

# INFLATION EXPECTATIONS, UNCERTAINTY, THE PHILLIPS CURVE, AND MONETARY POLICY

CHRISTOPHER A. SIMS

## I. A VIEW OF THE HISTORY OF THE PHILLIPS CURVE

The original observation by Phillips simply noted an empirical regularity: unemployment and inflation tended to be inversely related. This observation came at a time when Keynesian macroeconomic theory had a very simple and incomplete model of inflation. Keynesian theory treated wages as, if not fixed, then on an exogenously given time path. It was a theory of how nominal aggregate spending determined the level of output and employment, so long as supply-side limits on output and employment were not encountered. It was recognized that when aggregate demand exceeded supply-side limits, the result would be inflation, but the standard Keynesian theory had a discontinuity at the point where output hit “capacity”, and it had no quantitative predictions about the determination of the level of inflation once capacity limits were hit.

---

*Date:* August 1, 2008.

This research was supported by NSF grant SES-0719055. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation (NSF).

As macroeconomists began to think about quantitative modeling of the aggregate economy, the Phillips curve offered a way to make Keynesian inflation theory continuous and quantitative. The level of unemployment could be used to measure how far the economy was from capacity, and thereby to make quantitative predictions about how inflation would be affected by the level of aggregate demand. Policy, whether monetary or fiscal, was conceived as affecting inflation via a causal chain, from aggregate demand, to the level of output and employment (and thereby unemployment), to the rate of inflation. Through the 1960's and 1970's probably most economists thought about inflation-determination this way, and many still do. I am not arguing here that many economists think such a two-equation recursive model of the economy is the full story of inflation determination, but simple one and two equation models are part of the mental furniture of most macroeconomists, and this particular simple model remains influential.

Primiceri (2006) models inflation-determination in the 1960's and 70's and 80's as reflecting policy-makers' use of a model like this and learning over time about the value of its coefficients. One may be skeptical of his results because of his assumption that the model about which the policy-makers are learning is correct, with only the coefficient values uncertain. Nonetheless, the fact that Primiceri's interpretation of history works as well as it does may explain why this way of thinking still has a hold on policy-makers' thinking.

This is interesting, because we know that Lucas and Rapping in a series of papers in the late 60's and early 70's (1973; 1969b; 1969a) developed a model with some plausibility in which Phillips's empirical regularity could be misleading if used, as the Keynesian models were doing, to analyze the effects of policy. This new simple model arrived on the scene just as the US entered a period in the 1970's of simultaneous high unemployment and high inflation, making the data in unemployment-inflation plots jump off the historical Phillips curve. The simple rational expectations version of this theory, in which the causal direction is reversed, with inflation surprises causing changes in unemployment, did not fit the data any better than the deteriorating standard Phillips curve, but it provided a qualitative story about why a Phillips curve might first appear in the data, then disappear in the presence of Keynesian policy-making.

While a few of the early advocates of rational expectations modeling (Sargent reference) held out the hope that it would generate "cross-equation restrictions" that would lead to improved quantitative policy models, the new theory was more commonly interpreted as implying the entire enterprise of large-scale policy modeling was quixotic. Simultaneous equation econometrics began to disappear from economics PhD training in the US, while every new PhD could explain how the "Lucas critique" implied that Keynesian macro models would lead to policy errors. With the simple "Lucas supply curve" (the rational expectations, reversed-direction, Phillips curve) replacing the Phillips curve, there was furthermore no need for big policy

models. The best monetary policy could do was to avoid creating surprises. Milton Friedman's proposal of a fixed growth rate for the money stock (which he supported with a different set of arguments) fit well with the rational expectations policy analysis.

Meanwhile, those actually making monetary policy faced a continuing need to make decisions responsibly in the light of data emerging week by week. The Thatcher government's experiment in the UK with a simple monetary growth rate policy rule showed that the historical statistical relationships among various measures of the money stock, and between the money stock and inflation and output, could deteriorate when exploited for policy purposes in the same way, and for the same reasons, that the empirical Phillips curve had decayed. With academic economic research turned almost entirely away from large scale policy modeling, central bank economists developed their own solutions. They emerged with models that preserved many of the characteristics of the first generation of Keynesian models: equation-by-equation specification; emphasis on flow equilibrium; and Phillips curves as the locus for non-neutrality of monetary policy. Expectations now entered the models more pervasively, and the models, to sidestep the Lucas critique, made it at least formally possible to treat expectations as rational. The discipline of simultaneous equations econometric inference was entirely abandoned.<sup>1</sup>

---

<sup>1</sup>I discussed the state of central bank modeling in a 2002 Brookings paper.

For policy modeling, the simple Lucas supply curve was inadequate. Besides not fitting the data, its microeconomic underpinnings were either informal or, in formal models, highly abstract and unrealistic — for example models of “island economies” in which people had to infer the value of the economy-wide interest rate or money stock from the price level on their own island. The policy models began by simply adding an inflation expectations term to the right-hand-side of the original Phillips curve, but there was no satisfactory theory of how such a relationship arose out of individual economic behavior. Into this gap sprang the New Keynesian Phillips Curve.

## II. THE NEW KEYNESIAN PHILLIPS CURVE: IS IT A PHILLIPS CURVE? IS IT USEFUL?

The New Keynesian (NK) Phillips curve is not an empirical relation between unemployment and inflation. It nonetheless can play the same role as the Phillips curve in a policy model: it links a continuously varying, observable measure of “distance from capacity” to predictions about the rate of inflation. Furthermore, it provides a microeconomic story about how this relation emerges, a story in which people have rational expectations and have no money illusion. There are a number of reasons, though, to see the NK Phillips curve as a Pyrrhic victory.

The theory of the NK Phillips curve is well known and documented elsewhere, e.g. in Woodford (2003), so I will just summarize it here. A continuum of monopolistically competitive firms have control over their own prices, because of product differentiation, but have an incentive to keep their prices in line with those of other

firms, because there are competitive pressures. They face some friction in price-setting, however. There are a number of postulated forms of friction. One is that prices are set in contracts of fixed length, an idea first explored by John Taylor. Another, more convenient form is that prices are fixed for random periods, with the duration of the random period determined exogenously. (This latter is “Calvo pricing”.) There are further variations on the form of the friction, some of which we will discuss below. Because of the friction, when the aggregate price level moves, not all firms respond to the change at once, and this creates non-neutrality for monetary policy.

This theory sidesteps the Lucas critique, because it contains expectations explicitly and assumes that expectations are rational. But the Lucas critique is only one special case of a generic problem we face in econometric modeling: we make simplifications and approximations that we realize are contingent, so that some kinds of changes in policy, or in the nature of exogenous disturbances, will force us to change the model. The NK Phillips curve is clearly unstable under some kinds of policy change — indeed under exactly the same kinds of policy change that the Lucas critique claimed could undermine old Keynesian models. Though the agents in the NK model have rational expectations and no money illusion, the theory has simply moved the non-neutrality from agent behavior itself into the constraints the agent faces, the frictions. The contract lengths of Taylor and Calvo theory are clearly not constants of nature;

surely they will change systematically with the level, variability and forecastability of inflation.

But there is a perhaps more important problem with the NK theory: it props up the simple Phillips curve way of thinking about the link from monetary policy to inflation. Though it suggests a different way of measuring real tightness — the “output gap” in place of unemployment — it still provides an equation in which real tightness appears as the crucial determinant of inflation. Of course in principle once inflation expectations are admitted to a Phillips curve equation, new style or old, it becomes possible for disturbances anywhere in the model to impact inflation directly, without any intermediating move in the measure of real tightness. If such influences are small, or slow-moving, it may nonetheless be helpful to think of inflation as determined, via a Phillips curve, by real tightness. But it is also possible that the opposite is true — the impact of policy and other disturbances on inflation is mainly direct, through the expectation term in the Phillips curve, so that retaining the Phillips curve as the central focus of informal thinking about inflation determination is misleading. Orphanides (2001) has explained how the US inflation in the 70’s could have emerged from policy-makers’ difficulties in real-time measurement of the output gap. But these difficulties played such a central role in good part because of Phillips curve thinking — the notion that some measure based on real data, with no statistical input from inflation itself or inflation expectations, was the central determinant of inflationary or disinflationary pressure.

The NK theory gives a central role not to unemployment, but to the output gap. Recently the empirical literature (Sbordone, 2003), e.g., has recognized that the output gap is actually important in the theory because it measures marginal cost, and has moved toward more direct measures of this, in particular to looking at the labor share of output.

It is reasonable then, to ask whether we have any evidence on this issue: to what extent is some version of a Phillips curve central to the determination of inflation? We can look for an answer to monetary structural vector autoregressions, multivariate statistical models that distinguish monetary policy behavior, and disturbances to it, from other sources of variation in the economy, without imposing detailed interpretations of the estimated dynamics in terms of individual behavior. Such a model must allow for the very different monetary policy behavior during the 1979-82 period of Volcker's reserve targeting and for the substantial decline and shifts in relative sizes of volatility of disturbances between the pre-Volcker and post-1982 periods. The results we discuss here are from models fit separately to quarterly data from 1959:1 through 1979:2 and from 1983:1 through 2008:1.<sup>2</sup> The model was identified by assuming the pattern of zero restrictions shown in Table 1 for the matrix of contemporaneous coefficients in a model of the form

$$A(L)y_t = \alpha + \varepsilon_t, \quad (1)$$

---

<sup>2</sup>The data were all drawn from the FRED website of the Federal Reserve Bank of St. Louis, during the last week of May and the first week of June, 2008.

	FF	y	oph	w	p	m	pcrm
policy	X	0	0	0	0	X	X
pcrm	X	X	X	X	X	X	X
y	0	X	X	X	X	X	0
oph	0	0	X	X	X	X	0
w	0	0	0	X	X	X	0
p	0	0	0	0	X	X	0
m	0	0	0	0	0	X	0

TABLE 1. Identifying restrictions

where the  $y$  vector consists of the Federal Funds Rate, output, output per hour, employee compensation per hour, price, M1 money stock, and producer prices for crude materials, as labeled in that order across the top of Table 1. Output, output per hour, employee compensation, and price all refer to the non-farm business sector.<sup>3</sup>

The first equation, labeled “policy” in the table, represents monetary policy behavior. The zero restrictions in that row reflect the fact that gdp data are not available

---

<sup>3</sup>The model was estimated using a dummy observation “Minnesota prior”, shrinking toward an independent random walks prior mean. The R programs used to estimate the model are available via the subversion internet protocol at <svn://sims.princeton.edu/R>.

to policy-makers within the quarter.<sup>4</sup> The zero restrictions in the first and last column reflect an assumption that private-sector variables not set in auction markets (i.e. those other than *pcrm*) respond only with a delay to interest rates or to rapidly fluctuating commodity prices. The triangular pattern of zeros in the lower part of the central five columns are simply normalizations. The last five equations are interpreted as a block that determines the central five variables, with the individual equations having no distinct interpretations. Though identifying assumptions like these are often characterized as controversial, behavioral models with complete interpretations often embed the same or similar timing assumptions in much more restrictive frameworks, in which these sorts of timing assumptions pass by without comment.

This kind of model can let us examine two questions: Does the NK Phillips curve mechanism seem to be playing a central role in transmitting the effects of shifts in monetary policy? and, Does new information about the labor share variable and its predicted future values play an important role in changing inflation forecasts?

To answer the first question, Figures 1 and 2 show the estimated responses of all variables in the system to a monetary policy disturbance. All variables are measured in log units, except the interest rate, which is measured at an annual rate as a proportion (not percent). The last entry at the lower right at each figure, labeled “*lsh*”,

---

<sup>4</sup>Employment data are available monthly, and the data from before 1979:3 provide some indication that in that period a positive within-period response of the funds rate to employment changes should not have been ruled out.

**Pre-79:3 Responses to M Policy**




FIGURE 1.

is the response of the labor share variable, which is constructed as  $w - oph - p$ . The

**Post-83 Responses to M Policy**




FIGURE 2.

pattern of responses is largely similar for the two subperiods, and also broadly similar to estimated responses to monetary policy disturbances in models estimated with

other identifying assumptions. In particular, the funds rate rises, then returns to or falls below its original level; wages and prices both fall, though prices, in the earlier period especially, fall with more of a delay than for wages; output and output per hour fall. Note that the falls in prices and wages are persistent and predictable as soon as the monetary policy disturbance occurs and the changes in output are almost as persistent. This does not fit well with stories that surprise changes in the inflation rate are what generates real effects of monetary policy.

The most prominent difference between the two periods is that in the later period a monetary policy shock forecasts a hump-backed time path of further increases in the funds rate, followed by a later decline. In the earlier period the model estimates less of this interest-rate-smoothing behavior.

In both periods, the labor share variable moves very little in response to a monetary policy shock. In the earlier period, it moves up somewhat in response to a monetary contraction, though not by a lot relative to its overall standard deviation. This does not accord well with the notion that the NK Phillips curve can be the center of a causal chain in which restrictive monetary policy reduces inflation by reducing current and expected future marginal costs measured this way. This does not mean that a NK Phillips curve could not play an important role in aligning an equilibrium model with the data. Indeed it looks as if the NK Phillips curve might help explain why the general price level is so slow to respond to monetary contraction — monetary contraction may produce a fall in productivity, rising or slowly falling

marginal costs, and hence, via the NK Phillips curve mechanism, a tendency for price decreases to lag behind wage decreases.

We can also consider the second question — how important is the NK Phillips curve as a way to understand determinants of inflation other than policy disturbances. For this we can look at the responses of the price level to all seven sources of disturbance in the system, as in Figures 3 and 4. Since labor share, in logs, is wage minus output per hour minus price, the effects of a surprise change in labor share are the corresponding linear combination of lines on the graph — aqua minus purple minus blue. In the pre-79:3 graph, this nets out to close to zero. In the post-83 graph, there is a strong effect of productivity surprises, with high productivity leading to increased inflation. This is not due to any complicated behavior of productivity in response to its own shocks. Productivity shocks are the main source of variation in labor share, and they produce sustained, single-signed movements in labor share. What produces this pattern is not clear from this partially identified model; but it is clear that the unidirectional NK Phillips curve causal chain is not at work here, as declining costs are associated with increasing inflation.

My conclusion is that the data show a perhaps surprisingly stable pattern of influence on prices, wages and output. Monetary policy is not neutral. But thinking about this pattern in terms of the NK Phillips curve does not appear to be helpful.

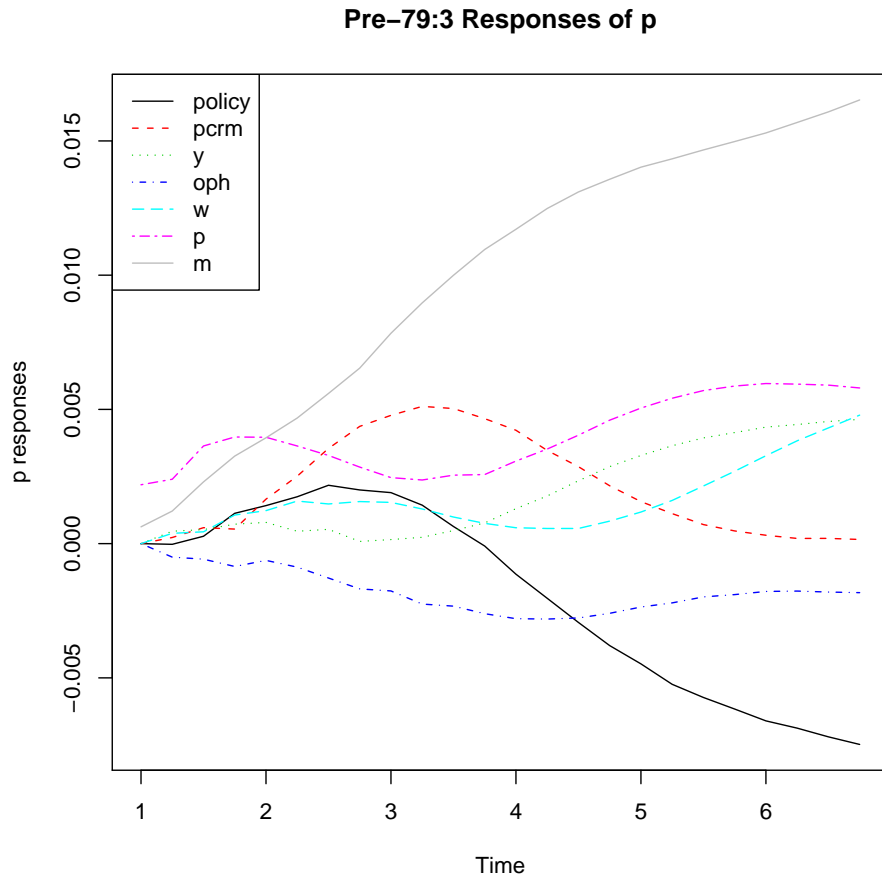


FIGURE 3.

### III. INFLATION-DETERMINATION WITHOUT A PHILLIPS CURVE

If we cannot rely on a single Phillips-curve like equation to organize our thinking about inflation, what is the replacement? There are two main directions to pursue, I think. One, which I will take up below, is to explore theories about deviations from the simple rational expectations paradigm. This may help us understand not only price stickiness and non-neutrality, but also sluggishness and inertia in economic




FIGURE 4.

behavior more generally. The other, which can be fruitfully pursued even within the rational expectations framework, is to be more explicit and systematic in taking a full dynamic general equilibrium approach to macro modeling, and in particular to model more carefully the interaction of monetary policy with asset markets and the interaction of asset markets with “the real economy”.

Current and expected future fiscal and monetary policy have immediate and strong impacts on asset markets. In a fully articulated dynamic equilibrium model with

rational agents, these impacts involve invoking transversality conditions. I have a colleague who interrupts every discussion of this kind of model with “Is this going to involve transversality conditions?”. His view is that few if any economists really understand transversality conditions (which is also my view) and that it is therefore unreasonable to entertain models that invoke transversality conditions to explain the behavior of actual human beings.

But transversality conditions apply even to less-than-hyper-rational agents. They are really just a name for wealth effects. If monetary policy raises the rate of return on government bonds, and if agents project that this rise in the relative return of government paper will be persistent, government paper becomes more attractive, people will tend to trade other assets for government paper, and there will therefore be downward pressure on the rate at which government paper trades for other goods — i.e. the price level. But there are conditions under which a rise in interest rates on government bonds, generated by the central bank, will not lead bond-holders to believe in persistently higher returns on government bonds. Higher real returns are possible, in general equilibrium, only if increased primary surpluses emerge in response to the higher interest rates. In an economy in which political economy or bureaucratic inefficiency makes increased primary surpluses impossible, the higher interest rates will only generate an increased rate of issue of government paper, with no increased rate of return — indeed with capital losses for holders of long nominal debt. It may take some time for bondholders to appreciate the nature of these fiscal

dynamics, so that the inflationary effects of increased interest rates do not take hold immediately. But this only makes the real value of the outstanding debt at current prices increase more rapidly, so that when the realization that the increased debt has no real backing sinks in, the eventual effects on demand are even larger. This kind of situation is widely acknowledged to have existed in some countries and some time periods, especially where interest expense has become a large fraction of the total government debt and nominal interest rates are high.

Most macroeconomists, though, think of this type of scenario as applying perhaps to Brazil in some periods, but not to the US, ever. My view is that we should reevaluate this possibility. Our recent history of a stock market boom, a housing price boom, then a commodity price boom and a decline in the value of the dollar, may be best understood as reflecting the evolution of thinking by bondholders about current and future US monetary and fiscal policy. In the 1970's when the US had its great burst of inflation, fiscal policy was by some measures much more unstable than monetary policy. On average over time any country that can issue debt must be running primary surpluses — the conventional surplus plus interest payments. The US ran primary surpluses in all but four of the years from 1962 through 1974, for example, but ran primary deficits every year from 1975 through 1994, except for two years of small primary surpluses. Then from 1995 through 2002 it ran large primary surpluses, to the point where it seemed the US government debt might essentially vanish. And now we are again in a period of primary deficits. What ended the long

period of primary deficits? What were bondholders thinking about future fiscal policy in this period? How did interest rate policy, which during the early 80's was causing large changes in the size of the interest expense component of the budget, interact with the political economy of fiscal policy?<sup>5</sup>

These issues are of course only one component of a full general equilibrium approach to assessing the effects of monetary and fiscal policy on inflation. Nonetheless, it seems to me that there may be high returns to focussing more of our attention on this component, even at the expense of less attention to the microeconomics of price and wage dynamics.

#### IV. DEPARTING FROM RATIONAL EXPECTATIONS: NEW IDEAS ABOUT MODELING THE EFFECTS OF UNCERTAINTY AND INERTIA

There is plenty of room for progress in integrating financial markets into our analysis of monetary and fiscal policy even within the framework of rational behavior and what Sargent calls the "communist" assumption that there is a single probability measure shared by nature and by all the economic agents interacting in a model. But increasingly economists are impatient with this assumption and interested in the implications of deviations from it. I will not try to catalog or discuss all the directions of deviation that economists have been exploring. Thomas Sargent, Giorgio Primiceri,

---

<sup>5</sup>In a 2008 paper I elaborate these points and present a model in which fiscal policy might have prevented the Fed from controlling inflation in the 1970's, even though the Fed was capable of creating recessions and corresponding temporary pauses in inflation.

Noah Williams, In-Koo Cho, George Evans, and Seppo Honkapohja, among others, have explored the implications of learning, both by policy-makers and private agents. Many economists, in the area labeled “behavioral economics”, have tried to incorporate insights from psychology and experimental evidence about deviations from rational expectations. Mankiw and Reis have proposed a theory I think of as Intermittent Observation theory, in which agents process information only at certain widely separated moments.

Here, though, I want to discuss two lines of thinking that I find particularly interesting and promising. One is rational inattention theory. This is a theory of why people do not use all of the information that lies in front of them “for free”. It invokes Shannon’s notion of a “channel” with finite “capacity” to process information, and assumes that people are such finite-capacity channels. This implies that there are limits on how quickly and precisely their behavior can react to information about a stochastically evolving economic environment. The attractive feature of Shannon’s theory for engineers is that it allows discussion of information flows and the capacity of information channels in a way that is quantitatively precise, yet abstracts from the physical characteristics of the channel and of the information. These days we are all familiar with the notion that our internet connections can be characterized by the “bits per second” figure that measures their Shannon capacity, and that this is a good measure of speed of transmission whether we are transmitting photos of grandchildren, spreadsheets of historical GDP data, or MP3 files downloading from

E-Music. The bits per second figure means the same thing for copper wire connections, fiber-optic connections, and cable connections.

This same independence of the hardware makes the theory attractive for modeling economic behavior, at least from the point of view of economists. It frees us from needing to know the details of the mental and physical limitations that prevent us from reacting at every moment to every bit of information impinging on us — we only need to know that the limitations exist, and to make the economist's usual assumption that information processing capacity, like other resources, is used optimally. I have explored these ideas in several papers (1998; 2003; 2006). The 2003 paper shows that the theory implies modifications in the permanent income model that bring it more closely in line with observed behavior. The 2006 paper considers a two-period savings model and shows that the theory can generate discretely distributed behavior, even in the face of continuously distributed information. By now a number of other economists have taken up these ideas, including Maćkowiak and Wiederholt (2005) and Matějka (2008), who show that some of the observed puzzling facts about microeconomic price behavior can be explained in the rational inattention framework.

The other area of recent research activity that I find interesting, though I have not contributed to it myself in published work, is exploring models in which rational agents, sharing the same information set and the same idea of the range of possible states of the world, have different probability distributions over those states. There

is no reason why rational optimizing agents need share the same probability distribution. When optimizing agents with different probability distributions interact in markets, they will be attracted to betting with each other, if not explicitly, then by borrowing, lending, and making speculative investments. Furthermore, if optimistic investors, having borrowed from pessimists, discover they were mistaken, there will be a rapid adjustment in asset prices, shifts in wealth between agents, and a high volume of transactions — all phenomena we see, and are concerned about, in actual asset markets.

## V. IMPLICATIONS FOR MONETARY POLICY

So what are the implications of these new strands of research for the Phillips curve, monetary policy, and macroeconomics more generally. I do not have space to consider all the implications here, but some interrelated implications are worth drawing out.

Rational inattention implies that people will behave as if they are observing market signals with error, and that agents with a bigger stake will invest more of their capacity in precise observation of a given signal. It therefore provides one rationale for why economic agents might have different probability distributions over the state of the economy, and for why the differences might persist despite the accumulation of “freely observable” evidence. Rational inattention and differences of opinion both may be related to why it is so hard, and yet so important, to model the interaction of asset markets with monetary policy and with the economy. Hard as

it may be to model how a set of “communist” rational agents would have modeled the future of fiscal policy in the 70’s and 80’s, it is harder still to imagine that every agent, whether he held bonds or not, whether she was 75 years old or 23, whether she was thinking of taking out a mortgage to buy a first home or had lived in the same house for 40 years and paid off her mortgage, had the same views about the future of fiscal policy and, therefore, the values of nominally denominated assets. Differences of views, learning, and rational inattention might explain why the interaction of monetary policy and fiscal policy with asset markets seems sometimes to work itself out on a long time scale. Not everyone will make the same assessment, at the same time, of the implications of transversality conditions. It may be that this can lead to wide swings in asset markets, and to delayed and unpredictable effects of monetary policy shifts.

We can’t model every person’s beliefs individually, and working formally with rational inattention theory at least at this point seems hard<sup>6</sup>. Nonetheless it seems important, especially in the light of the recent history of asset markets and their interaction with monetary policy, to get some working approximation of the effects of rational inattention and differences of beliefs into our models.

Where does this leave the Phillips curve? Something like the Phillips curve will continue to have a place in a general equilibrium model, as part of characterizing the

---

<sup>6</sup>Working formally with rational expectations theory also appeared hard at one point, though, so this may change.

interaction of costs, prices wages and output. But the rational inattention perspective suggests that locating stickiness and inertia in that one equation may be a mistake. The same limits on information processing may be at work in the slow adaptation of prices and wages to each other that are at work in sluggish reactions of consumption to income, or of investment to interest rates. Recognizing that sluggishness of various kinds may be related, through dependence on a common resource constraint, and that sluggishness represents conservation of a valuable resource, may lead to new modeling insights and to new ways of assessing the welfare implications of price stability and instability.

As with many important theories, the long run value of Phillips curve theories may lie in the new flames that are emerging from their dying embers.

[Note: In some sections of this version of the paper I discuss ideas drawn from work that I have not yet properly referenced.]

#### REFERENCES

**Lucas, Robert E. Jr.**, "Some International Evidence on Output-Inflation Tradeoffs,"

*American Economic Review*, jun 1973, 63 (3), 326–334.

\_\_\_\_\_ **and Leonard Rapping**, "Price Expectations and the Phillips Curve," *Ameri-*

*can Economic Review*, 1969.

\_\_\_\_\_ **and** \_\_\_\_\_, "Real Wages, Employment, and Inflation," *Journal of Political*

*Economy*, 1969.

**Maćkowiak, Bartosz and Mirko Wiederholt**, “Optimal Sticky Prices under Rational Inattention,” Technical Report, Humboldt University Berlin February 2005.

**Matějka, Filip**, “Rationally Inattentive Seller: Sales and Discrete Pricing,” Technical Report, PACM, Princeton University May 2008.

**Orphanides, Athanasios**, “Monetary Policy Rules, Macroeconomic Stability, and Inflation: A View from the Trenches,” Technical Report, Board of Governors of the Federal Reserve System May 2001.

**Primiceri, Giorgio**, “Why Inflation Rose and Fell: Policymakers’ Beliefs and US Post-war Stabilization Policy,” *Quarterly Journal of Economics*, August 2006, 121, 867–901.

**Sbordone, Argia**, “U.S. Wage and Price Dynamics: A Limited-Information Approach,” *International Journal of Central Banking*, September 2003, 2 (3).

**Sims, Christopher A.**, “Stickiness,” *Carnegie-rochester Conference Series On Public Policy*, 1998, 49 (1), 317–356.

\_\_\_\_\_, “The Role of Models and Probabilities in the Monetary Policy Process,” *Brookings Papers on Economic Activity*, 2002, 2002 (2), 1–62.

\_\_\_\_\_, “Implications of Rational Inattention,” *Journal of Monetary Economics*, April 2003, 50 (3), 665–690.

\_\_\_\_\_, “Rational Inattention: Beyond the Linear-Quadratic Case,” *American Economic Review*, May 2006, 96 (2), 158–163. Proceedings issue.

\_\_\_\_\_, “Stepping on a Rake: The Role of Fiscal Policy in the Inflation of the 1970’s,” Technical Report, Princeton University 2008.

**Woodford, Michael**, *Interest and Prices: Foundations of a Theory of Monetary Policy*,

Princeton University Press, 2003.