Declining Active Risk in Japanese Equity Portfolios

Since the collapse of the Internet bubble, many Japanese portfolio managers have observed a surprising contrast between trends in tracking error and market volatility: tracking errors have fallen dramatically for many portfolios, while the volatility of the TSE1 index has declined much more gradually. The decrease in tracking error is related to a phenomenon occurring in markets around the globe. The cross-sectional dispersion of asset returns within these markets is much smaller today than it was a few years ago. The low level of cross-sectional dispersion makes the art of active portfolio management more difficult than it was before. It demands a considered response from managers.

Section I of this paper examines the evolution of market and cross-sectional volatilities in the Japanese equity market. The impact of declining cross-sectional volatility on portfolio tracking error is the subject of Section II. Section III discusses how portfolio managers should adapt to this new environment. Should active managers take more aggressive positions in order to maintain a given level of expected return? Or should they accept smaller active returns — and perhaps smaller management fees? Section IV summarizes.

I. Declining Volatility of Japanese Equities

Figure 1 displays the history of risk in the Japanese TSE1 equity market index.

Figure 1: The Evolution of Market Volatility

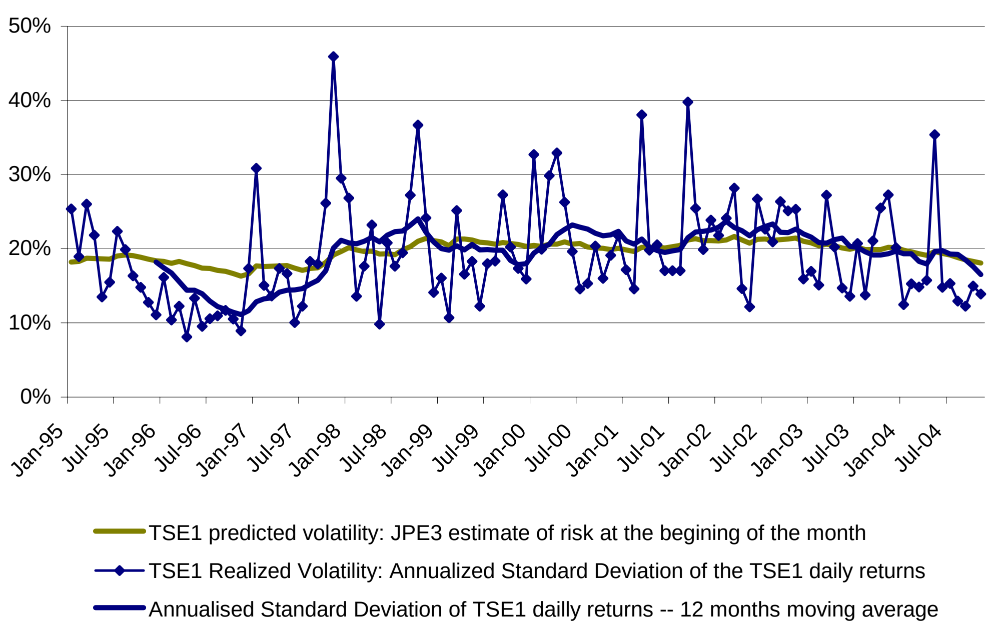

Trailing month-long windows of daily returns data are used to estimate the risk in the figure, expressed as an annualized standard deviation in return. The sudden rise in volatility that follows the 1997 Asian crisis is one of the most notable features of the history. Following this rise, the risk of the TSE1 Index hovered around 20% for several years. Since 2001, market risk has undergone a slow decline. Measurement error and volatility clustering in the history make it difficult to describe the monthly changes precisely, but a 12-month moving average shows a prolonged decline in market risk. The average risk for the index in 2004 was close to 17%, compared with 22% in 2001.

Over the same period, cross-sectional volatility experienced more extreme changes, as illustrated in Figure 2. Like market volatility, cross-sectional volatility jumped after the Asian crisis of 1997, and continued rising through the 1998 Russian crisis and the Internet bubble. It reached a high of 23% in February 2000. Cross-sectional volatility subsequently fell sharply, dwindling to 10% by mid-2000. By December 2004 it had sunk to 5%.

Figure 2: Cross-sectional Volatility of TSE1 Assets

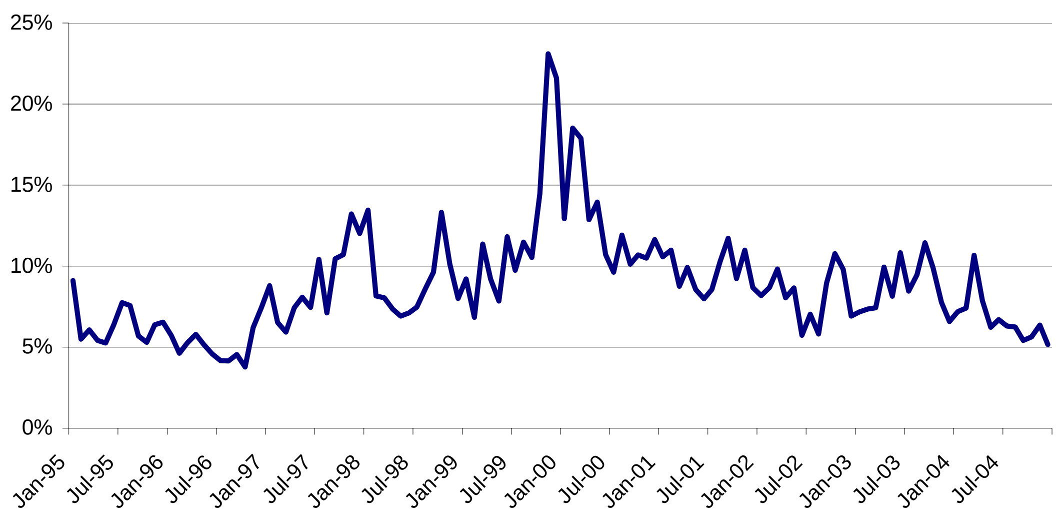

While the behaviors of market and cross-sectional volatility are linked, cross-sectional volatility has a special influence on active risk and return. This measure of returns dispersion is a good gauge of the opportunities available to managers for generating active returns. If all assets had the same return, their cross-sectional dispersion would be zero and there would be no opportunity to generate active returns. As the cross-sectional dispersion of asset returns rises, so does the opportunity to out-perform (or under-perform) a benchmark. In general, the active risk of a portfolio is expected decline as cross-sectional volatility falls, and active managers find that their task has become more difficult.

II. The Impact of Reduced Cross-sectional Volatility on Tracking Error

To see how the change in market conditions affects portfolio active risk, consider the behavior of two style indices benchmarked against the TSE1 Index. Figure 3 shows histories of forecast and realized active risks in value and growth tilts. The tilt portfolios are the Russell Nomura Large-cap Value and Growth indices. The risk forecasts are from JPE3 and JPE3S, Barra's models for Japanese equity risk. Realized risk is measured using a 12-month forward-looking standard deviation.

Figure 3: Evolution of Tracking Error for Japanese Growth and Value Indices

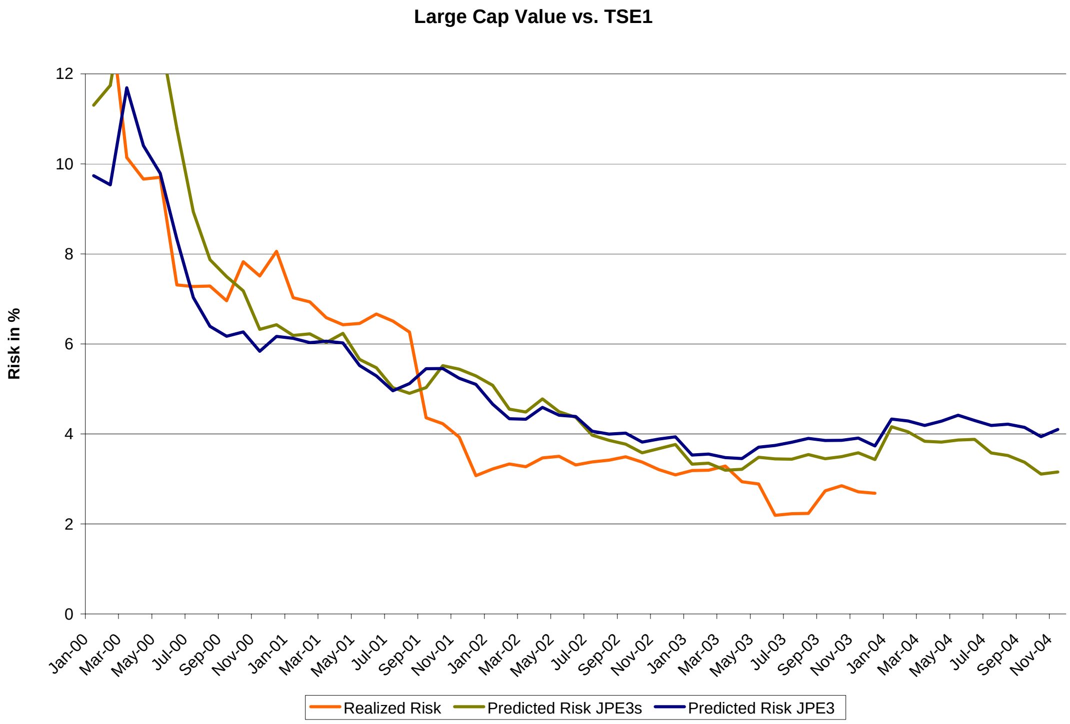

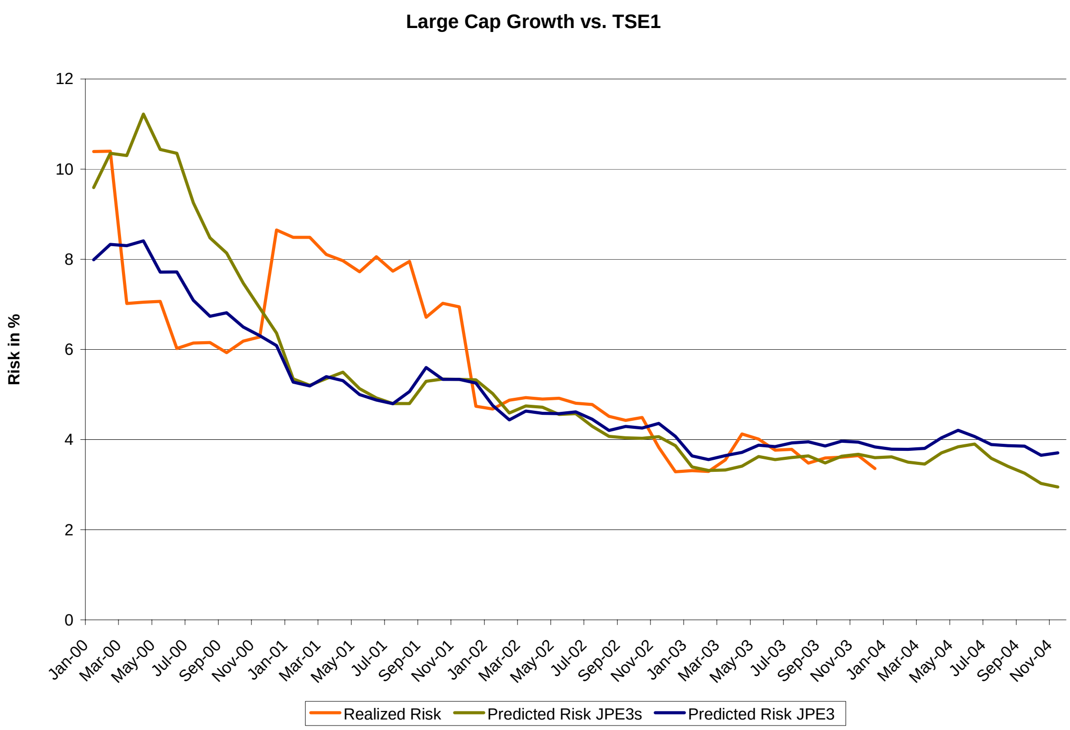

Active risk levels fell significantly with the collapse of the technology bubble. The realized tracking error of Russell Nomura Value Index plummeted from a peak value of approximately 12% in 2000 to 4% in 2004. Despite the severity of the change, the forecasts follow the realized tracking error reasonably well. Both risk models reacted satisfactorily to the changing risk environment.

Tracking error in the Value and Growth portfolios comes from a mixture of common factor and asset-specific risk. In style tilts such as these, common factors are usually the dominant source of active risk. While cross-sectional dispersion also reflects both sources of risk, it is more strongly influenced by specific risk. (Recall that about two-thirds of asset-level variance is specific risk and that market risk accounts for a significant fraction of common factor risk; c.f. the relative contributions to adjusted r-squared described in the JPE3 Research Notes, available on Barra's client website.) Thus it is most directly relevant to asset selection strategies, for which asset-specific returns play a leading role. Comparing the evolution of market-residual common factor risk with that of specific risk will help in understanding what is happening to opportunities for active management in the Japanese equity market.

Market-residual volatility is defined as the volatility of a given factor not explained by its market exposure:

$$\sigma_{residual\ factor\ f} = \sqrt{\sigma^2_{factor\ f} - \beta^2_{factor\ f}\sigma^2_{TSE1}}$$

Figures 3 and 4 show the average market-residual volatility of JPE3 style and industry factors from 1995 through 2004[1]. Both styles and industries exhibit strongly declining market-residual volatility after 2000. If portfolio positions (and hence factor exposures) remained similar throughout this period, common factor contributions to active risk would have fallen proportionately.

Figure 3: Equal Weighted Average Residual Volatility of Style Factors

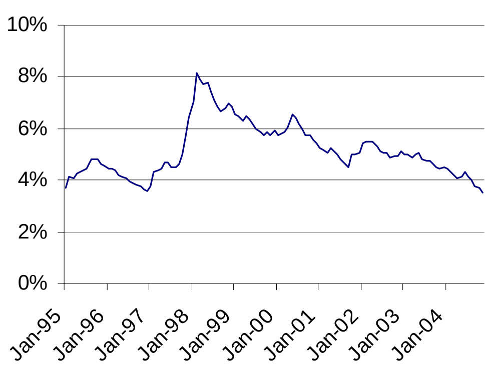

Figure 4: Equal Weighted Average Residual Volatility of Industry Factors

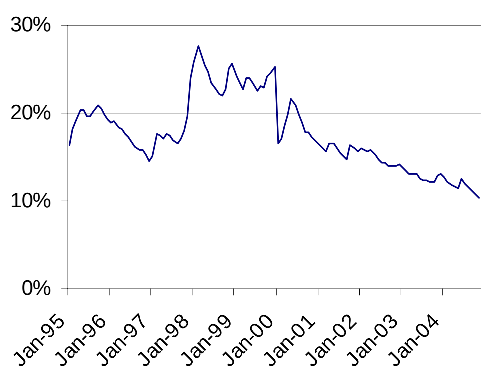

Figure 5 shows the history of predicted and realized average specific risk for Japanese equities. The downward trend for specific risk resembles the trend among

[1] The volatilities were extracted from our short term Japanese model, JPE3S. They are computed using daily factor returns and a 90-day half-life. For more information on JPE3S, consult http://www.barra.com/support/library/JPE3S_research_notes.pdf .

the residual common factors; both decline from their peaks to about half of their peak values over a 5-year period. This implies that strategies based on stock picking (sensitive to stock-specific risk) and on factor rotations (sensitive to common factor risk) have been affected to similar extents as the risk environment has changed.

Figure 5: Evolution of Average Specific Risk for Japanese Equities

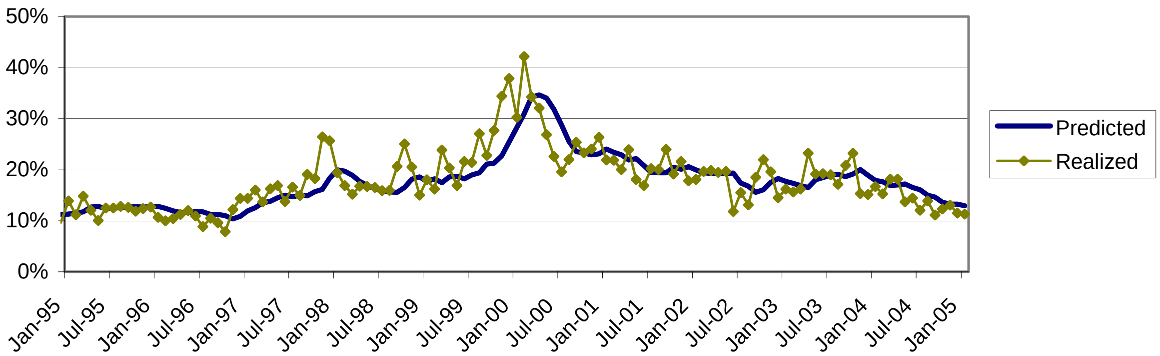

III. Implications for Active Portfolio Management

The historical decline in non-market risks has narrowed the range of active returns that any given portfolio might produce. Consequently it is now harder for an active manager to distinguish his or her portfolio from the benchmark. The active return a plan sponsor might expect from investing with a skilled portfolio manager has diminished. In this environment, the crucial thing for fund managers and plan sponsors alike is to understand how and to what extent opportunities can be preserved.

In Section II, we examined active risk in large-cap growth and value portfolios. By construction, these indices have roughly constant Growth and Value exposures. Also, although the weights of individual assets may vary over time, the population of active position sizes is relatively stable. The exposure to stock-specific risk can therefore be expected to be roughly constant. Because specific and common factor risk exposures do not respond to changing risk levels, risk in the portfolios is essentially unmanaged.

In order to maintain a constant level of active risk – and therefore preserve the potential for active return – a portfolio manager must increase active exposures as volatility declines. We use a simple example to illustrate this point.

Imagine a long-only portfolio manager who believes that large-cap stocks will out-perform the benchmark. To over-weight the Size factor, the manager uses each stock's Size exposure to construct asset-level alphas — the alpha of each stock is simply proportional to its Size exposure. The portfolio is rebalanced each month to maintain a constant forecast tracking error of 4%. Thus, every month optimization maximizes the portfolio's overall Size exposure while maintaining a constant tracking error.

This simple strategy, when implemented with JPE3S over the 5-year period 2000—2004, produced an annualized active return of 0.15% and a realized tracking error of 3.85% (15 bps below the target of 4%). To maintain a constant level of active risk, the portfolio's active Size exposure more than doubled over the study period to offset the precipitous decline in cross-sectional volatility. Figure 5 shows the historical development of the Size exposure.

Figure 5: Size Exposure History for a Constant Ex-ante Tracking Error Portfolio

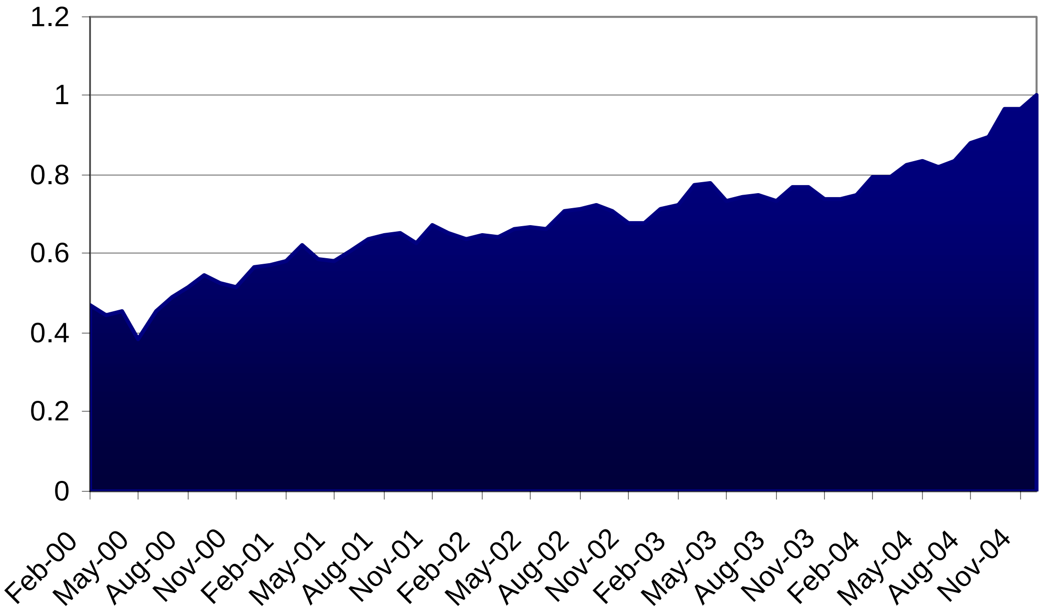

So far the discussion has focused on the need to manage risk rather than allowing it to drift amidst changing circumstances. Fixing the level of active risk implies control over risk, but it does not mean that the risk is being deployed in the best possible way. In addition to controlling active risk, a manager must justify that risk with expected active returns: the goal of active management is to maximize the risk-adjusted return of the portfolio, not to maintain a constant level of risk. In order to adapt a strategy to the new financial environment one must take into account not only the decline in residual volatility but also how expected returns have changed.

Consideration of a simplified investment problem provides some insight into how to best manage exposures. In general, a quantitative manager strives to maximize the risk-adjusted portfolio return,

$$U = h^T \alpha - \lambda h^T V h - TC .$$

Here α is the vector of forecast alphas, h is the vector of active asset weights in the portfolio, V is the residual returns covariance matrix, and TC is the trading cost of establishing and unwinding positions, expressed as a return — that is, as a fraction of portfolio wealth, amortized over the lifetime of the active positions. The "cost of risk" is described by the risk-aversion parameter λ. The manager chooses the active position sizes h to maximize the risk-adjusted active return U, subject to any policy constraints (e.g., long-only) that might be in force. Some of the optimized positions are clearly driven by forecast returns, while others help hedge the portfolio's risk. Computer-aided portfolio construction (as offered, for example, by

Barra's Aegis Portfolio Optimizer) is usually required to calculate optimal positions sizes.

In the absence of transactions costs and position constraints, the solution for the optimal positions can be written very simply:

$$h = \frac{1}{2\lambda} V^{-1}\alpha .$$

How do positions react when the financial environment changes? The risk aversion λ describes the risk preferences of the manager or asset owner; it can be expected to stay fixed. If risk (expressed as a standard deviation) halves then the variances and covariances in V will shrink to one quarter of their previous values, and the inverse covariance matrix V^{-1} will quadruple. If forecast alphas remain unchanged, position sizes should quadruple as well! The target tracking error of the portfolio would then double, justified by the superior risk-reward relationship in the new low-volatility environment. The active return of the portfolio would increase by a factor of four, and so would the risk-adjusted active return. The portfolio's information ratio, the ratio of portfolio active return to active risk, would double. But it is unlikely that alphas actually remain unchanged in the new environment. Few Japanese equity managers have seen their information ratios double as volatility has fallen.

It is more likely that the information coefficient, the ratio of forecastible alpha to risk, remains approximately constant. In this case, alpha and risk scale in the same way. When risks halve, asset-level alphas also decline by a factor of two. In response, the optimal position sizes should double. When this is done, the portfolio's tracking error target stays the same in both environments, and so does its expected risk-adjusted active return. Its information ratio therefore remains unchanged. For a portfolio with no constraints and negligible transactions costs, nothing has been lost, but neither has anything been gained. The simple policy of maintaining a fixed tracking error target has turned out to be precisely the right thing to do. Needless to say, if alphas had decreased more drastically than risk levels had the best course would be to actually reduce the tracking error target.

Maintaining a constant tracking error seems like a reasonable course of action, but policy constraints and transactions costs can make it less optimal than it appears at first sight. Before turning to the impact of transactions costs, consider first the effect of a long-only policy constraint. The long-only constraint is always damaging, since it reduces the set of information on which a manager can freely act (see for example Chapter 15 in Active Portfolio Management, 2nd Edition, by Richard Grinold and Ronald Kahn). It can become an even greater handicap in a low-volatility environment. Suppose asset volatilities decline by 50%, while asset alphas and residual return correlations remain unchanged. The long-only efficient frontier shifts to the left, as shown in Figure 6.

Figure 6: Impact of Declining Volatility and Constant Alphas on the Efficient Frontier

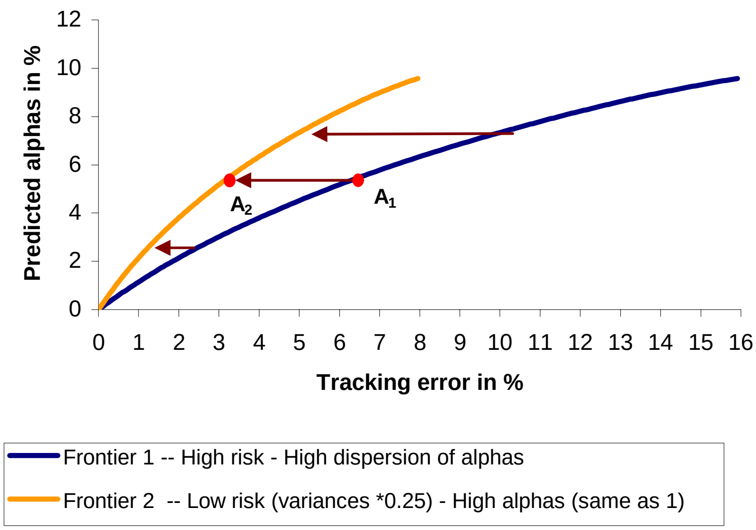

Because only the overall level of non-market risk has changed, the portfolios located on the efficient frontier remain the same. Each portfolio still yields the maximum return for a given amount of active risk, although its active risk is halved.

Suppose that a manager selected portfolio A when risk was high, and that at that time it was located at point A_1 in the risk-return space of Figure 6. After volatility has fallen, the same portfolio occupies point A_2 on Frontier 2, the new low-volatility efficient frontier. Although the risk-adjusted active return of portfolio A has improved, under the new circumstances it is no longer optimal. Since risk is now better rewarded, the manager should enlarge his exposures to increase his active returns. Keeping the same portfolio exposures would only make sense if the manager had become much more risk-averse. The manager therefore moves away from the old portfolio at A_2, choosing instead a higher risk portfolio on the new frontier.

In increasing the active positions and their associated risk, note that the manager can freely enlarge only the long active positions. Positions that are under-weighted relative to the benchmark are constrained, since the asset weight in the portfolio must be zero or greater. The expected return of portfolios on the long-only frontier does not increase in proportion to risk, but more slowly. The information ratio decreases with increasing risk. Thus, even if transactions costs are negligible and asset-level alphas remain the same, the information ratio will not double as it would without a long-only constraint. The long-only constraint is eliminating more information than it did in the high-volatility regime.

Now consider a more realistic example in which asset-level alphas and volatilities decline together (for simplicity we assume that residual return correlations remain constant). Figure 7 illustrates the resulting risk-reward tradeoff.

Figure 7: Impact of Declining Volatility and Proportionately Declining Alphas on the Efficient Frontier

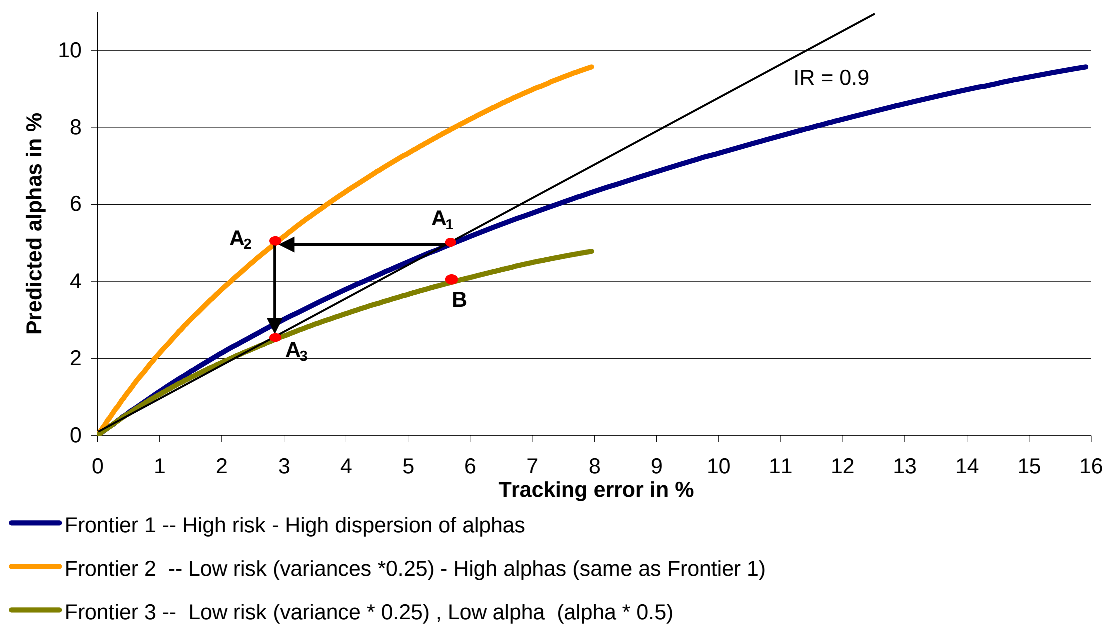

As we saw before, when risk levels are cut in half, the long-only efficient frontier moves leftward from Frontier 1 to Frontier 2. If asset-level alphas also halve, the efficient frontier shifts downward from Frontier 2 to Frontier 3. The initial portfolio A now occupies point A_3 on the risk-reward diagram. Because both active risks and active returns have been reduced, the information ratios at point A_1 and point A_3 are identical. Once again however, portfolio A is not optimal, given the manager's risk preferences. The manager should choose to increase risk-adjusted return by selecting a higher risk portfolio on Frontier 3. The optimal portfolio lies between points A_3 and B, and has a lower risk-adjusted return than in the high-volatility regime. Note that it will also have a lower information ratio than the one the manager previously enjoyed. Under the same scenario, a long-short manager would simply double the size of all active positions to maintain the same risk-adjusted return and information ratio. One again, there is less usable information in the low-volatility environment, because of the long-only constraint.

We now turn to the role of transactions costs in the new environment of lower risk.
It is simplest to discuss costs in the context of small orders, for which commissions and the loss on the spread remove a fixed fraction f from the value from each trade, independent of the exact size of the trade (typically for small orders, f might be a few dozen basis points; when orders become very large the fraction f will start to depend on order size and will increase as orders become larger). If an active position is held for a time τ, the act of setting up and unwinding an active position reduces the return rate on the position from α to the cost-adjusted value

$$\tilde{\alpha} = \alpha - \frac{2f}{\tau}\operatorname{sign}(h) .$$

(Unless a position is taken purely for hedging, the sign of h should be the same as the sign of α.) The cost-adjusted value can be used to estimate optimal active position sizes in the absence of policy constraints:

$$h = \frac{1}{2\lambda} V^{-1} \tilde{\alpha} \ .$$

Now suppose that residual risk levels and forecast returns fall to half of their previous values. If the trading cost rate f fails to decease along with the expected residual return α, the cost-adjusted residual return will decrease by more than half; indeed, many cost-adjusted alphas will become equal to zero.

As an example consider a very plausible situation in which, prior to the decrease in volatility, a manager used to lose a third of his forecast alpha to transactions costs. Then, the financial environment changes and both risk and asset-level alphas are halved. If commissions and spreads remain unchanged, the manager now loses two thirds of the asset-level alpha and so is left with cost-adjusted alphas that are *one-quarter* of their previous sizes (1/3 of 1/2 as opposed to 2/3 of 1). Taking transactions costs into account, the correct response on the part of the manager is to allow tracking error to fall and, roughly speaking, keep position sizes the same. Managers with less burdensome costs can afford to increase their position sizes somewhat, partially compensating for the loss in risk-adjusted return. Obviously transactions costs have a dramatic effect on the decision of how best to proceed in the new environment, and should be carefully considered.

IV. Conclusions

Although the risk of the TSE1 index has been relatively stable since the end of the Internet bubble, cross-sectional volatility levels have declined sharply. As a result, many Japanese equity managers have seen their strategies weaken as their tracking errors have decreased. Active management has become a harder game.

Taking more aggressive positions in order to compensate for the decline in tracking errors should help address the problem, but transactions costs and policy constraints imply that the compensation need not be complete; indeed, some reduction in targeted tracking error is probably optimal. The reduced tracking error is accompanied by a lower risk-adjusted return, an inevitable consequence of the harsher environment. Performance will suffer most in portfolios with long-only constraints or other strong constraints on position sizes, and in portfolios with high turnover. Long-short portfolios and portfolios with low turnover can compensate more completely and should be less severely affected.

Authors:
Edouard Senechal: Barra Applied Research Group
Xiaowei Li: Barra Research Group
Guy Miller: Barra Research Group