



Portfolio Choice with Path-Dependent Scenarios

Mark Kritzman, Ding Li, Grace (TianTian) Qiu & David Turkington

To cite this article: Mark Kritzman, Ding Li, Grace (TianTian) Qiu & David Turkington (2020): Portfolio Choice with Path-Dependent Scenarios, Financial Analysts Journal

To link to this article: <https://doi.org/10.1080/0015198X.2020.1841539>



Published online: 16 Dec 2020.



Submit your article to this journal [↗](#)



View related articles [↗](#)



View Crossmark data [↗](#)

Portfolio Choice with Path-Dependent Scenarios

Mark Kritzman, CFA , Ding Li , Grace (TianTian) Qiu ,
and David Turkington, CFA 

Mark Kritzman, CFA, is a founding partner and CEO of Windham Capital Management, Boston, Massachusetts. Ding Li is senior vice president at GIC Private Limited, Singapore. Grace (TianTian) Qiu is senior vice president at GIC Private Limited, Singapore. David Turkington, CFA, is senior managing director at State Street Associates, Cambridge, Massachusetts.

Sophisticated investors rely on scenario analysis to select portfolios. We propose a new approach to scenario analysis that enables investors to consider sequential outcomes. We define scenarios not as average values but as paths for the economic variables. And we measure the likelihood of these paths on the basis of the statistical similarity of the paths to historical sequences. We also use a novel forecasting technique called “partial sample regression” to map economic outcomes onto asset class returns. This process allows investors to evaluate portfolios on the basis of the likelihood that the scenario will produce a certain pattern of returns over a specified investment horizon.

Many investors select portfolios by using the technique of scenario analysis, which involves the following five steps. First, define each scenario as a set of values for relevant economic variables. Second, assign probabilities to prospective scenarios. Third, translate economic scenarios into expected asset class returns. Fourth, specify alternative portfolio choices. And fifth, compute portfolio performance metrics and select a portfolio.

The two most challenging aspects of scenario analysis are assigning probabilities to prospective scenarios and translating those scenarios into asset class returns. Czaronis, Kritzman, Pamir, and Turkington (2020) showed how to use the Mahalanobis distance, a statistical measure of similarity, to estimate the relative probabilities of prospective economic scenarios.¹ We extend their innovation in two important ways. First, we propose that investors define scenarios as sequences of values for the economic variables instead of single-period average values. Defining scenarios as paths rather than single-period averages has several advantages, which we discuss later. Second, we apply a novel forecasting technique called “partial sample regression” to translate economic values into asset class returns.² These techniques allow investors to extrapolate from data in a way that is not arbitrary or subject to bias, and they provide a language that enables investors to consider their subjective views within the context of an objective baseline. Importantly, the methodology we propose imposes an internal consistency across probability estimation and return forecasting by virtue of its application of the Mahalanobis distance to both challenges.

Note that our approach is distinct from Monte Carlo simulation, which is often used to model multiperiod outcomes. Simulation methods create paths of returns by stringing together random draws from prescribed probability distributions. Thus, the outcomes and probabilities they generate merely reflect the distributions they are given. Analysts still face the task of estimating probabilities from data or deriving them theoretically. Our approach is more direct, and probably more intuitive for most investors, because it starts with explicit scenario definitions and estimates their probabilities from the data in one step.

Disclosure: The authors report no conflicts of interest.

PL Credits: 2.0

Moreover, simulations produce a large—possibly infinite—number of paths. Our approach is aimed at investors who wish to analyze a more parsimonious set of scenarios. Finally, our use of partial sample regression provides a simple and intuitive link between economic variables and asset returns that is typically obscured in a simulation analysis.

We proceed as follows. We first discuss the merits of defining an economic scenario as a multistage outcome rather than an average outcome. We next describe how we define multistage scenarios and estimate their likelihood of occurrence. Then, we describe how we map economic scenarios onto asset class performance by using partial sample regression. We next illustrate our approach with a case study in which we generated a rich set of results. The case study captures not only the probable average performance of the alternative portfolios but also the pattern of their returns and the dispersion of their performance across scenarios and through time. We conclude with a summary.

Scenarios as Paths

Investors typically define economic scenarios with a set of single numbers that represent the average values or end-of-period values of relevant variables for the scenario horizon. We argue that this approach is unnecessarily vague and potentially harmful. We propose that investors, instead, define scenarios as paths in which each variable is assigned a sequence of values. Defining scenarios as paths is significantly advantageous both for measuring the relative likelihood of prospective scenarios and for mapping them onto asset class returns.

As we soon discuss, we assess a scenario's likelihood as a function of its statistical similarity to recent economic conditions. By characterizing a scenario as a sequence of values as opposed to a single average value, we gain a large set of observations to use in measuring statistical similarity. For example, the average outcomes for two sets of variables may represent alternative scenarios that occurred commonly throughout history, but when the two scenarios are specified as multiperiod paths, one scenario may be found to have been without precedent whereas the other was a usual occurrence.

Also, we use the novel forecasting technique of partial sample regression to convert economic scenarios into asset class returns. The essence of this methodology, which we later describe in detail,

is to estimate asset class returns from a subset of relevant observations in which relevance is defined, in part, by statistical similarity. Without getting ahead of ourselves, an example may be illustrative and helpful.

Suppose we have annual observations for real GDP growth from 1929 through 2019 and we want to find the three-year periods most statistically similar to the global financial crisis (GFC) of 2006–08. We construct one version of the data that consists of three variables (Year 1, Year 2, and Year 3) for the paths. We construct another version of the data that consists of average three-year growth rates. **Figure 1** shows the most similar three-year periods based on the year-by-year paths.

Figure 2 shows the most similar three-year periods based on average three-year growth rates.

Figure 1 and Figure 2 show that when we define scenarios as paths, the periods we identify as most like the GFC are different from those when we define scenarios as multiyear averages. Moreover, the periods we identified using paths bear a much stronger resemblance to the actual pattern of economic growth that occurred during the GFC.

The key motivation for defining scenarios as paths is that doing so enables us to measure statistical similarity with reliability, which improves our ability to assign probabilities to scenarios and to forecast asset

Figure 1. Most Similar Three-Year Periods Based on Paths

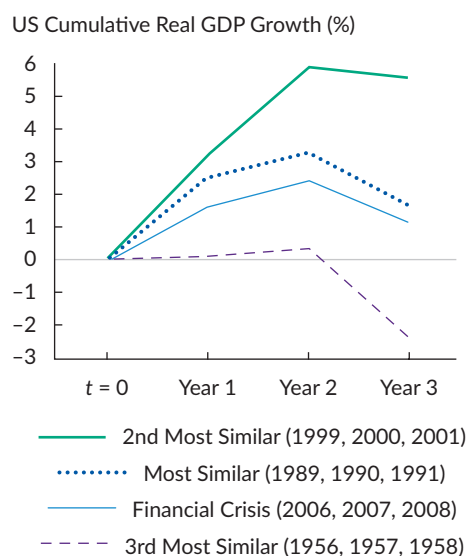
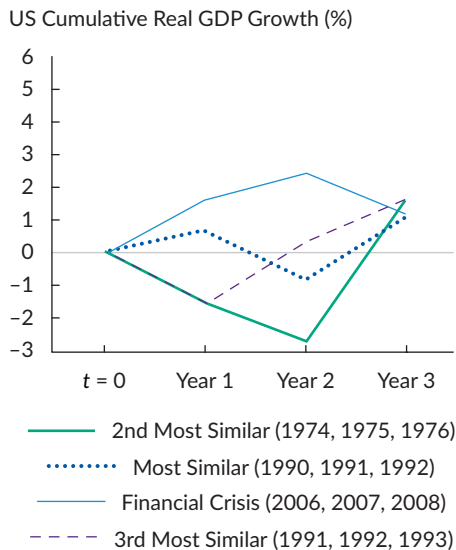


Figure 2. Most Similar Three-Year Periods based on Average Values



class returns. The quantitative measurement of paths also aligns with the logic and intuition of qualitative forecasting, in which, for example, analysts use economic narratives to relate various historical events to current circumstances.

We believe that this innovation significantly enhances the methodology proposed by Czaronis et al. (2020) for assigning probabilities to scenarios, and it enhances the application of partial sample regression for forecasting asset class returns.

Defining scenarios as paths confers additional benefits. It enables investors to make informed tactical shifts ranging from simple rebalancing to active tilts away from a steady-state investment posture. In addition, knowledge of the paths of alternative scenarios allows investors to consider a rich range of metrics by which to evaluate alternative portfolios.

One might suspect that defining scenarios as paths would require a complicated model to relate sequential observations of a variable to each other and to the sequential observations of the other variables. The Mahalanobis distance excels at precisely this task, however, so no additional complexity is needed to model sequences as opposed to averages of variables.

We next describe how we define scenarios and assign probabilities to them.

Scenarios and Their Likelihood of Occurrence

For the reasons stated previously, we propose that investors define scenarios as a set of multiple values for chosen economic variables representing the early, middle, and late stages of a pattern. The pattern might cover multiple months, quarters, or years. Or investors might choose more or less granular descriptions of a pattern. However the scenarios are defined, the multistage values for all relevant variables are to be collected and arranged into a single vector that describes the sequence of outcomes associated with that scenario.

To estimate the relative likelihood of the prospective scenarios, we compute the statistical similarity of the prospective scenarios to the most recent economic experience by using the Mahalanobis distance. Effectively, we are asking, Given the recent economic experience, how likely would it be for one scenario to prevail going forward versus an alternative scenario?

Of course, we could choose a different period from the most recent economic experience to anchor our measure of the Mahalanobis distance. For example, if we believe that the recent experience is highly unusual and that conditions will revert to a more normal experience, we might anchor the Mahalanobis distance to values representing a more typical pattern. Before we proceed, a useful step at this point is to review the Mahalanobis distance.

The *Mahalanobis distance* was introduced in 1927 and modified in 1936. The aim was to analyze resemblances in human skulls among castes in India.³ The measure is powerful and convenient because it characterizes in a single number the distance between two multivariate observations. In doing so, it accounts for the expected variation of each underlying variable from its average as well as the expected covariation of each pair of underlying variables from their respective averages. Thus, a large Mahalanobis distance may result from greater dispersion in the values of one underlying variable from one observation to the next or from pairs of deviations that are not particularly large but that depart from the typical pattern of covariation for the variables.

The usefulness of the Mahalanobis distance is evidenced by its application to a diverse set of challenges, including diagnosing liver disease (Su and Li 2002), sleep apnea (Wang, Su, Chen, and Chen 2011), and breast cancer (Nasief, Rosado-Mendez, Zagzebski, and Hall 2019) and detecting anomalies

in self-driving vehicles (Lin, Khalastchi, and Kaminka 2010). Within the field of investing, it has been applied to measure financial turbulence (Chow, Jacquier, Kritzman, and Lowry 1999), estimate the likelihood of single-period economic scenarios (Czasonis et al. 2020), improve the forecast reliability of linear regression analysis (Czasonis, Kritzman, and Turkington 2020a), and forecast the correlation between stocks and bonds (Czasonis, Kritzman, and Turkington 2020b).

In our application of the Mahalanobis distance, d is computed as shown in Equation 1:⁴

$$d = (\mathbf{x} - \boldsymbol{\gamma})' \boldsymbol{\Omega}^{-1} (\mathbf{x} - \boldsymbol{\gamma}), \quad (1)$$

where d is the Mahalanobis distance, \mathbf{x} is a vector comprising the multistage values of a set of economic variables used to characterize a future scenario, $\boldsymbol{\gamma}$ reflects the recent multistage values of the economic variables, $\boldsymbol{\Omega}$ is the historical covariance matrix of changes in values for those variables, and the prime indicates a vector transpose. We express all vectors as column vectors. When we construct the covariance matrix, we must take care to properly capture the lagged cross-relationships of the variables. If, for example, we divide our paths into three stages, then for each variable, we must line up periods 1 through $n - 2$ with periods 2 through $n - 1$ and periods 3 through n .

As stated previously, we should have more confidence in the Mahalanobis distances between paths than the Mahalanobis distances between the average values of the paths because paths impose an additional condition and thereby contain more information.

The Mahalanobis distance is closely related to a scenario's probability of occurrence. Specifically, an observation with a high Mahalanobis distance will tend to occur less frequently than one with a low Mahalanobis distance. If we assume that the economic variables follow a multivariate normal distribution, we can measure the relative likelihood of scenarios precisely. The likelihood of an observation decays as the Mahalanobis distance increases. It decays according to an exponential function, which gives rise to the normal distribution. We measure the likelihood that we would observe a given scenario as shown in Equation 2:⁵

$$\text{Likelihood} \propto e^{-d/2}. \quad (2)$$

The likelihoods we compute are in comparable statistical units across scenarios; they will not sum to 1, however, because we have specified only a subset of all possible outcomes. Therefore, we rescale the likelihoods to sum to 1 so that we may interpret them as probabilities.

The next step is to map these economic scenarios onto expected returns for each of the asset classes we wish to consider.

Scenarios and Asset Class Returns

We apply the novel forecasting technique of partial sample regression (Czasonis et al. 2020a) to convert our projected economic scenarios into estimates of expected returns for the asset classes we wish to consider. We then apply this technique in a case study.

Partial sample regression relies on a convenient but obscure mathematical equivalence. The prediction generated by a linear regression model may be written equivalently as a function of the weighted average of the past values of the dependent variable in which the weights are the relevance of the past observations for the independent variables. Relevance is equal to the sum of the statistical similarity of the past observations to the current values for the independent variables and the informativeness of the past observations. Both quantities are measured as Mahalanobis distances.

Equation 3 defines the multivariate similarity between \mathbf{x}_i and \mathbf{x}_t , which is the opposite (negative) of the Mahalanobis distance between them:

$$\text{Similarity}(\mathbf{x}_i, \mathbf{x}_t) = -(\mathbf{x}_i - \mathbf{x}_t)' \boldsymbol{\Omega}^{-1} (\mathbf{x}_i - \mathbf{x}_t), \quad (3)$$

where \mathbf{x}_i is a vector of the prior values of the independent variables, \mathbf{x}_t is a vector of the current values of the independent variables, the prime symbol indicates a matrix transpose, and $\boldsymbol{\Omega}^{-1}$ is the inverse covariance matrix of \mathbf{X} , where \mathbf{X} comprises all the vectors of the independent variables. This measure takes into account not only how independently similar the components of the \mathbf{x}_i 's are to those of the \mathbf{x}_t 's, but also the similarity of the co-occurrence of the \mathbf{x}_i 's to the co-occurrence of the \mathbf{x}_t 's. All else being equal, prior observations for the independent variables that are more like the current observations are more relevant.

However, relevance has a second component. Observations that are more distant from their historical averages are more unusual and, therefore, more likely to be driven by events. These event-driven observations are potentially more informative than the averages.⁶ Equation 4 defines the informativeness of a prior observation, \mathbf{x}_i , as its multivariate distance from its average value, $\bar{\mathbf{x}}$:

$$\text{Informativeness } (\mathbf{x}_i) = (\mathbf{x}_i - \bar{\mathbf{x}})' \mathbf{\Omega}^{-1} (\mathbf{x}_i - \bar{\mathbf{x}}). \quad (4)$$

The relevance of any observation \mathbf{x}_i is equal to the sum of its multivariate similarity and its informativeness:

$$\begin{aligned} \text{Relevance } (\mathbf{x}_i) \\ = \text{Similarity } (\mathbf{x}_i, \mathbf{x}_t) + \text{Informativeness } (\mathbf{x}_i). \end{aligned} \quad (5)$$

In summary, among prior periods that are like the current period, those that are different from the historical average are more relevant than those that are not.

Because linear regression is equivalent to a relevance-weighted average of the past values of the dependent variable, we generate our forecasts from Equations 6 and 7, which apply the same weights but only to a subset of the n most relevant observations.⁷

$$\hat{y}_t = \bar{y} + \frac{1}{2n} \sum_{i=1}^n [\text{Similarity } (\mathbf{x}_i, \mathbf{x}_t) + \text{Informativeness } (\mathbf{x}_i)] (y_i - \bar{y}) \quad (6)$$

and

$$\bar{y} = \frac{1}{n} \sum_{i=1}^n y_i. \quad (7)$$

As noted, we improve the reliability of our forecasts by specifying the alternative scenarios as paths rather than as single-period averages because this refinement gives us more information upon which to base our assessment of the relevance of historical observations. The earlier illustration showed how paths enhance our ability to assess statistical similarity. The same principle applies to informativeness. We should expect to forecast asset class returns more reliably by specifying scenarios as paths rather than averages.

Most people are inclined to extrapolate predictions from the most similar historical events as opposed to the most different ones. By censoring the least

relevant observations, partial sample regression aligns with this commonly accepted wisdom. Even if we set aside the prediction accuracy that is likely to be gained through this approach, it provides a significant benefit by identifying the historical periods with the greatest influence on a specific prediction that comes from a linear regression model. This perspective is essential if one wants to combine quantitative and qualitative assessments of the historical data.

The case study we next present illustrates our approach for assigning probabilities to multistage scenarios and for converting these scenarios into estimates of return sequences for the relevant asset classes.

Case Study

To illustrate our methodology, we consider six prospective economic scenarios, which we define as three-year patterns for economic growth and inflation. These scenarios pertain to the United States, and the variables are expressed in US dollars.⁸

We acknowledge that the choice of scenarios is not one-size-fits-all, but we believe that the following principles should apply.

- The scenarios should span a comprehensive range of economic outcomes.
- The scenarios should consider mean-variance analysis as a comparison. Mean-variance analysis implicitly accounts for all potential scenarios across a continuous distribution defined by an expected return and standard deviation. Scenario analysis, by construction, considers only a finite number of scenarios, but these scenarios should be spread relatively evenly across an imaginary continuous distribution. Specifying scenarios that reside on only one side of the imaginary distribution would not be helpful and could lead to biased decision making.
- The scenarios should be relatively distinct from one another. If they are more redundant than distinct, the probabilities derived from their respective Mahalanobis distances will be understated and unstable. Moreover, the probabilities assigned to the other scenarios may be understated. This principle is related to the first principle.
- The chosen economic variables representing the scenarios should span the key fundamental

drivers of future market behavior. The scenarios are representable by the economic variables chosen to define them. If the key features or narratives of a scenario are not captured by the variables, the inclusion of the scenario will not provide useful information. Investors may choose from a large universe and many combinations of economic variables. The choice of scenario variables should, however, be holistic, parsimonious, and (ideally) orthogonal. This principle is related to the first two principles.

- The scenarios should be conceptually and empirically plausible. The scenarios used to guide asset allocation should capture the plausible set of future outcomes. If we define a scenario whose combination of values is statistically contrary to historical precedent and general market intuition, its likelihood of occurrence will be close to zero. Although such extreme scenarios are useful for stress-testing purposes, they are less relevant for asset allocation, which should be based on a plausible range of outcomes.

We defined our case study scenarios with two macroeconomic variables, real GDP growth and inflation. Theoretically and empirically, these two variables are the key drivers of future market behavior. The choice of just two variables is parsimonious, and these variables are orthogonal, in the sense that they capture two distinct macroeconomic dimensions.

We constructed six distinct and plausible scenarios by defining the three-year paths of these two macro variables. The baseline estimates were taken from

the current Bloomberg consensus forecast for US GDP and inflation. The remaining scenarios were constructed by shocking GDP and inflation in predefined directions on the basis of economic narratives for the scenarios.

Two aspects are useful to keep in mind. First, the six scenarios were constructed following the beginning of the Covid-19 pandemic, but they are by no means unique to current conditions. They could just as well represent alternative recovery paths following any economic or financial shock. Second, these scenarios are meant to be illustrative. Investors may wish to construct their own scenarios based on their own views or the opinions of experts. Our focus is not to propose specific scenarios but to propose a comprehensive framework for conducting path-dependent scenario analysis.

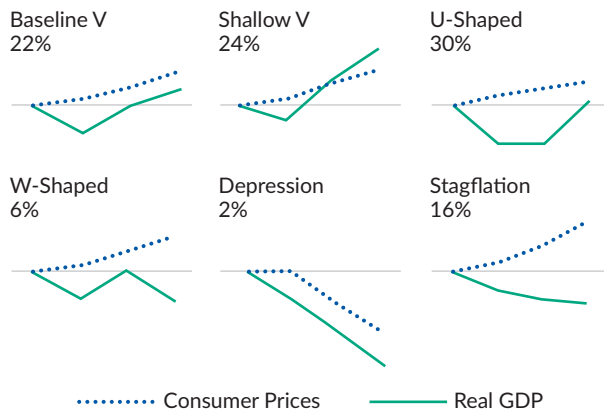
Table 1 shows the annual outcomes in the six scenarios we chose to consider, and **Figure 3** graphically shows the corresponding paths by plotting the cumulative values of the variables.

The column labeled “Current” in Table 1 shows the three-year paths ending in 2019 for growth and inflation. These values serve as the anchor for computing each scenario’s Mahalanobis distance. The columns to the right show the three-year paths of the prospective scenarios. The bottom row gives the relative probabilities of the scenarios.

These probabilities were computed as follows. We computed the time series of the yearly percentage change in real GDP per capita and the yearly

Table 1. Economic Scenarios

Scenario	Current	Baseline V	Shallow V	U-Shaped	W-Shaped	Depression	Stagflation
<i>Growth</i>							
Year 1	2.2%	−3.5%	−1.9%	−3.5%	−3.5%	−5.1%	−5.1%
Year 2	2.0	3.8	5.4	0.0	3.8	−5.9	−2.7
Year 3	1.8	2.3	3.9	3.9	−4.2	−7.4	−0.9
<i>Inflation</i>							
Year 1	2.1%	1.0%	1.0%	1.0%	1.0%	−0.3%	2.3%
Year 2	1.9	1.7	1.7	0.4	1.7	−5.9	4.2
Year 3	2.3	2.0	2.0	0.7	2.0	−5.6	5.8
Probability (as of 2019)		21.5%	24.0%	30.1%	5.9%	2.4%	16.1%

Figure 3. Economic Scenarios

percentage change in the US Consumer Price Index. We used the JST dataset from 1927 to 2015 and data from the Federal Reserve Economic Data (FRED) service of the Federal Reserve Bank of St. Louis from 2016 to 2019.

We formulated time series to represent a three-year path of growth and inflation. For each year from 1929 to 2019, we included growth and inflation from two years previous, growth and inflation from one year previous, and growth and inflation from the current year.

We computed the covariance matrix from the changes in the values of the economic variables from one three-year period to the next three-year period (in other words, the differences in the paths of the variables).

Using the three-year paths for these variables ending in 2019 as the anchor, we applied these inputs (three yearly values for growth and inflation for each scenario) in Equation 1 to compute the scenarios' Mahalanobis distances. We used Equation 2 to convert the Mahalanobis distances into probabilities and rescaled them to sum to 1. Next, we converted these scenarios into asset class returns. We considered three asset classes—US stocks, US bonds, and cash (T-bills)—and we proceeded as follows.

We obtained the time series of yearly total returns for stocks, bonds, and cash from the JST dataset for 1927–2015 and from the S&P 500, Bloomberg US government bond, and JP Morgan three-month cash indexes for 2016–2019. Next, we subtracted annual inflation from each return to arrive at historical real returns, and we subtracted the real returns of

T-bills from those of stocks and T-bonds to convert the stock and T-bond returns to excess returns. We computed the year-over-year change in the cash returns, which reflects a path of interest rate changes that we projected from today's current interest rate levels. For each year from 1929 to 2019, we formulated these three-year paths for each of the three asset classes in a vector (as we did for the economic variables) with nine elements.

We applied partial sample regression (with a subsample of the most relevant 25% of observations) to nine dependent variables and obtained the yearly changes in real cash returns for three years, stock return premiums above cash for three years, and bond return premiums above cash for three years. We cumulated the changes in cash returns to get the total cash returns for each year, and we added them to the premiums of stocks and bonds to get the total real returns for stocks and bonds.

Table 2 shows the expected real returns of stocks, bonds, and cash associated with each of the economic scenarios.

In terms of economic intuition, the asset class returns shown in Table 2 seem remarkably consistent with the scenario descriptions, which gave us confidence in our choice of scenarios and in our process for mapping the scenarios onto asset class returns. We attributed this outcome to our use of partial sample regression to estimate the asset class returns.

Next, we specified the portfolios we wished to consider, which are shown in **Table 3**.

By applying these portfolio weights to the asset class returns, we derived the return paths of the portfolios for each of the economic scenarios, as shown in **Table 4**.

The next question was how we should evaluate these alternative portfolios. In conventional scenario analysis, where we would have just a single average return for each scenario, we would compute the weighted average return across the scenarios and select the portfolio with the highest return. Or we could specify a utility function and select the portfolio with the highest expected utility. Because we specified the scenarios as paths, however, we have a richer set of data to evaluate the alternative portfolios. **Table 5** presents a variety of metrics by which to evaluate the portfolios. All the information is presented for each scenario, for a probability-weighted average across the scenarios, and for a worst-case scenario.

Table 2. Expected Asset Class Real Returns

Asset Class/Year	Baseline V	Shallow V	U-Shaped	W-Shaped	Depression	Stagflation
<i>Stocks</i>						
Year 1	16.3%	19.7%	-3.1%	19.5%	-22.2%	1.7%
Year 2	5.8	7.0	13.4	-7.2	-22.0	-1.9
Year 3	-0.5	1.1	7.2	-3.2	8.7	-6.0
<i>Bonds</i>						
Year 1	4.8%	4.9%	6.4%	-1.1%	3.9%	-0.8%
Year 2	-1.4	-0.7	2.5	-0.2	13.8	-10.6
Year 3	4.9	2.8	5.1	10.0	22.4	-10.0
<i>Cash</i>						
Year 1	1.4%	1.4%	1.3%	1.9%	3.2%	0.8%
Year 2	0.5	0.5	1.5	1.4	8.8	-0.5
Year 3	0.1	0.1	1.1	0.7	8.1	-2.1

The first panel of Table 5 shows the annualized cumulative return for each of the portfolios. This information is the same as we would have generated had we used single-period average values to define the scenarios. The remaining panels provide information that would not be known had we not defined the scenarios as paths. The second panel shows the maximum drawdown if we observed the portfolios only annually. The third panel shows within-horizon loss, which measures how far the portfolios might decline below their initial values during the three-year horizon. Within-horizon loss differs from maximum drawdown in that it does consider losses that occurred from a higher value than the portfolios' initial values. The final two panels show the worst annual loss and the number of annual losses for each portfolio during the three-year horizon.

If we cared only about cumulative return, we would select the aggressive portfolio. We might believe

that this metric reflects risk because the probability-weighted cumulative returns consider a wide range of both positive and negative outcomes. But cumulative return reveals nothing about the extremes in performance that occur within the horizons. Because we specified the scenarios as paths, we are better able to observe each portfolio's within-horizon exposure to loss, which might incline us more toward the moderate or conservative portfolio.

Another advantage to defining scenarios as paths arises if we value consistency. The returns displayed in Table 4 enable us to measure consistency. One approach would be simply to compute the spread between the highest and lowest returns across the scenarios and through time for each portfolio. Although somewhat informative, these spreads do not consider the relative likelihood of the scenarios. A large spread between two unlikely scenarios might not suggest the same level of inconsistency as would a tighter spread between two more likely scenarios. We can address this issue by computing, considering the scenarios' relative probabilities, the standard deviation of returns across the scenarios and across the years for each portfolio in Table 2. These spreads and standard deviations are shown in Table 6.

Not surprisingly, the conservative portfolio offers the greatest degree of consistency across scenarios and through time.

Table 3. Alternative Portfolios

Asset Class	Conservative	Moderate	Aggressive
Stocks	40%	60%	80%
Bonds	50	35	20
Cash	10	5	0

Table 4. Portfolio Return Paths

Portfolio/Year	Baseline V	Shallow V	U-Shaped	W-Shaped	Depression	Stagflation
<i>Conservative</i>						
Year 1	9.1%	10.5%	2.1%	7.4%	-6.6%	0.3%
Year 2	1.7	2.5	6.8	-2.9	-1.0	-6.1
Year 3	2.3	1.8	5.6	3.8	15.5	-7.6
<i>Moderate</i>						
Year 1	11.5%	13.6%	0.5%	11.4%	-11.8%	0.8%
Year 2	3.0	4.0	9.0	-4.3	-7.9	-4.9
Year 3	1.4	1.6	6.2	1.6	13.5	-7.2
<i>Aggressive</i>						
Year 1	14.0%	16.8%	-1.2%	15.4%	-17.0%	1.2%
Year 2	4.4	5.5	11.2	-5.8	-14.8	-3.7
Year 3	0.6	1.4	6.8	-0.6	11.5	-6.8

Table 5. Portfolio Metrics

Measure	Baseline V	Shallow V	U-Shaped	W-Shaped	Depression	Stagflation	Probability-Weighted Average	Worst Case
Probabilities:	21.5%	24.0%	30.1%	5.9%	2.4%	16.1%		
<i>Cumulative return</i>								
Conservative	13.4%	15.3%	15.1%	8.3%	6.8%	-13.0%	9.6%	-13.0%
Moderate	16.5	20.1	16.3	8.3	-7.8	-11.1	11.8	-11.1
Aggressive	19.7	24.9	17.4	8.1	-21.2	-9.2	13.9	-21.2
<i>Maximum drawdown</i>								
Conservative	0.0%	0.0%	0.0%	-2.9%	-7.6%	-13.3%	-2.5%	-13.3%
Moderate	0.0	0.0	0.0	-4.3	-18.8	-11.8	-2.6	-18.8
Aggressive	0.0	0.0	-1.2	-6.3	-29.3	-10.2	-3.1	-29.3
<i>Within-horizon loss</i>								
Conservative	0.0%	0.0%	0.0%	0.0%	-7.6%	-13.0%	-2.3%	-13.0%
Moderate	0.0	0.0	0.0	0.0	-18.8	-11.1	-2.2	-18.8
Aggressive	0.0	0.0	-1.2	0.0	-29.3	-9.2	-2.5	-29.3
<i>Worst loss</i>								
Conservative	0.0%	0.0%	0.0%	-2.9%	-6.6%	-7.6%	-1.6%	-7.6%
Moderate	0.0	0.0	0.0	-4.3	-11.8	-7.2	-1.7	-11.8
Aggressive	0.0	0.0	-1.2	-5.8	-17.0	-6.8	-2.2	-17.0
<i>Number of negative years</i>								
Conservative	0	0	0	1	2	2	0.43	2
Moderate	0	0	0	1	2	2	0.43	2
Aggressive	0	0	1	2	2	2	0.79	2

Table 6. Consistency

Portfolio	Spread	Standard Deviation
Conservative	23.1%	4.9%
Moderate	25.4	5.9
Aggressive	33.8	7.2

Summary

We introduced an enhancement to scenario analysis in which we define prospective economic scenarios as paths for economic variables rather than as single-horizon averages. We discussed several merits of this approach. The key benefit of defining scenarios as paths is that it enables us to estimate probabilities and forecast asset class returns more reliably than other approaches. We showed how to extend the methodology first introduced by Czaronis et al. (2020) to assign probabilities to these multistage scenarios. Next, we applied a novel forecasting technique called “partial sample regression” to map the multistage economic scenarios onto return paths for asset classes. We illustrated this new approach with a case study. We produced a variety of path-dependent metrics by which to evaluate alternative portfolios, which revealed that preferences based on path-dependent outcomes could lead investors to choose a different portfolio than they would choose in the absence of this rich set of results.

Although we would like to test the innovations that we propose in this article empirically, no observable

baseline exists by which to evaluate our recommended methodology. We cannot know the subjective probabilities assigned to scenarios by a representative group of investors. Neither can we observe a representative set of return forecasts. We offer our approach to scenario analysis as a data-driven mathematical framework in lieu of a subjective process, just as John Burr Williams did when introducing the dividend discount model and Harry Markowitz did with mean–variance analysis.

Editor’s Note

Submitted 28 May 2020

Accepted 21 October 2020 by Stephen J. Brown

Disclaimer: 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 State Street Global Markets, State Street Corporation and its affiliates, Windham Capital Management, or GIC, Singapore’s sovereign wealth fund, GIC Private Limited, and its affiliates (collectively, “GIC”). While State Street, Windham Capital Management and GIC have collaborated for purposes of conducting research and developing this article, State Street, Windham Capital Management, and GIC are not engaged in any joint venture, affiliated in any way, or collectively providing or offering any services or products.

Notes

1. We describe the Mahalanobis distance in more detail later in this article.
2. Czaronis, Kritzman, and Turkington (2020a) showed that this approach yields more reliable forecasts of factor returns than conventional linear regression analysis, and Czaronis, Kritzman, and Turkington (2020b) offered evidence that this approach improves the forecast reliability of the stock–bond correlation.
3. See Mahalanobis (1927, 1936).
4. The Mahalanobis distance is often multiplied by $1/N$ so that the average distance score across the dataset will equal 1. We excluded this simple scaling factor for purposes of our analysis. The measure is sometimes shown as the square root of this quantity, which is another form of scaling.
5. The probability density function for the multivariate normal distribution has a similar form but includes a constant term that ensures that the cumulative probability of all possible outcomes equals 1. The scaling is irrelevant to our analysis because we are interested in the relative probabilities of a discrete set of scenarios that we rescale to sum to 1.
6. For further discussion of noise-driven versus event-driven observations and their relationship to estimating risk, see Chow et al. (1999).
7. See Czaronis et al. (2020a) for a thorough discussion of this technique.
8. Most of our data come from the Jordà-Schularick-Taylor (JST) Macrohistory Database (Release 4 from May 2019), which is publicly available at <http://www.macrohistory.net/data>.

References

- Chow, G., E. Jacquier, M. Kritzman, and K. Lowry. 1999. "Optimal Portfolios in Good Times and Bad." *Financial Analysts Journal* 55 (3): 65–73.
- Czaronis, M., M. Kritzman, B. Pamir, and D. Turkington. 2020. "Enhanced Scenario Analysis." *Journal of Portfolio Management* 46 (4): 69–79.
- Czaronis, M., M. Kritzman, and D. Turkington. 2020a. "Addition by Subtraction: A Better Way to Forecast Factor Returns (and Everything Else)." *Journal of Portfolio Management* 46 (8): 98–107.
- . 2020b. "The Stock–Bond Correlation." MIT Sloan Working Paper 6108-20.
- Lin, R., E. Khalastchi, and G. Kaminka. 2010. "Detecting Anomalies in Unmanned Vehicles Using the Mahalanobis Distance." Working paper. The Maverick Group, Computer Science Department, Bar-Ilan University, Ramat-Gan, Israel.
- Mahalanobis, P. C. 1927. "Analysis of Race-Mixture in Bengal." *Journal of the Asiatic Society of Bengal* 23: 301–33.
- . 1936. "On the Generalised Distance in Statistics." *Proceedings of the National Institute of Sciences of India* 2 (1): 49–55.
- Nasief, H., I. Rosado-Mendez, J. Zagzebski, and T. Hall. 2019. "A Quantitative Ultrasound-Based Multi-Parameter Classifier for Breast Masses." *Ultrasound in Medicine and Biology* 45 (7): 1603–16.
- Su, C., and S. Li. 2002. "A Mahalanobis Distance Based Classifier for Diagnosing Diseases." *Journal of the Chinese Institute of Industrial Engineers* 19 (5): 41–47.
- Wang, P., C. Su, K. Chen, and N. Chen. 2011. "The Application of Rough Set and Mahalanobis Distance to Enhance the Quality of OSA Diagnosis." *Expert Systems with Applications* 38 (6): 7828–36.