Working Paper Series
Congressional Budget Office
Washington, D.C.

Inflation Expectations and Their Formation

Jeffrey Schafer
Congressional Budget Office
Jeffrey.Schafer@cbo.gov

Working Paper 2022-03

March 2022

To enhance the transparency of the work of the Congressional Budget Office and to encourage external review of that work, CBO's working paper series includes papers that provide technical descriptions of official CBO analyses as well as papers that represent independent research by CBO analysts. Papers in this series are available at http://go.usa.gov/xUzd7.

This paper has not been subject to CBO's regular review and editing process. The views expressed here should not be interpreted as CBO's.

I am grateful to Olivier Coibion of the University of Texas, Austin, and Mathieu Pedemonte and Randal Verbrugge of the Federal Reserve Bank of Cleveland for reviewing and providing detailed and valuable comments on this paper. (The assistance of external reviewers does not imply their responsibility for the paper's content.) The paper has also benefited greatly from the guidance of Bob Arnold, Mark Doms, John Kitchen, and Jeff Werling. Ron Gecan offered insightful comments on this paper. I also gratefully acknowledge the editorial assistance provided by Christine Bogusz.

www.cbo.gov/publication/57398

Abstract

This paper reviews theory and evidence on how consumers and firms form their expectations about inflation and how monetary policymakers might influence that process; both of those factors have implications for the Congressional Budget Office's baseline projections and policy analyses.

Inflation expectations are important because they affect decisions that determine actual inflation. Surveys of consumers and business managers provide especially useful data for studying the process of forming expectations. Empirical evidence suggests that economically meaningful information frictions exist in the process of forming expectations and that the structure of those information frictions is such that realizing more persistent inflation or increasing the incentive to collect information on inflation—perhaps as the result of ongoing supply disruptions, such as the ones precipitated by the coronavirus pandemic—would reduce them.

There is also empirical evidence of substantial biases in individual-level inflation forecasts, which possibly amplify trends in the realized inflation data. Surveys reveal that, before the onset of the pandemic, firm managers did not pay much attention to aggregate inflation dynamics and were not well informed about the objectives or actions of monetary policymakers. If the observed level of information frictions decreased enough at the same time that individual decisionmakers' inflation forecasts continued to amplify trends in the realized inflation data, then aggregate measures of inflation expectations might begin to amplify those trends as well.

Keywords: inflation, inflation expectations, noisy information, diagnostic expectations

JEL Classification: D83, D84, E17, E31, E37

Contents

Section 1: Introduction

Inflation expectations matter. Theoretical models, from expectations-augmented Phillips curves to the canonical New Keynesian Phillips curve, predict that inflation expectations play a crucial role in determining actual inflation. Empirical evidence suggests that changes in inflation expectations can, in fact, lead to changes in actual inflation. Measures of long-run inflation expectations can improve inflation forecasts by capturing trend dynamics.

Inflation expectations cannot be observed directly, although they can be gauged through market- and survey-based measures. Market-based measures are derived from the premium that financial market participants are willing to pay for a security that is indexed to inflation instead of one that is not—which implies what they expect inflation to be over the term of the security, on average. (An increase in that premium implies that financial market participants expect inflation to be higher over that duration, for example.) Market-based measures benefit from the financial incentive that financial market participants have to devote resources to being as accurate as possible, although research shows that the relative prices of the securities used to compute the measures may be distorted by other factors, such as the relative illiquidity of inflation-indexed securities.

Survey-based measures collect information on the inflation expectations of different groups of people, such as consumers, professional forecasters, and firm managers. Prices are determined by the decisions made by consumers and firms. Two downsides of survey-based measures are that surveys of consumers or firms are not always available, and data collected from them may cover a relatively short historical period.

Measures of inflation expectations computed from surveys of consumers provide especially useful data that are available in the United States. When included alongside other macroeconomic indicators, measures of inflation expectations from surveys of consumers improve the fit of inflation forecasting equations more than measures from surveys of professional forecasters. Unfortunately, data measuring the inflation expectations of firms in the United States are lacking in both quality and availability. Survey measures of the inflation expectations of firms in other countries provide not only another possible source of useful data but also confirmation that the expectations of firms closely mirror those of consumers. In this paper, which examines the literature on the process of forming expectations, I rely heavily on empirical evidence gathered from surveys of either firm managers or consumers.

Inflation expectations are an important part of the Congressional Budget Office's inflation projections. CBO forecasts inflation using a bottom-up approach in which it first creates forecasts for many component price indexes and then aggregates them, using an accounting identity, into forecasts for the overall price indexes. To forecast the component price indexes, CBO uses various models, several of which depend, in part, on a top-down Phillips curve equation that determines the change in aggregate prices as a function of demand, supply, and expectations. Over time, excess demand and other cyclical factors fade in CBO's projections. Therefore, to determine the rate of inflation in the long run, CBO mainly uses long-run expected inflation, which stems mostly from monetary policy in the long run.

Understanding how expectations about inflation are formed is critical to CBO's analysis of fiscal policy. CBO's Phillips curve equation incorporates the assumption that inflation expectations are partially accelerationist and partially anchored, meaning that they are a weighted average of past inflation and a constant. The extent to which expectations are anchored—the weight the constant receives in CBO's equation—is estimated using historical data. Estimating any effects that a change in policy might have on that relative weighting relies critically on an understanding of the underlying nature of the expectations formation process.

In this paper, I begin by analyzing the dominant approach to modeling the formation of economic expectations. Sticky price models, perhaps most notably the model in Guillermo Calvo's seminal 1983 paper, and their offspring, the New Keynesian Phillips curve (NKPC), cemented the full information rational expectations (FIRE) theory's popularity among academic macroeconomists. That theory imposes some rather strong assumptions. Among them is the assumption that decisionmakers' forecasts are optimal, meaning that errors cannot be predicted using any information available when the forecast was made. Although the FIRE theory has been popular in macroeconomics literature since the 1970s and 1980s, it has been consistently rejected by empirical evidence. One such piece of evidence is the positive correlation between consensus ex-post inflation forecast errors and ex-ante revisions to those consensus forecasts, which shows that the errors are predictable. That rejection of the FIRE assumption implies that forecasters are not fully informed, are not completely rational, or perhaps both.

Next, I analyze models that relax the FIRE theory's full information assumption. Strong evidence supports the notion that consumers and firms misperceive or completely ignore information that could otherwise be useful in forecasting future inflation. When built into a theoretical model, that misperception or ignoration is known as an information friction. Information frictions were first proposed during the earliest days of the rational expectations theory's rise to prominence but have recently become a more common assumption. Many models with information frictions are either sticky information models or noisy information models. A sticky information model reflects the assumption that there is a fixed cost to updating one's information. A noisy information model reflects the assumption that one only observes a noisy signal of the true information state. In this paper, I show that empirical evidence in the literature strongly supports noisy information models. Furthermore, I show that there is evidence implying that more persistent inflation or a greater incentive to collect information about aggregate inflation would reduce information frictions.

I then analyze two models that relax the assumption of full rationality. There is evidence of significant bias in inflation forecasts at the individual level, supporting the notion that those forecasts may not be fully rational. The first model is referred to as adaptive expectations. It was one of the earliest theoretical models for the expectations-formation process and remains widely used today. First proposed by Milton Friedman in the 1960s, the theory is based on the assumption that inflation expectations *adapt* to equal currently observed, or recently observed, rates of inflation. Although some of the implications of the theory are not supported by empirical evidence from the survey data on consumer inflation expectations, adaptive expectations remains a common assumption in many inflation forecasting models—including, at least partially, CBO's.

The second model is known as diagnostic expectations. Upon receiving new information, economic decisionmakers with diagnostic expectations tend to place too much weight on the probability of future events that have become *relatively* more likely, even though those future events may still be unlikely in an *absolute* sense. That dichotomy can cause decisionmakers to overreact to new information, a finding supported by empirical evidence and visible in individual-level forecasts of many key macroeconomic variables, including inflation. Another implication of diagnostic expectations, also supported by empirical evidence, is that the overreaction will amplify trends that already exist in the inflation data. (By contrast, adaptive expectations implies that expectations will be uncorrelated with those properties of the data.)

I then explore the literature on whether monetary policymakers' communications could directly influence the expectations-formation process. Evidence shows that firm managers significantly update their beliefs when presented with new information, and those beliefs greatly affect their economic decisions. In addition, firms may devote substantially more resources toward collecting information on inflation when their incentive to do so is higher, such as during recessions. (That evidence is consistent with noisy information in the expectations-formation process.) Those times may also coincide with a period when monetary policymakers' more traditional tools are constrained by the effective lower bound (ELB) on nominal interest rates.

The rest of this paper is organized as follows: Section 2 examines the importance of inflation expectations and discusses which measures are most useful for furthering researchers' understanding of inflation dynamics. Section 3 describes CBO's current framework for modeling inflation expectations. Section 4 critically reviews the literature on models of the expectations-formation process, and section 5 explores whether monetary policymakers might be able to manage that process directly. Finally, section 6 provides conclusions and discusses some possible future applications of this work toward improving CBO's inflation forecasting methods.

Section 2: The Importance and Measurement of Inflation Expectations

Inflation expectations are beliefs that an individual or group of individuals have about the growth rate of prices in the future. Those individuals or groups who believe that the price level will grow more rapidly have higher inflation expectations, and those who believe that the price level will grow more slowly have lower inflation expectations.

Consumers' and firms' inflation expectations are particularly important, for several reasons. First, prices are determined by the intersection of supply and demand, which stems from the decisions of consumers and firms. Theory predicts that their inflation expectations affect those economic decisions greatly. Second, empirical evidence confirms that inflation expectations have an effect on actual inflation. Third, inflation forecasting models that include expectations generate more accurate projections than those that do not.

Among measures of inflation expectations, surveys of consumers or business managers provide especially useful data that are available in the United States. Survey-based measures of the inflation expectations of consumers in the United States have been shown to be empirically useful, but data on the inflation expectations of firms in the United States are noticeably lacking in both quality and availability. In other countries where those data are more widely available, firm managers' expectations more closely resemble those of consumers than of professional forecasters or financial market participants.

The Importance of Inflation Expectations

According to economic theory, inflation expectations play a crucial role in determining actual inflation. The expectations-augmented Phillips curve (Friedman, 1968) suggests that expectations are critical to the relationship between inflation and unemployment in the long run. In the canonical New Keynesian Phillips curve model (see Gali, 2008, for an example), a firm's expectations of future prices affect its pricing decisions in the current period.

Correlational empirical evidence supports the predicted importance of inflation expectations. Long-run trends in actual inflation closely mirror changes in long-run inflation expectations (Yellen, 2015). However, that observed correlation does not necessarily imply a causal relationship in which changes in long-run inflation expectations determine long-run trends in actual inflation. Expected inflation instead could respond to changes in other macroeconomic conditions—such as actual inflation or growth in costs per unit of labor—which themselves determine long-run trends in actual inflation (Rudd, 2021).

Some empirical evidence suggests a causal relationship between expected inflation and actual inflation. A positive shock to expected aggregate inflation is estimated to produce a significant increase in actual aggregate inflation (Clark and Davig, 2008). In addition, higher inflation expectations lead firms to raise their prices (Coibion, Gorodnichenko, and Ropele, 2020). That evidence points to the self-fulfilling nature of inflation expectations; higher expected inflation can lead to higher actual inflation.

Including inflation expectations can improve inflation forecasting models. Dozens of those models exist, and they vary greatly in their structure and complexity. Empirical research comparing the performance of a number of them shows that forecasting models that account for the slowly varying trend component of inflation dynamics outperform those that do not (Faust and Wright, 2013). Because movements in trend inflation are highly correlated with long-run inflation expectations, models that include long-run inflation expectations perform better than those that do not.

The Measurement of Inflation Expectations

Inflation expectations cannot be observed directly, so researchers measure them indirectly.

Market-based measures of inflation expectations calculate the expected inflation that is implied by the prices of different financial assets. One such measure, known as the breakeven inflation rate, is derived from the difference between the yield on a nominal Treasury security and the yield on a Treasury inflation-protected security (TIPS). The breakeven inflation rate can be used to estimate the inflation expectations of financial market participants over various time periods for which those two types of securities are available.

Even though market-based measures of inflation expectations may benefit from the fact that financial market participants have a financial incentive to devote resources to being as accurate as possible, those measures have flaws. Empirical evidence suggests that yields on TIPS exceed the true inflation-adjusted risk-free interest rate, primarily owing to the lack of liquidity in the market for TIPS relative to other Treasury securities (D'Amico, Kim, and Wei, 2018). That liquidity premium, along with several other smaller issues, clouds measures of inflation expectations estimated using TIPS.

Prices are determined by the decisions of consumers and firms, and surveys of those groups in the United States and abroad provide insight into their inflation expectations. One example of a U.S. survey is the Michigan Survey of Consumers, which has asked consumers about their short-term inflation expectations since 1978 and their longer-term inflation expectations since 1990.

Unfortunately, data on the inflation expectations of firms in the United States are severely limited. One survey conducted by the Federal Reserve Bank of Atlanta's inflation project provides an example. That survey, which estimates businesses' inflation expectations, has several limitations: It is based on data collected from within a single Federal Reserve district, it uses a relatively small sample, and it dates back only to October 2011. Another limitation of that survey is that it asks firm managers what they expect will happen to their own prices, not what they think will happen to prices more generally. A new survey of firms' inflation expectations, introduced this year, addresses some of those limitations—it is a nationwide survey that uses a larger sample and that asks questions about a specific aggregate price index—but it dates back only to 2018 (Candia, Coibion, and Gorodnichenko, 2021). That lack of data constrains researchers' ability to study the impact of firms' inflation expectations on the dynamics of actual inflation in the United States.

To counter those data limitations, researchers commonly use measures that are more readily available in the United States as a proxy for the inflation expectations of firms. Frequently, the inflation expectations of professional forecasters are used as a proxy for the inflation expectations of firms (for example, see Ball and Mazumder, 2011). Surveys of professional forecasters, such as the one conducted by the Federal Reserve Bank of Philadelphia, collect data on inflation expectations. The popularity of those measures is due, in part, to the availability of long time-series data in the United States and to the belief that the expectations of business managers might be informed by professional forecasters. Research has shown that proxies may be of limited empirical usefulness, though. Empirical estimation of inflation-forecasting equations (including both consumers' and professional forecasters' measures of expectations) shows that, conditional on the other terms in the equations, only the inflation expectations of consumers add statistically useful information (Coibion and Gorodnichenko, 2015a; Chen, 2019).

Limited access to data on the inflation expectations of firms in the United States may not be an impediment to economic researchers, however, because international empirical evidence suggests that firms' inflation expectations closely mirror those of consumers. In New Zealand, where surveys of both consumers and firms have been conducted, firms' inflation expectations are, on average, similar in level to consumers' inflation expectations. Another similarity is that firms' and consumers' inflation expectations are extremely dispersed—much more so than those of professional forecasters (Kumar and others, 2015).

Surveys of firms abroad and surveys of consumers in the United States are two of the most useful sources for data on inflation expectations.

Section 3: CBO's Current Framework for Inflation Expectations

CBO's inflation forecasts are informed by both bottom-up and top-down methodologies. Under the bottom-up approach, CBO first breaks down the overall price indexes into a mutually exclusive and exhaustive set of component indexes, including, for example, those for food, housing services, and motor vehicles and parts. CBO next creates a forecast for each component index and then aggregates those forecasts to arrive at a forecast for the overall price index. In contrast, under the top-down approach, CBO seeks to forecast the overall price index directly, using aggregate macroeconomic variables.

When producing a forecast, the agency places particular emphasis on the predictions from the top-down approach for the medium- and longer-term portions and the bottom-up approach for the shorter-term portion. One empirical observation that motivated CBO to develop its bottom-up methodology was that over time certain components of overall inflation had become nearly completely insensitive to business-cycle fluctuations, whereas other components had remained relatively sensitive. The bottom-up approach allows the agency to operationalize that empirical observation to improve the portion of its forecast that covers the time period of a typical recession or recovery. Beyond that time period, though, that observation becomes less relevant. The top-down approach allows CBO to create a forecast that is logically consistent with the agency's long-run assumptions pertaining to a relatively small handful of explanatory variables.

Currently, CBO's top-down approach consists of a single-equation Phillips curve model.

CBO's Phillips Curve Model

The Phillips curve refers to the relationship between slack in the labor market and inflation. The original Phillips curve plotted the relationship between the unemployment rate and wage inflation in the United Kingdom, but the Phillips curve now more commonly refers to the relationship between unemployment—or perhaps another measure of labor utilization—and price inflation in goods and services. The idea behind that relationship is that as the unemployment rate falls, firms must offer higher wages to entice prospective workers from a shrinking pool of candidates to fill their job openings. Firms' cost of production goes up as a result, and a portion of that increased cost gets passed on to consumers in the form of higher prices.

CBO's Phillips curve model incorporates explicit variables for demand, supply, and inflation expectations.[1] The forecasting equation has the following form:

[1] This is sometimes referred to as a triangle model; see Gordon (1977) and Gordon (1982).

$$
\begin{aligned}
\text{Core PCE Price Inflation}_t \\
&= \beta_1 \cdot \text{Unemployment Gap}_t + \beta_2 \cdot \text{Unemployment Gap}_{t-1} + \beta_3 \\
&\cdot \text{Unemployment Gap}_{t-2} + \beta_4 \cdot \text{Unemployment Gap}_{t-3} + \beta_5 \\
&\cdot \text{Unemployment Gap}_{t-4} + \beta_6 \cdot \text{Relative Price of Imported Goods}_t \\
&+ \text{Expected Inflation}_t + \varepsilon_t\ ,
\end{aligned}
$$

where ε_t is a white noise Gaussian error term. The terms involving the contemporaneous unemployment gap and the lagged values of the unemployment gap are used to model shifts in demand, the next term involving the relative price of imported goods models shifts in supply, and the final term models expected inflation. I briefly explain those three elements in the following subsections.[2]

Labor Market Slack. The first few variables in CBO's Phillips curve model account for the effect of fluctuations in aggregate demand on inflation. Macroeconomic theory suggests that when an increase in aggregate demand moves the economy out of equilibrium, the quantity of goods and services demanded exceeds the quantity of goods and services produced at the current price level. The resulting temporary shortage of goods and services in the economy causes prices to rise, as the limited quantities are sold to consumers who are willing to pay more. Theory suggests that the opposite is also true—namely, a decrease in aggregate demand causes a temporary surplus of goods and services in the economy, which causes prices to fall. Those increases and decreases in aggregate demand are referred to as business-cycle fluctuations.

To model the effects of those business-cycle fluctuations on aggregate inflation, CBO's Phillips curve equation includes a measure of labor market slack. During or immediately following a recession, when aggregate demand is low, firms cut back on their production of goods and services. As a result, they cut back on their use of labor, which creates significant slack in the labor market. Conversely, during an economic boom, firms increase their demand for labor, and the amount of slack in the labor market diminishes. The measure of labor market slack included in the equation is CBO's estimate of the unemployment gap, which is defined as the difference between the unemployment rate and the noncyclical rate of unemployment—the rate of unemployment that results from all sources except fluctuations in aggregate demand.

The strength of the correlation between the unemployment gap and inflation has lessened over the past 30 years. As a result, the coefficients on those terms in the forecasting equation are estimated to be very small. According to CBO's most recent estimates, the sum of the coefficients on the five terms in the equation above is -0.076. This means that an increase of 1 percentage point in the unemployment rate, holding the noncyclical rate constant, would boost core inflation by less than one-tenth of 1 percent.

[2] For a complete description of CBO's Phillips curve model, see Chen (2019).

Import Prices. The next variable in CBO's Phillips curve model accounts for one effect of changes in short-run aggregate supply on inflation. Macroeconomic theory suggests that a sudden increase in short-run aggregate supply causes a temporary surplus of goods and services in the economy, pushing prices down. Conversely, a sudden decrease in short-run aggregate supply causes a temporary shortage of goods and services, lifting prices up. A sudden increase or decrease in short-run aggregate supply might be caused by a variety of factors—favorable weather for production of agricultural goods, an unforeseen decrease in the production of commodities (such as oil), or a surge in imported goods, for example.

CBO's Phillips curve equation includes a measure of the relative price of imported goods. The measure is the inflation rate for nonpetroleum, non-capital-goods imports, net of core PCE (personal consumption expenditures) price inflation, weighted by the nominal expenditures on those imports as a fraction of nominal GDP (gross domestic product). That measure is positive when imported goods inflation is relatively high and negative when it is relatively low. When inflation in the prices of imported goods is very low, then imported goods become more competitive with domestically produced goods, which puts substantial downward pressure on domestic price inflation (according to empirical evidence).

Other shocks to short-run aggregate supply have had a smaller effect on core inflation (which excludes prices for food and energy) in recent years. Chen (2019) estimated a Phillips curve with a term for the relative price of imports as well as a term for the relative price of energy goods but determined that the pass-through of energy price shocks to core inflation is currently low. As a result, the term for the relative price of energy goods is excluded from the model. According to CBO's most recent estimates, the coefficient on the term for the relative price of imported goods in the equation above is 0.82. Under the assumption that nominal expenditures on nonpetroleum, non-capital-goods imports are 7 percent of nominal GDP (which was approximately the case at the end of 2019), an increase in import prices that is 5 percentage points slower than core PCE price inflation would decrease core PCE price inflation by about three-tenths of 1 percent.

Expected Inflation. Expected inflation can affect current inflation through a change in either supply or demand. Inflation expectations can be a self-fulfilling prophecy—that is, an increase in expected inflation can cause an increase in actual inflation. If consumers expect inflation to be higher in the future, they may realize that their purchasing power will decrease more rapidly, so they increase their demand for goods and services today, creating inflationary pressure in the economy. And if firms expect inflation to be higher in the future, they may sense that it will be more profitable for them to produce goods or services later, when the selling price of their output will be higher but some of their input prices, like wages, remain fixed. As a result, they may choose to decrease their production today, creating inflationary pressure in the economy.

Modeling Expected Inflation

CBO currently models expected inflation as a function of past inflation and a fixed anchor. In CBO's forecasting equation, expected inflation takes this form:

$$\text{Expected Inflation}_t = \gamma_1 \cdot \text{Core PCE Price Inflation}_{t-1} + \gamma_2 \cdot \text{Core PCE Price Inflation}_{t-2} + \gamma_3 \cdot \text{Core PCE Price Inflation}_{t-3} + \gamma_4 \cdot \text{Core PCE Price Inflation}_{t-4} + \gamma_5 \cdot \text{PCE Price Inflation Target}\,,$$

where it is assumed that

$$\gamma_1 + \gamma_2 + \gamma_3 + \gamma_4 + \gamma_5 = 1\,.$$

The first four terms change over time as realized inflation varies; the fifth term is fixed over time.

Models in which expected inflation is a function solely of past inflation (meaning that γ_5 equals zero in the equation above) and changes in past inflation flow through completely to actual inflation (meaning that the sum of γ_1 through γ_4 equals one) are known as "accelerationist" models. In such models, a 1 percentage-point increase in actual inflation would cause a 1 percent increase in expected inflation, further accelerating the growth in prices. The theory behind that dynamic for inflation expectations, and the development of the so-called accelerationist Phillips curve, are examined later in this paper.

CBO's current framework models inflation expectations as partially accelerationist. Expected inflation depends partly on realized values of inflation—meaning that γ_5 is greater than zero in the equation above. As a result, movements in actual inflation flow through to expected inflation but do not cause a one-for-one movement in expected inflation because the sum of γ_1 through γ_4 is less than 1 in the equation above. Chen (2019) provides evidence that the Phillips curve has shifted away from an accelerationist form.

CBO's current framework also models inflation expectations as partially anchored. Expected inflation depends partly on a value that the agency assumes does not change over time. That fixed anchor is set equal to 2 percent, the Federal Reserve's long-run goal for inflation. If inflation expectations did not depend at all on realizations of inflation, then they would be considered completely anchored. Chen (2019) suggests that consumers' inflation expectations are not completely anchored, however.

CBO uses historical data to estimate the relative weight that the accelerationist terms and the anchoring term receive in its Phillips curve model. According to the agency's most recent estimates, the coefficients on the accelerationist terms in the equation above sum to 0.41. So, a 1 percent increase in core PCE price inflation over four quarters would raise expected inflation by about four-tenths of 1 percent. By construction, the coefficient on the anchoring term is equal to one minus the sum of the coefficients on the accelerationist terms, which is 0.59. One interpretation of those coefficient estimates is that inflation expectations in CBO's model are roughly 40 percent accelerationist and 60 percent anchored.

Simulation Exercises

To further illustrate the mechanisms at work in CBO's framework for estimating inflation expectations, I examine the effects of an unexpected increase in inflation and the effects of an unexpected decrease in the unemployment rate.

An Unexpected Increase in Inflation. The first simulation exercise that I explore is an unexpected increase in realized inflation. I start from a baseline forecast for all of the variables that appear in the Phillips curve model, then I assume that the quarterly growth rate (annualized) of the core PCE price index in each of the previous four quarters was 1 percentage point higher. As shown in the top panel of Figure 1, the initial sharp rise in the solid black line (which represents the difference in actual inflation from the baseline) displays an unexpected increase in inflation. Because the increase was unexpected, the dashed blue line (which represents the difference in expected inflation from the baseline) does not rise contemporaneously. However, as actual inflation remains elevated, the blue line rises.

Owing to the partially accelerationist nature of the expected inflation framework, that increase in actual inflation flows through partly into expected inflation. As shown in the top panel of Figure 1, actual inflation that was 1 percentage point higher for four quarters raises expected inflation by about four-tenths of 1 percentage point, equal to the sum of the coefficients on the accelerationist terms in CBO's Phillips curve model. Expected inflation remains a couple of tenths of 1 percentage point higher in several subsequent quarters, and it continues to remain a few basis points higher for the next few years. (A basis point is one one-hundredth of a percentage point.)

That higher expected inflation results in higher actual inflation (assuming all other variables remain unchanged from the baseline), which continues to push prices further above their baseline level. As shown in the bottom panel of Figure 1, actual inflation that was 1 percentage point higher for four quarters raises prices 1 percent above their baseline level. As the resulting increase in expected inflation feeds back into actual inflation, prices continue to rise higher above their baseline level. Because inflation expectations are not fully accelerationist, however, expected inflation is not permanently increased by the unexpected increase in inflation in this simulation exercise. As a result, the black line in the bottom panel of Figure 1 (which represents the percentage deviation in prices from their baseline level) does not continue to rise, plateauing at about 1.4 percent.

Figure 1. A Simulation of How an Increase in Actual Inflation in the Past Would Affect Actual and Expected Inflation in Subsequent Quarters

Difference From the Baseline, Annualized Quarterly Percentage Change

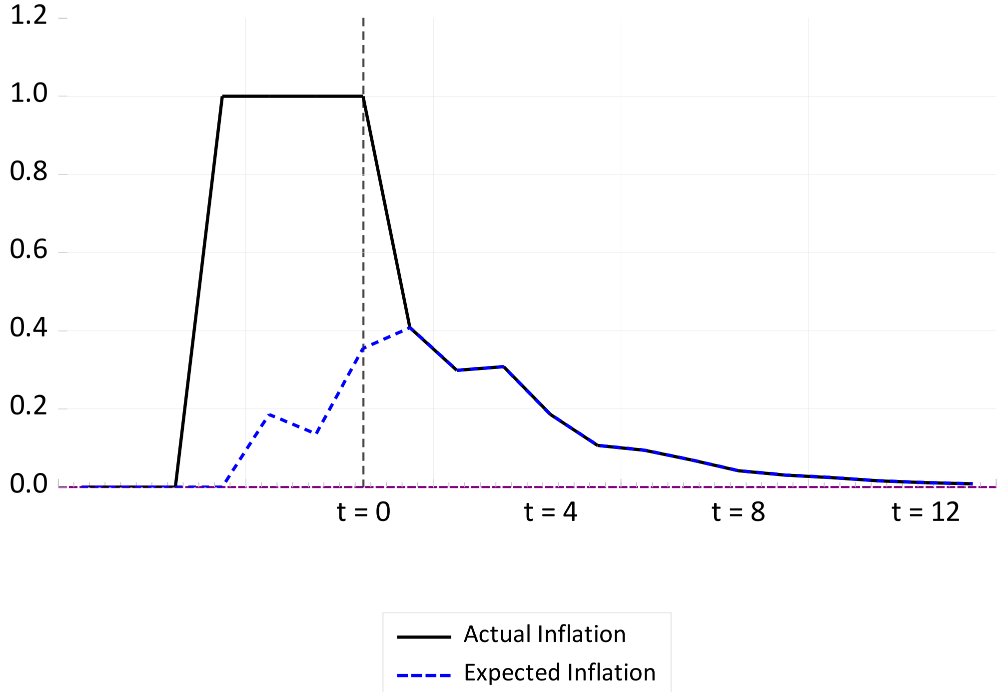

Percentage Difference From the Baseline

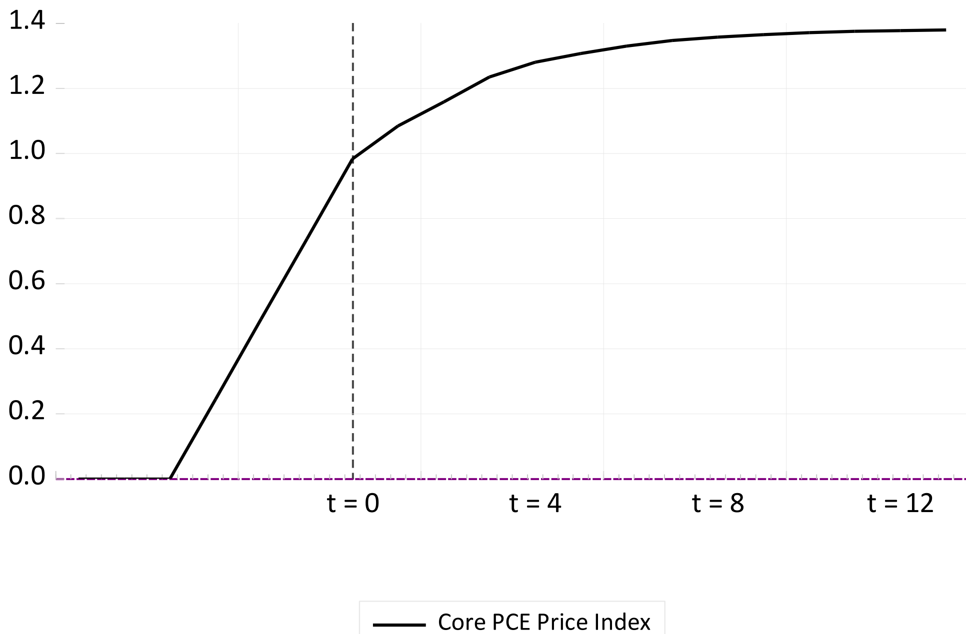

Data source: Author's calculations, using data from the Congressional Budget Office.

An Unexpected Decrease in the Unemployment Rate. The second simulation exercise I explore is an unexpected decrease in the unemployment rate. I start from a baseline forecast for all of the relevant variables, then I assume that four quarters ago the unemployment rate unexpectedly decreased to 1 percentage point below the baseline level. It remains 1 percentage point below the baseline level for an additional five quarters before rising to the baseline level over the following 10 quarters. Because I assume no change to the noncyclical rate of unemployment from its baseline forecast, that decrease in the unemployment rate results in a one-for-one change in the unemployment gap, the measure of labor market slack included in the Phillips curve model. As shown in the top panel of Figure 2, the fall, plateau, and subsequent rise in the black line displays that difference from the baseline forecast for the unemployment gap.

That decrease in labor market slack causes actual inflation to rise. In the middle panel of Figure 2, the solid black line (which represents the difference in actual inflation from its baseline forecast) immediately rises to about eight basis points, equal to the sum of the coefficients on the labor market slack terms in CBO's Phillips curve model. Over the following 10 quarters, the unemployment gap remains below its baseline level, placing upward pressure on prices and causing the black line to remain elevated.

Figure 2. A Simulation of How a Decrease in the Unemployment Rate Would Affect the Unemployment Gap, Actual Inflation, and Expected Inflation

Difference From the Baseline, Percent

Difference From the Baseline, Annualized Quarterly Percentage Change

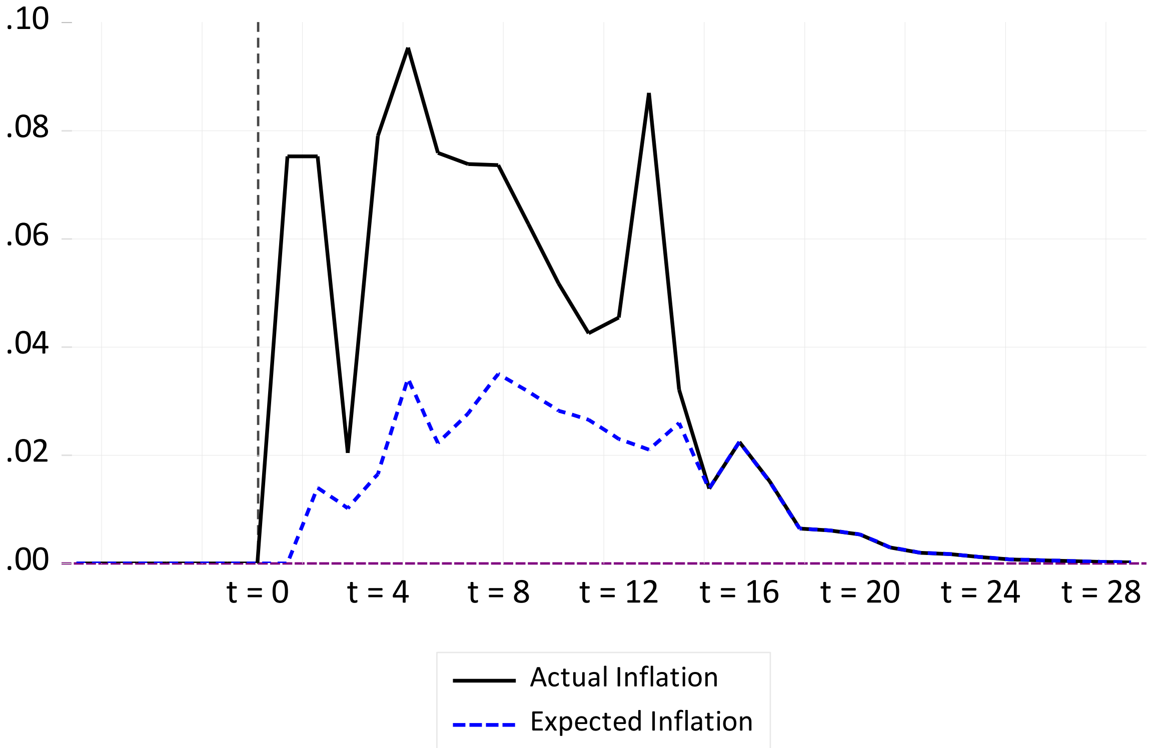

Percentage Difference From the Baseline

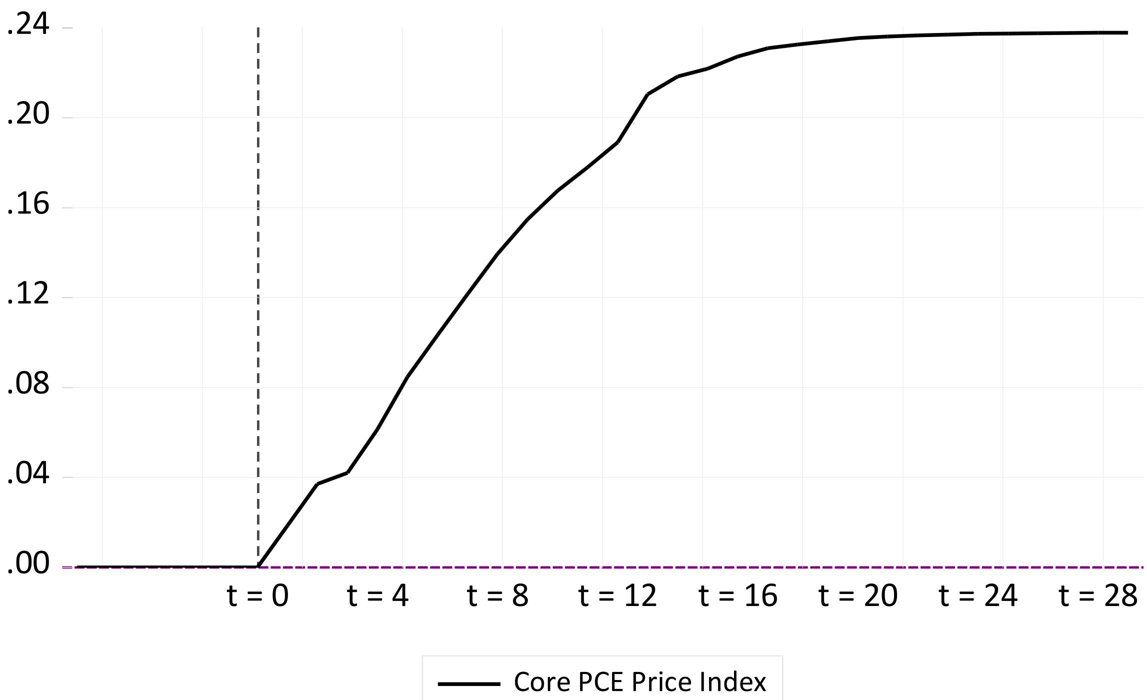

Data source: Author's calculations, using data from the Congressional Budget Office.

That increase in actual inflation causes expected inflation to rise in turn. In the middle panel of Figure 2, the dashed blue line (which represents the difference in expected inflation from its baseline forecast) subsequently rises to about three basis points, which is roughly 40 percent—approximately equal to the sum of the coefficients on the accelerationist terms in CBO's Phillips curve model—of the increase in the solid black line. The blue line remains elevated for several years before diminishing to zero.

Finally, the increase in expected inflation feeds back into actual inflation, which remains elevated even after the unemployment rate has returned to its baseline level. As shown in the bottom panel of Figure 2, the unexpected decrease in the unemployment rate causes prices to immediately rise above their baseline level by less than 0.1 percent. Prices continue to rise as inflation is boosted by decreased labor market slack and increased expected inflation, increasing by about 0.1 percent four quarters later. By the time the unemployment gap returns to its baseline forecast, prices have risen about 0.2 percent above their baseline level. Prices continue to rise thereafter because expected inflation remains elevated for several years. However, because expectations are not completely accelerationist, expected inflation returns to its baseline level, inflation is not permanently increased, and the black line plateaus at a level slightly above 0.2 percent.

Section 4: How Are Inflation Expectations Formed?

Empirical evidence suggests that consumers and firms form their expectations using simple rules and that their ability to process all the information available to form those expectations is limited. Those limits mean that many consumers and firms will misperceive or simply ignore relevant information. The simple rules that consumers and firms use to form their expectations result in significant biases—causing their expectations to shift too much or too little, respectively, as consumers and firms either overreact or underreact to new information.

Modeling elements that restrict the set of information that consumers or firms have access to when making their economic decisions are known as information frictions. Most models with information frictions can be classified as either sticky information models or noisy information models. Empirical evidence from surveys of consumers or firm managers supports the implications of noisy information but not sticky information. Both types of models imply that average forecasts underreact to new information, an implication supported by the data.

I explore two models in which consumers and firms form their expectations in a manner that is not fully optimal. Both adaptive expectations and diagnostic expectations can account for the fact that errors in individuals' inflation forecasts are correlated with certain information available when those forecasts were made. However, diagnostic expectations are also able to account for the fact that individuals' inflation expectations tend to exaggerate trends already present in the inflation data.

Both information frictions and forecasts that are less than fully optimal contradict the dominant approach in the literature on expectations, which I present first.

Full Information Rational Expectations

Although it imposes a strong set of assumptions, FIRE is the dominant approach to modeling the formation of expectations. Two key assumptions that are standard among FIRE models are the following:

- Economic decisionmakers perfectly perceive a complete set of information containing all the variables in the model and their entire history.
- Their forecasts for future variables are optimal, given all the information available. In other words, expectations do not differ systematically or predictably from the equilibrium predictions of the model itself.

Many economists believe that consumers and firms make forecasts using all publicly available information and that those forecasts might turn out to be wrong, but not systematically.

The dominance of FIRE sprung from the failure of previous modeling approaches. Before the 1980s, many economic models that were used to analyze policy changes failed to account for how consumers and firms would react. Those models therefore ignored the fact that this dynamic behavior would alter relationships estimated using only historical macroeconomic data.[3] To overcome that limitation, macroeconomic models began to be grounded in more microeconomic foundations. Throughout the 1970s and 1980s, many influential models were developed containing agents that optimized their choices according to their expectations about future economic conditions.[4]

The New Keynesian Phillips curve helped cement FIRE as the dominant theory of the formation of inflation expectations. In his seminal 1983 paper, Guillermo Calvo proposed a microfounded model with staggered price adjustments, known as a sticky-price model. In the model, firms can change their prices only in certain periods, which arrive randomly. When a firm has the opportunity to change its prices, it considers not only the price that will maximize profits today, but also the profit-maximizing price in each future period, until the next period during which it is allowed to change its price again. Thus, the firm's pricing decision relies critically on its expectations about future prices. In that type of model, after some algebraic manipulation, it can be shown that the overall inflation rate depends on expectations of future inflation and a measure of economic capacity utilization, a relationship known as the NKPC. Because it is derived from microeconomic foundations and has a clear economic interpretation, the NKPC remains the workhorse model for many inflation researchers today. And because the NKPC incorporates the

[3] This is commonly referred to as the Lucas critique; see Lucas (1976).

[4] See Lucas (1972) for one example; or see Stokey, Lucas, and Prescott (1989) for a textbook treatment.

assumption that inflation expectations are FIRE, the FIRE theory remains one of the most popular models of the expectations-formation process.

There are several implications of the NKPC and the FIRE theory. The first implication is that price setting should depend on firms' short-term inflation expectations—more specifically, on their forecast for the path of the price level until the next period when they can adjust their price, which is typically only a few months or quarters. The second implication is that the inflation rate would respond to shocks immediately. For example, an expansionary monetary policy shock would result in a large and immediate increase in the inflation rate. The third implication is that forecast errors would be completely unpredictable. If information was publicly available that would help decisionmakers improve their forecast, on average, they would use it.

Since the 1980s, a large body of empirical evidence has been gathered that rejects those implications, though. First, empirical evidence suggests that long-run trends in actual inflation closely mirror changes in long-run inflation expectations, not short-run inflation expectations (Yellen, 2015).[5] Second, empirical estimates of the response of the inflation rate to monetary policy shocks are small and slow (Christiano, Eichenbaum, and Evans, 1999). The response of the inflation rate to monetary policy shocks exhibits hump-shaped dynamics: It is very small at first, builds over time, and then declines again. Third, forecast errors are correlated with revisions to the previous forecast (Coibion and Gorodnichenko, 2015b). This means that forecast errors are, in fact, predictable.

Relaxing the Full Information Assumption

One possible reason for the empirical failure of the FIRE theory is that consumers and firms might misperceive or completely ignore information that could otherwise be useful when forecasting future inflation. Such information frictions were introduced during the early stages of the rational-expectations revolution—for example, Robert Lucas suggested that agents made decisions with less than perfect information in a 1977 paper—but they have become more common during the past 20 years. Many papers published since the early 2000s have specified models with information frictions; those papers have had varying degrees of success accounting for the behavior of aggregate inflation.

Many models containing information frictions can be classified as sticky information or noisy information. Sticky information models reflect the assumption that in each period a random fraction of firms obtains new information about the state of the economy (Mankiw and Reis, 2002). As a result, firms only compute their optimal price intermittently, but when they do, they

[5] That evidence may not directly contradict the sticky price assumption, however. Rudd (2021), for example, conceives of a model in which the inflation equation displays a dependence on long-run expected inflation, even though the underlying pricing equation depends on short-run expected inflation.

do so according to FIRE. The probability in each period that a firm does not obtain new information determines the amount of information friction in the model.

One might imagine that information spreads throughout the economy from those who have a large incentive to collect it (such as professional forecasters) to those who may not (the general public). Modeling the spread of information over time in that way using a model like the ones used to study the spread of an infectious disease results in a sticky information model (Carroll, 2003). That example provides one possible economic interpretation for the intermittent updating in a sticky information model. In contrast, noisy information models reflect the assumption that agents update their information in every period, *t*, but each individual, *i*, observes an imperfect signal, S_t^i, of the true information state:

$$S_t^i = X_t + w_t^i \text{, Eq. 1}$$

where X_t is the true unobserved information state and w_t^i are idiosyncratic white Gaussian noise error terms that are independent across agents (Woodford, 2003). As a result, in each period, agents update their estimate of the true unobserved information state by combining an old estimate with their new private signal. They use this equation:

$$F_t^i(X_t) = G_t \cdot S_t^i + (1 - G_t) \cdot F_{t-1}^i(X_t) \text{, Eq. 2}$$

where F_t^i is the individual's updated estimate and F_{t-1}^i is the individual's prior estimate made in the previous period using knowledge of the true economic system.[6] The updated estimate is a weighted average of the individual's signal of the true information state today and his or her previous estimate of the state today, with weights, G_t, determined by the information content of the individual's signal.[7] If the information content is low, meaning the variance of w^i is large, then G is small, and the individual puts very little weight on his or her signal. (In the extreme case in which G equals zero, long-run expectations are constant.) If the information content is high, then G is large, and he or she does the opposite. Therefore, the weight that an individual puts on his or her previous estimate of the state today, $(1 - G_t)$, measures the amount of information friction in the model. One popular economic interpretation for the noise in agents' observations of the true information state in a noisy information model is that it results from the fact that economic decisionmakers possess a limited capacity for processing information (Sims, 2003).[8]

---

[6] This equation is a version of the update step in the well-known Kalman filter algorithm. See Harvey (1989).

[7] G is commonly referred to as the Kalman gain.

[8] Rudd (2021) points out that households still appear to pay attention to aggregate inflation in other contexts, though. For example, consumer sentiment and actual inflation are strongly correlated. That is a valid criticism of a strict rational inattention model but perhaps not of the noisy information framework more generally.

Available measures of inflation expectations provide empirical evidence that does not seem to support several implications of sticky information models.

- The assumptions in a static sticky information model—in which the probability in each period that an agent obtains new information does not change over time—imply that the highest proportion of agents at any time have obtained new information in the current period. However, according to survey measures, most individuals tend to update their information somewhat infrequently (Branch, 2007).
- Sticky information models also imply that the dispersion of information (meaning the amount of disagreement) across agents increases following an economic shock. (In noisy information models, the amount of disagreement across agents is completely independent of those shocks.) The empirical evidence rejects the hypothesis that the amount of disagreement across agents increases after a shock (Coibion and Gorodnichenko, 2012).

In contrast, several implications of noisy information models do have empirical support.

- Noisy information models imply that the less persistent—in other words, the more *noisy*—a variable is, the higher that variable's amount of information friction is. Evidence shows that economic decisionmakers' expectations for less persistent macroeconomic variables display more information rigidity (Coibion and Gorodnichenko, 2015b).
- Noisy information models also imply that the more incentive agents have to reduce the noise in their observations of a variable, the lower that variable's amount of information friction is. There is empirical evidence that firms with fewer competitors or firms that sell more of their products abroad (meaning firms with less incentive to collect information about factors influencing domestic inflation) display more information rigidity in their inflation expectations (Kumar and others, 2015).

Both sticky and noisy information models imply that average errors made in past forecasts (ex-post, which can only be calculated after the true value is realized in a given period) are positively correlated with average revisions made to subsequent forecasts (ex-ante, which can only be calculated before the next realization). Consider the following regression equation:

$$Y_t = \beta_0 + \beta_1 X_t + u_t \text{, Eq. 3}$$

where Y_t is the ex-post forecast error (the realized inflation rate at the end of period *t* minus forecasted inflation for period *t*), X_t is the ex-ante forecast revision (the forecasted inflation for period *t* minus forecasted inflation for period *t-1*), and u_t is a white Gaussian noise error term. FIRE theory would imply that β_1 is zero, because forecast errors cannot be predicted using any information that was available at the beginning of the period. However, both sticky information models and noisy information models imply, when average forecast errors are regressed on average forecast revisions, that β_1 is positive and that the magnitude of β_1 depends only on the amount of information friction in the model (Coibion and Gorodnichenko, 2015a).

That fact implies that agents, on average, are underreacting to news; their average expectations are not getting updated sufficiently. For example, imagine that the average expected inflation rate is currently 10 percent. Then some news arrives that, had all the agents' inflation expectations been formed according to FIRE theory, would have lowered agents' average inflation expectations to 2 percent. However, the average expected inflation rate underreacts, falling to only 8 percent. In that case, the average ex-ante revision (X_t in equation 3) would be negative 2 percentage points. In the next period, the inflation rate is 3 percent, close to the FIRE forecast of 2 percent. That represents an average ex-post forecast error (Y_t in equation 3) of negative 5 percentage points. If average expectations continued to underreact, an econometrician would begin to observe a positive correlation ($\beta_1 > 0$ in equation 3) between the average ex-ante revisions and the average ex-post errors.

Empirical evidence finds that average forecasts do, in fact, underreact to news. If equation 3 was estimated using the average forecasts of inflation from surveys of professional forecasters, then β_1 would be statistically significantly larger than zero, and the amount of information friction implied by the estimate of β_1 would be economically meaningful (Coibion and Gorodnichenko, 2015a). That pattern is true not only for inflation, but also for consensus forecasts of many macroeconomic variables (Bordalo and others, 2018).

It is important to note that in both sticky information and noisy information models, individuals within the model are assumed to be fully rational; that is, they do not under- or overreact to any new public or private information that they receive. There is no correlation between the ex-post errors resulting from an individual's forecasts and the ex-ante revisions made by that same individual. The underreaction detected in the average forecasts thus results solely from the information structure and the process of aggregating the individual forecasts.

To understand that fact, consider a noisy information model like the one described earlier. The key assumption that leads to this result is that all signals are private; an individual does not observe anyone else's signal. When estimating the true unobserved information state, each individual places the weight on the private signal that he or she observes, G_t in equation 2, that is optimal because the individuals are assumed to be fully rational. That assumption ensures that there will be neither overreaction nor underreaction in the long run. However, an omnipotent econometrician—if he or she was able to observe everyone's signals—would benefit greatly from taking an average of all the signals when estimating the true information state. As the number of individuals in the model approached infinity, the sample average would approach the true value of the information state. Putting the same weight on that *average* signal when estimating the true information state, which is what the aggregation process in those models effectively does, is not optimal, though. It results in the underreaction observed in those models' average forecasts.

Consequently, learning more about how individuals form their inflation expectations requires analysis of individual-level data on those expectations. Even though the empirical evidence presented in this section clearly suggests that there are important information frictions in the macroeconomy, it does not confirm or reject the rationality of individual forecasts. Data on the inflation expectations of specific individuals are particularly limited, but recent studies have furthered researchers' understanding of how individuals form their expectations.

Relaxing the Full Rationality Assumption

There are many approaches, some of them overlapping, to relaxing the assumption that decisionmakers are fully rational. Some models assume different types of bounded rationality. In those models, decisionmakers are still "rational" in that their forecasts for future variables are optimal, but that rationality is "bounded" by the fact that they are optimizing according to an incorrect (usually simpler) understanding of how the economy in the model works. Other models assume various types of learning. When forecasting future uncertain states, decisionmakers do not use the true model or an incorrect version, they instead are constantly updating their understanding of how each variable of interest evolves as they "learn" from its latest observed value. Those differing approaches often have similar implications.

A full review of the literature on departing from full rationality is outside the scope of this paper. Instead, I analyze two models of the expectations-formation process that successfully account for several of the features of the data on inflation expectations.

Adaptive Expectations. This theory of how inflation expectations are formed was first developed in the 1960s and popularized in the 1970s. It incorporates the assumption that expected inflation is equal to the previous period's realized inflation rate, or an average of the past several periods' realized inflation rates. For example, if consumers or firms believed that inflation was going to be 2 percent and inflation turned out to be 3 percent (or the average inflation over the past several periods turned out to be 3 percent), they would update their beliefs to expect inflation to be 3 percent going forward. Consumers and firms are assumed to be purely backward-looking; they form their expectations for inflation solely on the basis of that variable's past performance.

Some early inflation models ignored the process of expectations formation altogether, implying a potentially appealing trade-off for monetary policymakers. The earliest models of the Phillips curve reflected the assumption that inflation expectations were constant (Phillips, 1958). That assumption implied an exploitable long-run trade-off between inflation and unemployment. If policymakers desired a lower unemployment rate, they need only tolerate a higher long-run rate of inflation. In his 1968 presidential address to the American Economic Association, Milton Friedman suggested that exploiting that trade-off would instead require policymakers to tolerate ever-increasing rates of inflation. He suggested—laying the foundation for the adaptive-expectations theory—that this was because inflation expectations would not remain constant but would *adapt* to changes in realized inflation.

Although the theory of adaptive expectations does not rule out a short-run trade-off, it does imply that an exploitable long-run trade-off between inflation and unemployment is not possible. If monetary policymakers pursued a higher rate of inflation in order to lower the unemployment rate, eventually consumers and firms would adapt and begin to expect that new higher rate of inflation to persist.

Inflation's behavior in the 1970s supported Friedman's adaptive-expectations theory (Perry, 1970; Perry, 1978). As policy actions designed to maintain low unemployment caused the rate of inflation to rise ever higher throughout the decade, no exploitable long-run trade-off seemed possible. That impasse led to the development of a new Phillips curve model, augmented by adaptive-inflation expectations, which became known as the accelerationist Phillips curve.

That extrapolative, purely backward-looking theory of expectations formation has several other implications. First, by incorporating the assumption that expectations depend only on realizations of aggregate inflation, the theory implies that the economic agent making the inflation forecast uses no other information to forecast future inflation. Second, it implies that forecast errors might be predictable if any other macroeconomic variables were useful for forecasting future inflation. For example, if the unemployment rate helped explain current inflation and the unemployment rate was rather persistent, then the unemployment rate would help forecast future inflation, and the forecast errors would probably be correlated with the unemployment rate. In other words, the forecasts would be wrong in a systematic way.

The adaptive-expectations theory remains very influential among economic forecasters and policymakers even though some of its implications are unsupported by the data on inflation expectations. For example, empirical data do not support the implication that expectations are not influenced by anything other than past realizations of aggregate inflation. Instead, consumers' inflation expectations are strongly influenced by changes in the prices of products they purchase more frequently, and they weigh those price changes more heavily than they are weighted in the expenditure-weighted aggregate inflation measures (D'Acunto and others, 2019). Nevertheless, the adaptive-expectations framework remains popular for policy analysis or inflation forecasting, most likely because it is simple to operationalize and accounts for some features of the data on inflation expectations. CBO's current inflation forecasting method employs a partially accelerationist Phillips curve equation.

Diagnostic Expectations. This approach to modeling is based on a heuristic principal called representativeness. (Just like adaptive-expectations theory, this theory also models departures from rationality.) At roughly the same time that macroeconomists were embracing the rational-expectations hypothesis, psychologists were amassing evidence that rejected it. In their seminal 1974 paper, Amos Tversky and Daniel Kahneman documented several ways in which our expectations about uncertain events are biased. They provided evidence suggesting that people rely on simple rules, called heuristics, when creating their beliefs about uncertain events, and as a result their beliefs suffer from biases. One such heuristic states that decisionmakers tend to overweight the likelihood of "representative" events (Tversky and Kahneman, 1974). For example, when respondents were asked to assess the likelihood that an individual who had been described as quiet and bookish had one of several occupations, they often placed too much weight on the likelihood that the individual was a librarian because those characteristics are representative of that group. In reality, the likelihood that the individual was a librarian was generally much lower than the respondent imagined.

Diagnostic expectations formalize that heuristic by assuming that the representativeness of a future state is equal to the proportional increase in its probability after observing some new information. In laboratory experiments in which decisionmakers were given a piece of news, they tended to overweight future states that had become *relatively* more likely given that news (Gennaioli and Shleifer, 2010). For example, consider a distribution of occupations that has very few librarians and many lawyers. If a decisionmaker is given new information that the individual whose occupation is in question is shy and quiet, then the decisionmaker will tend to overweight occupations that do not require public speaking, because those types of occupations are more "representative" of the group of occupations that a person who is shy and quiet might have. So occupations such as librarian will have become relatively more likely given the news, although they are still not very likely given the distribution of occupations. Diagnostic expectations were developed to account for a variety of judgmental biases that are well documented in laboratory experiments (Bordalo, Gennaioli, and Shleifer, 2018).

Empirical evidence from analyses of individual-level forecasts of aggregate inflation supports the implications of diagnostic expectations. Models whose decisionmakers have diagnostic expectations imply that their forecasts will exaggerate true patterns that exist in the data. In contrast, models with purely extrapolative forecasts, like the ones made by agents with adaptive expectations, bear no relationship to those patterns. Individual-level survey data show that the degree to which a decisionmaker updates his or her forecast in response to new information depends on the amount of persistence displayed by the variable in question (Bordalo and others, 2018). Decisionmakers update their forecasts less in response to new information when the variable in question is less persistent, as aggregate inflation has become over the past 20 years.

Models with diagnostic expectations imply that individual forecasters overreact to news, so their ex-post forecast errors are negatively correlated with their ex-ante forecast revisions. In other words, a regression of the individual-level forecast errors, pooled for all the individuals, on the individual-level forecast revisions (of the form described in equation 3) would yield an estimate of β_1 that was negative. For example, consider an individual whose prior forecast for inflation was 2 percent but who now observes that inflation is 3 percent. Influenced by the representativeness heuristic, the individual will overweight high-inflation states that have become *relatively* more likely (in other words, more likely than they were before) but remain unlikely, and thus will positively revise his or her forecast too greatly. That action makes it more likely that the individual's next forecast will overshoot actual observed inflation, resulting in a negative forecast error. Over time, and pooled for many decisionmakers, that tendency will cause an econometrician to observe the negative relationship described above. A regression of that type, using data from the forecasts of individual professional forecasters, yields the predicted negative correlation (Bordalo and others, 2018), providing further evidence in support of models with diagnostic expectations.

Using certain parameters, a model with both diagnostic expectations and noisy information can imply an overreaction to news of individual forecasts and an underreaction to news of average forecasts. The previous paragraph concluded that individuals with diagnostic expectations will overreact to news. If those individuals also have a limited ability to process information, then those diagnostic decisionmakers can be modeled as if they observe a noisy private signal of the true information state, as was done in the noisy information models examined earlier. The individuals in that model setting will produce an estimate of the current information state using a slightly modified version of equation 2,

$$F_t^{i,\alpha}(X_t) = \{1 + \alpha\} \cdot G_t \cdot S_t^i + (1 - \{1 + \alpha\} \cdot G_t) \cdot F_{t-1}^{i,\alpha}(X_t), \text{Eq. 4}$$

where $\alpha > 0$ is assumed for agents with diagnostic expectations, implying an overreaction to news. For certain parameter values, when α is not too large and G is small enough (meaning that the amount of information friction is large), this model implies that the correlation between average ex-post forecast errors and average ex-ante forecast revisions will be positive, and the correlation between individual ex-post forecast errors and individual ex-ante forecast revisions will be negative (Bordalo and others, 2018). That result accounts simultaneously for two important features documented in the data on inflation expectations. The intuition is that if information is very noisy, individuals will discount their news heavily, and G will be very small. Even if an individual's expectations are influenced by the representativeness heuristic, and as a result he or she discounts the news a little less, $\{1 + \alpha\} \cdot G_t$ may still be discounting the average news too heavily, resulting in an underreaction to that news.

This review of the literature suggests that economically significant information frictions, as well as significant biases, are part of the expectations-formation process. Decisionmakers overreact to new information, possibly because it is representative of something to them. Significant information frictions most likely arise from limits on decisionmakers' ability to process information; those limits cause average forecasts to underreact to new information. The evidence supports noisy information as opposed to sticky information in the expectations-formation process, which implies that if agents' incentive to collect information increased, the observed level of information friction would decrease. And if the observed level of information friction decreased enough at the same time that individual decisionmakers continued to overreact to new information, then average inflation expectations might begin to overreact to news as well.

Section 5: Can Inflation Expectations Be Managed?

The evidence on whether monetary policymakers can manage inflation expectations is mixed. Some evidence shows that firms will update their inflation expectations when presented with pertinent information, which will affect their pricing decisions accordingly. Other evidence suggests that conveying that pertinent information to firms is extremely difficult, however.

Traditionally, one of the main ways in which the Federal Reserve attempts to achieve its goals of low and stable inflation and maximum employment is by influencing nominal interest rates. Real interest rates—which are equal to nominal interest rates minus expected inflation—affect the macroeconomy by altering the incentives to consume or save income and by altering the cost of borrowing to finance investment. Most of the monetary policy tools that the Federal Reserve currently uses to stimulate the economy, including nontraditional tools such as quantitative easing (QE, or the purchase of long-term securities on the open market) and forward guidance (communications about the likely future course of monetary policy), largely aim to lower nominal interest rates or to reduce expected nominal interest rates in the future.

Could other tools be developed that would lower real interest rates by raising inflation expectations directly? The answer to that question would depend on two factors:

- Whether the information that monetary policymakers might want to communicate to economic decisionmakers would affect their decisions in any meaningful way, and
- Whether policymakers could overcome significant information frictions to ensure that the information was received in the first place.

Does New Information Affect Firms' Pricing Decisions?

When presented with new information that is relevant to their pricing decisions, firm managers will significantly alter their beliefs. In a causal study in which a random sample of firm managers were presented with some new information about recent inflation, those managers immediately and systematically updated their beliefs about future inflation (Coibion, Gorodnichenko, and Ropele, 2020). Also, the managers who were more uncertain about their beliefs before receiving the new information made larger revisions to their beliefs.

Firm managers' inflation expectations significantly affect their economic decisions. At least one study has found that an increase in a firm's expectations of future inflation caused that firm to raise its prices (Coibion, Gorodnichenko, and Ropele, 2020). More importantly for monetary policymakers, the increase in the firm's inflation expectations caused it to increase its employment level. Those findings suggest that the receipt of pertinent new information had an economically meaningful impact on firm managers' decisions. For that reason, there would appear to be scope for monetary policymakers to use their communications with the public to affect the macroeconomy by influencing inflation expectations.

Can Monetary Policymakers Overcome Information Rigidity?

Whether monetary policymakers could communicate new information to consumers and firms and influence their economic behavior remains unresolved as the current evidence is mixed. On the one hand, there is compelling evidence that consumers and firms are not paying much attention to information that would help them form more accurate expectations about inflation. In part, that is because firms do not view aggregate inflation as important to their business decisions. In surveys, firm managers say they devote few resources to tracking current inflation (Coibion, Gorodnichenko, and Kumar, 2018). And when firms are presented with new information relevant to their pricing decisions, that information's effect on their inflation expectations is transitory and mostly gone after only six months (Coibion, Gorodnichenko, and Kumar, 2018). There is also evidence, in line with the implications of a noisy information model, that low and stable recent inflation reduces a firm manager's incentive to collect information about inflation (Coibion and others, 2018). In addition, firms display little knowledge of monetary policies that have an impact on inflation. Surveys of firm managers reveal that they lack information regarding their central bank's policy objectives (Kumar and others, 2015). Firm managers are also largely unaware of their central bank's policy actions (Coibion, Gorodnichenko, and Kumar, 2018). In total, surveys of firm managers reveal that firms devote few resources to tracking aggregate inflation, supporting the conclusion that significant information frictions exist.

On the other hand, there is evidence that at a time when it could be critical for monetary policymakers to overcome those significant information frictions is precisely when decisionmakers would be likely to devote more resources to collecting information about inflation. The amount of effort that firm managers devote to acquiring new information about inflation depends on the state of the economy. Surveys of firm managers reveal that they put more effort into acquiring that information when they receive bad news about the economy (Coibion, Gorodnichenko, and Kumar, 2018). That finding is consistent with the implication from noisy information models that when a firm's incentive to collect information is higher, the observed amount of information friction will be lower. That fact could be critical for monetary policymakers, who at those times are more likely to be trying to communicate information to decisionmakers in order to affect their inflation expectations.

Some evidence shows that the Federal Reserve's previous attempts to communicate with greater transparency information about its policy objectives have had an effect on consumers and firms. One measure of uncertainty in consumers' inflation expectations, calculated using responses to the Michigan Survey of Consumers, declined from 2012 to 2016 (Binder and Verbrugge, 2016). In that survey, the inflation-expectations questions often elicit responses that are unrealistically high and tend to fall disproportionately on even or round numbers. That measure of uncertainty reflects the assumption that survey respondents who are particularly uncertain about future inflation tend to respond by giving integer values that are a multiple of five. Plotting a histogram of the individual responses to the survey's inflation-expectations questions reveals large numbers of responses for zero, 5, 10, or even 15 percent inflation, and the number of responses for those values is much larger than for adjacent values. Over the 2012–2016 period, the number of those responses decreased, whereas the number of responses for either 1 or 2 percent inflation increased. That decline in uncertainty roughly coincided with the Federal Reserve's announcement of its explicit long-run 2 percent goal for inflation.

Section 6: Conclusion

Empirical evidence from measures of inflation expectations computed using surveys of consumers or firm managers conducted before the pandemic reveals several important properties of the expectations-formation process. To begin with, strong evidence supports the existence of economically meaningful information frictions in that process. The structure of those information frictions is such that realizing more-persistent inflation or having more incentive to collect information on inflation would reduce them. That incentive to collect information on inflation is especially strong today, in light of inflation's currently high levels. Evidence also confirms significant biases in individual-level inflation forecasts, which possibly amplify trends already present in the inflation data. That fact could cause inflation expectations to over-react to news during the current period of high inflation.

According to those surveys, firm managers were not paying much attention to aggregate inflation dynamics, nor were they well informed on the objectives or actions of monetary policymakers before the pandemic. Firms may devote substantially more resources toward collecting information on inflation during recessions, though, a finding that is consistent with the evidence on the structure of information frictions in the expectations-formation process. That greater incentive to collect information during economic downturns may coincide with periods when monetary policymakers are constrained by the effective lower bound on nominal interest rates.

Those findings could guide future changes to CBO's inflation forecasting models.

- **Assume a different anchor value.** The agency's partially accelerationist, partially anchored Phillips curve model implies that average inflation expectations underreact to news, which is consistent with the findings in this paper. Average expectations are currently partially anchored at 2 percent, the Federal Reserve's long-run goal for inflation. However, evidence shows that firms are mostly unaware of the central bank's objective. CBO could consider other alternatives for the value of the partial anchor in its model.
- **Vary the degree of anchoring.** The degree to which average expectations are accelerationist or anchored is fixed throughout the 10-year projection period in CBO's current model. Given the evidence that the degree of information friction depends on the state of the economy, that

parameter might be better modeled as time varying. In the current period of more-persistent inflation, for example, it may be useful to assume that the degree to which expectations are anchored is now lower than its average over the past two decades.

- **Include survey measures.** CBO's models make use of purely statistical estimates of expectations. To provide the fullest information possible, CBO could include survey measures of inflation expectations in those models.

References

Laurence Ball and Sandeep Mazumder, “Inflation Dynamics and the Great Recession,” *Brookings Papers on Economic Activity*, vol. 42, no. 1 (2011), pp. 337–405.

Carola Conces Binder and Randal Verbrugge, “Digging Into the Downward Trend in Consumer Inflation Expectations,” *Federal Reserve Bank of Cleveland Economic Commentary* (November 2016).

Pedro Bordalo, Nicola Gennaioli, and Andrei Shleifer, “Diagnostic Expectations and Credit Cycles,” *Journal of Finance*, vol. 73, no. 1 (2018), pp. 199–227.

Pedro Bordalo and others, *Over-Reaction in Macroeconomic Expectations,* NBER Working Paper 24932 (National Bureau of Economic Research, 2018).

William A. Branch, “Sticky Information and Model Uncertainty in Survey Data on Inflation Expectations,” *Journal of Economic Dynamics and Control*, vol. 31, no. 1 (2007), pp. 245–276.

Guillermo A. Calvo, “Staggered Prices in a Utility-Maximizing Framework,” *Journal of Monetary Economics*, vol. 12, no. 3 (1983), pp. 383–398.

Bernardo Candia, Olivier Coibion, and Yuriy Gorodnichenko, *The Inflation Expectations of U.S. Firms: Evidence From a New Survey,* Working Paper 28836 (National Bureau of Economic Research, 2021).

Christopher D. Carroll, “Macroeconomic Expectations of Households and Professional Forecasters,” *Quarterly Journal of Economics*, vol. 118, no. 1 (2003), pp. 269–298.

Yiqun Gloria Chen, *Inflation, Inflation Expectations, and the Phillips Curve,* Working Paper 2019-07 (Congressional Budget Office, August 2019), www.cbo.gov/publication/55501.

Lawrence J. Christiano, Martin Eichenbaum, and Charles L. Evans, “Monetary Policy Shocks: What Have We Learned and To What End?” *Handbook of Macroeconomics,* vol. 1, part A (1999), pp. 65–148.

Todd Clark and Troy Davig, *An Empirical Assessment of the Relationships Among Inflation and Short- and Long-Term Expectations,* Working Paper 08-05 (Federal Reserve Bank of Kansas City, 2008).

Olivier Coibion and Yuriy Gorodnichenko, “What Can Survey Forecasts Tell Us About Information Rigidities?” *Journal of Political Economy,* vol. 120, no. 1 (2012), pp. 116–159.

Olivier Coibion and Yuriy Gorodnichenko, “Is the Phillips Curve Alive and Well After All? Inflation Expectations and the Missing Disinflation,” *American Economic Journal: Macroeconomics,* vol. 7, no. 1 (2015a), pp. 197–232.

Olivier Coibion and Yuriy Gorodnichenko, "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *American Economic Review,* vol. 105, no. 8 (2015b), pp. 2644–2678.

Olivier Coibion, Yuriy Gorodnichenko, and Saten Kumar, "How Do Firms Form Their Expectations? New Survey Evidence," *American Economic Review,* vol. 108 (2018), pp. 2671–2713.

Olivier Coibion, Yuriy Gorodnichenko, and Tiziano Ropele, "Inflation Expectations and Firm Decisions: New Causal Evidence," *Quarterly Journal of Economics,* vol. 135, no. 1 (2020), pp. 165–219.

Olivier Coibion and others, *Inflation Expectations as a Policy Tool,* NBER Working Paper 24788 (National Bureau of Economic Research, 2018).

Francesco D'Acunto and others, *Exposure to Daily Price Changes and Inflation Expectations,* CESifo Working Paper 7798 (2019).

Stefania D'Amico, Don H. Kim, and Min Wei, "Tips From TIPS: The Informational Content of Treasury Inflation Protected Security Prices," *Journal of Financial and Quantitative Analysis,* vol. 53, no. 1 (2018), pp. 395–436.

Jon Faust and Jonathan Wright, "Forecasting Inflation," in Graham Elliott and Allan Timmermann, eds., *Handbook of Economic Forecasting* (Elsevier, 2013), pp. 2–56.

Milton Friedman, "The Role of Monetary Policy," *American Economic Review,* vol. 58, no. 1 (1968), pp. 1–17.

Jordi Gali, *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework* (Princeton University Press, 2008).

Nicola Gennaioli and Andrei Shleifer, "What Comes to Mind," *Quarterly Journal of Economics,* vol. 125, no. 4 (2010), pp. 1399–1433.

Robert J. Gordon, "Can the Inflation of the 1970s Be Explained?" *Brookings Papers on Economic Activity* (1:1977), pp. 253–277.

Robert J. Gordon, "Inflation, Flexible Exchange Rates, and the Natural Rate of Unemployment," in Martin N. Baily, ed., *Workers, Jobs, and Inflation* (Brookings, 1982).

Andrew C. Harvey, *Forecasting, Structural Time Series Models, and the Kalman Filter* (Cambridge University Press, 1989).

Saten Kumar and others, “Inflation Targeting Does Not Anchor Inflation Expectations: Evidence From Firms in New Zealand,” *Brookings Papers on Economic Activity* (Fall 2015), pp. 151–225.

Robert E. Lucas, “Expectations and the Neutrality of Money,” *Journal of Economic Theory*, vol. 4, no. 2 (1972), pp. 103–124.

Robert E. Lucas, “Econometric Policy Evaluation: A Critique,” *Carnegie-Rochester Conference Series on Public Policy*, vol. 1, no 1 (1976), pp. 19–46.

Robert E. Lucas, “Understanding Business Cycles,” *Carnegie-Rochester Conference Series on Public Policy*, vol. 5, no. 1 (1977), pp. 7–29.

Gregory N. Mankiw and Ricardo Reis, “Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve,” *Quarterly Journal of Economics*, vol. 117, no. 4 (2002), pp. 1295–1328.

George L. Perry, “Changing Labor Markets and Inflation,” *Brookings Papers on Economic Activity*, vol. 1, no. 3 (1970), pp. 466–473.

George L. Perry, “Slowing the Wage-Price Spiral: The Macroeconomic View,” *Brookings Papers on Economic Activity*, vol. 9, no. 2 (1978), pp. 259–299.

A. W. Phillips, “The Relation Between Unemployment and the Rate of Change of Money Wage Rates in the United Kingdom, 1861–1957,” *Economica*, vol. 25, no. 100 (1958), pp. 283–299.

Jeremy B. Rudd, “Why Do We Think That Inflation Expectations Matter for Inflation? (and Should We?),” *Finance and Economics Discussion Series 2021–062* (Washington: Board of Governors of the Federal Reserve System).

Christopher A. Sims, “Implications of Rational Inattention,” *Journal of Monetary Economics*, vol. 50, no. 3 (2003), pp. 665–690.

Nancy L. Stokey, Robert E. Lucas, and Edward C. Prescott, *Recursive Methods in Economic Dynamics* (Harvard University Press, 1989).

Amos Tversky and Daniel Kahneman, “Judgement Under Uncertainty: Heuristics and Biases,” *Science*, vol. 185, no. 4157 (1974), pp. 1124–1131.

Michael Woodford, “Imperfect Common Knowledge and the Effects of Monetary Policy,” in Philippe Aghion and others, eds., *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps* (Princeton University Press, 2003), pp. 25–58.

Janet L. Yellen, "Inflation Dynamics and Monetary Policy," remarks at The Philip Gamble Memorial Lecture, University of Massachusetts, Amherst (2015).