

MIT Sloan School of Management

MIT Sloan School Working Paper 6416-21

History, Shocks and Drifts: A New Approach to Portfolio Formation

Mark Kritzman and David Turkington

This work is licensed under a Creative Commons Attribution-NonCommercial License (US/v4.0)



<http://creativecommons.org/licenses/by-nc/4.0/>

March 10, 2021

HISTORY, SHOCKS AND DRIFTS:
A NEW APPROACH TO PORTFOLIO FORMATION

Mark Kritzman
kritzman@mit.edu

David Turkington
dturkington@statestreet.com

THIS VERSION: MARCH 10, 2021

Abstract

Investors intuitively view future possibilities as a combination of historical outcomes, shocks that occur suddenly, and drifts that unfold gradually over several years. The authors show how to build portfolios based on such a view of the future. Their key innovation is to create a mixed-frequency return sample that properly balances short-term and long-term returns, and to form portfolios by considering all the returns of the sample instead of a statistical summary of them.

HISTORY, SHOCKS, AND DRIFTS: A NEW APPROACH TO PORTFOLIO FORMATION

When we build portfolios, it is both common practice and common sense to form our view of the future from what we observe in the historical record, at least as a point of departure. It might be prudent, however, to contemplate a future with a higher frequency of the shocks that occurred historically or with different shocks that have not previously occurred. And perhaps of greater importance, we should account for forces that emerge gradually and cause asset values to drift in one direction or another over a long horizon, but that are invisible from higher frequency observations.

This consideration of the future poses several challenges to portfolio formation. It challenges us to observe data at more than one frequency. It challenges us to consider important features of the data that we cannot summarize statistically. And it challenges us to think about investor preferences more realistically than implied by the utility functions we typically employ to choose portfolios.

We introduce a portfolio formation process that meets all these challenges. We first describe the future as a set of three types of scenarios: historical scenarios, shock scenarios, and drift scenarios. We next match these scenarios with return samples observed from history or produced by judgment and simulation. Then we combine these return samples into a composite sample that comprises the high frequency observations of the historical and shock

scenarios with the low frequency observations of the drift scenarios. We then find the asset mix that maximizes expected utility based on a realistic description of investor preferences.

Our key innovation is to blend the return samples in a non-arbitrary way that allows us to balance our concern for short-term outcomes with our concern for long-term outcomes. Our innovation is not especially complicated, but it requires greater explication than would be appropriate in an introduction. As a preview, though, our solution depends on a sum of a sum of utilities and utilities of sums.¹ Hopefully, it will be easier to understand than it is to say.

We organize the rest of the paper as follows: In Part 1, we describe three types of scenarios that we blend together to describe the future, and we explain how to match these scenarios with return samples. In Part 2, we describe our methodology for mixing short-term and long-term observations, and we demonstrate why our methodology is the logically correct approach. We also show how to identify the optimal portfolio. In Part 3, we apply our approach to real data and discuss its versatility. We also describe how it relates to mean-variance analysis. We summarize in Part 4.

Part 1: Scenarios

Investors view the future through the lens of history, but they recognize that the future is unlikely to repeat itself exactly. For example, it may include shocks that are unprecedented either in form or frequency. And it could be shaped by forces that emerge so gradually their effects are hidden from observation. The challenge for portfolio construction is to create a

unified return sample that simultaneously captures three types of scenarios: history, shocks, and drifts, keeping in mind that we observe these scenarios at different frequencies.

Historical Scenarios

A historical scenario is simply a replication of a large segment of history. Investors often rely exclusively on history to guide their view of the future, but they typically summarize the historical record statistically rather than account for each return individually. Our approach is to include a large historical sample of individual returns observed at a monthly frequency as a major component of our overall return sample. This is a sound point of departure, especially when one considers that the last half century included the energy crisis of the mid 1970s, the stagflation of the late 1970s and early 80s, the 1987 stock market crash, the dot com bubble, 9/11, the housing bubble, the Global Financial Crisis, the COVID pandemic, several wars, and an attempted coup against American democracy.

We have two choices for constructing the historical component of our return sample. We could construct it from available returns and limit its chronological reach to the earliest observation of the asset class with the shortest history, or we could use maximum likelihood estimation to estimate means, standard deviations, and correlations and simulate returns from these statistics reaching back to the earliest observation of the asset class with the longest history.

Shock Scenarios

We propose that investors augment the historical record with shock scenarios that are of special concern to them. These shocks might take the form of another pandemic, a geopolitical crisis along the lines of the invasion of Kuwait by Iraq, a financial crisis, or an infrastructure disaster similar to the Fukushima meltdown.

We could introduce these shock scenarios by repeating historical return observations that coincide with individual shocks or combinations of them, which would imply that we believe past shocks are likely to reappear at a higher frequency than their historical incidence. Alternatively, we could hypothesize unprecedented shocks and simulate their effects on returns using judgment. We could represent shock scenarios at the same frequency of the monthly scenario or at a higher frequency.

Drift Scenarios

Next, we should consider scenarios in which a phenomenon evolves gradually over a multi-year horizon. Because these phenomena evolve slowly, their effect on financial markets may not be apparent from monthly observations. Climate change is the prototypical example of a drift scenario. We might also consider a scenario in which technology displaces labor, potentially leading to a new social contract to determine how goods and services are distributed to society.

These long-term drift scenarios, to some extent, may already be encapsulated in the historical sample. However, we might not observe their effect on financial markets if we only observe returns at a monthly frequency. For example, assets that move together at a monthly frequency may drift apart over longer periods. And if an asset's returns are serially correlated,

its volatility of longer-term returns will differ from what is implied by the volatility of its monthly returns. We should therefore represent drift scenarios with cumulative long-term historical returns. Moreover, if we believe these long-term phenomena are driven by latent forces that have not yet surfaced in asset returns, we should add drift terms to the historical returns.

Our main point is that we intuitively view future possibilities as a combination of history, shocks, and drifts. We should therefore build a return sample to match this view of the future. But rather than form our portfolio from a statistical summary of this return sample, as would be the case with mean-variance analysis, we should form it directly from the individual returns to capture nuances that summary statistics would fail to detect.

Before we proceed to our discussion of methodology, we should first comment on the distribution of our composite return sample. One might assume that if we combine a variety of samples into a composite sample, especially some that are generated by simulating returns from a normal distribution, we will end up with a normally distributed composite sample, owing to the Central Limit Theorem. This would defeat our purpose to capture nuances in the data that mean-variance analysis would fail to detect. However, we need not worry about the Central Limit Theorem in this context. The Central Limit Theorem states that the sum or average of independent² observations that are not themselves normally distributed will tend toward normality as more and more variables are averaged or added together. However, when we assemble return samples into a composite sample, we do not average or add them. We combine them. To understand this distinction, imagine a sample of 50 annual returns in which the mean of the first 25 observations is -10% and mean of the last 25 observations is +10%. If

we were to add the first return to the 26th return and the second return to the 27th return and proceed in like fashion across all the observations, the resultant distribution would tend toward normality. However, the distribution of the 50-year sample would hardly be normal. It would be bimodal with some of the returns clustered around a mean of -10% and the rest clustered around a mean of +10%. Our point is that by including shock scenarios, for example, we preserve the effect on kurtosis and skewness that we would expect from these scenarios as well as other non-normal influences that might arise from these shocks.

Part 2: Methodology

Our methodology comprises three stages. The first stage is to construct a mixed-frequency return sample that properly blends short-term returns with long-term returns. The second stage is to specify utility, which likely differs depending on whether we are considering short-term or long-term returns. The final stage is to add the utilities across short-term and long-term returns for as many portfolios as necessary to identify the portfolio with the highest expected utility.

Mixed-Frequency Return Sample

Consider a path of returns that we observe at short intervals and at long intervals. These intervals exist along this path in a certain relation to one another because every long interval contains a specific number of short intervals. If, for example, we define our short intervals as months and our long intervals as five years, each long interval contains 60 short intervals. Therefore, for every five-year return in our sample, we must include 60 monthly returns. If we

include fewer than 60 monthly returns for every five-year return, we would assume implicitly that one or more five-year returns was composed from fewer than 60 monthly returns. And if we include more than 60 monthly returns for every five-year return, we would assume implicitly that one or more five-year returns was composed from more than 60 monthly returns. It follows from pure reason that the only sensible balance of short-term and long-term returns is for the number of long-term returns to equal the number of short-term returns divided by the number of short-term periods in the long-term period.

But in case there is doubt, we can invoke a well-known property of log-wealth utility to dispose on any such doubt. Paul Samuelson (1963)³ famously showed that investors with log-wealth utility who prefer a certain mix of risky and safe assets for a short horizon will prefer the same mix for a long horizon. This indifference to horizon applies to all power utility functions as long as returns do not exhibit any serial correlation. But even if returns mean revert or trend, log-wealth investors are still indifferent to horizon. Exhibit 1 illustrates this property of log-wealth utility.

Exhibit 1: Indifference to Horizon

	Initial Wealth	1st Period Distribution	2nd Period Distribution	3rd Period Distribution
				237.04 x .125
			177.78 x .25	133.33 x .125
		133.33 x .50		133.33 x .125
			100.00 x .25	75.00 x .125
	100.00			133.33 x .125
			100.00 x .25	75.00 x .125
		75.00 x .50		75.00 x .125
			56.25 x .25	42.19 x .125
Expected wealth	100.00	104.17	108.51	113.03
Expected utility	4.6052	4.6052	4.6052	4.6052

Exhibit 1 shows the change in expected wealth and utility for a \$100 investment that has an equal chance of gaining 1/3 or losing 1/4 each period. For a log-wealth investor, this investment conveys the same utility, 4.6052 as keeping the \$100 and not investing. Note that the natural log of \$100 is 4.6052, and $[\.5 \times \ln(\$133.33) + \.5 \times \ln(\$75)]$ also equals 4.6052. Therefore, \$100 is said to be the certainty equivalent of this risky investment. But notice what happens after two and three periods. Even though expected wealth increases with each period, expected utility remains the same. This outcome holds for a log-wealth investor because the lower probability of a cumulative loss is exactly offset by the greater amount one

could potentially lose over longer horizons. For our purposes, though, the salient fact is that log-wealth investors derive the same expected utility from a given investment irrespective of the number of periods over which they invest. We rely on this canonical example of utility to validate our assertion that a mixed-frequency return sample must include all the high-frequency returns, but no more, that make up every low-frequency return.

Consider, for example, Exhibit 2. It shows the sum of the utilities of the four one-period returns as well as the utility of the three overlapping two-period returns. Notice that the sum of the utilities of the one-period returns exactly equals the sum of utilities of the two-period cumulative returns. This equivalence must hold as long as we include the utility of every one-period return that goes into every two-period return, given the properties of log-wealth utility.

Exhibit 2: A Balanced Sample Illustrating Log-Wealth Horizon Indifference

Time Periods	Portfolio Returns	Log-Wealth Utility
1	5.0%	0.049
2	10.0%	0.095
1 and 2	15.5%	0.144
2	10.0%	0.095
3	15.0%	0.140
2 and 3	26.5%	0.235
3	15.0%	0.140
4	-20.0%	-0.223
3 and 4	-8.0%	-0.083
Sum of utilities		0.296
Utility of sums		0.296

Now consider Exhibit 3. Again, we show the sum of the utilities of the four one-period returns as well as the utility of the three overlapping two-period returns. However, notice that the sum of the utilities of the one-period returns is less than the sum of the utilities of the two-period cumulative returns. This difference occurs because we use period 2's return in both the cumulative return for periods 1 and 2 and for periods 2 and 3. We use it twice in the cumulative returns, but we only include it once as a one-period return. The same is true for period 3's return. It shows up twice in the cumulative returns but only once as a one-period return. In other words, two of the two-period cumulative returns are missing one-period returns because they were already used in the overlapping periods.

Exhibit 3: An Imbalanced Sample Violating Log-Wealth Horizon Indifference

Time Periods	Portfolio Returns	Log-Wealth Utility
1	5.0%	0.049
2	10.0%	0.095
3	15.0%	0.140
4	-20.0%	-0.223
1 and 2	15.5%	0.144
2 and 3	26.5%	0.235
3 and 4	-8.0%	-0.083
Sum of utilities		0.061
Utility of sums		0.296

Hopefully, this example convinces you that a valid mixed-frequency return sample, within the context of computing utility, must include all the short-term returns that make up long-term returns. This relationship applies whether the long-term returns are overlapping or

non-overlapping.

We can only show this equivalence of the sum of utilities and the utilities of sums for log-wealth utility, owing to its special property. But the correct mix of short-term and long-term returns that leads to this equivalence applies to all mixed-frequency return samples, no matter which utility function we apply to the return sample.

Investor Utility

Although we can apply any utility function to a mixed-frequency return sample, we are better able to appreciate its versatility if we use a utility function that we cannot adequately approximate by just mean and variance. We, therefore, assume a kinked utility function in which investors have greater aversion to losses below a chosen threshold than they do above the threshold. This utility function is described by Equation 1.

$$U_{kinked}(R) = \begin{cases} \frac{(1+R)^{1-\gamma} - 1}{1-\gamma}, & \text{for } r \geq k \\ \frac{(1+R)^{1-\gamma} - 1}{1-\gamma} - \omega(k - R), & \text{for } r < k \end{cases} \quad (1)$$

The term $U_{kinked}(R)$ is expected utility, R is the return of the portfolio, k is the location of the kink, γ determines the curvature of the function above the kink, and ω is the slope of the function below the kink. With this utility function, investor satisfaction drops precipitously when returns fall below the kink. When returns are above the kink, investor satisfaction

conforms to a power utility function. This utility function is designed to express a stronger aversion to losses below the kink than above the kink.

It is also worth noting that we are not bound to a single specification of utility. In fact, at the very least, it is reasonable to locate the kink at a lower threshold for short-term returns than for long-term returns based on the logic that investors are more tolerant of short-term losses that may be transitory than losses that persist over a more extended horizon. Also, we could use entirely different utility functions for short-term and long-term returns. This ability to form portfolios that apply different preferences to different investment horizons speaks to the versatility of our approach.

Portfolio Formation

We apply a technique called full-scale optimization to form portfolios.⁴ With this approach we calculate a portfolio's utility for every period in our sample considering as many asset mixes as necessary to identify the weights that yield the highest expected utility, given any utility function.

We start with any asset mix, but preferably one we believe, based on experience or judgment, is a reasonable expectation of the solution. We then compute the sum of the utilities associated with the composite historical, shock, and drift scenarios and store this value. Next, we substitute another asset mix, compute its utility, and store that value. We proceed in this fashion until we have computed utility for enough asset mixes that we are confident one of them yields the maximum utility or a level of utility that is sufficiently close to the maximum. We then rank the utilities and choose the asset mix with the highest sum. This approach

considers all features of the mixed-frequency return sample, including kurtosis, skewness, and any other peculiarities of the sample. Next, we demonstrate this methodology empirically.

Part 3: Empirical Analysis

We base our empirical analysis on the asset classes shown in Exhibit 4. We use their US dollar denominated total returns covering the period from February 1973 through December 2020.

We chose these asset classes because they span a wide range of risk and because we can observe their returns as far back as 1973.

Exhibit 4: Asset Classes and Indexes

Asset Classes	Indexes
US Stocks	MSCI USA
Non-US Stocks	MSCI Developed World ex USA
US Treasury Bonds	Bloomberg Barclays Treasury
US Corporate Bonds	Bloomberg Barclays Corporate
Real Estate	NAREIT
Cash	Ken French online data risk free rate

We construct the composite mixed-frequency return sample as a combination of one historical scenario, one shock scenario, and one drift scenario.

We use the entire sample of monthly returns to represent the historical scenario. However, rather than using the total returns for each asset class, we subtract the risk free rate that prevailed in each historical period to arrive at excess returns that make sense in the

context of today's interest rates. Moreover, we add drift terms to them to reflect anticipated drifts that may occur over long horizons.

We define the shock scenario as a repeat of the Global Financial Crisis; therefore, we represent it as the 60 monthly returns from January 2008 through December 2012. Although these returns are already included in the historical scenario, we repeat them to reflect our view that a financial crisis is twice as likely looking ahead as its single occurrence in our historical sample suggests.

We consider our drift scenarios to encompass climate change in addition to other tendencies that cause asset classes to covary over multi-year horizons. We represent these drift scenarios as the cumulative returns of the non-overlapping five-year periods ending 1980, 1985, 1990, 1995, 2000, 2005, 2010, 2015 and 2020, as well as the five-year period ending in 2012 given that we repeated the monthly returns during this period as our shock scenario. This maintains the correct proportion of monthly and five-year observations. And as we noted, to capture the anticipated effect of forces that remained latent throughout history, we add the following five-year drifts to each asset class. We convert these prospective drifts to monthly units and apply them to every monthly observation to maintain consistency:

- US Stocks: -10%
- Non-US Stocks: -5%
- US Treasuries: 0%
- US Corporate Bonds: -3%
- Real Estate: -10%
- Cash: 0%

At this point it is worth emphasizing how we reflect the impact of the drift scenarios on return and risk. We add drift terms to the monthly returns of the historical scenarios and shock scenarios to capture a hypothetical view of how climate change will affect asset class returns beyond what has been observed historically. In addition, we compound the monthly returns into five-year cumulative returns to capture the effect of serial correlation on the range of outcomes over five-year intervals both for each asset in isolation and for their co-variation. Because we include a five-year cumulative return for every 60 monthly returns, including the shock scenario returns, our composite mixed-frequency return sample is properly balanced between short-term and long-term returns.

We use the kinked utility function described by Equation 1 to compute the utility of alternative asset mixes. For the monthly returns of the historical and shock scenarios, we set the kink k equal to -10%, ω equal to 2 to set the slope of the utility function below the kink, and γ equal to 2 to determine the curvature of the utility function above the kink. For the five-year cumulative returns of the drift scenario, we set the kink k equal to 0%, ω equal to 5, and γ equal to 0.25.

We apply full-scale optimization to compute the utility of approximately 50,000 alternative asset mixes at increments of 5%, and we select the portfolio with the highest expected utility. Next, we use a local search algorithm to further maximize expected utility from this starting point and arrive at a final set of weights for the full-scale optimal allocation. We then compare this portfolio to two mean-variance portfolios with the same expected return as the full-scale portfolio: one with the covariance matrix estimated from monthly returns and the other with the covariance matrix estimated from five-year returns.

Exhibit 5 presents the results of this analysis.

Exhibit 5: Full-Scale versus Mean-Variance Optimal Portfolios

	Full-Scale with Mixed Sample	Mean-Variance (Monthly)	Mean-Variance (5 Year)
Portfolio Weights			
US Stocks	23.7%	30.9%	13.1%
Non-US Stocks	21.8%	0.0%	0.0%
US Treasury Bonds	49.6%	57.6%	0.0%
US Corporate Bonds	0.0%	0.0%	69.0%
Real Estate	4.9%	11.5%	17.8%
Cash	0.0%	0.0%	0.0%
Monthly Performance			
Average	0.3%	0.3%	0.3%
Volatility	2.3%	2.1%	2.4%
Skewness	-0.47	-0.37	-0.31
Frequency of loss	40.5%	42.0%	42.2%
Frequency of >10% loss	0.3%	0.0%	0.5%
Worst month	-10.3%	-9.0%	-11.9%
10th percentile	-2.6%	-2.2%	-2.3%
90th percentile	2.9%	2.5%	2.6%
Best month	7.9%	7.2%	11.1%
5 Year Performance			
Average	16.9%	16.9%	16.9%
Volatility	31.4%	20.5%	14.9%
Skewness	1.47	0.70	-0.53
Frequency of loss	0.0%	10.0%	10.0%
Worst 5 years	0.0%	-3.4%	-3.2%
10th percentile	5.3%	11.2%	14.0%
90th percentile	66.8%	51.3%	43.1%
Best 5 years	99.1%	67.4%	49.8%

Exhibit 5 reveals a variety of interesting insights. The full-scale optimal portfolio allocates roughly half to stocks and real estate and half to bonds. By contrast, the mean-variance portfolios, which target the same expected return, tilt more heavily toward bonds,

likely due to their high risk-adjusted returns as a result of secular declines in interest rates during the historical sample. Notably, both mean-variance portfolios have no allocation to foreign stocks because mean-variance analysis is blind to some of the nuanced diversification benefits foreign stocks provide, for reasons we will discuss in more detail shortly.

The full-scale portfolio seeks to balance monthly and five-year performance. Its monthly performance is broadly similar to that of both mean-variance portfolios, but its five-year performance is substantially better. Though its 10th percentile return falls below that of the mean-variance portfolios, the full-scale portfolio configures its allocation for much larger upside potential while simultaneously reducing exposure to the largest losses. At first glance, these benefits for the full-scale portfolio appear to come at the expense of higher volatility. But we must recognize that this higher volatility results from upside potential and not from increased exposure to loss; therefore, it is a flawed measure of suitability. Indeed, the long-horizon distributions reflect compounding which introduces positive skewness. Given a kinked utility function, mean-variance analysis is ill-equipped to cope with this feature of long-horizon returns.

The full-scale optimal portfolio is sensitive to nuances in correlation dynamics. Exhibit 6 shows the correlations for the monthly and five-year return samples. We see, for example, that Treasury bonds and corporate bonds are close substitutes in terms of long-run performance, but Treasury bonds offer better diversification for risky assets in the short run. This explains why the full-scale portfolio and the mean-variance monthly portfolio prefer Treasury bonds, while the mean-variance five-year portfolio prefers corporate bonds. Similarly, the five-year mean-variance portfolio's preference for real estate may be explained by its greater

diversification over long periods. However, in addition to balancing the tension between short- and long-term diversification, full-scale optimization also considers asymmetric correlations. Recall that we over-weighted the 2008-2009 financial crisis in our sample to reflect a shock in which real estate and stocks suffer in tandem. The full-scale portfolio seeks to avoid this negative shock that the five-year mean-variance portfolio fails to perceive. In view of these dynamics, the full-scale portfolio relies more on international diversification.

Exhibit 6: Short-Horizon and Long-Horizon Correlations

Monthly	US Stocks	Non-US Stocks	US Treasury Bonds	US Corporate Bonds
Non-US Stocks	0.71			
US Treasury Bonds	0.02	-0.03		
US Corporate Bonds	0.36	0.30	0.77	
Real Estate	0.63	0.55	0.08	0.41
5 Year	US Stocks	Non-US Stocks	US Treasury Bonds	US Corporate Bonds
Non-US Stocks	0.43			
US Treasury Bonds	-0.04	0.26		
US Corporate Bonds	0.04	0.29	0.94	
Real Estate	-0.22	-0.25	-0.31	-0.19

Intuition and Versatility

In our judgment, the application of full-scale optimization to mixed-frequency return samples has two overarching benefits. First, it enables investors to characterize the future as a blend of history, shocks, and drifts, which matches how many, if not most, investors consider the future. It therefore gives structure to our intuition.

Second, it is extraordinarily versatile. We have already shown in our empirical analysis that we can observe our view of the future at two different frequencies and that we can apply different preferences to these frequencies. We should note that we are not limited to only two frequencies when we form mixed-frequency return samples. For example, we could observe the historical scenarios monthly, the shock scenarios daily, and the drift scenarios over multiple-year horizons.

Our approach also allows us to include variations of our scenarios. We could, for example, repeat our drift scenarios with different drift terms. We only need to ensure that we preserve the correct balance between short-term and long-term returns. And we could modify the size of our scenario return samples to reflect our view of their relative likelihoods. These observations are but a few of the ways we can adapt our mixed-frequency return sample to conform to our view of the future.

Mean-Variance Analysis with Mixed Frequencies

It is also possible to adapt mean-variance analysis to consider different return frequencies. The objective of mean-variance analysis is to maximize expected utility defined as expected return minus risk aversion multiplied by portfolio variance. This process produces an efficient frontier comprised of portfolios that offer the highest expected returns for given levels of risk.

We could expand this objective function to include two estimates of portfolio variance: one that we estimate from the short-term returns in the mixed-frequency return sample and one that we estimate from the long-term returns. And we could apply different specifications

of risk aversion to each measure of portfolio variance in accordance with our preferences regarding short-term and long-term outcomes (see, for example, Kinlaw, Kritzman, and Turkington, 2014). However, this approach compares unfavorably to our proposed methodology, because it relies on a statistical summary of the mixed-frequency return sample, and because it is limited in how it considers expected utility. It therefore fails to consider nuances in the return distribution as well as nuances in investor preferences.

Summary

Investors care about short-term outcomes because significant short-term losses might disrupt ongoing spending plans or derail longer-term goals. Investors also care about outcomes that unfold over longer horizons because adverse cumulative outcomes affect future wellbeing. To address these dual concerns, we introduce a methodology for constructing a mixed-frequency return sample that enables investors to consider the sum of utilities associated with returns observed at a high frequency, together with the utilities of cumulative returns (the utilities of sums) observed at a low frequency. This innovation allows investors to contemplate the future as a composite of three types of scenarios: historical scenarios in which the past repeats itself; shock scenarios such as a pandemic, a geopolitical crisis, a financial crisis, or an infrastructure disaster; and drift scenarios like climate change or technological displacement.

We also show how to form portfolios from this mixed-frequency return sample using a technique called full-scale optimization. This approach to portfolio construction compares favorably to conventional mean-variance analysis in three important ways. It allows investors

to consider returns at more than one frequency. It accounts for every feature of the data, instead of relying on a statistical abstraction of a return sample. And it allows investors to describe their preferences more realistically than preferences approximated only by mean and variance.

Our key innovation is the way in which we combine short-term and long-term returns to form a mixed-frequency return sample. We justify the validity of our approach based on pure reason as well as a special property of log-wealth utility.

We illustrate our methodology by applying it to a realistic sample of returns that combines a historical sample, shock scenarios, and drift scenarios to demonstrate its practicality and versatility. And we provide in-sample metrics to compare it to mean-variance analysis. We believe this new approach to portfolio formation will better empower investors to match complex preferences with complex realities.

Notes

We are grateful to Jarrod Wilcox for helpful comments.

This material is for informational purposes only. The views expressed in this material are the views of the authors, are provided “as-is” at the time of first publication, are not intended for distribution to any person or entity in any jurisdiction where such distribution or use would be contrary to applicable law and are not an offer or solicitation to buy or sell securities or any product. The views expressed do not necessarily represent the views of Windham Capital Management, State Street Global Markets®, or State Street Corporation® and its affiliates.

References

W. Kinlaw, M. Kritzman, and D. Turkington. 2014. “The Divergence of High- and Low-Frequency Estimation: Causes and Consequences,” *Journal of Portfolio Management*, Vol. 40, No. 5 (40th Anniversary Issue).

P. A. Samuelson. 1963. “Risk and Uncertainty: A Fallacy of Large Numbers,” *Scientia*, 98 (April/May).

P.A. Samuelson. 1979. “Why We Should Not Make Mean Log of Wealth Big Though Years to Act Are Long,” *Journal of Banking and Finance*, 3. Pages 305-307.

¹ To be precise, we sum utilities of the high-frequency returns and add them to the utilities of the low frequency returns, which we can define as compounded high frequency returns, in which case they are products, or as the sum of continuous returns, in which case we must reconvert them to discrete returns.

² Some formulations of the Central Limit Theorem require observations to be identically distributed in addition to being independent, but there are also versions that do not require this.

³ Samuelson’s insight about the effect of horizon on utility provoked considerable debate. As part of the debate, Samuelson wrote an article to explain his position using words of only one syllable except for the last word which was syllable, at which point he noted that all the previous words had only one syllable. See Samuelson (1979).

⁴ Paul Samuelson coined the term “full-scale optimization” in a letter to Mark Kritzman which was in response to his request of Kritzman to test mean-variance analysis against direct utility maximization.