Comment on "Imperfect Knowledge, Inflation Expectations and Monetary Policy" by Athanasios Orphanides and John C. Williams

George W. Evans
University of Oregon

March 31, 2003

1 Introduction

This is a very nice paper. The main points are important, the structure is simple and clear, and I find the key arguments persuasive. In my comments I'm going to begin by summarizing the heart of the Orphanides-Williams argument. Then I will locate their paper within the rapidly growing literature on learning and monetary policy. Finally I will return to their paper and offer a number of specific comments on natural extensions or alternative approaches.

2 Summary of the Argument

Orphanides and Williams (OW) work with a simple two equation macro model. The first equation is an augmented Phillips curve with inertia:

$$\pi_{t+1} = \phi \pi^e_{t+1} + (1-\phi)\pi_t + \alpha y_{t+1} + e_{t+1},$$

where π_{t+1} is the rate of inflation between period t and period $t+1$, π^e_{t+1} is the rate of inflation over this period expected at time t, y_{t+1} is the level of the output gap in $t+1$ and e_{t+1} is a white noise inflation shock. The second equation is an aggregate demand relation that embodies a lagged policy effect,

$$y_{t+1} = x_t + u_{t+1}.$$

x_t is set by monetary policy at t and u_{t+1} is white noise. Through monetary policy it is assumed that policy makers are able one period ahead to control aggregate output up to the unpredictable random disturbance u_{t+1}.

The combination of this aggregate demand equation and the neoclassical (as opposed to neo-Keynesian) inflation equation yields a particularly tractable model for studying the effects of private agents learning. In particular, the timing assumptions are carefully crafted to yield simplicity.

Policy makers choose the x_t process to minimize

$$(1-\omega)Ey_t^2 + \omega E(\pi_t - \pi^*)^2.$$

This is a standard quadratic loss function. We can think of ω as reflecting policy makers preferences, which may (or may not) be derived from the preferences of the representative agent.

2.1 Optimal Policy under RE

Under rational expectations (RE) optimal policy takes the form of the feedback rule

$$x_t = -\theta^P(\pi_t - \pi^*),$$

where $\theta^P = \theta^P(\omega, \alpha/(1-\phi))$. This leads to an efficiency frontier, described by a familiar trade-off between σ_π and σ_y, shown in their Figure 1.

For this choice of feedback parameter, in the rational expectations equilibrium (REE) inflation follows the process

$$\begin{aligned} \pi_t &= c_0^P + c_1^P \pi_{t-1} + \text{noise}_t \\ E_t\pi_{t+1} &= c_0^P + c_1^P \pi_t, \end{aligned}$$

where c_0^P, c_1^P depend on $\theta^P \alpha/(1-\phi)$. Here noise$_t$ is white noise. The superscript "P" refers to "perfect knowledge," which OW use as a synonym for RE.

Thus under RE the problem is quite straightforward. How "aggressive" policy should be with respect to deviations of inflation from target depends in a natural way on the structural parameters ϕ, α and the policy maker preferences as described by ω.

2.2 Least Squares Learning

Now we make the crucial step of backing away from RE. Instead of assuming that agents are endowed *a priori* with RE we model the agents as forecasting in the same way that an econometrician might: by assuming a simple time series model for the variable of interest, and by estimating its parameters and using it to forecast. Specifically suppose that private agents believe that inflation follows an AR(1) process, as it does in an REE, but that they do not know c_0^P, c_1^P. Instead they estimate the parameters of

$$\pi_t = c_0 + c_1 \pi_{t-1} + v_t$$

by a least-squares-type regression, and at time t forecast

$$\pi_{t+1}^e = c_{0,t} + c_{1,t}\pi_t.$$

Over time the estimates $c_{0,t}, c_{1,t}$ are updated as new data become available. We consider two cases for this updating.

2.2.1 Infinite memory – "decreasing gain"

First we suppose that agents literally do Least Squares using all the data. We assume that policy-makers do not explicitly take account of private agent learning and follow the feedback rule with $\theta = \theta^P$. Then, with "infinite memory" (no discounting of observations), one can show (e.g. (Evans and Honkapohja 2001))

$$c_{0,t}, c_{1,t} \to c_0^P, c_1^P \text{ w.p.1},$$

so that asymptotically we get the optimal REE.

Technically the most convenient way to set up least squares learning by private agents is using the RLS (recursive least squares) algorithm.[1] In this algorithm the agents carry their parameter estimates (and an estimate of the second moment matrix of the regressors) into the next period. Updated estimates next period are then generated recursively using the most recent data point. Because each data point is counted equally by least squares the "gain" κ_t, i.e. the effective weight placed on the last data point, is given by

[1] The technique of formulating learning as a recursive algorithm, and then applying stochastic approximation tools to analyze convergence, was introduced by (Marcet and Sargent 1989).

$\kappa_t = 1/t$, i.e. by the inverse of the sample size. In the learning literature this is called the "decreasing gain" case, because $\kappa_t \to 0$ as $t \to \infty$.

I remark that convergence to the REE is not obvious. This is because the model is "self-referential." i.e. the evolution of the data depends on expectations and hence on the estimated coefficients and these in turn are updated using the data generated. Convergence to REE does take place because the equilibrium in this model satisfies the "E-stability" conditions that govern stability in such a system.

2.2.2 Finite memory – "constant gain"

OW make a small but significant change to the standard least-squares updating formula. Instead of assuming that all observations count equally, they discount or downweight past data. In terms of the RLS algorithm, this is accomplished technically by setting the gain, the weight on the most recent observation used to update estimates, to a small constant i.e. setting $\kappa_t = \kappa$ (e.g. 0.05).

Why would it be natural for agents to use a constant rather than decreasing gain? The main rationale for this procedure is that it allows estimates to remain alert to structural shifts. As economists, and as econometricians, we tend to believe that structural changes occasionally occur, and we might therefore assume that private agents also recognize and allow for this. Although in principle one might attempt to model the process of structural change, this typically unduly strains the amount of knowledge we have about the economic structure. A reasonable alternative is to adjust parameter estimators to reflect the fact that recent observations convey more accurate information on the economy's law of motion than do past data, and "constant gain" estimators are one very natural way of accomplishing this downweighting of past data.[2]

2.2.3 Implications of constant gain Least Squares

With constant gain procedures, estimates no longer fully converge to the REE. The estimators $c_{0,t}, c_{1,t}$ converge instead to a stochastic process. Be-

[2]Two remarks are in order. First, an alternative rationale for constant gain is that it can be an equilibrium in learning rules, even if structural change is not present – see Section 14.4 of (Evans and Honkapohja 2001). Second, there are other ways of allowing for strucrural change, e.g. through time varying gain sequences or explicit models of structural variation.

cause of this OW use the term "perpetual learning" to refer to the constant gain case.

If the gain parameter κ is very small, then estimators will be close to the REE values for most of the time with high probability, and output and inflation will be near their REE paths. Nonetheless, small plausible values like $\kappa = 0.05$ can lead to very different outcomes in the calibrations OW consider. In particular they find:

1. The standard deviations of $c_{0,t}$ and $c_{1,t}$ are large even though forecast performance remains good.

2. There is a substantial increase in the persistence of inflation, compared to the REE.

3. Most strikingly, the policy frontier shifts out very substantially and in a non-monotonic way (See their Figure 5).

2.2.4 Policy Implications

Under perpetual learning if policy makers keep to the same class of rules

$$x_t = -\theta^S(\pi_t - \pi^*),$$

then they should choose a different θ. Here the notation θ^S is meant to indicate that we restrict policy makers to choose from the same "simple" class of policy rules. There are four main implications for policy in the context of constant gain (perpetual) learning by private agents.

1. Naive policy choice can be strictly inefficient. This is illustrated in the second diagram of their Figure 5. By "naive" policy is meant the policy that assumes RE (perfect knowledge) on the part of agents, when in fact the agents are following perpetual learning with $\kappa > 0$. In particular there are cases in which increasing θ^S would decrease the standard deviations of *both* inflation and output.

2. In general, policy should be more hawkish, i.e. under perpetual learning the monetary authorities should pick a larger θ^S than if agents had RE.

3. Following a sequence of unanticipated inflation shocks, inflation doves (i.e. policy-makers with low θ reflecting a low ω) can do very poorly. This is illustrated in their Figure 3.

4. If the inflation target π^* is known to private agents, so that they need estimate only the slope parameter c_1, then the policy frontier is more favorable than when it is not known. This is illustrated in the first diagram of their Figure 9.

I'll return to a discussion of these and other specific results after discussing learning and monetary policy in a more general setting.

3 Learning in Monetary Policy

Recently considerable research has begun to focus on the implications for monetary policy when explicit account is taken of the literature on adaptive/econometric learning in macroeconomics.[3]

I will give a selective overview of this recent research and locate OW within this context. Then I will return to a discussion of OW. There are four main issues I'll use to group my general remarks: (i) the theoretical roles played by learning, (ii) the question of who or what group of agents is learning, (iii) the particular implications of constant gain learning, and (iv) some further (personal) thoughts on rationality.

3.1 Roles for Learning

There are three main types of result that can be delivered by incorporating learning into a monetary policy model.

3.1.1 Stability under private agent learning

Under learning an REE need not necessarily be stable under private agent learning. It is logically possible that if agents follow LS learning (with the usual decreasing gain) that the system fails to converge to an REE, even if their parameter estimates are initially close to the REE.

This theoretical possibility of instability turns out to be a genuine concern for monetary policy in New Keynesian/New PC models (as is the related but distinct issue of indeterminacy). (Bullard and Mitra 2002) show that

[3]For example, two recent workshops or conferences have considered this topic, one at the Cleveland Federal Reserve Bank, in February 2001, on "Learning and Model Misspecification," and a second at the Atlanta Federal Reserve Bank, in March 2003, on "Monetary Policy and Learning."

stability under private agent learning should not be taken for granted if policy makers follow Taylor-type rules. Depending on the specific formulation of the rule, instability can arise for certain choices of parameter settings. (Evans and Honkapohja 2003a), (Evans and Honkapohja 2002b), and (Evans and Honkapohja 2002a) examine this issue in the context of optimal monetary policy. They show that stability under learning is a pervasive problem when the interest rate rule is formulated as a reaction to fundamental shocks, but can be overcome when the rule reacts appropriately to private expectations. Recent work by (Preston 2003) has considered this issue in the context of long-horizon agents.

3.1.2 Selection criterion

In some models the phenomenon of indeterminacy, i.e. multiple REE, arises. In this setting, learning can provide a natural way of choosing between equilibria. A particular question of interest is the following. It is known that when a steady state of a linear model is indeterminate there exist "sunspot" equilibria, i.e. REE in which the solution is driven by extraneous noise. Such solutions, with economic fluctuations driven in a self-fulfilling way by extrinsic random variables, would usually be considered an unintended and undesirable by-product of economic policy. A particular question of interest, in cases of multiple equilibria, is whether the sunspot equilibria can be stable under learning.

It has been known for some time that it is possible in some cases for sunspot equilibria to be stable under learning. This was initially demonstrated by (Woodford 1990) in the context of the overlapping generations model of money. In general, whether a sunspot equilibrium is stable under learning depends on the model and the particular solution, see Chapter 12 of (Evans and Honkapohja 2001). There has been recent interest in whether stable sunspot solutions can arise in more realistic monetary models. In particular, (Evans, Honkapohja, and Marimon 2001) look at when this can occur cash-in-advance models, and (Honkapohja and Mitra 2001), (Carlstrom and Fuerst 2000) and (Evans and McGough 2003) examine the issue for New Keynesian models.

3.1.3 Non-REE learning dynamics

Finally, we move to the possibility that the economy under learning generates solutions that in some way go beyond rational expectations. Here it appears useful to group results into two broad categories. One possibility is that learning converges to a "Restricted Perceptions Equilibria." This arises if agents are endowed with an econometric model that is misspecified asymptotically, as discussed in Chapter 13 of (Evans and Honkapohja 2001)). For example, agents may omit some variables that help forecast the variables of interest or their forecasting model may fail to capture nonlinearities that are present.

Somewhat more radically, learning may generate "persistent learning dynamics" (see Chapter 14 of (Evans and Honkapohja 2001)) as a result of local instability of an REE under learning (as in (Bullard 1994)) or due to a learning rule that fails to fully converge to REE parameter values (as in constant gain learning rules). The OW paper falls into this last class: private agents use a learning rule in which parameter estimates never quite converge to REE values. This "perpetual learning" then turns out to have major policy implications, even when the deviation from REE might be thought not too large.

3.2 Who is learning?

The earliest literature on learning focussed on private agents, i.e. households and firms. In dynamic macroeconomic models private agents, in order to make optimal decisions, must make forecasts of relevant future variables. Clearly the expectations of households and firms do matter enormously for the actual evolution of the economy. The rational expectations revolution made the crucial advance of defining and analyzing what it means for expectations to be consistent with the economic structure and optimizing agents. However, this has had the potential disadvantage of demoting private expectations as an independent force. Consequently it was natural that the initial focus of the learning literature was on private agent learning. The OW paper follows the primary strand of the literature in this respect.

However, policy makers also need to form expectations and make forecasts and they too are not endowed with full knowledge of the economic structure or fully rational forecast functions. Some recent research has begun to tackle this issue. Most notably Thomas Sargent's book on the disinflation in the

1990s, (Sargent 1999), emphasized learning by policy makers about a (misspecifed) Phillips curve trade-off. Sargent's model incorporates a tantalizing combination of misspecification, learning and optimal policy formulation.

Obviously it is possible to allow for separate learning by private agents and policy-makers. In fact (Sargent 1999) actually allows for this in some cases, though much of his analysis, and that of (Cho, Williams, and Sargent 2002), focus on learning by policy makers with rational expectations assumed for private agents. Simultaneous learning by policy makers is also analyzed in (Honkapohja and Mitra 2002) and discussed in (Evans and Honkapohja 2002a).

There is an additional asymmetry that should be noted. Both private agents and policy makers need to make forecasts of future aggregate variables, but in addition, implementation of optimal policy may require simultaneous estimation of structural parameters. This issue is considered in (Evans and Honkapohja 2003a) and (Evans and Honkapohja 2002a).

3.3 Constant gain learning

As already emphasized, the use of constant gain (or "perpetual") learning plays a central role in OW. In general constant gain learning can lead to a number of phenomena. First, the work of (Sargent 1999), (Cho, Williams, and Sargent 2002), (Williams 2002) and (Bullard and Cho 2002) emphasize the possibility of "escapes," i.e. occasional big deviations from a unique REE. This is a surprising finding: for significant periods of time learning dynamics can drive the economy away from the REE, but in a predictable direction.

When there are multiple REE escapes can take a different form. The most widely examined case is the case of multiple distinct REE steady states. Here escapes take the form of periodic shifts between the different steady states as a result of large random shocks interacting with the learning dynamics. This phenomenon is seen in Chapter 14 of (Evans and Honkapohja 2001)), the hyperinflation model of (Marcet and Nicolini 1998), the exchange rate model of (Kasa 2002) and the liquidity trap model of (Evans and Honkapohja 2003b).

Finally, it turns out that, even in a quite standard model with a unique REE, and without the more exotic effects just described, constant gain learning has significant implications for optimal policy. This is the important new finding that is demonstrated in the current paper by OW.

3.4 Some further thoughts on rationality

In constructing economic models we have three kinds of agents: (i) private agents, (ii) policy makers and (iii) economists (us). In the bad old days of adaptive expectations, private agents made systematic mistakes, but we the economists were very smart. We told policy makers what to do, so they were smart too.

The rational expectations revolution changed all this. Now private agents became smart, and policy makers (and earlier economists) were mistaken, as shown by the Lucas critique. As theorists we were again smart (because we understood how private agents really formed expectations), but as econometricians we were not quite so smart. This is because as econometricians we had to estimate parameters that were known with certainty by the private agents and theorists.

The adaptive learning viewpoint has the enormous advantage over these earlier approaches that it (potentially) achieves greater cognitive consistency between these three kinds of agents. In particular, private agents are modeled as behaving like econometricians, i.e. like economists in our forecasting role. Of course as theorists we still typically analyze models with a specified structure that is effectively known only to us, but at least it can be consistently treated as unknown to private agents, policy makers and econometricians. Furthermore, the degree of smartness of each group is a matter of choice or judgement for us as theorists.

An important aspect of this "bounded rationality" approach is that many features of rational expectations do carry over to the adaptive learning approach. For example, the Lucas critique can apply under bounded rationality, as emphasized in (Evans and Ramey 2001). The Lucas critique will often arise if agents attempt to forecast in an optimal way, even if they are not perfectly rational in the sense of "rational expectations."

4 Back to Orphanides & Williams

Returning now to the OW paper, let me make some specific critical comments and suggest some extensions.

1. The inflation shocks experiment. My first point concerns the inflation shocks scenario shown in their Fig. 3. OW examine a sequence of unanticipated positive inflation shocks starting with $e_1 = 2\%$ and declining to zero

over 9 (semi-annual) periods. My main point is that this is more like a structural shift, and that the effects are the same as a decrease in potential output over 4 years. This raises several questions that would need to be explicitly addressed in a full treatment of this issue.

Suppose, for example, that e_{t+1} partly predictable, as seems appropriate for a structural shift, and that the loss function is

$$L = E_0\{\sum_{t=0}^{\infty}\left[(1-\omega)(y_t - y_t^*)^2 + \omega(\pi_t - \pi^*)^2\right]\}.$$

Depending on the source of the shock, policy makers may want to lower their output target y_t^* (to $y_t^* = -\alpha^{-1}e_t$). Even if policy-makers continue to set $y_t^* = 0$, policy should take into account expected $e_{t+1} > 0$.

This is perhaps a set-up in which it would be particularly fruitful also to incorporate policy maker learning.

2. Bias towards "hawkishness." OW show that policy makers should be more hawkish. The intuition for this result is fairly intuitive. A more hawkish (high θ) policy helps to keep inflation expectations π_{t+1}^e "in line," i.e. closer to rational expectations values. This gives policy an additional role, besides stabilizing y and π, and this additional role means that under perpetual learning it is optimal for policy makers to be more hawkish than they would be, for given policy maker preferences, under RE.

This observation leads naturally to the question of how robust this result is. In particular, in New Keynesian models y_{t+1}^e also matters. The structure is such models is

$$y_t = -\varphi(i_t - \pi_{t+1}^e) + y_{t+1}^e + g_t$$
$$\pi_t = \lambda y_t + \beta\pi_{t+1}^e + \gamma\pi_{t-1} + u_t.$$

Will the presence of y_{t+1}^e in the "IS" curve make the direction of bias for the policy-maker ambiguous? The answer is not clear *a priori* and would need to be explicitly analyzed.

3. Choice of gain parameter κ. The value of κ is taken as given and not explained. This is quite standard in the constant gain learning literature. In one respect this is convenient, since it can then be treated as a parameter to be estimated empirically.

However, one can think about the issue further from a theoretical viewpoint. The most typical rationale for introducing constant gain, as indicated above, is that it is a way of allowing for structural shifts. The choice of κ

can then be thought of a providing a balance between tracking and filtering: high values of κ allow the estimator to better track structural change, but with the disadvantage of yielding noisier estimators.

One possibility would then be to explicitly introduce structural shifts into the model and find the optimal value of κ. This type of exercise is done in Chapter 14 of (Evans and Honkapohja 2001) and in (Evans and Ramey 2001). In OW this would add complexity and is unlikely to matter. However, the issue of the optimal choice of gain is likely to become important in future work.

4. Smarter agents. Using the bounded rationality approach one can always ask: should the agents be smarter? less smart? This is always a matter of judgement. There are several possible ways in which the private agents in OW could be "smarter." For example, private agents could be modeled as estimating an AR(p), instead of an AR(1). Indeed, one could consider the possibility that the agents choose the lag length p in the same way as an applied econometrician. Similarly, agents might consider forecasting based on a vector-autoregression, perhaps using one of the standard statistics to choose the order of the VAR.

It seems likely that the qualitative results would be unaffected, but it would be of interest to know how the detailed results depend on such specification issues. It might appear unsatisfactory, compared to the lack of ambiguity in the RE approach, to be faced with questions about lag length and model specification. But this is really a strength of the adaptive learning framework. Econometricians dealing with forecasting and estimation problems inevitably in practice face precisely such issues. It seems absurd to assume that private agents and policy-makers have clear-cut answers to problems that in effect remain research issues for us as econometricians.

5 Conclusions

This is an important paper. Theoretically, Orphanides and Williams provide a new reason for studying adaptive learning, based on optimal policy when agents follow "perpetual learning" rules. From an applied viewpoint, the paper suggests another factor that can generate stagflation, and it provides policy recommendations that are intuitive and plausible. I hope (and confidently anticipate) that the authors (and others) will do more work along these lines.

References

BULLARD, J. (1994): "Learning Equilibria," *Journal of Economic Theory*, 64, 468–485.

BULLARD, J., AND I.-K. CHO (2002): "Escapist Policy Rules," Working paper.

BULLARD, J., AND K. MITRA (2002): "Learning About Monetary Policy Rules," *Journal of Monetary Economics*, 49, 1105–1129.

CARLSTROM, C. T., AND T. S. FUERST (2000): "Forward-Looking versus Backward-Looking Taylor Rules," Working paper, no. 09, Federal Reserve Bank of Cleveland.

CHO, I.-K., N. WILLIAMS, AND T. J. SARGENT (2002): "Escaping Nash Inflation," *Review of Economic Studies*, 69, 1–40.

EVANS, G. W., AND S. HONKAPOHJA (2001): *Learning and Expectations in Macroeconomics*. Princeton University Press, Princeton, New Jersey.

——— (2002a): "Adaptive Learning and Monetary Policy Design," Journal of money, credit and banking conference paper, www.valt.helsinki.fi/RAKA/seppo.hmtl and http://darkwing.uoregon.edu/∼gevans/.

——— (2002b): "Monetary Policy, Expectations and Commitment," mimeo, www.valt.helsinki.fi/RAKA/seppo.hmtl, and http://darkwing.uoregon.edu/∼gevans/.

——— (2003a): "Expectations and the Stability Problem for Optimal Monetary Policies," *Review of Economic Studies, forthcoming*.

——— (2003b): "Policy Interaction, Expectations and the Liquidity Trap," Paper prepared for federal reserve bank of atlanta conference on monetary policy and learning.

EVANS, G. W., S. HONKAPOHJA, AND R. MARIMON (2001): "Stable Sunspot Equilibria in a Cash in Advance Economy," mimeo, www.valt.helsinki.fi/RAKA/seppo.hmtl, and http://darkwing.uoregon.edu/∼gevans/.

Evans, G. W., and B. McGough (2003): "Monetary Policy, Indeterminacy and Learning," work in progress.

Evans, G. W., and G. Ramey (2001): "Adaptive Expectations, Underparameterization and the Lucas Critique," Working paper 237, University of Oregon.

Honkapohja, S., and K. Mitra (2001): "Are Non-Fundamental Equilibria Learnable in Models of Monetary Policy?," Working paper, www.valt.helsinki.fi/RAKA/seppo.hmtl.

——— (2002): "Performance of Monetary Policy with Internal Central Bank Forecasting," manuscript, www.valt.helsinki.fi/RAKA/seppo.hmtl.

Kasa, K. (2002): "Learning, Large Deviations, and Recurrent Currency Crises," *International Economic Review, forthcoming.*

Marcet, A., and J. P. Nicolini (1998): "Recurrent Hyperinflations and Learning," Working Paper 1875, CEPR.

Marcet, A., and T. J. Sargent (1989): "Convergence of Least-Squares Learning Mechanisms in Self-Referential Linear Stochastic Models," *Journal of Economic Theory*, 48, 337–368.

Preston, B. (2003): "Learning about Monetary Policy Rules when Long-Horizon Forecasts Matter," mimeo, Princeton University.

Sargent, T. J. (1999): *The Conquest of American Inflation.* Princeton University Press, Princeton NJ.

Williams, N. (2002): "Escape Dynamics in Learning Models," mimeo.

Woodford, M. (1990): "Learning to Believe in Sunspots," *Econometrica*, 58, 277–307.