JOURNAL OF INVESTMENT MANAGEMENT, Vol. 21, No. 3, (2023), pp. 29–41

www.joim.com

THE DETERMINANTS OF INFLATION

William Kinlaw[a], *Mark Kritzman*[b,c], *Michael Metcalfe*[d] *and David Turkington*[e]

The authors apply a Hidden Markov Model to identify regimes of shifting inflation and then employ an attribution technique based on the Mahalanobis distance to identify the economic variables that dynamically determine the trajectory of inflation. Their analysis enables policymakers to focus on the most effective tools to manage inflation, and it offers guidance to investors whose strategies might benefit from knowledge of the prevailing determinants of inflation. Their analysis reveals that as of February 2022, the most important determinant of the recent spike in inflation was spending by the federal government.

We apply a Hidden Markov Model to identify regimes of shifting inflation and then employ an attribution technique based on the Mahalanobis distance to measure the sensitivity of these regimes to different economic variables. This methodology reveals how the influence of economic variables on inflation shifts through time. Our analysis enables policymakers to focus on the most effective tools to manage inflation, and it offers guidance to investors whose strategies might benefit from knowledge of the prevailing determinants of inflation. Our analysis reveals that as of February 2022, the most important determinant of the recent spike in inflation was spending by the federal government.

We organize the paper as follows. First, we describe how we use a Hidden Markov Model to identify inflation regimes, and we provide descriptive information of these regimes. Next, we define the variables we consider as candidates for the determinants of inflation, and we describe our methodology for measuring their relative importance. We then present our results and conclude with a summary.

1 Inflation Regimes

We define an inflation regime at a given time as the shift from the rate of inflation during the past

[a]State Street, Boston, MA, USA.
E-mail: wbkinlaw@statestreet.com
[b]Windham Capital Management in Boston, MA, USA.
[c]MIT Sloan School of Management in Cambridge, MA, USA. E-mail: kritzman@mit.edu
[d]State Street, London, UK.
E-mail: mmetcalfe@statestreet.com
[e]State Street Associates, Cambridge, MA, USA.
E-mail: dturkington@statestreet.com

three years to the most recent one-year rate of inflation. We focus, therefore, not on the rate of inflation, but rather on the change in the rate of inflation, as shown in Equation (1). A positive value indicates that inflation accelerated, whereas a negative value means that inflation decelerated.

$$\textit{Inflation Shift} = \frac{CPI_t}{CPI_{t-12}} - \left(\frac{CPI_t}{CPI_{t-36}}\right)^{1/3} \quad (1)$$

In Equation (1), CPI is the Consumer Price Index for all urban consumers, not seasonally adjusted, and excluding food and fuel prices. We use monthly data from the Federal Reserve Bank of St. Louis covering the period January 1960 through February 2022.[1] Exhibit 1 shows how the rate of inflation has shifted during this period and underscores how exceptional the recent shift in inflation has been compared to its behavior during the past four decades.

Having defined our variable of interest, our next step is to uncover distinct regimes from our data. To do so, we employ a technique called a Hidden Markov Model. This technique assumes that the observed values come from multiple distributions that switch according to a hidden regime characteristic. Additionally, each regime has a degree of persistence from one period to the next. For example, if we are in an accelerating inflation regime this month, this regime might be more likely to prevail next month, as well. But there is also a chance that the regime will shift abruptly to a decelerating inflation regime. The term Markov refers to the assumption that the underlying regime characteristic follows a Markov process in that next period's regime probability depends only on the regime we are in today.

To fit a hidden Markov model to data, we select a characteristic to distinguish regimes, and we find the probability of transitioning from one regime to another, given our sample of historical values for the regime characteristic. Thankfully, the Baum–Welch algorithm turns this potentially laborious search into a computationally straightforward exercise.

To implement the Baum–Welch algorithm we first guess the probabilities of shifting from one regime

Exhibit 1 Inflation shift.

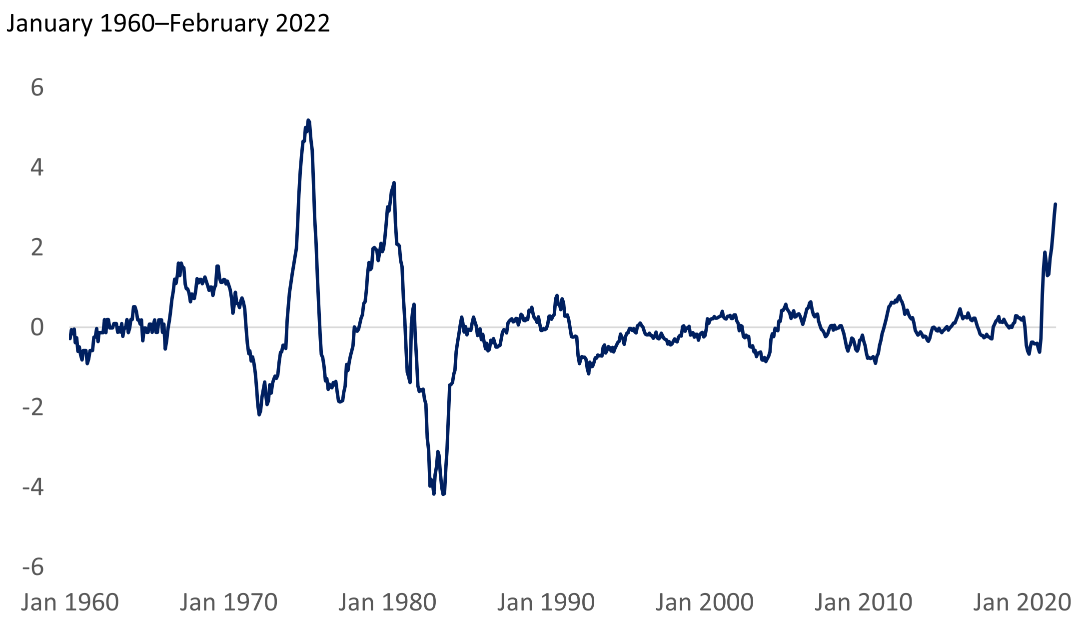

to another, along with the mean and standard deviation of the regime characteristic for each regime. These initial guesses are chosen arbitrarily. The algorithm then computes forward probabilities. For the first period, the algorithm evaluates the likelihood of each regime based on our initial guesses, together with that period's value for the regime characteristic and the distribution of the characteristic for each regime. For the next period, the algorithm evaluates the likelihood of each regime based on the new value of the characteristic and the same initial guesses, and accounting for the likelihood of each regime from the prior period. It iterates forward in this fashion until we have forward probabilities for every period. We now have a time series for each regime that tells us how likely it is that we would observe that value for the regime characteristic, given the distribution of the characteristic for each regime, based on everything that occurred previously. The algorithm captures the fact that some values for the regime characteristic are more likely to have come from one regime than another, given their distributions. It also captures the fact that a regime is more likely if it is highly persistent and believed to have prevailed in the preceding months.

Next, the algorithm follows the same procedure in reverse, to generate backward probabilities. It then combines the forward and backward probabilities into smoothed probabilities, which tell us the likelihood of each regime at each point in time, given what regimes were likely to have occurred before and after.[2] Using these regime likelihoods, we estimate the transition probabilities and distributions of outcomes for each regime, leading to a new iteration of the entire process which we repeat until the parameter estimates converge to stable values.

We applied this algorithm to uncover four regimes from our time series of the shift in the inflation rate, which we name based on their empirical characteristics.

Steady: Stable inflation, low volatility
Rising Stable: Rising inflation, low volatility
Rising Volatile: Rising inflation, high volatility
Disinflation: Sharply declining inflation

Exhibit 2 shows the transition frequencies from one regime to another throughout our sample, along with the mean value and standard deviation of our inflation shift variable. It reveals that all the regimes are highly persistent, and that the most persistent regime is rising inflation with high volatility. This observation does not bode well for the prospects of ameliorating the recent spike in inflation.

Exhibit 3 shows the historical periods corresponding to each regime. We assume that the regime with the highest likelihood from the Hidden Markov Model is the current regime.

Exhibit 2 Hidden Markov Model regimes for inflation shift.

Transition frequencies (%)	Steady	Rising Stable	Rising Volatile	Disinflation
From Steady to...	91.0	5.2	0.0	3.7
From Rising stable to...	6.5	92.4	1.1	0.0
From Rising volatile to...	0.0	0.6	97.7	1.7
From Big drop to...	6.1	0.0	2.3	91.6
Average (%)	−0.1	0.3	0.4	−0.6
Standard deviation (%)	0.1	0.2	2.1	0.2

Exhibit 3 Historical Hidden Markov Model regimes.

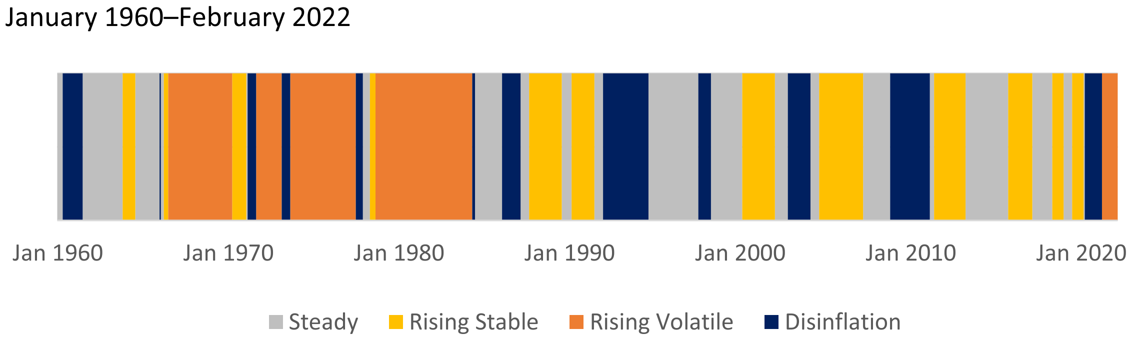

Exhibit 4 Economic variables.

Category	Variable	Data	Units	Frequency/ Start Date
Cost Push	Producer Prices	Producer price index*	Year-over-year change	Monthly, 1913
Demand pull	Wages & Salaries	Compensation of employees received (wage and salary disbursements)	Year-over-year change	Monthly, 1959
	Personal Consumption	Personal consumption expenditures	Year-over-year change	Monthly, 1959
Inflation Expectations	Inflation Expectations	Michigan Consumer Survey**	One-year change in one-year-ahead forecast	Monthly, 1978
Monetary Policy	Interest Rates	Federal funds rate	Year-over-year change	Monthly, 1954
	Yield Curve	10-year yield minus federal funds rate	One-year moving average	Monthly, 1954
	Money Supply	Money supply M2***	Five-year change	Monthly, 1960
Fiscal Policy	Federal Spending	Federal spending****	Five-year change	Quarterly, 1948

*We use the core PPI which excludes food and energy costs. Prior to 1974 the core PPI is not available so we use the headline PPI. Ideally, we would have also included a variable in the supply category to capture supply chain disruptions; however, no such series exists with sufficient historical data. **Prior to 1979 survey data is unavailable; for this period we use the most recent one-year change in headline CPI minus the prior one-year change in headline CPI. ***The one-year change in M2 becomes available in January 1960; we begin with the one-year change and increase the window until January 1964, then we roll forward a five-year change. We note that changes in M2 can be driven by changes in the Federal Funds Rate as well as by Quantitative Easing or Tightening. It may therefore be sensible to consider the attribution to these factors collectively as the impact of monetary policy. ****We repeat the most recent quarterly value for each month in our analysis. Notes on data sources: We download all data from the Federal Reserve Bank of Saint Louis ALFRED database. Specifically, we use the following series: WPUFD4131_PC1_20220413 for producer prices, A576RC1_PC1_20220331 for wages & salaries, PCE_PC1_20220331 for personal consumption, MICH_CH1_20220225 for inflation expectations, FEDFUNDS_CH1_20220401 for interest rates, FEDFUNDS_20220401 and GS10_20220401 for the level and slope of the yield curve, M2NS_PC1_20220322 for money supply and NA000283Q_PC1_20220428 for federal spending.

We next define the variables we consider as candidates for determining the path of inflation, and we describe the methodology we use to measure their influence.

2 Variables and Methodology

2.1 *Economic variables*

We define eight economic variables to help us understand the determinants of inflation, which fall into five categories: cost push, demand pull, inflation expectations, monetary policy, and fiscal policy. These variables are shown in Exhibit 4. We compile these variables into a set of monthly observations starting in January 1960 through February 2022 when the most recent data are available. We do not account for revisions, because our goal is to identify the determinants of inflation rather than to forecast inflation out of sample. For this purpose, we want the best data available at each point in time.

Exhibit 5 shows the average values of each variable during each of our four regimes.

Exhibit 5 reveals intuitive patterns. For example, we observe that producer prices, wages and salaries, and personal consumption expenditures rose most rapidly during the Rising Volatile regime and most slowly during the Disinflation regime. On the other hand, federal spending rose somewhat rapidly during the Disinflation regime, which tends to capture recessions.

Exhibit 6 shows conditional returns for major asset classes during each regime in both nominal and real terms.[3] We observe intuitive patterns in Exhibit 6 as well, which also presents average inflation during each regime. Rising and volatile inflation regimes showed the worst performance of the various regimes for both stocks and bonds in nominal and real terms; only cash offered some protection, albeit very little in real terms. This highlights the importance for investors to understand the type of inflation regime we are in at any given time and how that might change.

2.2 *Attribution methodology*

We now show how we measure the influence of the economic variables on the path of inflation, based on a statistic called the Mahalanobis distance.[4]

The Mahalanobis distance was introduced originally in 1927 and modified in 1936 to analyze

Exhibit 5 Averages of variables during each regime.

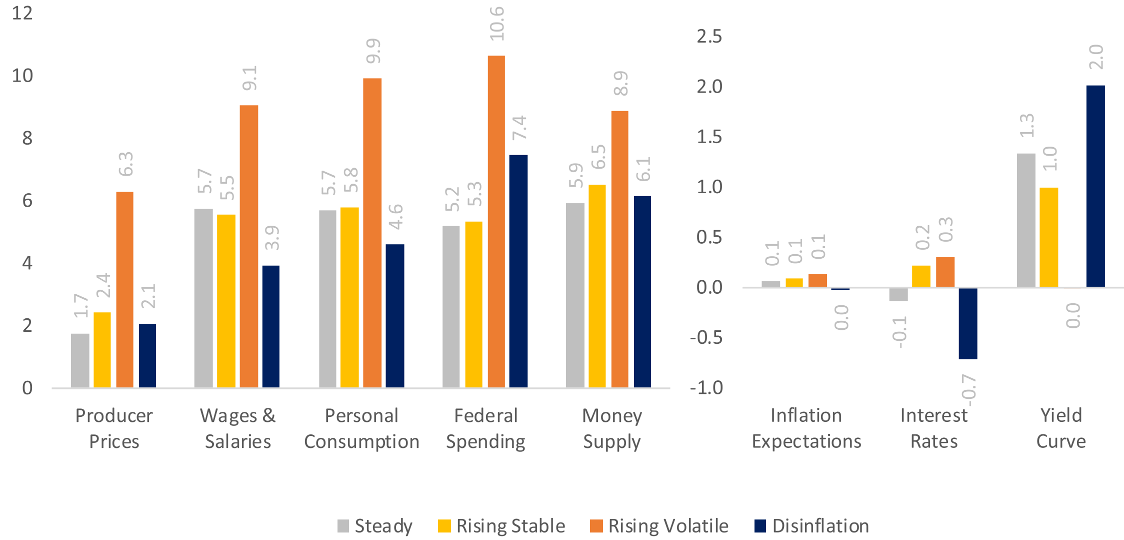

Exhibit 6 Conditional returns for major asset classes.

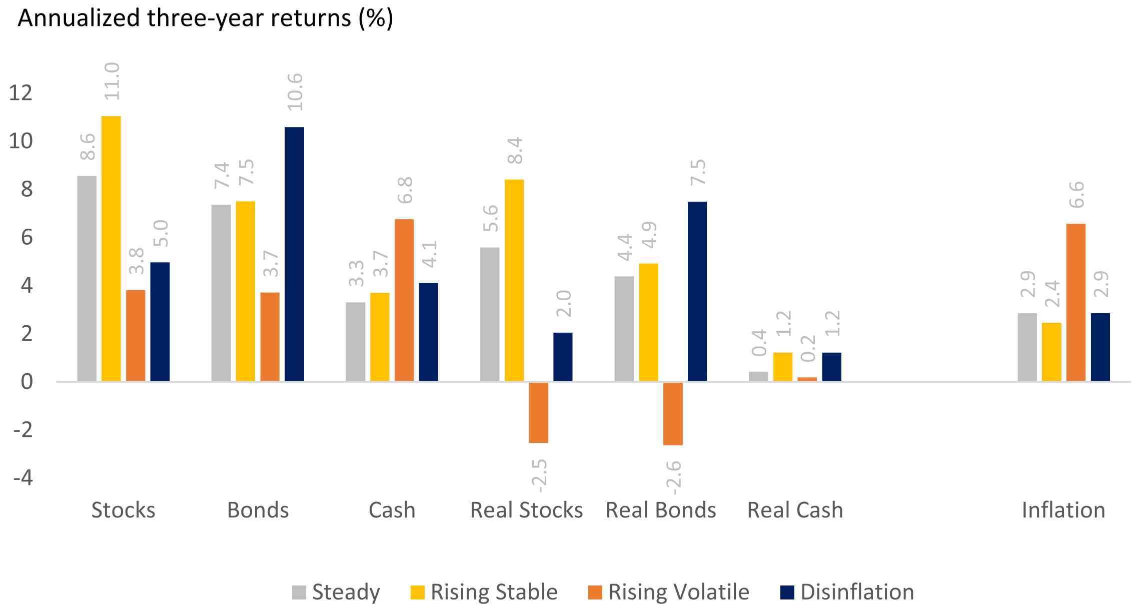

resemblances in human skulls among castes in India.[5] Mahalanobis compared a set of measurements for a chosen skull to the average of those measurements across skulls within a given caste. He also compared the co-occurrence of those measurements for a chosen skull with their covariation within the caste. He summarized these comparisons in a single number which he used to place a given skull in one caste or another.

The Mahalanobis distance has since been applied across many different fields. Chow *et al.* (1999) derived the Mahalanobis distance independently to measure turbulence in the financial markets. Su and Li (2002) applied the Mahalanobis distance to diagnose liver diseases. Wang *et al.* (2011) used the Mahalanobis distance to diagnose obstructive sleep apnea, and Nasief *et al.* (2019) used it to diagnose breast cancer. The Mahalanobis distance has also been applied to detect anomalies in self-driving vehicles (Li *et al.*, 2010), and to improve the forecast reliability of linear regression analysis (Czasonis *et al.*, 2022). Perhaps the application that is most like our analysis of inflation is Kinlaw *et al.* (2021), which used the Mahalanobis distance to create an index of the business cycle.

The Mahalanobis distance, as originally conceived to measure the statistical similarity of human skulls, is given in Equation (2).

$$d = (x - \mu)\Sigma^{-1}(x - \mu)' \tag{2}$$

In Equation (2), d equals the Mahalanobis distance, x equals a row vector of values for a set of dimensions used to characterize a skull, μ equals the average values from a chosen group of skulls, and Σ^{-1} equals the inverse of the covariance matrix of the group's dimensions, and $'$ denotes matrix transpose. The term $(x - \mu)$ captures how similar each dimension, by itself, is to the group's average values. By multiplying $(x - \mu)$ by the inverse of the covariance matrix, it captures how similar the co-occurrence of the dimensions is to their co-occurrence in the group. This multiplication also converts the variables into a common unit of variation which is defined as one standard deviation. This feature is especially important

Exhibit 7 Scatter plot of two hypothetical skull dimensions.

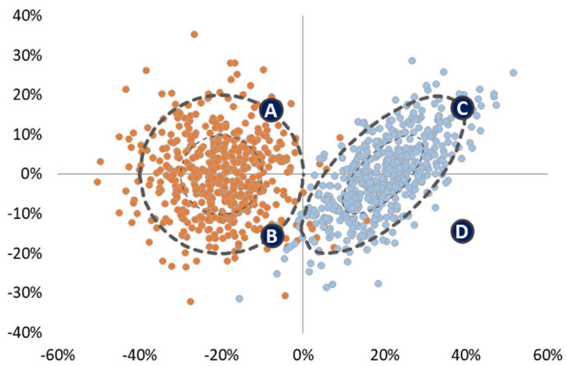

for economic variables, some of which may be measured as percentage changes whereas others may be measured as levels.

Exhibit 7 illustrates the features of the Mahalanobis distance in two dimensions using hypothetical data. Suppose the dots represent values of the two dimensions for various skulls. The cluster of observations on the left-hand side of the chart pertains to hypothetical skull measurements for a group where the attributes are uncorrelated and have equal standard deviations. When the correlation among the variables equals zero, the Mahalanobis distance of a given point reflects the average of its squared z-scores in the same fashion as the most common measurement of physical distance, the Euclidean distance. The two data points shown, A and B, have identical Euclidean distances from this group's average, and because the variables are uncorrelated these points also have identical Mahalanobis distances from the mean. Other points that fall on the same iso-distance curve, shown as a dotted line circle, also share the same Mahalanobis distance. The cluster of observations on the right-hand side of the chart is from a different sample with different characteristics: the attributes are positively correlated and have different standard deviations. Points C and D have identical Euclidean distances from their average for the group, but their Mahalanobis distances are not the same. Point C conforms to the typical correlation pattern of the data, so it is less unusual and less distant from the center of the distribution. Point D reflects an opposite alignment of the two variables, which is highly unusual, given the positive correlation. Thus, points C and D fall on different iso-distance curves. The iso-distance curves for the Mahalanobis distance are ellipses rather than circles, owing to the different standard deviations of the variables and their non-zero correlation. When there are more than two variables, the iso-distance curve is an ellipsoid within a higher-dimensional space. This is difficult to visualize, but the same intuitive interpretation applies for any number of dimensions.

In summary, the Mahalanobis distance accounts for two important features of statistical similarity. It scales each value of the chosen observation by the variability of the values in the group, which converts all values into common units. And it accounts for the co-occurrence of the values for a given observation.

We apply the Mahalanobis distance to our set of economic variables in the same manner Kinlaw *et al.* (2021) deployed it to construct a business cycle index. A distinguishing feature of their methodology, which is critical to our purpose in this research, is that it enables us to identify the determinants of inflation dynamically. We can do so because we consider the distribution of each inflation regime separately and measure their likelihoods as relative probabilities, which has the consequence of shifting the weights that are placed on the input variables based on their prevailing values. This feature critically distinguishes our approach from regression analysis, for example, which can only measure a variable's influence as a static coefficient that is applied across the entire sample. This distinction will become more apparent as we describe the details of our methodology.

First, we compute the Mahalanobis distance of a row vector of the current values of our economic variables, x, from the average of those variables' values during regime r, which is a row vector μ_r. In Equation (2), these values represented skull dimensions instead of values of economic variables.

$$d_r = (x - \mu_r)\Omega_r^{-1}(x - \mu_r)' \tag{3}$$

In Equation (3), Ω_r^{-1} is the inverse of the covariance matrix of the variables during regime r, and $'$ denotes the transpose of a vector (in this case, from a row to a column). Next, we convert the Mahalanobis distance into a statistical likelihood according to a normal distribution:

$$\xi_r(d_r) = (\det(2\pi\Omega))^{-1/2}e^{-d_r/2} \tag{4}$$

In this expression, det is the matrix determinant, and e is the base of the natural logarithm.

It is important to note that, although we assume normality for each regime separately, we consider the regimes collectively as a mixture of normal distributions, which would almost certainly be non-normal. Nonetheless, for those who are troubled by the assumption of regime-specific normality, there are simple techniques for tilting the distributions of the specific regimes away from normality toward distributions that are more aligned with empirical outcomes.[6]

Next, we rescale the likelihood of regime r by the sum of all regime likelihoods, which we interpret as a probability.

$$p_r = \frac{\xi_r}{\sum_{all\ regimes\ i} \xi_i} \tag{5}$$

To determine the importance of each variable at a given point in time, we first compute the sensitivity of each probability measure to a change in variable values, x, by taking its derivative using the chain rule of calculus: $\frac{\partial p_r}{\partial x} = \frac{\partial p_r}{\partial \xi_r}\frac{\partial \xi_r}{\partial d_r}\frac{\partial d_r}{\partial x}$. The resulting derivative indicating the sensitivity of the value for regime r is given by:

$$\frac{\partial p_r}{\partial x} = p_r\left(\left(\sum_{\substack{regimes\\ i\neq r}} p_i\frac{\partial d_i}{\partial x}\right) - (1-p_r)\frac{\partial d_r}{\partial x}\right) \tag{6}$$

In Equation (6), $\frac{\partial d_r}{\partial x}$ (and likewise $\frac{\partial d_i}{\partial x}$) is defined as:

$$\frac{\partial d_r}{\partial x} = \Omega_r^{-1}(x - \mu_r)' \tag{7}$$

The sensitivity $\frac{\partial p_r}{\partial x}$ is a vector containing the sensitivity of p_r to each variable contained in the vector x. In our empirical analysis we have four regimes, which we label as A, B, C, and D. To make this expression more explicit, the sensitivity of regime A, for example, is:

$$\frac{\partial p_A}{\partial x} = p_A\left(p_B\frac{\partial d_B}{\partial x} + p_C\frac{\partial d_C}{\partial x} + p_D\frac{\partial d_D}{\partial x} - (1-p_A)\frac{\partial d_A}{\partial x}\right) \tag{8}$$

We summarize the total sensitivity for a given point in time by taking the average of the absolute values of the four regime sensitivity vectors. We must take the absolute value, otherwise the result will always equal zero because the probabilities are rescaled to sum to one.

$$sensitivity = \frac{1}{4}\left(\left|\frac{\partial p_A}{\partial x}\right| + \left|\frac{\partial p_B}{\partial x}\right| + \left|\frac{\partial p_C}{\partial x}\right| + \left|\frac{\partial p_D}{\partial x}\right|\right) \tag{9}$$

We multiply each element of the sensitivity vector by the standard deviation for that variable measured over the full sample. The result represents the sensitivity of the collective probabilities to a standardized shock which we can compare across variables.

Finally, we rescale the sensitivity vector by the sum of its components to arrive at a measure of the relative importance for each variable.

$$variable\ importance = \frac{sensitivity \circ \sigma}{|sensitivity \circ \sigma|} \quad (10)$$

In this expression, σ is a vector of the standard deviations of the variables measured over the full sample, $\circ$ denotes element-by-element multiplication between two vectors, and $|\cdot|$ computes the absolute value norm of a vector.

3 Results

We first provide evidence in Exhibit 8 that the relative likelihoods of the respective regimes as given by Equation (5), together with our choice of economic variables, matches the occurrence of the regimes based on our application of the Hidden Markov Model.

Exhibit 8 Regime Likelihood Indices (lines) and Associated HMM Regimes (shading).

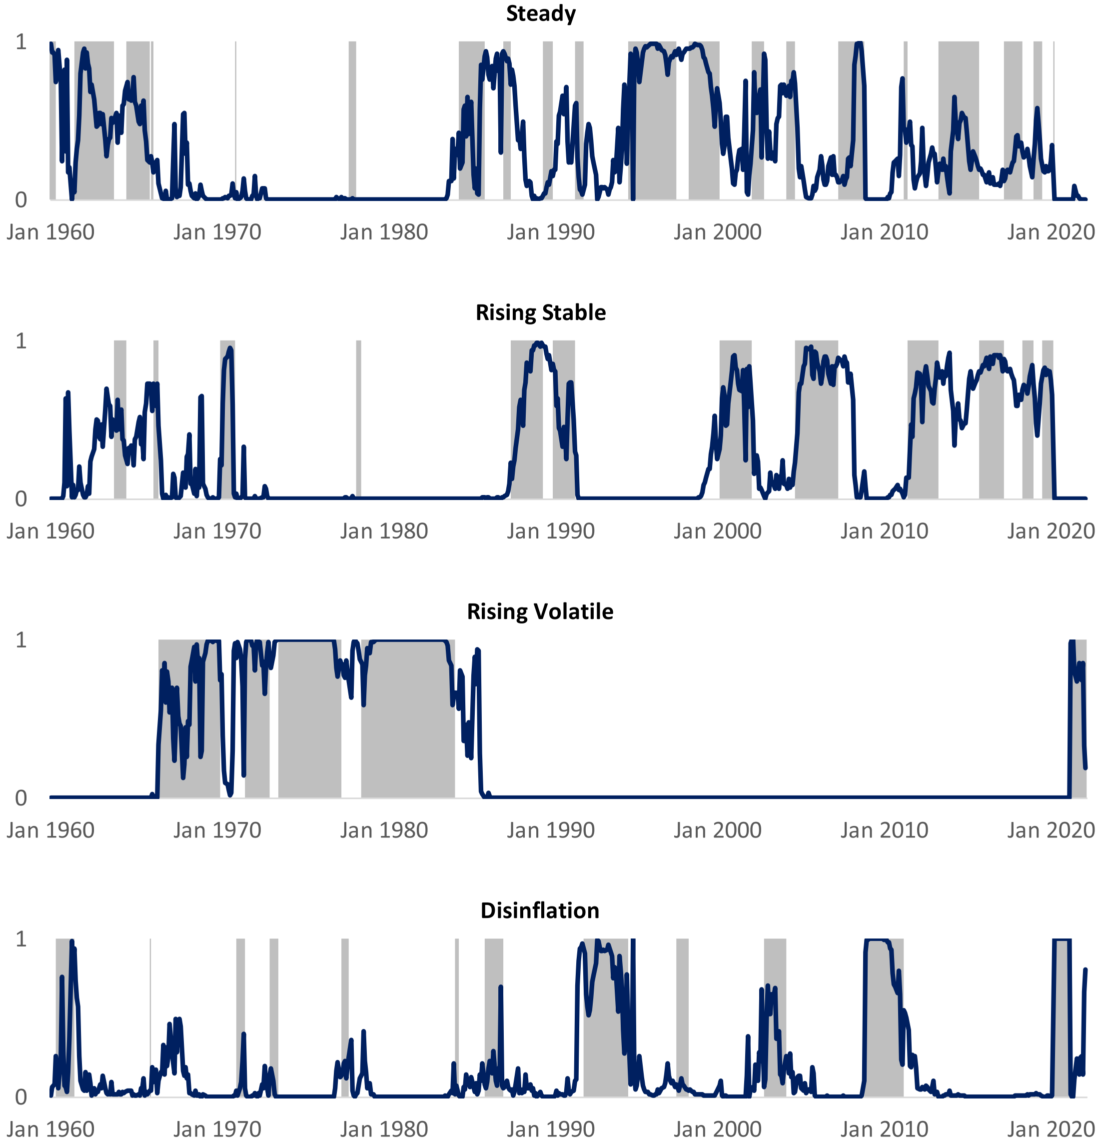

We observe from Exhibit 8 that the likelihoods estimated from Equation (5) track quite closely with the actual regimes. This tight association should give us confidence that the regimes, which we identified by the characteristic variable, inflation shift, and their relative probabilities, which we estimated with a separate set of economic variables, are well specified.

Exhibit 9 shows the time varying influence of each variable on the relative likelihoods of the different regimes, as determined by Equation (10).

It is evident from Exhibit 9 that the relative importance of the variables is not constant. It changes because we account for means and covariances that shift across regimes. It is important to keep in mind for Exhibits 9, 10, and 11 that these quantities reflect the influence of the variables and not the size of the variables. For example, if the influence of federal spending is large, it does not mean that federal spend was necessarily large. It could have had a large influence because it was larger than average, as was the case in the later months of the Covid pandemic, or because it was smaller than average, as it was late in the Global Financial Crisis. Federal spending could also have a large impact even if it conforms to its normal pattern. Our point is simply that the size of a variable's impact on the pace of inflation need not correspond to the magnitude of the variable.

In Exhibit 10, we highlight the changes in the determinants of inflation for the period 1967 through 1977, which was characterized mostly by rising and volatile inflation, as shown in Exhibit 3. In the mid-1960s, the inflation rate began to increase and remained at elevated levels for more than a decade, reaching a level of more than 14% in 1980. What caused inflation to last so

Exhibit 9 Historical attribution.

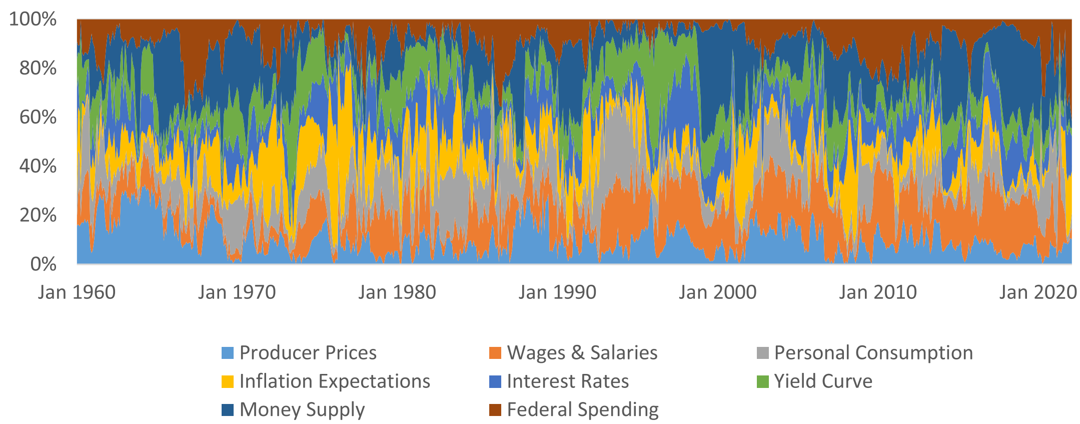

Exhibit 10 Attribution during the Late 1960s and 1970s.

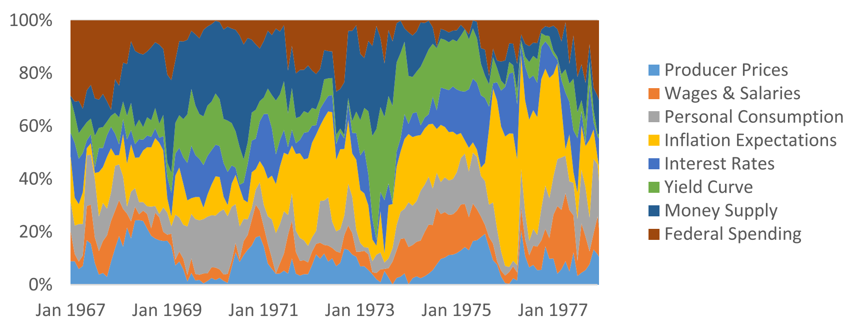

Exhibit 11 Attribution during the Global Financial Crisis and Covid pandemic.

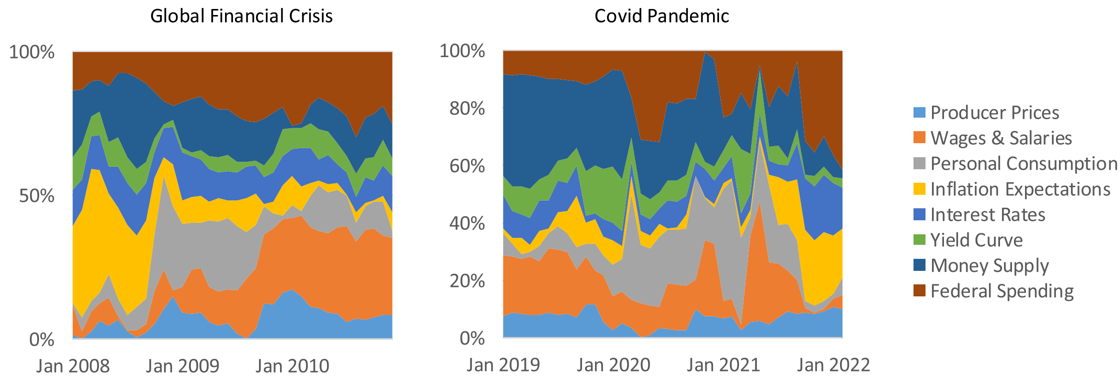

long is a subject of debate among economists, but there is little debate about its source: Federal Reserve policies that expanded the money supply. Our methodology reveals that the money supply was the most persistent driver in the late 1960s and early 1970s. It also shows that federal spending became important in the late 1960s as U.S. President Lyndon Jonson implemented his Great Society programs and increased spending on the war in Vietnam. In the summer of 1971, with rising pressure on the gold standard, then-President Richard Nixon exited the Bretton Woods system, untethering the U.S. dollar from gold. The dollar depreciated continuously for the next several years and inflation picked up further, which caused inflation expectations to soar (as shown in yellow in Exhibit 10). Nixon then introduced wage and price controls between 1971 and 1974, a period during which our methodology shows wages and personal consumption expenditures as important drivers. But these policies failed, and inflation expectations continued to rise. Only after Paul Volcker was appointed Chairman of the Federal Reserve system in August of 1979 and exerted tighter control over money supply growth—which drove the economy into recession—do we see inflation expectations begin to recede in importance.

Exhibit 11 focuses on the Global Financial Crisis as well as the Covid Pandemic. At the onset of the financial crisis at the beginning of 2008, the path of inflation was driven mostly by inflation expectations. Then in November of 2008 personal consumption became the primary determinant of the path of inflation as consumers retrenched, and it remained a dominant force, along with federal spending, through mid-2009 as the financial crisis ran its course. Then in the fourth quarter of 2009, wages and salaries became the dominant force as unemployment soared. These forces kept inflation low for an extended period, which was compounded by a retrenchment in federal spending as the government sought to control debt.

Exhibit 11 also shows that the path of inflation at the inception of the Covid pandemic in March 2020 was driven by demand factors, but by the end of 2021 inflation expectations and federal spending were the dominant forces driving inflation.

Finally, we show the relative influence of the economic variables on our most recent observation, February 2022, which corresponds with the recent burst in inflation. Exhibit 12 reveals that spending by the federal government stands out unambiguously as the most important determinant of inflation in this environment. So, what began looking like the inflation regime of the Global Financial Crisis, by February 2022 was akin to the

Exhibit 12 Current attribution (as of February 2022).

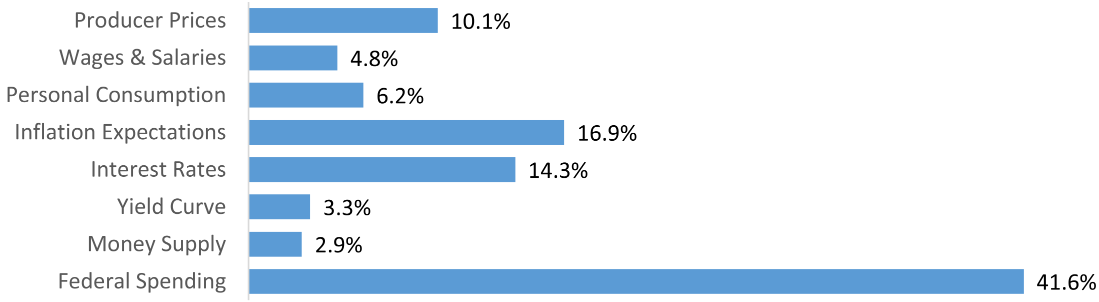

late 1960s when fiscal spending and then inflation expectations drove the path of inflation.

4 Summary

We deployed a Hidden Markov Model to identify regimes of shifting inflation based on inflation data from January 1960 through February 2022. This analysis uncovered four distinct regimes: a Steady regime in which inflation was stable and volatility was low; a Rising Stable regime in which inflation was rising and volatility was low; a Rising Volatile regime in which inflation was rising and volatility was high; and a Disinflation regime in which inflation was declining sharply.

We then proposed several economic variables as potential determinants of the path of inflation. We first computed the Mahalanobis distance of each month throughout our sample to each of the four regimes. We then converted these distances into the relative likelihood that a given month belonged to a particular regime. And we provided evidence confirming that our estimates of the relative likelihoods of the regimes matched the occurrence of the regimes as identified by the Hidden Markov model.

We then proceeded to compute the derivatives of the monthly shift in inflation throughout our sample with respect to the economic variables to measure the relative influence of each variable. Our analysis revealed that spending by the federal government stands out as the most important determinant of the recent spike in inflation.

Notes

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.

Acknowledgments

We thank Alberto Cavallo, Deborah Lucas, and Andrew Weisman for their helpful comments.

Endnotes

1 Specifically, we source the following series from the Federal Reserve Bank of Saint Louis's ALFRED database: CPILFENS_PC1. This series represents one-year percent changes in CPI. We derive also three-year changes from this series to compute the shift measure given by Equation (1).

2 We calibrate the HMM as follows. The initial guess for the transition matrix is 80% persistence for each regime and equal probability of switching to each of the other

three regimes. We select these values based on the intuition that economic regimes are typically quite persistent. The initial guess for the starting probabilities is 25% for each regime. The initial guess for mean and sigma is equal to the full-sample mean and sigma with random noise added to each regime's guess so that the four guesses are not identical. Due to the random noise, the answer can be different each time we re-run the algorithm. The Baum–Welch algorithm always converges to a local optimal solution, but is not guaranteed to find the global optimal. So we run the whole algorithm 50 times and use the best solution. This approach yields consistency across the runs and produces manifestly reasonable outputs; we therefore assume that we are at or near the global best fit. We initially experimented with three regimes (expecting them to capture high, medium, and low inflation) but discovered that the data are best fit with two rising regimes (one volatile and one stable) so expanded the model to incorporate four regimes. For more information about the Baum–Welch algorithm, see Baum *et al.* (1970). Also, Kritzman *et al.* (2012) offer an intuitive description with a simple example.

[3] We use the following data for each asset return series: stocks are S&P 500 returns from Robert Shiller's website, bonds are Bloomberg Long Government Bond Index returns from Datastream (and from Ibbotson prior to February 1973), and cash is the risk-free rate from Ken French's website. To compute real returns for each month, we subtract the monthly change in CPI from the nominal asset return for that month.

[4] For a thorough description of the close connections of the Mahalanobis distance to traditional statistics, see Czasonis *et al.* (2022).

[5] See Mahalanobis (1927, 1936).

[6] To tilt a normal distribution toward an empirical distribution, we raise the term d in Equation (4) to an exponent (j, for example) that is estimated as follows. We first calculate the Mahalanobis distance each period from the historical average. We simulate a larger sample of distances, $d_{simulated}$, using Monte Carlo simulation based on the mean and standard deviation of the distribution of empirical distances. We evaluate $d_{simulated}^{j}$ for a range of $j \in (0, 1)$ and calculate the resulting average simulated sample skewness and kurtosis. We select j that best matches the distribution of the empirical distances based on the minimum Euclidean distance of skewness and kurtosis. Or we can minimize the Kullback–Leibler divergence between the simulated and empirical distributions.

References

Baum, L., Petrle, T., Soules, G., and Weiss, N. (1970). "A Maximization Technique Occurring in Statistical Analysis of Probabilistic Functions of Markov Chains," *Annals of Mathematical Statistics* **41**(1).

Chow, G., Jacquier, E., Kritzman, M., and Lowry, K. (1999). "Optimal Portfolios in Good Times and Bad," *Financial Analysts Journal* **55**(3), 65–71 (May/June).

Czasonis, M., Kritzman, M., and Turkington, D. (2022). *Prediction Revisited: The Importance of Observation.* (John Wiley & Sons, Inc. Hoboken, New Jersey).

Kinlaw, W., Kritzman, M., and Turkington, D. (2021). "A New Index of the Business Cycle," *Journal of Investment Management* **19**(3) (July).

Kritzman, M., Page, S., and Turkington, D. (2012). "Regime Shifts: Implications for Dynamic Strategies," *Financial Analysts Journal* **68**(3) (May/June).

Lin, R., Khalastchi, E., and Kaminka, G. (2010). *Detecting Anomalies in Unmanned Vehicles Using the Mahalanobis Distance*. The Maverick Group, Computer Science Department, Bar-Ilan University, Ramat-Gan, Israel 52900.

Mahalanobis, P. C. (1927). "Analysis of Race-Mixture in Bengal," *Journal of the Asiatic Society of Bengal* **23**, 301–333.

Mahalanobis, P. C. (1936). "On the Generalised Distance in Statistics," *Proceedings of the National Institute of Sciences of India* **2**(1), 49–55.

Nasief, H., Rosado-Mendez, I., Zagzebski, J., and Hall, T. (2019). "A Quantitative Ultrasound-Based Multi-Parameter Classifier for Breast Masses," *Ultrasound in Medicine & Biology* **45**(7), 1603–1616 (July).

Su, C. and Li, S. (2002). "A MD Based Classifier for Diagnosing Diseases," *Journal of Chinese Institute of Industrial Engineers* **19**(5), 41–47 (September).

Wang, P., Su, C., Chen, K., and Chen, N. 2011. "The Application of Rough Set and Mahalanobis Distance to Enhance the Quality of OSA Diagnosis," *Expert Systems with Applications* **38**(6), 7828–7836 (June).

Keywords: Inflation; economics; business cycle; federal spending; fiscal policy; causes of inflation; mahalanobis distance and hidden markov models.