

Imperfect Common Knowledge and the Effects of Monetary Policy *

Michael Woodford
Princeton University

October 29, 2001

Abstract

This paper reconsiders the Phelps-Lucas hypothesis, according to which temporary real effects of purely nominal disturbances result from imperfect information about the nature of these disturbances. This explanation for the real effects of monetary policy is often dismissed on the ground that the Lucas (1972) model predicts only highly transitory effects on real activity, and none at all insofar as changes in monetary policy can be observed by the public prior to any measurable effect on aggregate nominal expenditure.

The present paper finds instead that an imperfect-information model can predict highly persistent effects on real activity if the assumptions of the Lucas model are modified in two crucial respects. First, a model of monopolistically competitive pricing is assumed in which optimal pricing decisions are highly dependent upon the prices set by others, so that higher-order expectations are crucial for aggregate inflation dynamics, as argued by Phelps (1983). And second, decisionmakers' subjective perceptions of current conditions are assumed to be of imperfect precision, owing to finite information processing capacity, as argued by Sims (2001).

Noisy subjective perceptions, even when rationally interpreted by the decisionmakers, lead to greater uncertainty about higher-order expectations (others' expectations of others' expectations ...) than about the aggregate disturbance itself, as stressed in the recent literature on "global games". This in turn makes higher-order expectations adjust only sluggishly to a disturbance, even when the public's average estimate of what has occurred adjusts fairly rapidly. If strategic complementarities in price-setting are strong enough, the real effects of a nominal disturbance may be both large and highly persistent.

*Prepared for the Festschrift Conference in Honor of Edmund S. Phelps, Columbia University, October 5-6, 2001. I would like to thank Greg Mankiw and Lars Svensson for comments, Charlie Evans for sharing his VAR results, Hong Li for research assistance, and the National Science Foundation for research support through a grant to the NBER.

The dynamic responses of inflation and output to a monetary disturbance implied by this model are compared to those implied by a standard sticky-price model, the model of staggered pricing proposed by Calvo (1983). Under certain circumstances, the models imply identical responses. However, when the effect of the disturbance on nominal GDP growth exhibits substantial persistence, the predictions of the models differ. The noisy information model, unlike the Calvo model, is consistent with an inflation response that lags the effects of the monetary disturbance on real activity, and in this respect matches the empirical evidence better.

1 Imperfect Information and Price Adjustment

A perennial question in macroeconomic theory is the reason for the observed real effects of changes in monetary policy. It is not too hard to understand why central-bank actions can affect the volume of *nominal* spending in an economy. But why should not variations in nominal expenditure of this sort, not associated with any change in real factors such as tastes or technology, simply result in proportional variation in nominal wages and prices, without any effect upon the quantities produced or consumed of anything? It has long been observed that wages and prices do not immediately adjust to any extent close to full proportionality with short-run variations in nominal expenditure, but again, why should not self-interested households and firms act in a way that brings about more rapid adjustment?

A famous answer to this question is that people are not *well enough informed* about changes in market conditions, at least at the time that these changes occur, to be able immediately to react in the way that would most fully serve their own interests. Phelps (1970) proposed the parable of an economy in which goods are produced on separate “islands,” each with its own labor market; the parties determining wages and employment on an individual island do so without being able to observe either the wages or production decisions on other islands. As a result of this informational isolation, an increase in nominal expenditure on the goods produced on all of the islands could be mis-interpreted on each island as an increase in the relative demand for the particular good produced there, as a result of which wages would not rise enough to prevent an increase in employment and output across all of the islands. Lucas (1972) showed that such an argument for a short-term Phillips-curve tradeoff is consistent with “rational expectations” on each island, *i.e.*, with expectations given by Bayesian updating conditional upon the market conditions observed on that island, starting from a prior that coincides with the objective *ex ante* probabilities (according to the model) of different states occurring.

This model of business fluctuations was, for a time, hugely influential, and allowed the development of a number of important insights into the consequences for economic policy

of endogenizing the expectations on the basis of which wages and prices are determined. However, the practical relevance of the imperfect-information model was soon subjected to powerful criticism. In the Lucas model, equilibrium output differs from potential only insofar as the *average estimate* of current aggregate nominal expenditure differs from the actual value. In terms of the log-linear approximate model introduced in Lucas (1973) and employed extensively in applied work thereafter, one can write

$$y_t = \alpha(q_t - q_{t|t}), \tag{1.1}$$

where $0 < \alpha < 1$ is a coefficient depending upon the price-sensitivity of the supply of an individual good. Here y_t denotes the deviation of aggregate (log) real GDP from potential, q_t denotes aggregate nominal GDP, and $q_{t|t}$ the average (across islands) of the expected value of q_t conditional upon information available on that island in period t .

Furthermore, all aggregate disturbances in period t — and hence the volume of aggregate nominal expenditure q_t — become public information (observable on all islands) by date $t+1$. This implies that

$$E_t[q_{t+1|t+1}(i)] = E_t[q_{t+1}]$$

in the case of each island i , where $E_t[\cdot]$ denotes an expectation conditional upon the history of aggregate disturbances through date t , and $q_{t+1|t+1}(i)$ the expectation of q_{t+1} conditional upon the information available on island i in period $t+1$. Averaging over i , it follows that

$$E_t[q_{t+1|t+1}] = E_t[q_{t+1}].$$

Then, taking the expectation of both sides of (1.1) for period $t+1$ conditional upon the history of aggregate disturbances through date t , it follows that

$$E_t[y_{t+1}] = 0. \tag{1.2}$$

Equation (1.2) implies that deviations of output from potential cannot be forecasted a period earlier by someone aware of the history of aggregate disturbances up to that time. This means that a monetary disturbance in period t or earlier cannot have any effect upon

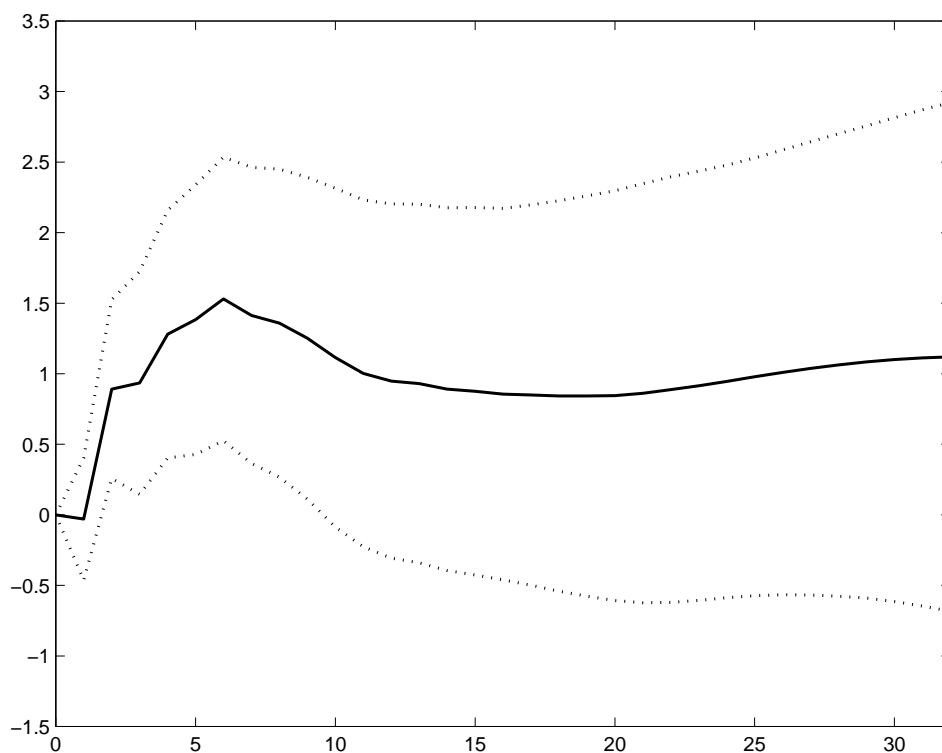


Figure 1: Estimated impulse response of nominal GDP to an unexpected interest-rate reduction. Source: Christiano *et al.* (2001).

equilibrium output in period $t + 1$ or later. But it follows that such real effects of monetary disturbances as are allowed for by (1.1) must be highly transitory: they must be present only in the period in which the shock occurs. The model was accordingly criticized as unable to account for the observed *persistence* of business fluctuations.

Of course, the degree to which the prediction of effects that last “one period” only is an empirical embarrassment depends upon how long a “period” is taken to be. In the context of the model, the critical significance of a “period” is the length of time it takes for an aggregate disturbance to become public information. But, many critics argued, the value of the current money supply is published quite quickly, within a few weeks; thus real effects of variations in the money supply should last, according to the theory, for at most a few weeks. Yet statistical analyses of the effects of monetary disturbances indicated effects lasting for many quarters.

Furthermore, the theory implied that monetary disturbances should not have even transitory effects on real activity, except insofar as these resulted in variations in aggregate nominal expenditure that *could not be forecasted* on the basis of variables that were already public information at the time of the effect on spending. But the VAR literature of the early 1980s (e.g., Sims, 1980) showed that variations in the growth rates of monetary aggregates were largely forecastable in advance by nominal interest-rate innovations, and that the monetary disturbances identified by these interest-rate surprises had no noticeable effect upon nominal expenditure for at least the first six months. This has been confirmed by many subsequent studies; for example, Figure 1 shows the impulse response of nominal GDP to an unexpected loosening of monetary policy in quarter zero, according to the identified VAR model of Christiano *et al.* (2001). (Here the periods on the horizontal axis represent quarters, and the dashed lines indicate the ± 2 s.e. confidence interval for the response.) Although the federal funds rate falls sharply in quarter zero (see their paper), there is no appreciable effect upon nominal GDP until two quarters later.

Thus given the estimated effects of monetary disturbances upon nominal spending — and given the fact that money-market interest rates are widely reported within a day — the Lucas model would predict that there should be *no effect* of such disturbances upon real activity at all, whether immediate or delayed. Instead, the same study finds a substantial effect on real GDP, as shown in Figure 2. Furthermore, the real effects persist for many quarters: the peak effect occurs only six quarters after the shock, and the output effect is still more than one-third the size of the peak effect ten quarters after the shock.

These realizations led to a loss of interest, after the early 1980s, in models of the effects of monetary disturbances based upon imperfect information — and indeed, in a loss of interest in monetary models of business fluctuations altogether, among those who found unpalatable the assumption of non-informational reasons for slow adjustment of wages or prices. However, this rejection of the Phelpsian insight that information imperfections play a crucial role in the monetary transmission mechanism may have been premature. For the unfortunate predictions just mentioned relate to the specific model presented by Lucas

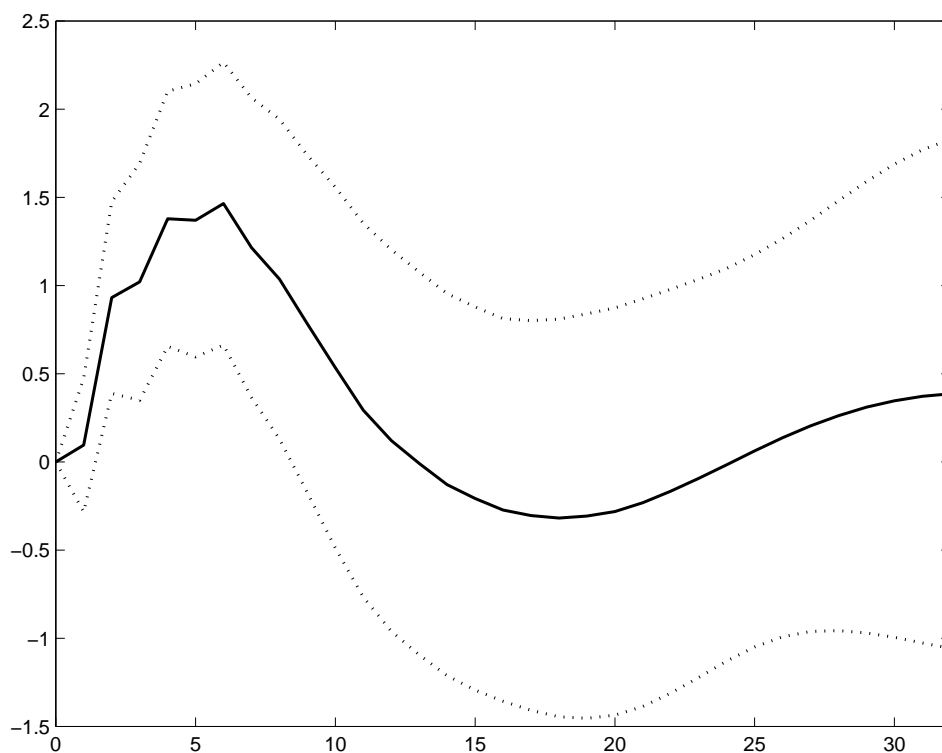


Figure 2: Estimated impulse response of real GDP to an unexpected interest-rate reduction. Source: Christiano *et al.* (2001).

(1972), but not necessarily to alternative versions of the imperfect-information theory.

Persistent effects of monetary disturbances on real activity can instead be obtained in a model that varies certain of Lucas’ assumptions. In particular, one may argue that the Lucas model does not take seriously enough the Phelpsian insight that informational isolation of the separate decisionmakers in an economy — captured by the parable of separate “islands” — is an important source of uncertainty on the part of each of them as to what their optimal action should be. For in the Lucas model, the only information that matters to decisionmakers, about which they have imperfect information, is the current value of an exogenous aggregate state variable: the current level of nominal GDP (or equivalently in that model, the current money supply). Instead, for the “isolated and apprehensive ... Pinteresque figures” in an economy of the kind imagined by Phelps (1970, p. 22), an important source of uncertainty is the unknowability of *the minds of others*.

Here I follow Phelps (1983) in considering a model in which the optimal price for any given supplier of goods to charge depends not only upon the state of aggregate demand, but also upon the average level of prices charged by *other* suppliers. It then follows that the price set by that supplier depends not only upon its own estimate of current aggregate demand, but also upon its estimate of the average estimate of others, and similarly (because others are understood to face a similar decision) upon its estimate of the average estimate of that average estimate, and so on. The entire infinite hierarchy of progressively higher-order expectations matters (to some extent) for the prices that are set, and hence for the resulting level of real activity.

This is important because, as Phelps argues, *higher-order* expectations may be even slower to adjust in response to economic disturbances. Phelps (1983) suggests that rational expectations in the sense of Lucas (1972) are a less plausible assumption when the hypothesis must be applied not only to estimates of the current money supply, but also to an entire infinite hierarchy of higher-order expectations. But here I show that higher-order expectations can indeed be expected to adjust more slowly to disturbances, even under fully rational expectations.¹ The reason is that even when observations allow suppliers to infer that aggregate demand has increased, resulting in a substantial change in their *own* estimate of current conditions, these observations may provide less information about the way in which *the perceptions of others* may have changed, and still less about others' perceptions of others' perceptions. Thus in the model presented here, a monetary disturbance has real effects, not so much because the disturbance passes unnoticed as because its occurrence is not *common knowledge* in the sense of the theory of games.

A second important departure from the Lucas (1972) model is to abandon the assumption that monetary disturbances become public information — and hence part of the information set of every agent — with a delay of only one period. Were we to maintain this assumption, it would matter little that in the present model output depends not only upon the discrepancy

¹Previous illustrations of the way that additional sources of persistence in economic fluctuations can be created when higher-order expectations matter include Townsend (1983a, 1983b) and Sargent (1991). These applications do not, however, consider the issue of the neutrality of money.

between q_t and $q_{t|t}$, but on the discrepancy between q_t and higher-order average expectations as well. For if the monetary disturbance at date t is part of every supplier's information set at date $t + 1$ (and this is furthermore common knowledge), then any effect upon q_{t+1} of this disturbance must increase not only $q_{t+1|t+1}$ but also all higher-order expectations by exactly the same amount. (The argument is exactly the same as in our consideration above of the effect upon first-order average expectations.) We would again obtain (1.2), and the criticisms of the Lucas model mentioned above would continue to apply.

Hence it is desirable to relax that assumption. But how can one realistically assume otherwise, given the fact that monetary statistics are reported promptly in widely disseminated media? Here it is crucial to distinguish between *public information* — information that is available *in principle* to anyone who chooses to look it up — and the information of which decision-makers are *actually aware*. Rather than supposing that people are fully aware of all publicly available information — a notion stressed in early definitions of “rational expectations”, and of critical importance for early econometric tests of the Lucas model² — and that information limitations must therefore depend upon the failure of some private transactions to be made public, I shall follow Sims (1998, 2001) in supposing that the critical bottleneck is instead the limited capacity of private decision-makers to *pay attention* to all of the information in their environment.³

In the model presented below, I assume that each decision-maker acts on the basis of his or her own subjective *perception* of the state of aggregate demand, that I model as observation of the true value with error (a subjective error that is idiosyncratic to the individual observer).⁴

²Lucas (1977, sec. 9), however, implicitly endorses relaxation of this position, when he suggests that it is reasonable to suppose that traders do not bother to track aggregate variables closely. “An optimizing trader will process those prices of most importance to his decision problem most frequently and carefully, those of less importance less so, and most prices not at all. Of the many sources of risk of importance to him, the business cycle and aggregate behavior generally is, for most agents, of no special importance, and there is no reason for traders to specialize their information systems for diagnosing general movements correctly.”

³A similar gap between the information that is publicly available and the information of which decisionmakers are actually aware is posited in the independent recent work of Mankiw and Reis (2001). The Mankiw-Reis model is further compared to the present proposal in section 4.3 below.

⁴The implications of introducing idiosyncratic errors of this kind in the information available to individual agents has recently been studied in the game-theoretic literature on “global games” (e.g., Morris and Shin, 2001). As in the application here, that literature has stressed that in the presence of strategic complemen-

That is, all measurements of current conditions are obtained through a “noisy channel” in the communications-theoretic sense (e.g., Ziemer and Tranter, 1995). Given the existence of private measurement error, agents will not only fail to immediately notice a disturbance to aggregate demand with complete precision, but they will continue to be uncertain about whether others know that others know that others know ... about it — even after they can be fairly confident about the accuracy of their *own* estimate of the aggregate state. Thus it is the existence of a gap between reality and perception that makes the problem of other minds such a significant one for economic dynamics.

Moreover, given the use of a limited “channel capacity” for monitoring current conditions, it will not matter how much and how accurate of information may be made “public” (e.g., on the internet). Indeed, in the model below I assume that all aggregate disturbances are “public information”, in the sense of being available in principle to anyone who chooses to observe them with sufficient precision, and in the sense of being actually observed (albeit with error) by every decision-maker in the entire economy. There is no need for the device of separate markets on different “islands” in order for there to be imperfect common knowledge. (Presumably, Phelps intended the “islands” as a metaphor for this sort of failure of subjective experience to be shared all along — though who can claim to know other minds?) Nor is there any need for a second type of disturbance (the random variations in *relative* demand of the Lucas model) in order to create a non-trivial signal-extraction problem. The “channel noise” generated by each decision-maker’s own over-burdened nervous system suffices for this purpose.

This emphasis upon the limited accuracy of private perceptions is in the spirit of recent interest in weakening the idealized assumptions of rational-decision theory in macroeconomics and elsewhere (e.g., Sargent, 1993). Limitations upon the ability of people (and animals) to accurately discriminate among alternative stimuli in their environments are better documented (and admit of more precise measurement) than most other kinds of cognitive

tarities, even a small degree of noise in the private signals can have substantial consequences for aggregate outcomes, owing to the greater uncertainty that is created about higher-order expectations.

limitations, having been the subject of decades of investigation in the branch of psychology known as “psychophysics” (e.g., Green and Swets, 1966). While it might seem that the introduction of a discrepancy between objective economic data and private perceptions could weaken the predictions of economic theory to the point of making the theory uninteresting, the type of theory proposed here — which assumes that agents correctly understand the characteristics of the noisy channel through which they observe the world, and respond optimally to the history of their subjective observations — is still relatively tightly parameterized. The proposed generalization here of a standard neoclassical model adds only a single additional free parameter, which can be interpreted as measuring the rate of information flow in the noisy channel, as in Sims (2001).

Section 2 develops a simple model of pricing decisions in an environment characterized by random variation in nominal spending and imperfect common knowledge of these fluctuations for the reason just discussed. It shows how one can characterize equilibrium output and inflation dynamics in terms of a finite system of difference equations, despite the fact that expectations of arbitrarily high order matter for optimal pricing policy. Section 3 then derives the implications of the model for the real effects of monetary disturbances, in the special case where erratic monetary policy causes nominal GDP to follow a random walk with drift, as in the Lucas model. It is shown that not only are deviations of output from potential due to monetary disturbances not purely transitory, but their degree of persistence may in principle be arbitrarily long. Indeed, arbitrarily long persistence of such real effects is possible (though less empirically plausible) even in the case of quite accurate individual perceptions of the current state of aggregate demand. The dynamics of higher-order expectations are also explicitly characterized, and it is shown that higher-order expectations respond less rapidly to a disturbance, as argued above.

Section 4 then compares imperfect common knowledge as a source of price inertia, and hence of real effects of monetary policy, to the more familiar hypothesis of “sticky prices,” in the sense of a failure of prices to be continuously updated in response to changing conditions. In the case of a random walk in nominal GDP, the predicted dynamics of output and inflation

are essentially the same in the model developed here and in the familiar Calvo (1983) model of staggered price adjustment — corresponding to any given assumed average frequency of price adjustment there is a rate of information acquisition that leads to the same equilibrium dynamics in the imperfect-information model, despite continuous adjustment of all prices. However, this equivalence does not hold for more generally stochastic processes for nominal GDP. In the case of positive serial correlation of nominal GDP growth (the more realistic specification as far as actual monetary disturbances are concerned), the predictions of the two models differ, and in a way that suggests that an assumption of incomplete common knowledge of aggregate disturbances may better match the actual dynamics of output and inflation following monetary disturbances. Section 5 concludes.

2 Incomplete Common Knowledge: A Simple Example

Here I illustrate the possibility of a theory of the kind sketched above by deriving a log-linear approximation to a model of optimal price-setting under imperfect information. The log-linear approximation is convenient, as in Lucas (1973) and many other papers, in allowing a relatively simple treatment of equilibrium with a signal-extraction problem.

2.1 Perceptions of Aggregate Demand and Pricing Behavior

Consider a model of monopolistically competitive goods supply of the kind now standard in the sticky-price literature. The producer of good i chooses the price p_t^i at which the good is offered for sale in order to maximize

$$E \left\{ \sum_{t=0}^{\infty} \beta^t \Pi(p_t^i; P_t, Y_t) \right\} \quad (2.3)$$

where period t profits are given by

$$\Pi(p; P, Y) = m(Y)[Y(p/P)^{1-\theta} - C(Y(p/P)^{-\theta}; Y)]. \quad (2.4)$$

Here Y_t is the Dixit-Stiglitz index of real aggregate demand, and P_t the corresponding price index, the evolution of each of which is taken to be independent of firm i 's pricing policy.

Firm i expects to sell quantity $y_t^i = Y_t(p_t^i/P_t)^{-\theta}$ if it charges price p_t^i , for some $\theta > 1$. Real production costs are given by $C(y_t^i; Y_t)$, where the second argument allows for dependence of factor prices upon aggregate activity. Finally, (2.4) weights profits in each state by the stochastic discount factor $m(Y_t)$ in that state, so that (2.3) represents the financial-market valuation of the firm's random profit stream. (See, e.g., Woodford, 2001.) The model here abstracts from all real disturbances.

I assume that the firm can choose its price independently each period, given private information at that time about the aggregate state variables. In this case, the pricing problem is a purely static one each period, of choosing p_t^i to maximize $E_t^i \Pi(p_t^i; P_t, Y_t)$, where E_t^i denotes expectation conditional upon i 's private information set at date t . The first-order condition for optimal pricing is then

$$E_t^i[\Pi_p(p_t^i; P_t, Y_t)] = 0. \quad (2.5)$$

In the absence of information limitations, each supplier would choose the same price (which then must equal P_t), so that equilibrium output would have to equal the *natural rate* of output \bar{Y} , defined as the level such that $\Pi_p(P; P, \bar{Y}) = 0$. (This is independent of P .)

To simplify the signal-extraction issues, I shall approximate (2.5) by a log-linear relation, obtained by Taylor-series expansion around the full-information equilibrium values $p_t^i/P_t = 1$ and $Y_t = \bar{Y}$.⁵ This takes the form

$$p_t(i) = p_{t|t}(i) + \xi y_{t|t}(i), \quad (2.6)$$

introducing the notation $p_t(i) \equiv \log p_t^i$, $p_t \equiv \log P_t$, $y_t \equiv \log(Y_t/\bar{Y})$, and letting $x_{t+j|t}(i) \equiv E_t^i x_{t+j}$ for any variable x and any horizon $j \geq 0$. Assuming that C is such that $C_y > 0$, $C_{yy} \geq 0$, and $C_{yY} > 0$, one can show that $\xi > 0$. I shall assume, however, that it satisfies

⁵We abstract here from any sources of real growth, as a result of which the full-information equilibrium level of output, or "natural rate" of output, is constant. Nothing material in the subsequent analysis would be different were we to assume steady trend growth of the natural rate of output. We abstract here from stochastic variation in the natural rate so that producers need *only* form inferences about the monetary disturbances. One advantage of this of this simplification is that it makes clear the fact that the present model, unlike that of Lucas (1972), does not depend upon the existence of both real and nominal disturbances in order for there to be real effects of nominal disturbances.

$\xi < 1$, so that the pricing decisions of separate producers are *strategic complements* (again see Woodford, 2001).

Finally, I specify the demand side of the economy by assuming a given stochastic process for aggregate nominal expenditure. A traditional justification for such an assumption is that the central bank determines an exogenous process for the money supply, and that there is a constant, or at any rate exogenous, velocity of money. Yet we need not assume anything as specific as this about the monetary transmission mechanism, or about the nature of monetary policy. All that matters for the analysis below is (i) that the disturbance driving the nominal GDP process is a monetary policy shock, and (ii) that the dynamic response of nominal GDP to such shocks is of a particular form. The assumption of a particular response of nominal GDP under historical policy is something that can be checked against time series evidence, regardless of how one believes that this response should best be explained. Direct specification of a stochastic process for nominal GDP eliminates the need for further discussion of the details of aggregate demand determination, and for purposes of asking whether our model is consistent with the observed responses of real activity and inflation to monetary disturbances, this degree of detail suffices.⁶

Letting q_t denote the exogenous process $\log(P_t Y_t / \bar{Y})$, and averaging (2.6) over i , we obtain

$$p_t = \xi q_{t|t} + (1 - \xi) p_{t|t}, \quad (2.7)$$

introducing the notation $x_{t+j|t} \equiv \int x_{t+j|t}(i) di$. The (log) price level is then a weighted average of the average estimate of current (log) nominal GDP (the exogenous forcing process) and the average estimate of the (log) price level itself.

Iterating (2.7) allows us to express p_t as a weighted average of the average estimate of q_t ,

⁶My point here is essentially the same as that of Christiano *et al.* (1998), who argue that it is possible to test the predictions of their model by computing the predicted responses to a given money-growth process, even if they do not believe (and do not assume, in their VAR strategy for identifying the effects of monetary policy shocks) that monetary policy is correctly described by an exogenous process for money growth. Of course, if one wanted to ask a question such as what the effect would be of an improvement in suppliers' information, it would be necessary to take a stand on whether or not the nominal GDP process should change. This would depend on how aggregate nominal expenditure is determined.

the average estimate of that average estimate, and so on. Introducing the notation

$$\begin{aligned} x_t^{(k)} &\equiv x_{t|t}^{(k-1)} && \text{for each } k \geq 1 \\ x_t^{(0)} &\equiv x_t \end{aligned}$$

for higher-order average expectations, we obtain

$$p_t = \sum_{k=1}^{\infty} \xi(1-\xi)^{k-1} q_t^{(k)}. \quad (2.8)$$

Thus the (log) price level can be expressed as a weighted average of expectations and higher-order expectations of the current level of (log) nominal GDP, as in Phelps (1983). Since $y_t = q_t - p_t$, it follows that

$$y_t = \sum_{k=1}^{\infty} \xi(1-\xi)^{k-1} [q_t - q_t^{(k)}]. \quad (2.9)$$

Thus output deviates from the natural rate only insofar as the level of current nominal GDP is not common knowledge. But this equation differs from (1.1), the implication of the Lucas model, in that higher-order expectations matter, and not simply the average estimate of current nominal GDP.

2.2 Equilibrium Inflation Dynamics

To consider a specific example, suppose that the growth rate of nominal GDP follows a first-order autoregressive process,

$$\Delta q_t = (1 - \rho)g + \rho\Delta q_{t-1} + u_t, \quad (2.10)$$

where Δ is the first-difference operator, $0 \leq \rho < 1$, and u_t is a zero-mean Gaussian white noise process. Here g represents the long-run average rate of growth of nominal GDP, while the parameter ρ indexes the degree of serial correlation in nominal GDP growth; in the special case that $\rho = 0$, nominal GDP follows a random walk with drift g . The disturbance u_t is assumed to represent a monetary policy shock, which therefore has no effect upon the real determinants of supply costs discussed above.

In the case of full information, the state of the economy at date t would be fully described by the vector

$$X_t \equiv \begin{bmatrix} q_t \\ q_{t-1} \end{bmatrix}.$$

That is to say, knowledge of the current value of X_t would suffice to compute not only the equilibrium values of p_t and y_t , but the conditional expectations of their values in all future periods as well. In terms of this vector, the law of motion (2.10) can equivalently be written

$$X_t = c + AX_{t-1} + au_t, \quad (2.11)$$

where

$$c \equiv \begin{bmatrix} (1-\rho)g \\ 0 \end{bmatrix}, \quad A \equiv \begin{bmatrix} 1+\rho & -\rho \\ 1 & 0 \end{bmatrix}, \quad a \equiv \begin{bmatrix} 1 \\ 0 \end{bmatrix}.$$

With incomplete information, however, average expectations and higher-order average expectations $X_t^{(k)}$ will also matter for the determination of prices and output and of their future evolution.

Suppose that the only information received by supplier i in period t is the noisy signal

$$z_t(i) = q_t + v_t(i), \quad (2.12)$$

where $v_t(i)$ is a mean-zero Gaussian white noise error term, distributed independently both of the history of fundamental disturbances $\{u_{t-j}\}$ and of the observation errors of all other suppliers. I shall suppose that the complete information set of supplier i when setting p_t^i consists of the history of the subjective observations $\{z_t(i)\}$; this means, in particular, that the person making the pricing decision does not actually observe (or does not pay attention to!) the quantity sold at that price.

Suppose, however, that the supplier forms optimal estimates of the aggregate state variables given this imperfect information. Specifically, I shall assume that the supplier forms minimum-mean-squared-error estimates that are updated in real time using a Kalman filter.⁷ Let us suppose that the supplier (correctly) believes that the economy's aggregate state

⁷This is optimal if the supplier seeks to maximize a log-quadratic approximation to his or her exact objective function; however, the exact objective function implied by the model above would not be log-quadratic.

evolves according to a law of motion

$$\bar{X}_t = \bar{c} + M\bar{X}_{t-1} + mu_t, \quad (2.13)$$

for a certain matrix M and vectors \bar{c} and m that we have yet to specify, where

$$\bar{X}_t \equiv \begin{bmatrix} X_t \\ F_t \end{bmatrix}$$

and

$$F_t \equiv \sum_{k=1}^{\infty} \xi(1-\xi)^{k-1} X_t^{(k)}. \quad (2.14)$$

Thus our conjecture is that only a particular linear combination of the higher-order expectations $X_t^{(k)}$ is needed in order to forecast the future evolution of that vector itself. Our interest in forecasting the evolution of this particular linear combination stems from the fact that (2.8) implies that p_t is equal to the first element of F_t . In terms of our extended state vector, we can write

$$p_t = e_3' \bar{X}_t, \quad (2.15)$$

introducing the notation e_j to refer to the j th unit vector (*i.e.*, a vector the j th element of which is one, while all other elements are zeros).

In terms of this extended state vector, the observation equation (2.12) is of the form

$$z_t(i) = e_1' \bar{X}_t + v_t(i). \quad (2.16)$$

It then follows (see, *e.g.*, Chow, 1975; Harvey, 1989) that i 's optimal estimate of the state vector evolves according to a Kalman filter equation

$$\bar{X}_{t|t}(i) = \bar{X}_{t|t-1}(i) + k[z_t(i) - e_1' \bar{X}_{t|t-1}], \quad (2.17)$$

where k is the vector of *Kalman gains* (to be specified), and the forecast prior to the period t observation is given by

$$\bar{X}_{t|t-1}(i) = \bar{c} + M\bar{X}_{t-1|t-1}(i). \quad (2.18)$$

Substituting (2.18) into (2.17), we obtain a law of motion for i 's estimate of the current state vector. Integrating this over i (and using (2.16) to observe that the average signal is

just $q_t = e'_1 \bar{X}_t$), we obtain a law of motion for the *average estimate* of the current state vector,

$$\begin{aligned}\bar{X}_{t|t} &= \bar{X}_{t|t-1} + ke'_1[\bar{X}_t - \bar{X}_{t|t-1}] \\ &= \bar{c} + ke'_1 M \bar{X}_{t-1} + (I - ke'_1) M \bar{X}_{t-1|t-1} + ke'_1 m u_t.\end{aligned}$$

Next we observe that (2.14) implies that

$$F_t = \bar{\xi} \bar{X}_{t|t}, \quad (2.19)$$

where

$$\bar{\xi} \equiv \begin{bmatrix} \xi & 0 & 1 - \xi & 0 \\ 0 & \xi & 0 & 1 - \xi \end{bmatrix}.$$

Substituting the above expression for $\bar{X}_{t|t}$, we obtain

$$F_t = \bar{\xi} \bar{c} + \hat{k} e'_1 M \bar{X}_{t-1} + (\bar{\xi} - \hat{k} e'_1) M \bar{X}_{t-1|t-1} + \hat{k} e'_1 m u_t, \quad (2.20)$$

where $\hat{k} \equiv \bar{\xi} k$.

We wish now to determine whether the laws of motion (2.11) and (2.20) for the elements of \bar{X}_t can in fact be expressed in the form (2.13), as conjectured. We note first that (2.11) implies that the matrices and vectors in (2.13) must be of the form

$$\bar{c} = \begin{bmatrix} c \\ d \end{bmatrix}, \quad M = \begin{bmatrix} A & 0 \\ G & H \end{bmatrix}, \quad m = \begin{bmatrix} a \\ h \end{bmatrix},$$

where c , A and a are defined as in (2.11), and the vectors d and h and the matrices G and H are yet to be determined.

Making these substitutions in (2.20), we then obtain

$$F_t = \hat{c} + \hat{k} A_1 X_{t-1} + [\xi A + (1 - \xi)G - \hat{k} A_1] X_{t-1|t-1} + (1 - \xi) H F_{t-1|t-1} + \hat{k} u_t, \quad (2.21)$$

where

$$\hat{c} \equiv \xi c + (1 - \xi) d, \quad (2.22)$$

and A_1 is the first row of A , i.e., the row vector $[1 + \rho \quad -\rho]$. Finally, we note that (2.19) for date $t - 1$ implies that

$$(1 - \xi) F_{t-1|t-1} = F_{t-1} - \xi X_{t-1|t-1}.$$

Using this substitution to eliminate $F_{t-1|t-1}$ from (2.21), we finally obtain

$$F_t = \hat{c} + \hat{k}A_1X_{t-1} + HF_{t-1} + [\xi A + (1 - \xi)G - \xi H - \hat{k}A_1]X_{t-1|t-1} + \hat{k}u_t. \quad (2.23)$$

This has the same form as the lower two rows of (2.13) if it happens that the expression in square brackets is a zero matrix.

In this case, we are able to make the identifications

$$d = \hat{c}, \quad (2.24)$$

$$G = \hat{k}A_1, \quad (2.25)$$

$$h = \hat{k}. \quad (2.26)$$

Given (2.22), (2.24) requires that $d = c$, and (2.25) and (2.26) uniquely identify G and h once we know the value of the gain vector \hat{k} . Using solution (2.25) for G , we observe that the expression in square brackets in (2.23) is a zero matrix if and only if

$$H = A - \hat{k}A_1. \quad (2.27)$$

Thus we have a unique solution for H as well. It follows that once we determine the vector of Kalman gains k , and hence the reduced vector \hat{k} , we can uniquely identify the coefficients of the law of motion (2.13) for the state vector \bar{X}_t . This then allows us to determine the equilibrium dynamics of p_t and y_t , using (2.15) and the identity $y_t = q_t - p_t$.

2.3 Optimal Filtering

It remains to determine the vector of Kalman gains k in the Kalman filter equation (2.17) for the optimal updating of individual suppliers' estimates of the aggregate state vector. Let us define the variance-covariance matrices of *forecast errors* on the part of individual suppliers:

$$\Sigma \equiv \text{var}\{\bar{X}_t - \bar{X}_{t|t-1}(i)\},$$

$$V \equiv \text{var}\{\bar{X}_t - \bar{X}_{t|t}(i)\},$$

Note that these matrices will be the same for all suppliers i , since the observation errors are assumed to have the same stochastic properties for each of them.

The Kalman gains are then as usual given by⁸

$$k = (\sigma_z^2)^{-1} \Sigma e_1, \quad (2.28)$$

where

$$\sigma_z^2 \equiv \text{var}\{z_t(i) - z_{t|t-1}(i)\} = e_1' \Sigma e_1 + \sigma_v^2. \quad (2.29)$$

Here $\sigma_v^2 > 0$ is the variance of the individual observation error $v_t(i)$ each period. Relations (2.28) – (2.29) then imply that

$$\hat{k} = (e_1' \Sigma e_1 + \sigma_v^2)^{-1} \bar{\xi} \Sigma e_1. \quad (2.30)$$

Thus once we have determined the matrix Σ , \hat{k} is given by (2.30), which allows us to solve for the coefficients of the law of motion (2.13) as above.

The computation of the variance-covariance matrix of forecast errors also follows standard lines. The transition equation (2.13) and the observation equation (2.16) imply that the matrices Σ and V satisfy

$$\begin{aligned} \Sigma &= M V M' + \sigma_u^2 m m', \\ V &= \Sigma - (\sigma_z^2)^{-1} \Sigma e_1 e_1' \Sigma, \end{aligned}$$

where σ_u^2 is the variance of the innovation term u_t in the exogenous process (2.10). Combining these equations, we obtain the usual stationary *Riccati equation* for Σ :

$$\Sigma = M \Sigma M' - (e_1' \Sigma e_1 + \sigma_v^2)^{-1} M \Sigma e_1 e_1' \Sigma M' + \sigma_u^2 m m'. \quad (2.31)$$

The matrix Σ is thus obtained by solving for a fixed point of the nonlinear matrix equation (2.31). Of course, this equation itself depends upon the elements of M and m , and hence upon the elements of G, H , and h , in addition to parameters of the model. These latter coefficients can in turn be determined as functions of Σ using (2.25) – (2.27) and (2.30).

⁸Add refs!!!

Thus we obtain a larger fixed-point equation to solve for Σ , specified solely in terms of model parameters.

Except in the special case discussed below, this system is too complicated to allow us to obtain further analytical results. Numerical solution for Σ in the case of given parameter values remains possible, however, and in practice proves not to be difficult.

3 The Size and Persistence of the Real Effects of Nominal Disturbances

We now turn to the insights that can be obtained regarding the effects of nominal disturbances from the solution of the example described in the previous section. In particular, we shall consider the impulse responses of output and inflation in response to an innovation u_t implied by the law of motion (2.13), and how these vary with the model parameters ρ, ξ , and σ_v^2/σ_u^2 .⁹

One question of considerable interest concerns the extent to which an unexpected increase in nominal GDP growth affects real activity, as opposed to simply raising the money prices paid for goods. But of no less interest is the question of the length of time for which any real effect *persists* following the shock. This is an especially important question given that the inability to explain persistent output effects of monetary policy shocks was one of the more notable of the perceived weaknesses of the first generation of asymmetric-information models.

3.1 The Case of a Random Walk in Nominal Spending

In considering the question of persistence, a useful benchmark is to consider the predicted response to an unexpected permanent increase in the level of nominal GDP. In this case, the subsequent dynamics of prices and output are due solely to the adjustment over time of

⁹It should be evident that it is only the *relative* size of the innovation variances that matters for the determination of the Kalman gains k , and hence of the coefficients M and m in the law of motion. It is also only the relative variance that is determined by a particular assumed rate of information flow in the “noisy channel” through which a supplier monitors current aggregate demand. See Sims (2001) for details of the computation of the rate of information flow.

a discrepancy that has arisen between the level of nominal spending and the existing level of prices, and not to any predictable further changes in the level of nominal spending itself. This corresponds to the computation of impulse response functions in a special case of the model of the previous section, the case in which $\rho = 0$, so that the log of nominal GDP follows a random walk with drift.

In this special case, the equations of the previous section can be further simplified. First, we note that in this case, the state vector X_t may be reduced to the single element q_t . The law of motion (2.11) continues to apply that now $c = g$, $A = 1$, and $a = 1$ are all scalars. The law of motion for the aggregate state can again be written in the form (2.13), where F_t is defined as in (2.14); but now F_t is a scalar, and the blocks G, H and h of M and m are each scalars as well. Equation (2.19) continues to apply, but now with the definition

$$\bar{\xi} \equiv [\xi \ 1 - \xi].$$

Equation (2.26) holds as before, but now \hat{k} is a scalar; equations (2.25) and (2.27) reduce to

$$G = \hat{k},$$

$$H = 1 - \hat{k}.$$

Substituting these solutions for the elements of $M(\hat{k})$ and $m(\hat{k})$, we can solve (2.31) for the matrix $\Sigma(\hat{k})$ in the case of any given reduced Kalman gain \hat{k} . The upper left equation in this system is given by

$$\Sigma_{11} = \Sigma_{11} - (\Sigma_{11} + \sigma_v^2)^{-1} \Sigma_{11}^2 + \sigma_u^2.$$

This equation involves only Σ_{11} , and is independent of \hat{k} . It reduces to a quadratic equation in Σ_{11} , which has two real roots, one positive and one negative. Since the variance Σ_{11} must be non-negative, the positive root is the only relevant solution. This is given by

$$\Sigma_{11} = \frac{\sigma_u^2}{2} \left\{ 1 + [1 + 4(\sigma_v^2/\sigma_u^2)]^{1/2} \right\}. \quad (3.1)$$

The lower left equation in the system (2.31), in turn, involves only Σ_{21} and Σ_{11} , and given that we have already solved for Σ_{11} , this equation can be solved for Σ_{21} . We obtain

$$\Sigma_{21}(\hat{k}) = \sigma_u^2 \frac{1 + 2(\sigma_v^2/\sigma_u^2) + [1 + 4(\sigma_v^2/\sigma_u^2)]^{1/2}}{(2/\hat{k}) - 1 + [1 + 4(\sigma_v^2/\sigma_u^2)]^{1/2}}. \quad (3.2)$$

Finally, (2.30) expresses \hat{k} as a function of Σ , which in fact depends only upon the elements Σ_{11} and Σ_{21} . Substituting expressions (3.1) – (3.2) into this relation, we obtain a quadratic equation for \hat{k} , namely

$$(\sigma_v^2/\sigma_u^2)\hat{k}^2 + \xi\hat{k} - \xi = 0. \quad (3.3)$$

It is easily seen that for any parameters $\xi, \sigma_u^2, \sigma_v^2 > 0$, equation (3.3) has two real roots, one satisfying

$$0 < \hat{k} < 1, \quad (3.4)$$

and another that is negative. Substituting our previous solutions for $M(\hat{k})$ and $m(\hat{k})$ into (2.13), we note that this law of motion implies that

$$q_t - F_t = (1 - \hat{k})(q_{t-1} - F_{t-1}) + (1 - \hat{k})u_t. \quad (3.5)$$

Law of motion (3.5) implies that $q_t - F_t$, which measures the discrepancy between the actual level of nominal spending and a certain average of higher-order expectations regarding current nominal spending, is a *stationary* random variable if and only if $|1 - \hat{k}| < 1$. This requires that $\hat{k} > 0$, and so excludes the negative root of (3.3). Thus if we are to obtain a solution in which the variances of forecast errors are finite and constant over time, as assumed above, it can correspond only to the root satisfying (3.4). This root is given by

$$\hat{k} = \frac{1}{2}\{-\gamma + [\gamma^2 + 4\gamma]^{1/2}\}, \quad (3.6)$$

where

$$\gamma \equiv \xi\sigma_u^2/\sigma_v^2 > 0. \quad (3.7)$$

3.2 Dynamics of Real Activity

Since in this special case, $p_t = F_t$, (3.5) immediately implies that (log) real GDP y_t evolves according to

$$y_t = \nu(y_{t-1} + u_t), \quad (3.8)$$

where $\nu = 1 - \hat{k}$ and \hat{k} is given by (3.6). Since $0 < \nu < 1$, this describes a stationary process with positive serial correlation. The implied effect of a monetary shock at date t upon current and expected subsequent real activity is given by

$$E_t(y_{t+j}) - E_{t-1}(y_{t+j}) = \nu^{j+1}u_t,$$

which holds for all $j \geq 0$. Thus the same coefficient ν determines both the size of the initial impact upon real activity of a monetary shock (y_t is increased by νu_t), and the degree of persistence of such an effect (the effect on output j periods later decays as ν^j).

While the model implies that the real effects of a monetary shock die out with time, output is *not* predicted to again equal the natural rate on average in any finite time, as in the Lucas model. Indeed, the degree of persistence of such real effects may be arbitrarily great. For (3.6) implies that \hat{k} may be an arbitrarily small positive quantity (so that ν is arbitrarily close to 1), if γ is small enough; and the half-life of output disturbances tends to infinity as ν approaches one.

More generally, the degree of persistence is observed to be a monotonically decreasing function of γ , which depends both upon ξ and upon σ_v^2/σ_u^2 . Not surprisingly, this implies that persistence is greater the larger is σ_v^2 relative to σ_u^2 ; that is, the less the information contained in the individual suppliers' subjective perceptions of the state of nominal GDP. And if this information is small enough, persistence may be arbitrarily great. This may seem little different from the conclusion in the case of the Lucas model that the output effects of a monetary disturbance may persist for a substantial time if it takes a long time for changes in the money supply to become public information. But because the bottleneck in our case is assumed to be the inaccuracy of individual subjective perceptions, rather than limitations of the statistics that are publicly available should people bother to pay attention, the mere fact

that monetary data quickly enter the public domain does not in itself imply that perceptions of the state of aggregate demand must be accurate.

Even more interestingly, persistence is predicted to be greater the smaller is ξ , which is to say, the greater the extent of “real rigidity” in the sense of Ball and Romer (1990), and hence the greater the degree of strategic complementarity in individual suppliers’ pricing decisions.¹⁰ In fact, the model implies that *regardless* of the degree of accuracy of the suppliers’ observations of the aggregate state — as long as they are not *perfect* — the degree of persistence of the real effects of a monetary policy shock can be *arbitrarily great*, if the degree of “real rigidity” is sufficiently great (*i.e.*, ξ is sufficiently small)!

This means that substantial real effects of monetary policy, and significant persistence of such effects, do not depend upon private parties being wholly ignorant of the occurrence of the disturbance to monetary policy. If σ_v^2/σ_u^2 is not too large, each individual supplier will have a fairly accurate estimate of current aggregate demand at the time of setting its price, and individual estimates $q_{t|t}(i)$ will quickly adjust by nearly as much as the permanent change in nominal spending that has occurred. Nonetheless, *prices* may be quite slow to adjust, owing to continuing uncertainty about *others’* estimates of current aggregate demand, and even greater uncertainty about others’ estimates of others’ estimates. Thus the sluggishness of higher-order expectations stressed by Phelps (1983) can play a critical role in explaining both the size and persistence of the real effects of monetary policy.

3.3 Dynamics of Higher-Order Expectations

This can be shown explicitly through an analysis of the impulse responses of higher-order average expectations following a monetary shock. While we have seen above that it is not necessary to solve for the complete hierarchy of expectations in order to solve for the equilibrium dynamics of output (only the particular average of higher-order expectations represented by F_t), consideration of the dynamics of expectations at different levels can

¹⁰See Woodford (2001) for further discussion of the interpretation of this parameter and various factors that can make it small in an actual economy.

provide further insight into the reason for the sluggishness of price adjustment in this model.

Similar Kalman-filtering techniques as in the previous section can be used to determine the dynamics of average expectations at each level of the hierarchy. Let $q_t^{(k)}$ denote the average k -th order expectation at date t regarding the current level of (log) nominal GDP, where $q_t^{(0)}$ is defined as q_t , and let us conjecture a law of motion of the form

$$q_t^{(k)} = \sum_{j=0}^k \alpha_{kj} q_{t-1}^{(j)} + a_k u_t \quad (3.9)$$

for each $k \geq 0$, where for $k = 0$ we have $\alpha_{00} = 1$ and $a_0 = 1$. We wish to determine the coefficients α_{kj} and a_k for higher values of k .

Supplier i 's estimate of the value of $q_t^{(k)}$ should evolve according to a Kalman filter equation of the form

$$q_{t|t}^{(k)}(i) = q_{t|t-1}^{(k)}(i) + \kappa_{k+1}(z_t(i) - z_{t|t-1}(i)),$$

where the $k + 1$ st order Kalman gain κ_{k+1} remains to be determined. Substituting the observation equation (2.12) for $z_t(i)$ and its forecast as before, and averaging over i , we obtain

$$q_t^{(k+1)} = q_{t|t-1}^{(k)} + \kappa_{k+1}(q_t^{(0)} - q_{t|t-1}^{(0)}).$$

Then substituting the average forecasts at date $t - 1$ implied by the assumed law of motion (3.9), and the law of motion itself for $q_t^{(0)}$, we obtain

$$q_t^{(k+1)} = \sum_{j=0}^k \alpha_{kj} q_{t-1}^{(j+1)} + \kappa_{k+1}(q_{t-1}^{(0)} - q_{t-1}^{(1)} + u_t).$$

This yields a law of motion for the next higher order of expectations of the desired form (3.9).

Identifying the coefficients $\alpha_{k+1,j}$ and a_{k+1} with the ones appearing in this last relation, we obtain equations that can be used to solve recursively for these coefficients at each order of expectations. For each $k \geq 1$, we find that

$$\begin{aligned} \alpha_{k0} &= \kappa_k, \\ \alpha_{kj} &= \kappa_{k-j} - \kappa_{k+1-j} \quad \text{for each } ; \quad 0 < j < k, \end{aligned}$$

$$\begin{aligned}\alpha_{kk} &= 1 - \kappa_1, \\ a_k &= \kappa_k.\end{aligned}$$

Thus once we determine the sequence of Kalman gains κ_k , we know the complete law of motion (3.9) for all orders of expectations.

The Kalman gains can also be determined using methods like those employed above. Letting

$$\sigma_{k0} \equiv \text{cov}\{q_t^{(k)} - q_{t|t-1}^{(k)}(i), q_t - q_{t|t-1}(i)\},$$

then the usual reasoning implies that the Kalman gains are given by

$$\kappa_{k+1} = (\sigma_z^2)^{-1} \sigma_{k0} \quad (3.10)$$

for each $k \geq 0$.¹¹ These covariances in turn satisfy a Riccati equation,

$$\sigma_{k0} = (1 - (\sigma_z^2)^{-1} \sigma_{00}) \sum_{j=0}^k \alpha_{kj} \sigma_{j0} + a_k \sigma_u^2. \quad (3.11)$$

for each $k \geq 0$. Note that once we know the value of σ_{00} , this is a linear equation in the other covariances; and we have already solved for $\sigma_{00} = \Sigma_{11}$ in (3.1).

Substituting the above solution for the α_{kj} and a_k coefficients as functions of the Kalman gains, and using (3.10) to replace each covariance σ_{k0} by a multiple of κ_{k+1} , it is possible to rewrite (3.11) in terms of the Kalman gains alone. We obtain the relation

$$\kappa_{k+1} = \frac{1 - \kappa_1}{1 - (1 - \kappa_1)^2} \left\{ \sum_{j=1}^k \kappa_j \kappa_{k+1-j} - \sum_{j=2}^k \kappa_j \kappa_{k+2-j} + \kappa_k \frac{\sigma_u^2}{\sigma_v^2} \right\} \quad (3.12)$$

for each $k \geq 1$. This relation allows us to solve recursively for each of the κ_k , starting from the initial value

$$\kappa_1 = \frac{-1 + [1 + 4(\sigma_v^2/\sigma_u^2)]^{1/2}}{2\sigma_v^2/\sigma_u^2}$$

implied by (3.10) using (3.1) for σ_{00} .

Figure 3 gives a numerical illustration of the implied dynamics of higher-order expectations in response to an immediate, permanent unit increase in nominal spending. The figure

¹¹Note that in the case $k = 0$, this equation is equivalent to the first row of (2.28.)

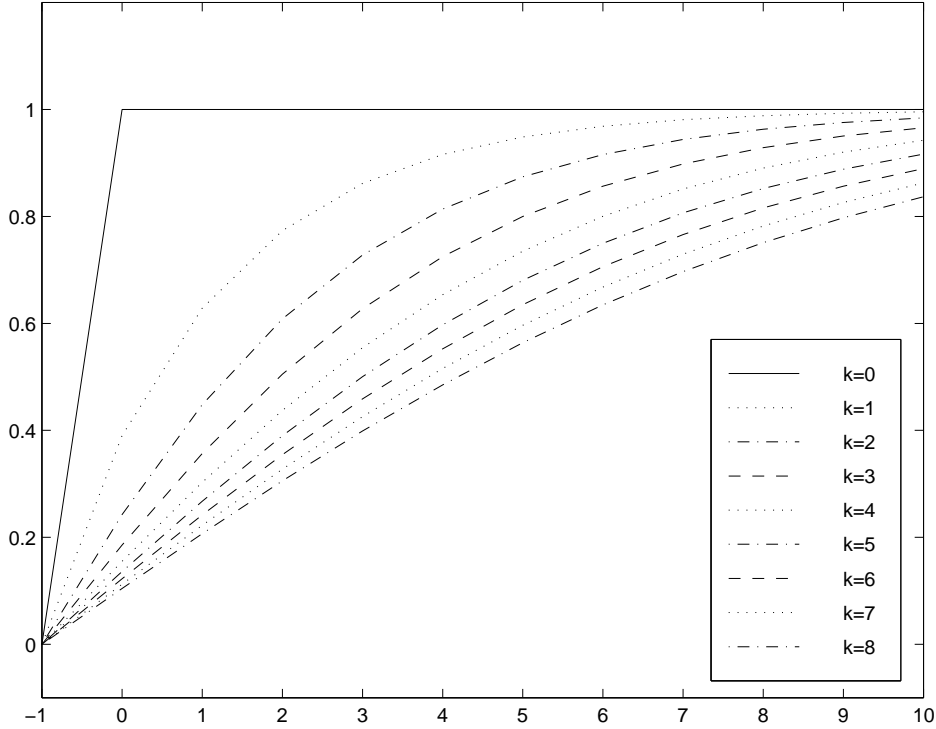


Figure 3: Impulse response functions for higher-order expectations $q_t^{(k)}$, for various values of k . The case $k = 0$ indicates the exogenous disturbance to log nominal GDP itself.

shows the impulse responses of $q_t^{(0)}$ (nominal GDP itself), $q_t^{(1)}$ (the average estimate of current nominal GDP), $q_t^{(2)}$ (the average estimate of the average estimate), and so on, up through the eighth-order expectation $q_t^{(8)}$, in the case of a relative innovation variance $\sigma_v^2/\sigma_u^2 = 4$. One observes that even with this degree of noise in subjective estimates of current nominal spending, the average estimate of current nominal GDP adjusts fairly rapidly following the disturbance: forty percent of the eventual adjustment occurs in the period of the increase in nominal GDP itself, and eighty percent has occurred within two periods later. Higher-order expectations instead adjust much more sluggishly. Eighth-order expectations adjust only a fifth as much as do first-order expectations during the period of the disturbance; even three periods later, they have not yet adjusted by as much as first-order expectations do in the period of the disturbance, and it is only nine periods after the disturbance that they have adjusted by eighty percent of the size of the disturbance.

The extent to which these different orders of average expectations matter for pricing depends, of course, on the degree of strategic complementarity between the pricing decisions of different suppliers. If ξ is near one, then the average price level will adjust at the rate that the average estimate $q_t^{(1)}$ does, and the real effects of the disturbance will be modest after the period of the shock, and the next period or so.¹² On the other hand, if ξ is small, so that strategic complementarity is great, the sluggishness of higher-order expectations can matter a great deal. Woodford (2001) suggests that $\xi = .15$ is an empirically plausible value for the U.S. In this case, the impulse response of the average price level would be a weighted average of those shown in Figure 1 (and the responses of still higher-order expectations, not shown), with a weight of only .15 on the response of first-order expectations. More than half the weight is put on expectations of order $k > 4$, and more than a quarter of the weight is put on expectations of order $k > 8$, *i.e.*, expectations that adjust more slowly than any that are shown in the figure. Thus the insight of Phelps (1983), that the dependence of aggregate outcomes upon higher-order expectations can be an important source of inertia in the response of prices to nominal disturbances, is born out.

4 Comparison with a Model of Sticky Prices

It may be worth briefly considering the extent to which the predictions of such a model resemble, and differ from, those of a model in which prices do not immediately adjust to nominal disturbances, not because price-setters are unaware of the adjustment that would best serve their interests at any of the times at which they actually consider changing their prices, but simply because they do not continuously reconsider their prices. This familiar hypothesis of “sticky prices” is clearly not entirely unrelated to the hypothesis of incomplete information. In particular, insofar as suppliers behave in the way assumed in models with sticky prices, they surely do so *not* primarily in order to economize on the cost of price changes themselves — literal “menu costs” are in most cases quite small — but rather in order to

¹²If ξ exceeds 1, as is theoretically possible (Woodford, 2001), then prices will adjust even more rapidly than does the average expectation of current nominal GDP.

economize on the cost of *having to make more frequent decisions* about whether their current prices are significantly out of line or not.¹³ And there is obviously a close relation between the hypothesis that there are substantial costs associated with constant close monitoring of current conditions (the hypothesis explored in this paper) and the hypothesis that there are substantial costs associated with constant reconsideration of how close one’s current prices are to those that are optimal under current conditions.

For this reason, it is interesting to ask how similar or different the implications of the hypothesis of incomplete common knowledge for aggregate dynamics are to those of a model with sticky prices. Here I show that the dynamics of aggregate output and the aggregate price index derived above in the case of a random walk in nominal GDP are indistinguishable from those predicted by a standard sticky price model, namely, a discrete-time version of the model proposed by Calvo (1983). Thus it need not be possible to distinguish among these models empirically, using aggregate data alone. Nonetheless, this does not mean that the models make identical predictions regardless of the nature of monetary policy, as consideration of a more general policy specification will show.

4.1 Dynamics of Real Activity under the Calvo Pricing Model

In the well-known Calvo (1983) model of staggered pricing, the price charged by each supplier is reconsidered only at random intervals of time, with the probability that any given price will be reconsidered within a particular time interval being independent of which price it is, how long ago it was last reconsidered, and the level of the current price (relative either to other prices or to other aspects of current market conditions). In this case (and proceeding directly to a log-linear approximation to the optimal pricing condition), (2.6) becomes instead

$$p_t(i) = (1 - \alpha\beta) \sum_{j=0}^{\infty} (\alpha\beta)^j E_t[p_{t+j} + \xi y_{t+j}] \quad (4.1)$$

¹³Zbaracki *et al.* (1999) document this in the case of a single industrial firm whose operations they study in detail. They find that the firm’s “managerial costs” of price adjustment are many times larger than the physical costs of price changes.

for any supplier i that reconsiders its price in period t , where $0 < \alpha < 1$ is the probability that any given price is *not* reconsidered during any given period, and $0 < \beta < 1$ is again the discount factor in (2.3). This says that the price chosen is a weighted average of the prices that would be optimal at the various dates and in the various states of the world in which the price chosen at date t has not yet been revised. Because we now assume full information, subjective expectations at date t are now replaced by an expectation conditional upon the history of disturbances through that date. If instead i does not reconsider its price in period t , then we have simply $p_t(i) = p_{t-1}(i)$.

This model of pricing results (see, *e.g.*, Woodford, 2001) in an aggregate supply relation of the form

$$\Delta p_t = \kappa y_t + \beta E_t \Delta p_{t+1}, \quad (4.2)$$

where

$$\kappa = \frac{(1 - \alpha)(1 - \alpha\beta)}{\alpha} \xi > 0. \quad (4.3)$$

This relation is sometimes called the “New Keynesian Phillips Curve.” Note that it holds regardless of the assumed evolution of nominal spending. Let us first consider the case of a random walk with drift in nominal GDP, as in section 3.

The rational expectations equilibrium associated with such a policy is then a pair of stochastic processes for the price level and real GDP that are consistent with both (4.2) and

$$\Delta p_t + \Delta y_t = g + u_t. \quad (4.4)$$

The unique solution in which inflation and output fluctuations are stationary is given by

$$y_t = \nu(y_{t-1} + u_t),$$

$$\Delta p_t = g + (1 - \nu)(u_t + y_{t-1}),$$

where $0 < \nu < 1$ is given by

$$\nu = \frac{1 + \beta + \kappa - [(1 + \beta + \kappa)^2 - 4\beta]^{1/2}}{2\beta}. \quad (4.5)$$

We observe that output fluctuations again follow a law of motion of the form (3.8), except that now the autoregressive coefficient ν depends upon the frequency of price adjustment among other parameters. Thus the impulse responses of both prices and real activity in response to a monetary disturbance are of the same form as in the noisy-information model. In fact, for given values of ξ and β , to any value of the variance ratio σ_v^2/σ_u^2 (or rate of information flow in the model with noisy information) there corresponds a particular value of α (or degree of price stickiness) that results in *identical* dynamics of prices and output. Thus in the case that nominal GDP evolves according to (4.4), and we treat both α and the variance ratio as free parameters (to be estimated from the dynamics of aggregate output and the aggregate price index), the predictions of the two models are *observationally equivalent*.

In the case that β is near one (a plausible assumption), we can go further, and obtain an equivalence between a particular value of the variance ratio and a particular value of α that holds *regardless of the value of ξ* . When we set β equal to one, (4.5) reduces to exactly the same expression for ν as in the noisy-information model (one minus the right-hand side of (2.30)), except that γ is equal to κ . Comparing expression (4.3) for κ (and setting $\beta = 1$) with expression (3.7) for γ , we see that the value of α required for the sticky-price model to imply the same dynamics as the noisy-information model is the one such that

$$\frac{\alpha}{(1-\alpha)^2} = \frac{\sigma_v^2}{\sigma_u^2}. \quad (4.6)$$

In this limiting case, the required value of α is independent of the value of ξ . This means that even if the structure of the economy were to shift in a way that changed the value of ξ , the predictions of the two models would *continue* to be identical.

4.2 Consequences of Persistence in the Growth of Nominal Spending

However, it would be a mistake to conclude more generally that the noisy-information model is observationally equivalent to the Calvo model of staggered pricing. The models cease to predict the same dynamics of output and inflation if nominal GDP does not follow a random walk with drift. This can be seen by considering the more general stochastic process for

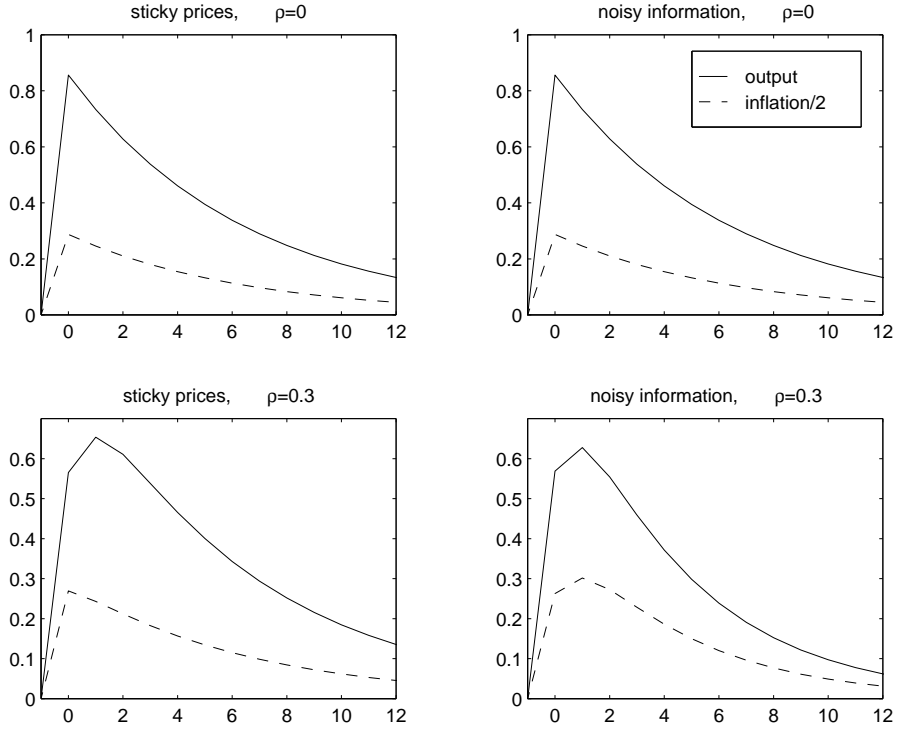


Figure 4: Comparison of impulse response functions predicted by the two models, for the cases $\rho = 0$ and $.3$.

nominal GDP (2.10) considered earlier, in the case that $\rho > 0$, so that the growth rate of nominal GDP exhibits serial correlation.¹⁴ In this case, we are unable to obtain an analytical solution to the nonlinear equation system (2.31), and so must resort to numerical solution for particular assumed parameter values.

Figures 4 and 5 plot the impulse responses of output and inflation¹⁵ to an innovation in nominal GDP growth at date zero, that eventually raises (log) nominal GDP by a unit amount. (The innovation at date zero is thus of size $u_0 = 1 - \rho$.) The two rows of Figure 4 consider nominal spending processes characterized by $\rho = 0$ and $\rho = .3$ respectively, while

¹⁴It is important to note that this is the case of practical interest, given that variations in nominal GDP growth do exhibit considerable persistence. More to the point, VAR estimates of the effects of monetary policy shocks indicate an effect on nominal GDP that takes many quarters to reach its eventual magnitude, rather than an immediate permanent increase of the kind implied by the random-walk specification.

¹⁵In these figures, “inflation” is defined as $4\Delta p_t$, corresponding to an annualized inflation rate if the model “periods” are interpreted as quarters.

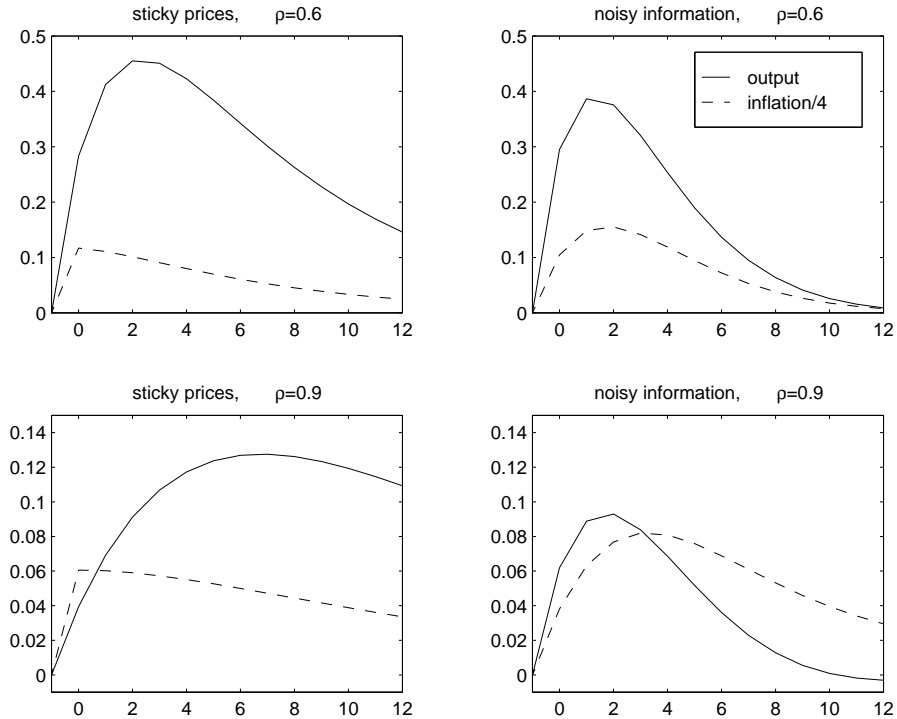


Figure 5: The comparison extended to the cases $\rho = .6$ and $.9$.

the two rows of Figure 5 consider the further cases $\rho = .6$ and $\rho = .9$. The two columns of both figures compare the predictions of two models for each case, the model with Calvo pricing (the left column) and the model with noisy information (the right column).

In each case, the value of ξ is fixed at $.15$, a value that is argued to be realistic for the U.S. economy in Woodford (2001). The sticky-price model is further calibrated by assuming $\beta = .99$, a plausible discount factor if the periods are interpreted as quarters, and $\alpha = 2/3$, so that one-third of all prices are revised each quarter. This implies an average interval between price changes of 9 months, consistent with the survey evidence of Blinder *et al.* (1998, Table 4.1). The noisy-information model is then calibrated by assuming that $\sigma_v^2/\sigma_u^2 = 6.23$, the value required in order for the predicted inflation and output dynamics of the two models to be identical in the case that $\rho = 0$.¹⁶

¹⁶This value differs slightly from the variance ratio of 6 that would be indicated by (4.6), because β is not exactly equal to one.

Comparing the two columns, we observe that the predicted impulse responses are the same for both models when $\rho = 0$ (as we have shown above analytically), but that they become progressively more different the larger the value assigned to ρ . Thus the two models are not observationally equivalent in the case of an *arbitrary* monetary policy, and will not give the same answers to a question about the consequences of changing the way in which monetary policy is conducted.

Furthermore, the failure of the predictions to agree in the case of substantial persistence in nominal GDP growth is not one that can be remedied by adjusting the value of α in the sticky-price model. The impulse responses predicted by the noisy-information model when $\rho > 0$ are ones that are not consistent with the Calvo model for *any* parameter values. This is because relation (4.2) can be “solved forward” to yield

$$\Delta p_t = \kappa \sum_{j=0}^{\infty} \beta^j E_t y_{t+j}, \quad (4.7)$$

which implies that the predicted path of inflation is a function solely of expected *subsequent* output gaps. It follows that a monetary disturbance with a *delayed* positive effect on output must increase inflation *earlier*. It is thus not an artifact of the particular parameter values assumed in Figures 2 and 3 that the inflation response is observed to peak sooner than the output response when $\rho > 0$. The noisy-information model can instead generate responses in which inflation peaks *later*, as is especially evident in the case $\rho = .9$. Such a response is plainly inconsistent with (4.7).

Further insight into the difference in the predictions of the two models may be obtained from Figure 6, which plots the impulse response functions for the price level implied by the two models alongside the impulse response for nominal GDP. (The case shown corresponds to the case $\rho = .9$ in Figure 5.) A monetary disturbance results in a gradual increase in the log of nominal GDP, to an eventual level that is higher by one than its level before the shock. The sticky-price model predicts that the average log price of goods will not rise as much as the increase in nominal GDP, and so real output is temporarily increased. But still, by comparison with the noisy-information model, the sticky-price model predicts relatively

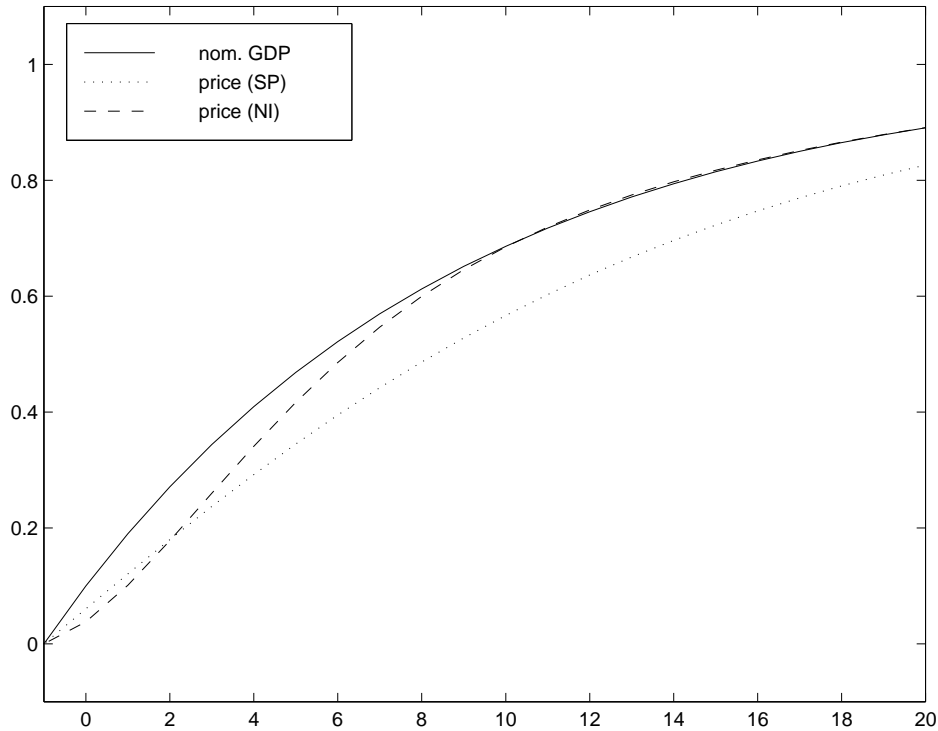


Figure 6: Impulse response function for the price level in the sticky-price model (SP) and the noisy-information model (NI), for the case $\rho = .9$.

strong price increases in the time immediately following the shock. The reason is that, under the assumption of full information, suppliers who revise their prices soon after the shock can already anticipate that further increases in nominal GDP are coming in the next few quarters. Then, because there is a substantial probability that the supplier’s price will not be revised again while those increases in aggregate demand, it is desirable to increase the price *immediately* in order to prevent it from falling too far behind its desired level before the next opportunity for revision arises.

In the noisy-information model, instead, there is no such need to “front-load” price increases in the case of a disturbance that is expected to result in persistent above-average growth in nominal spending. Suppliers who suspect that such a shock has occurred will increase prices some, but can *plan to increase prices more later* if their estimate of demand conditions has not changed in the meantime. In the absence of a need to “front-load,” initial

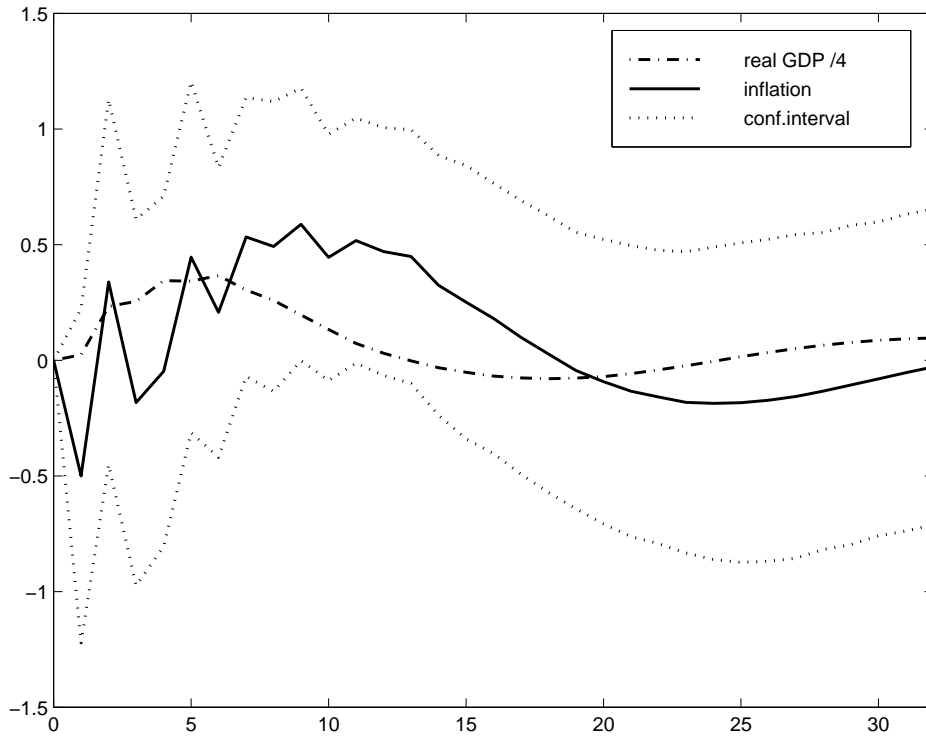


Figure 7: Estimated impulse responses of real GDP and inflation to an unexpected interest-rate reduction. Source: Christiano *et al.* (2001).

price increases are quite small, owing to uncertainty about whether others are expecting others ... to expect others to perceive the increase in demand. A few quarters later, instead, price increases are more rapid than in the sticky-price model. Once suppliers can become fairly confident that others expect ... others to have noticed the surge in spending, the fact that prices were not already increased earlier does not prevent them from being rapidly brought into line with the current volume of nominal spending. The result is a surge in inflation that occurs after the peak effect on output.

The Calvo pricing model has in fact come under extensive criticism for implying that the rate of inflation should be a purely “forward-looking” variable, and the relative timing of the output and inflation responses predicted by the noisy-information model are, at least qualitatively, more similar to those indicated by VAR estimates of the effects of monetary policy shocks. For the estimated responses generally indicate a stronger effect on inflation

in the quarters *after* the peak effect on output; see, for example, the responses in Figure 7, which are again taken from Christiano *et al.* (2001).¹⁷ The question of how well the precise quantitative predictions of the noisy-information model match empirical evidence of this kind is left for future work.¹⁸ But the model offers some promise of providing a more satisfactory explanation than a standard sticky-price model can.

4.3 Responses to Other Disturbances

The noisy-information model offers qualitatively different predictions from a sticky-price model in another respect as well. We have thus far only considered the predictions of the two models about the effects of a single kind of disturbance, a monetary policy disturbance that affects the path of nominal spending with no effect upon potential output (the constant \bar{Y} above). However, even when the two models predict identical effects of a disturbance of this kind, they need not predict identical effects for other types of disturbances as well.

In general, they will not, for a simple reason. In the sticky price model, the rate at which prices adjust following a disturbance depends on the rate at which various suppliers choose to reconsider their prices; but this rate, if taken as exogenous as in “time-dependent pricing” models like the Calvo model, will be the same *regardless of the type of disturbance* to which the economy must adjust. On the other hand, there is no reason why the rate of *flow of information* about different disturbances must be the same, in a noisy-information model. Some variables may be observed with more precision, and others with less; and as a result, prices may succeed better at bringing about an efficient response to some disturbances than to others.

A simple example can easily illustrate the way that this can result in predictions different

¹⁷This observation is related to what Mankiw and Reis (2001) call “the acceleration phenomenon,” though the evidence that they discuss relates to unconditional correlations between cyclical output and subsequent inflation acceleration, rather than to the co-movements of these variables that are associated with identified monetary policy shocks.

¹⁸We cannot address the question here, both because the estimated impulse response of nominal GDP shown in Figure 1 is plainly not consistent with the simple law of motion (2.10) for any value of ρ , and because our theoretical calculations have assumed that nominal GDP is affected only by monetary disturbances, while the identified VAR implies otherwise.

from those of a sticky-price model. Let us again assume a random walk in nominal GDP, the case in which the two models will (for appropriate parameter values) predict the same responses to a monetary disturbance u_t . But let us now generalize the above model, so that the log of the natural rate of output (\bar{y}_t) follows a random walk with drift — for example, as a result of a random walk with drift in a multiplicative technology factor¹⁹ — that is independent of the random walk in nominal GDP resulting from the actions of the central bank. We can write this process as

$$\bar{y}_t = \bar{g} + \bar{y}_{t-1} + \bar{u}_t,$$

where \bar{g} is the average rate of growth in the natural rate, and \bar{u}_t is a mean-zero i.i.d. disturbance, distributed independently of u_t .

The optimal price for any price-setter is still given by (2.6), if now y_t is interpreted as output *relative* to the time-varying natural rate. Similarly, optimal pricing policy in the sticky-price model continues to be described by (4.1), under the same reinterpretation. It follows that in the sticky-price model, the relation between inflation and the output gap continues to be described by (4.2). Equation (4.4) continues to hold as well, except that the right-hand side becomes

$$(g - \bar{g}) + u_t - \bar{u}_t.$$

Since the composite disturbance $u_t - \bar{u}_t$ is still completely unforecastable at any date prior to t , the stationary rational expectations equilibrium of the sticky-price model takes the same form as before, except with g replaced by $g - \bar{g}$ and u_t replaced by $u_t - \bar{u}_t$. In particular, the equilibrium output gap will evolve according to

$$y_t = \nu(y_{t-1} + u_t - \bar{u}_t). \tag{4.8}$$

The predictions of the noisy-information model will instead depend upon what we assume about the observability of the additional disturbance process \bar{y}_t . Suppose, for simplicity, that

¹⁹See Woodford (2001) for explicit analysis of how the natural rate of output is affected by technology shocks, and other real disturbances, in a model of monopolistic competition of the kind used here.

each supplier observes \bar{y}_t *precisely*, while still observing the state of aggregate demand only with noise. In this case, there is again only a single “hidden” state variable to estimate on the basis of the noisy observations. In fact, our previous calculations continue to apply, if we replace p_t throughout by \tilde{p}_t , the log of “natural nominal GDP” (*i.e.*, P_t times the natural rate of output). For \tilde{p}_t satisfies the identity $\tilde{p}_t + y_t = q_t$, given our reinterpretation of y_t ; and the perfect observability of the natural rate means that (2.6) may equivalently be written

$$\tilde{p}_t(i) = \tilde{p}_{t|t}(i) + \xi y_{t|t}(i).$$

With this reinterpretation of the price variable, our derivations go through as before. In particular, the equilibrium output gap will evolve according to

$$y_t = \nu(y_{t-1} + u_t).$$

For appropriately chosen parameter values, the coefficient ν here may take the same value as in (4.8). But even in that case, there remains an important difference in the predicted responses to the technology shock. In the sticky-price model, technology shocks produce deviations of output from potential that are exactly as long-lasting as those that result from monetary disturbances. Instead, in the noisy-information model (under our special assumption about the observability of \bar{y}_t), technology shocks have no effects upon the output gap (or upon \tilde{p}_t) at all. For while prices adjust only slowly to a change in demand conditions (owing to the assumed imperfect common knowledge regarding disturbances of this kind), they adjust immediately to a change in technology (as this is assumed to be common knowledge).

This differing prediction is not just another indication that the two models are not equivalent. It is again potentially of interest as an explanation for one of the more notable embarrassments for the sticky-price model. An extensive empirical literature dating back several decades²⁰ has found that prices respond more, and more rapidly, to increases in the marginal cost of supply resulting from increases in factor prices than to increases resulting from an increased scale of production as a result of increases in demand. Such a difference is not

²⁰See Bils and Chang (1999) for a review of this evidence, as well for further evidence for the same conclusion.

easily rationalized in terms of a standard sticky-price model. Some have argued that such evidence indicates that prices are set on the basis of considerations other than a constant desired markup over marginal supply cost. The noisy-information model suggests a different, and possibly simpler explanation. Prices are set in proportion to marginal cost, but it must be the supplier's subjective *estimate* of marginal cost; and if suppliers are better informed about certain disturbances that affect supply cost than about others, those disturbances will have a larger and more immediate effect on prices.

Of course, I have given no reason why one should assume that suppliers are better informed about variation in the natural rate of output than about variation in aggregate nominal spending. My point is simply that there is no reason why the logic of the noisy-information model should imply that the rate of information flow with regard to different shocks must be the same. Even if one supposes that, on grounds of theoretical parsimony, one should prefer to derive the degree of noise associated with the observation of various disturbances from a single underlying limitation on human information-processing capacity, one should not in general expect that the amount of scarce processing capacity allocated to monitoring different types of disturbances should be the same.

This possibility of explaining the differential responsiveness of prices to different types of disturbances is also an important advantage of the noisy-information model over the model recently proposed by Mankiw and Reis (2001), that is in some ways similar. Mankiw and Reis also argue for a pricing model in which each supplier's price at any given time is optimal conditional upon that supplier's information set, and in which price adjustment in response to a disturbance to aggregate demand is delayed owing to suppliers' not all having complete information about the disturbances that have already occurred. But rather than assuming continuous observation of demand conditions using a noisy channel, as is proposed here, Mankiw and Reis assume that suppliers obtain no new information at all except at random intervals. Yet on the occasions upon which a supplier updates its information, it acquires *complete* information about all disturbances that have occurred up until that time.²¹ This

²¹The Mankiw-Reis model is thus an example of what Sims (2001, sec. 8) calls "information-delay RE"

is a model in which the relevant cost of information flow is a fixed cost of logging on to the internet; on the occasions upon which one bears this cost, there is zero additional cost of downloading all of the available news with infinite precision.

A full comparison of these alternative types of incomplete-information models is beyond the scope of the present paper.²² But one disadvantage of the Mankiw-Reis approach is that it suggests that the rate at which suppliers (in aggregate) learn about particular events should be the same for all events, being determined by the single parameter that indicates the frequency of information updates. The noisy-information model instead makes it natural that learning should be more rapid about some events than about others.

5 Conclusions

We have seen that the Phelps-Lucas hypothesis, according to which temporary real effects of purely nominal disturbances result from imperfect information about the nature of these disturbances, deserves more continued interest than is often supposed. When one departs from the assumptions of the Lucas (1972) model in two crucial respects — introducing a monopolistically-competitive pricing framework in which the optimal pricing decisions of individual suppliers of goods depend crucially upon the prices that they expect others to set, and allowing individual suppliers’ subjective perceptions of current conditions to be contaminated by the noise that inevitably results from finite information-processing capacity — it is possible to explain not only real effects of purely nominal disturbances, but real effects that may persist for a substantial period of time.

We have shown that a model of this kind offers not only a potential explanation for the kinds of real effects that are usually mentioned as grounds for the assumption of substantial price stickiness, but also some prospect of an explanation of aspects of price dynamics that are

modelling, as opposed to “signal-extraction RE” modelling, the category to which the present paper would belong. See Sims for further discussion of the importance of this distinction.

²²Their implications are certainly not equivalent. For example, in the case of a random walk in nominal GDP, the Mankiw-Reis model does not imply inflation and output dynamics that are observationally equivalent to those predicted by the Calvo model, except in the special case that $\xi = 1$.

not easily reconciled with sticky-price models that assume optimization with full information, subject only to a constraint upon the frequency of price changes. Of course, there is no reason why the best model might not involve *both* sticky prices and noisy information — it may be most realistic to suppose that prices remain fixed for a time, but also that when revised they are adjusted on the basis of imperfect subjective perceptions of current conditions. But our preliminary investigation here suggests at least that there is an important cost to abstracting from the information limitations of price-setters.

While the model proposed here seeks to rehabilitate certain aspects of the explanation of the real effects of monetary policy advocated by Phelps and Lucas thirty years ago, acceptance of it would not necessarily lead to all of the conclusions emphasized in the earlier literature. The Lucas (1972) model was widely argued to imply that there should be little scope for the use of monetary stabilization policy to offset the fluctuations in output relative to potential that would otherwise be caused by other disturbances to the economy. For that model implied that monetary policy could have no effect on real activity that was systematically correlated with real disturbances unless the central bank were able to observe and respond to those disturbances, while the private sector could not also observe them and use them to predict the central bank's response. Successful monetary stabilization policy would then be impossible if the central bank's only information about real disturbances were also available to the general public.

But the interpretation proposed here of the nature of the relevant information limitations undermines this conclusion. If suppliers have an inaccurate estimate of current aggregate conditions not because of the unavailability of good data in the public domain, but because of paying insufficient *attention* to the available public-domain data, it is quite possible for the central bank to affect real activity in ways that are correlated with that public information. This should greatly increase the plausible scope of monetary stabilization policy. Analysis of the optimal conduct of policy in the presence of the kind of imperfect common knowledge described here should accordingly be an important topic for further study.

References

- Ball, Laurence, and David Romer, "Real Rigidities and the Non-Neutrality of Money," *Review of Economic Studies* 57: 183-203 (1990).
- Bils, Mark, and Yongsung Chang, "Understanding How Price Responds to Costs and Production," NBER working paper no. 7311, August 1999.
- Blinder, Alan S., Elie R.D. Canetti, David E. Lebow, and Jeremy B. Rudd, *Asking About Prices: A New Approach to Understanding Price Stickiness*, New York: Russell Sage Foundation, 1998.
- Calvo, Guillermo A., "Staggered Prices in a Utility Maximizing Framework," *Journal of Monetary Economics* 12: 383-398 (1983).
- Chow, Gregory C., *Analysis and Control of Dynamic Economic Systems*, New York: Wiley, 1975.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, "Modeling Money," NBER Working Paper no. 6371, January 1998.
- —, — —, and — —, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," NBER Working Paper no. 8403, July 2001.
- Green, D.M., and J.A. Swets, *Signal Detection Theory and Psychophysics*, New York: Wiley, 1966.
- Harvey, Andrew C., *Forecasting, Structural Time Series Models, and the Kalman Filter*, Cambridge: Cambridge Univ. Press, 1989.
- Lucas, Robert E., Jr., "Expectations and the Neutrality of Money," *Journal of Economic Theory* 4: 103-124 (1972).
- —, "Some International Evidence on Output-Inflation Tradeoffs," *American Economic Review* 63: 326-334 (1973).
- —, "Understanding Business Cycles," *Carnegie-Rochester Conference Series on Public Policy* 5: 7-29 (1977).
- Mankiw, N. Gregory, and Ricardo Reis, "Sticky Information versus Sticky Prices: A

- Proposal to Replace the New Keynesian Phillips Curve,” NBER working paper no. 8290, May 2001.
- Morris, Stephen, and Hyun S. Shin, “Global Games: Theory and Applications,” Cowles Foundation discussion paper no. 1275R, August 2001.
- Phelps, Edmund S., “Introduction: The New Microeconomics in Employment and Inflation Theory,” in E.S. Phelps *et al.*, *Microeconomic Foundations of Employment and Inflation Theory*, New York: Norton, 1970.
- —, “The Trouble with ‘Rational Expectations’ and the Problem of Inflation Stabilization,” in R. Frydman and E.S. Phelps, eds., *Individual Forecasting and Aggregate Outcomes*, Cambridge: Cambridge University Press, 1983.
- Sargent, Thomas J., “Equilibrium with Signal Extraction from Endogenous Variables,” *Journal of Economic Dynamics and Control*, 15: 245-273 (1991).
- —, *Bounded Rationality in Macroeconomics*, Oxford: Oxford University Press, 1993.
- Sims, Christopher A., “Comparison of Interwar and Postwar Business Cycles: Monetarism Reconsidered,” *American Economic Review* 70(2): 250-257 (1980).
- —, “Stickiness,” *Carnegie-Rochester Conference Series on Public Policy* 49: 317-356 (1998).
- —, “Implications of Rational Inattention,” unpublished, Princeton University, revised June 2001.
- Townsend, Robert M., “Forecasting the Forecasts of Others,” *Journal of Political Economy* 91: 546-588 (1983a).
- —, “Equilibrium Theory with Learning and Disparate Expectations: Some Issues and Methods,” in R. Frydman and E.S. Phelps, eds., *Individual Forecasting and Aggregate Outcomes*, Cambridge: Cambridge University Press, 1983b.
- Woodford, Michael, “Optimizing Models with Nominal Rigidities,” unpublished, Princeton University, revised July 2001.
- Zbaracki, Mark, Mark Ritson, Daniel Levy, Shantanu Dutta, and Mark Bergen, “The

Managerial and Customer Dimensions of the Cost of Price Adjustment: Direct Evidence from Industrial Markets,” unpublished, Univ. of Pennsylvania, November 1999.

Ziemer, R.E., and W.H. Tranter, *Principles of Communications*, fourth ed., New York: Wiley, 1995.