extent, other risk dimensions not assigned their own governance controls). Fund A has a 0.32 beta, but exhibits strong negative convexity, has tough lock-up terms and a potential 20% side pocket, and exhibits significant inherent leverage. Certainly we might have considered all these characteristics to some extent before developing a risk framework, but now these are front-and-center in how we think about, describe, and decide on adding this bundle of risks to our existing risk portfolio. The governance process has essentially forced us to think the way our investment committee wants us to think. This makes them much more comfortable not only with the resulting risk profile, but also with giving us the full authority we need to do our jobs as investment managers, as they have resolved a good bit of the inherent 'agency problem.' This further ties into the discussion of risk culture, below.

FIG. 1

Risk Governance

| Risk Drivers of TRIP Investment Returns | Governance Control |
|------------------------------------------|------------------------------------------|
| (Primary) Global Equity Risk | β target 0.77 (0.70–0.85 range) |
| Liquidity Risk (Illiquidity Premium) | Glide path to 35% 'Private Asset' target |
| Short Optionality Premium/Convexity Risk | Offset with TRIP Protection Allocation |
| Portfolio Leverage | No Explicit Leverage |

The following graphic shows, broadly, the extent to which the investment process is integrated with the risk process—one of the key precepts decided on early in the development of our risk process. At the top and middle is a representation of the governance approach we have just finished discussing. Other sections relate to other facets of the risk management framework, to give you a sense of the size and scope of our fairly new but extensive risk management program, as we have designed it. We'll turn now to some of these other aspects of implementing risk management within an endowment. (*Figure 2*)

Risk Infrastructure

People and systems are needed to implement the risk framework, track risk exposure, perform analysis, and execute the other steps to assess risk and better inform investment management. We put an experienced risk team in place in 2010—two senior risk management executives from Morgan Stanley, Mike Edleson and David Warn; and an experienced derivatives trader, Michael Suh at first, and more recently, Ari Paul. We went to work immediately on a database and data model, risk capture, risk mapping procedures, performance and risk model, scenario engine, enhanced reporting capabilities, and all other aspects of the risk framework. The goal was to nearly match the functionality and sophistication of the risk framework of the top hedge funds, but with a much smaller resource commitment.

The risk infrastructure should be robust and expandable, and able to address current and future governance concerns, and answer questions of importance to the investment decision process. Perhaps some of the lessons we have learned are obvious; these would include: knowing what you need and want before you build or buy something; don't underallocate time or resources to data—your data challenges will be bigger than your analytical challenges; all investment risk and portfolio management inputs and processes should be as single-sourced and straight-through as possible; and, while a spreadsheet might be a nice tool to get the ball rolling, never build your risk system on a spreadsheet, as it is too brittle to expand the scale or scope of your solution.

Risk Models & Metrics

There are two key decisions involved in putting together a risk model to describe, monitor, and aggregate the risks of an investment portfolio. First, how will you describe the risks of a portfolio asset—by its historical return profile, or by assigning (mapping) an asset to a like asset, or proxy? Each has merit, but we chose a proxy or mapped representation, rather than using the assets' own returns. Second, how will you assemble these asset risks into portfolio risk measurements? Two popular approaches here are *historical* (replacing each asset by a time series of returns or proxy returns) or *parametric* (describing each asset by summary statistics identifying its relationship to other assets, such as covariances). We use a historical approach, not only due to its increased flexibility, but also to better capture fat tails in returns, and other complex relationships (e.g., shifting correlations) that aren't addressed in a parametric approach.

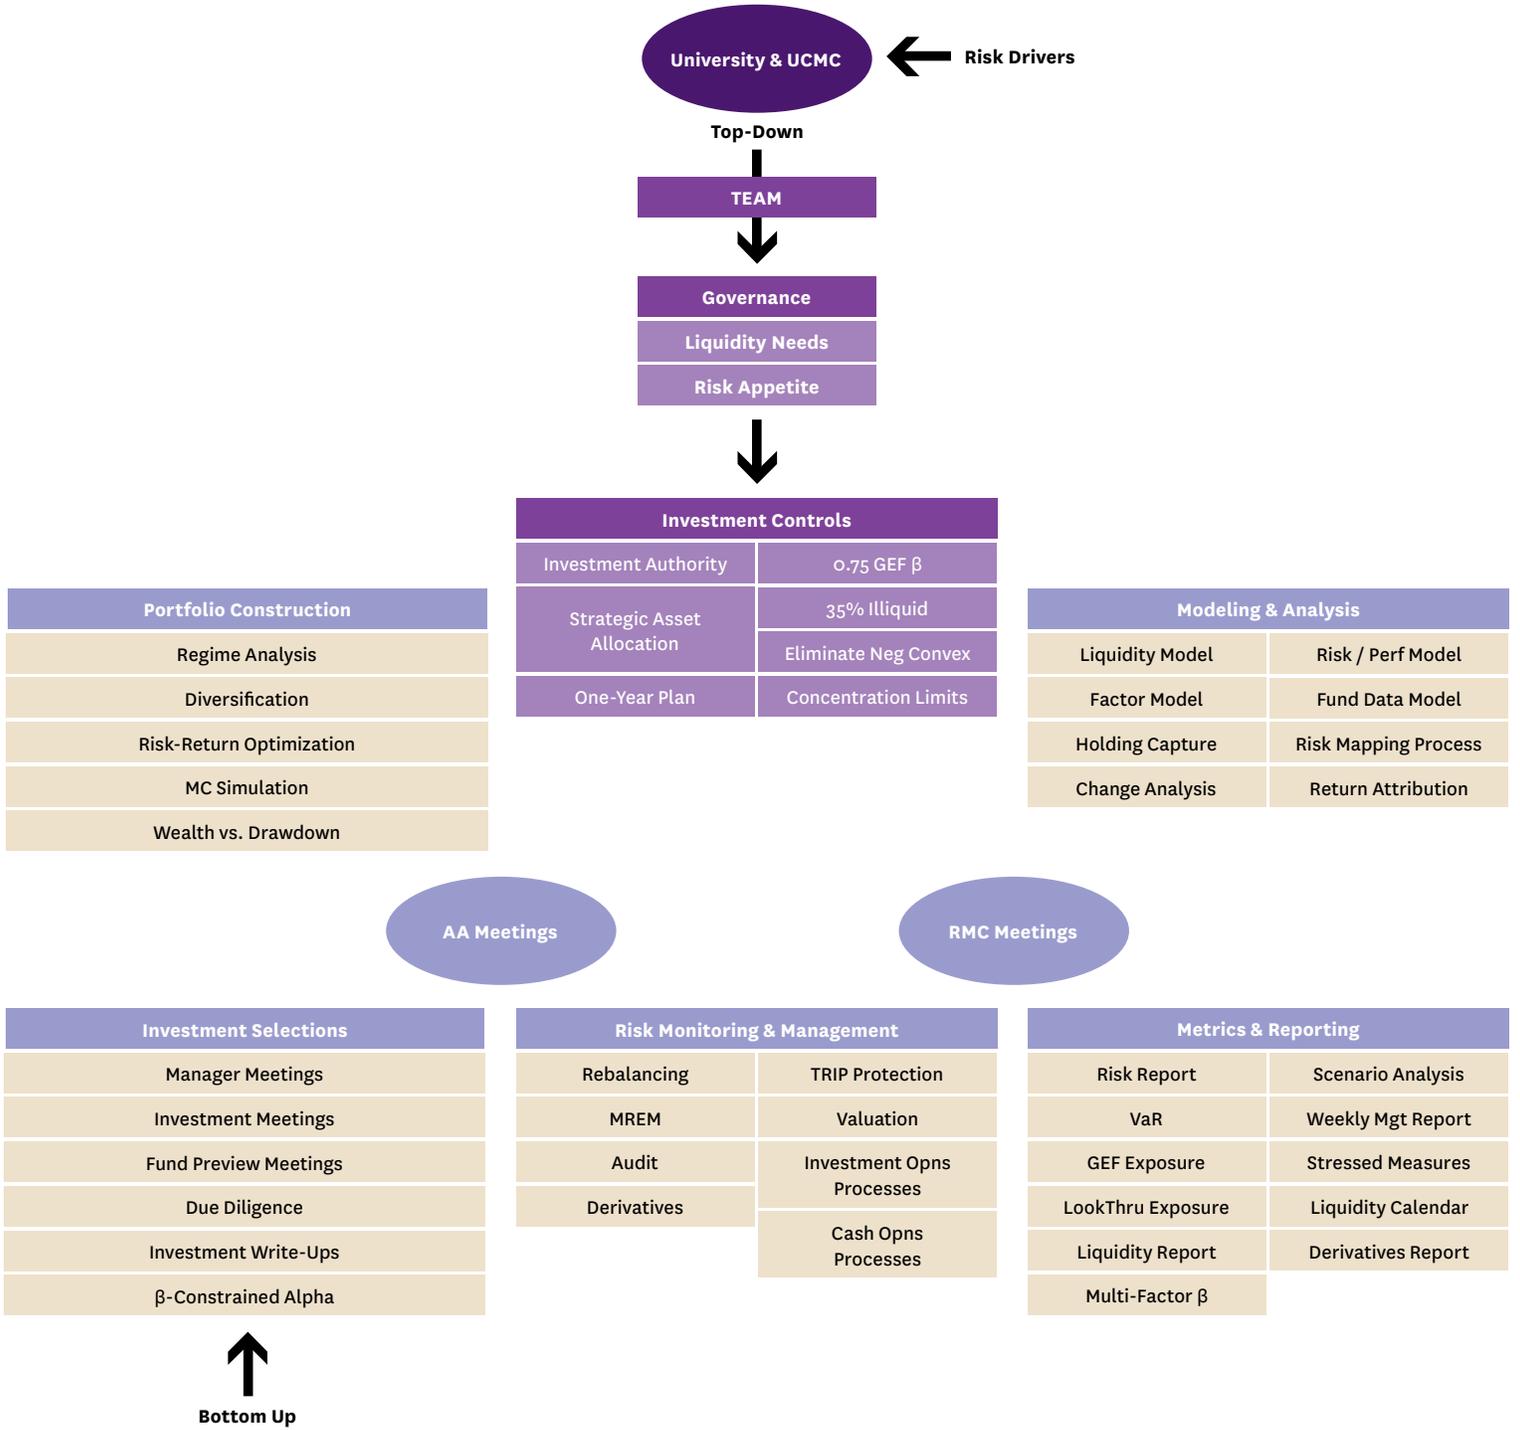
The mapping of proxy returns for risk representation is probably the hardest piece of the puzzle. Every investment we make needs to be represented by some public proxies that have a time series of daily returns. This might mean representing a public equity fund by a weighted group of country-sector-cap indexes, or even the stocks themselves, if we have that information; or a hedge fund by a weighted list of hedge fund strategy indexes, or more granular indexes to match a fund composition report. A private equity fund might be represented by indexes or comparable stocks that are traded publicly. The idea here is to pick up the current risks in our holdings as they relate to current market moves, and not rely on stale valuation data that might not be updated for a quarter or longer. We have formal portfolio capture and risk capture processes to support this, updating holdings' information on a regular basis. The risk team works closely with each asset class every quarter on the risk mappings for all its holdings, and summarizes and gets approval for those updates at both the director level and the full endowment level. This involves almost 4,000 holdings, mapped into the return vectors of over 1,000 proxy assets. Then the risk representation of the entire portfolio is easily aggregated as the dollar-weighted sum of all of those mapped proxy daily returns. While this is an intensive process, much of it can be outsourced (for a price). But we find that the deep dive into the portfolio's risk composition on a regular basis pays substantial dividends in terms of understanding and monitoring our portfolio.

Thus, a single time series emerges for endowment, which represents the hypothetical returns of the portfolio, if it had been structured through time as it is today. We use 10 years of data for each proxy asset, so we end up with a 10-year hypothetical daily time series as the endowment portfolio's risk profile. Using that roughly 2550-days of history, it is easy to compute just about any risk statistic(s) you might imagine. Regressions vs. general market or other factor indexes will give you a panoply of beta measures; VaR is calculated simply by ordering the returns vector; scenarios can be evaluated by selecting the appropriate subset of returns; by block-bootstrapping the return vectors, Monte Carlo analysis can be applied to the portfolio, asset, holding, or any other level; shorter periods can be broken out to investigate the risk sensitivity to changing regimes. This—historical hypothetical returns, based on mapped representations—is an extremely flexible approach, because the math can hardly be made any simpler.

Another benefit of this approach is real-time performance tracking or prediction. With a performance model built atop the data model and proxy risk mapping described above, we use proxy returns to give us current daily predictions of where our portfolio returns are. When actual returns are reported (later), the performance model provides a basis to identify and investigate 'return surprises,' and a useful means of deeper return attribution.

The risk analytics that we model are focused on supporting two important goals: aligning the portfolio with our governance guidance and supporting the investment decision process by better understanding the risk-return characteristics of our investments. We see the global equity factor beta (GEF β) on a regular basis, to ensure

How the Investment Office Manages Risk



our primary risk exposure is in line with our investment committee's stated risk appetite. But we can look at our risks in a variety of other helpful ways as well—some sample measures will be shown in the *Reporting* section below. We can look at metrics to track our liquidity position in a number of ways as well.

Liquidity modeling is another major initiative for the purpose of taking a 'risk view' of our investments. We built an extensive data model to describe, quantitatively, all of the liquidity characteristics of both the normal fund investments (like hedge funds) and the private investments (like venture funds). While the data model alone was a massive undertaking, we built and integrated two liquidity models (one focusing on private funds, one on public) on top of this, and also integrated these completely with our risk and performance models above. The ability to do liquidity planning and 'what-ifs' for next quarter and more than a decade out into the future provides considerable transparency and comfort at a governance level, and has led to better decision making in terms of capturing illiquidity premia.

Risk Monitoring & Management

While risk monitoring and management may be fairly obvious, it is also probably the most difficult facet of the risk framework. Actively monitoring and managing risk exposures requires us to pull together all the other aspects of the risk framework, effectively integrating everything mandated by our risk governance. Much of this involves ensuring that our risk exposures are in line with our governance controls, planning so that this will still be the case in the future, and taking action where needed to realign risks with appetite.

For our GEF β target, this means reviewing the risk model to see if we are still close to target, then estimating future changes to our β by integrating information from the performance and liquidity models, and deciding whether to use derivatives or rebalancing investments to move toward target. For liquidity, the short-term action is about mitigating risk with investment/redemption decisions, attempting to be proactive and not reactive. Longer-term liquidity risk management is tougher—trying to find the balance over the next quarter or year between capturing risky illiquidity premia, and staying on course towards our governance target. For convexity risk, we manage a strategic program with a small protection (positively-convex) allocation, and closely monitor whether this is helping to linearize our return profile to mitigate wrong-way risk, as mandated.

But keeping close tabs on our risks and ensuring they're always in line with our risk appetite and governance controls is not nearly enough; it's just the starting point. Qualitative risk management concerns are extremely important, and require focus and resources as well. Integrating the risk process and seamlessly weaving it into the investment decision process is where much of the risk management and monitoring effort goes. Thus, a good bit of our risk framework is 'operational'—and is extremely important both in allocating and mitigating our risk.

While operational risk management can be easily overlooked in a risk overview, it holds an important role in our investment and monitoring processes. We've created tools to help our asset classes navigate our process and hired third parties to help supplement our internal process. For our office, operational risk management starts

FIG. 3

Topic List from FY2011 Risk Management Committee Meetings

| | |
|-------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Endowment Risk status update—Risk Report review | Derivative Risk & performance reporting; update on Derivs' performance |
| Proposal for creating DB / Data Warehouse | Review of annual audit |
| Derivatives—Proposal, Use & Controls | Proposed changes to GEF β regression methodology (extend data time frame, move to longer periodicity, away from weekly blocks, and use of our new MegaRegressor) |
| Debrief of Operations Offsite | Detailed brief/discussion on case for TRIP Protection |
| FY2011 Manager Survey | Risk Framework for operational risk |
| Risk Drivers of Return, Principal Components, and GEF β intro | Review of total rebuild of Management Report automated process & tools |
| Development of Risk Framework | Business Continuity Planning (BCP) |
| Discussion of Manager Risk Evaluation & Monitoring, creation of investment action policy around MREM | Proposal to collar significant profitable single name position/concentration |
| Introduction to Tail Hedging | Review of new Operations Manual draft |
| GEF β analysis—Empirical vs. Mapped Betas | Proposal for Board for TRIP Protection (including volatility in strategic asset allocation), and approval for selection & implementation plan. |
| Risk-based Performance Attribution | Review of large manager's operational due diligence conference |
| Updated Factor Risk mappings for Real Assets | Manager Risk Evaluation & Monitoring, each asset class on a quarterly cycle |
| Changing risk mappings to incorporate fund-level leverage | |
| Moving to a Beta risk-based allocation governance control | |
| Time-matched risk & return, serial correlation and impact on volatility—comparing 'reported' performance risk to real economic risk | |

on an individual investment once an asset class makes the determination to potentially pursue an investment opportunity. At the onset, our operations group is included in the discussion and diligence process. During the investment due diligence process, third parties are engaged to assist us in confirming the integrity of a management team. Our operations group, lead by an experienced investment operations manager, Loryn Mischke, reporting to COO Pat O'Hara, stays engaged as part of the process and through updates to our tracking systems. We also have an outside consultant for deep-dive due diligence on many of our funds, and legal document review; both of these mitigate a variety of investment risks, and are a costly but value-added part of our process. The Operations team is responsible for the execution and data capture as an investment is completed and added to our portfolio.

For our investment holdings, our operational risk management includes cash flow management, monitoring, value and holdings' capture, and auditing. The cash flow management and monitoring occur on a daily basis, with the auditing of our investments effectively occurring quarterly with the production of our Investment Committee report book. Our final valuation and audit check occurs before completing the University financial statements. Annually our operations group keeps in contact with every manager by sending each of them a questionnaire. The questionnaire informs us of any changes (be they operational, legal, or otherwise) the manager might have made throughout the year. Changes that require attention are then followed up by our operations group and any changes are made to our processes and databases. The operations group also works closely with the risk group to ensure the integrity of the investment data, resolves issues caught by a host of error-capture routines in our performance and risk models, and maintains and improves a data warehouse to support a variety of investment and risk analyses.

In "A Framework for Manager Selection," by Matt Stone, Prakhar Bansal and Kate Carder in the November 2012 *NMS Exchange*, our public-side team described our investment process, which culminates in a final approval meeting and a formal investment recommendation of 40-50 pages, which includes an extensive section on risk management. As part of the monitoring of ongoing investments, we track and rate all of our managers as 'red-yellow-green' on over 20 investment risk dimensions, which emanate from our investment philosophy and risk governance. This 'manager risk evaluation & monitoring' cycles through each fund twice yearly at our monthly Risk Management Committee meeting, and we spend time to discuss any funds in depth that are rated red, or had a rating change. This helps to formalize active engagement with and monitoring of our fund managers, and has noticeably reduced surprises.

The RMC, or Risk Management Committee is the lynchpin of our risk framework, a key extension of our risk governance, and a forum for management discussion of important risk issues. The RMC includes our CIO and all Managing Directors, and is co-chaired by the Chief Risk Officer and Chief Operating Officer. Each 90-minute meeting begins with a review of current risk

reports and manager risk evaluation & monitoring for 1-2 asset classes. Alternate months focus on liquidity risk and planning. The other topics covered are quite varied; while the list from the 53 RMC meetings is too long to present, here is a sample of the topics from the first year's agenda. (*Figure 3*)

Risk Reporting & Communications

The communication of critical risk issues must happen across the investment organization, so that risks are reviewed and discussed as a standard part of every investment decision process. This starts with reporting, but it cannot end there. We have a modular, automated suite of reports that is completely integrated with our official operations data (single-sourced data is absolutely critical for analytic solutions, but very difficult to engineer). So, as long as data inputs are current, we can run any and all reports with almost no lag. An entire set of reports is ready every Tuesday (and each month-end), with the capability to generate answers to other ad hoc risk questions or custom reports same day. With so much analytical output, we find it important to use dashboards or 1-page reports to focus attention on the most cogent investment risk metrics.

Our primary dashboard, the weekly Management Report, is organized by theme. The top is performance focused, using our proxy model to provide a market-based estimate of what our asset class and portfolio returns will eventually be reported to be. The middle is liquidity focused, categorizing all of our cash flows and using our liquidity model to track current and future unfunded status, projected cash flows several years out, and estimating future progression of various liquidity metrics and coverage ratios, with a particular focus on our governance target for illiquid investments. The bottom is risk focused, providing a variety of volatility and beta measures, especially highlighting our positioning relative to our primary risk governance target, our GEF β . We also look at our risk exposure to a variety of risk factors (besides equity, we examine credit, real estate, commodities, interest rates and inflation), regressing our portfolio on the factors singly and in combination, and also providing a less-statistical 'look-through' measure of where our exposures really are, as opposed to an asset-class bucketing. This report goes out to our governing body and to the university's key leadership each Tuesday.

Other risk reports go much deeper, and are generally distributed monthly, even though they are available each week. We try to keep our reports theme-based, grouping related tables and graphics by the questions they are answering. These get reviewed at governance forums like the Investment Committee and Risk Management Committee meetings. Some of our most useful 'reports' are actually diagnostic tools, such as our 'difference report' which does a multi-dimensional analysis of how money was made and lost over the past week/month. Not only does this report allow us to catch nearly all data or model errors that might sneak through, but it also keeps all our asset class heads attuned to that week's developments, highlighting unexpected results. It tightens up the process of keeping in touch with our investments and

relating assets' risk to return, on a regular basis. We get to frequently correlate our preconceived notions of what we *think* should be driving our investment returns, with what we actually observe to be driving them—and we discuss what we learn from the differences.

Risk Culture

When risk management is truly integrated at all levels of decision-making in the investment process, and 'risk' becomes as common in the daily parlance as 'return,' it could be said that you have achieved a *risk culture*. Our experience is that a risk culture makes investment managers' lives a bit harder, but makes the investment process and results *much* better. Threading risk issues as a matter-of-course into discussion and decisions produces better awareness and results than checking a few risk boxes as a separate afterthought, and in turn should be much better than ignoring risk issues altogether. Strangely, though, a risk culture is not the natural state of humans and organizations; we'll leave it for the behaviouralists to figure out why that is so.

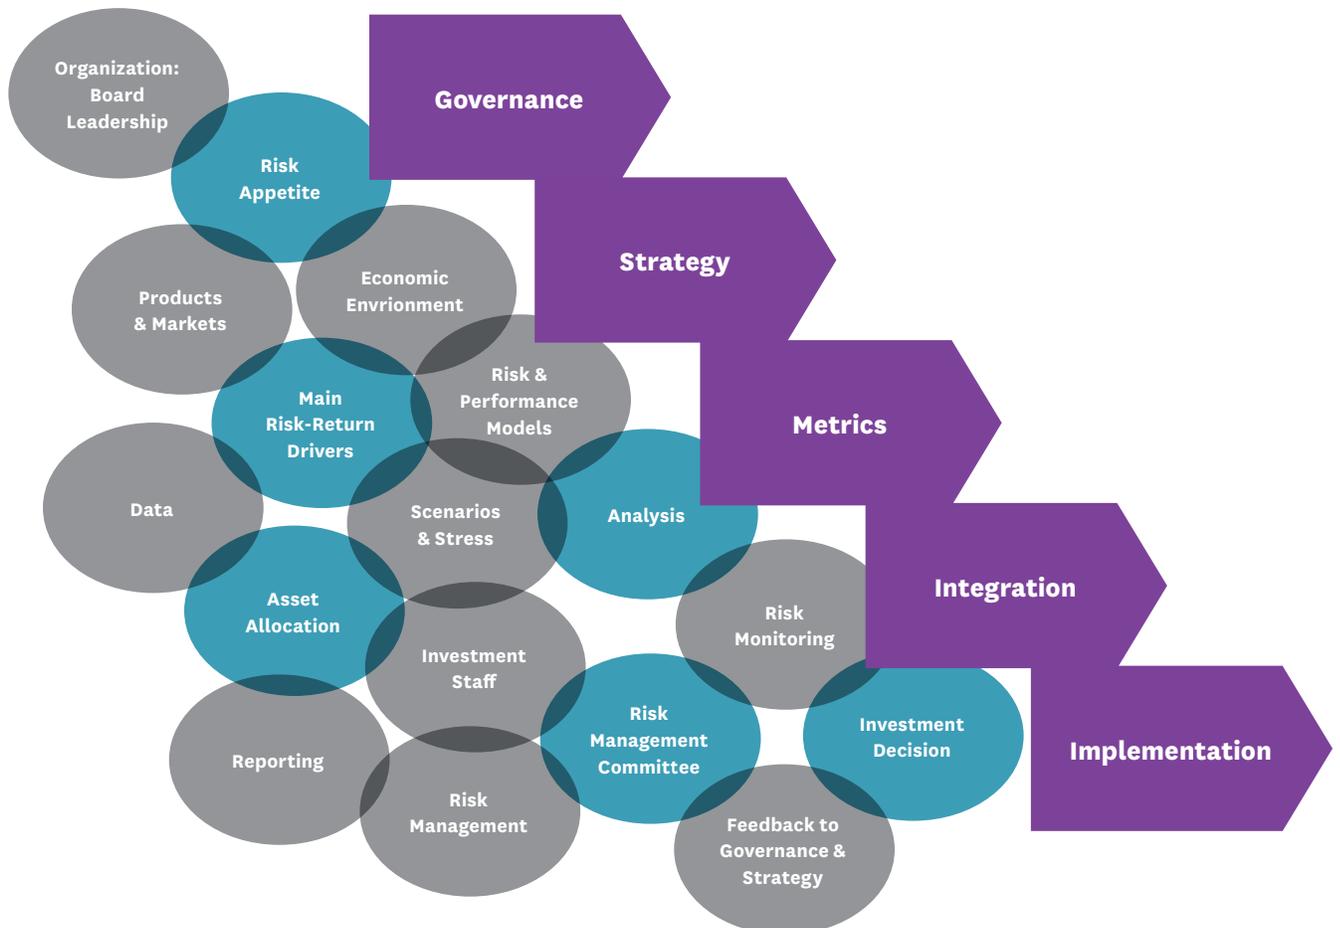
And it's not just about reducing mistakes and surprises. Armed with a risk tolerance, risk budget, and a risk framework, the investment process can hardly help assessing trade-offs between risk and return. Once risk is part of the calculus of every investment decision, these decisions

become much more rigorous, and naturally focused on squeezing as much return out of each 'unit' of risk as possible. Ironically, by focusing more on (limiting) beta, the end result becomes an intense focus on alpha. We actually think and talk in this manner about our investments now. Previously, many investments with slightly lower return targets but considerably lower risks were bypassed as a matter of course. But investing with a governance target/range around our risk has been strangely liberating—the board knows the portfolio will be in their comfort zone and trusts the staff more; and, it has as a practical matter opened up new, less-crowded investment opportunities to us. The discipline and trade-offs inherent in tacking to our risk tolerance have now pushed us into investments with higher ex ante return-per-unit-risk profiles. From our university's standpoint, this is increased value from better decision making. This is a fairly obvious (in retrospect) but highly underrated benefit of operating with a risk culture.

There is a subtle additional advantage, beyond the implicit selection of the highest risk-per-return investments. Suppose you selected a basket of investments that exceeded your risk tolerance. There are two ways to trim that to budget—you can do less, cutting back your marginal investments, or you can hedge using market instruments. This choice,

FIG. 4

Key Steps to Get to Enterprise Risk-Based Investing



always available to you, has you evaluating your active investment directly against the passive market hedge, which is essentially the risk-return benchmark. Effectively, this is properly transfer-pricing the cost of risk, even though there are no complex calculations to do so. You can get to your risk budget actively, or passively—which is better? Said differently, you are in essence evaluating your marginal investment against the market to see if your active investment choices are adding value, or destroying it. This is what we are supposed to be doing as managers; but only with a risk framework guiding our hand is this rigor made innate to the decision process. This approach will have you capturing as much active ‘alpha’ as you can harness. In this sense, risk management is very much an *offensive* weapon in your investment arsenal.

A team with a good risk culture, where risk is integrated into every key strategy and decision, should produce a high alpha portfolio with less negative convexity, compared to a traditional team with little risk focus. As long as the parent organization is perceptive and honest about its risk tolerance, this should in turn produce higher value for the assets’ owner as well.

Risk culture is not simply about “If you build it, they will come.” Participation and discussion is key to achieve these benefits, as is transparency about the process. While there are dedicated positions focused on risk, the risk culture results from integration, and not from an attitude of “You do your risk thing, and leave us to do our investment thing.” The risk team is involved as a key participant in all meetings with our fund managers, and in every step of the investment decision process. The investment and operation teams are not only part of the Risk Management Committee, but are also heavily involved in many practical facets of implementing the risk framework. Among our Managing Directors are investment professionals with differing philosophies, but the common language of risk focuses us in a unified fashion on the goal of adding as much value to the university as possible.

Building Risk Management

We know this isn’t for everyone. It is, admittedly, simpler to keep endowment goals separate from the university, and/or to focus on return without engaging too much on risk issues—and some will go that route. For those that would aspire to capture the benefits of a risk program, though, here are a few parting words on that prospective journey.

It seems daunting, almost scary to imagine building all this from scratch. It’s sort of like walking, though. You get an idea of where you’re going, and maybe why, and then it is just one-foot-in-front-of-the-other. It’s a journey, not a destination. OK, it’s a destination too, but you don’t have to get there in one day or even one year. And some of the benefits start right away—you don’t have to have a 100%-implemented risk framework to move toward an effective risk culture.

What we have found is that there are 5 major areas or steps that need significant attention in order to build a risk program. We’ve discussed most of them above, especially the most important and initial step, governance. For us, each step took about 9-12 months to complete, and you can’t skip a step. But fortunately, you can work on several steps at once; in fact it’s even better to do it that way. From a standing start, it took us a bit less than 2 years to develop our first cut at a complete, fully-implemented risk framework. You can try to speed this up by outsourcing some aspects, but there are still challenges to integrate, say, your data into an external risk system. This will become a part of your investment DNA, so you wouldn’t want to outsource the whole thing even if you could. But system vendors, data providers, consultants, custodians, and others, are all happy to help out with your risk program. (*Figure 4*)

Unlike building a house, when you are done building a risk program, you really aren’t done at all. Each of the 5 steps never actually ends, as the investment process it supports never ends. It gets easier once it is all built out, but risk is an area that is ripe for continual improvement, as it is so central to your key mission of continually adding value.

Summary

What have we learned in last five years of embedding risk management deeply within our investment process? It’s not easy, but it’s worth it. It supplements the rigor and discipline in your process with a keener sense of trade-offs and focus on ‘true’ alpha. A risk framework is not a one-size-fits-all solution; your approach must fit with your institution’s risk drivers and support its needs. As a result you are not going to always run with the pack, so your leadership must be comfortable with that. A sincere commitment by your governance body is imperative for success—it is difficult and ineffective to go halfway down this road while looking over your shoulder. Done well, a risk culture develops over time, with improved awareness, deeper discussion, better decisions, and a more effective investment process.