Barra

The Japan Short-Term Equity Model, JPE3S: A Highly Responsive Risk Model for Japan

Xiaowei Li, Research Consultant
Guy Miller, Senior Research Consultant

April 2004

This report introduces the Japan Short-Term Equity Model, JPE3S, a model for near-term (several month) Japanese equity risk. JPE3S responds more quickly to changes in risk levels than the Japan Equity Model, JPE3. In order to make more responsive risk forecasts, JPE3S employs daily returns data and accounts for their serial correlations. Daily data provide denser and more detailed intra-horizon volatility information than would be available from monthly returns, and allow the model to base its forecasts on a relatively short data history. Serial correlations are significant in aggregating daily factor returns to longer horizons; incorporating them substantially enhances model performance.

1 Introduction

This report introduces the Japan Short-Term Equity Model, JPE3S, a model for Japanese equity risk over terms of one to six months. In keeping with the extremely volatile nature of risk over monthly horizons, its forecasts respond to changing risk levels more aggressively than those of the Japan Equity Model, JPE3.

1.1 Risk Behavior and Investment Horizons

Risk levels in financial markets change relentlessly. The appropriate response from an investor to this fundamental fact depends partly on the investment horizon; i.e., on how much time typically elapses between major revisions in portfolio positions. The risk that matters to the investor is the average level of risk that the portfolio will experience over the investment horizon. Figure 1 shows two realized risks in Japan's Tokyo Stock Exchange First Section (TSE1) index from 1980 through 2003. The two risks shown are estimates, based on daily data (and corrected for autocorrelation in returns), of the risk the TSE1 experienced over the subsequent month and over the subsequent half-year. The figure shows that risk levels over the shorter horizon display more pronounced clustering and are far more volatile than levels over the longer horizon. This observation is hardly surprising, but it is important. Risk forecasts aimed at investors with shorter horizons should naturally be more volatile than risk forecasts aimed at longer-term investors. The short-horizon forecast must react more sensitively to rapid—and typically short-lived—changes in risk levels than its long-horizon counterpart.

Figure 1 also suggests a second point. Since risk over monthly horizons is both volatile and erratic, it will almost certainly be impossible to predict each peak and valley. Rather, a good risk forecast for the coming month will aim at an expected average level. It will be reactive, but still smoother than the realized volatility itself. Thus a good forecast for monthly risk should be acceptable over horizons from one to several months. The window of investment horizons over which JPE3S should perform well is one to six months.

The behavior seen in Figure 1 is not unique to market risk; it is generic. The realized monthly specific risk of assets in the estimation universe, as shown in Figure 2, also varies considerably over time, displaying short-lived fluctuations similar to those in the history of TSE1 risk.

FIGURE 1
Realized Volatility of TSE1
Standard Deviation of Daily Returns over Each Month

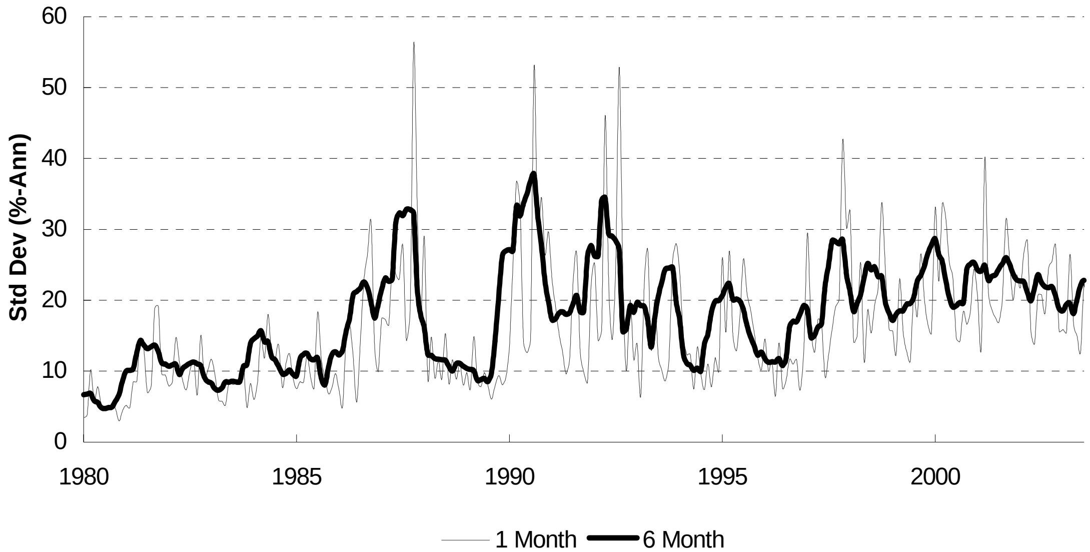

1.2 Implications for Short-Term Risk Forecasting

To respond quickly to changing risk levels, risk forecasts must emphasize more recent data. This effectively shortens the data history. In EWMA (exponentially weighted moving average) forecasts, the *half-life* specifies how the weights placed on historical observations decline through time, with more recent observations receiving greater weight. For example, a half-life of 12 months implies that observations from 12 months ago receive half the weight of the current month's observations. One way of achieving greater model responsiveness is therefore to retain monthly returns data, and simply to shorten the half-life to 12 or 24 months in constructing the factor covariance matrix. The disadvantage of such short half-lives with monthly data is that they reduce the *effective sample size* significantly, resulting in noisy and less precise forecasts. Thus, although using monthly factor returns with a short half-life would increase responsiveness, it is a poor modeling choice.

In order to build risk forecasts that are not only more responsive but also precise, we need to consider using returns that are sampled more frequently than once a month (e.g., daily or weekly). Higher frequency data permit us to shorten the half-life while maintaining a large enough sample of historical observations that precision is unimpaired. For example, if we use daily instead of monthly data with the same half-life, the improvement in precision is of order $\sqrt{T}$ [1], where T is the number of trading days in a month.

This section has reviewed some of the differences between near-term and longer-term risk, and delved into the main implications of these differences for near-term risk prediction. The rest of this paper explores the structure and capabilities of the JPE3S model. Section 2 describes its features, emphasizing points at which it differs from the monthly returns-based JPE3 model. The performance of the new model is documented in Section 3. Section 4 takes up the question of model choice, providing guidelines for who will be better served by a more responsive model like JPE3S, and who will find that JPE3 better meets their needs.

2 Building A More Responsive Risk Model

2.1 Daily Factor Returns and Serial Correlation

The JPE3 factors provide a detailed and robust description of sources of commonality in the returns to Japanese equities. The economic rationale behind the factors and their high statistical power have been documented elsewhere.[2] Here we simply note that the factor structure of JPE3 provides an excellent basis for the cross-sectional modeling of common factors over horizons of a month or more. We therefore adopt it without modification in constructing a more responsive model. In so doing, we preserve JPE3's cross-sectional detail. Our goal will be to refine its temporal resolution, and accelerate its "forecasting speed."

JPE3S risk forecasts are therefore based on a time series of daily factor returns to the JPE3 model factors. These are produced by regressing daily asset returns against factor exposures that are computed at the start of each month. The regression is identical in all other respects to JPE3's monthly factor return regression.

Daily returns naturally form the basis of more agile risk forecasts—but some technical sophistication is required to extract their full value. Daily returns are more susceptible than

[1] Assuming daily returns are independent.

[2] See "Japan Equity Model (JPE3) Research Notes" available at www.barra.com/support /library.

monthly returns to the influences of market microstructure (asynchronous trading, bid-ask spreads, limits on daily price movements, etc.) and other evanescent market phenomena (information diffusion effects, under- and over-reaction). These effects can induce lead-lag relationships in asset-level returns, which then manifest themselves as serial correlations in factor returns. When using daily data for forecasting risk over horizons that are long compared to the sampling frequency, it is important to account for serial correlations.

Appendix A contains a technical account of how serial correlations enter the JPE3S risk forecast. Qualitatively, their importance is easy to understand. If positive (negative) returns on one day are likely to be followed by positive (negative) returns on the next day, the returns reinforce one another and monthly volatility is greater than if the returns on each day were independent. If the returns on one day tend to be cancelled by subsequent returns, monthly volatility is subdued.

One of the JPE3 factors most dramatically affected by serial correlation is Momentum. The monthly variance of the Momentum factor tends to be strongly enhanced by positive serial correlations extending over several days. Although the enhancement can vary with time, throughout the period from 1990 to the present serial correlations have consistently increased the monthly variance of Momentum returns by a factor of 1.7 or more above the value obtained by simply scaling the daily variance.

2.2 Specific Risk

From Figure 2, we conclude that the specific risk system of JPE3 is already sufficiently reactive to track changing risk levels well over horizons of one month. Furthermore, our tests of its forecasts on monthly horizons over many time periods and a vast range of portfolios show that it performs very well. We have experimented with modifications to the system that depend explicitly on the use of daily returns data, but have not found any modifications that systematically improve its portfolio-level forecasts of specific risk. In constructing a more responsive version of JPE3, we therefore adopt the JPE3 specific risk forecasting system without modification.

2.3 Model Construction

We used the daily factor returns to JPE3 factors to generate monthly-horizon factor

FIGURE 2
Realized and Forecast Average Specific Risk*

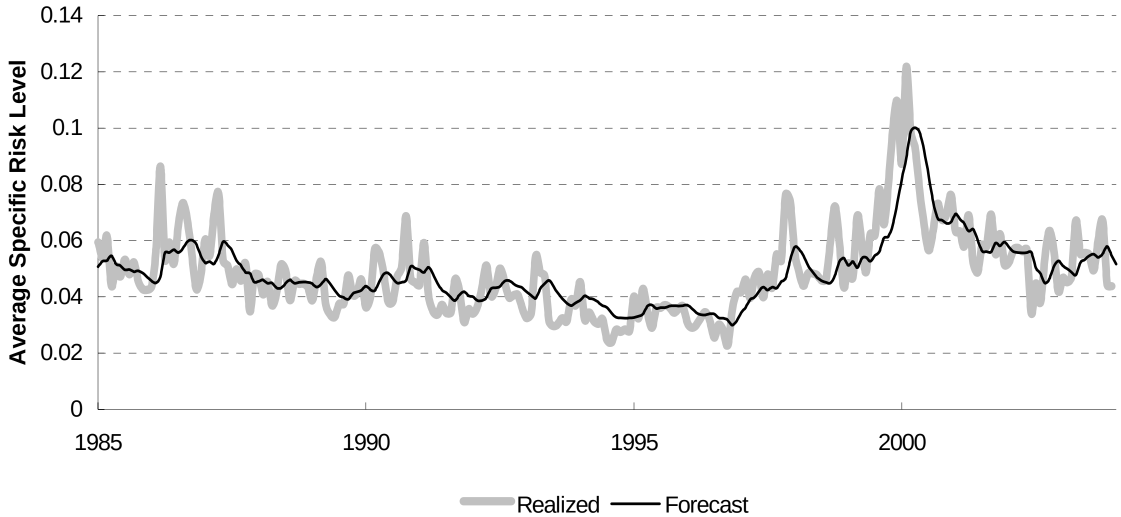

**Realized specific risk is the capitalization-weighted average of the absolute monthly specific return to each asset in the estimation universe for JPE3. Forecast is from JPE3 model.*

covariance matrices, using the method described in Appendix A. Forecasts under this method depend on three parameters. The first parameter is the half-life for factor variances. The forecasting method tracks daily levels of risk in each JPE3 factor independently; an exponentially weighted moving average (EWMA) is used to make a variance forecast for each model factor. (Alternatives to the simple EWMA forecasting scheme such as ARMA and GARCH were explored, but ultimately were rejected as less robust and reliable.) The second parameter is the correlation half-life, which controls the effective length of historical data applied to calculating serial correlations and correlations across factors. By treating variances and correlations differently, this forecasting methodology attempts to capture changes in risk levels promptly and sensitively, without unduly sacrificing accuracy in its description of relations between factors. The third parameter is the number N of successive trading days beyond which serial correlations are assumed to be negligible. If N=0, each day is assumed to be independent of the next and serial correlations are omitted from the forecasts. If N=10, then a return on one day is assumed to influence returns over the following 10 “lagged” days; relations within chains of 11 consecutive trading days enter the forecasts.

In order to identify the best forecasting parameters, we constructed a series of models accounting for different lengths N of serial correlation and using different half-lives. Specifically, we considered zero lagged correlations for all factors; five lags for all factors; ten lags for all factors; 20 lags for all factors; and a “mixed model” with 20 lags for risk factors and risk/industry cross factors, and five lags for industry factors. The half-lives we explored included 480, 750, and 1000 trading days for the correlations, and 30, 60, 90, 120, 150, 180, and 250 trading days for the variances.

3 Model Evaluation

In previous sections, we explained our motivation for leaving the set of factors and the specific risk model unchanged from JPE3’s. The choice of short-term model comes down to selecting a method of constructing the covariance matrix. In Section 2 we saw that factor covariance forecasts depend on three parameters: a serial correlation length N and two half-lives, a longer one for estimating factor correlations, and a shorter one for estimating factor variances. In Table 1 we show characteristics of three candidate models with covariance matrices constructed with different values for these parameters. Alongside the final choice of model, JPE3S, are two alternatives.

TABLE 1
Model Parameters

Model Name	*Half-Life for Variance σ^2 (Days)*	*Half-Life for Cross Correlation ρ & Serial Correlation ϕ (Days)*	*Lags (N) of Serial Correlation ϕ (Days)*	*Features*
JPE3S	90	480	10	• Shorter half-life for variances & correlations • Accounts for serial correlations
1000_250	250	1000	20	• Longer half-life for variances and correlations • Accounts for serial correlations
480_90nc	90	480	0	• Shorter half-life for variances & correlations • Ignores all serial correlations

Note: 480 days corresponds to approximately two years of trading days and 90 days to about four and a half months.

The forecasting parameters of JPE3S have been selected from a large set of possibilities to provide accurate and robust near-term risk forecasts over a wide range of market conditions.

This section describes the performance of JPE3S. It focuses on factor risk. The discussion begins with a qualitative, graphical comparison between JPE3S factor risk forecasts and realized risk levels. This graphical introduction is followed by quantitative comparisons between JPE3S, the two other daily-based models listed in Table 1, and JPE3. The section concludes with results from JPE3S performance tests in several historical periods

3.1 Risk Forecast Behavior

Figures 3a–3h illustrate the dynamic behavior of JPE3S risk forecasts. The first four figures in the series show forecasts and realized monthly risk for four representative industries. The second four display risk indices. From the figures, it is clear that JPE3S forecasts usually track realized volatility levels rather closely. Even in times of very rapid change, when the forecast necessarily trails realization, JPE3S aggressively pursues the changing level of realized risk. Over a half year, the JPE3S risk forecast can change substantially.

Plotted alongside the JPE3S risk prediction is the forecast of JPE3. The JPE3 forecast does not follow each rise and dip in the moving level of monthly risk as assiduously as JPE3S. Instead, its forecasts adjust slowly and conservatively. Although its expectation for the coming year is occasionally challenged by an unusually large and long-lived variation (e.g., Size in the mid-1990s), in general its strategy of conservative adjustment is successful in identifying and following risk level trends that persist for a year or longer.

3.2 The Bias Statistic as a Measure of Forecast Accuracy

The primary measure that we use to assess the accuracy of risk forecasts is the bias statistic. The bias statistic can be interpreted intuitively as the ratio of realized risk to forecast risk over a period of time. Appendix B discusses the construction of the bias statistic. If a model forecasts risk accurately over a monthly horizon within the period under examination, the bias statistic will be close to 1. If the bias statistic is significantly below 1, then the model forecast has been too high. Conversely, if the bias statistic is significantly above 1, then the model has under-predicted risk.

FIGURE 3A
Realized and Forecast Risk for Computer Industry

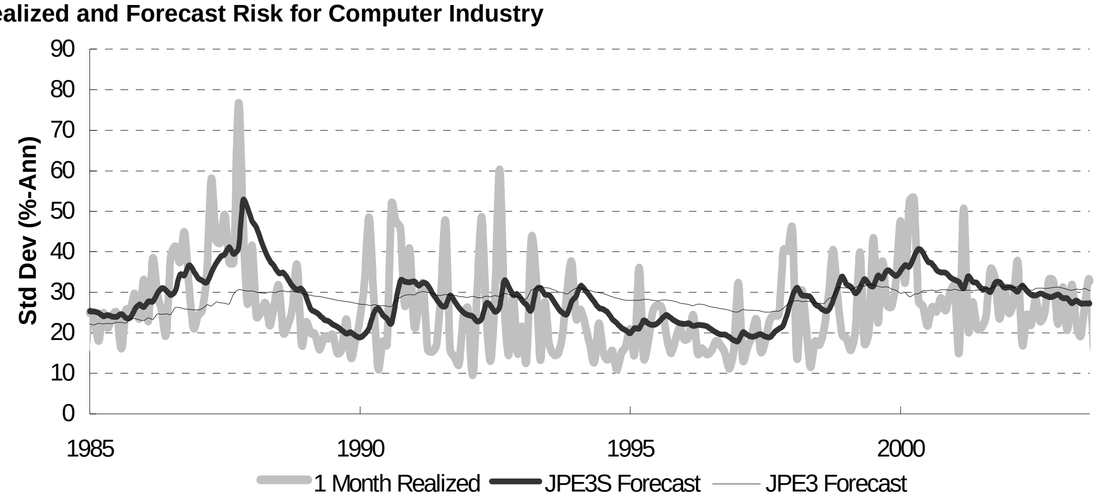

FIGURE 3B
Realized and Forecast Risk for Bank Industry

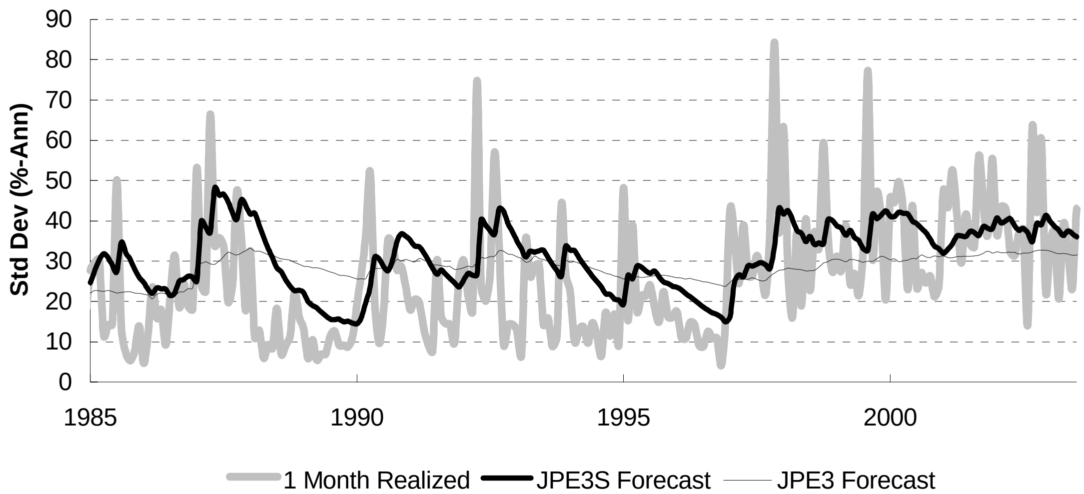

FIGURE 3C
Realized and Forecast Risk for Energy Industry

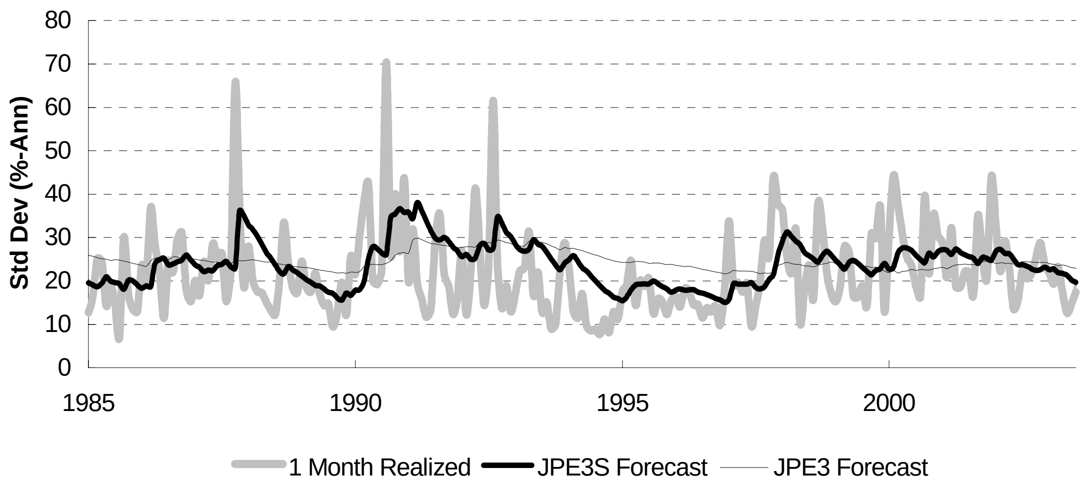

FIGURE 3D
Realized and Forecast Risk for Automobile Industry

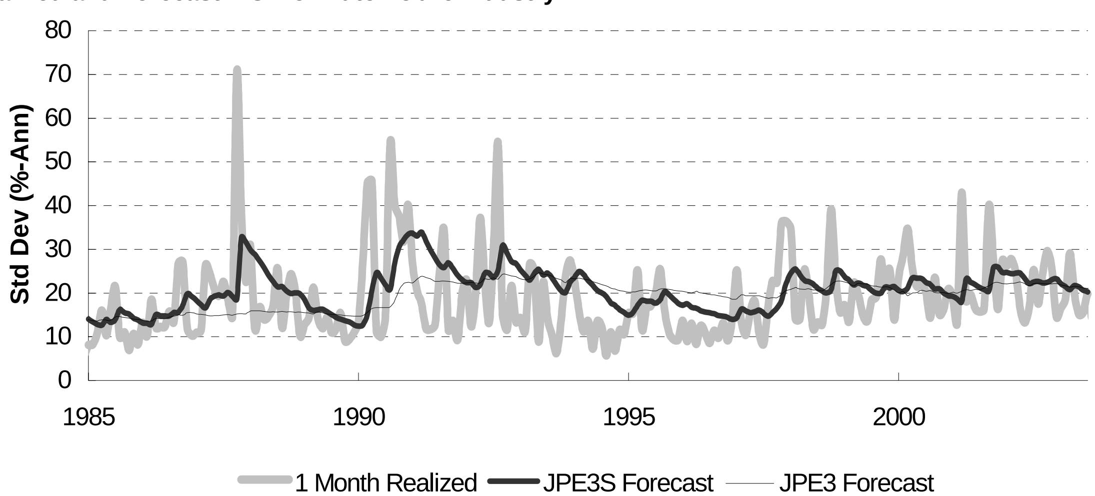

FIGURE 3E
Realized and Forecast Risk for Volatility Risk Index

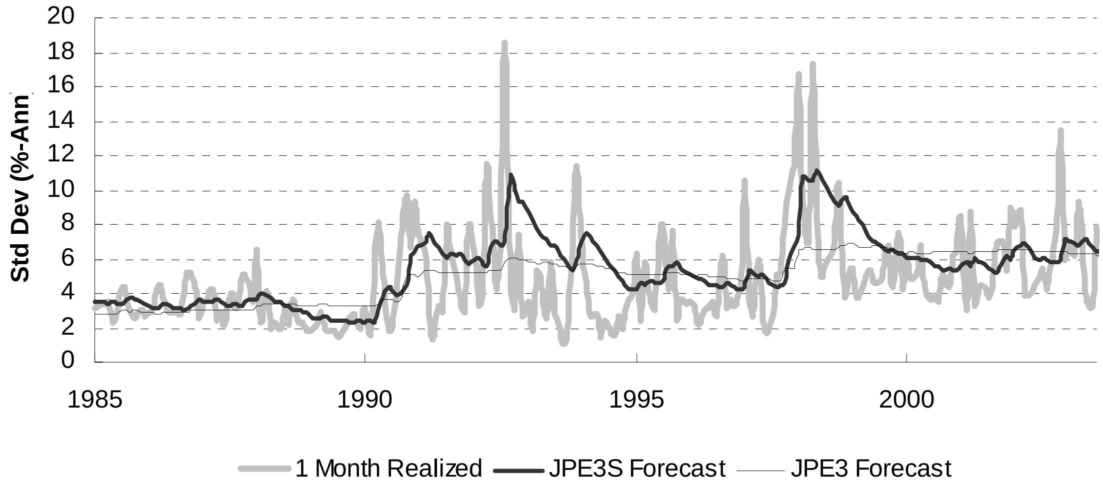

FIGURE 3F
Realized and Forecast Risk for Size Risk Index

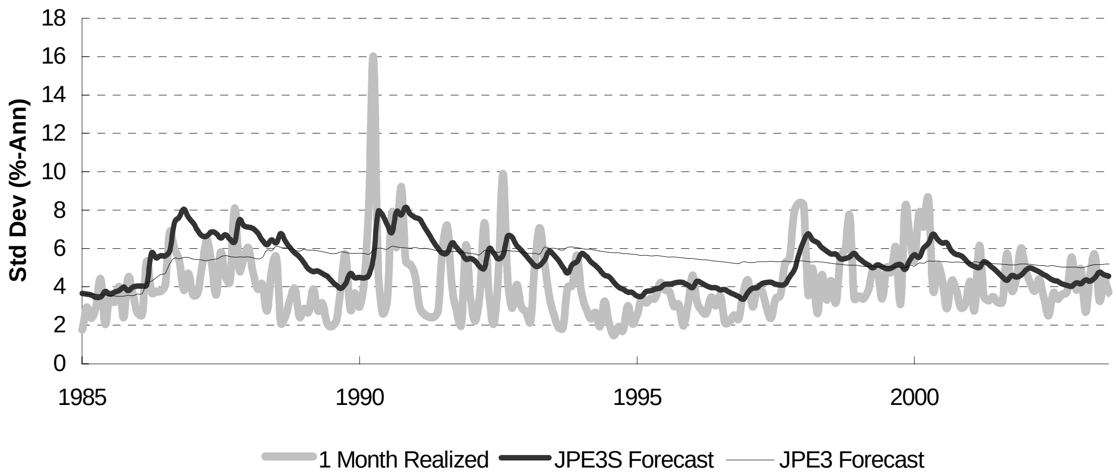

FIGURE 3G
Realized and Forecast Risk for Value Risk Index

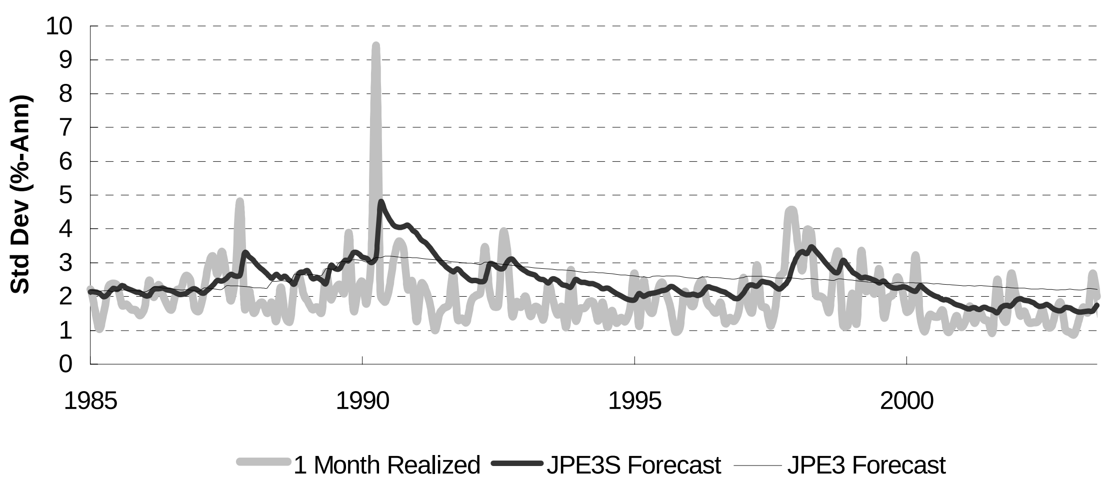

FIGURE 3H
Realized and Forecast Risk for Growth Risk Index

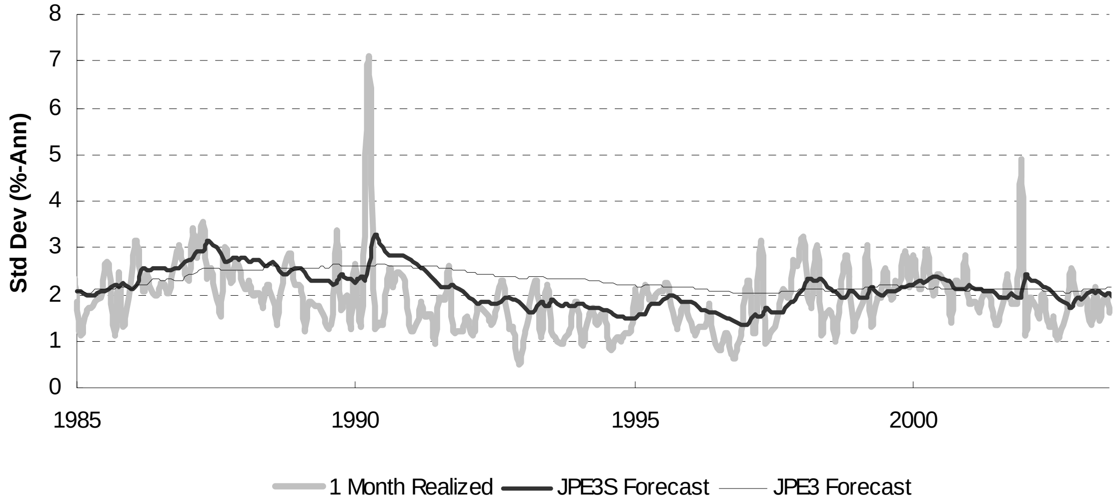

3.3 Bias Test Results

Tables 2a–2c show distributions of bias statistics for monthly factor risk forecasts over three historical intervals.

TABLE 2A
Distribution of Bias Statistics for Factors
January 1987 to December 1991

Model	*Mean*	< 0.65	0.65-0.75	0.75-0.85	0.85-0.95	0.95-1.05	1.05-1.15	1.15-1.25	1.25-1.35	1.35-1.45	1.45-1.55	1.55-1.65	1.65-1.75	1.75-1.85	1.85-1.95	> 1.95
JPE3	1.26	0	0	0	2	4	7	10	10	9	7	1	0	0	0	0
JPE3S	1.14	0	0	1	0	8	17	18	5	0	1	0	0	0	0	0
1000_250	1.15	0	0	1	2	13	6	20	5	2	0	0	1	0	0	0
480_90nc	1.18	0	0	0	3	4	15	18	5	1	2	1	1	0	0	0

Shaded cells indicate bias statistic is significantly different from 1.

TABLE 2B
Distribution of Bias Statistics for Factors
January 1992 to December 1997

Model	*Mean*	< 0.65	0.65-0.75	0.75-0.85	0.85-0.95	0.95-1.05	1.05-1.15	1.15-1.25	1.25-1.35	1.35-1.45	1.45-1.55	1.55-1.65	1.65-1.75	1.75-1.85	1.85-1.95	> 1.95
JPE3	0.90	0	4	10	24	6	5	2	0	0	0	0	0	0	0	0
JPE3S	1.00	0	0	0	10	28	10	3	0	0	0	0	0	0	0	0
1000_250	0.93	0	0	7	30	8	5	1	0	0	0	0	0	0	0	0
480_90nc	1.09	0	1	0	7	10	21	5	5	1	0	1	0	0	0	0

Shaded cells indicate bias statistic is significantly different from 1.

TABLE 2C
Distribution of Bias Statistics for Factors
January 1998 to December 2003

Model	*Mean*	< 0.65	0.65-0.75	0.75-0.85	0.85-0.95	0.95-1.05	1.05-1.15	1.15-1.25	1.25-1.35	1.35-1.45	1.45-1.55	1.55-1.65	1.65-1.75	1.75-1.85	1.85-1.95	> 1.95
JPE3	0.92	0	2	20	12	9	7	2	1	0	0	0	0	0	0	0
JPE3S	0.90	0	0	19	18	10	6	0	0	0	0	0	0	0	0	0
1000_250	0.89	0	2	20	19	7	4	1	0	0	0	0	0	0	0	0
480_90nc	0.89	1	9	13	16	6	4	2	1	1	0	0	0	0	0	0

Shaded cells indicate bias statistic is significantly different from 1.

From these tables, the following three points are clear. First, it is important to account for serial correlations. In almost all of the study periods, the daily-based model that omits serial correlations (480_90nc) is conspicuously inferior to JPE3S. In fact, it tends to make more severe under-forecasts than *any* of the other models, including JPE3. Second, after accommodating serial correlations, the daily-based models generally outperform the monthly-based model, JPE3. By effectively sampling denser data over a shorter period, these models react to changing monthly risk levels more nimbly than is feasible for a model based on monthly returns data. Third, a longer variance half-life (as in the 1000_250 model) is too long for risk forecasts over monthly horizons. An overly long half-life prevents the risk model from adequately exploiting recent volatility data. On the other hand, a very short half-life (e.g., 30 trading days) causes the model to overreact to the short-lived fluctuations always present in daily data. A half-life of several months represents a "happy medium," with forecasts that are responsive, but not oversensitive.

3.4 Factor Covariance Forecast Quality

The relatively short factor variance half-life in JPE3S allows it to react quickly to changes in factor risk. Factor correlations in the model are estimated with a longer 480-day half-life. What can we say about the quality of the model's cross-factor forecasts?

A practical test of the model's overall covariance structure can be performed by comparing the forecast common factor risk for a large number of randomly generated portfolios with the realized monthly factor returns attributed to each portfolio. Of special interest is the active return to each portfolio (the difference between the portfolio return and the return to the TSE1 benchmark), since it excludes much of the market risk and allows forms of non-market risk to take on a larger role. The results appear in Table 3. The overwhelming success of the active risk forecasts strongly suggests that the model has captured relations between different sources of risk accurately.

TABLE 3
Distribution of Common Factor Bias Statistics for Random Active Portfolios

Period	*Mean*	< 0.65	0.65-0.75	0.75-0.85	0.85-0.95	0.95-1.05	1.05-1.15	1.15-1.25	1.25-1.35	1.35-1.45	1.45-1.55	1.55-1.65	1.65-1.75	1.75-1.85	1.85-1.95	> 1.95
1987-1991	1.01	0	0	0	5	22	3	0	0	0	0	0	0	0	0	0
1992-1997	1.00	0	0	0	0	26	4	0	0	0	0	0	0	0	0	0
1998-2003	0.97	0	0	0	10	20	0	0	0	0	0	0	0	0	0	0

Shaded cells indicate bias statistic is significantly different from 1.

3.5 Overall Forecast Quality

How does JPE3S perform on monthly horizons with real portfolios? We address this question by selecting a group of portfolios and directly comparing their realized monthly returns with JPE3S risk forecasts.

Table 4 shows the list of portfolios, each accompanied by bias statistics in three historical intervals. For active portfolios, the return tested is the portfolio return minus the return to a benchmark portfolio. Asterisks (*) mark forecast biases that are significant at the 95% confidence level, under the conservative assumption of normally distributed normalized returns. Overall, JPE3S forecasts of monthly portfolio risk perform well. The least successful period for the model is 1987–1991, but within this period model performance is better than the bias statistics would suggest. With the exception of the TSE2 portfolio, the significant forecast biases all originate from two shocks, the first in March 1990 and the second in October of the same year.

4 Investment Horizon, Forecast Variability, and Model Choice

Which investors should prefer JPE3S to JPE3? Figures 3a–3h suggest that JPE3S forecasts, although nominally tuned to a 1–6 month investment horizon, often perform well over longer horizons. Nevertheless, as the investment horizon lengthens its advantage in accuracy over JPE3 becomes less clear. Its *disadvantage* with respect to JPE3 then becomes a consideration. Because it is a more reactive model that emphasizes very recent data, its risk

TABLE 4
Historical Portfolio Bias Statistics

Portfolio	*1987–1991*	*1992–1997*	*1998–2003*
TOPIX Core 30	NA	1.117	0.893
TOPIX Core 30 vs. TSE1	NA	1.029	0.937
FTSE Japan	1.207*	1.087	0.876
FTSE Japan vs. TSE1	1.006	0.992	0.888
TOPIX Mid 400	NA	1.072	0.861
TOPIX Mid 400 vs. TSE1	NA	0.958	1.091
MSCI Japan	1.191*	1.103	0.862
MSCI Japan vs. TSE1	0.940	0.975	1.080
Nikkei 225	1.151	1.107	0.843
Nikkei 225 vs. TSE1	0.943	0.728*	0.764*
Nikkei 300	NA	1.109	0.853
Nikkei 300 vs. NK225	NA	0.812	0.834
Nikkei 500	1.236*	1.086	0.862
Nikkei 500 vs. NK225	0.917	0.700*	0.718*
TOPIX 100	NA	NA	0.886
TOPIX 100 vs. TSE1	NA	NA	0.955
TOPIX 500	NA	NA	0.865
TOPIX 500 vs. TSE1	NA	NA	0.954
TSE First Section	1.259*	1.081	0.879
TSE Second Section	1.366*	1.007	1.141
TSE Second Section vs. TSE1	1.136	0.838	0.843
Top 50 by cap vs. TSE1	0.902	0.977	0.970
Top 50 by cap	1.093	1.099	0.916

**Bias statistic is significantly different from 1.*

forecasts tend to change more vigorously from month to month than those of JPE3. This variability makes model selection a matter of judgement for managers with semiannual horizons.

Figure 4 displays the realized risk of the Computer industry versus forecasts from JPE3 and JPE3S. Realized risk is represented by a forward-looking six-month average of daily factor volatility, corrected for serial correlation. It is important to note that when the realized volatility level flies upward, it does so six months *before* the actual volatility appears in returns. As such, it is a prescient ideal target for real risk models, which can only base their forecasts on events that already have taken place.

Comparing Figure 4 with Figure 3a, one sees that the history of realized six-month risk is similar to the history of one-month risk, but smoother and flatter. As the investment horizon lengthens, the resemblance of realized risk to the JPE3 forecast tends to increase. Thus, although JPE3S still appears to have an advantage in accuracy over JPE3 at six-month horizons, its advantage is less clear-cut than on shorter horizons; there are only a few prolonged intervals over which its advantage is uncontestable (e.g., the mid-1990s). A manager with a semiannual investment horizon whose trading is strongly affected by risk forecast variability should prefer JPE3 to JPE3S, since the advantage of JPE3S in accuracy occurs sporadically, while the disadvantage in forecast variability is always present. A manager who carefully minimizes trading, either by incorporating transactions costs in portfolio optimization or by simply constraining turnover, may favor JPE3S. Managers with horizons of a year or longer will generally prefer JPE3.

FIGURE 4
Forecast and Realized Volatility of Computer Industry

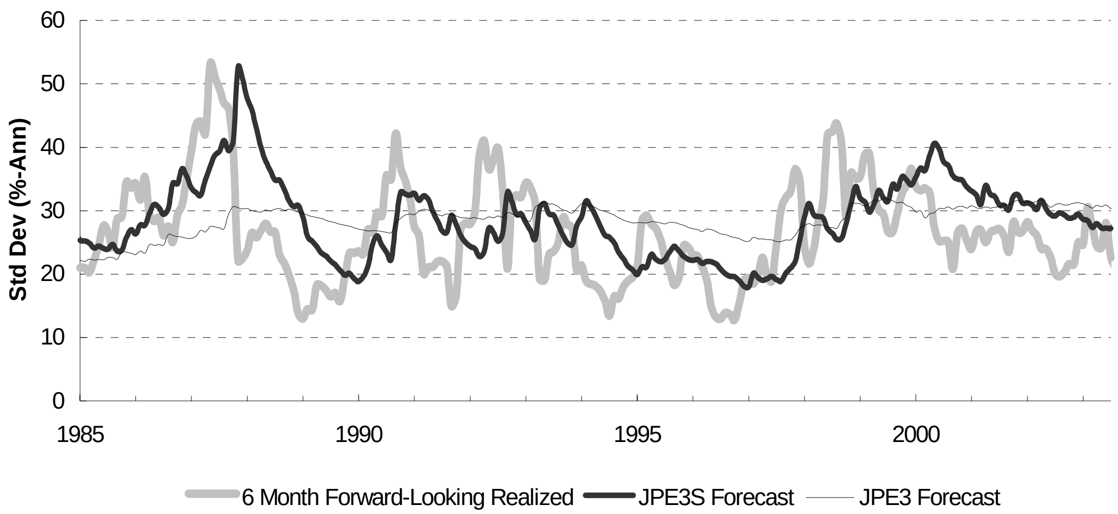

Table 5 contains JPE3S and JPE3 forecasts in June and December of 2003. The forecasts are for total risk in the TOPIX 100 index, and for its active risk against the TSE1 benchmark. The risks shown are standard deviations, expressed as annualized percentages. Over the six-month period separating the June and December forecasts, JPE3 shows little change. Changes in the JPE3S forecasts, although not unreasonable in size, are considerably larger.

TABLE 5
Risk Forecast Comparison
(Std Dev, %-Ann)

	JPE3		*JPE3S*	
Portfolio	*Jun-03*	*Dec-03*	*Jun-03*	*Dec-03*
TOPIX 100	21.62%	21.61%	19.43%	21.13%
TOPIX 100 vs. TSE1	4.15%	3.91%	3.67%	3.46%

A convenient way to quantify forecast variability is to use the *average relative forecast change*, FC. This is defined as the average over a stated historical period of the absolute changes in forecast standard deviation from one month to the next, each divided by the forecast for the earlier month. For example, if a risk forecast changed from 3.00% in June to 3.06% in July, the absolute relative change would be 0.02, or 2%. If the forecast then changed from 3.06% to 2.96% in August, the absolute relative change would be 0.04. The average over the period of June through August would be FC=0.03.

Table 6 compares average relative forecast changes between JPE3 and JPE3S in each of three historical intervals. The forecast changes are averaged over industry factors and risk indices separately. Changes in the forecast average level of specific risk (cf. Figure 2) are also shown. The data in the table show that the monthly factor risk forecast variability in JPE3S is roughly 2 or 3 times larger than that seen in JPE3. Since the two models share the same specific risk system, their specific risk forecast variabilities are identical.

TABLE 6
Average Relative Forecast Changes

	JPE3			*JPE3S*		
Component	*1987–91*	*1992–97*	*1998–2003*	*1987–91*	*1992–97*	*1998–2003*
Risk Indices	0.0200	0.0168	0.0178	0.0536	0.0510	0.0523
Industries	0.0303	0.0277	0.0260	0.0842	0.0647	0.0538
Specific Risk	0.0311	0.0301	0.0415	0.0311	0.0301	0.0415

Variability in a risk forecast can drive undesirable turnover. The increase in turnover a manager experiences in moving from JPE3 to JPE3S will depend on how much of the manager's risk comes from common factors (as opposed to specific risk), and on whether the manager adjusts position sizes every month or less frequently. To get an idea of the magnitude of the expected turnover change in a severe case, consider a manager who manages against a benchmark and who carries mainly factor risk. He adjusts his active positions every month to meet a targeted tracking error. When the risk forecast increases, the manager must reduce the active position sizes to return to the target. When the forecast decreases, the manager increases the active position sizes. The amount of turnover in the active portfolio (half of the cash value traded annually, divided by the summed absolute cash values of all active positions) driven by forecast variability is approximately TO = 6×FC.

If the manager goes from JPE3 with an average relative forecast change of FC=0.02 to JPE3S with FC=0.05, turnover in the active portfolio will increase by roughly 18%, all other things being equal. Similar estimates apply to a manager who uses mean-variance optimization for portfolio selection, unless trading is constrained or transactions costs are explicitly included in the optimization.

5 Conclusion

JPE3S uses daily returns data to make accurate and responsive near-term risk forecasts. Since risk levels over several month horizons are very volatile, near-term risk prediction requires that a model monitor risk levels closely, and adjust quickly to changes. Daily data permit JPE3S to emphasize information from the most recent quarter without relying on an unacceptably sparse set of returns. The factor structure that it inherits from JPE3 helps JPE3S to interpret these data robustly and efficiently.

An important feature of the new model is its inclusion of serial correlation effects—relations between factor returns on different days. Serial correlations determine how risk aggregates from single days to longer intervals. As is documented in Section 3, models that omit serial correlation effects consistently fail to realize the full benefits of daily data. Their forecast performance is appreciably impaired compared to “serial correlation-aware” models.

In keeping with its character as a model for near-term risk, JPE3S forecasts vary more from month to month than forecasts made by JPE3. For managers who hold investment positions less than six months, the ability of JPE3S to follow risk level changes on these timescales will make it their model of choice. Over these relatively short horizons it has an indisputable advantage in accuracy over JPE3. Managers who hold positions for a year or longer are likely to prefer the more conservative forecast adjustments of JPE3, which suppress the effects of short-lived fluctuations, and aim at capturing more persistent changes in the risk environment.

Appendix A

Constructing a Covariance Matrix in the Presence of Serial Correlation

Let $f_{t,t+H}$ denote the factor return to a particular factor over a month, where H is the number of trading days in a month. In terms of daily factor returns, it can be approximated as

$$f_{t,t+H} \approx \sum_{i=0}^{H-1} f_{t+i,t+i+1}\,,$$

where $f_{t,t+i+1}$ denotes the factor return over day $t+i+1$. The monthly factor variance for factor f is given by:

$$\sigma^2_{f,t,t+H} = H\sigma^2_{f,t,t+1}\left[1+2\left(1-\frac{1}{H}\right)\phi_{ff,1}+2\left(1-\frac{2}{H}\right)\phi_{ff,2}+\ldots\right]. \qquad \text{(A1)}$$

Here $\sigma^2_{f,t,t+H}$ is the H-period variance, $\sigma^2_{f,t,t+1}$ is the daily variance and $\phi_{ff,k}$ is the k^{th} order serial correlation for factor f, i.e., it is the correlation between $f_{t,t+1}$ and f_{t+k+1}. Note that in this formula for the monthly factor variance, serial correlations of order up to H-1 appear. For example, if the number of trading days in a month is 21, then serial correlations with 20 successive trading days must be evaluated.

If we believe that after N days further serial correlations are unimportant, we can modify equation A1 to become

$$\sigma^2_{f,t,t+H} = H\sigma^2_{f,t,t+1}\left[1+2\left(1-\frac{1}{N+1}\right)\phi_{ff,1}+2\left(1-\frac{2}{N+1}\right)\phi_{ff,2}+\ldots\right]. \qquad \text{(A2)}$$

Essentially, we are aggregating over N+1 days instead of over a month. We then treat successive N+1 day intervals as statistically independent, and aggregate the N+1 day variance to a monthly horizon. In JPE3S, N=10.

The monthly horizon covariance between two factors *f* and *g* is given by:

$$\sigma_{fg,t,t+H} = H\sigma_{f,t,t+1}\sigma_{g,t,t+1}\left[\rho_{fg}+\left(1-\frac{1}{N+1}\right)\left(\phi_{fg,1}+\phi_{gf,1}\right)+\left(1-\frac{2}{N+1}\right)\left(\phi_{fg,2}+\phi_{gf,2}\right)+\ldots\right] \qquad \text{(A3)}$$

Here, $\rho_{fg} = corr(f_{t,t+1}, g_{t,t+1})$ and $\phi_{fg,k} = corr(f_{t,t+1}, g_{t+k,t+k+1}), k > 0$. Note that $\phi_{fg,k}$ is not necessarily equal to $\phi_{gf,k}$. In the JPE3S model, the daily standard deviations $\sigma_{f,t,t+1}$ and $\sigma_{g,t,t+1}$ are calculated from daily factor return variances. These variances are exponentially weighted moving averages, evaluated with a 90-day half-life.

The remaining ingredient needed to construct the covariance forecast is the “pseudo-correlation” C_{fg}:

$$C_{fg} = \rho_{fg} + \left(1 - \frac{1}{N+1}\right)\left(\phi_{fg,1} + \phi_{gf,1}\right) + \left(1 - \frac{2}{N+1}\right)\left(\phi_{fg,2} + \phi_{gf,2}\right) + \ldots \qquad \text{(A4)}$$

If serial correlations across factors are generally positive, then C_{fg} is greater than ρ_{fg}. Conversely, if serial correlations across factors are generally negative, C_{fg} is less than ρ_{fg}. The pseudo-correlation C_{fg} is estimated with an exponentially weighted moving average, but in contrast with the factor standard deviations uses a 480-day half-life. Thus, JPE3S adjusts very rapidly to changes in factor risk levels, but preserves a larger historical sample when estimating the effects of cross-factor and serial correlations.

Appendix B

Measuring the Accuracy of Risk Forecasts Using the Bias Statistic

The technique that we use for evaluating the risk model is the bias test. To briefly recap this test, suppose we want to evaluate the model's ability to forecast risk. We collect a history of returns, $\{r_t, t = 1, \cdots, T\}$, and a history of risk forecasts, $\{\hat{\sigma}_t, t = 1, \cdots, T\}$. Note that the return, r_t, is measured *over* the period t, while the forecast risk, $\hat{\sigma}_t$, is measured at the *start* of period t. The standardized return, denoted Z_t, is computed as follows:

$$Z_t = \frac{r_t}{\hat{\sigma}_t}$$

If the model forecasts risk accurately, then the standard deviation of the time series of $\{Z_t\}$ will be close to one, subject to sampling variability. If the standard deviation of $\{Z_t\}$ is significantly below one, then the model has over-predicted risk. Conversely, if the standard deviation of $\{Z_t\}$ is significantly above one, then the model has under-predicted risk.

It can be shown more formally that under the null hypothesis that the model is unbiased, assuming that the sample size is large and returns are normally distributed, the standard deviation of $\{Z_t\}$ will lie in the interval $\left[1-\sqrt{2/T}, 1+\sqrt{2/T}\right]$ approximately 95% of the time. This interval represents the 95% confidence interval for the bias statistic. When the computed bias statistic lies within this interval, the null hypothesis that the model is unbiased cannot be rejected. Conversely, when the computed bias statistic lies outside this interval, the null hypothesis that the model is unbiased is rejected.