Assessing the Risk in Risk Assessments

by

Stephen Figlewski
Professor of Finance
New York University Stern School of Business

Presentation for IAFE / PRMIA
April 23, 2003

Overview of Presentation

Introduction - Measuring Financial Risk Exposure

Forecasting the Volatility of Returns

Predicting the Tails of the Returns Distribution

Simulating the Effect of Model Drift

Other Simulation Results

Concluding Comments

Common Measures of Market Risk Exposure

Probability of Loss -- The tail of the returns distribution gives the probability α for a given loss level L that the realized return will be L or worse: $\alpha = P(r \leq L)$.

Value at Risk (VaR) -- Given the holding period (1 day, 10 days) and probability level α, Value at Risk is the return such that the probability of a worse return over that period is no more than α.

Conditional Value at Risk (C-VaR) -- C-VaR is the expected value of the return, conditional on it being a larger loss than the VaR. (Also called "Expected Shortfall" or "Excess of Loss")

All of these measures of risk exposure require selecting a probability distribution and estimating (forecasting) its parameters.

©2003 Figlewski

Plain Vanilla Risk Estimation

Suppose we have a sample of returns $\{r_1, r_2, ..., r_T\}$ and we want to estimate the risk exposure of r_{T+1}.

The most common assumption is that the generating function for security returns is normal with mean μ and volatility σ (prices are lognormal). These parameters must be estimated.

Classical Statistics

$$\hat{\mu} = \frac{1}{T} \sum_{t=1}^{T} r_t$$

$$\hat{\sigma} = \sqrt{\frac{1}{T-1} \sum_{t=1}^{T} (r_t - \hat{\mu})^2}$$

Plain Vanilla Risk Estimation, p.2

Unfortunately, the classical estimator for the mean is grossly inaccurate.

The standard error of the sample mean as an estimate of the true mean does not depend upon the frequency of observation (e.g., using daily returns does not give a more accurate estimate than monthly returns).

The standard error of estimate is equal to $\sigma / \sqrt{T}$, where σ is the annual volatility and T is the time span of the sample data.

Example: Suppose the (annualized) true μ and σ are: $\mu = 10\%$ and $\sigma = 20\%$.
The standard error of the sample mean from a 3-month sample (1/4 year) of daily data will be $20\% / \sqrt{(1/4)} = 40\%$.

All the sample data tell us is that about 95% of the time the sample mean will be somewhere between -70% and +90%.

Plain Vanilla Risk Estimation, p.3

The Not-So-Classical Real World Solution: Set the mean to zero

$$\hat{\sigma} = \sqrt{\frac{1}{T}\sum_{t=1}^{T} r_t^2}$$

Plain Vanilla Risk Estimation, p.4

The standard Value at Risk calculation: Treat the mean as being zero. Multiply the estimated volatility of the returns distribution by the 5% or 1% critical value for the tail of a normal distribution:

$$5\%\ \text{VaR} = 1.645\ \sigma_{est}$$

$$1\%\ \text{VaR} = 2.326\ \sigma_{est}$$

Estimation error: This ignores the fact that σ_{est} is only an estimate of the true σ, subject to estimation error.

- Under the classical statistics estimator, the distribution of r_{T+1} given $\{r_1, r_2, ..., r_T\}$ is not normal. It is proportional to a Student-t with T-1 degrees of freedom, but the variance is (T+1)/T times greater because of the sampling error.
- Under the zero mean constraint, the forecast error is also Student-t if the true mean is zero.
- But if the true mean is not zero, the estimator is biased (even though it may be more accurate in terms of RMSE) and the distribution is not simply a Student-t.

©2003 Figlewski

Accurately Forecasting Volatility is Not Easy

RMSE of forecast of 1 month volatility using 3, 12, and 60 months of past data, and an exponentially weighted moving average (EWMA) with decay factor fitted to past data.

all figures in annualized percent

	S&P 500 Index	3-Month LIBOR	10 Year US Treasury yield	Deutschemark exchange rate
Data Sample	Dec.1975 - Jan. 1996	Nov. 1979 - Jan. 1996	Jan. 1976 - Jan. 1996	Jan. 1976 - Jan. 1996
Average realized volatility	12.9	20.8	12.7	10.0
3 months	7.2	10.2	4.6	3.9
12 months	7.5	10.2	4.8	3.8
60 months	7.8	13.0	5.7	4.0
EWMA	7.0	9.8	4.7	4.0
average decay	.979	.973	.976	.981

Source: Green and Figlewski. "Market Risk and Model Risk for a Financial Institution Writing Options," Journal of Finance Aug. 1999.

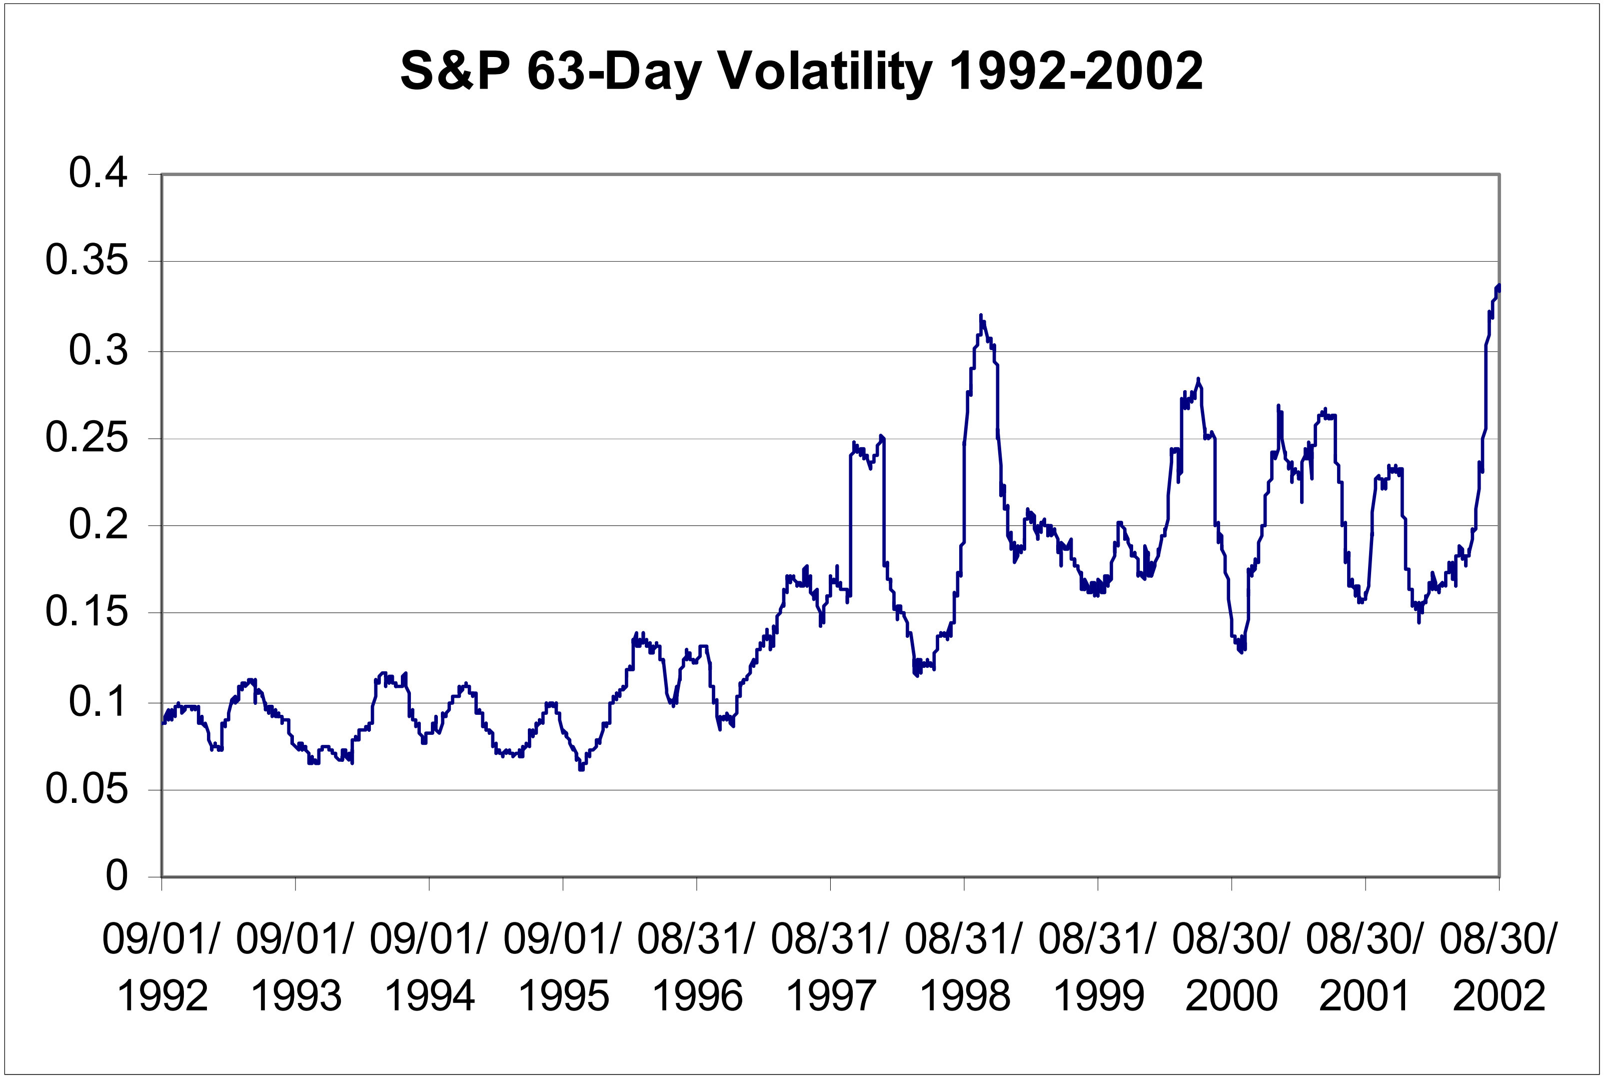
S&P 63-Day Volatility 1992-2002
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
09/01/ 1992
09/01/ 1993
09/01/ 1994
09/01/ 1995
08/31/ 1996
08/31/ 1997
08/31/ 1998
08/31/ 1999
08/30/ 2000
08/30/ 2001
08/30/ 2002

Simulating Estimated Risk Exposure under Ideal Conditions

We assume returns are generated by the model

$$\frac{dS}{S} = \mu\, dt + \sigma\, dz$$

For the simulation, we discretize this as

$$r_{t+1} = \ln\left(S_{t+1}/S_t\right) = \mu\,\Delta t + \sigma\,\tilde{z}_t\sqrt{\Delta t}\ ; \qquad \tilde{z} \sim N(0,1)$$

- Simulation: Sequential returns are simulated for a period of 2,500,000 days (10,000 years).
- Estimation sample: K = 63 days; 39,682 non-overlapping intervals.
- True Volatility: $\sigma = 0.20$; True mean: $\mu = 0$; Sample mean is not estimated
- All tail statistics are reported in standard deviations.

Table 3: Constant Volatility Baseline Simulation

Tail statistics are reported in standard deviations.

RMSE of volatility estimate = 0.0178

Probability α	α-Tail cutoff for Normal	Actual α-tail cutoff	True Prob / Predicted Prob
0.10	-1.282	-1.285	1.01
0.05	-1.645	-1.649	1.01
0.01	-2.326	-2.384	1.17
0.005	-2.576	-2.660	1.23
0.002	-2.878	-3.015	1.46
0.001	-3.090	-3.297	1.76
0.0005	-3.290	-3.514	2.03
0.0002	-3.540	-3.777	2.21
0.0001	-3.719	-4.144	2.37

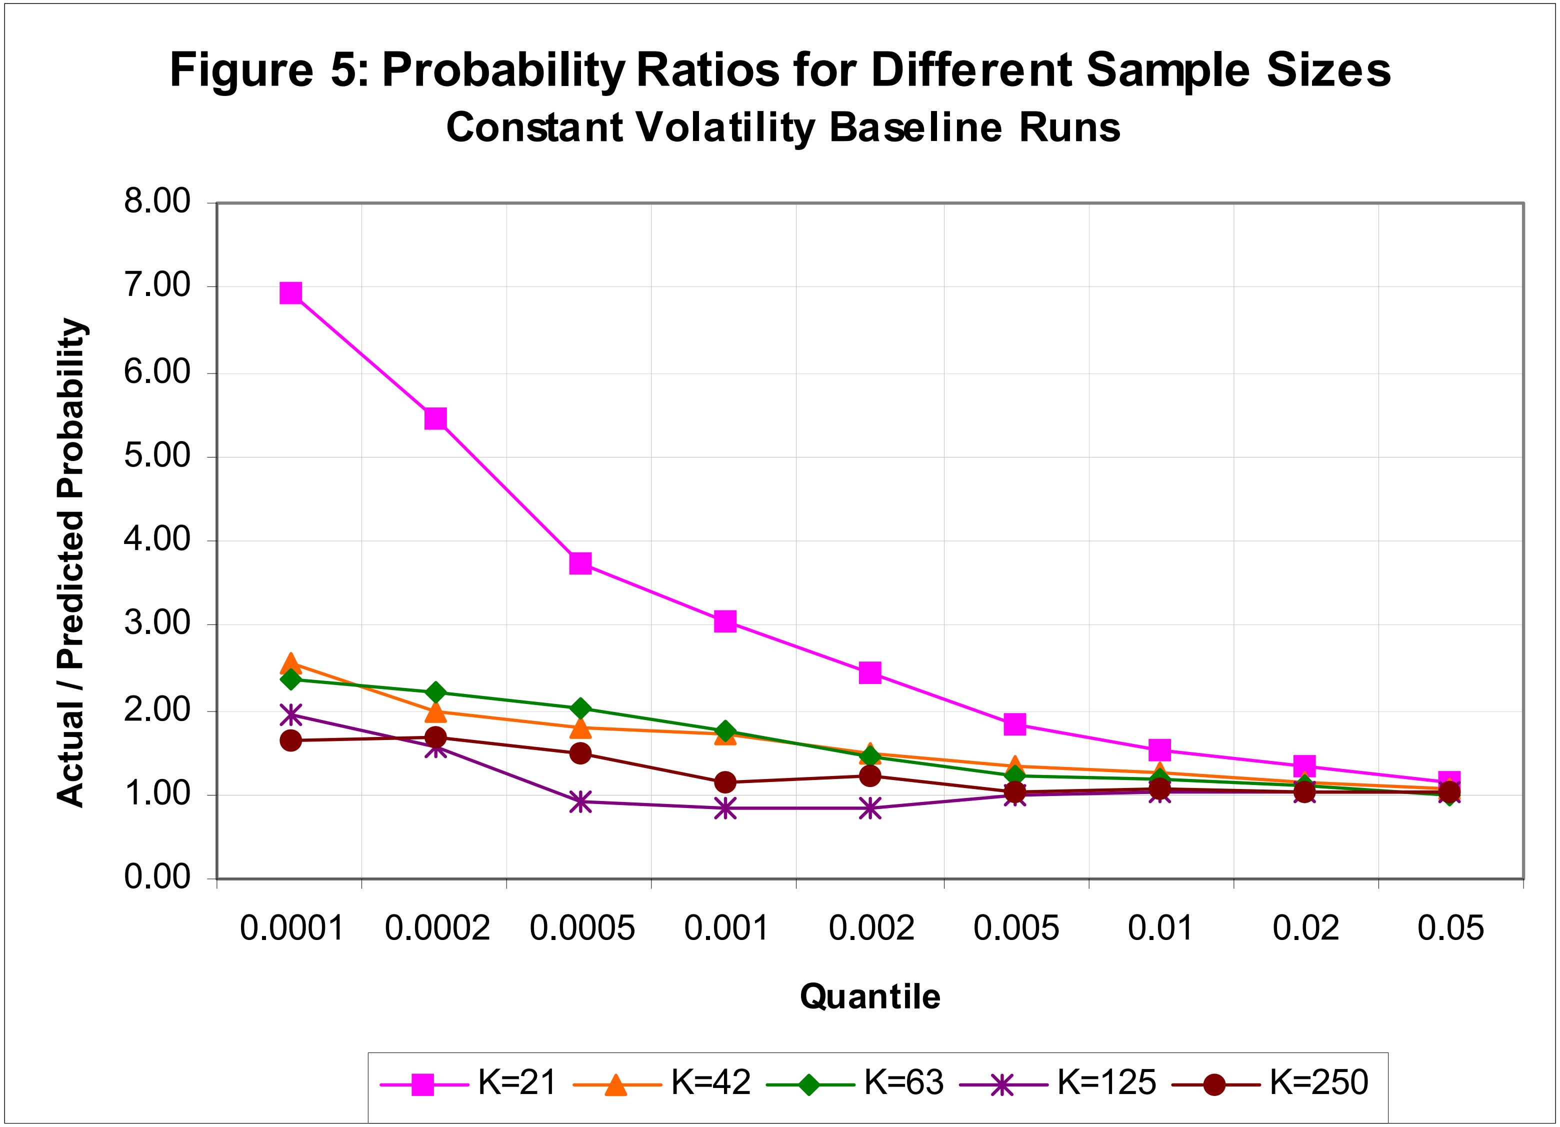
Figure 5: Probability Ratios for Different Sample Sizes
Constant Volatility Baseline Runs
Actual / Predicted Probability
8.00
7.00
6.00
5.00
4.00
3.00
2.00
1.00
0.00
0.0001
0.0002
0.0005
0.001
0.002
0.005
0.01
0.02
0.05
Quantile
K=21
K=42
K=63
K=125
K=250

Realized Tail Events for the Standard and Poor's 500 Index

We simulate the same risk estimation strategy as examined above using historical returns on the S&P 500 stock index, from July 2, 1962 - Aug. 30, 2002. We tabulate the occurrence of tail events and compute the realized / predicted probability ratios for various critical values.

Sample		Tail Probability								
		.05	.02	.01	.005	.002	.001	.0005	.0002	.0001
21-day	Events predicted	505	202	101	50	20	10	5	2	1
	Actual events	587	304	195	128	85	65	50	42	30
	Probability ratio	1.16	1.51	1.93	2.54	4.21	6.44	9.91	20.81	29.73
63-day	Events predicted	502	201	100	50	20	10	5	2	1
	Actual events	548	283	184	130	75	59	43	30	22
	Probability ratio	1.09	1.41	1.83	2.59	3.73	5.87	8.56	14.93	21.89
250-day	Events predicted	493	197	99	49	20	10	5	2	1
	Actual events	507	256	171	114	77	56	43	29	22
	Probability ratio	1.03	1.30	1.73	2.31	3.90	5.68	8.72	14.70	22.31

©2003 Figlewski

Simulating Risk Exposure When Volatility is Stochastically Time-Varying

We assume volatility varies according to a standard square root process, but it is estimated for risk assessment purposes using the zero mean standard (constant volatility) estimator.

Theoretical Model:

$$\frac{dS}{S} = \mu\, dt + \sigma_t\, dz$$

$$\sigma_t = \sqrt{V_t}$$

$$dV = \kappa(\bar{V} - V)\, dt + \theta \sqrt{V}\, dw$$

$\bar{V}$ is the long run variance,

κ is the rate of reversion of the variance toward that long run value,

$\theta\sqrt{V_t}$ is the volatility of the variance process, and

dw is a second Brownian motion, independent of dz.

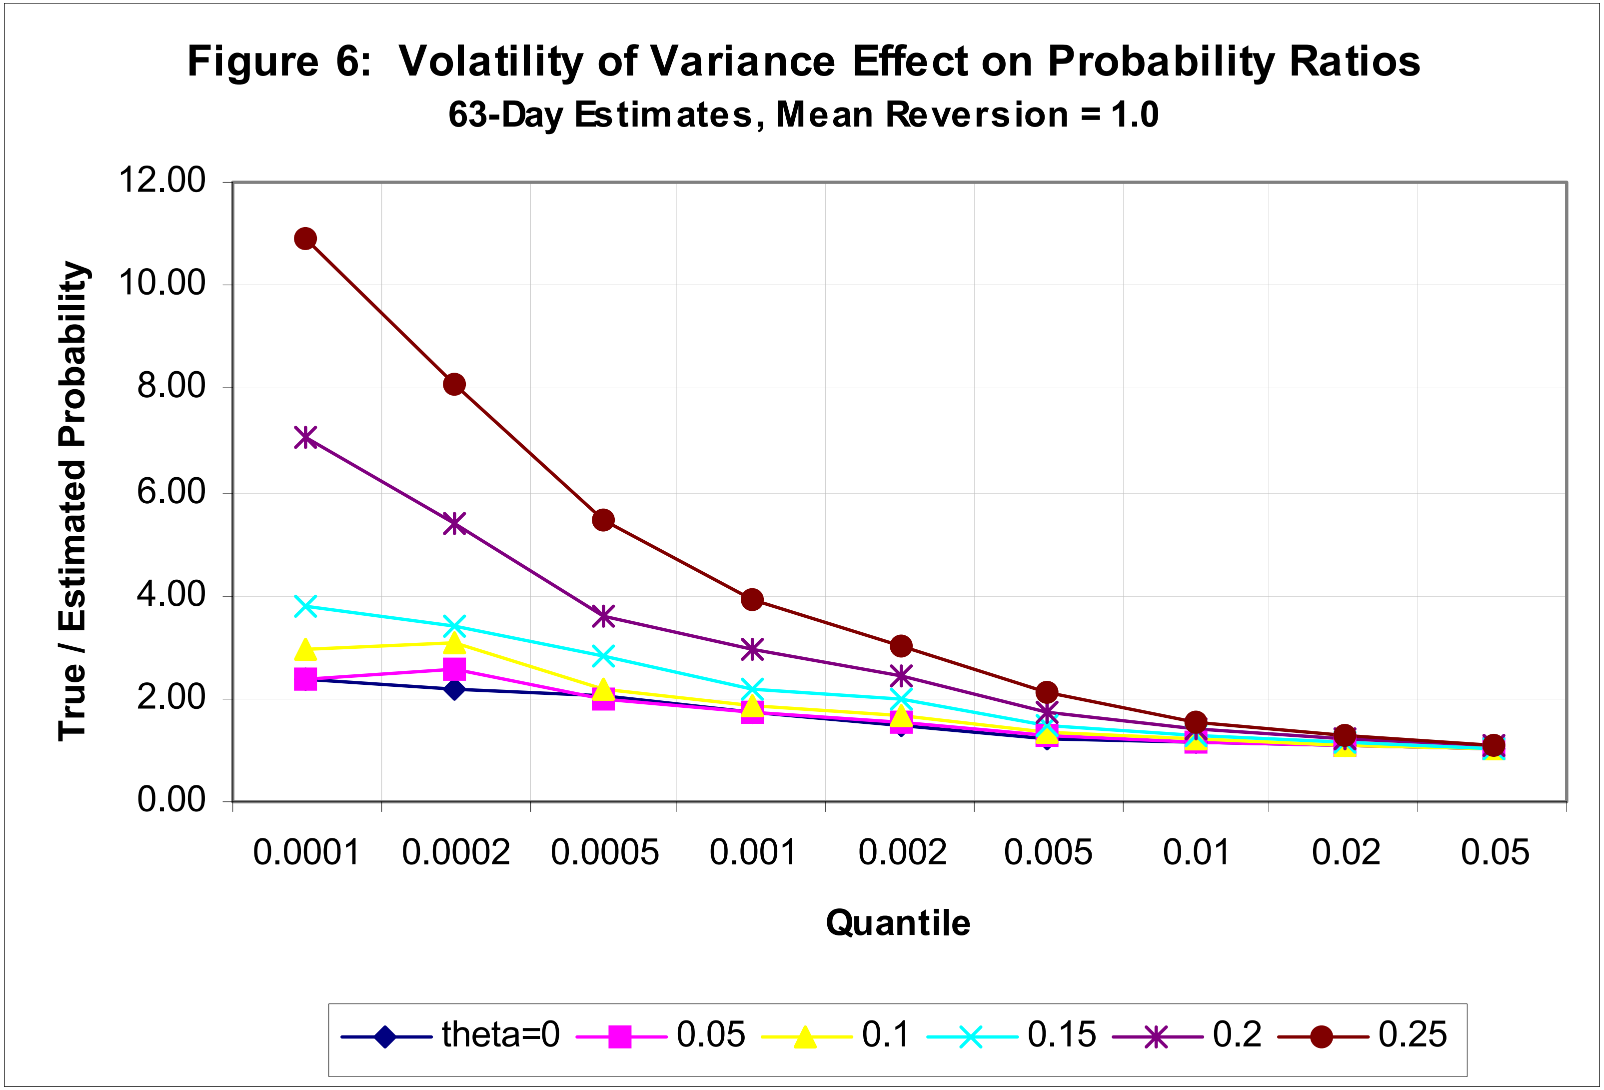
Figure 6: Volatility of Variance Effect on Probability Ratios
63-Day Estimates, Mean Reversion = 1.0
True / Estimated Probability
12.00
10.00
8.00
6.00
4.00
2.00
0.00
0.0001
0.0002
0.0005
0.001
0.002
0.005
0.01
0.02
0.05
Quantile
theta=0
0.05
0.1
0.15
0.2
0.25

Simulating the Effect of Model Drift

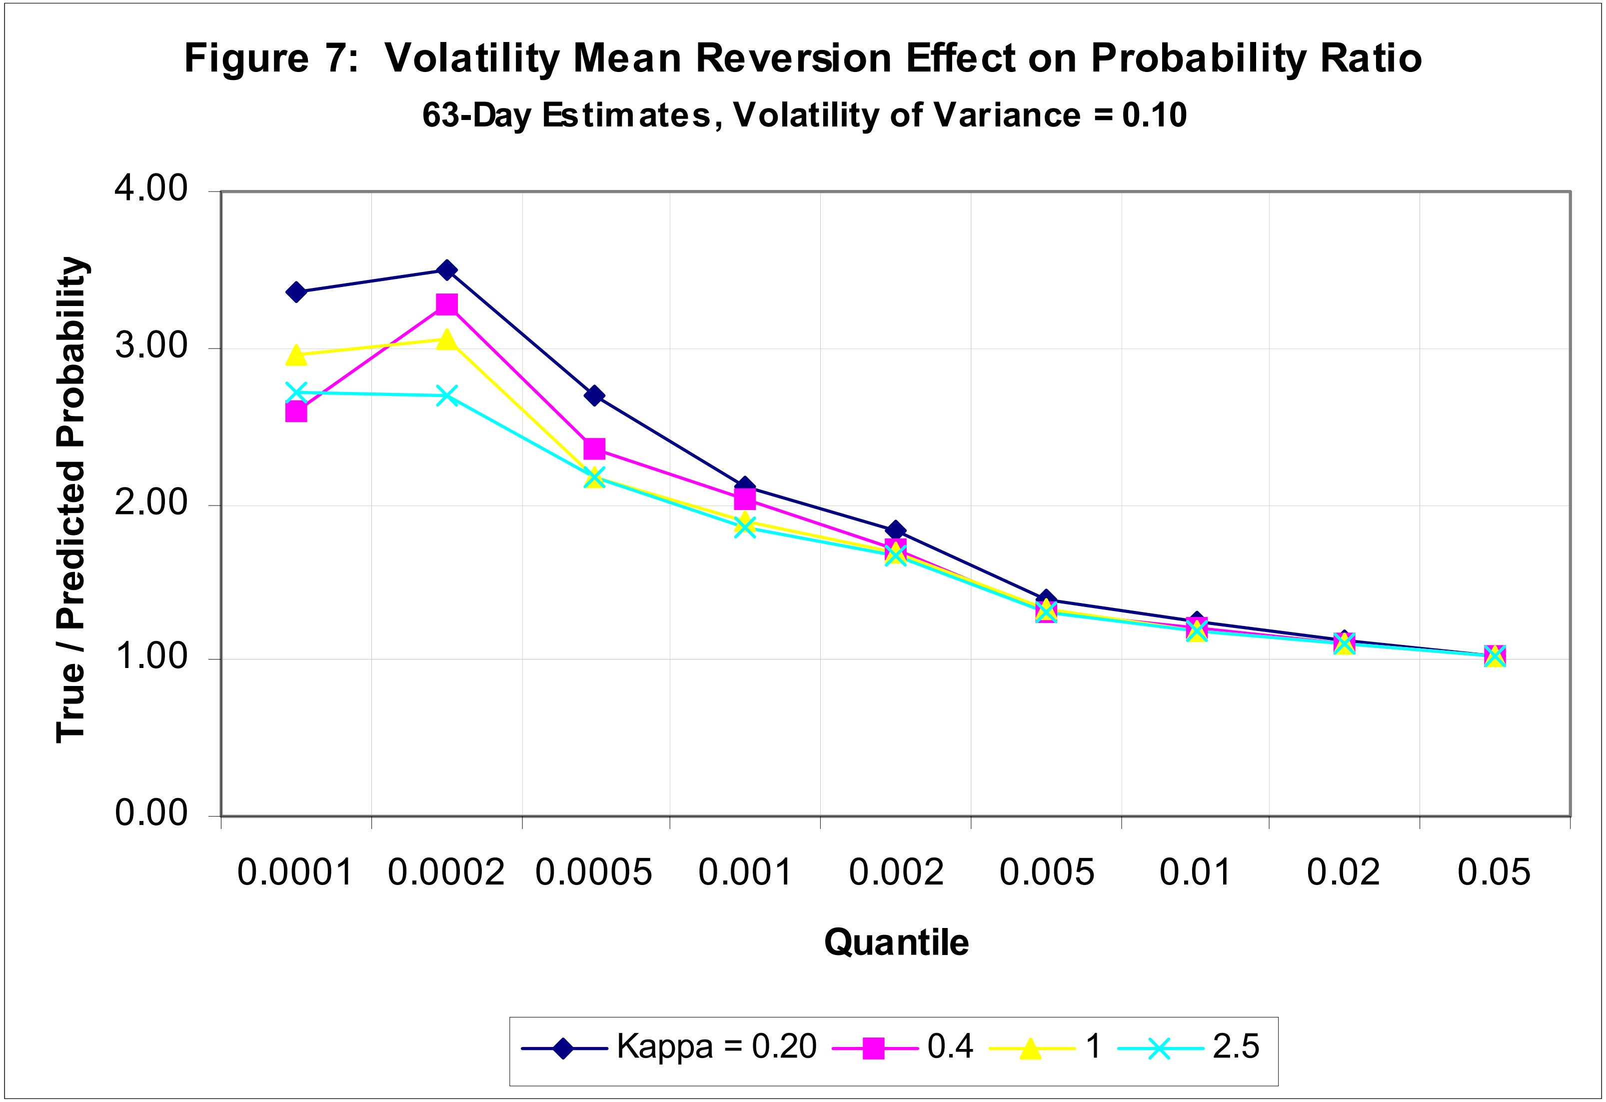

Simulating the Effect of Model Drift

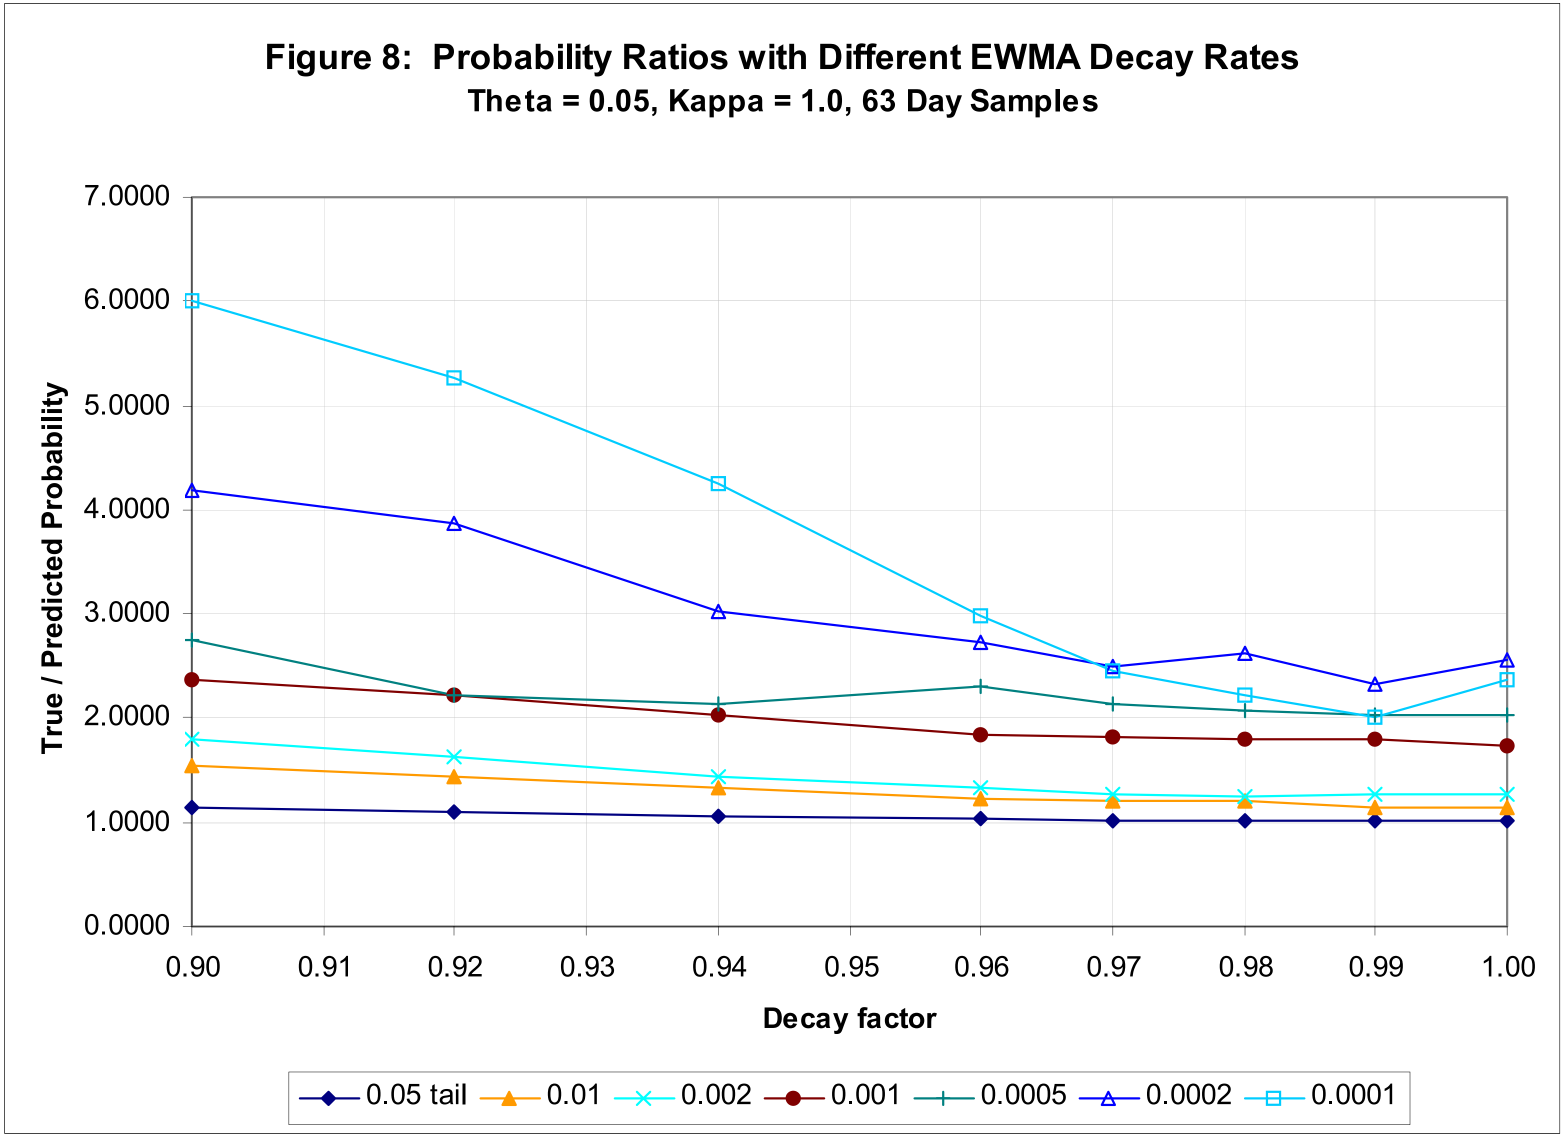

Simulating the Effect of Model Drift

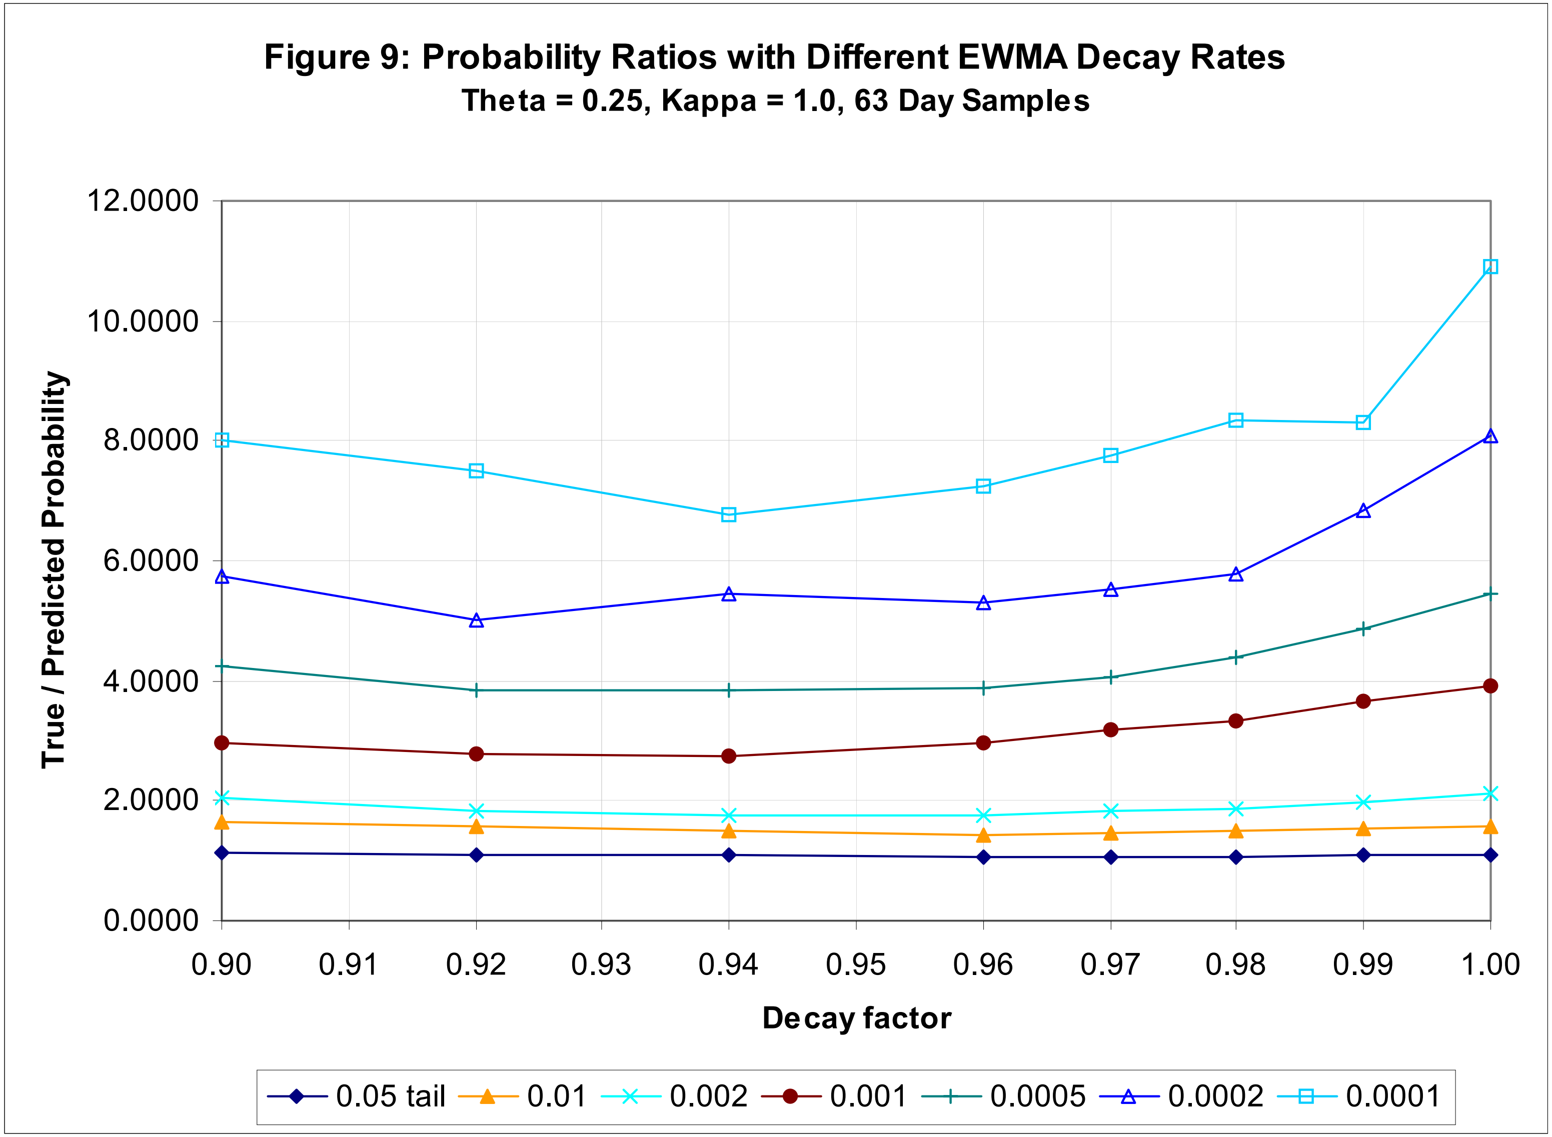

Other Important Issues to be Explored

Fat-tailed shocks

Rolling sample estimation

Historical simulation as an alternative way to fit tail probabilities

Jumps / regime shifts

Sample probability ratios in Baseline runs using t(7) and t(4) return shocks

Tail prob	Normal	t(7)	t(4)
0.05	1.01	0.98	0.91
0.01	1.17	1.61	1.83
0.002	1.46	3.43	4.64
0.001	1.76	5.39	7.68
0.0005	2.03	8.49	12.75
0.0002	2.21	14.72	26.84
0.0001	2.37	23.21	45.79

©2003 Figlewski

Other Simulation Results

Rolling Sample Results

Location of specified tail of the returns distribution (in standard deviations).

2,500,000 simulated days

Nonoverlap: volatility is estimated once every 63 days, from independent 63-day samples;

Overlap: volatility is estimated daily, using overlapping 63-day moving windows

		Normal	Baseline	Run1	Run2	Run3	Run4	Run5	Run6	Run7	Run8	Run9
Theta		--------------	0.00	0.05	0.10	0.05	0.10	0.10	0.20	0.10	0.20	0.40
Kappa		--------------	0.00	0.20	0.20	0.40	0.40	1.00	1.00	2.50	2.50	2.50
	Nonoverlap	--------------	0.0178	0.0193	0.0230	0.0192	0.0229	0.0227	0.0331	0.0222	0.0318	0.0545
RMSE	Overlap	--------------	0.0177	0.0188	0.0233	0.0192	0.0223	0.0224	0.0328	0.0219	0.0313	0.0540
	Nonoverlap	-1.645	-1.649	-1.654	-1.663	-1.654	-1.659	-1.660	-1.684	-1.658	-1.674	-1.728
0.0500	Overlap	-1.645	-1.660	-1.802	-2.169	-1.704	-1.987	-1.721	-2.039	-1.699	-1.797	-2.232
	Nonoverlap	-2.326	-2.384	-2.388	-2.414	-2.385	-2.398	-2.302	-2.504	-2.392	-2.455	-2.702
0.0100	Overlap	-2.326	-2.384	-2.707	-3.889	-2.513	-3.199	-2.560	-3.390	-2.248	-2.720	-3.890
	Nonoverlap	-2.878	-3.015	-3.036	-3.107	-3.036	-3.091	-3.065	-3.239	-3.057	-3.183	-3.693
0.0020	Overlap	-2.878	-2.998	-3.523	-6.482	-0.201	-4.504	-3.315	-4.967	-3.142	-3.621	-5.737
	Nonoverlap	-3.090	-3.297	-3.295	-3.401	-3.294	-3.383	-3.627	-3.580	-3.324	-3.448	-4.000
0.0010	Overlap	-3.090	-3.200	-3.907	-7.924	-3.483	-5.133	-3.909	-5.710	-3.388	-4.054	-6.686
	Nonoverlap	-3.290	-3.514	-3.538	-3.632	-3.534	-3.596	-3.750	-3.810	-3.587	-3.715	-4.397
0.0005	Overlap	-3.290	-3.458	-4.306	-9.454	-3.749	-3.770	-4.282	-6.501	-3.622	-4.385	-7.499
	Nonoverlap	-3.540	3.777	-3.789	-3.912	-3.780	-3.770	-3.750	-4.225	-3.776	-3.944	-4.846
0.0002	Overlap	-3.540	-3.719	-4.786	-12.070	-4.108	-6.687	-4.282	-7.884	-3.969	-4.871	-8.884
	Nonoverlap	-3.872	-4.144	-3.928	-4.103	-3.942	-3.933	-3.915	-4.472	-3.957	-4.520	-5.261
0.0001	Overlap	-3.719	-3.949	-5.102	-14.445	-4.925	-7.125	-4.662	-8.760	-4.223	-5.310	-10.283

Conclusions

Estimation risk is true risk. It increases effective risk exposure.

Estimation risk is unavoidable

- Finance is not physics
- Economic relationships are not stationary
  - technological change
  - evolution of financial institutions
  - economic agents learn from experience
  - they may even learn from financial research

Conclusions

Suggestions for Dealing with Estimation Risk

- Be aware of parameter uncertainty as a significant risk factor
- Use robust methods and models
- Out of sample testing is essential
- Critically examine stability of model parameters
- Worry about estimation techniques that require a lot of data from a stationary distribution to achieve asymptotic performance
- Worry about estimation techniques that depend heavily on properties of an assumed returns distribution
- Don't be afraid to be Bayesian

 ©2003 Figlewski

Risk Evaluation Practices in the Real World

What are the most common approaches to estimating market risk exposure?

- parametric VaR (how?)
- historical simulation (how?)
- stress testing (where do scenarios come from?)

How do you adjust for nonlognormal returns?

- fat tails
- jumps
- regime changes
- nonlinear payoffs
- etc.

Do you try to adjust for estimation error and time-varying parameters? How?

Where are the biggest problems with estimating market risk exposures in practice?

©2003 Figlewski