Finance and Economics Discussion Series
Divisions of Research & Statistics and Monetary Affairs
Federal Reserve Board, Washington, D.C.

The Role of Expectations in Changed Inflation Dynamics

Damjan Pfajfar and John M. Roberts

2018-062

Please cite this paper as:
Pfajfar, Damjan, and John M. Roberts (2018). "The Role of Expectations in Changed Inflation Dynamics," Finance and Economics Discussion Series 2018-062. Washington: Board of Governors of the Federal Reserve System, https://doi.org/10.17016/FEDS.2018.062.

The Role of Expectations in Changed Inflation Dynamics*

Damjan Pfajfar[§]
Federal Reserve Board

John M. Roberts[‡]
Federal Reserve Board

August 17, 2018

Abstract

The Phillips curve has been much flatter in the past twenty 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—in the context of the canonical Calvo model, firms are adjusting prices less often. The other is that the expectations of firms and households about future inflation are now less well informed by macroeconomic conditions; because expectations are important in the setting of current-period prices, inflation is therefore less sensitive to macroeconomic conditions. To distinguish between our two hypotheses, we bring to bear information on inflation expectations from surveys, which allow us to distinguish changes in the sensitivity of inflation to economic conditions conditioning on expectations from changes in the sensitivity of expectations themselves to economic conditions. We find that, with some measures, expectations are less tied to economic conditions than in the past, and thus that this reduced attentiveness can account for a significant portion of the reduction in the sensitivity of inflation to economic conditions in recent decades.

Keywords: Inflation dynamics; Phillips curve; Survey Inflation Expectations.

JEL classification: E31; E37.

*We would like to thank Fernanda Nechio and Sergey Slobodyan for their comments. We would also like to thank seminar participants at Federal Reserve Board, IAAE 2018 Annual Conference, System Macro Meetings at the SF Fed, and Bank of Russia Conference Inflation: New Insights for Central Banks for their comments. Kevin Starnes 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.

‡ *Address:* Board of Governors of the Federal Reserve System, 2000 Constitution Ave NW, Washington, DC 20551, U.S.A. *E-mail*: john.m.roberts@frb.gov. *Web:* https://www.federalreserve.gov/econres/john-m-roberts.htm.

§ *Address:* Board of Governors of the Federal Reserve System, 2000 Constitution Ave NW, Washington, DC 20551, U.S.A. *E-mail*: damjan.pfajfar@frb.gov. *Web:* https://sites.google.com/site/dpfajfar/.

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 et al. (2014), and Blanchard (2016). We explore two hypotheses about the origin of the flatter Phillips curve. One is that prices at the firm level are "stickier" than they used to be. Ball et al. (1988), for example, argued that lower inflation would lead to less-frequent adjustment of prices; because inflation has been lower in the past twenty 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 prediction 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 first 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. To distinguish between our two main hypotheses, we then bring to bear information on inflation expectations, taken from surveys. As a number of authors (Roberts, 1997; Mavroeidis et al., 2014; Fuhrer, 2017; Coibion et al., 2017) argue, survey measures of inflation expectations bring valuable additional information to the analysis of aggregate inflation. In our case, introducing direct measures of expectations helps us distinguish changes in the sensitivity of inflation to economic conditions conditional on expectations from changes in the sensitivity of expectations to economic conditions directly.

We find that, across surveys and time periods, survey measures of inflation expectations react more sluggishly than the rational benchmark would predict, in a manner similar to the epidemiological model of Carroll (2003). As we show in a simplified model, departures of expectations from perfect rationality have similar predictions to nominal rigidities (Mankiw and Reis, 2002 also make this point). Consistent with that

[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).

intuition, the sensitivity of inflation to economic conditions is estimated to be greater when we condition on survey expectations than in the model with model consistent expectations (MCE)—and is perhaps more in line with estimates on nominal rigidity from microeconomic studies.

Results on our central hypothesis are sensitive to the measure of inflation expectations. With various measures of expectations, the reduction in the sensitivity of inflation to economic activity across our two subsamples is considerably less than in the MCE case. However, it is only with the University of Michigan's survey of household inflation expectations that we also find a large reduction in attentiveness. The other measures of expectations we examine are surveys of forecasters, such as the Survey of Professional Forecasters. These measures display show some reduction in the longer-run convergence of expectations to the MCE benchmark, but not enough to account for much of the overall sensitivity of inflation to economic activity. 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 (2015) argue that household expectations may be closer to those of actual decision makers than are forecasts from economists and thus that results based on the Michigan survey should be favored. In particular, they find that the Michigan survey does a better job of accounting for the behavior of inflation and that household expectations appear to be closer to the expectations of the firms that are setting prices than are forecasts from professionals. 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.

Our use of survey information allows us to derive a direct measure of sunspots. As discussed in Lubik and Schorfheide (2004), sunspots are movements in inflation expectations that are not justified by fundamentals. In our structural modeling of the survey measures of inflation expectations, we allow for both sunspots and measurement error. 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. We find that sunspot shocks are much more important in the first part of the sample than in the second part. That result is consistent with the finding of many authors (Roberts, 2006 and citations therein) that inflation has been better controlled by the central bank in recent years; as Lubik and Schorfheide (2004) note, inadequate inflation control can be a source of sunspots.

We cross-check our findings with available 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. New data collected by Nakamura et al. (2016), 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 such the change in the price-change frequency documented by Nakamura et al. (2016). We find that while the 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 (MCE).

Mavroeidis et al. (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 as well as 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 et al. (2014), and their results suggest that we are entering treacherous waters. In the Appendix to this paper, we use single-methods akin to those of Mavroeidis et al. (2014) to assess, in a reduced-form sense, our ability to identify the key relationships of our model. We find that we may be able to obtain well-identified and precise estimates despite the warnings of Mavroeidis et al. (2014). There are two reasons. First, our focus is different than 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 et al. (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. Our findings in the Appendix confirm that result, and we provide a conjecture as to why that may be the case. The question we focus on is different. In particular, we 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 be perfectly rational, and we ask to what extent these expectations may differ from perfect rationality. In addition, our full-system estimation, in which we use information on both expectations and inflation to inform the structural relationship between economic activity and inflation, should allow us to identify separately the degree of the departure from complete rationality as well as the impact of economic activity on inflation conditional on expectations.

2. Theory

2.1. A stylized model

In this subsection, we use a very simple model to illustrate the role that imperfect expectations formation can play in determining the reduced-form relationship between inflation and aggregate economic activity. We start with a simple version of the New Keynesian Phillips curve:

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

where Δp_t is inflation and y_t is the output gap. $E_t \Delta p_{t+1}$ represents the expectations of agents setting prices. It is typically assumed that expectations are model-consistent. In that case,

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

where $M_t \Delta p_{t+1}$ represents model-consistent expectations. Suppose instead that agents only have access to a noisy signal of expectations:

$$E_t \Delta p_{t+1} = M_t \Delta p_{t+1} + \eta_t. \tag{3}$$

In that case, as discussed, for example, in Sims (2003), agents' expectations will be related to the true, model-consistent expectations by:

$$E_t \Delta p_{t+1} = \mu M_t \Delta p_{t+1}, \tag{4}$$

where $0 \leq \mu \leq 1$. That is, when agents receive only a noisy signal of expectations, their actual expectations move only partially with the ideal, model-consistent expectations. As discussed in Sims (2003), it can be costly for agents to pay attention to the future course of inflation. With greater effort, agents can improve the quality of their inflation forecasts, implicitly reducing the variance of η in Equation 3. With sufficient effort, expectations will be close to the rational ideal and the value of μ will approach 1. Gabaix (2017) posits a similar specification for imperfectly rational inflation expectations.

Greenspan (2001) argued that a consequence of low and stable inflation is that agents would then need to pay less attention to inflation, and to forecasts of inflation. In the context of our model, an implications of Greenspan's hypothesis is that in an environment of low and stable inflation, the value of μ will fall.

If we substitute Equation 4 into Equation 1, we obtain:[2]

$$\Delta p_t = \beta \mu M_t \Delta p_{t+1} + \kappa y_t + \epsilon_t. \tag{5}$$

Equation 5 can be referred to as the "discounted" New Keynesian Phillips curve, in analogy to the "discounted Euler equation" proposed by McKay et al. (2015), McKay et al. (2016) (see also Gabaix, 2017).[3]

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:

$$y_t = \rho y_{t-1} + \zeta_t. \tag{6}$$

[2] Mavroeidis et al. (2014) argue that if we take literally the microfoundations of the New Keynesian Phillips curve, it is inappropriate to make the simple substitution we make here. However, the discussion in Gabaix (2017) suggests that a specification like Equation (5) can be derived formally.

[3] Of course, the term β in Equation 1 means that expectations are already discounted to some extent in the standard version of the New Keynesian Phillips curve. The value of β, however, is tied to concepts such as the real return on capital and household's time-preference rate and is therefore not likely to differ very much from one. The term μ introduces the possibility of additional discounting.

With this assumption, Equation 5 can be solved forward as:

$$\Delta p_t = \frac{\kappa}{1 - \beta\mu\rho} y_t + \epsilon_t, \tag{7}$$

assuming ϵ is i.i.d. As can be see, both κ and μ affect the reduced-form Phillips curve slope in this simple case. In particular, smaller values of either κ—the structural slope—or μ—the attentiveness of agents—would predict a reduced sensitivity of inflation to fluctuations in output. 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.

2.2. Empirical model

In our empirical work, we nest the limited attention hypothesis about expectations formation in a broader framework that includes other potential explanations.

One such alternative hypothesis is a variant of the epidemiological model of Carroll (2003):

$$E_t \Delta p_{t+1} = (1 - \lambda) M_t \Delta p_{t+1} + \lambda E_{t-1} \Delta p_t, \tag{8}$$

Under Carroll's hypothesis, expectations adjust only gradually toward a well-informed value.[4] When agents are fully rational and learn right away about model-consistent expectations, $\lambda = 0$.

We consider two sources of error in survey expectations. As noted in the introduction, Lubik and Schorfheide (2004) suggest that an error in expectations that has implications for actual inflation can be interpreted as a sunspot. We also allow for measurement error, which will affect inflation expectations but not actual inflation. One possible source of measurement error is sampling error.[5] In addition, survey respondents may report a different number to the survey taker than they use when they

[4] 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.

[5] 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.

actually make decisions. This latter source of error could become larger when survey respondents are less attentive.

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

$$E_t\Delta p_{t+1} = \mu M_t\Delta p_{t+1} + \lambda E_{t-1}\Delta p_t + \nu_t, \tag{9}$$

$$S_t\Delta p_{t+1} = E_t\Delta p_{t+1} + u_t, \tag{10}$$

$$u_t = \rho_u u_{t-1} + \omega_t, \tag{11}$$

where $S_t\Delta p_{t+1}$ is the survey measure of expectations. Equation 9 nests the limited attention and epidemiological models, and also allows for a fundamental error, ν. Equation 10 allows for measurement error in survey measures of expectations, and the specification in Equation 11 allows that measurement error to be serially correlated.[6]

In our empirical work, we detrend inflation and inflation expectations using an estimate of long-run inflation expectations. Such detrending puts our focus on cyclical movements in inflation, which lines up with the emphasis of the theoretical models. It also allows us to exploit the greater frequency of cyclical movements, which should allow us to better identify our key parameters. 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). Because it involves low-frequency behavior, however, there is inherently less that can be said about it, given the limited time series data available and so we focus on cyclical variations, about which we have more information.[7]

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

$$\Delta p_t - \gamma\Delta p_{t-1} = \beta(E_t\Delta p_{t+1} - \gamma\Delta p_t) + \kappa y_t + \epsilon_t. \tag{12}$$

In this case, there is partial indexation to lagged inflation.

[6] 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.

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

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

$$y_t = \phi_1 y_{t-1} + \phi_2 y_{t-2} + \phi_3 \Delta p_{t-1} + \phi_4 \Delta p_{t-2} + \eta_t. \tag{13}$$

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.[8]

3. Data

Key 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 household expectations of inflation over the coming twelve 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 expectations over the coming year from these surveys. In Appendix B, 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.[9]

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

[8] 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 13 serves that role well.

[9] 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 forecasts, which are available over a longer sample. Both approaches yielded similar results; we report the results from the former method.

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.

As we discuss in Section 2.2, we subtract from our measures of inflation and year-ahead inflation expectations a measure of the inflation trend, in particular, a measure of longer-run inflation expectations from the Federal Reserve's FRB/US model.[10] For our measure of the output gap, we use the measure from the Congressional Budget Office.

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. As the results in the next two sections will illustrate, the responsiveness of inflation to fluctuations in economic activity is very different in our two subsamples.

4. VAR Results

In this section, we use some simple VARs to establish stylized facts about the evolution of the dynamics of inflation and inflation expectations.

Figure 1 illustrates the reduction in the cyclical sensitivity of inflation over time. The figure shows the response of inflation and the output gap to a shock to the output gap in a two-variable VAR.[11] Our two estimation periods are shown, 1978 to 1996 and 1997 to 2015. The top panels show the responses of output to the identified aggregate demand shock. Both periods reflect the typical hump-shaped and persistent pattern,

[10] 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, PTR is based on expectations 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.

[11] The variables in the VAR are the CBO output gap and an inflation gap that is the difference between (annualized) inflation measured by the consumer price index for all items excluding food and energy and the measure of long-term inflation expectations described in Section 3. Three lags are included in the VAR. The aggregate demand shock is identified by assuming that output affects inflation in the current period but inflation does not affect output. In the event, the contemporaneous correlation between these variables is very small so the identifying assumption has very little effect on the results.

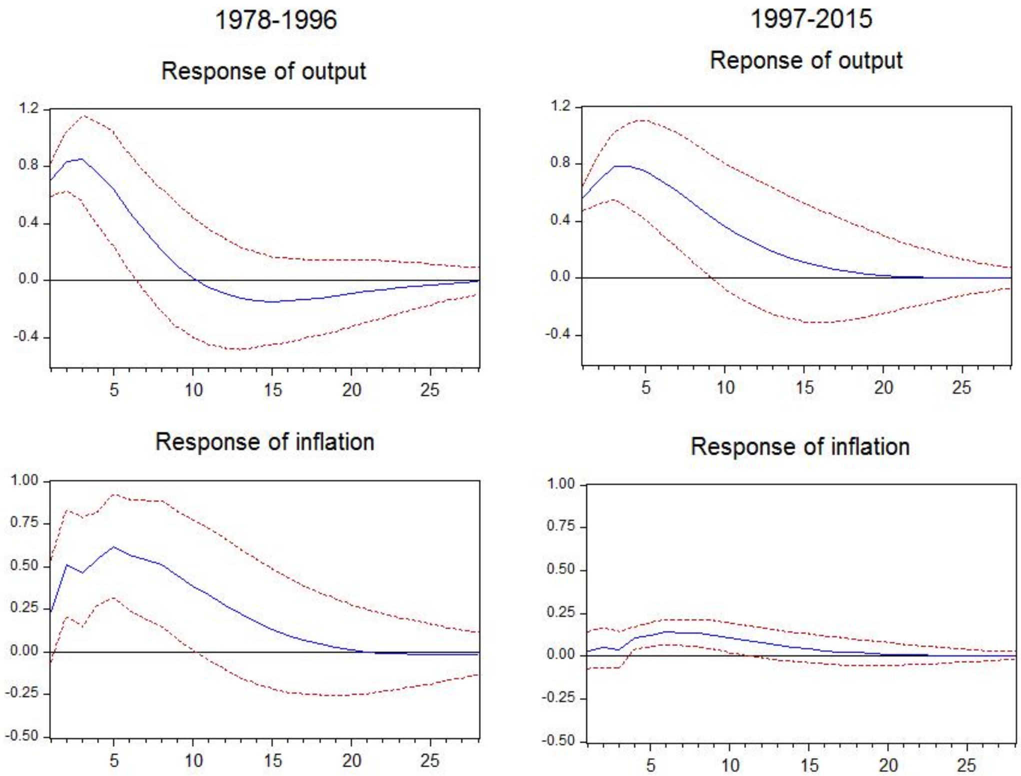

Source: Authors' calculations
Dashed red lines represent limits of 90 percent confidence intervals.

Figure 1. Impact of an aggregate demand shock in two periods

with the persistence somewhat greater in the later period. The bottom panels show the responses of inflation. As might be expected in response to an aggregate demand shock, inflation rises in both periods. But the increase in inflation is several times greater in the earlier period: Inflation has become much less sensitive to movements in aggregate demand.

Figure 2 adds the Michigan survey's measure of household inflation expectations over the coming year to the model.[12] The responses of the output gap and inflation are similar to those in Figure 1. In the early period, expectations reflect a large response to the aggregate demand shock; in the latter period, the responses are more muted. Note that with only this information in hand, it is difficult to determine which of our two

[12] As with inflation, we subtract from the Michigan survey our measure of long-run inflation expectations. For identification, we assume that neither output nor inflation affect inflation expectations in the current period. This identification is consistent with the view that expectations are formed on the basis of published inflation and output data, which are only available with a lag.

hypotheses may account for this change in behavior: It could be that a more-muted response of expectations to aggregate conditions could lead to a smaller increase in actual inflation; alternatively, the formation of inflation expectations may have adapted to greater structural nominal rigidity. We attempt to disentangle these two alternatives in our structural estimation.

5. Identification

Mavroeidis et al. (2014) emphasize that it may be difficult to estimate the parameters of the New Keynesian Phillips curve with much confidence. Since that's what we're attempting to do, we need to address the issues raised by Mavroeidis et al. (2014). In our context, we seek to estimate the parameters of the New Keynesian Phillips curve, γ and κ, as well as the parameters of the process for inflation expectations, μ and λ, a closely related problem. One of the key issues raised by Mavroeidis et al. (2014) is the lack of instruments—in particular, the lack of instruments for expected inflation. In the Appendix, we use single-equation, reduced-form analysis to explore the extent to which we can be confident we have good instruments for future inflation.

Our Appendix analysis clarifies that there is a strong reduced-form relationship between inflation, the output gap, and lagged inflation, especially in our early, pre-1997, sample. That suggests that we will have valid instruments for the model-consistent component of expectations in our specification of the surveys. We also find that the strong reduced-form results relating inflation to the output gap can become weaker when we condition on the survey measure of inflation expectations. Because we find that the survey itself is strongly related to the output gap, this result shouldn't be too surprising: Much of the correlation between the output gap and inflation is already captured by the survey. This result is related to the structural interpretation of the simple reduced-form Phillips curve from Section 2.1: When aggregate demand is the main driver of inflation, the relationship between inflation and the output gap cannot distinguish both the role of aggregate demand in driving expectations and in driving inflation conditional on expectations.

In our view, this interpretation may help explain the results of Mavroeidis et al. (2014), who found that they could not successfully estimate both a coefficient on the lead of inflation and the coefficient on a measure of economic activity: Under our interpretation, that's because economic activity is itself the best instrument for future

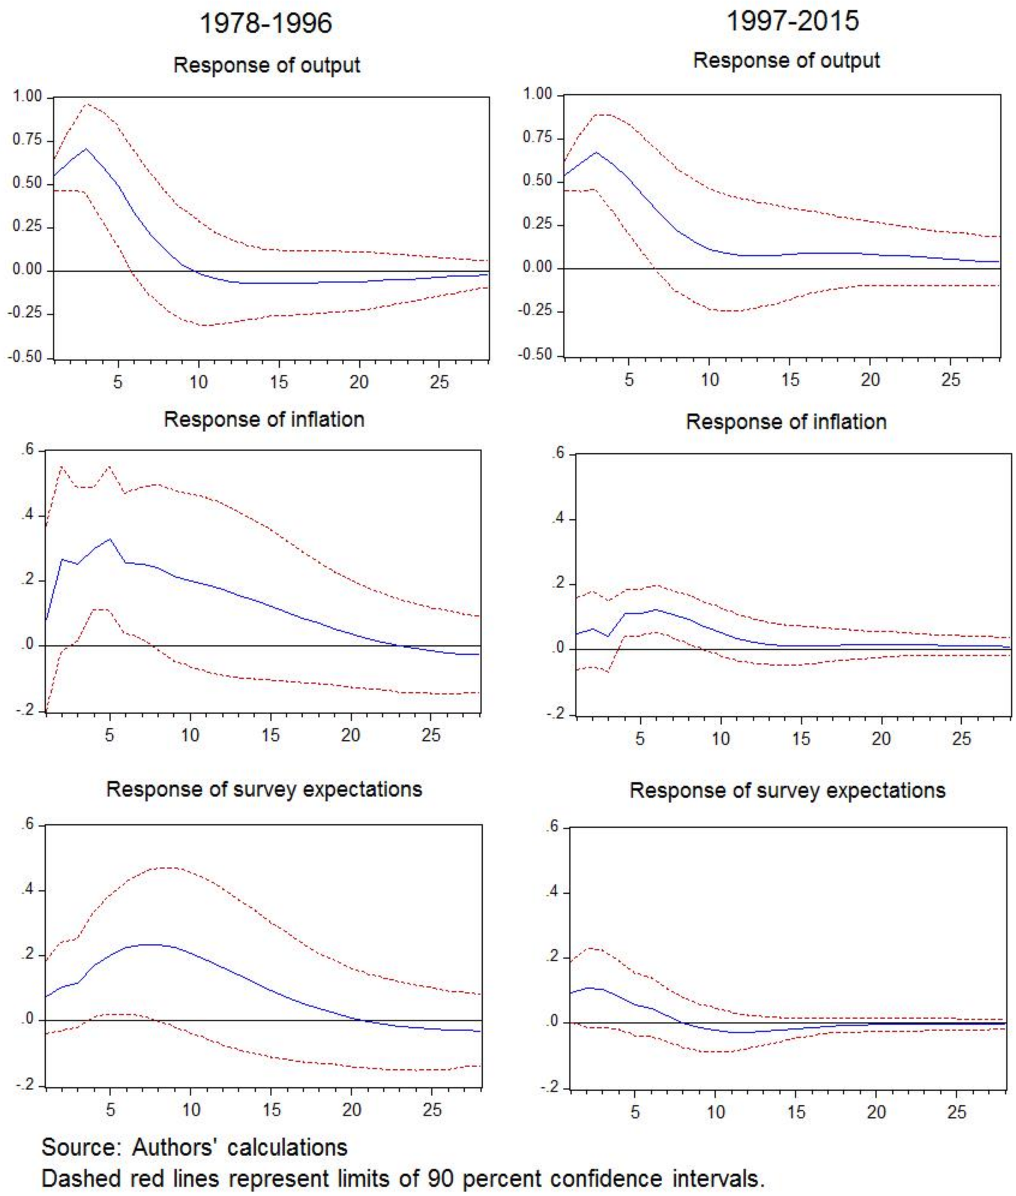

Figure 2. Adding inflation expectations

inflation; thus, it is not possible to identify separately the two channels through which economic activity is affecting inflation. From the perspective of our stylized model, the relationship between inflation and the output gap cannot provide estimates of both κ and μ. As both our work and that of Mavroeidis et al. (2014) indicates, introducing survey information does not necessarily help identify the coefficient on economic activity in a single-equation model. However, our full-information estimation may allow us to exploit the impact of aggregate demand on the survey measure of expectations to help us identify κ; according to the theory underlying our model, aggregate demand should only affect expectations because of its affect on inflation itself.

To be clear, we are not claiming to have resolved the tension emphasized by Mavroeidis et al. (2014). Mavroeidis et al. (2014) ask whether expectations belong in a structural model of inflation. For the reasons just noted, we agree with Mavroeidis et al. (2014) that that question cannot be answered. We are asking a different question. In particular, we assume that expectations belong in the structural model of inflation (as in our Equation 12). But we do not require those expectations to be perfectly rational, and we estimate the extent of the deviation from perfect rationality.

6. System Estimation Results

In this section, we turn to estimates of the full system of equations outlined in Section 2.2. We compare estimates of the hybrid New Keynesian Phillips curve under two hypotheses about expectations: the model-consistent expectations (MCE) assumption that is common in the literature and our model of expectations formation that relaxes the MCE assumption. In the latter, we use survey measures of expectations to inform the estimation.

6.1. Model estimates: MCE

Table 1 presents estimates of the system of equations consisting of the hybrid New Keynesian Phillips curve, Equation 12, and the reduced-form output-gap equation, Equation 13, under the assumption of fully model-consistent expectations. Column 1 shows results over the 1978-1996 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 roughly eight. The degree of indexation, γ, is also notably smaller in the latter sample.

	(1) MCE 1978-1996	(2) MCE 1997-2015
γ	0.689	0.287
	[0.662, 0.716]	[0.260, 0.314]
κ	0.0546	0.00710
	[0.0467, 0.0625]	[0.00618, 0.00802]
σ_ϵ	1.9181	0.5148
	[0.1847, 10.3868]	[0.0438, 11.7534]
ϕ_1	1.22	1.28
	[1.19, 1.25]	[1.25, 1.30]
ϕ_2	-0.285	-0.280
	[-0.314, -0.255]	[-0.307,-0.252]
ϕ_3	-0.0790	-0.0566
	[-0.0904, -0.0676]	[-0.0876, -0.0256]
ϕ_4	-0.0727	-0.204
	[-0.086, -0.0594]	[-0.235, -0.173]
σ_η	0.6976	0.554
	[0.0577, 12.0895]	[0.0449, 12.3499]

Table 1: Estimates of model with quarterly data, model-consistent expectations are assumed Estimated using Bayesian methods; the priors are described in Appendix C.

The bottom rows of the table show results for the reduced-form process for the output gap, Equation 13. The parameters ϕ_1 and ϕ_2 suggest that the process in both periods has the "hump-shaped" pattern typical of output impulse responses (see figures 1 and 2), with a coefficient greater than one on the first lag of the gap and a negative coefficient on the second lag. ϕ_3 and ϕ_4 show the sensitivity of output gap to lagged inflation. In each sample period, the sum of the two coefficients is negative, as expected.

6.2. Model estimates: Imperfectly rational expectations

Table 2 presents results for the model we introduced in Section 2.2, in which the assumption of model-consistent expectations is relaxed and survey expectations are added as an observable. Recall from Equation 9 that the model of expectations has several key features: It allows expectations to react less than predicted by the MCE hypothesis to incoming information ($\mu < 1$); it allows for gradual adjustment of expectations to

	(1) Michigan 1978-1996	(2) Michigan 1997-2015	(3) SPF CPI 1978-1996	(4) SPF CPI 1997-2015
γ	0.291	0.484	0.476	0.371
	[0.086, 0.500]	[0.269, 0.714]	[0.290, 0.666]	[0.158, 0.571]
κ	0.209	0.128	0.213	0.093
	[0.0979, 0.309]	[0.0620, 0.192]	[0.106, 0.321]	[0.0502, 0.135]
σ_ϵ	1.810	0.7346	2.223	0.677
	[1.403 2.212]	[0.567, 0.902]	[1.849, 2.620]	[0.5325, 0.8184]
λ	0.762	0.610	0.701	0.563
	[0.665, 0.857]	[0.406, 0.807]	[0.602, 0.797]	[0.443, 0.686]
μ	0.207	0.0555	0.199	0.247
	[0.104, 0.309]	[-0.0953, 0.208]	[0.122, 0.267]	[0.152, 0.339]
σ_ν	0.340	0.1928	0.263	0.062
	[0.215 0.460]	[0.0801, 0.303]	[0.163, 0.368]	[0.023, 0.102]
ρ	0.802	0.612	0.481	0.283
	[0.430, 1.00]	[0.406, 0.822]	[1e-04, 0.889]	[1e-04, 0.500]
σ_ω	0.287	0.386	0.251	0.143
	[0.135 0.431]	[0.305, 0.470]	[0.127, 0.378]	[0.110, 0.175]
ϕ_1	1.22	1.32	1.21	1.34
	[1.03, 1.42]	[1.15, 1.49]	[1.04, 1.40]	[1.17, 1.51]
ϕ_2	-0.293	-0.330	-0.267	-0.353
	[-0.485, -0.0916]	[-0.512, -0.157]	[-0.448, -0.089]	[-0.529, -0.176]
ϕ_3	-0.0883	-0.0444	-0.0984	-0.0761
	[-0.168, -0.0078]	[-0.261, 0.175]	[-0.169, -0.028]	[-0.288, 0.142]
ϕ_4	-0.0522	-0.190	-0.051	-0.149
	[-0.144, 0.0398]	[-0.404, 0.0275]	[-0.127, 0.029]	[-0.363, 0.0685]
σ_η	0.736	0.579	0.656	0.581
	[0.634 0.840]	[0.499, 0.657]	[0.565, 0.746]	[0.502, 0.661]

Table 2: Estimated using Bayesian methods; the priors are described in Appendix C.

their fully model-consistent level ($\lambda > 0$); 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 else as shocks to correctly-measured expectations, which can go on to affect actual inflation ($\sigma_\eta > 0$).

Results for the Michigan survey are in columns (1) and (2). Focusing first on the estimates for inflation expectations, the results suggest considerable deviation from purely model-consistent expectations in both samples. In particular, μ is considerably smaller than one. The results in the early sample provide some support for Carroll (2003)'s epidemiological model. In particular, Carroll (2003)'s model suggests that the sum of μ and λ should equal one: Although households do not react immediately to new information about future inflation, under the epidemiological interpretation, the knowledge would eventually spread. The evidence for Carroll (2003)'s hypothesis is quite strong in the early sample, when the sum of the coefficients is 0.97. In the later sample, however, the sum drops to 0.67, suggesting less of a tendency for expectations to converge to the model consistent expectations than in the earlier sample. Thus, in the latter sample, news would never be anticipated to spread fully among households.

Both μ and λ are smaller in the latter sample. In the early sample, μ is significantly greater than zero, while in the latter sample, it is not. In the early sample, there is also stronger evidence of a gradual diffusion of knowledge about future inflation among the public, with λ taking on a larger value than in the later sample.

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 is considerably smaller than in the MCE case, with the estimate of κ dropping from 0.21 to 0.13, a decline of about 40 percent, compared with a decline of more than 85 percent in Table 1.

On our preferred interpretation, a key reason for the smaller decline in the slope coefficient when the MCE assumption is relaxed is that there has been a sharp drop in the degree of rationality—that is, the estimates for the expectations process suggest that households pay much less attention to the fundamentals in forming expectations than was the case in the earlier period. Coupled with our finding that the degree of inattention is greater in that latter sample, 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 7.1, the decline in rationality predicts about as great a reduction in the overall sensitivity of inflation to economic activity as the reduction in κ.

The results suggest a significant degree of inflation persistence, with $\gamma > 0$ in both periods. In the early sample, $\gamma = 0.29$; it rises to 0.48 in the latter sample, although there is considerable overlap in the confidence intervals across the two samples.

Both measurement error and fundamental expectational shocks are important sources of variation in the Michigan survey. The structural shock to expectations is statistically significant in both subsamples. However, it is much less important in the latter sample, with a standard deviation that is about half as large. As well, λ is smaller in the latter sample, so that any given shock will be carried forward with much less persistence in the latter sample. Measurement error is large and statistically significant in both samples; it also displays considerable serial persistence, especially in the early sample. Serial persistence of the measurement error is somewhat smaller in the later sample.

As discussed earlier, we interpret our structural expectations shocks as an indication of sunspot equilibria. As discussed in Lubik and Schorfheide (2004), sunspot equilibria 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 et al., 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 (3) and (4) present results for the CPI version of the SPF. For the Phillips curve, the estimates are broadly similar to those for the Michigan survey: The point estimates of κ and γ are similar, and the drop in the slope coefficient κ across the two samples is similar, with κ falling by a bit more than half. For expectations, as with the Michigan survey, the hypothesis of rational expectations is strongly rejected, with $\mu = 0.2$ in the early sample. Also like the Michigan survey, the sum of λ and μ is fairly close to one in the early sample ($= 0.92$) and falls in the latter sample.

One key difference with the Michigan results is the evolution of μ over time: For the SPF, the point estimate of μ actually rose somewhat in the latter period, in contrast to the sharp drop for the Michigan survey. That would suggest 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 considerably less attention in the later sample. We return to an interpretation of this finding in section 6.3.

As with the Michigan survey, there is evidence of both measurement error and fundamental shocks to expectations. The results suggest, however, that measurement error is considerably less important for the SPF than for the Michigan survey: In the early sample, the point estimate of σ_ω is similar across the two surveys. However, the persistence of measurement error, ρ, is much smaller for the SPF, indicating that the overall contribution of measurement error to the variability of the SPF is considerably smaller. It is smaller still in the latter sample, with point estimates of both ρ and σ_ω dropping off. The fundamental expectations shock, ν, is also less important for the SPF than for the Michigan survey and, as with measurement error, it is much 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 Table 1—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 et al. (2014) also find that estimates of κ are larger when expectations are proxied using surveys. As discussed in the previous section, the larger value of κ is consistent with the predictions of the simple model in Section 2: When agents are less focused on the future, expectations of future inflation move less for any given change in current-period marginal cost. To explain the same current-period change in inflation, the model ascribes a larger role to current-period marginal cost.

6.3. A case for preferring the Michigan-based results

Coibion et al. (2017) 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. Thus, Michigan survey expectations may be closer to the expectations of decision makers in the economy than are the expectations of professional forecasters. Coibion et al. (2017) go on to cite earlier work by two of them (Coibion and Gorodnichenko, 2015) 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 2 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 less than 40 percent, consistent with the view that a larger decline in rationality 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 et al., 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.

6.4. 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 3 present a formal variance decomposition of the Michigan survey in the two samples, based on the results in columns 1 and 2 of Table 2. In the early sample, the variation of the Michigan survey is importantly influenced by the business cycle: The cycle accounts for 31 percent of the variability of the Michigan survey in the earlier sample. The fundamental shock to expectations also accounts for a substantial portion of the variability of the Michigan survey while the measurement error shock accounts for 14 percent of the total variability of the Michigan survey. By contrast, measurement error accounts for 77 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 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 almost 95 percent).

Columns 3 and 4 of Table 3 show the model's view of the contributions of the model's various shocks to the variance of core CPI inflation in the two subsamples. The business-cycle shock, η, accounts for about a quarter 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, again falling by almost 95 percent. So in the latter sample, the business cycle is explaining 27 percent of a very small number. In the early sample, the structural shock to inflation expectations accounts for almost 20 percent of the variability of inflation. Thus, "sunspots" make a non-trivial contribution to the variability of inflation in this period. By assumption, survey measurement error makes no contribution to the variability of actual inflation.

	(1) Michigan 1978-1996	(2) Michigan 1997-2015	(3) Core CPI 1978-1996	(4) Core CPI 1997-2015
ϵ	2	0	60	62
η	31	1	23	27
ν	53	11	17	11
ω	14	77	0	0
Memo: Variance	5.4	.3	7.8	.5
Standard deviation	2.3	.6	2.8	.7

Table 3: Variance decomposition for the Michigan Survey and core CPI inflation based on parameter estimates in Table 2, columns 1 and 2. Table entries are the percent of variance of each variable explain by each of the model's shocks.

	(1) SPF CPI 1978-1996	(2) SPF CPI 1997-2015	(3) Core CPI 1978-1996	(4) Core CPI 1997-2015
ϵ	10	8	80	62
η	29	58	14	37
ν	48	11	6	2
ω	13	23	0	0
Memo: Variance	.62	.10	3.7	.52
Standard deviation	.79	.31	1.9	.72

Table 4: Variance decomposition for the SPF and core CPI inflation based on parameter estimates in Table 2, columns 3 and 4. Table entries are the percent of variance of each variable explain by each of the model's shocks.

Table 4 shows variance decompositions based on the model estimates for the SPF, from columns 3 and 4 of Table 2. The results are qualitatively similar to those for the Michigan survey: The overall variability of both inflation and inflation expectations is much lower in the latter sample. For expectations, that reduction in variability comes much more heavily from measurement error. As with the Michigan survey, the shock to the cycle—η—accounts for a substantial portion of the variability of inflation in both samples.

7. Structural Interpretation

7.1. Implications for inflation dynamics

Figures 1 and 2 illustrated the fall in the sensitivity of inflation to economic activity using impulse responses from a reduced-form model. Here, we explore the ability of our model of imperfectly rational expectations to account for these changes in inflation dynamics.

Figure 3 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 2, which used data from the 1978-1996 period data. In the dashed red line, the estimates for the equations for inflation and inflation expectations use the results from column 4 of Table 2, which are based on the 1997-2015 sample. To isolate the effects of the change in inflation dynamics, these results use the same output gap equation as in the solid blue simulations. Consistent with the results in Figures 1 and 2, the effects of the aggregate-demand shock are much larger in the earlier sample; the peak effect is about three times greater.

To isolate the contribution of the change in inflation dynamics, the dot-dashed green line shows a simulation using the estimated expectations-formation equation from the latter period along with the inflation equation from the early period. In this simulation, the change in the inflation expectations process explains a bit more than half of the reduced effect of the aggregate demand shock on inflation, with shifts in the structural Phillips curve parameters, κ and γ, accounting for the rest.

7.2. Micro evidence on the frequency of price change

As mentioned in the introduction, Nakamura et al. (2016) 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):

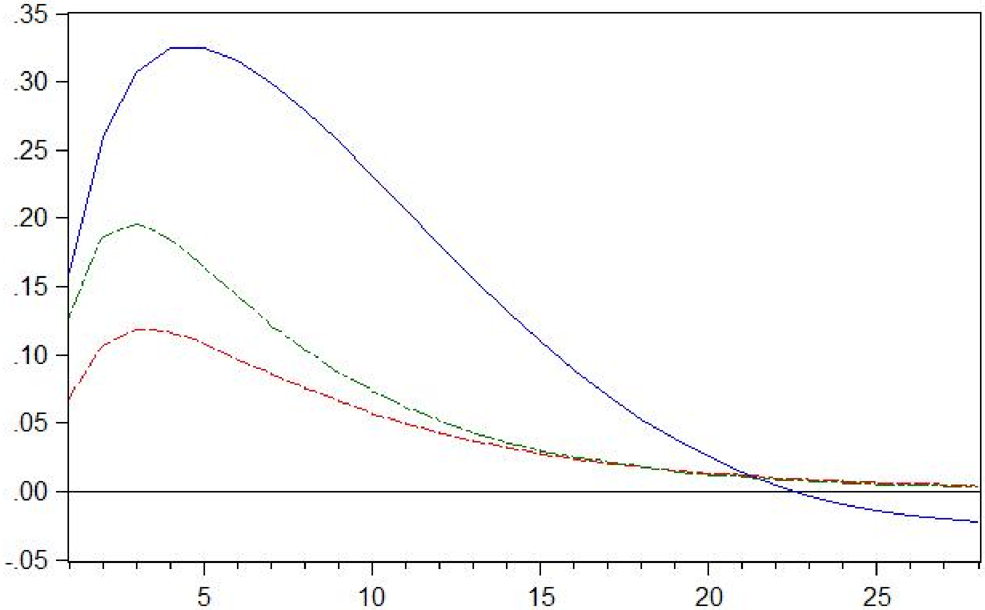

Figure 3. Impact of an aggregate demand shock in the structural model. Solid blue line presents results from model in column 3, Table 2, using the 1978-1996 period and the Michigan Survey. Red dashed line presents results that substitute the inflation and inflation expectations results from column 4 of Table 2, using the 1997-2015 period and the Michigan Survey. In the green dot-dash line, only the inflation expectations equation estimates from column 4 are used. See text for further details.

$$\kappa = \frac{(1-\alpha)(1-\beta\alpha)}{\alpha}\zeta, \tag{14}$$

where, recalling Equation 1, κ is the slope of the New Keynesian Phillips curve:

$$\Delta p_t = \beta E_t \Delta p_{t+1} + (1-\beta)\bar{\pi} + \kappa y_t + \epsilon_t. \tag{15}$$

According to the Ball et al. (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. (2016): They find that about 15 percent of prices changed each month in the 1978-to-1981 period, compared with about 10 percent per month, on average, in the 1983-to-2014 period. Inserting these values into Equation 14 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. (2016)'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 Table 1, the drop in the estimate of κ was considerably larger than the 15-to-50 percent range suggested by the results of Nakamura et al. (2016). 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 2, is only about 40 percent, within the range predicted by Nakamura et al. (2016).

8. Sensitivity: Alternative Processes for Inflation Expectations

In this section, we provide some sensitivity with respect to the inflation expectations process by relaxing Equation 9 and allowing a more general process of expectation formation:

$$E_t \Delta p_{t+1} = \mu M_t \Delta p_{t+1} + \lambda E_{t-1} \Delta p_t + \theta_3 p_{t-1} + \theta_4 y_t + \nu_t. \tag{16}$$

This process for inflation expectations 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. Equation 16 thus also nests additional processes for expectation formation, such as naive expectations and adaptive expectations. The implications of "'rationality"' in the setting of expectations are now expanded to include $\theta_3 = \theta_4 = 0$ in addition to $\lambda = 0$ and $\mu = 1$.

Table 5 presents the results. The estimates of the key parameters of the structural Phillips curve, γ and κ, are virtually identical to those in Table 2. For the survey models, as before, we can strongly reject rationality for both surveys and for both time periods.[13] Furthermore, we can also reject that expectations are formed using simple rules—like naive expectations or a simple AR(1) process—as $\theta_3 \neq 1$ and $\mu \neq \lambda \neq \theta_4 \neq 0$. As in Carroll (2003) and Pfajfar and Santoro (2013), the coefficient on the past inflation rate, θ_3, is not statistically significant for the Michigan survey. For the SPF, θ_3 is not significant in the early sample but is in the latter period. Results regarding variances are almost identical to those in Section 6.2—in particular, the variance of the shock to the Phillips curve significantly decreases in the later sample.

In Table 5, μ is significantly lower with this expectation formation mechanism than in Table 2 and in most cases is not significantly different from zero. At the same time, θ_4 is significantly positive in two of the four cases. These results suggest an alternative interpretation of Equation 16—namely, as a Phillips curve relationship that allows economic activity to influence the process for inflation indirectly through inflation expectations. This alternative interpretation allows us to study the relative importance of the direct and indirect channels through which the output gap may affect inflation.

For the Michigan survey, θ_4 is smaller in the later sample compared to the earlier sample. As in Table 2, μ also declines, and in the latter sample, neither θ_4 not μ is statistically significant. We view this result as broadly consistent with the finding in Table 2, namely, that in the latter sample, respondents to the Michigan survey paid less attention to economic fundamentals than in the pre-1997 period, and indeed may not have paid meaningful attention to fundamentals at all. For the SPF, we can see that θ_4 actually increases slightly, suggesting that, in contrast to the households that respond to the Michigan survey, participants in the SPF continue use the information on overall economic conditions to forecast inflation.[14] We also check whether oil price shocks play an important role in forming inflation expectations. Thus, we further augment Equation

[13]This result is consistent with Nunes (2009).

[14]Dräger et al. (2016) show that only about one-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. 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.

	(3) Michigan 1978-1996	(4) Michigan 1997-2015	(5) SPF CPI 1978-1996	(6) SPF CPI 1997-2015
γ	0.343	0.506	0.486	0.450
	[0.124, 0.543]	[0.286, 0.727]	[0.292, 0.681]	[0.222, 0.669]
κ	0.147	0.0969	0.185	0.0494
	[0.0231, 0.273]	[-0.0034, 0.197]	[0.052, 0.314]	[-0.0076, 0.104]
σ_ϵ	1.918	0.739	2.246	0.748
	[1.497, 2.333]	[0.573, 0.902]	[1.841, 2.650]	[0.579, 0.903]
λ	0.866	0.631	0.724	0.467
	[0.714, 1.01]	[0.386, 0.873]	[0.584, 0.870]	[0.320, 0.620]
μ	0.0807	0.0414	0.161	-0.0209
	[-0.0654, 0.225]	[-0.167, 0.251]	[0.045, 0.275]	[-0.157, 0.119]
θ_3	-0.006	-0.0924	0.0047	0.0937
	[-0.0803, 0.0682]	[-0.244, 0.0593]	[-0.046, 0.054]	[0.0319, 0.153]
θ_4	0.0661	0.032	0.020	0.0495
	[0.0096, 0.121]	[-0.025, 0.0903]	[-0.022, 0.063]	[0.0256, 0.0739]
σ_ν	0.383	0.204	0.277	0.094
	[0.261, 0.521]	[0.0857, 0.324]	[0.163, 0.392]	[0.040, 0.146]
ρ	0.653	0.651	0.495	0.280
	[0.167, 1.00]	[0.415, 0.913]	[0.250, 1.00]	[1e-04, 0.517]
σ_ω	0.269	0.383	0.250	0.113
	[0.118, 0.411]	[0.296, 0.472]	[0.118, 0.376]	[0.067, 0.158]
ϕ_1	1.21	1.32	1.21	1.31
	[1.03, 1.41]	[1.15, 1.49]	[1.03, 1.39]	[1.15, 1.49]
ϕ_2	-0.285	-0.330	-0.262	-0.326
	[-0.478, -0.0881]	[-0.510, -0.156]	[-0.435, -0.077]	[-0.496, -0.146]
ϕ_3	-0.085	-0.0482	-0.098	-0.049
	[-0.166, -0.0032]	[-0.261, 0.170]	[-0.168, 0.027]	[-0.268, 0.161]
ϕ_4	-0.054	-0.189	-0.053	-0.174
	[-0.151, 0.0382]	[-0.400, 0.0321]	[-0.132, 0.024]	[-0.388, 0.0412]
σ_η	0.737	0.580	0.656	0.580
	[0.633, 0.833]	[0.499, 0.657]	[0.563, 0.747]	[0.500, 0.657]

Table 5: Estimated using Bayesian methods; the priors are described in Appendix C.

16 by including oil price shocks calculated as in Hamilton (1996). Our results indicate that oil price shocks are insignificant in all regressions.[15]

[15]The data on oil price shocks are obtained from the on-line files of Coibion and Gorodnichenko (2012). Results are available upon request from the authors.

9. 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. Changes in expectations formation were less marked for the the Survey of Professional Forecasters. It is perhaps not surprising that professional forecasters 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 (2015), 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 more than half the decline in the overall sensitivity of inflation to aggregate demand shocks in the past couple of decades. 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. (2016).

It may be that the reduction in the frequency of price change and the reduction of attention paid by price setters are not completely 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 et al. (1988) predicted that lower inflation would lead to a step-down in the frequency of price-setting, and the results of Nakamura et al. (2016) 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 et al. (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

Aoki, K. (2001). Optimal monetary policy responses to relative-price changes. *Journal of Monetary Economics*, 48(1):55–80.

Atkeson, A. and Ohanian, L. E. (2001). Are Phillips curves useful for forecasting inflation? *Quarterly Review*, pages 2–11.

Ball, L., Mankiw, N. G., and Romer, D. (1988). The new keynsesian economics and the output-inflation trade-off. *Brookings Papers on Economic Activity*, 19(1):1–82.

Bils, M. and Klenow, P. (2004). Some evidence on the importance of sticky prices. *Journal of Political Economy*, 112(5):947–985.

Blanchard, O. J. (2016). The US Phillips Curve: Back to the 60s? Policy Briefs PB16-1, Peterson Institute for International Economics.

Carroll, C. D. (2003). Macroeconomic expectations of households and professional forecasters. *The Quarterly Journal of Economics*, 118(1):269–298.

Carvalho, C. and Nechio, F. (2014). Do people understand monetary policy? *Journal of Monetary Economics*, 66:108–123.

Cecchetti, S. G., Feroli, M. E., Hooper, P., Kashyap, A., and Schoenholtz, K. L. (2017). Deflating inflation expectations: The implications of inflation's simple dynamics. Working paper, U.S. Monetary Policy Forum.

Clarida, R., Galı, J., and Gertler, M. (2000). Monetary policy rules and macroeconomic stability: Evidence and some theory. *Quarterly Journal of Economics*, 115(1):147–80.

Coibion, O. and Gorodnichenko, Y. (2012). What can survey forecasts tell us about information rigidities? *Journal of Political Economy*, 120(1):116 – 159.

Coibion, O. and Gorodnichenko, Y. (2015). 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., Gorodnichenko, Y., and Kamdar, R. (2017). The Formation of Expectations, Inflation and the Phillips Curve. NBER Working Papers 23304, National Bureau of Economic Research, Inc.

Dräger, L., Lamla, M. J., and Pfajfar, D. (2016). Are survey expectations theory-consistent? The role of central bank communication and news. *European Economic Review*, 85(C):84–111.

Erceg, C. J. and Levin, A. T. (2003). Imperfect credibility and inflation persistence. *Journal of Monetary Economics*, 50(4):915–944.

Fuhrer, J. (2017). Expectations as a source of macroeconomic persistence: Evidence from survey expectations in a dynamic macro model. *Journal of Monetary Economics*, 86:22–35.

Gabaix, X. (2017). A behavioral new keynesian model. Discussion Paper N.A., Harvard University.

Gagnon, E. (2009). Price setting during low and high inflation: Evidence from mexico. *The Quarterly Journal of Economics*, 124(3):1221–1263.

Greenspan, A. (2001). Transparency in monetary policy. Technical report, Speech at the 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):215–220.

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

Mankiw, N. G. and Reis, R. (2002). Sticky information versus sticky prices: A proposal to replace the new keynesian phillips curve. *The Quarterly Journal of Economics*, 117(4):1295–1328.

Mavroeidis, S., Plagborg-Møller, M., and Stock, J. H. (2014). Empirical Evidence on Inflation Expectations in the New Keynesian Phillips Curve. *Journal of Economic Literature*, 52(1):124–188.

McKay, A., Nakamura, E., and Steinsson, J. (2015). The Power of Forward Guidance Revisited. NBER Working Papers 20882, National Bureau of Economic Research, Inc.

McKay, A., Nakamura, E., and Steinsson, J. (2016). The Discounted Euler Equation: A Note. NBER Working Papers 22129, National Bureau of Economic Research, Inc.

Melosi, L. (2016). Signaling effects of monetary policy. Working Paper 2016-14, Federal Reserve Bank of Chicago.

Nakamura, E. and Steinsson, J. (2008). Five facts about prices: A reevaluation of menu cost models. *The Quarterly Journal of Economics*, 123(4):1415–1464.

Nakamura, E., Steinsson, J., Sun, P., and Villar, D. (2016). The elusive costs of inflation: Price dispersion during the u.s. great inflation. NBER Working Papers 22505, National Bureau of Economic Research, Inc.

Nunes, R. (2009). On the epidemiological microfoundations of sticky information. *Oxford Bulletin of Economics and Statistics*, 71(5):643–657.

Pfajfar, D. and Santoro, E. (2013). News on inflation and the epidemiology of inflation expectations. *Journal of Money, Credit and Banking*, 45(6):1045–1067.

Roberts, J. M. (1997). Is inflation sticky? *Journal of Monetary Economics*, 39(2):173–196.

Roberts, J. M. (1998). Inflation expectations and the transmission of monetary policy. Finance and Economics Discussion Series 1998-43, Board of Governors of the Federal Reserve System (U.S.).

Roberts, J. M. (2006). Monetary Policy and Inflation Dynamics. *International Journal of Central Banking*, 2(3).

Sims, C. A. (2003). Implications of rational inattention. *Journal of Monetary Economics*, 50(3):665–690.

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

Tulip, P. (2009). Has the Economy Become More Predictable? Changes in Greenbook Forecast Accuracy. *Journal of Money, Credit and Banking*, 41(6):1217–1231.

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

Appendix A. Reduced-Form Results

This Appendix presents reduced-form results that form the basis of the assessment in Section 5 that we likely do have sufficient econometric power to identify the key parameters of our model.

A.1. Reduced-form results for the early sample

Table 6 shows reduced-form results for the early sample, 1978-1996. The first column shows results for a simple reduced-form Phillips curve, relating inflation to the current-period output gap and lagged inflation.

As we saw in Section 2.1, this specification would be structural if the output gap follows an AR(1) process and expectations are formed according to the simplest version of the rational inattention model. While these assumptions are extreme, and are dropped in our full-information estimation, they nonetheless give us a helpful point of departure for interpreting the model's coefficients. In particular, the coefficient on the output gap in Equation 7 is $\frac{\kappa}{(1-\beta\rho\mu)}$. A typical value for the discount factor β is 0.98; a simple regression of the CBO output gap on one lag of itself implies $\rho = 0.92$. Under the extreme assumption that $\mu = 0$—that is, expectations are not at all forward-looking—the estimate of κ is simply the coefficient on the output gap, 0.334. Under the other extreme assumption of $\mu = 1$—that is, perfectly rational expectations—$\kappa = 0.039$. Thus, the degree of rationality can have a large effect on the estimate of κ.

Column 2 assesses the ability to forecast inflation using lags of the output gap and inflation. Again, the $\bar{R}^2$ is quite high and the t-statistics on lagged output and inflation are also high. These results are consistent with the findings of Atkeson and Ohanian (2001) and of Tulip (2009), who find that over this period, there was an important predictable component of inflation.

In columns 3 and 4, we switch dependent variables and assess whether there is sufficient information to estimate our model of survey expectations. This exercise will allow us to address one of the key concerns emphasized by Mavroeidis et al. (2014), the lack of valid instruments for future inflation. In particular, in the empirical model of survey expectations that we lay out in equations 9 and 10, it is important that we have good instruments for future inflation. The results in column 2 indicate that current-period output and inflation should provide such instruments. That would suggest that in a reduced-form model of survey expectations, these two instruments ought to have im-

Dependent var:	(1) Δp_t	(2) Δp_t	(3) $S_t\Delta p_{t+1}$	(4) $S_t\Delta p_{t+1}$	(5) Δp_t	(6) Δp_t	(7) Δp_t
Est. method:	OLS	OLS	OLS	IV	OLS	IV	IV
y_t	.334	–	.104	–	.115	.095	.080
	(.095)		(.032)		(.095)	(.099)	(.080)
y_{t-1}	–	.308	–	–	–	–	–
		(.106)					
Δp_t	–	–	.065	–	–	–	–
			(.041)				
Δp_{t-1}	.535	.492	–	–	.166	.133	.0
	(.088)	(.098)			(.109)	(.119)	–
Δp_{t+1}	–	–	–	.474	–	–	–
				(.187)			
$S_t\Delta p_{t+1}$	–	–	–	–	.775	.846	.98
					(.161)	(.188)	–
$S_{t-1}\Delta p_t$	–	–	.824	.494	–	–	–
			(.052)	(.186)			
$\bar{R}^2$	.469	.443	.913	.777	.592	.588	.592
SER	1.53	1.56	.470	.753	1.34	1.35	1.34

Table 6: Reduced-form estimates, 1978-1996. Uses CBO output gap and the Michigan survey of inflation expectations. In column 4, the output gap, inflation, and the lagged survey are used as instruments. In columns 6 and 7, the output gap, lagged inflation, and the lagged survey are used as instruments

portant explanatory power. In particular, column 3 presents a reduced-form model for survey expectations, in which the model-consistent expectation of inflation, $M_t\Delta p_{t+1}$, is replaced by its two instruments. These results show that survey expectations are related to the output gap and current-period inflation, which is what we would expect if survey respondents are using information about these variables in forming their expectations.

In column 4, we replace the two instruments with expected inflation, effectively projecting future inflation on the two instruments. Consistent with our conjecture, the coefficient on expected inflation—μ—is significantly different from zero. Overall, the results in column 4 line up with the predictions of Carroll's epidemiological model—and are sharply at odds with the predictions of model-consistent expectations. In particular, μ is much smaller than one, and λ—the coefficient on lagged survey expectations—is much larger than zero. Consistent with Carroll's conjecture that expectations of the broader public gradually converge to a well-informed value, the sum of μ and λ is close to one.[16]

In the final three columns of the table, we return to inflation as the dependent variable and include the survey as an explanatory variable. We view these results as a bridge to our full-information estimates; they will also allow us to assess the results of Mavroeidis et al. (2014), who focus on single-equation models. The estimates in column 5 simply adds the Michigan survey as another explanatory variable. The results indicate that over this sample, there is information in the Michigan survey that is valuable for explaining inflation. In particular, the t-statistic on the survey is very large, indicating that the Michigan survey is useful for explaining inflation over this sample. It is worth noting that the coefficient on the output gap is smaller in this case, and the t-statistic has dropped to 1.2. This change is perhaps not too surprising: Theory would suggest that part of the role of the output gap in the reduced-form setting is as a proxy for expectations, and the results in column 3 are consistent with such an interpretation.

While the coefficient on the Michigan survey is statistically significant, it is considerably smaller than the value that would be predicted by theory, $\frac{\beta}{(1+\beta\gamma)}$. As the discussion of the empirical model in Section 2.2 indicated, there may be considerable measurement error in survey expectations; as a consequence, the estimate of the parameter is likely to be biased downward. To assess whether this may be the source of the small size of the coefficient on the Michigan survey, the next column presents results using instrumental variables, with, notably, the survey lagged one period as an instrument. As can be seen, the coefficient on the survey is now much larger, and closer to what the theory would predict.

It remains the case that the coefficient on the output gap has a small t-statistic. Column 7 assesses whether two restrictions could boost this coefficient: First, the coefficient on lagged inflation is small and not statistically significant, suggesting that indexation is not a needed additional feature.[17] Consistent with that interpretation, the coefficient on the survey is restricted to be 0.98. Even with these restrictions, however, the t-statistic on the output gap remains small. This result lines up with the those

[16]Our complete empirical model of survey expectations allows for measurement error in expectations, which would complicate the reduced-form estimation of the model. Our full-information estimation in the next section will address this complication. We nonetheless view the results in column 4 as important in helping establish that, in this period, it appears that respondents to the Michigan survey were at least somewhat informed in forming their inflation expectations.

[17]Roberts (1998) also found that, conditional on survey expectations, there is little evidence of indexation in the Phillips curve.

of Mavroeidis et al. (2014): Conditional on survey expectations, the slope of the Phillips curve is not precisely estimated. Recall, however, that consistent with our theoretical interpretation, the survey itself is strongly related to the output gap. If our interpretation is correct, and the sensitivity of the survey to the output gap is anchored in an imperfectly informed view of the economy's evolution, then there should be information in the evolution of the survey about the sensitivity of inflation to aggregate demand. The full-information estimation we carry out in Section 6 exploits this relationship by imposing the cross-equation restrictions implied by the theory and so may allow us to estimate κ more precisely.

In table 7, we assess the sensitivity of our results to the measure of survey expectations, repeating the estimation in columns 3 through 7 of table 6, now using the Survey of Professional forecasters one-year-ahead inflation expectations.[18]

Column 1 shows the estimate for the reduced-form model of inflation expectations. As in column 3 of Table 6, the coefficient on the output gap is statistically significant. However, the coefficient is about half as large as it was for the Michigan survey, indicating somewhat less responsiveness of the SPF to aggregate demand. The other coefficients are about the same as in Table 6. Notably, however, the standard error of the regression is considerably smaller than for the Michigan survey. One possible reason is that the SPF is a less noisy measure of expectations than is the Michigan survey. Column 2 shows the estimates of our "structural" model of expectations. These results suggest an even larger departure from rationality than with the Michigan survey: Both the coefficients on expected future inflation and the coefficient on the lagged survey are smaller. These results suggest that the larger response of the Michigan survey to the output gap may be closer to the fully model-consistent response.[19] Nonetheless, the sum of the coefficients is 0.9, reasonably close to the value of one predicted by Carroll's epidemiological model.

In columns 3, 4, and 5, inflation is the dependent variable. In column 3, the SPF is included as an explanatory variable. As in column 5 of table 6, survey expectations are strongly statistically significant. The coefficient is very large—in fact somewhat bigger than the value predicted by theory, although not significantly so. Also consistent with the results with the Michigan survey, the coefficient on lagged inflation is small and not

[18] As with all measures of inflation and expectations, the SPF is detrended using the PTR measure of longer-run inflation expectations.

[19] As noted in Roberts (1997), an older literature often found that the Michigan survey suggested smaller violations of MCE than did surveys of professional forecasters.

	(1)	(2)	(3)	(4)	(5)
Dependent var:	$S_t\Delta p_{t+1}$	$S_t\Delta p_{t+1}$	Δp_t	Δp_t	Δp_t
Est. method:	OLS	IV	OLS	IV	IV
y_t	.046	–	.199	.194	.265
	(.021)		(.083)	(.085)	(.078)
y_{t-1}	–	–	–	–	–
Δp_t	.062	–	–	–	–
	(.028)				
Δp_{t-1}	–	–	.046	.030	.0
			(.114)	(.124)	–
Δp_{t+1}	–	.201	–	–	–
		(.069)			
$S_t\Delta p_{t+1}$	–	–	1.192	1.232	.98
			(.211)	(.242)	–
$S_{t-1}\Delta p_t$	.846	.703	–	–	–
	(.046)	(.093)			
$\bar{R}^2$	.930	.898	.626	.623	.611
SER	.318	.375	1.28	1.29	1.31

Table 7: Reduced-form estimates for inflation and survey expectations, 1978-1996. Uses CBO output gap and the SPF survey of inflation expectations. In column 4, the output gap, inflation, and the lagged survey are used as instruments. In columns 6 and 7, the output gap, lagged inflation, and the lagged survey are used as instruments

statistically significant. In contrast to the results in table 6, however, the output gap is statistically significant. The size of the coefficient is also considerably larger than was the case for the Michigan survey. In this case, estimating with instrumental variables as in column 4 makes little difference for the results. That's consistent with the view that the SPF is a less noisy measure of expectations than is the Michigan survey. Imposing the New Keynesian theoretical restrictions as in column 5 has little effect on the fit of the model: The coefficient on the output gap remains statistically significant and is even larger in this column than in column 6.

One reason the coefficient on the output gap is larger and more highly statistically significant when the measure of expectations is taken from the SPF may be is that the SPF is less strongly related to the output gap than is the Michigan survey. As we noted earlier, the correlation of the Michigan survey with the output gap may be one reason why the output gap has such a small coefficient in columns 5, 6, and 7 of table 6. Because the SPF is less tightly linked to the output gap, the direct coefficient on the output gap is larger.

A.2. Reduced-form results for the later sample

Table 8 repeats the series of reduced-form estimations using the Michgan survey as in Table 6, this time in the latter sample, 1997 to 2015. As can be seen in column 1, in the very simple reduced-form Phillips curve, the coefficient on the output gap is much smaller than in the earlier sample. This result is consistent with the conventional wisdom discussed in the introduction and with the VAR results presented in Section 4: The reduced-form relationship between inflation and the business-cycle is now much weaker than in the past. The coefficient is nonetheless statistically significant. The coefficient on the output gap is also statistically significant in column 2, suggesting that, at least with this dataset, there is the potential for the output gap to serve as an instrument for future inflation.

The results in columns 3 and 4 foreshadow one of the key findings of our paper, here through the lens of the reduced-form analysis. In column 3, we estimate a reduced-form process for the Michigan survey, including the output gap and inflation as regressors—as before, because of their potential for serving as instruments for future inflation. In contrast to the results in Table 8, the coefficient on the output gap is very small and statistically insignificant. The coefficient on core inflation is also statistically insignificant. As might be expected given these results, when these variables are used as instruments for future inflation as in column 4, the coefficient on the instrumented variable is also statistically insignificant. These results can be interpreted as indicating that, in contrast to the early sample, there is very little relationship between inflation expectations and overall macroeconomic conditions in this sample. That is, the households represented by the Michigan survey have become inattentive to the output gap and core inflation in forming their inflation expectations.

In column 5, we add the Michigan survey as an observable to the reduced-form inflation model. The coefficient is much smaller than in the early sample and the *t*-statistic is only 1.1. In contrast to the early sample, the coefficient on the output gap is about the same size as in the version without the Michigan survey (column 1) and remains strongly statistically significant. Lagged inflation has a larger *t*-statistic than in the similar result in table 6 but the coefficient is not very large in economic terms.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Dependent var:	Δp_t	Δp_t	$S_t\Delta p_{t+1}$	$S_t\Delta p_{t+1}$	Δp_t	Δp_t	Δp_t
Est. method:	OLS	OLS	OLS	IV	OLS	IV	IV
y_t	.120	–	.007	–	.119	.120	.137
	(.032)		(.029)		(.031)	(.032)	(.039)
y_{t-1}	–	.136	–	–	–	–	–
		(.032)					
Δp_t	–	–	.087	–	–	–	–
			(.101)				
Δp_{t-1}	.253	.199	–	–	.245	.247	.0
	(.109)	(.111)			(.109)	(.109)	–
Δp_{t+1}	–	–	–	.146	–	–	–
				(.143)			
$S_t\Delta p_{t+1}$	–	–	–	–	.117	.083	.98
					(.103)	(.167)	–
$S_{t-1}\Delta p_t$	–	–	.606	.597	–	–	–
			(.092)	(.095)			
$\bar{R}^2$	.334	.341	.366	.340	.318	.317	-.334
SER	.482	.473	.430	.439	.482	.482	.674

Table 8: Reduced-form estimates for inflation and survey expectations, 1997-2015. Uses CBO output gap and the Michigan survey of inflation expectations. In column 4, the output gap, inflation, and the lagged survey are used as instruments. In columns 6 and 7, the output gap, lagged inflation, and the lagged survey are used as instruments

Column 6 uses the lagged Michigan survey as an instrument for the contemporaneous value; the results are little changed. Column 7 restricts the coefficient on lagged inflation to be zero and imposes the theory-based coefficient of 0.98 on the Michigan survey. In this case, the coefficient on the output gap remains statistically significant and of about the same size as in column 1. Such a finding is not surprising given that households are inattentive: Inattention implies that $\mu = 0$; in that case, the simple Phillips curve as in column 1 gives a direct reading on κ. Adding the Michigan survey as in column 7 adds noise but does not disturb the relationship between inflation and the output gap.

Table 9 shows the results using the Survey of Professional Forecasters one-year-ahead inflation expectations for the later sample. In contrast to the results with the Michigan survey, the output gap is statistically significant in the regression explaining the SPF. An interpretation of this finding is that, unlike households, professional fore-

	(1)	(2)	(3)	(4)	(5)
Dependent var:	$S_t\Delta p_{t+1}$	$S_t\Delta p_{t+1}$	Δp_t	Δp_t	Δp_t
Est. method:	OLS	IV	OLS	IV	IV
y_t	.070	–	.094	.067	.035
	(.014)		(.048)	(.096)	(.028)
y_{t-1}	–	–	–	–	–
Δp_t	.046	–	–	–	–
	(.037)				
Δp_{t-1}	–	–	.210	.168	.0
			(.123)	(.180)	–
$M_t\Delta p_{t+1}$	–	.388	–	–	–
		(.107)			
$S_t\Delta p_{t+1}$	–	–	.267	.534	.98
			(.361)	(.890)	–
$S_{t-1}\Delta p_t$	.401	.473	–	–	–
	(.090)	(.126)			
$\bar{R}^2$	.736	.272	.311	.306	.296
SER	.155	.238	.484	.486	.493

Table 9: Reduced-form estimates for inflation and survey expectations, 1978-1996. Uses CBO output gap and the SPF survey of inflation expectations. In column 4, the output gap, inflation, and the lagged survey are used as instruments. In columns 6 and 7, the output gap, lagged inflation, and the lagged survey are used as instruments

casters have continued to pay attention to economic conditions in forming their inflation expectations. This interpretation is also consistent with the results in column 2, which suggest that the "rational" expectation of future inflation continued to be an important determinant of the forecasts captured by the SPF. Indeed, the coefficient is actually larger in the latter period for the SPF, suggesting that professional forecasters did a better job assessing the impact of aggregate demand on inflation. The sum of the two coefficients in column 2 is 0.86, not far from the sum of 0.90 in Table 7.

In columns 3, 4, and 5, inflation is the dependent variable and the SPF is included as an explanatory variable. In contrast to the early-sample results for the SPF, and in line with the late-sample Michigan results, the survey is not statistically significant in the latter sample. Using IV as in column 4 boosts the coefficient but it remains statistically insignificant. Moreover, the coefficient on the output gap falls and is not statistically significant. This result is consistent with the early-sample Michigan survey results: The SPF is strongly related to the output gap and so an estimate with a large coefficient on this variable, as in column 4, will have a smaller coefficient on the output gap itself. In column 5, the coefficient on the survey is constrained to an even larger value and the direct coefficient on the output gap is smaller still.

Appendix B. Sensitivity: Additional Survey Measures

B.1. Estimation results with additional survey measures

Table 10 includes results for two additional survey measures: The Livingston survey of forecasters and the SPF measure of forecasts for inflation as measured by the GDP deflator. The Livingston Survey is a semi-annual survey of economists from industry, government, banking, and academia that is since 1990 conducted by the Federal Reserve Bank of Philadelphia. An advantage of the Livingston survey is that it is available on a consistent basis over a long sample (indeed, it's been collected since the 1940s), for consumer prices. Similarly, the SPF for GDP prices has been collected on a consistent basis over our sample. One peculiarity of the Livingston survey is that it is collected semiannually. In this section, we estimate our model for the Livingston survey using semiannual data; in the next section, we introduce the Livingston survey into a model using quarterly data and allow the Kalman filter to fill in missing information.[20] Because the frequency is semiannual, the coefficient magnitudes will be different for the Livingston survey. Also, when we measure expectations using the SPF for GDP prices, we use the corresponding measure of inflation.

[20]For the output gap, we average the two quarterly values. For inflation, we average the two quarterly inflation rates, making the inflation measure effectively the (annualized) percent change between the final quarters of each half year.

	(1) Livingston 1978-1996	(2) Livingston 1997-2015	(3) Michigan 1978-1996	(4) Michigan 1997-2015	(5) SPF CPI 1978-1996	(6) SPF CPI 1997-2015	(7) SPF PGDP 1978-1996	(8) SPF PGDP 1997-2015
γ	0.545	0.551	0.291	0.484	0.433	0.371	0.716	0.516
	[0.333, 0.753]	[0.326, 0.776]	[0.086, 0.500]	[0.269, 0.714]	[0.231, 0.625]	[0.158, 0.571]	[0.516, 0.919]	[0.317, 0.708]
κ	0.212	0.0866	0.209	0.128	0.215	0.093	0.167	0.153
	[0.0884, 0.337]	[0.0317, 0.141]	[0.0979, 0.309]	[0.0620, 0.192]	[0.106, 0.321]	[0.0502, 0.135]	[0.0666, 0.264]	[0.0619, 0.243]
σ_ϵ	1.6613	0.6111	1.8103	0.7346	2.375	0.6769	1.5785	1.8855
	[1.2813, 2.0199]	[0.4512, 0.7696]	[1.4031 2.2124]	[0.5671, 0.9021]	[1.9261, 2.8031]	[0.5325, 0.8184]	[1.3042, 1.8632]	[1.5693, 2.2037]
λ	0.657	0.501	0.762	0.610	0.719	0.563	0.715	0.690
	[0.513, 0.801]	[0.306, 0.706]	[0.665, 0.857]	[0.406, 0.807]	[0.621, 0.815]	[0.443, 0.686]	[0.628, 0.804]	[0.580, 0.804]
μ	0.244	0.214	0.207	0.0555	0.198	0.247	0.205	0.106
	[0.0967, 0.384]	[0.0415, 0.390]	[0.104, 0.309]	[-0.0953, 0.208]	[0.119, 0.267]	[0.152, 0.339]	[0.136, 0.276]	[0.0631, 0.148]
σ_ν	0.4205	0.1705	0.3395	0.1928	0.1919	0.062	0.1941	0.0849
	[0.2495, 0.5871]	[0.0548, 0.2836]	[0.2153 0.4597]	[0.0801, 0.3034]	[0.1027, 0.2854]	[0.0225, 0.1015]	[0.1243, 0.2655]	[0.0412, 0.1279]
ρ	0.538	-0.0525	0.802	0.612	0.643	0.283	0.442	0.616
	[0.153, 0.972]	[-0.366, 0.235]	[0.430, 1.00]	[0.406, 0.822]	[0.208, 1.00]	[1e-04, 0.500]	[5e-04, 0.818]	[0.228, 0.970]
σ_ω	0.2921	0.2846	0.2874	0.3856	0.1822	0.1427	0.1645	0.0946
	[0.0969, 0.4809]	[0.1935, 0.3806]	[0.1346 0.4314]	[0.3045, 0.4700]	[0.0728, 0.2851]	[0.1096, 0.1747]	[0.0773, 0.2485]	[0.0497, 0.1394]
ϕ_1	1.36	1.41	1.22	1.32	1.14	1.34	1.29	1.34
	[1.13, 1.6]	[1.19, 1.63]	[1.03, 1.42]	[1.15, 1.49]	[0.942, 1.33]	[1.17, 1.51]	[1.11, 1.48]	[1.16, 1.52]
ϕ_2	-0.480	-0.450	-0.293	-0.330	-0.217	-0.353	-0.354	-0.372
	[-0.702, -0.245]	[-0.686, -0.206]	[-0.485, -0.0916]	[-0.512, -0.157]	[-0.415, -0.0212]	[-0.529, -0.176]	[-0.547, -0.162]	[-0.548, -0.200]
ϕ_3	-0.336	-0.283	-0.0883	-0.0444	-0.0704	-0.0761	-0.110	0.0241
	[-0.514, -0.150]	[-0.837, 0.260]	[-0.168, -0.0078]	[-0.261, 0.175]	[-0.157, 0.020]	[-0.288, 0.142]	[-0.263, 0.0471]	[-0.0656, 0.112]
ϕ_4	0.198	-0.0901	-0.0522	-0.190	-0.079	-0.149	0.0640	-0.0606
	[-0.0041, 0.401]	[-0.615, 0.449]	[-0.144, 0.0398]	[-0.404, 0.0275]	[-0.175, 0.0138]	[-0.363, 0.0685]	[-0.0904, 0.224]	[-0.147, 0.028]
σ_η	0.8782	0.8264	0.7361	0.5785	0.8211	0.5814	0.794	0.5827
	[0.7012, 1.0395]	[0.6699, 0.9858]	[0.6344 0.8397]	[0.4985, 0.6572]	[0.6963, 0.9461]	[0.5021, 0.661]	[0.6835, 0.9023]	[0.5006, 0.6598]

Table 10: Estimated using Bayesian methods; the priors are described in Appendix C.

Looking first at the results for the Phillips curve in the upper panel, the estimates of κ with the Livingston survey are similar to those from the Michigan survey and SPF/CPI—in particular, the estimates of κ drop by roughly half across the two samples. The results for the SPF/GDP are different, however: The point estimate of κ is lower than for the other measures in the earlier period and higher in the latter period and, as a consequence, doesn't fall as much. There is a similar pattern for the estimated standard error of the shock to this equation, σ_ϵ: For the Livingston sruvey, there is a decline in σ_ϵ across the two samples whereas for the SPF/GDP, there is little change, and a value that is notably higher than for the other measures in the latter sample.

Turning to the results for the model of inflation expectations in the middle panels, the results for the Livingston survey more closely resemble those for the Michigan survey, with μ remaining large and statistically significant in the latter sample. By contrast, for the SPF/GDP, the point estimate of μ fell by factor of two across the two samples, more in line with the results for the Michigan survey, albeit not as extreme. In all four cases, the shock to the structural expectations equation, σ_ν, fell across the two samples. The point estimates in each sample are similar for the two sets of SPF results, and are smaller than for the other two surveys, which are broadly similar to each other.

Table 11 shows results for the additional measures of expectations using the more general model discussed in Section 6. As in Table 10, the results for the two additional surveys of forecasters—the Livingston survey and the SPF for GDP prices—are broadly similar to the results for the SPF for the CPI.

	(1) Livingston 1978-1996	(2) Livingston 1997-2015	(3) Michigan 1978-1996	(4) Michigan 1997-2015	(5) SPF CPI 1978-1996	(6) SPF CPI 1997-2015	(7) SPF PGDP 1978-1996	(8) SPF PGDP 1997-2015
γ	0.519	0.562	0.343	0.506	0.436	0.450	0.699	0.530
	[0.304, 0.730]	[0.347, 0.795]	[0.124, 0.543]	[0.286, 0.727]	[0.237, 0.635]	[0.222, 0.669]	[0.507, 0.898]	[0.325, 0.727]
κ	0.182	0.0568	0.147	0.0969	0.173	0.0494	0.192	0.107
	[0.0481, 0.321]	[-0.0071, 0.118]	[0.0231, 0.273]	[-0.0034, 0.197]	[0.0399, 0.306]	[-0.0076, 0.104]	[0.0856, 0.300]	[-6e-04, 0.221]
σ_ϵ	1.649	0.607	1.918	0.739	2.399	0.748	1.578	1.906
	[1.263 2.011]	[0.448, 0.771]	[1.497, 2.333]	[0.573, 0.902]	[1.964, 2.851]	[0.579, 0.903]	[1.322, 1.836]	[1.581, 2.229]
λ	0.604	0.523	0.866	0.631	0.724	0.467	0.616	0.515
	[0.389, 0.813]	[0.295, 0.752]	[0.714, 1.01]	[0.386, 0.873]	[0.558, 0.889]	[0.320, 0.620]	[0.497, 0.733]	[0.347, 0.681]
μ	0.106	0.0972	0.0807	0.0414	0.104	-0.0209	0.178	0.0464
	[-0.0816, 0.288]	[-0.119, 0.321]	[-0.0654, 0.225]	[-0.167, 0.251]	[-0.0208, 0.230]	[-0.157, 0.119]	[0.0766, 0.280]	[-0.0038, 0.0979]
θ_3	0.123	-0.111	-0.006	-0.0924	0.0444	0.0937	0.0712	0.0403
	[-0.0035, 0.257]	[-0.299, 0.0744]	[-0.0803, 0.0682]	[-0.244, 0.0593]	[-0.0121, 0.0977]	[0.0319, 0.153]	[0.0082, 0.136]	[0.0198, 0.0606]
θ_4	0.0899	0.0433	0.0661	0.032	0.0348	0.0495	-0.0121	0.0256
	[-0.0034, 0.182]	[0.0037, 0.0825]	[0.0096, 0.121]	[-0.025, 0.0903]	[-0.0098, 0.0798]	[0.0256, 0.0739]	[-0.0428, 0.0178]	[0.001, 0.0488]
σ_ν	0.439	0.193	0.383	0.204	0.211	0.094	0.216	0.078
	[0.251, 0.626]	[0.0736, 0.302]	[0.261, 0.521]	[0.0857, 0.324]	[0.100, 0.310]	[0.040, 0.146]	[0.158, 0.271]	[0.036, 0.118]
ρ	0.582	-0.0598	0.653	0.651	0.671	0.280	0.00	0.795
	[0.217, 0.977]	[-0.366, 0.225]	[0.167, 1.00]	[0.415, 0.913]	[0.250, 1.00]	[1e-04, 0.517]	[0, 0]	[0.640,0.993]
σ_ω	0.309	0.284	0.269	0.383	0.179	0.113	0.121	0.091
	[0.101, 0.501]	[0.188, 0.382]	[0.118, 0.411]	[0.296, 0.472]	[0.066, 0.284]	[0.067, 0.158]	[0.0539, 0.187]	[0.053, 0.129]
ϕ_1	1.35	1.4	1.21	1.32	1.13	1.31	1.29	1.34
	[1.12, 1.58]	[1.18, 1.62]	[1.03, 1.41]	[1.15, 1.49]	[0.937, 1.34]	[1.15, 1.49]	[1.11, 1.48]	[1.16, 1.51]
ϕ_2	-0.476	-0.444	-0.285	-0.330	-0.206	-0.326	-0.357	-0.364
	[-0.704, -0.250]	[-0.684, -0.201]	[-0.478, -0.0881]	[-0.510, -0.156]	[-0.405, -0.0046]	[-0.496, -0.146]	[-0.554, -0.170]	[-0.538, -0.191]
ϕ_3	-0.329	-0.273	-0.085	-0.0482	-0.0647	-0.049	-0.107	0.0304
	[-0.513, -0.145]	[-0.808, 0.276]	[-0.166, -0.0032]	[-0.261, 0.170]	[-0.154, 0.024]	[-0.268, 0.161]	[-0.261, 0.0494]	[-0.0592, 0.118]
ϕ_4	0.190	-0.0876	-0.054	-0.189	-0.0869	-0.174	0.059	-0.0672
	[-0.0216, 0.390]	[-0.604, 0.453]	[-0.151, 0.0382]	[-0.400, 0.0321]	[-0.182, 0.0072]	[-0.388, 0.0412]	[-0.0954, 0.219]	[-0.154, 0.0194]
σ_η	0.881	0.828	0.737	0.580	0.823	0.580	0.793	0.582
	[0.708, 1.049]	[0.667, 0.987]	[0.633, 0.833]	[0.499, 0.657]	[0.692, 0.943]	[0.500, 0.657]	[0.685, 0.905]	[0.501, 0.659]

Table 11: Estimated using Bayesian methods; the priors are described in Appendix C.

B.2. Pooling the information from the surveys

In this appendix, we take the position that three of the measures of expectations of consumer prices—from the SPF, the Livingston survey, and Michigan survey—are noisy indicators of the same underlying process. Thus, we assume that the same "'true"' measure of expectations is a common factor driving all three surveys, as in Equation 9, reproduced here for convenience:

$$E_t \Delta p_{t+1} = \mu M_t \Delta p_{t+1} + \lambda E_{t-1} \Delta p_t + \nu_t. \tag{17}$$

Each survey, however, has its individual measurement error:

$$S_t^i \Delta p_{t+1} = E_t \Delta p_{t+1} + u_t^i; \tag{18}$$

$$u_t^i = \rho_u^i u_{t-1}^i + \omega_t^i. \tag{19}$$

This specification has the advantage of bringing more information to bear on the estimation of Equation 9 at the risk of potential mis-specification if the fundamental component of expectations in fact differs across survey measures.

Table 12 shows results for this joint estimation The results are qualitatively similar to those shown in Table 2. For the structural Phillips curve, the results suggest that the slope parameter is smaller in the later sample, by a factor of two. As in the earlier results, that step-down is much smaller than the decline in the Phillips curve slope assuming model-consistent expectations reported in Table 1. The estimation degree of indexation (γ) is about the same in the two periods.

The model of inflation expectations, which is common across the three expectations measures, also shows some shift between the two sample periods. The parameter λ falls somewhat, while the parameter μ rises somewhat. The sum of the two coefficients declines, from 0.95 to 0.80. These results are similar to those for the SPF. The results also suggest that structural shocks to inflation expectations are less important in the later sample: The estimated standard deviation of the parameter σ_ν falls by a factor of two, and the shock is propagated forward less because λ is also smaller.[21]

[21] We also estimate an alternative variant of pooling, where we allow for different loading factors, $\Lambda_{j,t}$, for different survey measures of expectation formation. These results are shown in columns 3 and 4 of the table.

	(1) Combination/Equal 1978-1996	(2) Combination/Equal 1997-2015	(3) Combination/Factors 1978-1996	(4) Combination/Factors 1997-2015
γ	0.355	0.358	0.282	0.360
	[0.165, 0.544]	[0.132, 0.590]	[0.089, 0.474]	[0.140, 0.596]
κ	0.211	0.094	0.196	0.099
	[0.111, 0.311]	[0.050, 0.138]	[0.101, 0.291]	[0.053, 0.145]
σ_ϵ	1.958	0.705	1.799	0.703
	[1.60, 2.32]	[0.542, 0.857]	[1.44, 2.16]	[0.539, 0.856]
λ	0.766	0.537	0.752	0.540
	[0.688, 0.841]	[0.410, 0.667]	[0.670, 0.832]	[0.401, 0.672]
μ	0.185	0.256	0.218	0.236
	[0.115, 0.259]	[0.152, 0.356]	[0.135, 0.306]	[0.131, 0.344]
σ_ν	0.268	0.120	0.324	0.110
	[0.218, 0.317]	[0.092, 0.148]	[0.255, 0.389]	[0.069, 0.151]
Λ_{SPF}	-	-	0.760	1.143
	-	-	[0.645, 0.951]	[0.770, 1.497]
Λ_{Liv}	-	-	0.762	1.200
	-	-	[0.645, 0.880]	[0.777, 1.623]
ρ_{Mich}	0.852	0.770	0.815	0.751
	[0.748, 0.970]	[0.636, 0.907]	[0.686, 0.951]	[0.607, 0.901]
$\sigma_{\omega 1}$	0.335	0.386	0.306	0.390
	[0.282, 0.387]	[0.332, 0.440]	[0.251, 0.364]	[0.333, 0.445]
ρ_{SPF}	0.476	-0.366	0.0001	-0.366
	[0.004, 0.868]	[-0.366, -0.366]	[0.0001, 0.0001]	[-0.366, -0.366]
$\sigma_{\omega 2}$	0.066	0.068	0.074	0.069
	[0.021, 0.108]	[0.041, 0.095]	[0.033, 0.115]	[0.041, 0.098]
ρ_{Liv}	0.356	0.020	0.426	0.019
	[0.001, 0.626]	[-0.298, 0.315]	[0.049, 0.734]	[-0.317, 0.317]
$\sigma_{\omega 3}$	0.268	0.242	0.284	0.245
	[0.201, 0.334]	[0.191, 0.291]	[0.212, 0.358]	[0.194, 0.299]
ϕ_1	1.21	1.34	1.21	1.34
	[1.03, 1.39]	[1.17, 1.52]	[1.04, 1.39]	[1.17, 1.51]
ϕ_2	-0.261	-0.352	-0.261	-0.352
	[-0.444, -0.081]	[-0.531, -0.179]	[-0.444, -0.083]	[-0.528, -0.174]
ϕ_3	-0.100	-0.080	-0.101	-0.074
	[-0.170, -0.029]	[-0.295, 0.139]	[-0.171, -0.030]	[-0.294, 0.132]
ϕ_4	-0.0467	-0.138	-0.048	-0.140
	[-0.126, 0.029]	[-0.346, 0.079]	[-0.124, 0.031]	[-0.360, 0.069]
σ_η	0.657	0.583	0.658	0.582
	[0.564, 0.743]	[0.499, 0.662]	[0.566, 0.748]	[0.500, 0.659]

Table 12: Estimated using Bayesian methods; the priors are described in Appendix C.

In these pooled results, the results more closely line up with those from the SPF than from the Michigan survey. In particular, the sharp reduction in μ for the Michigan survey is not evident here.

Appendix C. Priors

The following table summarizes the priors we use for the Bayesian estimation of the model we introduce in Section 2.2.

	mean	sd	distribution	min	max
γ	0.5	0.15	Normal	-1	3
κ	0.1	0.1	Normal	-10	10
σ_ϵ	1	0.5	Beta	0	5
λ	0.75	0.15	Normal	-2	2
μ	0.25	0.15	Normal	-10	10
θ_3	0	0.5	Normal	-10	10
θ_4	0	0.5	Normal	-10	10
σ_ν	1	0.5	Beta	0	5
ρ	0.5	0.5	Normal	0	1
σ_ω	1	0.5	Beta	0	5
c_1	0.1	2	Normal	-10	10
c_2	0.1	2	Normal	-10	10
ϕ_1	1.3	0.5	Normal	-10	10
ϕ_2	-0.5	0.5	Normal	-10	10
ϕ_3	0.1	1	Normal	-10	10
ϕ_4	0.1	1	Normal	-10	10
σ_η	1	0.5	Beta	0	5

Table 13: Priors for Bayesian estimation.