

INFLATION EXPECTATIONS, UNCERTAINTY, AND MONETARY POLICY

CHRISTOPHER A. SIMS

I. INTRODUCTION

Monetary economics as practiced by central bank modelers has made a great deal of progress in recent years. In a 2002 paper I interviewed research economists at four central banks and surveyed the models in use at those banks. I criticized the models for having lost all touch with statistical inference and with its connection to decision theory. I also criticized them for not following the rational expectations literature by jointly specifying and estimating the equations in their systems. And I pointed out that none of the models had a consistent treatment of asset markets. Since then many central banks, taking advantage of the new computational methods for Bayesian inference that economists are learning to use, have made substantial progress toward meeting the first two of these criticisms. They have still for the most part done little about the third. And academic economists are beginning to question some of the standard assumptions in the rational expectations framework that underlies these models.

Recent events in financial markets, and the difficulties that they raise for central banks, make it painfully clear that even the frontier Bayesian DSGE models like that

Date: October 10, 2008.

in use at the Swedish Riksbank do not model asset markets in any depth. But the problem goes beyond that: these models, and most academic macro models as well, assume a standard rational expectations framework: there is only one probability measure in play, the “true” probability measure from which nature draws realizations. Agents in the model form expectations using this true distribution, conditioning on information sets that consist of all information in the model dated t and earlier. It is well documented that people do not actually behave this way, and in the literature on behavioral finance there is some suggestion that deviations from this standardized assumption of rational behavior given a common probability distribution may be important.

The recent events in financial markets — the dotcom boom, the US house price boom, perhaps the continuing commodity price boom — look to some observers like bubbles that must have fed off some sort of irrational behavior. Many observers think that monetary policy might have somehow fueled these bubble-like episodes in asset markets. These are important questions for monetary policy, and it is disturbing that the monetary policy models in use cannot even be used to pose these questions.

In this paper I focus on two particular, and related, deviations from the assumption that all agents have the same probability distribution and that they optimally process all information available up to some date t . I consider the implications of

agents' being able to process information only at a limited rate, and the implications of agents' assuming differing probability distributions.

II. RATIONAL INATTENTION

Rational inattention theory explains 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 make 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 people from reacting at every moment to every bit of information impinging on them — 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.

In addition to its ability to predict sluggish, noisy, and discontinuous reactions of rational agents to information, rational inattention theory suggests that they will have persistent differences of opinion, due to the fact that they are all economizing, in different ways, on their use of information.

III. CAN MONETARY POLICY FUEL SPECULATION?

Savage's axioms for decision-making under uncertainty imply that a rational economic agent making decisions under uncertainty will act as if he is maximizing expected utility under some probability distribution over the uncertain states of the

world. But nothing in these axioms implies that every rational agent must have the *same* probability distribution over uncertain states. There is relatively little economic theory that considers the case where opinions, in the sense of probability distributions over states, differ. One reason for this neglect is that if people start with differing opinions, but view the same stream of evidence and process it optimally, their opinions will tend to converge. Differing opinions are thus seen as rare, one-time situations, not characteristic of a dynamic, stochastic, steady state. But the rational inattention theory we have discussed above provides a rationale for something we all know to be true: in the stochastic steady state we actually live in, the real world, there are lots of differences of opinion.

Rationally inattentive people in a stochastically evolving environment who have the same flow of data available to them at no cost will have persistently differing opinions for two reasons. One is that, even if they have the same objective functions and constraints, they can have unrelated signal-processing error. The error in their implicit signals means that their actions and opinions contain a random component at every date, and thus that their beliefs will differ. The theory does not imply that this must be true; it implies only that there will be processing error. It could be identical across individuals, or it could be completely independent across individuals. It seems likely that neither extreme is usually correct, that people filter and simplify their information streams in part through common mechanisms — reading news sources, imitating what others are doing — but also in part idiosyncratically.

The idiosyncratic part will lead to persistent differences in probability distributions across agents. Perhaps more important is that not all data is equally useful to everyone. People contemplating taking out a fixed rate mortgage will likely follow news about interest rates closely around the time of the transaction, while people living in houses with paid-up mortgages and living off social security payments might easily totally ignore news about interest rates, even though the information would be of some value to them — it might simply not be valuable enough to displace attention paid to other aspects of life.

Rational inattention is not the only reason for differences of opinion, though. In periods where genuinely new phenomena are arising, or when policy seems to be on a new and unpredictable path, the argument that a long history of repeated observation leads to agreement loses its force. For example in the period 1975-2000, the wide swings in US fiscal policy (discussed below) could easily have led to differing views about the implications of those swings for future inflation. And in the late 90's in the US, when unemployment and interest rates stayed persistently low, there were differences of view even among specialist economists about the long term implications for the inflation rate.

It is sometimes suggested that low interest rates in the US fueled the dot-com boom in the stock market, the house price boom, the recent commodity price boom, or all three. It seems impossible to support this suggestion in a standard equilibrium

model with rational agents, except by assuming some form of irrationality or friction in the market. In a standard model, the monetary authority controls only the evolution of the price level. If the model has no money illusion and flexible prices, the monetary authority has little or no influence over real activity. Its actions control the nominal interest rate and the path of the price level, not any real asset price.

There are some theoretical models that study markets in which agents with differing opinions interact. Scheinkman and Xiong (2003a) provides a useful survey. The idea that differences of opinion can raise asset prices, at least if short sales are not possible, goes back at least to Miller (1977). Harrison and Kreps (1978) showed that this result holds in a dynamic model in which agents have no risk aversion and in which short sales are not possible. Wahhaj (1999) showed that when short sales are possible and agents are risk-averse, the classic result no longer holds in general. Recently Brunnermeier and Julliard (2008) have studied irrational beliefs based on money illusion (i.e. a difference between “true” beliefs and those held by borrowers) and Scheinkman and Xiong (2003b) have studied a model in which difference of belief emerge from the documented psychological tendency for people to exaggerate the precision of their own beliefs. Of these papers, only that of Brunnermeier and Julliard considers monetary policy explicitly, and they find an effect of monetary policy by postulating simple money illusion.

Here we develop a model in which no agent has money illusion, markets are frictionless, short sales are allowed, real investment produces returns according to

a non-stochastic production function known by all, and nonetheless differences of opinion about the course of inflation generate overinvestment in the real asset. The mechanism is fairly easy to understand once it is laid out. Suppose one group of agents believes that inflation is likely to be high and the return on nominal bonds therefore low. Another group believes that inflation is likely to be lower. Both face the same real return on investment, which depends on the aggregate level of investment. The return on real investment in equilibrium must match what each agent sees as the expected discounted return on the risky bonds. Since they agree on the return on real investment, it must be that the agent who expects low inflation wants to hold a lot of bonds, perhaps even lending to (i.e. buying nominal bonds issued by) the other agent. He may also want to sell the stock short. He will need a portfolio of this sort so that he perceives the high expected return on bonds as being offset by the fact that they are highly correlated with his total portfolio return, justifying a risk premium on them. The other agent, on the other hand, perceives selling bonds and borrowing from the first agent as a source of cheap capital to invest in the real asset. Because of the cheap source of capital, the second agent will, if he is not very risk averse, invest more in the real asset than he would if all agents shared his beliefs, and indeed so much more that the economy's total investment is higher than it would be if all agents shared the same beliefs. This all occurs only because of the heterogeneity of opinions. If all agents in this model have the same beliefs

about monetary policy, whatever the beliefs may be, the amount of real investment is invariant to their beliefs.

Because the model is meant only to make a point, not to be quantitatively realistic, it is extremely simple. Agents live two periods. They each begin life with an endowment of nominal bonds B_0 , and in the first period of life they each have an endowment Y of goods. They can consume in the first period, and the amount of their consumption is C_1 . They can finance their first period consumption and their investment S in the real asset from their real endowment or by selling some of their bonds. They can also purchase more bonds.

There are two types of agents $i = a, b$. There are two possible states of the world in the second period, states $j = f, m$. In the f state, the tax backing for bonds is low, and hence prices are high. In the m state taxes are high and prices are therefore lower. The problem of the agent of type i can therefore be written as

$$\max_{C_{i1}, B_i, S_i, C_{i2f}, C_{i2m}} U(C_1) + \beta(p_i U(C_{i2f}) + (1 - p_i)U(C_{i2m})) \quad \text{subject to} \quad (1)$$

$$C_{i1} + S_i + \frac{B_i - B_0}{P_1} = Y \quad (2)$$

$$C_{i2j} = \rho S_i + RB_i / P_{2j} - \tau_j + \delta, \quad j = f, m \quad (3)$$

$$(4)$$

Here ρ is the rental price of capital in the second period, τ_j is the lump-sum tax rate in the second period in state j , and δ is the profit dividend. We think of both types of agents as being endowed with half the ownership rights in the technology, so they

get dividends from the pure profits of the representative firm that are unaffected by the amount of capital they set aside in the first period to rent out in the second. Because this mode has to be solved numerically, we assume specific, convenient functional forms for U and the production function: $U(C) = C^{1-\sigma}/(1-\sigma)$ (with $\log C$ as a limiting case as $\sigma \rightarrow 1$) and $g(S) = S^{1-\alpha}$. Profit maximizing representative firms will then require

$$\rho = (S_a + S_b)^{-\alpha} \quad (5)$$

The government fixes R , the gross nominal interest rate, as well as τ_i , $i = f, m$. Its second period budget constraints are

$$\frac{RB_0}{P_{2j}} = \tau_j \quad j = f, m. \quad (6)$$

The government does no taxing, spending, or debt sales in the initial period, so market clearing requires $2B_0 = B_a + B_b$.

The first-order conditions for the agents lead to

$$\partial S : \quad C_{i1}^{-\sigma} = \rho \cdot (p_i C_{if}^{-\sigma} + (1-p_i) C_{im}^{-\sigma}) \quad (7)$$

$$\partial B : \quad \frac{1}{C_{i1}^\sigma P_1} = R\beta \left(\frac{p_i R}{P_{2f}} + \frac{(1-p_i)R}{P_{2m}} \right) \quad (8)$$

These equations, though I think not soluble analytically, are numerically tractable. To make this section's main point, I display two solutions, differing only in that in one $p_a = p_b = .5$, so beliefs are the same across the two types of agent, while in the other $p_a = .3$, $p_b = .7$, so the type a agent believes it is more likely that the tax backing

Y	R	τ_f	τ_m	α	β	A	σ	B_0
1.60	1.10	1.10	1.65	0.30	0.90	1.20	0.50	1.50

TABLE 1. Common parameter values

for the debt will be strong, and inflation therefore low, while the type b agent believes the probabilities are the reverse. The parameters that stay the same across the two solutions are shown in Table 1. The two solutions are shown in Table 2. Note that the total real investment in the solution with differing opinions is .96, while in the symmetric solution it is .88. Observe also that when opinions differ, there is a great deal of borrowing and short selling, with the agent who thinks nominal bonds are the better investment buying the entire initial stock from the other agent, and then lending him nearly as much again. The agent who thinks real capital the better investment buys promises to pay the capital return in an amount more than triple the actual amount of capital, while the other agent short sells a large amount of capital. This pattern, in which differences of opinion lead to large amounts of short selling and lending, is robust. The finding that differences of opinion about real interest rates lead to excess investment in real capital depends on the low value of σ we have chosen. To get this result, we must have agents who want to buy more capital when they perceive its return is relatively high. If $\sigma > 1$, Agents who perceive a high return reduce their current saving, and indeed in that case differences of opinion about real bond interest rates, while still leading to large amounts of lending and short selling, reduce rather than increase investment in real capital. The log-utility boundary case

	$p_a = .3, p_b = .7$	$p_a = p_b = .5$
C_{a1}	1.1189	1.1590
C_{a2f}	0.2356	0.7143
C_{a2m}	1.2828	0.7143
C_{b1}	1.1189	1.1590
C_{b2f}	1.2828	0.7143
C_{b2m}	0.2356	0.7143
B_a	4.3559	1.5000
B_b	-1.3559	1.5000
P_1	0.9270	0.9515
P_{2f}	1.5000	1.5000
P_{2m}	1.0000	1.0000
S_a	-2.5996	0.4409
S_b	3.5618	0.4409
ρ	0.8497	0.8722
δ	0.3504	0.3296

TABLE 2. Two solutions

makes heterogeneity of opinions unimportant for determining the aggregate amount of investment.

These exercises are not meant to be quantitatively realistic. Most economists think that asset market behavior suggests that investors tend to have $\sigma > 1$, for example. However in this model, real capital is the only way to generate future income. If we

had a more realistic model, with many types of investment good, large proportional changes in investment in any one good could be financed with smaller changes in bond holdings. This might make the low risk aversion in this example a better approximation. Also, the case where differences of opinion about bond returns lead to decreases in real investment might correspond to a situation that is perceived as speculative excess. Though the total amount of real investment is reduced, this happens, under high risk aversion, because all the investment is being done by the agent who sees bond returns as low; that agent's investment is much higher than it would be in the symmetric equilibrium. In fact the total investment is only lower because this type of agent, being risk averse, holds back due to the risky leveraged position his portfolio puts him in.

The point here is that the notion that some aspect of monetary policy might be related to distorted speculative excess in asset markets does not rest on invoking imprecise notions of asset market imperfection or irrationality of agents. Agents with differing views of probabilities will use asset markets to bet against each other, and in the process can push real allocations in directions that would not have arisen if either agent's belief were common across all agents in the economy.

IV. 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.

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.¹

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.

V. 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

¹I discussed the state of central bank modeling in a 2002 Brookings paper.

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? In an earlier, related paper 2008a I showed that structural VAR estimates of fairly strong effects of monetary policy on real activity, prices, and wages do not appear to be mediated by the marginal cost variable most commonly used in the recent NK Phillips curve literature, the share of labor in total costs. This does not suggest that the NK Phillips curve is refuted, or that it should not appear in the DSGE models where it is widely used. Indeed, it may play an important role in explaining why

consumption good prices respond considerably more slowly to a monetary contraction or expansion than do wages or commodity prices. But thinking of monetary policy as acting on inflation by first changing some measure of real tightness, like labor share, then affecting prices, seems to be missing the central part of the story.

VI. 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, already mentioned above, 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 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 1972 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?²

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.

VII. 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 they 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

²In a 2008b 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 it was capable of creating recessions and corresponding temporary pauses in inflation.

of asset markets with monetary policy and with the economy. Hard as it may be to model how a set of rational agents with a single probability distribution 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.

Recognition that diverse opinions about the course of the price level can be important, and that agents display rational inattention, has some immediate implications, it seems to me, about central bank communications with the public. On the one hand, rational inattention theory suggests that when monetary policy has been going well, one of its benefits is that people will pay little attention to it — and therefore may misperceive or ignore policy changes. This is a benefit because attention is a scarce resource. I think it likely that one of the main costs of high and variable inflation is that it forces people to spend a considerable fraction of their limited

information-processing capacity on tracking the price level and the exchange rate. Thus it is not a problem that the public pays little attention to monetary policy, most of the time. But there may be periods when policy has to change, and misperception of the change by the public could be costly. Rational inattention theory suggests that people will, no matter how information is presented to them, find ways to process it optimally. They will, therefore, try to be sure that they pay attention to monetary policy when it is important, perhaps relying on the services of information filters like newspapers (or, these days, internet news sites). But they will not be able to do this unless the information is there. It is a mistake, therefore, to take the evidence that in quiet times people ignore or misperceive monetary policy pronouncements as a reason to limit the flow of information about monetary policy.

This conclusion is amplified when we recognize that diversity of views about future monetary and fiscal policy can be a source of distortions of the behavior of real asset markets. If agents are forced to infer monetary policy from the time series of policy rate changes and from terse and cryptic summaries of the rationale for the rate changes, they will introduce their own signal processing errors and thereby make diversity of views more likely. The ideal communication strategy might then be multi-tiered. Very detailed and analytical descriptions of policies and the changes in it like those produced in the inflation reports of inflation-targeting central banks might be accompanied by more easily tracked simplified characterizations of policy. The point of the simplified presentations is not to hide detail from the public, but

to shape the simplified view that the public is bound to form, even if given the detailed information flow. It is worthwhile to try to move the public toward a common simplified view of monetary and fiscal commitments, rather than having them form views idiosyncratically and then bet with each other in asset markets.

REFERENCES

- Brunnermeier, Markus and Christian Julliard**, "Money Illusion and Housing Frenzies," *Review of Financial Studies*, 2008, 21 (1), 135–180.
- Harrison, Michael and David Kreps**, "Speculative investor behavior in a stock market with heterogeneous expectations," *Quarterly Journal of Economics*, 1978, 92, 323–336.
- 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," *American 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.

Miller, Edward, "Risk, uncertainty and divergence of opinion," *Journal of Finance*, 1977, 32, 1151–1168.

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).

Scheinkman, Jose and Wei Xiong, "Heterogeneous Beliefs, Speculation and Trading in Financial Markets," Paris-Princeton Lectures in Mathematical Finance, Springer 2003.

_____ and _____, "Overconfidence and Speculative Bubbles," *Journal of Political Economy*, 2003, 111, 1183–1219.

Sims, Christopher, "Inflation Expectations, Uncertainty, the Phillips Curve, and Monetary Policy," Technical Report, Princeton University June 2008. Presented at the June, 2008 Boston Federal Reserve Bank monetary policy conference.

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.

Wahhaj, Zaki, "Trading in Financial Markets with Differences of Beliefs," Undergraduate honors essay, Yale College 1999.

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