Stability-Adjusted Portfolios

Mark Kritzman and David Turkington

Mark Kritzman
is the CEO at Windham Capital Management in Boston, MA, and a senior lecturer at MIT Sloan School of Management in Cambridge, MA.
kritzman@mit.edu

David Turkington
is a senior vice president at State Street Associates in Cambridge, MA.
dturkington@statestreet.com

The instability of risk is one of the most daunting challenges in forming optimal portfolios. Investors typically attempt to improve risk estimates by compressing them toward a cross-sectional average or some other prior belief. We propose a different approach. We introduce a protocol for quantifying the degree of estimation error in covariances, and we show that it varies across assets. We argue that rather than compressing estimates of covariance toward each other, investors should incorporate their relative stability directly into the process of forming portfolios—which will likely render these estimates less similar to each other, though not necessarily.

We proceed by describing the sources of instability in estimating covariances. We then explain how to generate a multivariate return distribution that reflects the relative stability of covariances. Next, we describe how to derive optimal portfolios from a stability-adjusted return distribution. Finally, we show how these stability-adjusted portfolios differ from portfolios that are blind to estimation error, as well as those that rely on Bayesian shrinkage.

SOURCES OF ESTIMATION ERROR

Investors often rely on multidecade samples of historical data to forecast covariances over shorter future periods, such as a few years. These forecasts are subject to three sources of estimation error.[1] First, small-sample error arises when covariances of a long sample are used to forecast covariances of a specific smaller sample. Even though the true covariances of a long sample are known, the realization of those covariances in shorter subsamples may be meaningfully different. Second, independent-sample error arises when known covariances from one sample are projected onto a separate, independent sample. Third, interval error arises when covariances of high-frequency returns, such as monthly returns, differ from covariances of longer-period returns, such as five-year returns.[2] Exhibit 1 illustrates these three sources of error.

CONSTRUCTING A STABILITY-ADJUSTED RETURN DISTRIBUTION

We now describe how to construct a return distribution that accounts for a composite measure of estimation error comprising small-sample error, independent-sample error, and interval error.

1. First, we select a large sample of returns for the assets under consideration.
2. We then select a subsample from this large sample and compute its covariance matrix based on returns of the same interval as our investment horizon.[3]

Exhibit 1
Sources of Estimation Error

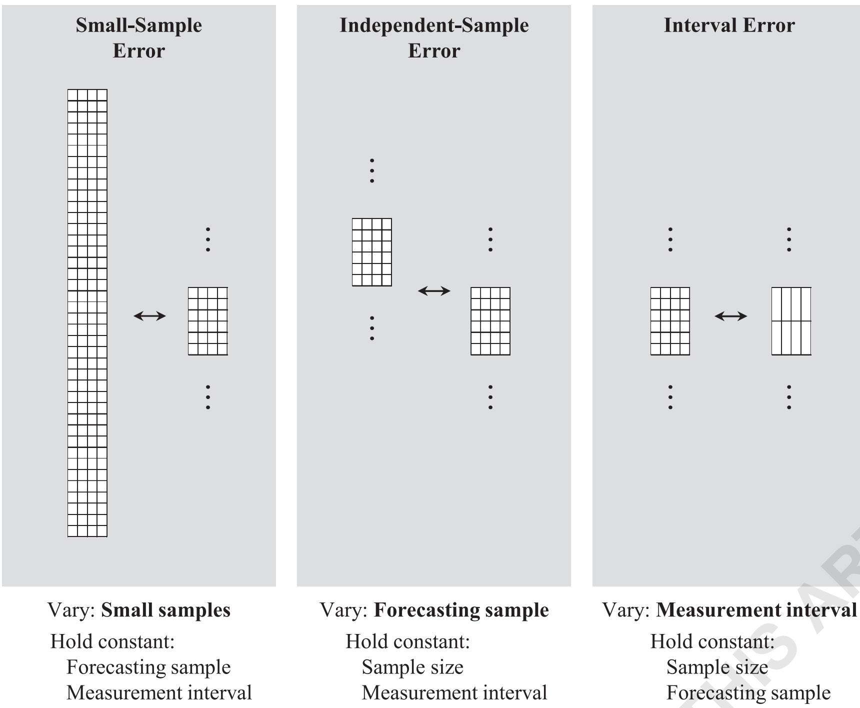

3. We then subtract the subsample covariances from the covariances estimated from the remaining observations of the large sample. These differences represent a composite error comprising small-sample error, independent-sample error, and interval error.
4. Next, we select a new subsample that partly overlaps with the first subsample, and we again compute the differences between the subsample covariances and the covariances estimated from the remaining observations of the large sample.[4]
5. We continue in this fashion until we have computed errors in covariances from all overlapping subsamples.[5]
6. Next, for all subsamples, we add the errors to the covariances of a base-case sample, which, for example, could be the median subsample.[6,7] Then, assuming normality, we generate simulated return samples from each error-adjusted covariance matrix.
7. Finally, we combine these return samples into a new large sample, which can be thought of as a stability-adjusted return sample.

We should note several features of this process. First, the composite errors incorporate all three sources of error. They reflect small-sample error, because the subsamples are smaller than the original sample. They reflect independent-sample error, because each subsample is distinct from the remaining observations in the large sample. And they also capture interval error, because the subsample covariances are estimated from longer-interval returns than those used to estimate the large-sample covariances.

We should also note that the resultant return distribution will not be normal—despite the distributions of the subsamples as well as the Central Limit Theorem. We should expect the stability-adjusted return distribution to have fatter tails than a normal distribution. The Central Limit Theorem states that the sum of independent random variables, which themselves do not need to be individually normally distributed, will approach normality as the quantity of random variables increases.[8] But we are not summing random variables. We are combining distributions.

For example, suppose a particular asset's daily returns for a given month are approximately normally distributed around a mean of 0.5%. And suppose their returns in the following month are again approximately normal, but this time around a mean of −0.5%. If we sum these daily returns for the first day of the two months, the second day of the two months, etc., the 20 summed observations will also be normally distributed, but around a mean of 0.0%. However, the 40 daily returns for the two-month period will not be normally distributed. They will have a bimodal distribution with some observations clustering around a peak of 0.5% and others clustering around a peak of −0.5%.

Finally, we should note that in contrast to Bayesian approaches that compress covariances toward the same prior belief, thereby discounting estimation error, we embrace estimation error and use it to inform a portfolio's composition.

CONSTRUCTING PORTFOLIOS FROM STABILITY-ADJUSTED RETURN SAMPLES

As we have just shown, the process of combining many return distributions, which themselves may or may not be normal, will result in a composite distribution that is not normal—nor is it likely to be elliptical. This poses a challenge for portfolio construction. Mean–variance optimization assumes that either 1) returns are elliptically distributed (of which the normal distribution is a special case) or 2) investors have preferences that can be well approximated by mean and variance. Although most power utility functions, such as the log-wealth utility function, can be reasonably approximated by mean and variance, utility functions that have kinks or inflection points cannot. We therefore employ a technique called *full-scale optimization*[9] to form portfolios.

We proceed as follows:

1. We select a particular utility function, which need not be amenable to approximation by mean and variance.
2. We choose a particular asset mix and apply it every period to the asset returns in the stability-adjusted return sample to compute the utility associated with that asset mix for every period.
3. We then sum utility across all periods and record this value.
4. We then choose a different asset mix and again apply it to the returns to compute its total utility across all periods.
5. We continue in this fashion until we arrive at the portfolio composition that yields the highest utility across all periods.

This full-scale approach to optimization may be computationally expensive;[10] nonetheless, it accounts for every feature of the data, even beyond kurtosis and skewness. It is thus suitable for non-elliptical distributions and for utility functions that cannot be described by mean and variance.

RESULTS

We next test stability-adjusted optimization in two different settings: asset allocation and index replication. We use two utility functions in our full-scale optimizations: 1) a power utility function with a curvature

EXHIBIT 2
Power and Kinked Utility Functions

Panel A: Power Utility Function

$$U(r) = \frac{1}{1-\theta}\left((1+r)^{1-\theta} - 1\right)$$

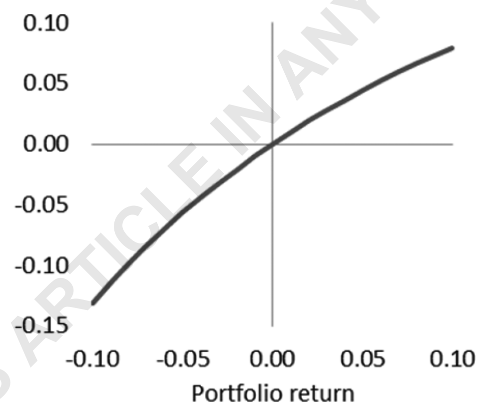

Panel B: Kinked Utility Function

$$U(r) = \begin{cases} \frac{1}{1-\theta}\left((1+r)^{1-\theta} - 1\right), if\ r \geq k \\ \frac{1}{1-\theta}\left((1+k)^{1-\theta} - 1\right) - \omega(k-r), if\ r < k \end{cases}$$

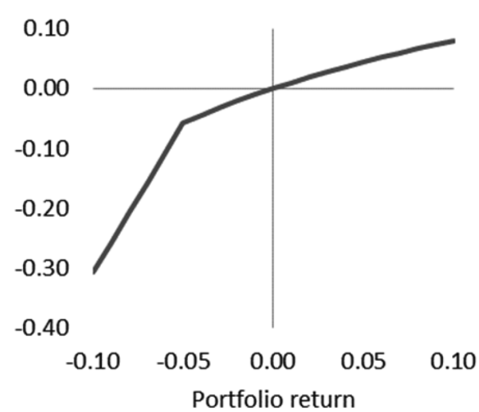

parameter θ equal to 5 and 2) a kinked utility function with a curvature parameter above the kink θ equal to 5, a kink threshold k equal to −5.0% for asset allocation and 0.0% for index replication, and a linear slope ω beneath the kink equal to 5. A kinked utility function is suitable for investors who face thresholds, such as funding requirements or performance hurdles. It is important to note that even if all assets' expected returns are equal and our objective is to minimize portfolio risk, expected returns are still relevant given a kinked utility function because they affect the likelihood of breaching the threshold. These utility functions are shown in Exhibit 2.

Stability-Adjusted Asset Allocation

For our asset allocation example, we search for the optimal portfolio composed of four asset classes: U.S. stocks, U.S. Treasuries, U.S. corporate bonds, and commodities. We assume that the expected returns of these asset classes, shown in Exhibit 3, are known. We set our small-sample window to equal 60 months, and we assume an investment horizon of one year. Hence our interval error will equal the difference between monthly covariances and one-year covariances. We generate a sample of 1,000 returns for each error-adjusted small-sample covariance matrix.

Exhibit 3 also reports the relative stability of the standard deviations and correlations, which make up the covariances. The stability measures are presented as inter-quartile ranges, reflecting small-sample error, independent-sample error, and interval error. We standardize the inter-quartile ranges for standard deviations by dividing them by the standard deviation of the asset class returns. Exhibit 3 reveals that the stability of these

EXHIBIT 3
Asset Class Assumptions and Inter-Quartile Range of Volatility and Correlation

Data and Assumptions			Inter-Quartile Range			
Asset Classes	Data: February 1973–December 2015	Expected Return	Standard Deviations	Correlations		
U.S. Stocks	S&P 500	9%	0.61			
U.S. Treasuries	Barclays U.S. Treasuries	4%	0.66	1.22		
U.S. Corporates	Barclays U.S. Corporate Bonds	5%	0.68	0.93	0.16	
Commodities	S&P/GSCI Commodities	5%	0.83	0.69	0.44	0.74

EXHIBIT 4
Portfolio Weights and Performance

	Power Utility			Kinked Utility		
	Ignoring Errors	Bayesian Shrinkage	Accounting for Stability	Ignoring Errors	Bayesian Shrinkage	Accounting for Stability
Optimal Allocation						
U.S. Stocks	40%	40%	50%	25%	25%	35%
U.S. Treasuries	10%	10%	50%	25%	25%	65%
U.S. Corporates	45%	45%	0%	40%	40%	0%
Commodities	5%	5%	0%	10%	10%	0%
	Annual Volatility			Annual Downside Volatility		
Dispersion of Risk						
90th–10th percentile	9.3%	9.3%	7.2%	11.3%	11.3%	7.3%
Maximum–minimum	11.7%	11.7%	10.0%	15.4%	15.4%	11.6%

Exhibit 5
Realized Volatility for Historical Five-Year Periods

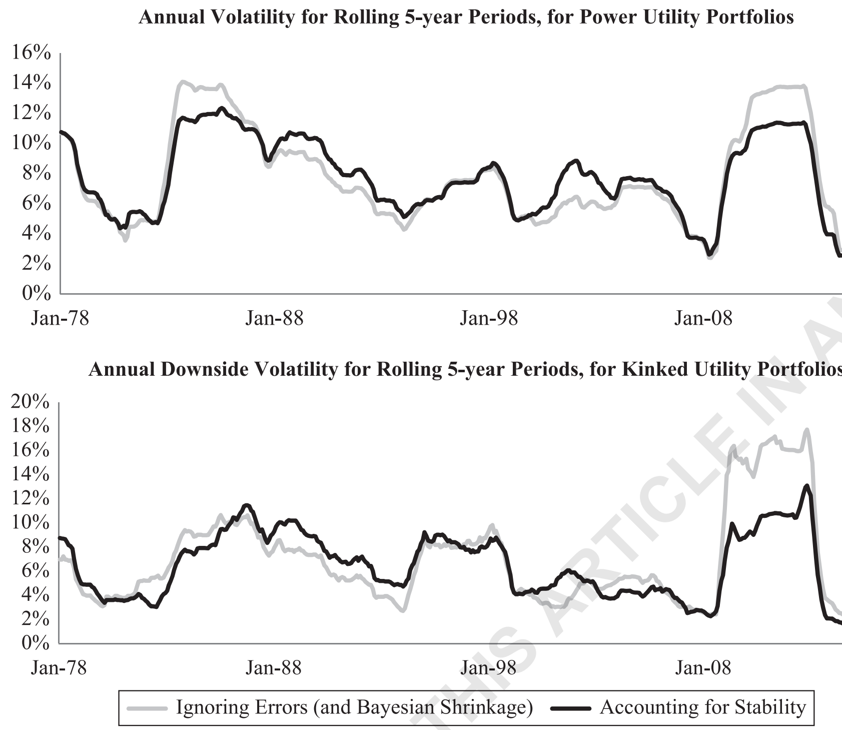

Exhibit 6
Replicating Stocks

Two Randomly Selected Stocks from each Sector	January 6, 2006–January 8, 2016
Stock Name	**Sector**
CARNIVAL	CONSUMER DISCRETIONARY
EXPEDIA	CONSUMER DISCRETIONARY
CHURCH & DWIGHT CO.	CONSUMER STAPLES
BROWN-FORMAN ‘B’	CONSUMER STAPLES
MURPHY OIL	ENERGY
ONEOK	ENERGY
MOODY'S	FINANCIALS
AMERICAN EXPRESS	FINANCIALS
UNITEDHEALTH GROUP	HEALTH CARE
DENTSPLY INTL.	HEALTH CARE
HUNT JB TRANSPORT SVS.	INDUSTRIALS
EQUIFAX	INDUSTRIALS
APPLIED MATS.	INFORMATION TECHNOLOGY
COGNIZANT TECH.SLTN.‘A’	INFORMATION TECHNOLOGY
INTL.FLAVORS & FRAG.	MATERIALS
VULCAN MATERIALS	MATERIALS
CENTURYLINK	TELECOMMUNICATION SERVICES
AT&T	TELECOMMUNICATION SERVICES
EDISON INTL.	UTILITIES
NEXTERA ENERGY	UTILITIES

risk measures varies across asset classes and significantly among correlations.

Exhibit 4 highlights the differences in optimal portfolio weights and performance given both a power utility function and a kinked utility function, depending on whether we ignore errors, shrink them, or account for their relative stabilities.[11] The Bayesian shrinkage approach blends each standard deviation and correlation equally with their cross-sectional averages. To compare these three methods of portfolio construction, we require the optimizations that are blind to errors and Bayesian shrinkage to have the same expected return as the optimal portfolio that accounts for errors. To measure the risk stability of each portfolio, we form a distribution of annual volatilities across all overlapping five-year windows in the historical sample and compute the spread between the maximum and minimum, as well as the spread between the 90th percentile and the 10th percentile of this risk distribution. We focus on downside volatility (derived from semivariance) for the kinked utility because that utility function implies a much larger aversion to negative deviations than positive deviations.

The first notable observation is that accounting for stability has a substantial impact on optimal portfolio weights, whereas Bayesian shrinkage yields the same weights as ignoring errors. Exhibit 4 also reveals that the portfolios that account for the relative stability of covariances have the least dispersion in risk outcomes.[12] Exhibit 5 shows the distribution of realized volatility for each five-year subsample. During the global financial crisis, the risk of Bayesian or error-blind portfolios increased dramatically, but the stability-adjusted portfolios suffered a much smaller increase in risk.

Stability-Adjusted Index Replication

We now evaluate the effect of relative stability on index replication. For this experiment, we randomly select two securities from each sector according to the Global Industry Classification Standard (GICS),[13] and we seek to allocate to these securities for the purpose of minimizing tracking error relative to the S&P 500 Stock Index. We use stability-adjusted weekly returns estimated over the period beginning January 2006 and ending January 2016. We set the subsample windows to equal one year, and we evaluate results over horizons of one quarter. We assume that each security has a known

EXHIBIT 7
Inter-Quartile Range of Volatility and Correlation

		Standard Deviations	Correlations																			
			a	b	c	d	e	f	g	h	i	j	k	l	m	n	o	p	q	r	s	t
a	S&P 500	0.79																				
b	CARNIVAL	0.74	0.73																			
c	EXPEDIA	0.67	0.66	0.56																		
d	CHURCH & DWIGHT CO.	0.47	0.57	0.65	0.77																	
e	BROWN-FORMAN ‘B’	0.44	0.62	0.54	0.71	0.48																
f	MURPHY OIL	0.80	0.55	0.90	0.92	0.49	0.79															
g	ONEOK	0.62	0.65	0.59	0.57	0.63	0.49	0.47														
h	MOODY’S	0.73	0.49	0.84	0.63	0.75	0.57	0.72	0.60													
i	AMERICAN EXPRESS	0.79	0.29	0.47	0.69	0.49	0.44	0.72	0.68	0.48												
j	UNITEDHEALTH GROUP	0.81	0.56	0.51	0.63	0.47	0.48	0.69	0.65	0.78	0.58											
k	DENTSPLY INTL.	0.66	0.53	0.69	0.84	0.54	0.53	0.45	0.69	0.92	0.51	0.72										
l	HUNT JB TRANSPORT SVS.	0.80	0.53	0.65	0.68	0.45	0.71	0.74	0.72	0.85	0.42	0.84	0.61									
m	EQUIFAX	0.78	0.46	0.80	0.65	0.42	0.61	0.86	0.70	0.45	0.56	0.72	0.68	0.47								
n	APPLIED MATS.	0.47	0.40	0.90	0.88	0.76	0.46	0.89	0.76	0.70	0.50	0.51	0.78	0.52	0.66							
o	COGNIZANT TECH.SLTN.‘A’	0.61	0.56	0.34	0.85	0.63	0.68	0.41	0.70	0.84	0.71	0.83	0.46	0.51	0.63	0.72						
p	INTL.FLAVORS & FRAG.	0.34	0.34	0.91	0.73	0.56	0.74	0.61	0.63	0.68	0.49	0.60	0.77	0.68	0.53	0.29	0.57					
q	VULCAN MATERIALS	0.60	0.55	0.52	0.83	0.44	0.70	0.69	0.48	0.75	0.54	0.84	0.76	0.42	0.44	0.78	0.35	0.44				
r	CENTURYLINK	0.47	0.49	0.73	0.65	0.57	0.54	0.80	0.62	0.60	0.51	0.64	0.65	0.78	0.59	0.52	0.68	0.39	0.53			
s	AT&T	0.36	0.50	0.61	0.65	0.61	0.36	0.67	0.53	0.68	0.51	0.59	0.49	0.48	0.64	0.50	0.72	0.42	0.87	0.55		
t	EDISON INTL.	0.38	0.70	0.85	0.74	0.43	0.49	1.05	0.82	1.07	0.48	0.53	0.53	0.52	0.72	0.60	0.71	0.37	0.63	0.72	0.40	
u	NEXTERA ENERGY	0.85	0.66	0.57	0.66	0.40	0.78	0.98	0.74	0.87	0.49	0.61	0.89	0.70	0.72	0.76	0.85	0.46	0.70	0.55	0.58	0.41

EXHIBIT 8
Index-Replication Weights and Performance

	Power Utility			Kinked Utility		
	Ignoring Errors	Bayesian Shrinkage	Accounting for Stability	Ignoring Errors	Bayesian Shrinkage	Accounting for Stability
Optimal Allocation						
CARNIVAL	0%	10%	0%	10%	10%	0%
EXPEDIA	0%	0%	0%	0%	0%	0%
CHURCH & DWIGHT CO.	0%	0%	0%	0%	10%	0%
BROWN-FORMAN 'B'	0%	0%	10%	0%	0%	10%
MURPHY OIL	10%	10%	10%	10%	10%	10%
ONEOK	0%	0%	10%	0%	0%	10%
MOODY'S	10%	10%	0%	10%	10%	0%
AMERICAN EXPRESS	0%	0%	10%	0%	0%	10%
UNITEDHEALTH GROUP	0%	0%	10%	0%	0%	10%
DENTSPLY INTL.	10%	10%	10%	10%	10%	10%
HUNT JB TRANSPORT SVS.	0%	0%	0%	0%	0%	0%
EQUIFAX	10%	10%	10%	10%	0%	10%
APPLIED MATS.	10%	10%	10%	10%	10%	10%
COGNIZANT TECH.SLTN.'A'	0%	0%	0%	0%	0%	0%
INTL.FLAVORS & FRAG.	0%	0%	10%	0%	0%	10%
VULCAN MATERIALS	10%	10%	0%	10%	10%	0%
CENTURYLINK	10%	0%	0%	0%	10%	0%
AT&T	10%	20%	10%	20%	10%	10%
EDISON INTL.	0%	0%	0%	0%	0%	0%
NEXTERA ENERGY	20%	10%	0%	10%	10%	0%
	Quarterly Tracking Error			Quarterly Downside Tracking Error		
Dispersion of Risk						
90th–10th percentile	2.0%	2.1%	0.8%	2.4%	2.1%	1.0%
Maximum–minimum	5.5%	5.0%	1.7%	6.4%	4.8%	2.6%

EXHIBIT 9
Realized Tracking Error for Historical One-Year Periods

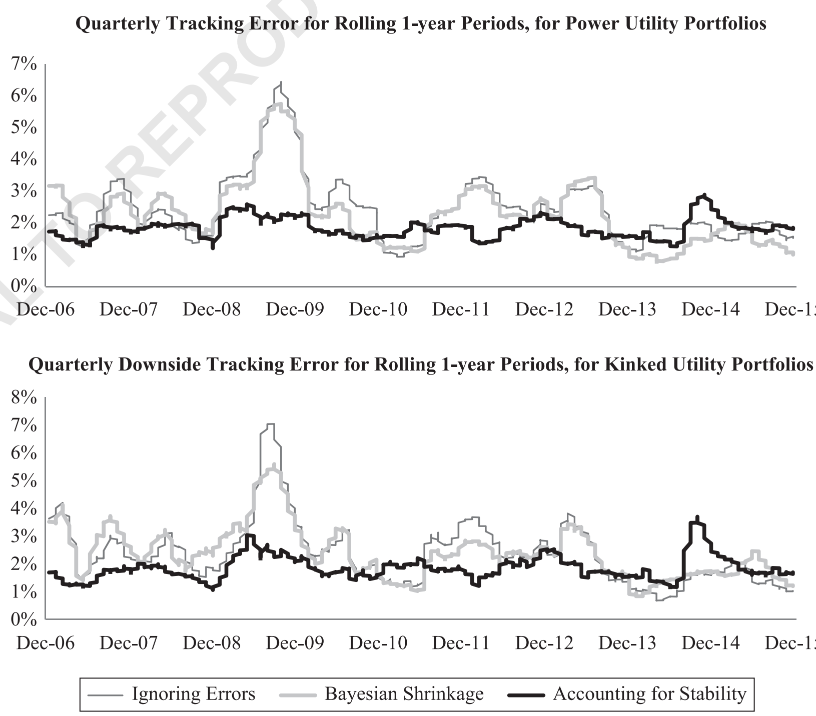

expected relative return equal to 0%, that the securities' weights sum to 100%, and that the weight of the S&P 500 Index is fixed at −100%.[14] We generate a sample of 1,000 returns for each error-adjusted, small-sample covariance matrix. Exhibit 6 lists the stocks selected for this experiment, and Exhibit 7 shows the relative stability of their volatilities and correlations.

Exhibit 8 shows the optimal allocations and dispersion of risk for portfolios that ignore errors, apply Bayesian shrinkage, and account for stability. Exhibit 9 shows realized tracking error for all overlapping one-year periods. It is interesting to note that the portfolios accounting for stability are identical for both power utility and kinked utility functions. In both cases, these portfolios have substantially less dispersion in volatility than the Bayesian shrinkage portfolios or those that ignore errors.

SUMMARY

We have introduced a methodology for incorporating estimation error in covariances into the portfolio formation process. In contrast to Bayesian approaches that attempt to mitigate estimation error by making the estimates more similar to each other, we propose that investors measure the relative stability of covariances and form portfolios that explicitly account for this feature. Our approach will most likely cause covariance estimates to be less similar to each other.

We compute covariances from all independent subsamples of a chosen size and measure composite errors in these subsamples. These composite errors comprise small-sample error, independent-sample error, and interval error. We then add these errors to a base-case covariance matrix and, assuming normality, generate stability-adjusted return distributions for all subsamples. We then combine these distributions into a stability-adjusted return distribution, which we show to be non-normal.

We apply full-scale optimization, which accommodates non-normality as well as complex utility functions, to the stability-adjusted return sample in order to derive optimal portfolios that account for asset-specific estimation error.

Because we account for the relative stability of errors in this way, these portfolio allocations differ significantly from those of portfolios that are blind to estimation error and portfolios that are modified by Bayesian approaches to estimation error. Moreover, portfolios that explicitly account for relative stability tend to have more stable risk than those that ignore errors or shrink them.

APPENDIX

If we are concerned only with errors in standard deviations (and not correlations), it is possible to derive some useful and intuitive analytical results for mixture distributions. Assume that asset *A*'s returns are distributed according to a mixture of normal distributions with a mixing density (distribution of the variance parameter) that follows a scaled inverse chi-squared distribution ($SI\chi^2$). The $SI\chi^2$ distribution is a natural assumption for the distribution of the variance parameter because it does not permit negative values but has a shape that is roughly similar to that of a normal distribution—especially if its own variance is small relative to its mean. The distribution has two parameters: ν and τ^2, which both must be greater than zero. The parameter ν determines the shape of the distribution, while τ^2 scales the distribution. The mean and variance of the $SI\chi^2$ distribution are given by the following:

$$Mean(\sigma^2) = \frac{\nu\tau^2}{\nu - 2}$$

$$Variance(\sigma^2) = \frac{\nu^2\tau^4}{(\nu-2)^2(\nu-4)}$$

We can calibrate the parameters of the $SI\chi^2$ distribution to historical data for asset *A* by computing the ratio of the standard deviation of σ^2 to the mean of σ^2 as a standardized measure of the instability of asset *A*'s variance:

$$\frac{StandardDeviation(\sigma^2)}{Mean(\sigma^2)} = \gamma = \frac{1}{\sqrt{\nu - 4}}$$

Therefore

$$\nu = \frac{1}{\gamma^2} + 4$$

Given these assumptions, it can be shown that the unconditional return distribution for asset *A* will be a *t*-distribution with ν degrees of freedom. When we apply well-known formulas for the variance and kurtosis of a *t*-distribution, it follows that

$$Variance(A) = \frac{\nu}{\nu - 2} Mean(\sigma^2)$$

$$Excess\ Kurtosis(A) = \frac{6}{\nu - 4} = 6\gamma^2$$

We can compute a multiplier for asset A's variance equal to

$$Variance\ Multiplier = \frac{Variance(A)}{Mean(\sigma^2)} = \frac{\nu}{\nu - 2} = \frac{\frac{1}{\gamma^2} + 4}{\frac{1}{\gamma^2} + 2}$$

Taking the square root, we can obtain an approximation of the multiplier for the standard deviation (or volatility) of asset A as follows:

$$Volatility\ Multiplier = \frac{Standard\ Deviation(A)}{Mean(\sigma)}$$

$$\approx \frac{Standard\ Deviation(A)}{\sqrt{Mean(\sigma^2)}} = \sqrt{\frac{\frac{1}{\gamma^2} + 4}{\frac{1}{\gamma^2} + 2}}$$

Exhibit A1 shows the volatility multiplier and excess kurtosis as a function of the degree of standard deviation instability, γ.

This analysis shows that accounting for uncertainty in variance will have two effects on the unconditional returns of an asset's distribution, compared to the normal distribution corresponding to its average variance. In particular, the distribution reflecting uncertain variance will have higher volatility and also fatter tails (positive excess kurtosis). We could approximate these effects by simply increasing the standard deviation of each asset within a multivariate normal distribution to reflect its degree of instability. To account for the excess kurtosis, however, we would have to relax the assumption of normality for asset returns. Nevertheless, adjusting the standard deviation for each asset in order to derive a new normal distribution that accounts for instability is a convenient approximation.

Unfortunately, this approach does not work for capturing instability in the correlations between assets. The distributions that result from a mixture of multivariate normal distributions with different correlations may not be well-represented by any adjustment to the correlation coefficients. For example, consider two assets that are sometimes 0.8 correlated and sometimes −0.8 correlated. The unconditional correlation may be zero; however, a correlation of zero obscures the important fact that the assets sometimes move dramatically in the same direction, and sometimes move dramatically in opposite directions. This feature of returns is important for portfolio construction, because it implies that portfolios consisting of these two assets are likely to experience extreme tail events when the assets move in tandem. For this reason, we prefer to use an empirical approach to capture instability in the entire covariance matrix.

EXHIBIT A1
Predicted Volatility Multiplier and Excess Kurtosis

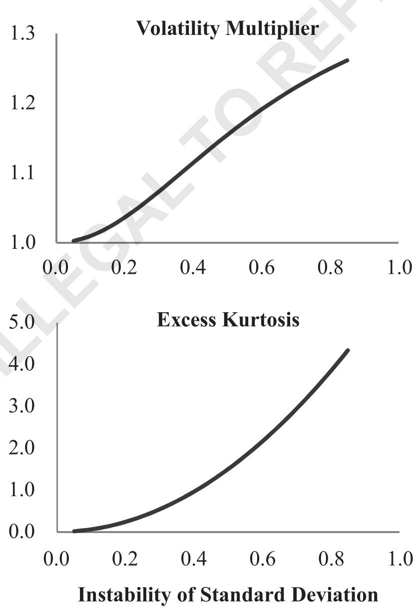

ENDNOTES

The material presented is for informational purposes only. The views expressed in this material are the views of the authors and are subject to change based on market and other conditions and factors. Moreover, they do not necessarily represent the official views of MIT, Windham Capital Management, State Street Global Exchange, or State Street Corporation and its affiliates.

We thank Megan Czasonis for her helpful comments.

[1]If the risk measures pertain to units that are not directly investable, such as factors, then investors face an additional source of error referred to as mapping error. See Cocoma et al. [2015] for more detail about this issue.

[2]See Kinlaw, Kritzman, and Turkington [2014, 2015] for more details about this issue.

[3]For example, our original sample may comprise monthly returns, but our investment horizon may be five years. Therefore, we would estimate the covariance matrix using five-year overlapping returns. We use log returns to calculate covariance matrixes in order to remove the effect of compounding. In particular, we transform each return observation by taking the natural logarithm of one plus the return. The multiperiod compounded returns of a normally distributed asset will be highly skewed due to compounding and therefore not nor-

mally distributed; however, the logarithms of the long-period returns will be normally distributed.

[4]We use overlapping samples to mitigate the distortion that could be caused by choosing a particular start date with independent samples. For example, it could be that a particular period has very high risk and the subsequent period has very low risk. If we were to choose a start date such that we combined half of the first period with half of the subsequent period, we would not capture these extreme episodes of risk.

[5]We remove any strong directional bias from the distribution of errors by subtracting the median error from each individual subsample error.

[6]We should not use the full-sample covariance matrix as our base case, because the full sample embeds the small-sample error of all the subsamples.

[7]Some of the subsample covariance matrixes may not be positive semi-definite. We therefore apply standard corrections to render all covariance matrixes invertible.

[8]In addition to independence, the Central Limit Theorem also assumes finite variances.

[9]See Cremers, Kritzman, and Page [2005] for more details on the full-scale optimization approach.

[10]Although it would be prohibitively challenging to test every possible asset mix in small increments, there are search algorithms that yield a reasonably reliable solution in a few seconds. A particular algorithm based on evolutionary biology initiates several searches simultaneously and iteratively terminates those searches that are sure to fail, thus transferring the search energy to the remaining feasible searches.

[11]We identify these optimal portfolios by searching for the weights in increments of 5% in order to reduce the search cost. We allow the weights to range from 0% to 100%, and we require that they sum to 100%.

[12]The descriptive performance statistics we present are based on the same asset return history used to generate inputs to the optimization process. Therefore, these results do not constitute a true out-of-sample performance test. However, the results are substantially out-of-sample in the sense that our optimization inputs do not directly capture many features of the data. In particular, the optimization inputs capture the distribution of covariance matrixes that occurred in the data, but they do not capture information on the sequence of these covariance matrixes, nor do they capture the particular features of a distribution within any given subsample.

[13]GICS was developed by and is the exclusive property of MSCI Inc. and Standard & Poor's. GICS is a service mark of MSCI and S&P and has been licensed for use by State Street.

[14]To reduce the search cost, we identify these optimal portfolios by searching for the weights in increments of 10%. We allow the weights to range from 0% to 20%, and we require that they sum to 100%.

REFERENCES

Cocoma, P., M. Czasonis, M. Kritzman, and D. Turkington. "Facts about Factors." Working paper, Revere Street, 2015.

Cremers, J.H., M. Kritzman, and S. Page. "Optimal Hedge Fund Allocations: Do Higher Moments Matter?" *The Journal of Portfolio Management*, Vol. 31, No. 3 (2005), pp. 70-81.

Kinlaw, W., M. Kritzman, and D. Turkington. "The Divergence of High- and Low-Frequency Estimation: Causes and Consequences." *The Journal of Portfolio Management,* Vol. 40, No. 5 (2014), pp. 156-168.

——. "The Divergence of High- and Low-Frequency Estimation: Implications for Performance Measurement." *The Journal of Portfolio Management,* Vol. 41, No. 3 (2015), pp. 14-21.