The Role of Expectations in Changed Inflation Dynamics*

Damjan Pfajfar and John M. Roberts
Federal Reserve Board

The Phillips curve has been much flatter in the past 20 years than in the preceding decades. We consider two hypotheses. One is that prices at the microeconomic level are stickier than they used to be. The other is that expectations of firms and households about future inflation are now less well informed by macroeconomic conditions. We use inflation expectations from surveys to help distinguish between our two hypotheses empirically. We find that reduced attentiveness can, in some cases, account for three-fourths of the reduction in the sensitivity of inflation to economic conditions in recent decades.

JEL Codes: E31, E37.

1. Introduction

As many authors have noted, the Phillips curve is much flatter than it used to be; a sampling includes Atkeson and Ohanian (2001), Roberts (2006), Mavroeidis, Plagborg-Møller, and Stock (2014), and Blanchard (2016). We explore two hypotheses about the origin of the

*We would like to thank the editor Boragan Aruoba for his comments and for sharing his ATSIX data. We would also like to thank two anonymous referees, Fernanda Nechio, Sergey Slobodyan, and Ed Knotek for their comments. In addition, we would like to thank seminar participants at Federal Reserve Board, IAAE 2018 Annual Conference, System Macro Meetings at the Federal Reserve Bank of San Francisco, 2019 CEBRA conference in New York, 2019 CFE conference in London, and Bank of Russia Conference Inflation: New Insights for Central Banks for their comments. Kevin Starnes and Samuel Messer provided excellent research assistance. The views expressed in this paper are those of the authors and do not necessarily reflect those of the Federal Reserve Board. Author contact: Board of Governors of the Federal Reserve System, 2000 Constitution Avenue NW, Washington, DC 20551, U.S.A. E-mail: damjan.pfajfar@frb.gov and john.m.roberts@frb.gov. Website: https://sites.google.com/site/dpfajfar/ and https://www.federalreserve.gov/econres/john-m-roberts.htm.

flatter Phillips curve. One is that prices at the firm level are "stickier" than in the past. Ball, Mankiw, and Romer (1988), for example, argued that lower inflation would lead to less-frequent adjustment of prices; because inflation has been lower in the past 20 years than in the decades before, we would expect less-frequent price adjustment and therefore, in the logic of sticky-price models, a flatter Phillips curve. The other conjecture is that firms and households pay less attention to macroeconomic conditions when setting wages and prices now than in the past. This conjecture was articulated in 2001 by then Federal Reserve Chairman Alan Greenspan, who expressed the hope that lower inflation would imply less of a need for firms and households to pay attention to inflation in making their economic decisions.[1] Although Greenspan did not express his hypothesis this way, it is similar in spirit to the rational inattention hypothesis of Sims (2003), who argued that when an economic decision becomes less salient, rational agents with limited bandwidth will devote less attention to it.

We document that from the perspective of the New Keynesian Phillips curve under model-consistent expectations, the sensitivity of inflation to economic activity has been markedly lower in the period starting 1997 than in the preceding two decades. When interpreted through the lens of the canonical Calvo model of staggered price setting, the frequency of price change fell dramatically. While Nakamura et al. (2018) document some reduction in the frequency of price change at the firm level, a much greater increase in nominal rigidity would be needed to account for the change in the slope of the New Keynesian Phillips curve found when the possibility of inattention is not entertained.

Central to our efforts to distinguish between our two main hypotheses, we bring to bear information on inflation expectations taken from surveys. As a number of authors (Roberts 1997; Mavroeidis, Plagborg-Møller, and Stock 2014; Fuhrer 2017; Coibion, Gorodnichenko, and Kamdar 2018) argue, survey measures of inflation expectations bring valuable additional information to the empirical analysis of aggregate inflation. In particular, Coibion,

[1] "Price stability is best thought of as an environment in which inflation is so low and stable over time that it does not materially enter into the decisions of households and firms" (Greenspan 2001).

Gorodnichenko, and Kamdar (2018) find that the coefficients of a reduced-form Phillips curve shift very little when the estimation is conditioned on household inflation expectations. Roth (2013) also finds little evidence of important shifts in Phillips-curve parameters when conditioned on survey expectations. Roberts (1997) and Fuhrer (2017) argue that the excessive persistence of survey forecasts helps explain inflation dynamics in a structural model. At the same time, Dräger and Lamla (2012), Coibion et al. (2018), and Mertens and Nason (2018) have argued that survey expectations in the past two/three decades have become less responsive to macroeconomic developments.[2] Furthermore, Dräger and Lamla (2018) and Eusepi et al. (2019) point out that after 1996, inflation expectations became more anchored. Our premise is that if survey forecasts are less responsive to economic conditions now than before, and if the inflation process involves expectations that depart from the simple benchmark of model-consistent expectations, then it stands to reason that a misspecified Phillips curve estimated assuming simple model-consistent expectations (MCE) would spuriously indicate that the Phillips curve has flattened in recent years.

We find that, across surveys and time periods, survey measures of inflation expectations react more sluggishly than the simple MCE benchmark would predict. These results are similar to those of Carroll (2003), Coibion et al. (2018), and Mertens and Nason (2018), who also find that expectations as captured by surveys adjust sluggishly.

Results on our central hypothesis are sensitive to the measure of inflation expectations. We find a large reduction in attentiveness across our two subsamples with the University of Michigan's survey of household inflation expectations. Based on the estimates of our model using the Michigan survey as an indicator of expectations, we find that a reduction in attentiveness can account for 75 percent of the reduction in the reduced-form sensitivity of inflation to an identified aggregate demand shock. The remaining 25 percent is explained by changes in the coefficients of the New Keynesian Phillips curve

[2]It could be argued that this is due to increased credibility of central banks. Christelis et al. (2016) and Lamla, Pfajfar, and Rendell (2019) study the relationship between confidence (trust) in the central bank and inflation expectations.

that is part of our structural model, which would reflect changes in the frequency of price change.

When we examine surveys of forecasters, such as the Survey of Professional Forecasters (SPF), we do not find the same sharp change in behavior evident in the Michigan survey. Participants in the SPF appear to be about equally attentive in the two periods we examine, and, conditional on that roughly constant degree of attentiveness, the reduction in the slope of the New Keynesian Phillips curve is about as large as in the estimates assuming model-consistent expectations. It is thus only for the Michigan survey of households that we find support for the Greenspan (2001) conjecture that firms and households would become less attentive in the formation of their expectations of inflation. It is possible that the Michigan survey results present a more accurate picture of the changes in the economy. Coibion and Gorodnichenko (2015b), for example, argue that household expectations may be closer to those of actual decision-makers than are forecasts from economists (such as the SPF) and thus that results based on the Michigan survey should be favored.

Ball and Mazumder (2011, 2019) also study the stability of the Phillips curve.[3] Like us, Ball and Mazumder (2011) posit that changes in the frequency of price adjustment and in expectations formation may have played a role in explaining the shift in Phillips-curve parameters. Ball and Mazumder (2011), however, only look at reduced-form evidence and they do not examine measures of expectations.[4] As a consequence, they are not able to distinguish the contribution of expectations formation to the flattening of the Phillips curve from other factors, such as the frequency of price change. By introducing survey measures of expectations into a structural model that considers inflation and expectations formation jointly, we are able to evaluate the contribution of each factor to the changes in the sensitivity of inflation to economic activity.

[3]As do many other authors, Ball and Mazumder (2011) find a break in the parameters of a reduced-form Phillips curve—in their case, in 1985. Ball and Mazumder (2019) argue that the Phillips curve has been stable since 1985. Dräger and Lamla (2018) and Eusepi et al. (2019), however, show that there was a break in anchoring of inflation expectations around 1996, after the preemptive tightening by the Greenspan-era Federal Open Market Committee.

[4]Ball and Mazumder (2019) consider long-run inflation expectations from the SPF to inform the level at which inflation expectations are anchored.

We cross-check our findings with results from microeconomic studies. Until recently, the available evidence had suggested that, at levels of inflation that have prevailed in the United States, there had been little variation in the frequency of prices change. That was the conclusion, for example, of Bils and Klenow (2004) and Nakamura and Steinsson (2008) in the United States. Examining Mexican data, Gagnon (2009) concluded that at very high levels of inflation (above 15 percent), the frequency of price change was sensitive to the prevailing rate of inflation, but that at levels of inflation below the 10 to 15 percent range (which is at the high end of the U.S. inflation experience), there was little sensitivity of the frequency of price change to inflation. Additional data collected by Nakamura et al. (2018), however, shows that in the late 1970s and early 1980s—a period of relatively high inflation in the United States—firms changed prices more frequently than in the subsequent period. We explore the potential macroeconomic implications of the changes in the price-change frequency documented by Nakamura et al. (2018). As we noted earlier, while we find that this microeconomic evidence predicts some reduction in the slope of the Phillips curve, it cannot fully account for the very large reduction we find in the conventional New Keynesian Phillips curve estimated under model-consistent expectations.

Mavroeidis, Plagborg-Møller, and Stock (2014) conduct an extensive analysis attempting to relate inflation, inflation expectations, and measures of economic activity from a single-equation perspective. Their conclusions are pessimistic: They find that it is not possible to estimate both the relationship between inflation and economic activity and the degree of forward-looking behavior. While their results are somewhat stronger when they introduce survey measures of expectations, they still were not able to estimate the key parameters of interest with any precision.

The questions we address are similar to those of Mavroeidis, Plagborg-Møller, and Stock (2014), and their results suggest that we are entering treacherous waters. However, our focus is different from theirs. Both our work and theirs assess the empirical validity of the canonical hybrid New Keynesian Phillips curve with model-consistent expectations. The specific question of Mavroeidis, Plagborg-Møller, and Stock (2014) is whether expectations belong in a structural model of inflation. They conclude that there is not enough information in the macro data to permit an answer to that question. We instead assume that expectations belong in the structural model of inflation, as in the canonical New Keynesian Phillips curve. But we do not require those expectations to obey the simple MCE benchmark, and we ask to what extent these expectations may differ from that benchmark. Perhaps most importantly, our full-system estimation, in which we use information on both expectations and inflation to inform the structural relationship between economic activity and inflation, allows us to identify separately the degree of the departure from the MCE benchmark as well as the impact of economic activity on inflation conditional on expectations.

2. Theory

2.1 Model

We'll first lay out our model of expectations formation. We designate inflation by Δp_t, and $E_t \Delta p_{t+1}$ represents the expectations of agents setting prices. It is typically assumed that expectations are model consistent. In that case,

$$E_t \Delta p_{t+1} = M_t \Delta p_{t+1}, \tag{1}$$

where $M_t \Delta p_{t+1}$ represents model-consistent expectations.[5] We explore two departures from the simplest version of model-consistent expectations based on informational frictions.[6] The first is motivated by the rational inattention model of Sims (2003). In Sims's model, agents receive only a noisy signal of (future) inflation. Thus, we assume that agents' expectations will be related to the true, model-consistent expectations by

$$E_t \Delta p_{t+1} = \mu M_t \Delta p_{t+1}, \tag{2}$$

[5] By "model-consistent expectations," we mean expectations based on the full structure of the model, including knowledge of all the parameters and contemporaneous shocks of the model. These expectations will be unbiased predictors of inflation.

[6] See Coibion, Gorodnichenko, and Kamdar (2018) for a comparison of informational frictions with other departures from full-information rational expectations.

where $0 \leq \mu \leq 1$. That is, agents' actual expectations move only partially with the ideal, model-consistent expectations. In Sims's model, it is costly for agents to pay attention to the future course of inflation. With greater effort, agents can improve the quality of their inflation forecasts, raising the value of μ. With sufficient effort, expectations will be close to the MCE benchmark and the value of μ will approach 1.

Another hypothesis about expectations formation we consider is related to the epidemiological model of Carroll (2003):

$$E_t\Delta p_{t+1} = (1 - \lambda)M_t\Delta p_{t+1} + \lambda E_{t-1}\Delta p_t. \tag{3}$$

Under Carroll's hypothesis, expectations adjust only gradually toward a well-informed value.[7] When agents learn right away about model-consistent expectations, $\lambda = 0$.

We nest these two conjectures about expectations formation to obtain

$$E_t\Delta p_{t+1} = \mu M_t\Delta p_{t+1} + \lambda E_{t-1}\Delta p_t. \tag{4}$$

The key parameter controlling the degree of attentiveness is μ: If μ is smaller, the fully informed expectational benchmark, $M_t\Delta p_{t+1}$, plays a smaller role in the determination of inflation expectations. Conditional of the value of μ, the coefficient λ determines the extent to which expectations eventually adjust to the fully informed benchmark. Under Carroll's model, $\mu + \lambda = 1$, and if $M_t\Delta p_{t+1}$ were to remain stable, $E_t\Delta p_{t+1}$ would eventually move to it.

In our empirical model of expectations formation, we consider two sources of error in survey expectations. One is measurement error, which will affect inflation expectations but not actual inflation. One potential source of measurement error is sampling error.[8]

[7]Carroll (2003) assumes that expectations of households gradually converge toward expectations of professional forecasters. We instead assume that expectations gradually converge to their model-consistent value. In this sense the specification is more similar to Pfajfar and Zakelj (2014), as the expectations gradually converge to MCE forecast. Coibion and Gorodnichenko (2015a) consider a similar model for the evolution of expectations, Equation 3, p. 2649. The main difference with Carroll's model concerns the final term, which in Coibion and Gorodnichenko's model is $\lambda E_{t-1}\Delta p_{t+1}$. Carroll (2003) provides conditions under which it is appropriate to use a specification such as our Equation (3).

[8]Sampling error is a significant issue in the Michigan Survey of Consumers, as the divergence of views about future inflation across households is very wide.

In addition, survey respondents may report a different number to the survey taker than they use when they actually make decisions (for example, they may report a rounded number; see Binder 2017). This latter source of error could become larger when survey respondents are less attentive.

We also allow for a structural shock to expectations, which, through its impact on expectations, can also affect actual inflation. We interpret this structural shock to expectations as a kind of sunspot, in line with Lubik and Schorfheide (2004), who suggest that an error in expectations that has implications for actual inflation can be interpreted as a sunspot. Thus, our model allows for survey measures to be affected by both sunspots and measurement error, and these innovations are distinguished by their effects on inflation: Measurement error affects the measure of expectations only, whereas the sunspot affects both expectations and actual inflation.

Putting together these various elements gives us our empirical model for survey measures of inflation expectations:

$$E_t\Delta p_{t+1} = \mu M_t\Delta p_{t+1} + \lambda E_{t-1}\Delta p_t + \nu_t, \tag{5}$$

$$S_t\Delta p_{t+1} = E_t\Delta p_{t+1} + u_t, \tag{6}$$

$$u_t = \rho_u u_{t-1} + \omega_t, \tag{7}$$

where $S_t\Delta p_{t+1}$ is the survey measure of expectations. Equation (5) generalizes Equation (4) to allow for a structural shock to expectations, ν. Equation (6) allows for measurement error in survey measures of expectations, and the specification in Equation (7) allows that measurement error to be serially correlated.[9]

Our empirical model of inflation is the hybrid New Keynesian Phillips curve that has been used widely:

$$\Delta p_t - \gamma\Delta p_{t-1} = \beta(E_t\Delta p_{t+1} - \gamma\Delta p_t) + \kappa y_t + \epsilon_t, \tag{8}$$

[9]Melosi (2016) also uses survey expectations as an observable to help identify a structural model of inflation expectations. Melosi (2016), however, uses a different structural model than we do, based on imperfect common knowledge. Fuhrer (2017) includes survey expectations as an observable in a structural macroeconomic model but in a reduced-form fashion; he does not specify a structural model for expectations. Neither paper addresses the possible contribution of changes in expectations formation to the flattening of the Phillips curve.

where y_t is the output gap. As discussed in, for example, Calvo (1983) and Woodford (2003), the parameter κ is related to the frequency of price change: The less frequently prices are changed, the smaller κ will be. Thus, a flattening of the Phillips curve caused by less-frequent price change would manifest as a smaller value of κ. Following the empirical literature (for example, Galí and Gertler 1999 and Christiano, Eichenbaum, and Evans 2005), we allow for partial indexation to lagged inflation.[10]

We complete our model with a reduced-form model of the output gap:

$$y_t = \phi_1 y_{t-1} + \phi_2 y_{t-2} + \phi_3 \Delta p_{t-1} + \phi_4 \Delta p_{t-2} + \eta_t. \qquad (9)$$

This specification allows the lagged inflation gap to capture the empirical regularity of some predictive power of lagged inflation for the output gap. We would expect ϕ_3 and ϕ_4 to be less than zero, reflecting the effect of tighter monetary policy in response to inflation shocks.[11]

We assume that the shocks to the model—ν, ω, ϵ, and η—are mutually uncorrelated white noise. The assumption that survey measurement error is unrelated to the other shocks should be relatively uncontroversial—indeed, that is essentially the definition of measurement error. In addition, we allow for an additional source of variation in expectations that is, effectively, correlated with movements in inflation: the structural expectations shock, ν. The assumption that

[10] Mavroeidis, Plagborg-Møller, and Stock (2014) argue that if we take literally the microfoundations of the New Keynesian Phillips curve, it is inappropriate to use survey expectations in Equation (8). However, Adam and Padula (2011) show that in certain cases, survey expectations can be used. Because our model is more complex than the case Adam and Padula (2011) consider, it is an open question whether their result applies in our case.

[11] Many New Keynesian models make the output gap a function of the real interest rate and include a monetary policy reaction function. We do not take this approach, because the U.S. economy has spent a substantial fraction of the time during our sample period at the effective lower bound (ELB) for nominal interest rates. Taking due account of the ELB would introduce considerable complication and would require taking stands on controversial topics such as the effect of forward guidance and the degree to which asset purchase programs were an adequate substitute for conventional monetary policy. Because our interest is in the inflation process, all that is needed is a simple forecasting equation for the output gap, and we believe Equation (9) serves that role well.

this shock is uncorrelated with the shock to the structural Phillips curve is essentially true by definition—the structural expectations shock is intended to reflect movements in expectations that are not related to other shocks affecting the economy. Moreover, because this shock is allowed to affect inflation simultaneously, the assumption of orthogonality is not restrictive.

Specifically, we assume that the Phillips curve and aggregate demand shocks (ϵ and η) are not correlated. Thus, we assume that the shock in the Phillips curve does not affect output contemporaneously but will affect the output gap through the lags of inflation in the output gap equation. We address the possibility of contemporaneous correlation with two different exercises in the robustness section (Section 5).

2.2 An Illustrative Model Solution

In this subsection, we use a simplified version of our model to illustrate how, in the absence of information on expectations, it can be difficult to distinguish a reduction in nominal rigidity from a reduction in attentiveness. Let's assume that expectations are formed according to Sims's rational inattention model, as in Equation (2). Suppose further that inflation is determined according to a simplified version of the New Keynesian Phillips curve, without indexation. In that case, our models for inflation and expectations are

$$E_t \Delta p_{t+1} = \mu M_t \Delta p_{t+1} \tag{10}$$

$$\Delta p_t = \beta E_t \Delta p_{t+1} + \kappa y_t + \epsilon_t. \tag{11}$$

If we substitute Equation (10) into Equation (11), we obtain

$$\Delta p_t = \beta \mu M_t \Delta p_{t+1} + \kappa y_t + \epsilon_t. \tag{12}$$

Equation (12) can be referred to as the "discounted" New Keynesian Phillips curve, in analogy to the "discounted Euler equation" proposed by McKay, Nakamura, and Steinsson (2016, 2017) (see also Gabaix 2017).

To aid in developing intuition about the possible implications of noisy expectations for empirical estimates of the slope of the Phillips curve, it is instructive to assume a simple AR(1) process for the output gap:

$$y_t = \rho y_{t-1} + \zeta_t. \tag{13}$$

With this assumption, Equation (12) can be solved forward as

$$\Delta p_t = \frac{\kappa}{1 - \beta\mu\rho} y_t + \epsilon_t, \tag{14}$$

assuming ϵ is i.i.d. As can be seen, both κ and μ affect the reduced-form Phillips-curve slope. In particular, a greater degree of nominal rigidity, and thus a smaller value of κ, would predict a reduced sensitivity of inflation to fluctuations in output. And so would a smaller degree of attention—that is, a smaller value of μ. Thus, estimating the Phillips curve with only information on inflation and the output gap would not allow us to distinguish between these two hypotheses. Of course, this is a very stylized model. But we will show later that in more realistic settings, a similar result holds: Shifts in either κ or μ lead to changes in the response of inflation to an aggregate demand shock. An implication is that if in fact the attentiveness of agents has fallen, then assuming $\mu = 1$, as is done in most estimation of New Keynesian models, will lead to a mistaken finding that κ has fallen. The purpose of the present paper is to bring additional information to bear, in the form of data on survey expectations, to help distinguish between these hypotheses.

3. Data and Estimation Details

3.1 Data

Central to our analysis are measures of inflation expectations. One measure is from the Survey of Consumer Attitudes and Behavior conducted monthly by the Survey Research Center at the University of Michigan. This measure of expectations has been collected on a consistent basis since 1978. It measures median household expectations of inflation over the coming 12 months. We also look at surveys of professional forecasters—in particular, the Survey of Professional Forecasters that is currently conducted by the Federal Reserve Bank of Philadelphia. The Survey of Professional Forecasters has several questions about inflation expectations, including forecasts of the

CPI, that are available for most of our sample. For consistency with the Michigan survey, we focus on median expectations over the coming year from these surveys. In the online appendices (available at http://www.ijcb.org), we consider additional measures of inflation expectations, including forecasts of GDP prices from the Survey of Professional Forecasters and the Livingston survey, an alternative survey of inflation forecasters.[12]

In most of our work, we use the CPI for items other than food and energy as the basis for our measure of inflation. We focus on a "core" measure, excluding food and energy, because the New Keynesian model is a model of sticky prices; food and energy prices are relatively volatile and thus the underlying model is not as appropriate for them (see Aoki 2001 for a discussion). We look at the CPI for two reasons. First, it is explicitly the variable that respondents to the SPF are asked to forecast. Second, it is the most widely cited measure of consumer prices and so is likely to line up with the views of respondents to the Michigan survey of households. For our measure of the output gap, we use the measure from the Congressional Budget Office (CBO).

In our empirical work, we detrend inflation and inflation expectations using an estimate of long-run inflation expectations. Specifically, we subtract from our measures of inflation and year-ahead inflation expectations a measure of longer-run inflation expectations that is available in the database for the Federal Reserve's FRB/US model.[13] Such detrending puts our focus on cyclical movements in inflation, which lines up with the emphasis of the theoretical models.

[12]The SPF only began asking about the CPI in 1981. We examined two techniques for extending the sample back to 1978. In one, we relied on the Kalman filter underlying our Bayesian estimation method to fill in the missing values. In the other, we projected the SPF's CPI forecasts on the survey's GDP deflator forecasts, which are available over a longer sample. Both approaches yielded similar results; we report the results from the former method.

[13]Data from the FRB/US model are available at https://www.federalreserve.gov/econres/us-models-about.htm. Specifically, we use the FRB/US variable PTR. Over most of its history, this measure of long-run expectations is based on forecasts of longer-run inflation from surveys of professional forecasters. An alternative approach to estimating trend inflation relies on statistical filters—see, for example, Stock and Watson (2007). We believe that a survey-based measure is more appropriate for our purposes. In particular, it allows us to rely on surveys for both short- and longer-term expectations, removing a possible source of discrepancy.

It also allows us to exploit the greater frequency of cyclical movements, which should allow us to better identify our key parameters. Because our focus is on short-term expectations, we do not address the question of greater "anchoring" of long-run inflation expectations that has recently received some attention.[14] The evolution of the central bank's inflation target, and its implications for the public's expectations for inflation over the longer run, is discussed, for example, in Erceg and Levin (2003).[15]

Here and throughout our empirical work, we will compare estimates over two periods, 1978 to 1996 and 1997 to 2015. The start of the sample is determined by the availability of quantitative measures of year-ahead inflation expectations in the Michigan survey of households. We then divide the sample roughly in half. Dräger and Lamla (2018) and Eusepi et al. (2019) have also noted that after 1996, inflation expectations became more anchored. As our results will demonstrate, the responsiveness of inflation to fluctuations in economic activity is very different in our two subsamples.[16]

3.2 Estimation Approach

While we do not make the benchmark assumption of model-consistent expectations throughout, model-consistent expectations nonetheless play a role in our model. Because our approach is uncommon, it is worthwhile explaining in a bit more detail.

In our model, agents' expectations $E_t \Delta p_{t+1}$ are determined by Equation (5). In particular, these expectations appear in the structural Phillips curve, Equation (8). As can be seen in Equation (5), model-consistent expectations $M_t \Delta p_{t+1}$ play a role in expectations formation, as discussed in Section 2. Crucially, where model-consistent expectations appear, they are solved using the full structure of the model, in the usual way. In particular, the solver for the model (specifically, Dynare) uses the full structure of the model in determining $M_t \Delta p_{t+1}$. This cross-equation aspect of our estimation approach enhances our ability to obtain precise

[14] See, for example, Dräger and Lamla (2018) and Eusepi et al. (2019).

[15] Our empirical equations also include constant terms, which could pick up, for example, biases in trend inflation or the output gap.

[16] In Section 5.2 we explore an alternative split between two subsamples.

estimates of both the Phillips-curve slope κ and the parameters of the expectations process, μ and λ.

We estimate our model using Bayesian methods, implemented using Dynare. The priors for our Bayesian estimation are laid out in the online appendices and are relatively uninformative and the same for both our subsamples. For each estimation, we use two blocks of 500,000 Markov chain Monte Carlo (MCMC) draws using the Metropolis–Hastings algorithm; in total, 1 million draws. In tables, we report the mean posterior estimates together with 90 percent confidence intervals. All estimations include constants in both the Phillips curve and the inflation expectations equation to account for potential differences in the measure of inflation reported by survey respondents and the core CPI inflation.[17]

4. System Estimation Results

In this section, we turn to estimates of the full system of equations outlined in Section 2.1. We compare estimates of the hybrid New Keynesian Phillips curve under two hypotheses about expectations: the model-consistent expectations assumption that is common in the literature and our model of expectations formation that relaxes MCE.

4.1 Model Estimates: MCE

Columns 1 and 2 of Table 1 present estimates of the system of equations consisting of the hybrid New Keynesian Phillips curve, Equation (8), and the reduced-form output gap equation, Equation (9), under the assumption of fully model-consistent expectations. Column 1 shows results over the 1978–96 period; column 2, over the 1997–2016 period. The slope of the Phillips curve, κ, is considerably smaller in the latter sample, by a factor of six-and-a-half. The degree of indexation, γ, is also notably smaller in the latter sample. Posterior mean estimates of γ and κ in the post-1996 sample are outside the 90 percent credible set of its estimate for the earlier sample.

The bottom rows of the table show results for the reduced-form process for the output gap, Equation (9). The parameters ϕ_1 and ϕ_2

[17]For brevity, we do not report the constants in the tables shown in the paper.

Table 1. Estimates of Model under Assumption of Model-Consistent Expectations and with Survey Inflation Expectations

	(1) MCE 1978–1996	(2) MCE 1997–2015	(3) Michigan 1978–1996	(4) Michigan 1997–2015	(5) SPF CPI 1978–1996	(6) SPF CPI 1997–2015
γ	0.739	0.320	0.139	0.566	0.432	0.211
	[0.534, 0.933]	[0.119, 0.522]	[–0.102, 0.375]	[0.133, 0.993]	[0.093, 0.761]	[–0.008, 0.422]
κ	0.072	0.011	0.271	0.143	0.331	0.077
	[0.015, 0.124]	[0.003, 0.018]	[0.111, 0.423]	[0.065, 0.222]	[0.155, 0.504]	[0.033, 0.120]
σ_ϵ	1.991	0.565	1.568	0.796	2.159	0.600
	[1.676, 2.282]	[0.478, 0.648]	[1.136, 1.986]	[0.528, 1.069]	[1.550, 2.741]	[0.467, 0.722]
λ	—	—	0.786	0.390	0.778	0.454
			[0.682, 0.898]	[–0.013, 0.802]	[0.683, 0.874]	[0.299, 0.621]
μ	—	—	0.198	–0.046	0.159	0.336
			[0.087, 0.303]	[–0.340, 0.244]	[0.087, 0.230]	[0.196, 0.472]
σ_ν	—	—	0.322	0.218	0.163	0.064
			[0.172, 0.458]	[0.075, 0.352]	[0.081, 0.241]	[0.021, 0.104]
ρ	—	—	0.879	0.640	0.837	0.256
			[0.743, 1.000]	[0.452, 0.847]	[0.496, 1.000]	[0.000, 0.470]
σ_ω	—	—	0.319	0.379	0.175	0.135
			[0.172, 0.458]	[0.285, 0.475]	[0.083, 0.261]	[0.099, 0.171]
$\varnothing_1$	1.219	1.292	1.214	1.315	1.229	1.331
	[1.039, 1.398]	[1.128, 1.461]	[1.039, 1.394]	[1.142, 1.490]	[1.050, 1.415]	[1.161, 1.502]
$\varnothing_2$	–0.273	–0.290	–0.260	–0.328	–0.277	–0.343
	[–0.455, –0.095]	[–0.459, –0.119]	[–0.441, –0.079]	[–0.510, –0.152]	[–0.454, –0.093]	[–0.517, –0.165]
$\varnothing_3$	–0.106	–0.115	–0.099	–0.035	–0.101	–0.082
	[–0.177, –0.037]	[–0.331, 0.097]	[–0.168, –0.026]	[–0.253, 0.179]	[–0.173, –0.031]	[–0.296, 0.138]
$\varnothing_4$	–0.055	–0.251	–0.046	–0.179	–0.040	–0.120
	[–0.132, 0.023]	[–0.465, –0.044]	[–0.125, 0.031]	[–0.393, 0.032]	[–0.119, 0.042]	[–0.335, 0.101]
σ_η	0.655	0.584	0.659	0.581	0.661	0.581
	[0.563, 0.743]	[0.503, 0.665]	[0.568, 0.746]	[0.500, 0.659]	[0.569, 0.752]	[0.500, 0.660]

Note: Estimated using Bayesian methods: see text for details; the priors are described in online appendix B. 90 percent confidence intervals are in square brackets.

suggest that the process in both periods has the "hump-shaped" pattern typical of the response of output to identified aggregate demand shocks in estimated structural VARs and DSGE models; there is a posterior mean estimate greater than one on the first lag of the gap and a negative estimate on the second lag. ϕ_3 and ϕ_4 show the sensitivity of the output gap to lagged inflation. In each estimation period, the sum of the two estimates is negative, as expected.

4.2 Model Estimates: Generalized Model of Expectations

Columns 3 to 6 of Table 1 present results for the model we introduced in Section 2.1, in which the assumption of model-consistent expectations is relaxed and survey expectations are added as an observable. Recall from Equations (5) and (6) that the model of expectations has several key features: It allows expectations to react to incoming information by less than predicted by the MCE hypothesis ($\mu < 1$); it allows for gradual adjustment of expectations ($\lambda > 0$); and it allows for shocks to the process for inflation expectations, either in the form of measurement error in the survey ($\sigma_\nu > 0$) or as shocks to correctly measured expectations, which can go on to affect actual inflation ($\sigma_\eta > 0$).

Results with the Michigan survey are in columns 3 and 4. Focusing first on the parameters related to inflation expectations, the results suggest a substantial departure from purely model-consistent expectations in both samples. In particular, the posterior mean estimates of μ are smaller than one in both samples, and by a large margin. The results in the early sample provide strong support for Carroll (2003)'s epidemiological model: The sum of the estimates μ and λ is 0.98, very close to the value of one suggested by Carroll's model. Thus, although households do not react immediately to new information about future inflation, under the epidemiological interpretation, the knowledge would eventually spread.

Estimates of both μ and λ are smaller in the later sample. In the early sample, the credible set for μ lies above zero, while in the latter sample, the point estimate is actually negative, albeit close to zero. Moreover, the estimates in the latter sample are outside the credible set for the early sample. A value of $\mu = 0$ would imply that households no longer pay attention to macroeconomic fundamentals in setting their inflation expectations.

Turning next to the estimates for the slope of the Phillips curve, κ, we find that, as with the MCE estimates, κ is smaller in the 1997–2015 sample. However, the extent of the decline in κ is considerably smaller than in the MCE case, with the estimate of κ dropping from 0.27 to 0.14, a decline of about 50 percent, compared with a decline of 85 percent in the first two columns. The mean posterior estimate of κ in the first sample lies outside the credible set for the second sample. According to the canonical Calvo model, this reduction in the value of κ corresponds to less-frequent price adjustment in the post-1996 sample.

On our preferred interpretation, a key reason for the smaller decline in the slope coefficient when the MCE assumption is relaxed is that households pay considerably less attention to the fundamentals in forming expectations than was the case in the earlier period, as indicated by the smaller value of μ. Thus, the smaller reduction in κ is consistent with the view that, at least in part, the reduction in the sensitivity of inflation to economic activity can be explained by a reduction in the attention paid to inflation by firms and households. As we will see in Section 6.1, the rise in inattention explains about three times more of the reduction in the overall sensitivity of inflation to economic activity than the reduction in κ.

In the early sample, there is little evidence of inflation persistence: $\gamma = 0.14$, where the credible set includes zero. This finding is consistent with the results of Roberts (1997) and Fuhrer (2017), who also found that conditioning on survey expectations led to reduced evidence of other sources of inflation persistence. γ rises to 0.57 in the latter sample.

Both measurement error and structural shocks to expectations are important sources of variation in the Michigan survey. The structural shock to expectations (σ_ν) is somewhat less important in the latter sample, with a standard deviation that is about two-thirds as large—although each estimate lies within the credible set of the other. In addition, λ is considerably smaller in the latter sample, so that any given shock will be carried forward with less persistence in the latter sample and, as we will see in Section 4.4, ν accounts for much less of the variability in inflation in the latter sample. Measurement error is large in both samples; it also displays considerable serial persistence, especially in the early sample.

As noted in Section 2, we interpret our structural expectations shocks as a form of sunspot. As discussed in Lubik and Schorfheide (2004), sunspots are more likely to arise when central bank control of inflation is weak. Our early sample includes the late 1970s and early 1980s, a period that a number of authors have identified as a period of transition from weak inflation control (Clarida, Galí, and Gertler 2000, Lubik and Schorfheide 2004, Roberts 2006). On this interpretation, it is not surprising that sunspot-related shocks are more prevalent in our early sample.

Columns 5 and 6 present results for the SPF.[18] Starting with the results for the expectations process, as with the Michigan survey, the simple MCE hypothesis is strongly rejected, with μ far from the MCE-implied value of one in both estimation periods. Also like the Michigan survey, the sum of λ and μ is fairly close to one in the early sample (= 0.94) and falls in the latter sample. One key difference from the Michigan results is the evolution of μ over time: For the SPF, the point estimate of μ actually rises somewhat in the latter period, in contrast to the sharp drop for the Michigan survey. That result suggests that professional forecasters continued to pay attention to fundamentals, albeit imperfectly, in the post-1996 sample, in contrast to the households captured by the Michigan survey, who, apparently, paid essentially no attention in the latter sample. We return to an interpretation of this finding in Section 4.3.[19]

[18]The SPF potentially exhibits a structural break in 1990:Q2, when the Federal Reserve Bank of Philadelphia started administrating the survey (before it was conducted by the American Statistical Association and the National Bureau of Economic Research). When we discuss robustness in Section 5.2, we perform the same analysis using alternative measures of expectations of professional forecasters.

[19]For the latter sample, our results on the attentiveness of professional forecasters are similar to those of Coibion and Gorodnichenko (2015a) and Mertens and Nason (2018). Those authors find that from about the mid-1990s onward, the professional forecasters captured by the SPF were relatively inattentive. Our results differ from these authors for the pre-1997 sample, however: Coibion and Gorodnichenko (2015a) and Mertens and Nason (2018) find that from the early 1970s through the early 1990s, SPF respondents were relatively attentive, with a weight on $M_t \Delta p_{t+1}$ that is close to one. There are many differences in statistical approach across these three papers. One is that we are imposing a structural New Keynesian Phillips curve as the model underlying the inflation process, which the other papers do not do. However, as we discuss in Section 5.1, we obtain similar results using a single-equation approach. We leave a resolution of these differences to future research.

For the Phillips curve, the early-sample estimates are broadly similar to those for the Michigan survey: The point estimate of κ is similar and, while the estimate for γ is larger with a credible set excluding zero, it is nonetheless of modest size. The results for the latter sample are different, however—the slope coefficient κ falls by about 75 percent, considerably more than for the Michigan survey—and are reminiscent of the MCE results presented in Section 4.1. The point estimates in each sample lie outside the credible set in the complementary sample.

As with the Michigan survey, there is evidence of both measurement error and structural shocks to expectations. The results suggest, however, that measurement error is less important for the SPF than for the Michigan survey, as both σ_ω and the persistence of measurement error, ρ, are smaller for the SPF, especially in the later sample. The structural expectations shock, ν, is also less important for the SPF than for the Michigan survey and, as with measurement error, is even less important in the latter sample.

For both the SPF and the Michigan survey, the point estimates of the Phillips-curve slope parameter are very different from those in columns 1 and 2—in particular, they are much larger than in the MCE case, in both samples. In their overview of empirical work on the New Keynesian Phillips curve, Mavroeidis, Plagborg-Møller, and Stock (2014) also find that estimates of κ are larger when expectations are proxied using surveys. The larger value of κ is consistent with the intuition provided by the simple model introduced in Section 2.2: When agents are less attentive, expectations of future inflation move less for any given change in the current-period output gap. To account for the same observed change in inflation, the model ascribes a larger role to the current-period output gap.

4.3 A Case for Preferring the Michigan-Based Results

Coibion, Gorodnichenko, and Kamdar (2018) argue strongly that the Michigan survey is to be preferred as a measure of inflation expectations. They argue that the preferred measure of inflation expectations is that of firms, as it is their expectations that most matter for pricing decisions. They cite their own work with inflation expectations in New Zealand, which suggests that expectations of firms are similar to those of households. As Carroll (2003) emphasizes, professional forecasters can be expected to be better informed than others about the state of the economy.[20] Thus, Michigan survey expectations may be closer to the expectations of decisionmakers in the economy than are the expectations of professional forecasters. Coibion, Gorodnichenko, and Kamdar (2018) go on to cite earlier work by two of them (Coibion and Gorodnichenko 2015b) that found that the Michigan survey was particularly helpful in explaining the lack of disinflation in the Great Recession. And they present new evidence suggesting that the Michigan survey performs better in empirical inflation models.

The results in Table 1 are broadly consistent with this view: consistent with the Greenspan hypothesis, the estimate of μ based on the Michigan survey is very small in the latter sample. In addition, the decline in κ is relatively modest in this case, with a decline of about 50 percent, consistent with the view that a larger decline in attentiveness implies a smaller decline in κ.

Cecchetti et al. (2017) argue that in recent decades, survey measures of inflation expectations have not been helpful in explaining actual inflation dynamics, in contrast to results in earlier (Roberts 1997) and longer (Mavroeidis, Plagborg-Møller, and Stock 2014; Fuhrer 2017) samples. Our results help explain why this might be the case. First, in our latter sample, μ is very small, notably so for the Michigan survey. Hence, the survey conveys less useful information about expectations of future economic conditions. Second, the structural (sunspot) shock to expectations is less important in the latter sample. So surveys are bringing less independent information to bear in the latter sample. As noted earlier, this outcome is consistent with predictions that in periods with greater inflation control by the central bank, sunspot equilibria are less likely to arise.

[20]Indeed, to the extent that professional forecasters make their living providing accurate assessments of the economy's evolution, it is perhaps not surprising that they would continue to pay appropriate attention to the relation between macroeconomic conditions and inflation. It is therefore possible that while these forecasts are more accurate, they are at the same time less relevant to price setting.

Table 2. Variance Decomposition for the Michigan Survey and Core CPI Inflation Based on Parameter Estimates in Table 1, Columns 3 and 4

	(1) Michigan 1978–1996	(2) Michigan 1997–2015	(3) Core CPI 1978–1996	(4) Core CPI 1997–2015
ε (Shock to the PC)	4	0	52	66
η (Aggregate Demand Shock)	37	0	32	27
ν (Structural Shock to Exp.)	39	19	15	7
ω (Measurement Error to Exp.)	21	81	0	0
Variance	2.2	.30	4.1	.47
Standard Deviation	1.5	.55	2.0	.68

Note: Table entries are the percent of variance of each variable explained by each of the model's shocks.

4.4 Variance Decompositions

In this section, we examine the contributions of the model's structural shocks to the variability of inflation and inflation expectations. Columns 1 and 2 of Table 2 present a formal variance decomposition of the Michigan survey in the two samples, based on the results in columns 3 and 4 of Table 1. In the early sample, the variation of the Michigan survey is importantly influenced by the business cycle: The cycle (η) accounts for 37 percent of the variability of the Michigan survey (column 1). The structural shock to expectations (ν) also accounts for a substantial portion of the variability of the Michigan survey, while the measurement error shock accounts for about one-fifth of the total variation.

In the latter sample, column 2, the coefficient μ is set equal to its theoretical lower bound of zero. As a consequence, neither of the economy's fundamental shocks, ϵ and η, account for any part of the Michigan survey's variation. Measurement error (ω) accounts for 80 percent of the variability of the Michigan survey in the latter sample. Thus, fluctuations in the Michigan survey are largely noise in the post-1996 period. It is worth noting, however, that the total

Table 3. Variance Decomposition for the SPF and Core CPI Inflation Based on Parameter Estimates in Table 1, Columns 5 and 6

	(1) SPF CPI 1978–1996	(2) SPF CPI 1997–2015	(3) Core CPI 1978–1996	(4) Core CPI 1997–2015
ε (Shock to the PC)	8	5	73	60
η (Aggregate Demand Shock)	56	65	24	38
ν (Structural Shock to Exp.)	28	10	3	2
ω (Measurement Error to Exp.)	9	20	0	0
Variance	.67	.10	4.7	.49
Standard Deviation	.82	.32	2.2	.70

Note: Table entries are the percent of variance of each variable explained by each of the model's shocks.

amount of noise in the Michigan survey is actually smaller in the latter sample. That's because the total variation in the Michigan survey is dramatically lower (the variance falls by around 85 percent).

Columns 3 and 4 of Table 2 show the contributions of the model's shocks to the variance of core CPI inflation in the two subsamples. The business cycle shock, η, accounts for around 30 percent of the variability of core CPI inflation in both samples. As with the results in columns 1 and 2, it is important to remember that the variance of inflation is much smaller in the latter sample, here falling by almost 90 percent. So in the latter sample, the business cycle is explaining 27 percent of a small number. In the early sample, the structural shock to inflation expectations accounts for about 15 percent of the variability of inflation. Thus, "sunspots" make a nontrivial contribution to the variability of inflation in this period. By assumption, survey measurement error makes no contribution to the variability of actual inflation.

Table 3 shows variance decompositions based on the model estimates for the SPF, from columns 5 and 6 of Table 1. As with the Michigan survey, the shock to the cycle—η—accounts for a substantial portion of the variability of inflation in both samples. One key difference in the results is that the fundamental shocks ϵ and, especially, η continue to account for a large portion of the variability of the SPF in the latter sample, as, according to the results in Table 1, SPF respondents continue to put important weight on the fundamentals in forming their inflation expectations.

5. Robustness

In this section, we provide various sensitivity analyses. We first look at single-equation methods. We then check the robustness of our multi-equation results with respect to different processes of inflation expectations, different measures of inflation expectations, different process and measure of the output gap, allowing for correlation of shocks in the estimation, regarding our detrending method, and to the definition of the two subsamples.

5.1 Single-Equation Estimates

In this section, we present single-equation, instrumental-variable estimates of our central equations of interest, the New Keynesian Phillips curve, Equation (8), and our model of expectations formation, Equation (5). These single-equation estimates provide a check on the robustness on the multi-equation approach of Section 4, in particular relaxing the assumptions that underlie our structural model.

We derive our equations for single-equation estimation by reorganizing our structural model in Section 2 as orthogonality conditions in which only observable variables appear. We then look for instruments that will be correlated with the observables and uncorrelated with the equation residuals. We begin by combining Equations (5) and (6) to obtain the following orthogonality condition related to survey expectations:

$$\nu_t + u_t - \lambda u_{t-1} + \mu(M_t \Delta p_{t+1} - \Delta p_{t+1}) = S_t \Delta p_{t+1} - \lambda S_{t-1} \Delta p_t - \mu \Delta p_{t+1}. \tag{15}$$

We seek instruments that are correlated with the terms on the right-hand side of Equation (15) and orthogonal to terms on the left-hand side. The first left-hand-side term, ν, will be correlated with any variable directly affected by the contemporaneous structural shock to expectations; importantly, ν is itself not serially correlated. In this case, both contemporaneous inflation and the survey itself are excluded as instruments. Note that, from Equation (13), the next two terms, u_t and u_{t-1}, which capture survey measurement error, are serially correlated. Thus, because any lagged values of the observable survey expectation $S_t\Delta p_{t+1}$ that appears in Equation (15) are excluded. Finally, the term $\mu(M_t\Delta p_{t+1} - \Delta p_{t+1})$ reflects the forecast error of inflation in period $t+1$ for forecasts made in period t. Because this forecast error is by assumption a rational one, it should not be serially correlated. The presence of this term would exclude any future variable as an instrument.

We can combine Equations (8) and (6) to obtain the following orthogonality condition related to the Phillips curve:

$$\epsilon_t - \beta u_t = (1+\beta\gamma)\Delta p_t - \gamma\Delta p_{t-1} - \beta S_t\Delta p_{t+1} - \kappa y_t. \tag{16}$$

Any variable not directly affected by contemporaneous inflation will be uncorrelated with the first term, ϵ. Thus, contemporaneous inflation itself and contemporaneous survey expectations would be excluded; lagged values of inflation are acceptable. Given our structural model, the contemporaneous output gap would be an acceptable instrument. In this analysis, however, we relax this assumption and only allow lagged values of the output gap to serve as instruments. As with the estimation of Equation (15), the presence of u_t excludes current or lagged values of the survey $S_t\Delta p_{t+1}$ appearing in Equation (16) as instruments.

We include the following variables as instruments in the estimation of Equations (15) and (16): In both equations, we include two lags of the output gap and of the inflation gap as instruments. As noted above, because measurement error is serially correlated in our model, lagged values of the survey that is included in each equation are not valid instruments. Note, however, that the survey universes of the two main surveys we examine—the Michigan survey of households and the Survey of Professional Forecasters—are entirely independent of one another. Thus, the measurement errors of these two surveys are not correlated. In Equation (15), we use two lagged values of the (complementary) survey; in the Phillips curve, we use the contemporaneous and one lagged value.

An important consideration in the estimation of models using instrumental variables is the explanatory power of the instruments for the observables. In a recent survey article, Andrews, Stock, and Sun (2019) recommend the use of the "effective F" statistic proposed by Montiel-Olea and Pflueger (2013) as the preferred robust test for instrument power. Unfortunately, the effective F statistic has only been worked out in the case of one instrumented variable in an equation and each of our equations includes two such variables. As an alternative, we report results from another robust weak-instrument test, the Kleibergen-Paap test, which is suitable for the case of multiple endogenous variables. For comparison, we also report the robust F statistic from the first-stage regression for each variable.[21]

A final consideration is that, because instruments can be weak not only when their correlation with the endogenous variable is literally zero but also when it is in a neighborhood around zero, conventional F distributions are not used. Instead, more conservative thresholds have been proposed: Montiel-Olea and Pflueger (2013) advocate a threshold of 23, while in their recent survey Andrews, Stock, and Sun (2019), drawing on earlier work by Staiger and Stock (1997), suggest that 11 may be conservative enough.[22]

Table 4 shows our estimates of Equations (15) and (16).[23] Many of the results in Table 4 are similar to those from our system estimation. For the equation explaining the surveys (columns 1, 2, 5, and 6), the coefficient on expected inflation drops notably across samples in the equation for the Michigan survey, but not in the equation explaining the SPF, just as in our structural results. For the Phillips curve (columns 3, 4, 7, and 8), the coefficient on the output gap is roughly similar in both samples when expectations are captured by the Michigan survey, but the coefficient falls dramatically in the latter sample when expectations are captured by the SPF—again, as in our structural estimation. And as before, the coefficient on the

[21] As a cross-check, we also computed the effective F statistic for each endogenous variable; the results were qualitatively similar to those from the reported robust F statistics.

[22] These thresholds compare with a conventional-F 5 percent rejection threshold of around 2 in the case of moderate overidentification.

[23] For estimation, we use GMM with an HAC weighting matrix that assumes a Bartlett kernel and a Newey-West serial-correlation adjustment computed with a fixed bandwidth of 4.

Table 4. Single-Equation Estimates with GMM

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
	Michigan Survey				**SPF Survey**			
Dep. Var.→	$S_t\Delta p_{t+1}$		$\Delta p_t - \beta S_t\Delta p_{t+1}$		$S_t\Delta p_{t+1}$		$\Delta p_t - \beta S_t\Delta p_{t+1}$	
Sample→	**1978–1996**	**1997–2015**	**1978–1996**	**1997–2015**	**1978–1996**	**1997–2015**	**1978–1996**	**1997–2015**
$S_{t-1}\Delta p_t$	.689	.417	—	—	.744	.757	—	—
	(.082)	(.300)			(.094)	(.247)		
	[43.4]	[4.6]			[48.8]	[50.9]		
Δp_{t+1}	.246	.053	—	—	.121	.150	—	—
	(.094)	(.119)			(.042)	(.180)		
	[9.9]	[5.9]			[14.6]	[6.6]		
$\Delta p_{t-1} - \beta\Delta p_t$	—	—	.287	–.036	—	—	.071	–.015
			(.141)	(.130)			(.172)	(.108)
			[6.2]	[23.5]			[9.4]	[26.2]
y_t	—	—	.122	.155	—	—	.186	.025
			(.083)	(.030)			(.109)	(.030)
			[83.2]	[335.0]			[117.6]	[340.7]
$P(j\text{-}stat)$	.22	.27	.72	.39	.73	.22	.51	.58
F statistics	5.6	20.5	27.7	31.8	2.9	3.0	16.1	66.8

Note: In columns 1–2 and 5–6, two lags of the inflation gap, two lags of the output gap, two lags of the other survey (SPF for Michigan, Michigan for SPF) are used as instruments. In columns 4–5 and 7–8, two lags of inflation gap, two lags of output gap, current value, and one lag of other survey (SPF for Michigan, Michigan for SPF) are used as instruments. Coefficient standard errors are in parentheses. Brackets include first-stage robust F statistics for instruments for that variable. The F statistic for the overall estimation is the Kleibergen-Paap F statistic.

term related to lagged inflation is economically small and, in three of four cases, not significantly different from zero.

While broadly consistent with the results from the system estimation, these results come with a number of caveats. First, and as might be expected, the parameter estimates are less precise than in our system approach, which benefits from imposing restrictions implied by theory. Second, weak-instrument tests often fail. This is especially true for the survey equation, where for three of four specifications, the Kleibergen-Paap statistic is below even the more generous Staiger-Stock criterion of 11. The instruments are stronger in the case of the Phillips curve, where in three of four cases the Kleibergen-Paap statistic is beyond the strict Montiel-Olea and Pflueger (2013) threshold of 23. One bright spot is that the over-identifying restrictions, captured by the J-statistic, are not rejected in any of the eight cases.

5.2 Sensitivity of Full-System Estimates

We examine the robustness of our multi-equation estimates along a number of dimensions.[24] We begin by assessing the sensitivity of the inflation expectations process. We relax Equation (5) by allowing a more general process of expectation formation that allows the key features of a reduced-form Phillips-curve relationship—the output gap and lagged inflation—to enter the specification for inflation expectations directly (see Equation (A.1) in online appendix A). Table A.1 presents the results. As in Carroll (2003) and Pfajfar and Santoro (2013), the credible set of the estimate on the past inflation rate includes zero for the Michigan survey. For the early-sample SPF estimates of the same coefficient, the credible set also includes zero, while for the later period it does not.

The credible set on the estimate for the output gap excludes zero in three of the four cases. From the standpoint of our structural model of expectations formation, these results suggest that inflation expectations are more sensitive to economic activity than would be justified by the model structure. Of course, Equation (A.1) can itself be viewed as a Phillips-curve relationship. This is the approach taken by Dräger, Lamla, and Pfajfar (2016), who examine the individual

[24]The results in this section are available in the online appendices to this paper.

responses underlying the surveys and find that only about a third of participants in the Michigan survey forecast unemployment and inflation consistent with the Phillips curve trade-off, while the share for the SPF is about one-half.[25]

For the SPF, the estimates of the key parameters of the Phillips curve, γ and κ, are similar to those in Table 1. For the Michigan survey, however, the estimate of κ is substantially smaller, and credible sets in both samples include zero. As we discuss in the online appendices, this may be because the output gap has a very strong direct effect on expectations in this model.

We also check whether oil-price shocks play an important role in forming inflation expectations. Coibion and Gorodnichenko (2015a), for example, show the importance of oil prices for formation of inflation expectations. We thus further augment Equation (A.1) by including oil-price shocks calculated as in Hamilton (1996). Our results indicate that oil-price shocks are insignificant in all regressions and do not directly influence the formation of inflation expectations in either of our samples. Results are presented in Table A.4.[26]

Third, we check the robustness of our estimates with respect to the measure of inflation expectations. In Table A.2, we present the results for the Livingston survey and the SPF forecast of the GDP price deflator. Results are discussed in detail in the online appendices. In general, they confirm our baseline results.[27]

In our fourth exercise, we take the position that three of the measures of expectations of consumer prices—from the SPF, the Livingston survey, and the Michigan survey—are noisy indicators of the same underlying process. Thus, we assume that the same "true"

[25]Similarly, Carvalho and Nechio (2014) find that only some households—in particular, those with at least a college degree—have interest rate expectations that are broadly consistent with the Taylor rule.

[26]Our results differ from those of Coibion and Gorodnichenko (2015a), who find an important role of oil prices in determining inflation expectations. Our approach differs from theirs in a number of dimensions, most notably that we focus on oil-price shocks and on core CPI rather than oil prices directly and headline CPI.

[27]We also check the robustness of our results for the Michigan survey by considering only forecasts of those who have a college degree or high-income households. Binder (2015) points out that forecasts of these individuals are actually closer to the forecasts of firms. Results are presented in Table A.5 and are indeed very similar to our baseline results using the Michigan survey.

measure of expectations is a common factor driving all three surveys. Table A.3 shows results for this joint estimation. The results are similar to those for the SPF shown in Table 1, with the point estimate of κ falling substantially and μ rising across sample periods.

Fifth, we consider a more general process for the output gap. Specifically, to Equation (9), we add additional two lags of inflation and the output gap to check whether the process for output gap would importantly alter our results. As shown in Table A.6, the main results are virtually identical. The additional lags are sometimes important for the early sample, while for the latter sample they are always close to zero.

Sixth, we replace the CBO output gap in Equation (9) with the CBO short-run unemployment rate gap.[28] Results in Table A.7 again suggest that our results are unaffected by this change and the estimates of the slope of the Phillips curve are similar to the one that we would get if we were to multiply the estimates in Table 1 with an Okun's law coefficient of –0.5.

Seventh, we check the robustness of our estimates in the baseline Table 1 by allowing for various shocks in our system estimation to be correlated. Table A.8 shows the results when we allow $corr(\sigma_\eta, \sigma_\epsilon)$—that is, the correlation between Phillips-curve shocks and the output disturbances—to be nonzero. The estimates suggest that this correlation is small, with its credible set always including zero, and thus our main results are unchanged. In Table A.9, we additionally estimate $corr(\sigma_\nu, \sigma_\epsilon)$ (the correlation between the Phillips-curve shock and the shock to inflation expectations). For the Michigan survey, the results are again virtually unchanged, as this additional correlation is also small. In the early sample SPF $corr(\sigma_\eta, \sigma_\epsilon)$ is positive, while in the latter sample $corr(\sigma_\nu, \sigma_\epsilon)$ is negative. Nevertheless, the estimates of μ and κ are similar to those in Table 1, although λ is lower in the late sample and γ is higher.

Eighth, we study sensitivity with respect to the detrending procedure, estimating the model without detrending. Estimates are available in Table A.10. Results are again very similar to the baseline results in Table 1.

[28]The CBO differentiates between long-run and short-run natural rate of unemployment, while the latter is designed to be used for forecasting inflation.

Ninth, we check the robustness of our results to a different definition of the break between the early and late sample. The choice of a break in 1997 is in line with one of the two breaks considered in Ball and Mazumder (2019). Dräger and Lamla (2018) and Eusepi et al. (2019) also show that after 1996 inflation expectations became more anchored. Alternatively, the literature that studies breaks in inflation process suggest that there may be a break somewhere between 1991 and 1993 (see, for example, Cecchetti and Debelle 2006). Thus, as a robustness test, we impose a break in the middle of 1992. Results are presented in Table A.11 and are very similar to those for our baseline subsamples.

In the last robustness check, we study the sensitivity of our late-sample SPF CPI results to a measure of one-year-ahead inflation expectations (ATSIX) derived from SPF CPI forecasts and Blue Chip forecasts that was proposed by Aruoba (2020) and is updated by the Federal Reserve Bank of Philadelphia. This measure has the advantage that it allows for a computation of inflation expectation at time t for forecast between t and $t + 4$.[29] The results using the ATSIX measure are reported in Table A.12, where for a comparison we reproduce our results from column 6 of Table 1. The results are very similar to the baseline.

6. Structural Interpretation

6.1 Implications for Inflation Dynamics

In this section, we use impulse responses to evaluate our two hypotheses and to explore the ability of changes in expectations formation to account for these changes in inflation dynamics.

Figure 1 shows the effects of a one-standard-deviation shock to the output gap equation in different variants of the model that uses the Michigan survey as its measure of inflation expectations. The solid blue line shows the results from the estimates in column 3 of Table 1, which used data from the 1978–96 period data.[30] In

[29] Strictly speaking, due to timing of the SPF, the inflation forecast that we use in our baseline analysis is at time $t-0.5$ for forecast between t and $t + 4$, which means that the information set in this forecast is less consistent with the model compared to the ATSIX one-year-ahead inflation expectations.

[30] For figures in color, see the online version of the paper at http://www.ijcb.org.

Figure 1. Impulse Response Function of Inflation to an Aggregate Demand Shock in the Structural Model

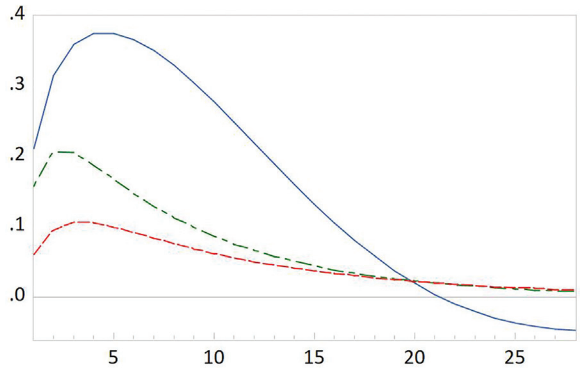

Note: Solid blue line presents results from model in column 3, Table 1, using the 1978–96 period and the Michigan survey. Red dashed line presents results that substitute the inflation and inflation expectations results from column 4 of Table 1, using the 1997–2015 period and the Michigan survey. In the green dot-dashed line, only the inflation expectations equation estimates from column 4 are used. See text for further details.

the dashed red line, the estimates for the equations for inflation and inflation expectations use the results from column 4 of Table 1, which are based on the 1997–2015 sample. To isolate the effects of the change in inflation dynamics, the dashed red version uses the same output gap equation as in the solid blue simulation. Consistent with the notion that aggregate demand shocks have a smaller impact on inflation than in the past, the effects of the aggregate demand shock are much larger in the earlier sample; for example, the peak effect is more than three times greater and the average effect over the first 12 quarters is similarly greater.

To isolate the contribution of the change in inflation dynamics, the dot-dashed green line shows a simulation that uses the estimated expectations-formation equation from the latter period along with the inflation equation from the early period. According to this simulation, the change in the inflation expectations process explains about 75 percent of the reduced effect of the aggregate demand shock on inflation over the first 12 quarters. Shifts in the structural Phillips-curve parameters, κ and γ, account for the rest.

6.2 Micro Evidence on the Frequency of Price Change

As mentioned in the introduction, Nakamura et al. (2018) find that prices changed more often in the United States in the high-inflation period of the late 1970s and early 1980s than in the subsequent period. In this section, we assess the potential implications of such a reduction in the frequency of price change for estimates of the slope of the Phillips curve.

As discussed in, for example, Woodford (2003), the slope of the Phillips curve in the New Keynesian model can be thought of as composed of two components, one related to the frequency of price change (α) and the other to the sensitivity of marginal cost to the state of the economy (ζ):

$$\kappa = \frac{(1-\alpha)(1-\beta\alpha)}{\alpha}\zeta, \tag{17}$$

where, recalling Equation (11), κ is the slope of the New Keynesian Phillips curve:

$$\Delta p_t = \beta E_t \Delta p_{t+1} + (1-\beta)\bar{\pi} + \kappa y_t + \epsilon_t. \tag{18}$$

According to the Ball, Mankiw, and Romer (1988) hypothesis, we would expect the frequency of price change α to be lower in recent years than in the high-inflation period, as there is less need to change prices in a low-inflation environment. This prediction lines up with the findings of Nakamura et al. (2018): They find that about 15 percent of prices changed each month in the 1978–81 period, compared with about 10 percent per month, on average, in the 1983–2014 period. Inserting these values into Equation (17) would imply that a reduction in the frequency of price change from 15 percent per month to 10 percent per month would lead to a reduction of about 50 percent in the Phillips-curve parameter κ, assuming other parameters are unchanged.

The 50 percent reduction in κ should probably be viewed as an upper bound for comparison with our estimates, however. Our empirical work compares the period 1978 to 1996 with 1997 to 2015. That suggests that our early sample mixes periods of relatively frequent and infrequent price changes. If we take an average of Nakamura et al. (2018)'s results over the 1978 to 1996 period, that suggests that, on average, the frequency of price change was about 11 percent per month, which would imply a reduction in the slope of the Phillips curve of only about 15 percent when comparing our early and late samples.

In our estimation of the conventional hybrid New Keynesian Phillips curve in the first two columns of Table 1, the drop in the estimate of κ was considerably larger than the 15–50 percent range suggested by the results of Nakamura et al. (2018). Thus, it appears that while the microeconomic evidence on the frequency of price change suggest some reduction in the Phillips-curve slope, it cannot fully account for the reduction in κ under the assumption of model-consistent expectations. By contrast, the drop in κ when we assume that expectations are well captured by the Michigan survey, as in columns 3 and 4 of Table 1, is just below 50 percent, within the range predicted by Nakamura et al. (2018).

7. Policy Implications

We begin by pointing out some key policy implications of a flatter Phillips curve regardless of its source and then focus on implications that are unique to reduced attentiveness. A first observation concerns economic performance at the effective lower bound on nominal interest rates (ELB). Because of the ELB, central banks are limited in their ability to reduce interest rates during an economic downturn. As Kiley and Roberts (2017), among others, have noted, the ELB can cause a significant deterioration in economic performance, and is a greater concern in the low-interest-rate environment that has prevailed in recent decades. However, when the Phillips curve is flatter, the adverse effects of the ELB are reduced. That's because inflation will fall by less in an economic downturn, and if nominal interest rates are bounded by the ELB, a smaller drop in inflation will mean that real interest rates will be lower than otherwise, providing greater support to economic activity. So, in a low-interest-rate environment, a reduced sensitivity of inflation can be a blessing.

Another important implication of a flatter Phillips curve concerns the effects of forward guidance about interest rates. As Carlstrom, Fuerst, and Paustian (2015), among others, have emphasized, aggressive forms of forward guidance, in which the central bank promises to cut interest rates in the far-distant future, can have perverse effects in standard New Keynesian models; these counterintuitive implications have been called the "forward-guidance puzzle." Other things equal, a flatter Phillips curve will tend to make the forward-guidance puzzle less severe because the puzzling implications of forward guidance turn crucially on the reaction of inflation.

Carlstrom, Fuerst, and Paustian (2015) have noted that in models that weaken the assumption of strictly model-consistent expectations, the forward-guidance puzzle is less problematic. They cite in particular the sticky information model of Mankiw and Reis (2002), which they show attenuates considerably the more puzzling aspects of forward guidance. Our empirical work finds strong support for alternatives to the standard MCE model, for all measures of expectations and in all periods. Our alternative is similar in many respects to the sticky information model of Mankiw and Reis (2002) that Carlstrom, Fuerst, and Paustian (2015) consider; in earlier work, Chung, Herbst, and Kiley (2015) had also found that alternatives to MCE—in this case, the Mankiw-Reis model that Carlstrom, Fuerst, and Paustian (2015) consider—also outperformed MCE empirically. Thus, the most extreme forms of the puzzle are unlikely to be features of real-world economies. In the limit, if expectations formation is as inattentive as our estimates with the Michigan survey in our post-1996 sample suggest, the forward-guidance puzzle does not arise at all.[31]

8. Conclusion

We examine the role that changes in the attentiveness of households and professional forecasters may have played in the reduction in the sensitivity of inflation to aggregate demand in the past couple of decades. Our most dramatic results are from the Michigan survey of households, where it appears that households now pay very little attention to macroeconomic conditions in setting their inflation expectations. In contrast, there is little evidence of a reduction in attentiveness among the respondents to the Survey of Professional Forecasters. It is perhaps not surprising that professional forecasters

[31] Relatedly, Beqiraj, Bartolomeo, and Pietro (2019) show that when inflation expectations are formed adaptively, the forward-guidance puzzle does not arise.

would continue to stay appropriately attuned to economic conditions in their forecasts; after all, that is their bread and butter. But as argued by Coibion and Gorodnichenko (2015b), it is plausible that the expectations of the firms that actually set prices are closer to those of households than of professional forecasters. Simulation results suggest that the reduced attentiveness in our Michigan survey results can account for the bulk of the decline in the overall sensitivity of inflation to aggregate demand shocks in the past couple of decades—around three-fourths. The remaining shift—which is ascribed in the New Keynesian model to a reduction in the frequency of price change—is in the range predicted by the microeconomic evidence on shifts in the frequency of price adjustment documented by Nakamura et al. (2018).

It may be that the reduction in the frequency of price change and the reduction of attention paid by price setters are not entirely distinct phenomena. The Volcker disinflation set off a number of important changes in U.S. monetary policy. First and foremost, average inflation has been lower. Ball, Mankiw, and Romer (1988) predicted that lower inflation would lead to a step-down in the frequency of price setting, and the results of Nakamura et al. (2018) confirm that such a change may have occurred. In addition, low inflation has historically tended to be more stable. In the U.S. case, one reason may have been that after the Volcker disinflation, monetary policy became more focused on inflation control; see, for example, Clarida, Galí, and Gertler (2000). The greater stability of inflation may have led firms and households to dedicate less bandwidth to monitoring inflation, as predicted by Greenspan. In addition, greater inflation control may have contributed to a reduction in sunspot-type fluctuations, as discussed by Lubik and Schorfheide (2004).

References

Adam, K., and Padula, M. 2011. “Inflation Dynamics and Subjective Expectations in the United States.” *Economic Inquiry* 49 (1): 13–25.

Andrews, I., J. Stock, and L. Sun. 2019. “Weak Instruments in Instrumental Variables Regression: Theory and Practice.” *Annual Review of Economics* 11: 727–53.

Aoki, K. 2001. "Optimal Monetary Policy Responses to Relative-Price Changes." *Journal of Monetary Economics* 48 (1): 55–80.

Aruoba, S. B. 2020. "Term Structures of Inflation Expectations and Real Interest Rates." *Journal of Business and Economic Statistics* 38 (3): 542–53.

Atkeson, A., and L. E. Ohanian. 2001. "Are Phillips Curves Useful for Forecasting Inflation?" *Quarterly Review* (Federal Reserve Bank of Minneapolis) (Winter): 2–11.

Ball, L., N. G. Mankiw, and D. Romer. 1988. "The New Keynesian Economics and the Output-Inflation Trade-off." *Brookings Papers on Economic Activity* 19 (1): 1–82.

Ball, L., and S. Mazumder. 2011. "Inflation Dynamics and the Great Recession." *Brookings Papers on Economic Activity* 42 (1) (Spring): 337–405.

———. 2019. "A Phillips Curve with Anchored Expectations and Short-Term Unemployment." *Journal of Money, Credit and Banking* 51 (1): 111–37.

Beqiraj, E., G. D. Bartolomeo, and M. D. Pietro. 2019. "Beliefs Formation and the Puzzle of Forward Guidance Power." *Journal of Macroeconomics* 60 (June): 20–32.

Bils, M., and P. Klenow. 2004. "Some Evidence on the Importance of Sticky Prices." *Journal of Political Economy* 112 (5): 947–85.

Binder, C. C. 2015. "Whose Expectations Augment the Phillips Curve?" *Economics Letters* 136 (November): 35–38.

———. 2017. "Measuring Uncertainty Based on Rounding: New Method and Application to Inflation Expectations." *Journal of Monetary Economics* 90 (October): 1–12.

Blanchard, O. J. 2016. "The US Phillips Curve: Back to the 60s?" Policy Brief No. 16-1, Peterson Institute for International Economics.

Calvo, G. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics* 12 (3): 383–98.

Carlstrom, C. T., T. S. Fuerst, and M. Paustian. 2015. "Inflation and Output in New Keynesian Models with a Transient Interest Rate Peg." *Journal of Monetary Economics* 76 (November): 230–43.

Carroll, C. D. 2003. "Macroeconomic Expectations of Households and Professional Forecasters." *Quarterly Journal of Economics* 118 (1): 269–98.

Carvalho, C., and F. Nechio. 2014. "Do People Understand Monetary Policy?" *Journal of Monetary Economics* 66 (September): 108–23.

Cecchetti, S. G., and G. Debelle. 2006. "Has the Inflation Process Changed?" *Economic Policy* 21 (46): 311–52.

Cecchetti, S. G., M. E. Feroli, P. Hooper, A. Kashyap, and K. L. Schoenholtz. 2017. "Deflating Inflation Expectations: The Implications of Inflation's Simple Dynamics." Working Paper, U.S. Monetary Policy Forum.

Christelis, D., D. Georgarakos, T. Jappelli, and M. van Rooij. 2016. "Trust in the Central Bank and Inflation Expectations." DNB Working Paper No. 537, Netherlands Central Bank.

Christiano, L., M. Eichenbaum, and C. L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy* 113 (1): 1–45.

Chung, H., E. Herbst, and M. T. Kiley. 2015. "Effective Monetary Policy Strategies in New Keynesian Models: A Reexamination." In *NBER Macroeconomics Annual 2014*, Vol. 29, ed. J. A. Parker and M. Woodford, 289–344 (chapter 5). Chicago: University of Chicago Press.

Clarida, R., J. Galí, and M. Gertler. 2000. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory." *Quarterly Journal of Economics* 115 (1): 147–80.

Coibion, O., and Y. Gorodnichenko. 2015a. "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts." *American Economic Review* 105 (8): 2644–78.

———. 2015b. "Is the Phillips Curve Alive and Well after All? Inflation Expectations and the Missing Disinflation." *American Economic Journal: Macroeconomics* 7 (1): 197–232.

Coibion, O., Y. Gorodnichenko, and R. Kamdar. 2018. "The Formation of Expectations, Inflation, and the Phillips Curve." *Journal of Economic Literature* 56 (4): 1447–91.

Coibion, O., Y. Gorodnichenko, S. Kumar, and M. Pedemonte. 2018. "Inflation Expectations as a Policy Tool?" NBER Working Paper No. 24788.

Dräger, L., and M. J. Lamla. 2012. "Updating Inflation Expectations: Evidence from Micro-data." *Economics Letters* 117 (3): 807–10.

———. 2018. "Is the Anchoring of Consumers' Inflation Expectations Shaped by Inflational Experience?" CESifo Working Paper No. 7042, CESifo Group Munich.

Dräger, L., M. J. Lamla, and D. Pfajfar. 2016. "Are Survey Expectations Theory-Consistent? The Role of Central Bank Communication and News." *European Economic Review* 85 (June): 84–111.

Erceg, C. J., and A. T. Levin. 2003. "Imperfect Credibility and Inflation Persistence." *Journal of Monetary Economics* 50 (4): 915–44.

Eusepi, S., E. Moench, B. Preston, and C. Viana de Carvalho. 2019. "Anchored Inflation Expectations." CEPR Discussion Paper No. 13900.

Fuhrer, J. 2017. "Expectations as a Source of Macroeconomic Persistence: Evidence from Survey Expectations in a Dynamic Macro Model." *Journal of Monetary Economics* 86 (April): 22–35.

Gabaix, X. 2017. "A Behavioral New Keynesian Model." Discussion Paper, Harvard University.

Gagnon, E. 2009. "Price Setting during Low and High Inflation: Evidence from Mexico." *Quarterly Journal of Economics* 124 (3): 1221–63.

Galí, J., and M. Gertler. 1999. "Inflation Dynamics: A Structural Econometric Analysis." *Journal of Monetary Economics* 44 (2): 195–222.

Greenspan, A. 2001. "Transparency in Monetary Policy." Speech, Federal Reserve Bank of St. Louis, Economic Policy Conference, St. Louis, Missouri, October 11.

Hamilton, J. D. 1996. "This is What Happened to the Oil Price-Macroeconomy Relationship." *Journal of Monetary Economics* 38 (2): 215–20.

Kiley, M. T., and J. M. Roberts. 2017. "Monetary Policy in a Low Interest Rate World." *Brookings Papers on Economic Activity* (Spring): 317–96.

Lamla, M. J., D. Pfajfar, and L. Rendell. 2019. "Inflation and Deflationary Biases in Inflation Expectations." BIS Working Paper No. 789.

Lubik, T. A., and F. Schorfheide. 2004. "Testing for Indeterminacy: An Application to U.S. Monetary Policy." *American Economic Review* 94 (1): 190–217.

Mankiw, N. G., and R. Reis. 2002. "Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics* 117 (4): 1295–1328.

Mavroeidis, S., M. Plagborg-Møller, and J. H. Stock. 2014. "Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve." *Journal of Economic Literature* 52 (1): 124–88.

McKay, A., E. Nakamura, and J. Steinsson. 2016. "The Power of Forward Guidance Revisited." *American Economic Review* 106 (10): 3133–58.

———. 2017. "The Discounted Euler Equation: A Note." *Economica* 84 (336): 820–31.

Melosi, L. 2016. "Signaling Effects of Monetary Policy." Working Paper No. 2016-14, Federal Reserve Bank of Chicago.

Mertens, E., and J. M. Nason. 2018. "Inflation and Professional Forecast Dynamics: An Evaluation of Stickiness, Persistence, and Volatility." Working Paper No. 713, Bank for International Settlements.

Montiel-Olea, J. L., and C. Pflueger. 2013. "A Robust Test for Weak Instruments." *Journal of Business and Economic Statistics* 31 (3): 420–32.

Nakamura, E., and J. Steinsson. 2008. "Five Facts about Prices: A Reevaluation of Menu Cost Models." *Quarterly Journal of Economics* 123 (4): 1415–64.

Nakamura, E., J. Steinsson, P. Sun, and D. Villar. 2018. "The Elusive Costs of Inflation: Price Dispersion during the U.S. Great Inflation." *Quarterly Journal of Economics* 133 (4): 1833–1908.

Pfajfar, D., and E. Santoro. 2013. "News on Inflation and the Epidemiology of Inflation Expectations." *Journal of Money, Credit and Banking* 45 (6): 1045–67.

Pfajfar, D., and B. Zakelj. 2014. "Experimental Evidence on Inflation Expectation Formation." *Journal of Economic Dynamics and Control* 44 (July): 147–68.

Roberts, J. M. 1997. "Is Inflation Sticky?" *Journal of Monetary Economics* 39 (2): 173–96.

———. 2006. "Monetary Policy and Inflation Dynamics." *International Journal of Central Banking* 2 (3, September): 193–230.

Roth, M. 2013. "Inflation Expectations and the Stability of the Phillips Curve Relationship." Working Paper, Goethe University Frankfurt.

Sims, C. A. 2003. "Implications of Rational Inattention." *Journal of Monetary Economics* 50 (3): 665–90.

Staiger, D., and J. Stock. 1997. "Instrumental Variables Regression with Weak Instruments." *Econometrica* 65: 557–86.

Stock, J. H., and M. W. Watson. 2007. "Why Has U.S. Inflation Become Harder to Forecast?" *Journal of Money, Credit and Banking* 39 (s1): 3–33.

Woodford, M. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, NJ: Princeton University Press.