RiskMetrics Group

Research Monthly

Language of risk

Christopher C. Finger
chris.finger@riskmetrics.com

February 2007

As risk managers, we think of ourselves as traders, physicists, statisticians, and technologists. But it is important also to think of ourselves also as linguists. Traders often have their own languages, unique to the markets in which they participate; investors, albeit in the same markets, may have a different language altogether and supervisory functions yet another. It is the task of the risk manager to bridge the language gaps across these different participants.

But language gaps between traders and risk control are nothing more than regional dialects compared to the gap between altogether different types of risk. So beyond presenting equity factor models as a means of bridging the small language gap between equity portfolio managers and risk managers, our greater challenge here is to take on a type of risk—governance risk—that is beyond our traditional focus, and to consider how this might be integrated with trading risk under a common language.

Factors for risk and more

In a recent article,[1] we presented an investigation of equity factor models for risk and an articulation of our adopted model framework. Our goals are twofold. First, the model should provide good empirical results: it should identify factors which historically explain a significant amount of the common movements in stock prices; moreover, the relationships between stock returns and factors should be robust through time, so that the model is useful not only for describing history, but also for forecasting future risk. Second, the model should provide economic and financial intuition, decomposing equity risk along dimensions that are useful for an investor constructing a single portfolio or for a risk manager aggregating exposures across an institution.

Straumann and Garidi state a third goal, that of parsimony, or a minimal number of model parameters. We might argue that parsimony is a necessary part of both of the goals already stated. Relying on a large number of factors can certainly improve our ability to explain historical price moves, but can in some cases actually impair our forecasts.[2] And there is a limit to how many factors a human investor or risk manager can make sense out of; part of providing intuition, then, is distilling the universe of equities into a manageable number of key drivers. Beyond these key drivers, we attribute the remaining risk as idiosyncratic, or specific, to the individual firms.

We should remark that as a general approach, this is not unlike how we approach bond or derivative

[1] See Straumann and Garidi (2007), who also provide extensive references to the literature on factor models.

[2] Straumann and Garidi present a case where a factor model including seventeen factors, among them ten sector indices, produces inferior forecasts to those of a model with only four basic factors.

markets, though in these cases market conventions play a greater role in how we build our language. Bonds trade off yield curves, with the movements of the curve describing most of what happens to a portfolio, and the balance of the risk attributed to issue-level spreads, or bases. Derivatives on a common underlying trade based on a curve or surface of implied volatility, again with some small amount of risk particular to specific options. In either case, it would be unnatural to develop a language of risk that departed from these conventions. In the equity market, less structure exists, in that there are fewer fundamental links between securities, and consequently language has never been as standardized.

For our model, we have elected to begin with factor-mimicking portfolios. The idea here is that if there is some underlying market dynamic such that stock returns cluster according to an attribute (such as market capitalization or price-to-book ratio), then we can construct a proxy for this dynamic by forming portfolios according to the relevant attribute.

For example, rather than trying to hypothesize how market capitalization is punished or rewarded in the markets, we propose to mimic the underlying size factor. The mimicking is achieved by forming one portfolio of the smallest stocks, and a second portfolio of the largest stocks, and then calculating the difference in return on the two portfolios.[3]

This produces a factor that rises when small firms outperform big ones, and falls otherwise. Moreover, since the stocks in the two portfolios should have more or less the same sensitivity to the broad market, the market sensitivity should cancel out when we form the factor, and the factor itself should be relatively uncorrelated to the broad market. Finally, the mechanism can be applied using any attribute with which we can partition our universe of stocks into long and short portfolios: price-to-book ratio, past performance, or the number of coffee machines in the company headquarters.

Factors formed in this way provide intuition if the attribute used to define them is sensible: a risk manager is certainly interested in the degree to which his bets rely on the outperformance of small stocks relative to large stocks, but perhaps less so in the outperformance of the most highly caffeinated companies. This is where language is important. If a factor represents a type of strategy an investor may execute, or is something a trader may have a view on, then the factor is a useful way to express risk, and a piece of language we should consider adopting.

As to our other goal, since we have not relied on any fundamental model of the market, we have to perform empirical tests to decide whether our prospective factor does have explanatory power. We will come back to this point later, and Straumann and Garidi provide more detail on such tests.

To the risk manager, the two empirical questions we raised earlier (do factors describe historical comovements in stocks? is the relationship between stocks and factors robust?) are paramount. For an investor, there is a more fundamental empirical concern: does the factor carry a risk premium? In other words, does the factor produce an expected return in excess of the market? For all of their use as risk factors, the true reason that the size and value/growth factors have become so ingrained in how we talk about equity investments is that Fama and French, and others, established that these factors do in fact carry risk premia, producing consistent excess returns.

[3] Of course, the sizes of stocks will change over time. We thus select new portfolios at some regular interval, typically annually.

For our basic set of factors, we choose those that have at once succeeded as a common language and demonstrated their worth empirically: the broad market return less the risk-free rate (RMRF); the aforementioned size factor (SMB); the value/growth (HML) factor, formed by constructing portfolios based on price-to-book ratio; and the momentum factor (MOM), formed by constructing portfolios based on prior six-month returns. The basic model decomposes the risk of a portfolio into risks deriving from each of these attributes, and the specific risk deriving from the particular companies held.

Enter governance

Recently, RiskMetrics Group acquired Institutional Shareholder Services (ISS), a provider of corporate governance risk solutions. Our new colleagues have made us acutely aware of the importance of governance risk in the investment process, and has prompted all of us to look for a language to bridge governance and trading risk. A logical first step is through a governance factor in the framework we have presented. The question is whether such a factor is a useful addition to our language, which takes us back to our criteria: intuition, return, and risk.

From the point of view of intuition, it is hard to argue (especially with our new colleagues) that investors and risk managers would not be keen to know how much they are exposed to governance. They might ask themselves if they have a view on governance: are positively governed firms likely to outperform negatively governed ones over our investment horizon? Alternately stated, is the portfolio manager seeking to align their portfolio with firms which, by virtue of good governance, will produce sustainable returns? If the answer to either of these questions is yes, then both the portfolio and risk manager should ask whether this explicit view is consistent with the implicit views they have expressed through their choice of portfolio.

So there is little question that it is interesting to look at portfolios along the dimension of governance quality. The challenge lies in quantifying "good governance", a decidedly qualitative concept.

Gompers et al (2003) introduce a simple governance index, with the goal of ranking firms along the balance of power between shareholders and management. They identify 24 governance variables, and count, for each firm, the number of variables that restrict shareholder rights and increase managerial power. For example, one of their governance variables is the existence of a classified board, that is, a board of directors with staggered terms, such that not all board members can be replaced in a given year. As a classified board detracts from the ability of a large shareholder to influence the firm's directors, this variable is counted as one that restricts shareholder rights, to the benefit of management.

For our analysis, we will use the Corporate Governance Quotient (CGQ) provided by ISS. The CGQ is similar in spirit to, but more comprehensive than, the score above. The CGQ utilizes a greater number of governance variables, ranging from those describing a firm's charter or bylaws (which constitute the bulk of the variables considered by Gompers) to the actual composition and practices of the board and its audit and compensation committees. Across this large set of variables, ISS has determined weights according to correlations with a number of profitability and valuation metrics. Among the most heavily weighted variables are:

- Composition of the audit committee. The CGQ is impacted positively if the committee is comprised solely of independent outsiders.
- Burn rate, or the cost of granting equity to management. A burn rate that significantly exceeds a firm's industry average has a negative impact on the CGQ.
- Composition of the board. The CGQ is impacted positively if the board is controlled by a supermajority[4] of independent outsiders.

Building a governance factor

Given a governance scoring mechanism, a naive approach to building a factor is to apply the methods cited before for the SMB, with the CGQ replacing market capitalization as the means to sort the universe. Intuitively, we create a factor that represents the outperformance of "good" governance firms to "bad" ones. As always, though, there are more details to think through.

First, we must decide on our "breakpoints", that is, the percentiles at which we apply cutoffs to form our good and bad governance portfolios. Whereas for size, we used the median of the sample as the cutoff, our intuition with governance is that the material difference in returns will come from avoiding firms with very low scores. We opt, then to form our bad governance portfolio from the lowest 10% of firms, and our good governance portfolio from the highest 10%. This is consistent with the approach of Gompers et al, who examined the differences between firms identified as "democracies" and "dictatorships". Looking back at scandalous names, Global Crossing, Tyco, Parmalat, and Ahold, just prior to their troubles, were all in the bottom ten percent, relative to their respective index peers; Adelphia was at the sixteenth percentile; and Enron was just below the median at 42%.

Another detail is to ascertain the relationship between CGQ and other attributes used to form equity factors, to be sure that governance affords a novel sorting of stocks. If the CGQ is perfectly related to, say, the price-to-book ratio, then sorting by CGQ will produce the same long and short portfolios, and the governance factor will be indistinguishable from our value/growth (HML) factor.

Our data sample consists of approximately 5000 US-listed firms. Of these, we have mapped 2800 to CGQ scores. We examine the cross-section of firms on 1 January 2006, and compute the rank correlation[5] of CGQ with three other attributes—market capitalization, price-to-book ratio, and prior six-month return—from which we build equity factors. We present the results in Table 1. It is encouraging that most of the pair-wise correlations are weak, meaning that each of our basic equity factors should describe a different market dynamic. The one pair we should concern ourselves with is size and CGQ. There are methods to explicitly correct for this relationship when building our factor, but these do not seem altogether warranted given the moderate relationship we observe.

In the end, we create our governance factor (GOV) by following our factor construction steps as with any other attribute. In assessing this factor, however, we should bear in mind that the portfolio of

[4] Of at least 75%, with greater impact if the supermajority is over 90%

[5] The rank (or Spearman) correlation is a measure of how similar is the ranking produced by two attributes, without regard to the particular form (linear, exponential, etc.) of the relationship between the two.

Table 1: Cross-sectional rank correlations (in %) of equity factor attributes, 1 January 2006

		1	2	3	4
1	Market Cap	100.00	31.43	-2.51	47.67
2	Price-to-book	31.43	100.00	-1.31	7.68
3	Prior 6m return	-2.51	-1.31	100.00	-2.46
4	CGQ	47.67	7.68	-2.46	100.00

Table 2: Average monthly returns (in %) on US equity factors, January 2002-June 2006

RMRF	SMB	HML	MOM	GOV
0.42	1.09	0.93	0.39	0.10

the highest CGQ firms will have a somewhat higher capitalization than the portfolio of the lowest CGQ firms. Thus, any analysis we perform on the governance factor should account for its moderate bias to be long large and short small firms.

Searching for returns

In their influential study, Gompers et al examined the relationship between governance and both valuation (that is, do better governance scores imply higher market multiples?) and performance (that is, are better governance scores predictors of better future returns?) Our focus is on the second of these.

We first examine the average returns for all our equity factors; these returns are presented in Table 2. Here, the governance factor is not impressive, with a lower return than any of the other factors. Note however that the SMB return is the strongest of all the factors, indicating that small firms outperformed large ones significantly over the time period in question. We should ask whether the poor performance of the governance factor is a result of its bias toward large stocks.

Our next step, then, is to regress the governance factor on our other equity factors. The weights on the other factors will describe in a sense the overlap between the governance attribute and the others, while the constant (or alpha) will represent the excess return that is truly afforded by governance. We present the results of this regression in Table 3.

We note first that the governance excess return is an impressive 76 basis points per month, and is statistically significant at the 95% confidence level. Second, the coefficient on the SMB factor is negative and also statistically significant. This is consistent with our intuition: our governance factor was, by construction, somewhat biased toward large firms. Thus, the governance factor was in a sense dragged down by the weak relative performance of large stocks; after correcting for this, its somewhat meager average return becomes much more impressive.

These results are remarkably consistent with others in the literature. Gompers et al examined a different

Table 3: Regression of governance versus other factors, January 2002-June 2006. For regression coefficients, ⋆ indicates significance at 90% confidence, ⋆⋆ significance at 95% confidence.

	Regression coefficients (%)				
R^2	Constant	RMRF	SMB	HML	MOM
0.32	0.76⋆⋆	2.18	-26.65⋆⋆	-34.56⋆	-13.03

Table 4: Factor regressions for 2895 US stocks, January 2002-June 2006

	Proportion (%) significant at 95%					
	Const	Mkt	SMB	HML	MOM	GOV
Without GOV	7.1	51.1	23.5	12.0	16.6	–
With GOV	6.0	51.1	22.4	11.4	16.3	8.2

time period (1990–1999), and a smaller sample (between 500 and 1000 of the largest US stocks). As explained before, their governance index was more narrow than the CGQ, but their breakpoints in creating their governance factor correspond well with ours. In a similar regression exercise to ours, they saw an excess return of 71 basis points per month, which was significant at the 99% confidence level. Further, they observed the same significant negative coefficients on the HML and SMB factors.

Explaining risk

As encouraged as we are by the evidence of governance-related risk premia, our interest does run deeper: we need to assess whether a governance factor adds to our ability to explain broad cross-sections of equity returns. As we mentioned earlier, we do not model the fundamentals of the equity market, and so it is our task to test empirically whether governance adds value as a factor. The low goodness-of-fit in Table 3 is cause for optimism. It is clear that the governance factors describe something different from the standard factors; the question for now is whether this something is useful.

We examine approximately 2900 US firms for which we have four years of return history. This is not the same universe as we used to construct the governance factors; some of the firms tested here do not have CGQ scores. Importantly, it is not necessary for a firm to have a CGQ score here, since we are only examining the empirical link between returns and the CGQ factor we have already created.

For each firm in our sample, we regress the equity returns against our four basic factors, and then against these factors plus the governance factor. We summarize our results in Table 4. We note first that when the governance factor is included in the regression, it appears statistically significant for over eight percent of the stocks in our sample; this is less of-

[6] In all of our analyses, we also considered a governance factor where we implemented an explicit adjustment for size, as outlined

ten than for the other factors, but is comparable at least to the frequency with which the well-accepted HML factor is chosen.[6] We should take some care in interpreting these proportions: since we have performed a large number of regressions and only asked for 95% confidence, we would have expected five percent of the regressions to show up as significant even if we were analyzing random noise. Over this many regressions, however, the difference between the 8.2% proportion of stocks with a significant governance loading and the 5% we would expect from noise is great enough for us to conclude that the governance factor does add to our analysis.

We also examine how the inclusion of the governance factor influences the significance of the basic factors. The size factor was most impacted, as it appeared significant for roughly one percent fewer stocks. Interestingly, there was a similar reduction in the proportion of stocks for which the constant appeared significant. To assert that the constant is significant is to say that a specific stock produces returns in excess of those explained by the common factors. Thus, for a number of stocks in our sample, what appeared to be an excess return is in fact a return explained by exposure to the governance factor.

Lastly, we tabulate the number of firms for which the increase in explanatory power obtained with the governance factor is more than we would expect from adding any arbitrary new variable to the regression. Examining the adjusted R^2 measures for the individual regressions, this is in fact the case for 38% of the stocks in our sample.

At the portfolio level, the factors give us the ability to track where risk is coming from. In Figure 1, we analyze a portfolio of 100 stocks, and track the decomposition of risk through time. For this particular case, we see that the market and the size factor dominate the risk, but that there is a significant portion of the risk attributed to the governance factor, particularly in the latter part of the time period. A portfolio manager might be cognizant of his bets on both the market and size factors, but may be surprised to find, in this case, that a non-trivial amount of his risk is attributable to a bet on the market rewarding poorly governed firms.[7] Seeing such an exposure, a risk manager might consider probing for "headline risk" by stressing the governance factor, thereby assuming that the worst governed firms fell substantially.

So what have we done?

From a narrow point of view, we have added a useful factor to our model of equity risk. That there is a marked discrepancy in the returns of the firms with the best and worst governance scores is reason enough to be interested in this dimension of a portfolio. And the gains in explanatory power that a new factor brings mean that risks or returns that were previously attributed specifically to a firm are now appropriately identified with governance.

From a broader point of view, we have identified a case where two seemingly unrelated types of risk—trading risk and governance risk—can be assessed in a common framework. This is a small step, but one on a path to what we hope is a more comprehensive language of risk.

in Straumann and Garidi. The size-adjusted governance factor performed slightly worse in tests for excess return, with a constant of 20 basis points, which was not statistically significant, but slightly better in the individual stock regressions, where it appeared significant for 9.5% of stocks.

[7]The sign of the coefficient on the governance factor is not indicated in the figure, but is in fact negative.

Figure 1: Portfolio volatility decomposition based on rolling 24-month regressions

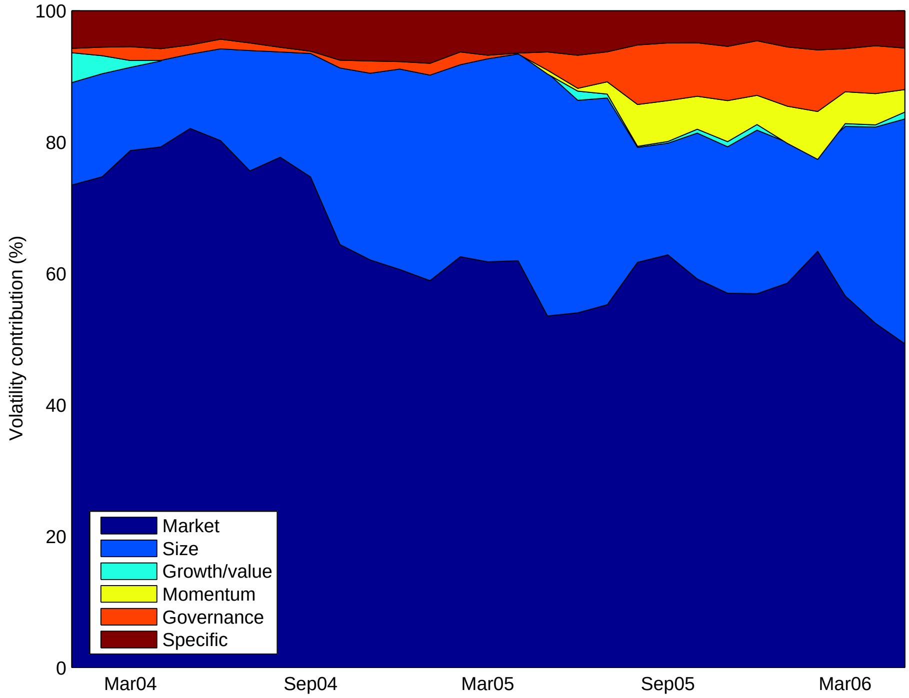

Further reading

- Cheng, D. and Wu, Y. (2006). Evolving corporate governance and equity prices: The recent evidence. Working paper. Institutional Shareholder Services.
- Gompers, P., Ishii, J., and Metrick, A. (2003). Corporate governance and equity prices. *Quarterly Journal of Economics*, **118**: 107–155.
- Straumann, D. and Garidi, T. (2007). Developing an equity factor model for risk. *RiskMetrics Journal*, **7**(1): 89–128.