

Fuhrer, Jeffrey C.

Working Paper
Inflation persistence

Working Papers, No. 09-14

Provided in Cooperation with:
Federal Reserve Bank of Boston

Suggested Citation: Fuhrer, Jeffrey C. (2009) : Inflation persistence, Working Papers, No. 09-14, Federal Reserve Bank of Boston, Boston, MA

This Version is available at:
<https://hdl.handle.net/10419/55645>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Inflation Persistence

Jeffrey C. Fuhrer

Abstract:

This paper examines the concept of inflation persistence in macroeconomic theory. It begins with a definition of persistence, emphasizing the difference between reduced-form and structural persistence. It then examines a number of empirical measures of reduced-form persistence, considering the possibility that persistence may have changed over time. The paper then examines the theoretical sources of persistence, distinguishing “intrinsic” from “inherited” persistence, and deriving a number of analytical results on persistence. It summarizes the implications for persistence from the literatures on “sticky-information” models, learning models, and so-called “trend inflation models,” providing some new results throughout.

Keywords: inflation persistence, Phillips curve, autocorrelation

JEL Classifications: E31, E52

Jeffrey C. Fuhrer is the director of research and an executive vice president in the research department of the Federal Reserve Bank of Boston. His e-mail address is jeff.fuhrer@bos.frb.org.

This paper, which may be revised, is available on the web site of the Federal Reserve Bank of Boston at <http://www.bos.frb.org/economic/wp/index.htm>.

The author thanks Fabià Gumbau-Brisa, Timothy Cogley, Denny Lie, Giovanni Olivei, Scott Schuh, Oz Shy, Raf Wouters, and participants at the Federal Reserve Bank of Boston’s lunchtime workshop for helpful comments.

The views and opinions expressed in this paper are those of the author and do not necessarily represent the views of the Federal Reserve Bank of Boston or the Federal Reserve System.

This version: November 30, 2009

1 Introduction

What is meant by “persistence” and why is it important to macroeconomists and policymakers? In broad terms, persistence is the economic analogue of inertia in physics. Inertia may be defined as the resistance of a body to changing its velocity (direction and rate of speed) unless acted upon by an external force. This law is often paraphrased as “a body at rest will remain at rest unless acted upon by an external force,” which is one example of the principle. Newton’s second law of motion, $F = ma$ or $a = \frac{F}{m}$, captures this idea algebraically: The magnitude of the force required to produce a given change of velocity (acceleration) is proportional to the body’s mass.

While no analogy is perfect, this one works reasonably well at an intuitive level. An economic variable is said to be persistent if, other things being equal, it shows a tendency to stay near where it has been recently, absent other economic forces that move it elsewhere. In the case of inflation, the rate of change of the price level tends to remain constant (inflation tends to be persistent) in the absence of an economic “force” to move it from its current level. This paper provides more precise definitions of inflation persistence in the sections below, but this physical analogy may provide motivating intuition.

1.1 Early inflation models and the empirical necessity of lagged inflation

For many decades, economists assumed that inflation was an inertial or persistent economic variable. The concept of the sacrifice ratio—the number of point-years of elevated unemployment required to reduce inflation by a percentage point—implies that inflation does not move freely, requiring significant economic effort in the form of lost output to reduce its level.¹ The early incarnations of the accelerationist Phillips curve modeled the apparent inertia in inflation by including lags of inflation. A canonical example of such specifications is Gordon’s “triangle model” of inflation, here replicated in simplified form

¹See Gordon, King, and Modigliani (1982) for the first study that uses the term “sacrifice ratio.” This study followed on the work of Arthur Okun (1977).

(Gordon 1982):²

$$\pi_t = \sum_{i=1}^k a_i \pi_{t-i} - b(U_t - \bar{U}) + cx_t + \epsilon_t. \quad (1.1)$$

Inflation, π_t , depends on its own lags (normally constrained to sum to one to reflect the Friedman-Phelps accelerationist principle), a measure of real activity (here the deviation of unemployment, U_t , from the non-accelerating inflation rate of unemployment or “NAIRU”), and supply-shifters such as key relative price shifts, summarized in x_t . In such a model, inflation moves gradually, partially anchored by its recent history, in response to real activity and supply shocks. These variables may themselves be persistent, in which case inflation will “inherit” some of their persistence. A key question is whether and why inflation has its own or “intrinsic” persistence, beyond that inherited from U_t and x_t (or perhaps ϵ_t , if it is also serially correlated). If inflation exhibits intrinsic persistence, then a model of inflation may require the equivalent of the lags in equation 1.1 above.

In this early literature, the theoretical justification for including lags of inflation was as a proxy for expected inflation and a proxy for contracting and other price-setting frictions. As an empirical matter, the lags helped the model fit the data. To see this last point simply, consider the R^2 for estimates of Gordon-style Phillips curves with and without the lags of inflation in table 1. The specification is $\pi_t = \sum_{i=1}^4 \alpha_i \pi_{t-i} + \sum_{j=1}^2 \beta_j U_{t-j} + \sum_{k=1}^2 \gamma_k \Delta r p_{t-k}^o + C$, where π_t is the quarterly percentage change in the core CPI, U is the civilian unemployment rate, and $r p^o$ is the relative price of oil. The sum of the α_i 's is constrained to one in some of the estimates.³

The points to take from this table are that (1) the lags of inflation are empirically critical, whatever they may represent structurally; and (2) as a consequence, it is critical to understand what these lags represent structurally.⁴

²See Friedman (1968) and Phelps (1968) for the earliest explications of the accelerationist Phillips curve.

³Note that this constraint is not statistically significant in these regressions—the R^2 's are identical in the constrained and unconstrained cases, and the p -value for the test of these restrictions exceeds 0.8 for both samples.

⁴The results in the table are completely invariant to the choice of the lag length for lagged inflation. For lag lengths up to 24, we obtain nearly identical coefficient sums, R^2 's, and non-binding unit sum constraints. However, if we begin the later sample in 1997, results

Table 1: R^2 for Gordon-style Phillips curves


Model	R^2
Core CPI, 1966:Q1-1984:Q4	
with lags, $\sum \alpha_i = 1$	0.74
with lags, $\sum \alpha_i \neq 1$	0.74
without lags	0.24
Core CPI 1985:Q1-2008:Q4	
with lags, $\sum \alpha_i = 1$	0.79
with lags, $\sum \alpha_i \neq 1$	0.79
without lags	0.09
Core PCE, 1966:Q1-1984:Q4	
with lags, $\sum \alpha_i = 1$	0.76
with lags, $\sum \alpha_i \neq 1$	0.77
without lags	0.39
Core PCE 1985:Q1-2008:Q4	
with lags, $\sum \alpha_i = 1$	0.72
with lags, $\sum \alpha_i \neq 1$	0.72
without lags	0.16

1.2 Rational expectations and inflation persistence: An introduction to some of the issues

The introduction of Muth’s (1961) theory of rational expectations into the macroeconomics literature and the consequent move toward explicit modeling of expectations posed considerable challenges in modeling prices and inflation. In the earliest rational expectations models of Lucas (1972) and Sargent and Wallace (1975), the price level was a purely forward-looking or expectations-based variable like an asset price, which in these models implied that prices were flexible, and could “jump” in response to shocks. It was difficult at first to reconcile the very smooth, continuous behavior of measured aggregate price indexes such as the consumer price index with the flexible-price implications of these early rational expectations models. Figure 1 displays the consumer price index for the post-war period.

A number of economists recognized the tension between the obvious persistence in the price level data and the implications of these early rational expectations models. Fischer (1977), Gray (1977), Taylor (1980), Calvo (1983), and change dramatically. The constraint on the sum of the lag coefficients is now binding, and the R^2 ’s drop to 0.25 or below, a result highlighted in Williams (2006). This apparent shift in reduced-form persistence is discussed in more detail below.

Figure 1: The consumer price index



Rotemberg (1982, 1983) developed a sequence of models that rely on nominal price contracting in attempts to impart a data-consistent degree of inertia to the price level in a rational expectations setting. The overlapping contracts of Taylor and Calvo/Rotemberg were successful in doing so, allowing contracts negotiated in period t to be affected by contracts set in neighboring periods, which would remain in effect during the term of the current contract.⁵ The subsequent trajectory of macroeconomic research drew heavily on these seminal contributors, who had neatly reconciled rational expectations with inertial (or persistent) macroeconomic time series.

However, in the early 1990s, a number of authors discovered that these rational expectations formulations had less satisfying implications for the *change* in the price level, that is, the rate of inflation. Ball (1994) demonstrated that such models could imply a counter-factual “disinflationary boom”—the central bank could engineer a disinflation that caused output to rise rather than contract. Fuhrer and Moore (1992, 1995) showed that Taylor-type contracting models implied a degree of inflation persistence that was far lower than was apparent in

⁵For example, a four-period contract negotiated in period t would be influenced by the contracts negotiated in the previous three periods, as well as by the contracts expected to be negotiated in the following three periods.

inflation data of the post-war period to that point. To build intuition, compare equation 1.1 with a Calvo- or Taylor-type inflation equation⁶

$$\pi_t = E_t \pi_{t+1} + \gamma \tilde{y}_t + \epsilon_t. \quad (1.2)$$

Implicitly, this makes inflation, π_t , a function of all expected future output gaps, \tilde{y}_t , (assuming that $E_t \epsilon_{t+i} = 0 \forall i > 0$, or equivalently that ϵ_t has no autocorrelation).⁷ Inflation will indeed inherit the persistence in output, but nothing beyond that. The lags in the Gordon triangle model are gone. Inflation is completely forward-looking: Following a shock to output, inflation can jump immediately in response. If output exhibits no persistence (that is, there are no “real rigidities”), then neither will inflation. In contrast, equation 1.1 implies that inflation depends on all *past* output gaps and shocks. As in the forward-looking model, inflation inherits the persistence in output, but the lagged inflation terms mean that inflation cannot jump in response to a shock to output—inflation exhibits additional persistence (or inertia) in the sense that in response to shocks, inflation has a tendency to remain near its most recent values.

To explore the dynamic implications of this forward-looking specification, we embed equation 1.2 in a skeletal macro model. To complete the model, we include a rudimentary I–S curve and a simplified policy rule

$$\begin{aligned} \tilde{y}_t &= \rho \tilde{y}_{t-1} - a(r_t - E_t \pi_{t+1}) \\ r_t &= b(\pi_t - \bar{\pi}), \end{aligned} \quad (1.3)$$

where r_t is the short-term policy rate controlled by the central bank, the inflation target is denoted $\bar{\pi}$, and the equilibrium real rate is suppressed for convenience.⁸ We will consider cases with zero and non-zero output persistence, that is, $\rho = 0$ or $\rho \neq 0$. The policy response to inflation gaps is calibrated by parameter b , which is set to 1.5 throughout.

Consider first the case in which output is not persistent. Inflation is at its target and the output gap is zero. For comparison with the large literature that

⁶See Roberts (1997) for a derivation that shows the approximate equivalence of these formulations.

⁷This is made explicit below in equation 3.2, but one can obtain the result by inspection through repetitive substitution of the definition of future inflation into equation 1.2.

⁸Note that the derivation of the Calvo Phillips curve of equation 1.2 assumes a zero steady-state inflation rate, whereas equations 1.3 allow for the possibility of a non-zero steady-state inflation rate. Cogley and Sbordone (2009) derive the appropriate Phillips curve for the case in which the central bank pursues a non-zero inflation rate. For the purposes of this exercise, the inaccuracy in the approximation is likely small.

attempts to measure the economy’s response to an identified monetary policy shock, consider a one unit positive shock to the policy rate r_t . With no inertia in output or inflation, both output and inflation are perturbed below their steady state for one period. In period 2, both return to their steady state. The solid lines in figure 2 display the results of this rather uninteresting exercise.

When output is persistent (here $\rho = 0.9$), the dynamics are a bit more complex. Starting from the same steady state, consider a one unit positive shock to the policy rate. The results of this simulation are depicted in the dashed lines in the same figure. Output is depressed below its steady state in the first period, and because of its own persistence, remains below its steady state for some time. In order for equation 1.2 to hold, it must be that expected inflation lies above current inflation whenever the output gap is negative. Because output is persistent, the expected change in inflation must be positive for as long as output remains negative. But ultimately, inflation will have to return to its original steady state. Thus, inflation must immediately jump down, and then rise to its equilibrium from below. Inflation exhibits dynamics that are reminiscent of the exchange rate in the famous Dornbusch (1976) overshooting model.⁹

For the sake of comparison, the dotted line in figure 2 shows the outcome when the inflation equation includes a lag of inflation. Here, the equation is a “hybrid” equation that mixes both forward- and backward-looking elements, of the form

$$\pi_t = \mu\pi_{t-1} + (1 - \mu)E_t\pi_{t+1} + \gamma\tilde{y}_t + \epsilon_t \quad (1.4)$$


with $\mu = 0.5$ in the figure. Now, in response to the monetary policy shock, output behaves approximately as before, but inflation declines gradually over several quarters, exhibiting the more typical “hump-shaped” response found in the literature on monetary VAR’s.¹⁰ It is no longer the case that expected inflation must exceed current inflation for as long as the output gap is negative.

Perhaps as striking are the differences in behavior in a fully credible disinflation. In the solid and heavy dashed lines of figure 3, inflation is forward-looking ($\mu = 0$). The central bank announces a permanent and fully credible reduction in its target inflation rate from 2 percent to 0 at time 1. As the figure

⁹This overshooting property is examined in Fuhrer and Moore (1995), and in Estrella and Fuhrer (2002).

¹⁰See Christiano, Eichenbaum, and Evans (2005) for representative results. Note that the output response differs modestly from the dashed line because the path of inflation, and thus the path of the short-term real rate, differs.

Figure 2: Response of inflation to a monetary policy shock in the simple Calvo model




indicates, regardless of how persistent is output, the inflation rate jumps to its new equilibrium in the period after the announcement, with no disruption to output, so that the solid and heavy dashed lines coincide in the figure.¹¹ When lagged inflation is added to the inflation equation, inflation declines gradually to its new long-run equilibrium, with a concomitant decline in output during the transition.

Many would view the dynamics of the purely forward-looking specification as strikingly counter-factual. Counter-factual or not, one needs to understand the dynamics of inflation to pursue appropriate monetary policy. A knowledge of the reduced-form behavior of inflation is not sufficient. The central bank needs to understand the sources of inflation dynamics—in these examples, whether it arises from the persistence of output, which may in turn arise from the behavior of monetary policymakers, or from persistence intrinsic to the price-setting process. This is one of the reasons that the issue of persistence is of more than passing interest to macroeconomists and policymakers.

To be sure, understanding why and when inflation may be persistent can

¹¹If the inflation target were pre-announced, the inflation rate would jump to its new target in the first period, again with no disruption to output.

Figure 3: Credible disinflations in the simple Calvo model



be more complicated than these simple examples suggest. For instance, the examples leave out the possibility that the policy shock is caused by a monetary authority with imperfect credibility; they also abstract from imperfect knowledge of the economy that might lead to learning on the part of private agents. Both of these situations would alter the dynamic implications of the simple examples. We return to some of these complications below.

2 Defining and measuring reduced-form inflation persistence

The discussion above suggests that it will be useful to distinguish reduced-form from structural inflation persistence. Reduced-form persistence will refer to an empirical property of an observed inflation measure, without economic interpretation. Structural persistence will refer to persistence that arises from identified economic sources. A key objective of recent inflation research has been to map observed or reduced-form persistence into the underlying economic

structures that produce it. To a significant extent, this challenge remains.

2.1 Defining reduced-form persistence

There is no single definitive measure of reduced-form persistence. As the sections below discuss, researchers have employed a variety of measures to capture the idea that inflation responds gradually to shocks, or remains close to its recent history.¹² Most of the measures of inflation persistence derive from the autocorrelation function for inflation, so it will be useful to define it here. The i^{th} autocorrelation, ρ_i , of a stationary variable, x_t —the correlation of the variable with its own i^{th} lag, x_{t-i} —may be expressed as¹³

$$\rho_i = \frac{E(x_t x_{t-i})}{V(x)}, \quad (2.1)$$

where $V(x)$ is the variance of x . The correlations are of course bounded between -1 and 1 . The variable's autocorrelation function is correspondingly defined as the vector of correlations of current period x with each of its own lags x_{t-i} from $i = 1$ to k :

$$A = [\rho_1, \dots, \rho_k]. \quad (2.2)$$

A time series will be said to be relatively persistent if its correlations with its own past decay slowly. Thus, a graphical depiction of this measure of a series' reduced-form persistence is provided by a plot of the variable's autocorrelation function. In figure 4 below, the correlation of the persistent series, P_t , with its own lags declines gradually from 0.9 to 0.3 over 12 periods. In contrast, the less persistent time series, N_t , shows a more rapid decline in its autocorrelation function from 0.5 to 0 in about eight periods. To a large extent, the alternative measures of persistence surveyed in section 2.2 below provide alternative ways of quantifying the rate at which inflation's autocorrelations decay.


The analytical representation of the autocorrelation function will also be useful in this discussion. For example, if the variable x_t is defined as a first-order autoregressive process

$$x_t = ax_{t-1} + \epsilon_t \quad (2.3)$$

¹²This definition implicitly assumes that inflation is *positively* correlated with its own lags, an assumption that holds up well over most of post-war history. More generally, a time series may be deemed persistent if the absolute value of its autocorrelations is high, so that a strongly *negatively* autocorrelated series would also be characterized as persistent.

¹³For convenience, this definition assumes that x_t is a mean zero series.

Figure 4: Hypothetical autocorrelation functions



$-1 < a < 1$, then its autocorrelation function is simply

$$A = [a, a^2, \dots, a^k]. \tag{2.4}$$

From equation 2.4 it is clear that the autocorrelations of x_t die out geometrically at the rate determined by the autoregressive parameter a . The analytical expressions for the autocorrelations of inflation in more complex structural models with rational expectations are derived below.

Some researchers define persistence as the extent to which shocks in the past have an effect on current inflation. This concept is related to the autocorrelation function. The more correlated is inflation with its distant past, the more shocks that perturb inflation in the distant past will be reflected in current inflation. More formally, adopting a simple first-order autoregressive model for inflation from equation 2.3, one can iterate the equation backward to obtain the moving-average representation of inflation

$$\begin{aligned} \pi_t &= \epsilon_t + a(a\pi_{t-2} + \epsilon_{t-1}) \\ &= \epsilon_t + a\epsilon_{t-1} + a^2\epsilon_{t-2} + a^3\epsilon_{t-3} + \dots, \end{aligned} \tag{2.5}$$

which shows that the larger is a , and thus the more slowly inflation's autocor-

relations decay, the larger is the influence of past shocks on current inflation—equivalently, shocks have a more persistent effect on inflation.¹⁴

2.2 Measuring reduced-form inflation persistence

There is little agreement in the extant literature on how best to measure persistence. Thus, we examine a battery of measures that attempt to capture the persistence in inflation:

- Conventional unit root tests;
- The autocorrelation function of the inflation series (as defined in equation 2.4 above);
- The first autocorrelation of the inflation series;
- The dominant root of the univariate autoregressive inflation process (defined below);
- The sum of the autoregressive coefficients for inflation;
- Unobserved components decompositions of inflation that estimate the relative contributions of “permanent” and “transitory” components of inflation (for example, the IMA(1,1) and related models proposed by Stock and Watson (2007)).

Because the autocorrelation function summarizes much of the information in a time series, it may be the best overall measure of persistence. But researchers often desire a single number that captures the overall persistence implied by the full autocorrelation function; hence, many of the other scalar measures itemized above. One can find overlap across the results from these tests, but, to be sure, the results are neither uniform nor unambiguous. Throughout, we examine these measures of inflation persistence across a number of subsamples, providing suggestive evidence as to whether persistence has changed over time.

¹⁴This algebra provides a more rigorous justification for the argument that old-style Phillips curves like those in equation 1.1 add persistence to inflation beyond that inherited from output—the lags of inflation imply persistent effects of past shocks on current inflation.

2.2.1 The data

For the purposes of this paper, we will focus on three key inflation measures. Each will be defined as 400 times the log change in the corresponding price index. The three indexes are the GDP deflator, the consumer price index (CPI), and the personal consumption expenditures chain-type price index (PCE). In some cases, we will examine the so-called “core” versions of the CPI and PCE, that is, the price indexes that exclude food and energy prices. We denote these series throughout as “CPI-X” and “PCE-X.” These series abstract from the high-frequency noise that can be introduced by volatile food and energy price series.¹⁵

Figure 5 displays the three overall series in the top panel, with corresponding core series in the bottom panel. Table 2 provides summary statistics for the three inflation series over several sample periods. Evident in the figure and echoed in table 2 are the drop in the level of all inflation measures since the early 1980s (the top panel of the table); the decline in the variance of all measures, consistent with the so-called “Great Moderation” (the second panel of the table); the high correlation among the three series (the third panel of the table); and the lesser volatility of the core series. Less evident in the figure but clear in the bottom panel of table 2 is the decline in the correlations across the series. In particular, while the CPI and PCE remain highly correlated, the correlations between the core and overall measures have declined, as have the correlations between the GDP deflator and the consumer price measures. Despite the relatively high correlations among these series, the measures of persistence presented below will show some noticeable differences across the series.

2.2.2 Unit root tests

Perhaps the first test of persistence should be a unit root test. If inflation contains a unit root, its persistence is unquestionably large (infinite) and its variance is unbounded.¹⁶ Many papers test for a unit root in inflation (see Barsky (1987), Ball and Cecchetti (1990)); prior to the 1990s, the results tend to suggest a unit root in inflation. In more recent years, researchers are more

¹⁵As indicated in figure 5, the core price series are available only starting in the late 1950s.

¹⁶A series with a unit root has infinite “memory,” in the sense that a shock in period t has influence on all periods $t+k$, $k > 0$. More formally, if $x_t = x_{t-1} + \epsilon_t$, then $x_t = \sum_{i=0}^{\infty} \epsilon_{t-i}$. Thus, any shock to a series with a unit root persists forever.

Figure 5: Inflation data


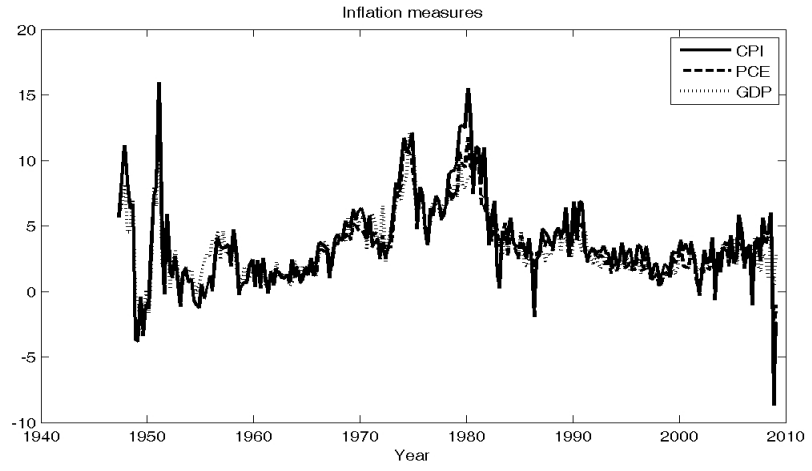


Table 2: Summary statistics, inflation measures

Series	Means			
	59-08	59-84	85-08	95-08
CPI	4.0	5.0	3.0	2.5
PCE	3.6	4.5	2.6	2.1
GDP	3.6	4.6	2.5	2.2
CPI-X	4.0	4.9	3.0	2.3
PCE-X	3.5	4.4	2.5	1.8
Series	Variances			
	59-08	59-84	85-08	95-08
CPI	9.4	12.7	3.7	4.3
PCE	6.4	8.6	2.2	2.2
GDP	5.7	7.8	0.95	0.83
CPI-X	7.2	10.9	1.2	0.31
PCE-X	4.7	6.4	1.2	0.25

Correlation matrix, 1959-2008					
	CPI	PCE	GDP	CPI-X	PCE-X
CPI	1.00				
PCE	0.96	1.00			
GDP	0.86	0.92	1.00		
CPI-X	0.86	0.85	0.86	1.00	
PCE-X	0.81	0.90	0.92	0.91	1.00
Correlation matrix, 1985-2008					
	CPI	PCE	GDP	CPI-X	PCE-X
CPI	1.00				
PCE	0.94	1.00			
GDP	0.57	0.70	1.00		
CPI-X	0.48	0.56	0.53	1.00	
PCE-X	0.40	0.62	0.65	0.87	1.00

likely to be unable to reject stationarity. Most monetary models would suggest that the more vigorous attention to inflation on the part of central banks around the world in recent decades is responsible for this change.

Table 3 provides univariate tests for the null that the inflation series contains a unit root, for a long sample (1966–2008), and for a “Great Moderation” sample (1985–2008). The results in the table are somewhat ambiguous. For the most recent decades, one can develop strong rejections of the null of a unit root, although this varies somewhat depending on the inflation series and on the test employed. For the longer sample, the Phillips-Perron test rejects the null, although not always very strongly. The ADF test fails to reject for all three inflation series.

Table 3: Unit Root Tests for Inflation Measures

<i>(p-values, null = series has unit root)</i>		
1966-2008		
Series	ADF	Phillips-Perron
CPI	0.14	0.0069
PCE	0.37	0.045
GDP	0.33	0.025
1985-2008		
Series	ADF	Phillips-Perron
CPI	0.00	0.00
PCE	0.07	0.00
GDP	0.31	0.00


What should one make of these tests? Certainly for the past 25 years, the U.S. central bank has behaved as if it had a specific, low inflation goal, although it has so far chosen not to announce that goal. If it does indeed have such a goal, then theory suggests that the U.S. inflation rate will not have a unit root—or that if it does have a unit root, the variance of that component of inflation is quite small.¹⁷ This reasoning, combined with the more frequent rejections of the unit root null in recent decades, suggests that one may safely assume that inflation does not contain a unit root, at least not under current monetary policy.

¹⁷This statement is somewhat too strong, as a very gradual drift downward in the target inflation rate could manifest itself as a unit root component of inflation, albeit with a relatively small variance. See Stock and Watson (2007) for an inflation model that formalizes this reasoning.

2.2.3 First-order autocorrelations

If we can proceed under the assumption that inflation does not contain a unit root, then section 2 above suggests the first-order autocorrelation coefficient for the series as a simple measure of persistence, a measure used, for example, in Pivetta and Reis (2007). Figure 6 extends and expands their figure 2a, presenting rolling-sample estimates of the first-order autocorrelation coefficient for three aggregate inflation measures.¹⁸ The figure shows that all three inflation

Figure 6: Rolling-sample estimates of the first-order autocorrelation coefficient




measures display similar time-variation in their autocorrelation. All rise in the 1970s to 0.8 or higher, and remain there until the mid-1990s, at which point the correlation drops to 0.5 or 0.6. A third decline is evident in the mid-2000s, with first-order autocorrelation dropping to very low levels indeed, between 0 and 0.4. These simple measures support the conclusion that inflation is not currently well-characterized as a process with a unit root. The autocorrelations also suggest that there has been noticeable time-variation in inflation's reduced-form persistence over time.

¹⁸The rolling-sample estimates presented here employ a 14-year window, the same as in Pivetta and Reis.

2.2.4 Autocorrelation functions

Extending the previous section’s results, figure 7 displays the full autocorrelation function for five of the key measures of inflation (CPI-X, PCE-X, CPI, PCE and the GDP deflator) for two sample periods, 1966:Q1–1984:Q4 and 1985:Q1–2008:Q4.¹⁹ As the figure suggests, over the first half of the past 43 years,

Figure 7: Autocorrelations of inflation data, various measures and samples



inflation exhibited considerable persistence in this reduced-form sense, according to all measures. Since the mid-1980s, roughly corresponding to the onset of the “Great Moderation,” the persistence of inflation has declined for some, but not all, measures.²⁰ The ambiguity of these results and the source of any reduction

¹⁹I use both core and total measures here because, as discussed in the next subsection, the influence of key relative prices on the reduced-form persistence of inflation in the latter half of the sample can be significant. The beginning of the early sample is chosen to coincide with the first use of the federal funds rate as the Federal Reserve’s policy instrument. Extending the sample back to 1959:Q2, the earliest date for which the PCE price series is available, produces results that are qualitatively the same.

²⁰See Benati (2008) for an empirical investigation into changing inflation persistence.

that may have occurred is of considerable interest, and section 2.4 discusses this in some detail.

2.2.5 Dominant root of the univariate timeseries process

An alternative measure of inflation persistence is the dominant root implied by the univariate autoregressive process for inflation. In particular, if the autoregressive representation of inflation is of lag length k

$$\pi_t = c_1\pi_{t-1} + \dots + c_k\pi_{t-k} + \epsilon_t, \quad (2.6)$$

then the companion matrix for the state-space representation of π_t is

$$C \equiv \begin{bmatrix} c_1 & \dots & c_k \\ I_{k-1 \times k-1} & & 0_{k-1 \times 1} \end{bmatrix}, \quad (2.7)$$

and the root of C with the largest magnitude is the (dominant) root of interest.

Table 4 summarizes the results for our three measures of inflation for a variety of samples. The table shows a high degree of persistence over the past 25 years, with a modest decline as the earlier decades are dropped from the sample. The results are somewhat dependent on the inflation measure; studies that focus on the GDP deflator (such as Pivetta and Reis 2007) may well uncover less evidence of a decline in reduced-form persistence. Both the CPI and the PCE measures show a more pronounced decline in persistence, particularly in the post-1995 subsample.²¹ Another widely used measure of reduced-form persistence is the sum of the autoregressive coefficients, which approximates the long-run impulse response to a unit shock $c(1) \equiv \sum_{i=1}^k c_i$. This measure is provided in the right-hand-most column of table 4. As the table shows, the mapping from the sum of the c_i 's to the dominant root is not perfect, particularly when the coefficients imply a dominant complex pair (indicated by the *), as is the case for the CPI and the PCE in the latter subsample. In this case, even though the sum of autoregressive coefficients may be small, the oscillatory behavior implied by the complex pair of roots may die out only slowly, implying significant persistence.

Because persistent relative price shifts can influence the persistence of measures of overall inflation, particularly in shorter samples, the bottom panel of

²¹Standard errors are computed from a Monte Carlo that draws 100,000 permutations of the estimated residuals for each inflation measure and sample, creating a new inflation series for each such residual permutation given the original estimated autoregression coefficients, and re-estimating the autoregression and the dominant root for each permutation.

Table 4: Dominant root of autoregressive process for inflation

Measure	Sample	Dominant Root (Std. Error)	Sum of AR coeff.'s ($c(1)$) (Std. Error)	
CPI	66-08	0.94 (0.032)	0.89 (0.055)	
	66-84	0.92 (0.034)	0.88 (0.068)	
	85-08	0.70 (0.12)	0.41 (0.18)	
	95-08	0.64* (0.11)	-0.039 (0.32)	
	PCE	66-08	0.95 (0.029)	0.91 (0.046)
PCE	66-84	0.90 (0.034)	0.87 (0.065)	
	85-08	0.81 (0.094)	0.61 (0.14)	
	95-08	0.68* (0.11)	0.22 (0.26)	
	GDP	66-08	0.95 (0.027)	0.92 (0.039)
		66-84	0.87 (0.033)	0.84 (0.075)
85-08		0.86 (0.074)	0.72 (0.11)	
95-08		0.83 (0.12)	0.64 (0.17)	
CPI-X		66-08	0.93 (0.027)	0.90 (0.042)
	66-84	0.86 (0.041)	0.81 (0.081)	
	85-08	0.96 (0.039)	0.92 (0.054)	
	95-08	0.75 (0.0641)	0.62 (0.16)	
	PCE-X	66-08	0.95 (0.023)	0.93 (0.033)
66-84		0.86 (0.032)	0.84 (0.066)	
85-08		0.95 (0.044)	0.91 (0.061)	
95-08		0.72 (0.097)	0.49 (0.18)	

* denotes complex roots

table 4 presents the same results for the core CPI and PCE inflation measures. The core-based measures suggest less evidence of a decline in persistence in recent years. Neither CPI nor PCE inflation measures show a decline in the dominant root or the sum of the AR coefficients for the past 25 years. The dominant root estimate declines modestly for the period since 1995, but the standard error of this estimate is correspondingly larger. For the more recent and relatively short subsamples, these results suggest that large movements in key relative prices may well distort the extent to which underlying inflation persistence has changed, particularly if persistence is measured over a relatively short sample.

2.3 Evidence of changing reduced-form persistence in the United States

Recognizing the reduced-form nature of inflation persistence, a number of authors have looked for evidence that changes in the underlying determinants of inflation may have given rise to a change in reduced-form persistence. The leading cause of a change in inflation behavior is thought to be a change in the systematic behavior of the central bank. To make matters simple, consider the stylized, backward-looking model of inflation below

$$\begin{aligned}\pi_t &= \pi_{t-1} + ax_t \\ x_t &= -bf_t \\ f_t &= c\pi_t.\end{aligned}\tag{2.8}$$

The first equation is a skeletal “Phillips curve,” in which the change in inflation is positively related to a variable x_t , which we will take here to be the output gap. The output gap in turn depends negatively on the short-term policy rate f_t (for federal funds rate), and the policy rate is a positive function of inflation (with an implicit target inflation rate of 0). The solution for inflation is

$$\begin{aligned}\pi_t &= \alpha\pi_{t-1} \\ \alpha &\equiv \frac{1}{1+abc}.\end{aligned}\tag{2.9}$$

Inflation will follow a first-order autoregression, and will be less persistent—the coefficient α will be smaller—the larger is the policy response to inflation (c), the more responsive is the output gap to the policy rate (b), and the more responsive is inflation to the output gap (a). In this simple framework, a central bank that behaves more aggressively in moving inflation towards its target will reduce the

persistence of inflation. The intuition from this skeletal model generalizes to a number of more sophisticated models that include rational expectations and a richer description of the key elements sketched above.²²

2.3.1 Pre-War inflation persistence under the gold standard

If persistence can be expected to change as a result of changes in the monetary regime, a natural question is the extent to which inflation persistence changed in the United States under considerably different monetary regimes. The contrast in the United States between the post-World War II fiat money system and the pre-World War I gold standard provides a natural experiment. Barsky (1987) finds that while U.S. inflation persistence was very high from 1960 to 1979, it was virtually non-existent prior to World War I—indeed, Barsky’s ARIMA modeling of inflation during this period suggests that inflation is white noise.²³

2.3.2 A parameterized characterization of reduced-form persistence

Stock and Watson (2007) posit a relatively straightforward time series model of inflation that captures many of the features of reduced-form persistence discussed so far. The model can be expressed as an integrated moving-average process of order one, or IMA(1,1)

$$\Delta\pi_t = a_t - \Theta a_{t-1} \tag{2.10}$$

or equivalently as an unobserved components model with stochastic trend and stationary components²⁴

$$\pi_t = \tau_t + \eta_t \tag{2.11}$$

$$\tau_t = \tau_{t-1} + \epsilon_t, \tag{2.12}$$

²²A separate line of research examines the contribution of a time-varying inflation target to measures of inflation persistence. The paper will return to this topic in more detail in section 3.5 below.

²³Barsky uses wholesale price data prior to World War I. To the extent that such data reflect movements of commodity prices rather than consumer goods and services prices, this result may be confounded by an inherent difference between commodities and final goods prices, rather than a difference in persistence across the monetary regimes. Note that Barsky also makes the link between persistence and forecastability, a link exploited in Cogley, Primiceri, and Sargent (2007), discussed in section 2.3.3 below.

²⁴One can see the equivalence by substituting the definition for τ_t $\tau_t = \frac{\epsilon_t}{(1-L)}$ into the equation for inflation to obtain $\pi_t(1-L) = \epsilon_t + \eta_t(1-L)$, which, after re-arrangement, yields an equation like equation 2.10, with inflation an integrated process with a moving average error term.

where η_t and ϵ_t are uncorrelated with one another, and are serially uncorrelated with mean zero and variances σ_η^2 and σ_ϵ^2 , respectively. One can think of τ_t as reflecting the “permanent” or trend component of inflation, and η_t as capturing the stationary component.

Interestingly, Stock and Watson find that the decline in the variance of inflation is largely due to a marked decline in the variance of the permanent component, that is, a reduction in σ_ϵ^2 . It is still not possible, in their methodology, to reject the null of a unit root in inflation, but the variance of the shock that drives τ_t is currently at historic lows.

2.3.3 Multivariate evidence of changes in reduced-form inflation persistence

For the most part, the measures of persistence discussed so far are univariate measures; that is, they use only the information in an inflation time series to draw inferences about its persistence. However, some authors have argued that one can draw more accurate inferences using a multivariate approach. In part, the intuition behind this claim is connected to the “trend inflation” models discussed in section 3.5 below. In those models, the slow-moving or trend component of inflation accounts for much of the persistence in inflation. That trend in turn is most commonly associated with the central bank’s target rate of inflation. As a consequence, including variables that reflect the central bank’s inflation-targeting behavior, such as short-term policy rates, may help to identify both trend inflation and its persistence, and thus the persistence of inflation.

Cogley, Primiceri, and Sargent (2007) use a time-varying VAR to estimate the trend component of inflation, which they associate with the central bank’s inflation target. They find continued persistence in inflation, but they associate it strongly with trend inflation. This implies that the Federal Reserve’s implicit target for inflation continues to have a unit root (or near-unit root), although Cogley *et al.* estimate the variance of that component to have declined, consistent with the findings in Stock and Watson (2007). The persistence of the “inflation gap”—the difference between actual inflation and its trend—appears to have declined in recent years.

Methodologically, Cogley *et al.* introduce a new measure of persistence that is related to the predictability of near-term movements in the variable of interest. Formally, persistence is calibrated by the R^2 of the j -step-ahead forecast of the

variable. The higher is the R^2 , the more predictable it is, and thus the more persistent, precisely because past shocks have a persistent influence on future inflation. They examine the R^2 s for 1-, 4- and 8-quarter-ahead forecasts.

They find economically and statistically significant changes in the j -step-ahead R^2 s for the inflation gap from 1960 to 2006, with the 1-quarter ahead R^2 peaking at over 90 percent in the 1970s and early 1980s, falling to about 50 percent by the mid-1980s and through the end of their sample. The 4-quarter ahead R^2 s peak at 50 to 75 percent during the Great Inflation, and decline to about 15 percent more recently; the 8-quarter ahead R^2 s peak at 20 to 35 percent in the same period, falling to 10 percent more recently. All of these changes appear to be quite significant statistically, judged by the estimated joint posterior distribution of the R^2 s in earlier and later periods.

2.4 International evidence of changing reduced-form persistence

A number of authors have developed empirical evidence on changes in the reduced-form persistence of inflation for samples of developed countries. Benati (2008) surveys the evidence from a broad array of developed countries over long samples. His empirical work focuses on differences in estimated persistence across different monetary regimes, and his key hypothesis is that regimes that clearly anchor inflation (or the price level, as in the gold standard) induce less persistence in the inflation rate. Benati examines a number of European countries pre- and post-EMU, the United Kingdom, Canada and Australia pre- and post-inflation targeting, and the United States pre- and post-Volcker disinflation.

A key finding in the paper, summarized in table 5 reproduced from Benati (2008), is that reduced-form persistence has declined in recent years for all of the aforementioned countries that have adopted an inflation-targeting regime. The stand-out is the United States, which has not formally adopted such a monetary regime.²⁵ The conclusion from Benati's results is that all of the inflation-targeting countries exhibit a marked decline in inflation persistence during their inflation-targeting period, and that lower persistence is statistically quite different from the persistence exhibited prior to inflation targeting. The United

²⁵See Benati (2008), tables I–VIII.

Table 5: Estimates of reduced-form inflation persistence (from Benati 2008)

Country	Early sample	Late sample	Test of difference
U.K. (RPI)	Bretton Woods to inflation targeting 0.95	Inflation targeting -0.07	0.00
Canada (CPI)	1971 to inflation targeting 0.90	Inflation targeting -0.33	0.00
Euro area (GDP defl.)	Bretton Woods to EMU 1.01	EMU 0.35	0.00
U.S. (CPI) (PCE)	Great Inflation 0.77 0.74	post-Volcker 0.49 0.81	0.046 0.59

States and Japan (not displayed in the table) are stand-outs, and they are also the countries that have not adopted a formal inflation-targeting regime. Benati also reports similar results using the degree of “indexation” in a structural New-Keynesian model of the economy, that is, the estimates of μ in equation 3.7 below.²⁶

Levin and Piger (2004) also examine inflation persistence across a number of countries, focusing on the possibility that the reduced-form process for inflation has changed in recent years, perhaps owing to changes in the central bank’s inflation objective. Their results, which employ unknown breakpoint methods in both the classical and Bayesian traditions, show that simply allowing for a change in the mean of inflation appears to reduce estimated reduced-form persistence for many of the countries in their sample. Other international evidence develops mixed results about changing inflation persistence. Ravenna (2000) documents a large post-1990 drop in Canadian inflation persistence. O’Reilly and Whelan (2005) employ methods very similar to those in Levin and Piger, but find that for the Euro-area price indexes (as compared with Levin and Piger’s individual-country price indexes) there has been no discernible change in inflation persistence.

2.5 Conclusions from the reduced-form evidence

From both theoretical and empirical perspectives, it seems likely that the contribution to inflation from its unit root component has diminished significantly

²⁶Benati (2009) explores the stability of parameters that reflect intrinsic persistence across developed countries. In general, he finds that these parameters, whether explicit or implicit, are not stable across changes in monetary regime.

in recent decades. In most macroeconomic models, inflation would contain an important unit root (in terms of contribution to variance) if the central bank were not acting to keep inflation low and stable, consistent with either an implicit or explicit inflation target. From a practical perspective, this suggests that most macroeconomists can think of inflation as a stationary series that will (in normal times) return to the central bank's inflation goal in finite time, and that the central bank's inflation goal, while not written in stone (or anywhere else at present in the United States) is unlikely to vary significantly over time. Minor time-variation in the inflation goal could add a small unit root component to inflation, but its contribution to the variance of inflation will likely be small.

With regard to the specific autocorrelation properties of a stationary inflation rate, the picture is considerably murkier. All authors agree that in the United States and many other developed countries, inflation exhibited considerable persistence from the 1960s through the mid-1980s. After that time, the statistical evidence is mixed. For both the United States and other countries, studies fall on both sides of the argument about the possibility of declining reduced-form persistence. On a methodological note, for the United States, the evidence on changing persistence from so-called "core" measures appears to differ substantially from the evidence from so-called "headline" or total inflation measures. As a rule, the evidence of a change in persistence from core measures is less compelling than the evidence from headline measures.

Weighing all of the evidence, it seems reasonable to conclude that the persistence of inflation has declined somewhat in recent years. But how much it has declined and whether in the extreme case inflation is now a non-persistent series, remain issues for further study. At the time of writing, we are in the midst of an economic environment characterized by large relative price swings and significant changes in common estimates of the output gap and marginal cost, factors that are commonly thought to influence inflation. These conditions should provide data that will help economists to test a number of aspects of inflation dynamics, including its persistence.

3 Structural sources of persistence

While establishing the degree of reduced-form persistence in inflation is an important first step, knowledge of the degree of reduced-form persistence is of

limited use to a policymaker unless she can understand the underlying sources of reduced-form persistence. As she contemplates potential policy actions, the policymaker must be able to determine whether or not the persistence is structural and thus may be taken as a stable feature of the economic landscape. In order to know this, she must be able to parse the sources of persistence into three types: (1) those generated by the driving process, (2) those that are a part of the inflation process “intrinsic” to inflation (that is, persistence that is imparted to inflation independent of the driving process), and (3) those that are induced by her own actions or communications. With respect to the last source, the research cited above suggests that central banks that are more explicit about their inflation goal—and act in accordance with that commitment—may enjoy less persistence in their nations’ inflation rates.

Disentangling these sources of inflation persistence is extraordinarily difficult in relatively short aggregate time series. The paper will return to this issue later. To begin, it is important to distinguish theoretically among the potential sources of persistence in inflation. Significant differences in theoretical models will imply somewhat different ways of dissecting inflation persistence. We begin with the most widely used model of price-setting.

3.1 Persistence in the Calvo/Rotemberg model

Many modern models of inflation derive from the seminal contributions of Calvo (1983) and Rotemberg (1982,1983) and typically imply an Euler equation for inflation π_t that takes the form

$$\pi_t = \beta E_t \pi_{t+1} + \gamma x_t + \epsilon_t. \quad (3.1)$$

E_t denotes the mathematical expectation using information available in period t , and β denotes the discount rate. The variable x_t represents a measure of the output gap or marginal cost, depending on details of the model. The role of the shock term, denoted ϵ_t , will be explored in greater detail below. The parameter γ is a function of the underlying frequency of price adjustment and the discount factor.²⁷ By iterating expectations forward and assuming that the expectation

²⁷As Galí and Gertler (1999) demonstrate, $\gamma = \frac{(1-\theta)(1-\beta\theta)}{\theta}$. The more frequent is price adjustment, the larger is θ , and the smaller is the coefficient on the driving process.

at time t of future shocks $\epsilon_{t+i} = 0$, equation 3.1 can be expressed as

$$\pi_t = \gamma \sum_{i=0}^{\infty} \beta^i E_t x_{t+i} + \epsilon_t. \quad (3.2)$$

As this rendering makes clear, inflation is the sum of two components, the discounted sum of expected marginal cost (say) and a shock that is by assumption *iid*, but that can in principle be serially correlated. This formulation clarifies the motivation for this inflation specification: In a Calvo world in which prices are expected to be fixed for some time, price-setters who can re-set their prices set them equal to the discounted average of marginal cost that is expected to prevail over the expected life of the contract price.

3.2 The analytics of inflation persistence: “inherited” and “intrinsic” persistence

The expression above for inflation affords a natural taxonomy of the sources of inflation’s persistence. First, equation 3.2 implies that the inflation rate directly “inherits” the persistence in the variable x_t . If the output gap is a persistent series in the sense defined in section 2 above, then other things equal, inflation will inherit some of that persistence.²⁸

3.2.1 The baseline case

In the simplest case, inflation is given by equation 3.1 above, and the process for x_t is a univariate first-order autoregression with autoregressive parameter ρ :

$$\begin{aligned} \pi_t &= \beta E_t \pi_{t+1} + \gamma x_t \\ x_t &= \rho x_{t-1} + u_t \\ Var(u_t) &= \sigma_u. \end{aligned} \quad (3.3)$$

For simplicity, no shock perturbs the inflation Euler equation ($\epsilon_t = 0 \forall t$). In this case, one can show that the autocorrelation function for inflation is²⁹

$$A_i = \rho^i.$$

That is, inflation inherits exactly the autocorrelations of the first-order autoregressive process describing x_t . In this simple version of the model, the effects of

²⁸The reasons for persistence in output gap or marginal cost series, that is, the source of so-called “real rigidities,” is the subject of a number of papers. See, for example, Blanchard and Galí (2007), Fuhrer (2000), Smets and Wouters (2003). Section 3.2.6 provides some empirical results on the persistence of widely used driving processes for the canonical inflation models.

²⁹See Fuhrer (2006) for derivations.

monetary policy, real rigidities in consumption, or real wages—in short, the behavior of inflation arising from any aspect of the economy—must enter through their effects on x_t .³⁰

3.2.2 More complex cases

The autocorrelations of inflation become more complex when one allows for

- Non-zero shocks to the Euler equation ($\epsilon_t \neq 0$)
- Variation in the size of γ , given non-zero shocks, and
- The possibility of some “backward-looking” element to inflation, as in Fuhrer and Moore (1995) or Christiano, Eichenbaum, and Evans (2005).

These added complexities make the identification of underlying sources of persistence correspondingly complex, both theoretically and empirically. The following subsections derive the analytical results for these cases.

3.2.3 Shocks to the Euler equation

The augmentation of the Euler equation with an *iid* shock changes the interpretation of inflation persistence quite significantly, and in a way that is not well recognized in much of the literature on this subject. Modifying equations 3.3 to include this disturbance

$$\begin{aligned} \pi_t &= \beta E_t \pi_{t+1} + \gamma x_t + \epsilon_t \\ x_t &= \rho x_{t-1} + u_t \\ \text{Var}(e_t, u_t) &= \Sigma \equiv \begin{bmatrix} \sigma_e^2 & 0 \\ 0 & \sigma_u^2 \end{bmatrix} \end{aligned} \tag{3.4}$$

makes a subtle but important difference in the autocorrelation function for inflation. The presence of the *iid* shock ϵ_t now makes the behavior of inflation a mixture of its inherited persistence from current and expected future x_t , with weight γ , and the nonpersistent shock process (with implicit weight of one). The larger is the variance of the shock process, the more inflation looks like white noise, with zero persistence. The smaller is γ , the smaller will be the importance of x_t in determining the autocorrelation of inflation.

³⁰Models in which price changes are state- rather than time-dependent allow for inherited persistence as well. Dotsey, King, and Wolman (1999) and Burstein (2006) examine cases in which variations in the size and persistence of money shocks result in more or less persistent inflation responses to monetary shocks.

Normalizing the variance of the shock to the x_t process to 1 for convenience, one can express the inflation autocorrelations in this case as

$$\begin{aligned} A_i &= \frac{\rho^i \gamma^2}{a\sigma_e^2 + \gamma^2} \\ a &= (1 - \rho^2)(1 - \rho\beta)^2. \end{aligned} \tag{3.5}$$

Note that equation 3.5 indicates that the autocorrelations decay at rate ρ (the expression is premultiplied by ρ^i , and this is the only expression that varies with i). More generally, the expression suggests that inflation will be more autocorrelated

- The higher is ρ —that is, the greater is the persistence of the real driving variable x_t ;
- The higher is γ —the larger is the coefficient on the driving variable x_t , and thus the more of x_t 's persistence is inherited by inflation;
- The smaller is the variance of the shock e_t that disturbs inflation from the Euler equation.³¹

Table 6 below provides the first autocorrelation of inflation for various values of the key parameters in the simple inflation model.³² For values of γ that correspond to those estimated in the literature (generally below 0.1), a modest relative variance for ϵ_t can imply a fairly low first autocorrelation. For example, if γ is estimated to be 0.05, and the variances of the two shocks are the same, the first autocorrelation is 0.44. As the autocorrelation of the driving process approaches 1, the persistence of the driving process begins to dominate the white noise of the error term, as shown in the bottom panels of the table. Thus even in this very simple model, it is clear that a persistent driving process need not impart any persistence to inflation, depending on the sizes of γ and σ_e .

3.2.4 The pivotal role of the coefficient on x_t

As the analytical results of the preceding subsection suggest, the influence of x_t on inflation—the size of the parameter γ —is pivotal both in interpreting the sources of inflation persistence and in identifying equation 3.4 as a Phillips curve or aggregate supply relation. If $\gamma = 0$, then (a) the Euler equation can no

³¹Because we have normalized σ_u to 1, this should be interpreted as the smaller is σ_e relative to σ_u .

³²The full derivation for the results in the table is provided in Fuhrer (2006).

Table 6: Value of A_1 for selected values of σ_e^2 and γ

σ_e^2	γ				
	.01	.03	.05	.1	.2
$\rho = 0.9$					
0	0.90	0.90	0.90	0.90	0.90
0.1	0.25	0.70	0.81	0.88	0.89
0.3	0.10	0.48	0.68	0.83	0.88
0.5	0.06	0.36	0.59	0.79	0.87
1	0.03	0.23	0.44	0.71	0.84
3	0.01	0.09	0.22	0.50	0.75
5	0.01	0.06	0.14	0.39	0.68
$\rho = 0.95$					
.5	0.29	0.76	0.87	0.93	0.94
3	0.06	0.37	0.61	0.83	0.92
$\rho = 0.99$					
0.5	0.91	0.98	0.99	0.99	0.99
3	0.65	0.93	0.97	0.98	0.99

longer be interpreted as a Phillips curve; (b) the equation becomes decoupled from marginal cost or the output gap, and thus from monetary policy; (c) in many models, this decoupling will lead to indeterminacy for inflation, as monetary policy can no longer determine the steady-state value of inflation; and (d) inflation no longer inherits any persistence from x_t .³³

Given the centrality of the parameter γ , it is of interest to determine how well identified this parameter is in the data. The answer varies from study to study, but a brief empirical exercise may help to illuminate the potential problems. Consider generalized method of moments (GMM) estimates of an Euler equation that follows the format of Galí and Gertler (1999) in allowing for “rule of thumb” price-setters in addition to the Calvo price setters:

$$\pi_t = \lambda_b \pi_{t-1} + \lambda_f E_t \pi_{t+1} + \gamma x_t + \epsilon_t \quad (3.6)$$

The parameters λ_b and λ_f calibrate the amount of backward- and forward-looking price-setting behavior that influences inflation. Following Galí and Gertler (1999), we employ an instrument set that consists of four lags each of inflation, real marginal cost, a measure of the output gap, wage inflation, the

³³To see point (c), consider the simplified model in equations 2.8, and its solution in equation 2.9. When the Phillips parameter $a \rightarrow 0$, the solution for π_t becomes $\pi_t = \pi_{t-1}$. Inflation is a random walk, and thus fails to converge to any value in particular. This logic transfers to more complex specifications with explicit expectations.

spread of the 10-year Treasury constant maturity rate over the federal funds rate, and oil prices. Table 7 summarizes the results.

Only when allowing for 12th-order correlation in the weighting matrix does the

Table 7: Estimates of parameters in equation 3.4

Sample period: 1960–1997			
GMM estimates			
(HAC Standard errors in parentheses)			
Number of terms in weight matrix	λ_b	λ_f	γ
ma=4	-	0.99	-0.00067
	-	(0.012)	(0.0060)
ma = 12	0.34	0.65	0.0042
	(0.044)	(0.045)	(0.0042)
	-	0.99	0.0020
	-	(0.0088)	(0.0046)
	0.36	0.63	0.0065*
	(0.024)	(0.025)	(0.0033)
ML estimates			
(BHHH Standard errors in parentheses)			
	λ_b	λ_f	γ
	0.51	0.48	0.0047
	(0.027)	(0.031)	(0.0028)
Bayesian estimates			
(Max. of posterior, estimated posterior sd's in parentheses)			
	λ_b	λ_f	γ
	0.51	0.48	0.0184
	(0.044)	(0.044)	(0.014)

coefficient on marginal cost enter with the correct sign and significantly at the 5 percent level. These results are provided as suggestive of the difficulties in identifying γ ; they are broadly consistent with the aggregated results found in Galí and Gertler (1999) and Rudd and Whelan (2006).³⁴

For comparison, the lower panel of the table provides maximum likelihood and Bayesian estimates of the same model, augmenting equation 3.6 with vector autoregressive equations for the variables employed as instruments above.³⁵

³⁴See Mavroeidis (2005) for a careful treatment of the difficulties in identifying New-Keynesian Phillips curves.

³⁵The Bayesian priors for the three parameters are conventional, with generalized beta densities for the λ 's, and a gamma density for γ . The prior distributions for the three param-

The VAR coefficients are taken as fixed from OLS estimates over the same sample, and the ML estimates of the λ 's (the weights on lagged and expected inflation in the hybrid model) and γ are presented in the table, along with BHHH standard errors. Once again, the estimates of γ are quite small and not significantly different from zero. Note that the ML estimate of the backward-looking component is somewhat larger than the GMM estimate, and is quite precisely estimated. The likelihood-ratio test for the restriction that λ_b and λ_f take the GMM values (with γ freely estimated) has p -value 0.0000. As we will see in the section on “trend inflation” models below (section 3.5), a pattern is emerging in which more tightly constrained models provide larger estimates of intrinsic inflation persistence.

3.2.5 Hybrid models of inflation and “intrinsic” persistence: Including lagged inflation

The debate over the empirical success of the basic specification summarized in equation 3.1 continues, with the ability of the specification to replicate the reduced-form persistence of inflation an important focus. A number of authors have proposed rationales for the presence of a lagged inflation term in their aggregate supply relation (*aka* intrinsic persistence), through indexation of price contracts (see Christiano, Eichenbaum, and Evans (2005)), “rule-of-thumb” behavior (see Galí and Gertler (1999)), alternative contract assumptions (see Fuhrer and Moore (1995)), alterations to the Calvo framework that assume a rising, rather than a constant hazard for the ability to reset prices (see Mash (2004) and Sheedy (2007)), or alternatives to rational expectations (see section 3.7, Orphanides and Williams (2004) and Roberts (1997)). Woodford (2007) provides a very helpful summary of the state of modeling intrinsic inflation persistence.³⁶

An augmented Phillips curve specification that allows for the influence of

eters are centered on [0.5, 0.5, 0.05], respectively, with standard deviations of [0.2, 0.2, 0.02], respectively. The posterior distributions are estimated using a Markov-Chain Monte Carlo algorithm, with four simulation blocks of 200,000 draws each.

³⁶Disaggregated price data provide limited support for many of these theoretical rationales for intrinsic persistence. The data examined in Bils and Klenow (2004) and Nakamura and Steinsson (2008), for example, provide little evidence that prices rise at a roughly constant rate between more significant resets, as might be implied by firms following a rule of thumb, or indexation.

lagged inflation takes the form

$$\pi_t = \mu\pi_{t-1} + (\beta - \mu)E_t\pi_{t+1} + \gamma x_t + \epsilon_t, \quad (3.7)$$

with the rest of the specification as detailed in equation 3.4. In a sense that is central to this paper, the presence of lagged inflation provides an augmented channel for what might be called “intrinsic” inflation persistence, that is, persistence that is not inherited from the driving process x_t . In the simpler model of equation 3.4, the *iid* shock ϵ_t provided a trivial source of intrinsic persistence; depending on the relative variance of that shock and the coefficient on x_t , the persistence inherent in x_t would be more or less inherited by inflation.

With a lag of inflation added in equation 3.7, any shock to the Euler equation will persist for longer, other things equal, independent of the evolution of the driving variable. A shock to inflation becomes part of the history of inflation, independent of shocks to x_t . In addition, the forward-looking component of the model incorporates the direct dependence on history, augmenting the direct effect on persistence. One can think of the model as comprising both rule-of-thumb and sophisticated forward-looking price-setters. The forward-looking price-setters, in forming an expectation for future inflation, must take into account the behavior of the rule-of-thumb price setters, who set current prices based on lagged inflation. Thus the sophisticated price-setters’ behavior reinforces the behavior of the rule-of-thumb price-setters.

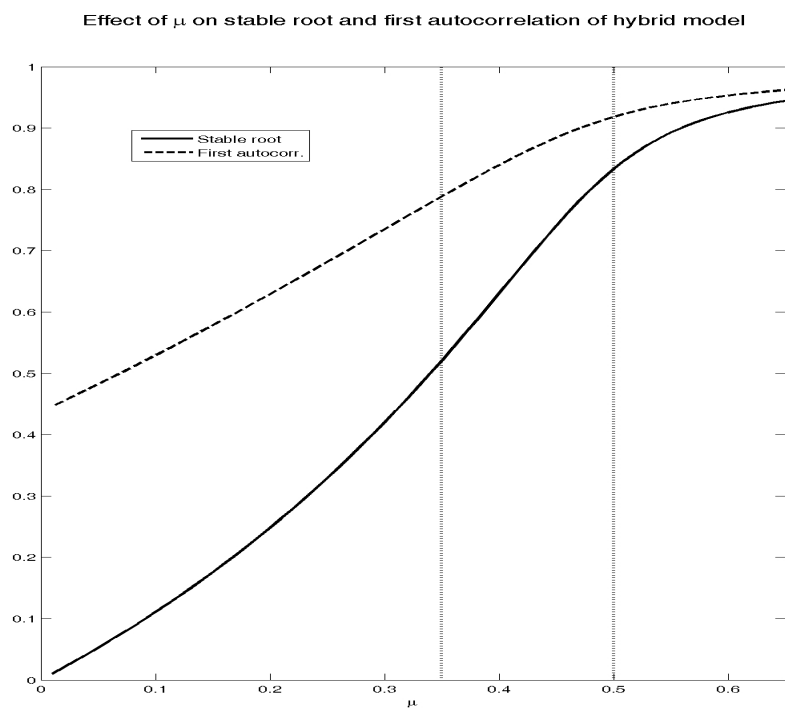
To see this algebraically and graphically, consider the analytic expression for the autocorrelations of the model augmented with a lag:

$$A_1 = \frac{a}{b\sigma_e^2 - c\rho\mu}d + \lambda_s, \quad (3.8)$$

where $[a, b, c, d]$ are functions of the stable root λ_s (in turn a function of β and μ) and the other underlying structural parameters of the model. Fuhrer (2006) shows that the additive term in λ_s dominates A_1 , and that λ_s and thus A_1 rise monotonically with μ . Figure 8 plots the stable root and the first autocorrelation of inflation as a function of μ .³⁷ In generating this figure, we set $\beta = 0.98$, $\gamma = 0.05$, $\rho = 0.9$, and $\sigma_e^2 = 1$. The figure shows that the first autocorrelation rises rapidly from about 0.4 to above 0.9 as μ increases from 0

³⁷Taken from Fuhrer (2006), Figure 2. Note that in this case, while the algebra is a bit messier, it can also be shown that the autocorrelations decay at rate ρ , so ρ and the first autocorrelation are sufficient statistics for the autocorrelation function.

Figure 8: Dependence of stable root and first autocorrelation on μ



to 0.6. This figure thus emphasizes the role that forward-looking behavior plays in augmenting the direct effect of lagged inflation in the model.

Table 8 provides the first autocorrelation for this model for a variety of parameter settings. In particular, the size of μ varies from 0.1 to 0.9, and the relative variance of ϵ varies from 0 to 5. In addition, the table displays the sensitivity of A_1 to the value of ρ .

Table 8: Value of A_1 for selected values of σ_ϵ^2 and μ

$\gamma=0.05, \beta=0.98, \rho=0.9$						
σ_ϵ^2	μ					
	0.0	0.1	0.3	0.5	0.7	0.9
0	0.90	0.92	0.96	0.99	1.00	1.00
0.3	0.68	0.74	0.86	0.96	0.98	0.99
0.5	0.59	0.66	0.82	0.94	0.97	0.98
1	0.44	0.53	0.74	0.92	0.97	0.98
2	0.29	0.39	0.64	0.89	0.96	0.98
3	0.22	0.32	0.59	0.88	0.96	0.98
5	0.14	0.25	0.54	0.86	0.96	0.98
$\rho = 0.5$						
σ_ϵ^2	μ					
	0.0	0.1	0.3	0.5	0.7	0.9
0	0.5	0.58	0.76	0.94	0.99	0.99
0.3	0.02	0.13	0.44	0.84	0.96	0.98
0.5	0.012	0.12	0.43	0.84	0.96	0.98
1	0.0063	0.12	0.43	0.84	0.96	0.98
2	0.0032	0.11	0.42	0.83	0.96	0.98
3	0.0021	0.11	0.42	0.83	0.96	0.98
5	0.0013	0.11	0.42	0.83	0.96	0.98

As shown in the top panel of the table, for relatively high values of σ_ϵ^2 that significantly lower the first autocorrelation in the purely forward-looking model (the first column of the table, $\mu = 0$), a modest lag coefficient dramatically raises the persistence of inflation. The lower panel of the table suggests that even with very little inherited persistence— $\rho = 0.5$ —a moderate value of μ implies a high degree of inflation persistence. An important implication of this table is that it may be difficult to distinguish among sources of inflation persistence, as the reduced-form implications of an inflation process that inherits a highly persistent driving process can be nearly the same as those of an inflation process that inherits little persistence from the driving process, but has a modest amount of

intrinsic persistence.

3.2.6 The persistence of the driving process

Most researchers will agree that the observed persistence of inflation is determined at least in part by the inherited persistence of the driving process. Thus, in thinking about potential structural causes of a change in reduced-form persistence, a natural (but so far largely unexplored) question is to what extent there have been changes in the persistence of the driving process.³⁸ In this section, we employ many of the same persistence measures that are used above for inflation. We consider three candidates for the driving process: a measure of real marginal cost, proxied by the labor share (or equivalently, real unit labor costs) of the nonfarm business sector; and two measures of the output gap, the first a Hodrick-Prescott detrended log GDP gap, and the second the log deviation between real GDP and the Congressional Budget Office’s estimate of potential GDP. We look at several subsamples. The first autocorrelation, the sum of the autoregressive coefficients, and the dominant root of the autoregressive process are displayed for each driving process and each subsample. The results are summarized in table 9. As the table indicates, there is remarkable stability

Table 9: Estimated persistence of driving variables

Driving variable	1966–2008	66–83	84–08	95–08
First autocorrelation				
Real mc	0.92	0.80	0.92	0.90
HP gap	0.85	0.86	0.84	0.80
CBO gap	0.92	0.91	0.92	0.90
Sum of the autoregressive coefficients				
Real mc	0.94	0.78	0.96	0.94
HP gap	0.78	0.77	0.80	0.80
CBO gap	0.91	0.90	0.95	0.97
Dominant Root of the autoregressive process				
Real mc	0.95	0.69	0.97	0.96
HP gap	0.81	0.82	0.79	0.79
CBO gap	0.78	0.78	0.82	0.87

in persistence measures across subsamples in all three proxies for the driving

³⁸As demonstrated above, a change in the coefficient on the driving process or in the relative variances of the inflation shock and the shock to the driving process may also affect the persistence of inflation.

process, for all three measures of persistence. This table suggests little evidence of a change in persistence for the driving variables most commonly associated with inflation.³⁹ It also suggests that the stronger hypothesis that inflation has lost all its autocorrelation—that is, inflation is an *iid* time series—is hard to justify. Unless the driving process is utterly decoupled from inflation, inflation must inherit some of its persistence. If inflation were completely decoupled from output or marginal cost, we would need an entirely new theory of inflation that steps outside the historical tradition of Phillips, Gordon, Calvo, and Rotemberg.

This de-coupling between the (still ambiguous) evidence of declining reduced-form inflation persistence and a relatively stable and persistent driving process may help guide the search for structural interpretations of possibly changing inflation persistence. One simple interpretation is that the evidence for changes in reduced-form persistence is weak, and the stable, high persistence of the driving variables is consistent with that observation. Another is that while the inherited persistence from the driving process may not have changed much, the importance of the lagged inflation term in Phillips curves has diminished, leading to diminished intrinsic persistence in the face of unchanged inherited persistence. Of course, it may also be that the Phillips slope parameter and relative error variances have changed, which will affect the extent to which the unchanged persistence in the driving process is inherited by inflation. The next section examines a number of these possibilities from the perspective of a standard dynamic stochastic general equilibrium (DSGE) model.

3.3 Using a DSGE model to interpret structural sources of persistence

The skeletal model summarized in equations 3.4 and 3.7 provides important insights into some of the structural sources of inflation persistence. However, the model leaves implicit the determination of real output and the role of monetary policy in influencing output and inflation. In theory, both the systematic component of monetary policy and the nature of the transmission of policy through the real side can have significant effects on the dynamic properties of inflation.

³⁹Of course, a number of authors have suggested that this standard proxy for real marginal cost is imperfect, and others have derived model-based measures of the output gap that can differ significantly from the simple measures used here. The results presented above are suggestive, and further research is warranted.

In this section, we explore the quantitative effects on inflation of various aspects of an articulated dynamic stochastic general equilibrium (DSGE) model.

The model remains relatively simple. It comprises the “hybrid” inflation model discussed above; an optimizing I-S curve that links real output to expected short-term real interest rates, allowing for real rigidity in the form of a lagged output term that can be motivated by the presence of habits in the consumer’s utility function;⁴⁰ and a canonical policy or Taylor (1993) rule that makes the short-term policy interest rate a function of deviations of inflation and output from their desired levels. The last also allows for the possibility of interest-rate smoothing.

The model can be summarized in the three equations⁴¹

$$\begin{aligned}
 \pi_t &= \mu\pi_{t-1} + (1 - \mu)E_t\pi_{t+1} + \gamma\tilde{y}_t + \epsilon_t \\
 \tilde{y}_t &= \mu_y\tilde{y}_{t-1} + (1 - \mu_y)E_t\tilde{y}_{t+1} - y_\rho(r_t - E_t\pi_{t+1}) + u_t \\
 r_t &= \rho r_{t-1} + (1 - \rho)[a_\pi\pi_t + a_y\tilde{y}_t].
 \end{aligned}
 \tag{3.9}$$

The two shocks to the system, ϵ_t and u_t , are assumed to be independent and *iid* with diagonal covariance matrix $\begin{bmatrix} \sigma_\epsilon^2 & 0 \\ 0 & \sigma_u^2 \end{bmatrix}$. The baseline parameters for the model are displayed in the second column of table 10.

Table 10: Baseline and alternative parameter sets for DSGE model

Parameter	Baseline value	Alternate value
μ	0.50000	0
π_y	0.10000	0.025
μ_y	0.50000	0
y_ρ	0.10000	
ρ	0.80000	0
a_π	1.50000	5
a_y	0.50000	
σ_ϵ^2	0.5	0.1
σ_u^2	0.5	

We vary the values of the key parameters in the model to gauge the effect of changes in the behavior of monetary policy, as well as changes in the price-setting

⁴⁰See, for example, Fuhrer (2000).

⁴¹While the model affords a more structural decomposition of inflation persistence than is possible using the model in equations 3.4 and 3.7, it still abstracts from some potentially important influences on inflation dynamics. The model does not allow for “trend inflation”; see section 3.5 below. The model uses the output gap, rather than the more current marginal cost measure, and thus ignores the role of wages and productivity in the inflation process.

and output sectors, on the autocorrelations of inflation. The goal is to determine the extent to which the persistence of inflation—whether fixed or changing over time—can plausibly be attributed to specific underlying structural features, or to changes in those features.

Figure 9 displays inflation’s autocorrelation function for a variety of parameter configurations. The autocorrelation function corresponding to the baseline parameters in table 10 is displayed in the solid line in both panels. It mirrors the properties of the autocorrelation functions displayed in figure 7: The autocorrelations are high for the first several quarters, decaying gradually toward zero and turning negative for several quarters thereafter.


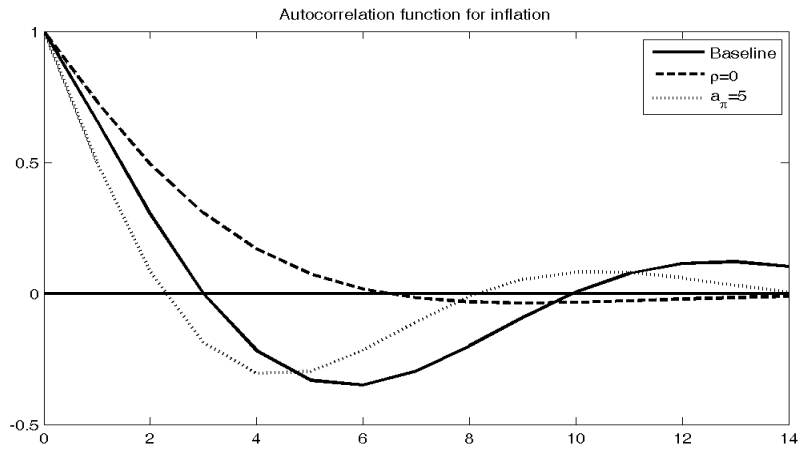
The top panel displays the autocorrelation function when the parameters governing the central bank’s behavior are altered. A dramatic shift in the emphasis on inflation, a_π —from the conventional 1.5 to 5.0—reduces the autocorrelations of inflation noticeably (the dotted line in the figure).⁴² The first autocorrelation of inflation is reduced from about 0.65 to 0.5, and subsequent autocorrelations drop a bit more rapidly below zero. But the difference arising from more aggressive inflation-fighting is not dramatic, especially given the threefold increase in the policy response to inflation. Thus, while a dramatic change in monetary policy might well account for some of the reduction in reduced-form inflation persistence, one would not want to overstate this structural source of changes in inflation persistence.

A somewhat larger change is induced by removing interest-rate smoothing in the policy rule ($\rho = 0$, the dashed line in the top panel of the figure.) Now the autocorrelations decay smoothly toward zero after about six quarters. The first autocorrelation of inflation is above that of the baseline case, but the effect on inflation of shocks after about five quarters is essentially zero. This suggests that a less inertial central bank could have removed some of the longer-lived, second-order oscillations that characterized inflation in the 1960s and 1970s.

The bottom panel of the figure displays inflation autocorrelations when other aspects of the model economy are altered from the baseline. The most striking change occurs when the backward-looking components of the inflation and output equations are eliminated, as shown in the dotted line. With purely forward-looking output and inflation equations ($\mu = \mu_y = 0$), the inflation au-

⁴²A similar result, not shown, is obtained for a large shift in the emphasis on output in the policy rule.

Figure 9: Effect of key structural parameters on inflation persistence, DSGE model



to correlations jump quickly to zero (reminiscent of the disinflation simulations in section 1.2), hovering near zero for all periods after the first. The results are almost identical if we only shut off the backward-looking component of the inflation equation ($\mu = 0$), as shown in the dashed line in the bottom panel. Smaller effects are achieved by reducing the size of the output effect on inflation ($\pi_y = 0.025$), and smaller still is the effect of reducing the relative size of the inflation shock (“small inflation shock”), the circles in the bottom panel.

Overall, these simulations suggest that while all of the aspects of the economy captured in the stylized model likely contribute to the persistence of inflation, the most potent effects arise from the lag coefficients in the inflation equation. More inflation-responsive monetary policy has some effect, as does the slope of the Phillips curve (π_y), but these effects are small compared to the effect from eliminating the intrinsic persistence in the Phillips curve. While these conclusions may well vary somewhat depending on details of the DSGE specification, the results suggest that it may be inappropriate to attribute too much persistence—or too much of the change in persistence—to the behavior of monetary policy. The largest effects may be attributed to the “hybrid” portion of the aggregate supply equation.

3.4 Persistence in state-dependent models of inflation

Most of the models in this paper employ a time-dependent pricing convention, that is, the probability that a firm will adjust its price is solely a function of time, not of economic conditions. An important and appealing alternative is that the time for adjusting prices is endogenously chosen by firms in response to economic conditions. In general, the literature on state-dependent pricing (SDP) has focused relatively little attention on the issue of inflation persistence. While the early models of Caplin and Leahy (1991) and Dotsey, King, and Wolman (1999) focused on the search for “persistence mechanisms,” that focus centered largely on the difficulty in developing a non-zero and persistent output response to a one-time monetary shock. The intuition behind this difficulty is straightforward. In a classic (S,s) model of SDP motivated by a fixed cost to changing prices, modest to large monetary shocks could push all firms to their (S,s) boundary, and consequently all prices would adjust one-for-one with money: Money is neutral. For smaller monetary shocks, a smaller fraction of

firms would adjust prices, so money could have a small aggregate effect. Thus, it can be difficult to avoid monetary neutrality in these models; persistent effects of monetary shocks on prices or inflation are even harder to come by.

Burstein (2006) suggests a variant of the SDP paradigm in which firms choose a price *path*, rather than a fixed price level, when they hit their (S,s) boundary. Equivalently, the firm faces a fixed cost to adjusting its price path rather than the level of prices.⁴³ Thus, a monetary shock that forces a firm to its boundary will lead the firm to set a sequence of price changes that in turn implies a persistent response of inflation and output to a monetary shock. Burstein’s model is capable of producing inflation responses to changes in the money growth rate that exhibit the hump shape typically found in the VAR literature.⁴⁴

A recent paper by Bakhshi, Khan, and Rudolf (2007) develops a Phillips curve from the Dotsey, King, and Wolman SDP framework, and examines the dynamics of inflation. Interestingly, that Phillips curve includes lagged inflation terms, reflecting the fact that optimal relative prices set in period t depend on optimal relative prices set in previous price vintages (see their equations 9–11). They find that the persistence of inflation implied by the SDP-based Phillips curve is significantly lower than that implied by the time-dependent New-Keynesian Phillips curve, largely because the persistence induced by lagged inflation in the SDP Phillips curve is more than offset by the number of price-setters who reset following a shock.

While recent theoretical developments in SDP pricing models such as those in Burstein (2006) and Bakhshi *et al.* are promising, the empirical literature based on state-dependent models is considerably less well developed. As a consequence, there are few empirical results on inflation persistence that map well into the reduced-form and structural persistence concepts developed in this paper. Overall, it seems fair to conclude that the theoretical results suggest that SDP models are likely to provide a less compelling explanation of reduced-form persistence than hybrid versions of the time-dependent models.

⁴³Calvo, Celasun, and Kumhoff (2002) develop a related, time-dependent pricing model in which firms also choose price paths when they are allowed to reset prices.

⁴⁴A fuller empirical examination of Burstein’s specification would be of interest. A key question in this regard is how the price paths set by resetting firms are allowed to respond to different shocks—mark-up shocks, shocks to the central bank’s inflation target, and so on. In the model of Calvo, Celasun, and Kumhoff (2002), it can be difficult to obtain data-consistent impulse responses to all shocks without making different assumptions about when price paths can and cannot respond to shocks. For more on this issue, see Fuhrer (2008).

3.5 Persistence in models of “trend inflation”

A series of papers beginning with Cogley, Primiceri, and Sargent (2007) and Cogley and Sbordone (2009) emphasizes the importance in modeling inflation of recognizing the slowly moving component of inflation that they dub “trend inflation.” The Cogley-Sbordone paper introduces two innovations. First, because long-run inflation is not a constant, the typical simplifications that give rise to the log-linearized Calvo model of equation 3.1 no longer apply. The standard log-linearization depends on the constancy of the long-run value of inflation.⁴⁵ Cogley and Sbordone derive the log-linear approximation that is appropriate when the long-run value of inflation has a trend. Hatted variables in the next equation denote deviations from steady-state values; for inflation, this implies the deviation of inflation from trend inflation

$$\hat{\pi}_t = \tilde{\rho}_t(\hat{\pi}_{t-1} - \hat{g}_t^\pi) + \zeta_t \widehat{mc}_t + b_{1t} \tilde{E}_t \hat{\pi}_{t+1} + b_{2t} \tilde{E}_t \sum_{j=2}^{\infty} \varphi_{1t}^{j-1} \hat{\pi}_{t+j} + b_{3t} \tilde{E}_t \sum_{j=0}^{\infty} \varphi_{1t}^j (\hat{Q}_{t+j,t+j+1} + \hat{g}_{t+j+1}^y) + u_t \quad (3.10)$$

where \hat{g}_t^π and \hat{g}_t^y are the innovation to trend inflation and the growth rate of real output, respectively; \widehat{mc} is the deviation of real marginal cost from its steady state; and $\hat{Q}_{t+j,t+j+1}$ is the one-period discount factor between periods $t+j$ and $t+j+1$. A key parameter in this specification is $\tilde{\rho}_t$, which calibrates the degree to which a lagged inflation term is required to match the autoregressive properties of inflation, once trend inflation is accounted for. Regardless of the estimate of $\tilde{\rho}$, this specification is useful for researchers who wish to allow for time variation in the steady-state value of inflation (such as a time-varying inflation target).

Second, Cogley and Sbordone find that the point estimate for $\tilde{\rho}_t$, the coefficient on lagged inflation in this specification, centers on zero. That is, once the model has accounted for the slow-moving variation in trend inflation, there is no need for a lag of inflation to account for the reduced-form persistence of inflation. While this empirical finding is controversial, the concept of trend inflation, which the authors associate with the central bank’s time-varying inflation goal, is an important contribution to the inflation literature.

⁴⁵See, for example, Woodford (2003) for the derivation.

3.5.1 Cogley and Sbordone’s measure of trend inflation

Table 11 displays the first autocorrelation for actual and de-trended inflation, using Cogley and Sbordone’s measure of trend inflation, for the subsamples in table 1 of their paper.⁴⁶ The table shows that the autocorrelations of detrended

Table 11: First autocorrelation of inflation

Sample	Detrended π	Raw Data
60:Q3–03:Q4	0.81	0.89
60:Q3–83:Q4	0.83	0.88
84:Q1–03:Q4	0.36	0.56

inflation are somewhat lower than those of the raw inflation data. However, in contrast to the table in Cogley and Sbordone (2009), the differences are small for the first two samples. For the most recent sample, the autocorrelation declines *both* for the detrended and for the raw series. Most of the decline in the detrended data’s autocorrelation can be explained by a corresponding decline in that of the raw data. This suggests that there are other, equally important factors at work in explaining the persistence of inflation, and in explaining changes in persistence over time.

While not presented in Cogley and Sbordone (2009), it is of interest to examine their model’s implication for the first autocorrelation of inflation. Using values of their key parameters that center on the median estimates over time presented in figure 4 of their paper, that is, $\tilde{\rho}_t \approx 0$, $\zeta_t = 0.03$, $b_1 = 0.9$, $b_2 = 0.02$, and $b_3 = 0$, and assuming a diagonal covariance matrix with variances estimated from a VAR over their sample period, I obtain a first autocorrelation for inflation of 0.22.⁴⁷ This estimate differs markedly from their data’s first autocorrelation over the same sample, 1960–2003. Matching the autocorrelation of inflation from the data requires a significantly different set of parameters—for example,

⁴⁶Inflation is measured as four times the log change in the GDP deflator, as in Cogley and Sbordone (2009). The measure of trend inflation was kindly provided by the authors and replicates that in Figure 1 of their paper.

⁴⁷Note that setting b_1 to 1.05 as in their figure 4 implies multiple solutions. I reduce the value of b_1 to 0.9 to keep it as high as possible, while still obtaining a unique solution to the model. The VAR employed in the exercise includes four lags each of the inflation rate from the GDP deflator, real marginal cost defined as in Cogley-Sbordone, the federal funds rate, and an output gap defined as the log difference between real GDP and Hodrick-Prescott filtered real GDP. The VAR equations for marginal cost, the funds rate, and the output gap are used in conjunction with equation 3.10 in order to compute stability conditions, and to compute autocorrelations given the estimated covariance matrix of the shocks.

setting $\tilde{\rho}$ near unity raises the first autocorrelation toward 0.8.

A recent paper by Barnes, Gumbau-Brisa, Lie, and Olivei (2009) examines the robustness of the finding that the detrended inflation model implies a value for $\tilde{\rho}_t$ of 0. Details of Cogley and Sbordone’s estimation procedure make a significant difference to the estimation results. Barnes *et al.* show that simply changing the form of the Euler equation—which implicitly imposes an additional constraint that is implied by the model—completely reverses the finding on $\tilde{\rho}_t$. They develop a precise estimate of this key “intrinsic persistence” parameter of about 0.8. This finding suggests caution in interpreting the rather striking implications of trend inflation models for inflation persistence.

3.6 Persistence in sticky-information models

Mankiw and Reis (2002) propose a model in which information, rather than prices, is sticky. In essence, the model applies the Calvo machinery to the updating of information, rather than that of prices. In the Mankiw and Reis (2002) model, a fraction of price setters get to update their information in each period in a manner analogous to that of the Calvo model. As a consequence, the age or vintage of price-setters’ information sets will be described by a geometric distribution, as is the case for the duration of price contracts in the Calvo setting.

The model implies a Phillips curve that links inflation to output and a geometric weighted average of lagged expectations of inflation and output (see Mankiw and Reis (2002), p. 1300):

$$\pi_t = \left[\frac{\alpha\lambda}{1-\lambda} \right] y_t + \lambda \sum_{j=0}^{\infty} (1-\lambda)^j E_{t-1-j}(\pi_t + \alpha\Delta y_t). \quad (3.11)$$

As the authors emphasize, in this Phillips curve it is past expectations of current conditions, rather than current expectations of future conditions, that determine inflation. The Mankiw-Reis Phillips curve is thus a close cousin to those of Fischer (1977) and Koenig (1996). The Mankiw-Reis paper documents a number of desirable features of their model, based on simulations in response to a permanent change in the level of demand, the growth rate of demand (a stylized “disinflation”), and an anticipated drop in the growth rate of demand.


In this model, one can see by inspection (and the authors verify) that inflation will inherit the persistence of the output process. The authors compute the

autocorrelations of inflation under the assumption that a simple quantity equation holds, trivially linking money to output, and that money growth follows an autoregressive process

$$\begin{aligned} m_t &= p_t + y_t \\ \Delta m_t &= 0.5\Delta m_{t-1} + \varepsilon_t. \end{aligned} \tag{3.12}$$

Under these assumptions, the autocorrelations for inflation are indeed quite high. Figure 10 displays the autocorrelations taken from table I of their paper (the solid line). The dashed line in the figure displays the autocorrelations

Figure 10: Autocorrelations of inflation, Mankiw-Reis model



when one allows for supply shocks (shocks to the Phillips curve). As discussed in section 3.2.3 above, the autocorrelation in the presence of supply shocks will depend on the variance of supply shocks relative to the shocks to the driving process. Here, we have set the ratio of supply shocks to money growth shocks at 0.25. As the figure indicates, supply shocks with relatively modest variance dramatically alter the implications for inflation’s autocorrelations in this model, as in the Calvo/Rotemberg model.

3.7 Persistence in learning models

When the agents that “populate” the models discussed above know the structure and parameters of the model, they can use that knowledge to form expectations by taking the mathematical expectation of the variable of interest. That is the essence of rational expectations, and up to this point, the structural models in this section have assumed rational expectations.

However, a long tradition dating back at least to Bray (1982) and Marcat and Sargent (1989) examines the dynamics of macroeconomic variables when the agents in the model lack perfect knowledge of the model and consequently must learn about their environment. Orphanides and Williams (2004) and Williams (2006) explore monetary economies in which agents must learn about their economic environment. Learning can significantly alter the dynamics of an otherwise standard macroeconomic model. Consider the simple case in which inflation is governed by a two-equation model comprising a Calvo-like Phillips curve and a reduced-form equation for output, as in equations 3.3 in section 3.2.1 above:

$$\begin{aligned}\pi_t &= F_{t-1}\pi_{t+1} + \gamma x_t \\ x_t &= \rho x_{t-1} + u_t.\end{aligned}\tag{3.13}$$

$$\tag{3.14}$$

The key difference is that expectations of inflation are not the mathematical expectation, denoted by the E_t operator, but are reasonable forecasts given the information available to the agents at time $t-1$, denoted by F_{t-1} . One plausible way to formalize learning, posited in Williams (2006), is for agents to estimate the reduced-form of the model using an adaptive estimation rule. The unique and stable reduced-form solution of this model under rational expectations is given by

$$\begin{bmatrix} \pi_t \\ x_t \end{bmatrix} = \begin{bmatrix} 0 & \rho \frac{\gamma}{1-\rho} \\ 0 & \rho \end{bmatrix} \begin{bmatrix} \pi_{t-1} \\ x_{t-1} \end{bmatrix}.\tag{3.15}$$

As discussed earlier, the model implies no dependence on lagged inflation, although inflation will indeed exhibit persistence to the extent that x_t does. But if agents are not endowed with knowledge of the solution coefficients under rational expectations, they may instead attempt to estimate a reduced-form equation for inflation and x_t such as

$$\begin{bmatrix} \pi_t \\ x_t \end{bmatrix} = \hat{A}_t \begin{bmatrix} \pi_{t-1} \\ x_{t-1} \end{bmatrix} + e_t.\tag{3.16}$$

If their initial estimate for \hat{A}_t coincides with the solution in equation 3.15, then the model will behave as under rational expectations. But in general, the model will exhibit different dynamics, governed in part by agents' current estimate \hat{A}_t in equation 3.16. If e_t is unforecastable from period t , then equation 3.16 implies that

$$F_{t-1}\pi_{t+1} \equiv e_\pi \hat{A}^2 \begin{bmatrix} \pi_{t-1} \\ x_{t-1} \end{bmatrix}, \quad (3.17)$$

where e_π selects the row of \hat{A} corresponding to the inflation equation in the reduced form. Substituting this expectation for inflation in equation 3.13, we obtain⁴⁸

$$\pi_t = \hat{a}_{11}\pi_{t-1} + \hat{a}_{12}x_{t-1} + \gamma x_t. \quad (3.18)$$

Depending on the current estimated values of \hat{a}_{ij} , inflation may now exhibit some intrinsic persistence, and will inherit—more or less, depending on the value of \hat{a}_{12} —the persistence of x_t .⁴⁹

This very stylized example makes it clear that learning can add another layer of dynamics to inflation. If agents use a forecasting rule that differs from the rational expectations solution, they can add intrinsic persistence to an otherwise forward-looking model that implies no such persistence.

4 Inference about persistence in small samples: “Anchored expectations” and their implications for inflation persistence

Implicit in the analysis of Benati (2008) and explicit in a paper by Williams (2006) is the suggestion that inflation expectations that are “well anchored” by the central bank’s explicit commitment to an inflation target may have altered the persistence and, more generally, the overall dynamics of inflation. Theoretically, this can, of course, be true to a degree. Consider once again the simple model of equations 2.8, modified slightly to make the central bank’s inflation target explicit:

$$\begin{aligned} \pi_t &= \pi_{t-1} + ax_t \\ x_t &= -b(f_t - \bar{\pi}) \\ f_t &= \bar{\pi} + c(\pi_t - \bar{\pi}). \end{aligned} \quad (4.19)$$

⁴⁸The time subscripts on the elements of \hat{A}_t are dropped for notational convenience.

⁴⁹Additional dynamics would, in principle, be added as the agents’ estimates of the \hat{a}_{ij} evolve over time. The specifics will vary depending on the estimation rule assumed, and a detailed investigation of this issue lies outside the scope of this paper.

As long as $c \neq 0$, the steady state for this model is given by⁵⁰

$$\begin{aligned}\pi_t &= \bar{\pi} \\ x_t &= 0 \\ f_t &= \bar{\pi}.\end{aligned}\tag{4.20}$$

But when $c = 0$, the central bank does nothing to move inflation toward a particular target, the steady-state for inflation becomes indeterminate, and the solution for inflation becomes

$$\pi_t = \pi_{t-1}.\tag{4.21}$$


In contrast to the stable autoregressive process in equation 2.9, the inflation rate in this case follows a random walk. This simple model demonstrates the importance in this class of models of keeping central bank actions consistent with pursuing an inflation target. When central banks do so, inflation is stationary. When they do not, it is not. In most models with explicit expectations, a similar proposition holds, as is well known. The intuition for a model with rational expectations is relatively straightforward. Consider a simple model based on the Calvo (1983) specification of inflation

$$\begin{aligned}\pi_t &= \beta E_t \pi_{t+1} + \gamma x_t \\ x_t &= -a(r_t - E_t \pi_{t+1}) \\ r_t &= \bar{r} + b(\pi_t - \bar{\pi}).\end{aligned}\tag{4.22}$$

As demonstrated in equation 3.2 above, the solution for π_t is the weighted sum of future x_t , which in turn depends on the future short-term real interest rates, $r_t - E_t \pi_{t+1}$. The central bank sets the short-term nominal interest rate r_t . Well-anchored inflation expectations require two things from the central bank. First, the central bank must have an inflation goal that is known to the private agents in the economy, and second, the central bank must move its policy rate in a way that systematically pushes the inflation rate toward that goal. Put differently, the ‘‘Taylor Principle’’ (Taylor 1999) operates in this model: As long as the central bank moves the policy rate by more than the inflation rate, so that it is increasing the short-term real rate when inflation is above its target and conversely when it is below, then the expected path of x_t will be consistent with returning inflation to its target under arbitrary initial conditions. In this sense, inflation expectations will be well anchored, and inflation will have a

⁵⁰The equilibrium real rate of interest is set to 0 for convenience.

Figure 11: Errors for Williams (2006) and random walk inflation models in recent years



determinate solution and be stationary. As is discussed in the introduction and in more detail below, in a purely forward-looking model such as this one, well-anchored expectations can imply not only determinacy and stationarity, but an inflation rate that follows a white noise process.⁵¹

Williams (2006) suggests that in recent years expectations may have become so well anchored that inflation may be well characterized by random deviations around a constant

$$\pi_t = c + \varepsilon_t .$$

As a description of the data, this simple model does reasonably well in recent years, as figure 11 suggests. The figure shows the errors made by a model that forecasts the four-quarter log change in prices as equal to its sample mean, versus one that sets the forecast equal to the previous four-quarter change.⁵² The random walk model has been advocated by Atkeson and Ohanian (2001) as

⁵¹This implication also relies on the assumption that all the shocks disturbing the equations above are *iid*.

⁵²Of course, the sample mean forecast uses information not available in real time to the forecaster.

an alternative to poorly performing Phillips curves. The errors over Williams’s sample are smaller for the constant-based forecast, with a root-mean-squared error of 0.29 versus 0.37 for the random walk model.⁵³

The diversity of results on reduced-form inflation persistence presented in the preceding subsections together should suggest some caution in arriving at conclusions about a possible change in inflation persistence in the past decade or two. We examine a simple Monte Carlo exercise to highlight the difficulty in inferring changes in inflation dynamics and their implications for persistence in relatively short samples.

Inflation in the exercise is generated by three simple equations: A backward-looking Phillips curve with lagged coefficients that sum to 1, an output equation that makes output a function of its own lag and the real interest rate, and a simple policy rule in which the policy rate responds only to the gap between inflation and the inflation target.

$$\begin{aligned}
 \pi_t &= \sum_{i=1}^4 a_i \pi_{t-i} + b \tilde{y}_t + \varepsilon_t \\
 \tilde{y}_t &= c \tilde{y}_{t-1} - d(r_{t-1} - \pi_t) + u_t \\
 r_t &= e r_{t-1} + (1 - e)(f(\pi_t - \bar{\pi})).
 \end{aligned}
 \tag{4.23}$$


The variances of the two shocks are calibrated from vector autoregressive equations estimated from 1966 to 2006, or split at 1984 to allow for changes in the estimated variance due to the so-called Great Moderation.⁵⁴ The baseline values for b and d are 0.1, for c 0.8, for e 0.8, and for f 1.5. The inflation target is set to 2.0. Fifty thousand samples of length 40 quarters are created using draws of the shocks to the inflation and output equations, assuming the individual shocks are random normal draws scaled by the estimated variances of ε and u . Initial conditions are randomized for each draw, rather than using a “burn-in” period.⁵⁵

⁵³Adding the last three years to the sample diminishes the advantage of the constant-based model. The RMSEs for the constant-based and random-walk models for the extended sample are 0.32 and 0.35, respectively. The constant-based model, with an estimated mean of about 1.8, consistently under-forecasts inflation over the past three years.

⁵⁴The equations estimate a bivariate VAR in Hodrick-Prescott-filtered real GDP and the annualized rate of inflation in the core consumer price index. The three samples are considered at 1966:Q1 to 2006:Q4, 1966:Q1 to 1984:Q4, and 1985:Q1 to 2006:Q4.

⁵⁵Tests for convergence of the distribution of estimated parameters suggest that 50,000 samples are more than adequate to assure convergence for this exercise.

Figure 12: Distribution of estimated coefficients



For each simulated sample, the following simple regression models are estimated. The first is designed as a simple test that inflation exhibits no persistence, perhaps because inflation expectations are well anchored at the central bank’s target, and can thus be modeled well as white noise around a constant c . The null for this model is that $\alpha = 0$, which would replicate Williams’s (2006) result. The second allows for the influence of output. Both are clearly misspecified, but the issue is whether one could be misled with a relatively short sample into believing that the properties of inflation had changed when in fact the underlying inflation process still exhibits persistence and a correlation with output.

$$\pi_t = \alpha\pi_{t-1} + c$$

$$\pi_t = \alpha\pi_{t-1} + by_{t-1} + c. \tag{4.24}$$

$$\tag{4.25}$$

The resulting regression estimates are displayed in the summary table below and the histograms in the panels of figure 12. For the first model, figure 12 and table 12 show that the standard error on the estimated lag coefficient is large; the median t -statistic for the hypothesis that $\alpha = 0$ has a p -value above

Table 12: Distribution of estimated coefficients

Model 1		
Coefficient	Median	Std. deviation
α	0.47	0.24
c	1.06	0.16
Model 2		
Coefficient	Median	Std. deviation
α	0.41	0.25
b	0.069	0.33
c	1.2	0.17

0.05. The estimate of c implies a mean for π_t of almost exactly 2.⁵⁶ One could be forgiven for inferring from these estimates that inflation had little or no persistence, and was well anchored around the central bank’s inflation target.⁵⁷

The results of the second regression model, displayed in the bottom panel of table 12, suggest that matters are even worse when one tests for the presence of a Phillips-like relationship in the misspecified model for a small sample. Now the median estimate of α is smaller still, it is imprecisely estimated (the median p -value for the test of $\alpha = 0$ is greater than 0.10), and the effect of output is biased downward and very imprecisely estimated. Once again, the mean of the estimated inflation target matches the true value quite precisely. One might conclude from simple regressions such as these that inflation has little or no persistence and that the Phillips correlation is absent. These small samples allow one to recover the average value of inflation over the period, but the rest of inflation dynamics may be very poorly estimated.

5 Microeconomic evidence on persistence

5.1 Persistence in micro data: U.S. evidence

Up to this point, the paper has focused on macroeconomic, aggregate evidence bearing on inflation persistence. Yet the dominant models in the literature aim

⁵⁶The standard error of the intercept, at 0.16, is quite small. The distribution of the mean of inflation, which is a nonlinear function of α and c ($\bar{\pi} = \frac{\alpha}{(1-c)}$), implies a relatively large standard error for the mean of 2.3.

⁵⁷A good portion of the downward bias evident in the estimates in figure 12 arises from the truncation of the lags of inflation compared with the lags in the data-generating process. Still, in large samples, the regression should retrieve the sum of the underlying AR coefficients in the single autoregressive coefficient. As the sample size in the exercise above increases from 10 years to 25 to 50, the median estimate of α in Model 1 increases from 0.47 to 0.70 to 0.84.

to provide microeconomic foundations for inflation, based on the price-setting decisions of individual firms. In this regard, it is striking that the now large literature that examines micro-price data has emerged relatively recently, led by the work of Bils and Klenow (2004). Bils and Klenow employ an unpublished dataset of about 350 CPI expenditure categories that account for about 70 percent of consumer expenditures. In examining the persistence and volatility of inflation, they use a subset of 123 spending categories that account for about 63 percent of spending.

For each of the CPI components that they examine p_i , they estimate a simple AR(1) process in order to assess the degree of persistence and volatility for its inflation rate dp_i (defined as the log change)

$$dp_{i,t} = \rho dp_{i,t-1} + \varepsilon_{i,t} .$$

They find that the average persistence of inflation, which they define as the arithmetic average of the ρ_i 's, is slightly negative (-0.05) for their shorter sample (1995–2000) and only modestly positive (0.26) for their longer sample (1959–2000). Interestingly, for the shorter sample, they find that the degree of persistence across price categories is *positively* correlated with the frequency of price change, clearly contrary to predictions from the Calvo/Rotemberg and Taylor models. Over the longer sample period, the correlation is positive but not statistically different from zero.

The authors estimate that adjusting for the presence of “temporary sales”—explicitly transitory reductions in prices that revert within days or weeks and would likely lead to an understatement of the persistence of non-sale prices—would have only a small impact on their estimates of persistence. Of course, the fact that temporary sales occur regularly casts some doubt on the underlying assumption of time-dependent pricing models, in which prices are simply fixed for long periods. It also raises the question as to whether the relevant price measures should exclude or include temporary sales. This issue is discussed in more detail in Nakamura and Steinsson (2008), who employ a dataset that allows them to study prices at a more disaggregated level. They find that temporary sales have a larger effect on the measured frequency of price changes than in the Bils and Klenow data; excluding temporary sales reduces the frequency of price changes by about one-half. They do not estimate the effect of temporary

sales on inflation persistence in their paper.

The results from these seminal papers suggest that the micro data exhibit behavior that is at odds with the prevailing time-dependent pricing models' description of underlying price behavior. Nakamura and Steinsson's paper suggests that several key features are also at odds with a menu-cost model. But whether the micro data are consistent with the estimates of persistence in aggregate data is less clear, for two reasons. First, aggregate price series may exhibit quite different properties from the individual price series. Second, what one observes in individual price changes likely reflects the combined influence of firm- or industry-specific shocks and macro shocks. The two points are related, and the relative importance of the two is an empirical question, but if individual prices respond differently to macro versus micro shocks, then it will be important to sort out these influences in evaluating the relevance of micro-data evidence for aggregate inflation persistence.

Section 5.3 discusses some recent work bearing on the first point. Boivin, Giannoni, and Mihov (2009) provide results bearing on the second point. They estimate a small number of common factors (principal components) from a large number of macroeconomic variables. They then relate individual price changes to these common factors in order to decompose individual inflation rates into idiosyncratic, sector-specific fluctuations and macroeconomic fluctuations. Denoting the matrix of common factors by C_t and the log change in the individual price series p_{it} as dp_{it} , they estimate regressions

$$dp_{it} = \lambda_i C_t + \epsilon_{it} .$$

The R^2 of each regression indicates the fraction of individual price changes that may be attributed to the common macroeconomic factors; one minus this R^2 is the fraction of variation attributable to sector-specific sources. Their baseline results suggest that about 85 percent of the variation in the individual price changes may be attributed to sector-specific shocks.

Using the decomposed individual inflation series, Boivin *et al.* estimate simple autoregressions for each of the series and their two components, measuring persistence by the sum of the autoregressive coefficients.⁵⁸ They find, like Bills and Klenow (2004), that individual inflation series exhibit relatively

⁵⁸Their data are observed at the monthly frequency, and they estimate autoregressions with 13 lags.

little persistence. The idiosyncratic components of the individual inflation series ϵ_{it} exhibit essentially no persistence, while the common components ($\lambda_i C_t$) vary in persistence from negative for some series to above 0.95 for some health-care components and tenant room and board. These findings imply that the aggregate inflation measures, which are quite persistent in their data, inherit their persistence from the common macroeconomic components of the individual price series, particularly from those that exhibit the highest persistence. The non-persistent idiosyncratic components essentially wash out in aggregation.

5.2 Persistence in micro data: Euro-area evidence

Altissimo, Ehrmann, and Smets (2006) examine both aggregate and disaggregated data to explore the properties of Euro-area inflation. Their conclusions on aggregate data echo those of others—inflation has been persistent; its reduced-form persistence has declined in recent years so that it now exhibits moderate persistence, although how moderate depends on how it is estimated; its decline may be attributable to a stable and well-focused monetary regime that anchors long-run inflation expectations. Their disaggregated sectoral data suggest that individual price series exhibit less persistence on average than their corresponding aggregates.

The studies of Angeloni, Aucremanne, Ehrmann, Galí, Levin, and Smets (2006) and Alvarez, Dhyne, *et al.* (2006) find ample evidence of infrequent price changes at the micro level. The former estimates that while there is substantial heterogeneity across sectors, prices on average are quite sticky, exhibiting a four to six quarter duration. Seen through the lens of the Calvo model, these estimates suggest that the “Calvo parameter,” which indexes the frequency of price changes, is large, implying a small effect of marginal cost on inflation. From equation 3.5 above, the smaller is the effect of marginal cost on inflation, the less of the persistence in marginal cost is inherited by inflation. Neither of these studies examines the persistence of disaggregated price series, nor do the studies explore the complications in aggregating disaggregated series with heterogeneous dynamics. But to the extent that they find relatively infrequent price changes, and to the extent that one feels comfortable mapping these into aggregate Calvo parameters, the studies imply less inherited persistence, other things equal.

5.3 More on aggregation and persistence

Bils and Klenow find some difference between the persistence properties of inflation based on individual price series and their expenditure-share-weighted aggregate inflation measure. For their longer sample, they estimate an autoregressive coefficient of 0.63, with a standard error of 0.03, in contrast to the much smaller average of the autoregressive coefficients for the individual price series. Boivin *et al.* find that the low-persistence, idiosyncratic components of disaggregated inflation series appear to wash out in aggregate inflation measures, leaving the persistent common macroeconomic components to dominate the persistence of aggregate inflation. These observations suggest that aggregation of price series may play an important role in determining the degree of aggregate inflation persistence.⁵⁹

Several recent papers examine the role of aggregation in inflation persistence. Mumtaz, Zabczyk, and Ellis (2009) employ a methodology that draws importantly on Boivin *et al.*, studying disaggregated price data for the United Kingdom. They also find that the persistence of the aggregate inflation measure is biased upwards relative to the persistence of the underlying price series and that the bias is driven by the macro components of those series. Altissimo, Mojon, and Zaffaroni (2009) delve more deeply into the aggregation process, using existing results on aggregation of time series (see, for example, Granger (1980) for parametric results and Zaffaroni (2004) for non-parametric results that are more closely related to the work in Altissimo *et al.*). Similarly to Boivin *et al.*, they assume that individual price series may be characterized by an unobserved components model that makes each price change series $dp_{i,t}$ a function of its own idiosyncratic persistence, a common shock u_t , and an idiosyncratic shock $\varepsilon_{i,t}$:

$$dp_{i,t} = \alpha_i dp_{i,t-1} + u_t + \varepsilon_{i,t} .$$

By assuming a particular form for the distribution of the persistence parameters $f(\alpha)$, as in Granger (1980), one can derive results for the persistence of the simple aggregate of the individual price changes $dP_t = (1/n) \sum_{i=1}^n dp_{i,t}$. The relative contributions of the common and idiosyncratic shocks will be important

⁵⁹On a note of theoretical counterpoint, Carvalho (2006) develops a multisector version of the Calvo model with heterogeneity in price stickiness that implies an aggregate inflation rate with *less* persistence than that of the standard Calvo model.

in understanding the relationship between individual and aggregate prices, but so too will be the differences in the way in which the common shock is propagated in the individual prices—that is, the differences in the α_i 's. Price series that perpetuate the common shock through a larger α_i will have a larger effect on the persistence of the aggregate than series with a smaller α_i .

They find that disaggregated price changes exhibit significantly less persistence than the aggregate and that the preponderance of the variance of the individual price series is accounted for by idiosyncratic volatility, in agreement with all of the studies cited above. Like Boivin *et al.*, they find that a single principal component of the individual price series dominates in explaining the low-frequency variation in the individual price changes. It follows that this factor must account for the common persistence among the disaggregated series. Estimation of a model like 5.3 reveals in addition that the propagation of the common shock in individual series is indeed quite varied. The high persistence of the common shock in services prices, in particular, combined with the relatively high weight of services in the aggregate price index, accounts for a substantial part of the persistence in aggregate inflation.

6 Conclusions

It may be early to draw firm conclusions about the structural sources of inflation persistence, or about the extent to which these sources have changed and manifested themselves in changes in reduced-form inflation persistence. In the first case, it may be premature because there is not yet widespread agreement about the appropriate mapping between micro data or reduced-form aggregate data and our structural models. In the second case, we have a fairly short sample from which to draw inferences about potential changes (see section 4 for more details).

Still, the research to date allows one to draw some conclusions. First, to the extent that reduced-form persistence has changed, policymakers need to gain clarity about the source of the change. This paper discusses a number of structural channels through which persistence may have changed. There may have been a change in the “intrinsic” persistence in inflation—the importance of lagged inflation in the structural Phillips curves. Alternatively, the amount of inherited persistence may have changed. In principle, this could arise because

the persistence of the driving process has changed, or because the coefficient on the driving process has changed, or because the relative variances of the shocks to inflation and the driving process have changed.

The analysis in this paper suggests that it is unlikely that any change in persistence has arisen from a change in the persistence of the driving process, as this has remained remarkably stable throughout the period. In addition, a DSGE model-based analysis suggests that while changes in the systematic component of monetary policy likely have led to less-persistent inflation, the largest changes in persistence are most likely due to changes in the so-called intrinsic sources of inflation persistence—whether those arise from indexation, rule-of-thumb price-setters, or a rising price reset hazard. Finally, the models that depart from the standard Calvo framework suggest that other aspects of the economy that impinge upon inflation persistence may be responsible for changes in inflation persistence. These may include smaller or less-frequent changes in “trend inflation” or a smaller role for learning, as central bank transparency about its goals has increased.

Second, we have now accumulated an impressive and growing body of evidence on the behavior of price- (and wage-)setting at the disaggregated level. This evidence strongly suggests that some of the inferences drawn from micro data about the frequency of price changes, as well as the degree of inflation persistence, may pertain largely to price responses to industry- or firm-specific shocks. The response to aggregate shocks by the aggregate component common to the individual price series may well have quite different properties from the responses of individual firms to idiosyncratic shocks. Integrating this evidence into our structural models, perhaps along the lines of models of “rational inattention” (see Sims (2003), Gorodnichenko (2008) and Maćkowiak and Wiederholt (2009)) seems a promising avenue for research.

Finally, we are currently accumulating additional evidence that should allow us to take a firmer stance on whether reduced-form persistence has changed, and to discern the structural sources of any such changes. The upheaval created by the 2007–2009 financial crisis and recession, with the concomitant prospect of a prolonged period of elevated unemployment, suggest that over the next decade we will have accumulated evidence that will allow us to test more fully the hypothesis of a decline in reduced-form inflation persistence and to test competing

theories that attribute the structural sources of persistence.

References

- Altissimo, Filippo, Benoit Mojon, and Paolo Zaffaroni (2009). “Can aggregation explain the persistence of inflation?” *Journal of Monetary Economics* **56**: 231–241.
- Altissimo, Filippo, Michael Ehrmann, and Frank Smets (2006). “Inflation persistence and price-setting behaviour in the Euro area: A summary of the IPN evidence.” National Bank of Belgium working papers No. 95, October.
- Alvarez, Luis, E. Dhyne, M. Hoeberichts, C. Kwapil, H. Le Bihan, P. Lunnemann, F. Martins, R. Sabbatini, H. Stahl, P. Vermeulen, and J. Vilmu, (2006). “Sticky prices in the euro area: a summary of new micro-evidence.” *Journal of the European Economic Association* **4**: 575–584.
- Angeloni, Ignazio, Luc Aucremanne, Michael Ehrmann, Jordi Galí, Andrew Levin, and Frank Smets (2006). “New evidence on inflation persistence and price stickiness in the Euro area: Implications for macro modeling.” *Journal of the European Economic Association*, **4**: 562–574.
- Atkeson, Andrew and Lee Ohanian (2001). “Are Phillips curves useful for forecasting inflation?” Federal Reserve Bank of Minneapolis *Quarterly Review* **25**, No. 1: 2–11, Winter.
- Bakhshi, Hasan, Hashmat Khan, and Barbara Rudolf (2007). “The Phillips curve under state-dependent pricing.” *Journal of Monetary Economics* **54**: 2321–2345.
- Ball, Laurence (1994). “Credible disinflation with staggered price-setting.” *American Economic Review* **84**(1): 282–289, March.
- Ball, Laurence, and Stephen G. Cecchetti (1990). “Inflation and uncertainty at short and long horizons.” *Brookings Papers on Economic Activity*: 215–54.
- Barnes, Michelle, Fabià Gumbau-Brisa, Denny Lie, and Giovanni Olivei (2009). “Closed-form estimates of the New Keynesian Phillips curve with time-varying trend inflation.” Working paper, Federal Reserve Bank of Boston.
- Barsky, Robert B. (1987). “The Fisher hypothesis and the forecastability and persistence of inflation.” *Journal of Monetary Economics* **19**: 3–24.
- Benati, Luca (2008). “Investigating inflation persistence across monetary regimes.” *Quarterly Journal of Economics* **123**(3): 1005–1060, August.
- Benati, Luca (2009). “Are intrinsic inflation persistence models structural in the sense of Lucas (1976)?” European Central Bank Working paper series No. 1038, March.
- Bils, M. and P. Klenow (2004). “Some evidence on the importance of sticky prices.” *Journal of Political Economy* **112**(5).

- Blanchard, Olivier and Jordi Galí (2007). “Real wage rigidities and the new Keynesian model.” *Journal of Money, Credit and Banking*, Supplement: Conference on “Quantitative Evidence on Price Determination,” September 29 and 30, 2005, co-sponsored by the JMCB and the Federal Reserve Board, **39**(1): 35–65.
- Boivin, Jean, Marc P. Giannoni and Ilian Mihov (2009). “Sticky prices and monetary policy: Evidence from disaggregated U.S. data.” *American Economic Review* **99**(1): 350–384.
- Bray, Margaret (1982). “Learning, estimation, and the stability of rational expectations.” *Journal of Economic Theory* **26**: 318–39, April.
- Burstein, Ariel T. (2006). “Inflation and output dynamics with state-dependent pricing decisions.” *Journal of Monetary Economics* **53**: 1235–1257.
- Calvo, Guillermo (1983). “Staggered prices in a utility-maximizing framework.” *Journal of Monetary Economics* **12**(3).
- Calvo, G.A., Celasun, O., and M. Kumhof (2002). “A theory of rational inflationary inertia.” in: P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford, eds., *Knowledge, Information and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*. Princeton: Princeton University Press.
- Caplin, Andrew and John Leahy (1991). “State-dependent pricing and the dynamics of money and output.” *Quarterly Journal of Economics* **106**(3): 683–708.
- Carvalho, Carlos (2006). “Heterogeneity in price stickiness and the real effects of monetary shocks.” *Frontiers of Macroeconomics* Volume 2, Issue 1, Article 1
- Christiano, Lawrence, Martin Eichenbaum, and Charles Evans (2005). “Nominal rigidities and the dynamic effects of a shock to monetary policy.” *Journal of Political Economy*.
- Cogley, T. and A. Sbordone (2009). “Trend inflation, indexation, and inflation persistence in New Keynesian Philips curve.” *American Economic Review*.
- Cogley, Timothy, Giorgio E. Primiceri, and Thomas J. Sargent (2007). “Inflation-gap persistence in the U.S.” Working paper, December.
- Dornbusch, Rudiger (1976). “Expectations and exchange rate dynamics.” *Journal of Political Economy* **84**(6): 1161–1176, December.
- Dotsey, Michael, Robert King, and Alexander Wolman (1999). “State-dependent pricing and the general equilibrium dynamics of money and output.” *Quarterly Journal of Economics*, **114**: 655–690, May.
- Estrella, Arturo and Jeffrey Fuhrer (2002). “Dynamic inconsistencies: Counterfactual implications of a class of rational expectations models.” *American Economic Review* **92**(4): 1013–1028.
- Fischer, Stanley (1977). “Long-term contracts, rational expectations, and the optimal money supply rule.” *Journal of Political Economy* **85**(1): 191–205.
- Friedman, Milton (1968). “The role of monetary policy.” *American Economic Review* **58**(1): 1–17.

- Fuhrer, Jeffrey (2000). "Habit formation in consumption and its implications for monetary-policy models." *American Economic Review* **90**(3): 367–390.
- Fuhrer, Jeffrey (2006). "Intrinsic and inherited inflation persistence." *International Journal of Central Banking* **2**(3): 49–86.
- Fuhrer, Jeffrey (2008). "Special issue comment on optimal price setting and inflation inertia in a rational expectations model." *Journal of Economic Dynamics and Control* **32**: 2536–2542.
- Fuhrer, Jeffrey and George Moore (1992). "Monetary policy rules and the indicator properties of asset prices." *Journal of Monetary Economics* **29**: 303–336.
- Fuhrer, Jeffrey and George Moore (1995). "Inflation persistence." *Quarterly Journal of Economics* **110**(1): 127–159, February.
- Galí, Jordi and Mark Gertler (1999). "Inflation dynamics: A structural econometric analysis." *Journal of Monetary Economics* **44**: 195–222.
- Gordon, Robert J. (1982). "Price inertia and policy ineffectiveness in the United States, 1890–1980." *Journal of Political Economy* **90**(6): 1087–1117.
- Gordon, Robert J., Stephen R. King, and Franco Modigliani (1982). "The output cost of disinflation in traditional and vector autoregressive models." *Brookings Papers on Economic Activity* **1982**(1): 205–244.
- Gorodnichenko, Yuri (2008). "Endogenous information, menu costs and inflation persistence." NBER Working Paper 14183.
- Granger, Clive (1980). "Long memory relationships and the aggregation of dynamic models." *Journal of Econometrics* **14**: 227–238.
- Gray, Jo Anna (1977). "Wage indexation: A macroeconomic approach." *Journal of Monetary Economics* **2**(2): 221–235.
- Koenig, Evan (1996). "Aggregate price adjustment: The Fischerian alternative." Federal Reserve Bank of Dallas Working Paper No. 9615.
- Levin, Andrew and Jeremy Piger (2004). "Is inflation persistence intrinsic in industrial economies?" European Central Bank working paper No. 334, April.
- Levin, A. and J. Williams (2003). "Robust monetary policy with competing reference models." *Journal of Monetary Economics* **50**(5).
- Lucas, Robert E. Jr. (1972). "Expectations and the neutrality of money." *Journal of Economic Theory* **4**: 103–124.
- Mackowiak, Bartosz and Mirko Wiederholt (2009). "Optimal sticky prices under rational inattention." *American Economic Review* **99**(3): 769–803.
- Mankiw, N.G. and R. Reis (2002). "Sticky information versus sticky prices: A proposal to replace the New Keynesian Phillips curve." *Quarterly Journal of Economics* **117**(4).
- Marcet, Albert and Thomas J. Sargent (1989). "Convergence of least-squares learning in environments with hidden state variables and private information." *Journal of Political Economy* **97**(6): 1306–1322, December.
- Mash, Richard, (2004). "Optimizing microfoundations for inflation persistence." Oxford University Dept. of Economics Discussion Paper No. 183, January.

- Mavroeidis, Sophocles (2005). "Identification issues in forward-looking models estimated by GMM, with an application to the Phillips curve." *Journal of Money, Credit, and Banking* **37**(3): 421–448, June.
- Mumtaz, Haroon, Pawel Zabczyk and Colin Ellis (2009). "What lies beneath: What can disaggregated data tell us about the behaviour of prices?" Bank of England Working Paper No. 364, March.
- Muth, John (1961). "Rational expectations and the theory of price movements." *Econometrica* **29**(3): 315–335.
- Nakamura, Emi and Jon Steinsson (2008). "Five facts about prices: A reevaluation of menu cost models." *Quarterly Journal of Economics* **112**: 1415–64, November.
- Okun, Arthur M. (1977). "Efficient disinflationary policies." *American Economic Review* **68** (May 1978, Papers and Proceedings (1977)): 348–52.
- O'Reilly, G. and K. Whelan (2005). "Has Euro-area inflation persistence changed over time?" *Review of Economics and Statistics* **87**(4): 709–720.
- Orphanides, A. and J.C. Williams (2004). "Imperfect knowledge, inflation expectations, and monetary policy." in B. Bernanke and M. Woodford, eds., *The Inflation Targeting Debate*.
- Phelps, Edmund S. (1968). "Money-wage dynamics and labor-market equilibrium." *Journal of Political Economy* **76**(4, Part 2): 678–711.
- Pivetta, Frederic and Ricardo Reis (2007). "The persistence of inflation in the United States." *Journal of Economic Dynamics and Control* **31**: 1326–1358.
- Ravenna, F. (2000). "The impact of inflation targeting in Canada: A structural analysis." manuscript, New York University.
- Roberts, J. (2006). "Monetary policy and inflation dynamics." *International Journal of Central Banking* **2**.
- Roberts, J. (1997). "Is inflation sticky?" *Journal of Monetary Economics* **39**(2).
- Rotemberg, Julio J. (1982). "Sticky prices in the United States." *Journal of Political Economy* **90**(6): 1187–1211.
- Rotemberg, Julio J. (1983). "Aggregate consequences of fixed costs of price adjustment." *American Economic Review* **73**(3): 433–436.
- Rudd, J. and K. Whelan (2006). "Can rational expectations sticky-price models explain inflation dynamics?" *American Economic Review* **96**.
- Sargent, Thomas J. and Neil Wallace (1975). "'Rational' expectations, the optimal monetary instrument, and the optimal money supply rule." *The Journal of Political Economy* **83**(2): 241–254.
- Sheedy, Kevin D. (2007). "Intrinsic inflation persistence." Centre for Economic Performance Discussion Paper No. 837, November.
- Sims, Christopher A. (2003). "Implications of rational inattention." *Journal of Monetary Economics* **50**: 665–690.
- Smets, F., and R. Wouters (2003). "Monetary policy in an estimated stochastic dynamic general equilibrium model of the Euro area." *Journal of European Economic Association* **1**(5).

- Smets, F. and R. Wouters (2003). “An estimated dynamic stochastic general equilibrium model of the Euro area.” *Journal of European Economic Association* **1**(5).
- Stock, J.H. and M. Watson (2007). “Has inflation become harder to forecast?” *Journal of Money, Credit, and Banking* **39**: 3–34.
- Taylor, John (1980). “Aggregate dynamics and staggered contracts.” *Journal of Political Economy* **88**.
- Taylor, John (1993). “Discretion versus policy rules in practice.” *Carnegie-Rochester Conference Series on Public Policy* **39**: 195–214.
- Taylor, John (1999). “A historical analysis of monetary policy rules.” in John B. Taylor, ed., *Monetary Policy Rules*. Chicago: University of Chicago Press: 319–41.
- Williams, John (2006). “The Phillips curve in an era of well-anchored inflation expectations.” Federal Reserve Bank of San Francisco, unpublished working paper, September.
- Woodford, Michael (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton: Princeton University Press.
- Woodford, Michael (2007). “Interpreting inflation persistence: Comments on the conference on Quantitative Evidence on Price Determination.” *Journal of Money, Credit and Banking*, Supplement to **39**(1): 203–210, February.
- Zaffaroni, Paolo (2004). “Contemporaneous aggregation of linear dynamic models in large economies.” *Journal of Econometrics* **120**: 75–102.