Evolution and Intelligent Design

Thomas J. Sargent[*†]

September 21, 2007

Abstract

This paper discusses the relationship between two sources of ideas that influence monetary policy makers today. The first is a set of analytical results that impose the rational expectations equilibrium concept and do 'intelligent design' by solving Ramsey and mechanism design problems. The second is a long trial and error learning process that constrained government budgets and anchored the price level with a gold standard, then relaxed government budgets by replacing the gold standard with a fiat currency system wanting nominal anchors. Models of out-of-equilibrium learning tell us that such an evolutionary process will converge to a self-confirming equilibrium (SCE). In an SCE, a government's probability model is correct about events that occur under the prevailing government policy, but possibly wrong about the consequences of other policies. That leaves room for more mistakes and useful experiments than exist in a rational expectations equilibrium.

Keywords: Rational expectations equilibrium, mechanism design, model misspecification, learning, evolution, observational equivalence, self-confirming equilibrium. (JEL).

[*]New York University and Hoover Institution. Email: ts43@nyu.edu.

[†]This is a draft of my presidential address to the American Economic Association in January 2008. It continues a long conversation I have had with Chris Sims (see Sims (1982)). After Bob Lucas read Sargent (1984), he wrote me that "With friends like you, Chris doesn't need enemies." Maybe it is more complicated than that. I thank Marco Bassetto, In-Koo Cho, Timothy Cogley, Lars Peter Hansen, Kenneth Kasa, Narayana Kocherlakota, Larry Jones, Athanasios Orphanides, Carolyn Sargent, Hyun Shin, Christopher Sims, François Velde, Peyton Young, and Tao Zha for helpful comments. I thank the National Science Foundation for research support.

1 Introduction

> ... the ideas of economists and political philosophers, both when they are right and when they are wrong, are more powerful than commonly understood. Indeed the world is ruled by little else. Practical men, who believe themselves to be quite exempt from any intellectual influences, are usually the slaves of some defunct economist. Keynes (1936, p. 383)[1]

Today the leading practical men at the most important institutions in my field, the Federal Reserve System and the Bank of England, are distinguished academic economists. They have used prevailing academic ideas about macroeconomics to choose policy actions within existing institutions. Some of them created new institutions. For example, in 1997 Mervyn King and a few others designed decision making protocols for Britain's monetary policy committee virtually from scratch in a matter of days.

This essay is about two important sources of prevailing ideas in macroeconomics, how they came to influence thinking at leading central banks, their accomplishments and limits, and the occasional tensions between them. The first is a collection of powerful results that apply the rational expectations equilibrium concept. The second is an evolutionary process embedded in an historical economic record littered with discarded ideas and policies. I consider how rational expectations equilibria and evolutionary processes are linked. An optimist can hope that an evolutionary process will converge to a rational expectations equilibrium, but that is usually hoping for too much. A system of adaptive agents converges to a self-confirming equilibrium in which all agents have correct forecasting distributions for events that occur often enough along an equilibrium path, but possibly mistaken views about policies and outcome paths that will not be observed. This matters because intelligent macroeconomic policy design within rational expectations equilibria hinges on knowing and manipulating expectations about events that will not be observed. Self-confirming equilibria allow government models to survive that imply mistaken policies even though they match historical data well.

I devote sections 2, 3, and 4, to describing powerful and useful ramifications of the assumption of rational expectations for macroeconomic theory and empirical work. When I was 30 years old, rational expectations was a new hypothesis whose appropriateness for macroeconomic modeling was hotly disputed (that made it fun). Those disputes ended long ago with the outcome that the rational expectations equilibrium concept is now adopted automatically in most applications, usually without discussion.[2] By equating all subjective distributions to an equilibrium objective distribution, the rational expectations hypothesis makes agents' beliefs disappear as extra components of a theory. That is what modelers mean when they say they abide by the discipline imposed by rational expectations.

The common beliefs assumption underpins influential rational expectations doctrines about whether observed inflation-unemployment dynamics can be exploited by policy mak-

[1]A younger Keynes was less optimistic about the influence of economists' ideas:

> Financiers of this type [Lord Rothschild, Lord Avebury, Lord Swaythling] will not admit the feasibility of anything until it has been demonstrated to them by practical experience. It follows, therefore, that they will seldom give their support to what is new. Keynes (1913, pp. 24-25)

[2]For exceptions, see for example, Brunnermeier and Parker (2005), Brunnermeier et al. (2007).

ers, the time inconsistency of benevolent government policy, the capacity of reputation to substitute for commitment, the incentives for a policy maker of one type to emulate another, the wisdom of making information public, trading volume, and asset prices. The rational expectations hypothesis carries a self-contained econometric theory that has changed how applied workers think about parameter identification (less emphasis on zero-one exclusion restrictions, more on cross-equation restrictions). The rational expectations assumption that nature, all agents inside a model, and the economist outside a model share a unique set of beliefs is stressed especially in modern theories of optimal macroeconomic policy that focus on how a government can beneficially shape a system of expectations. I view this 'intelligent design' approach to macroeconomics as the zenith of rational expectations macroeconomics. It has already been influential. For example, it defines the terms in which the most persuasive defenses of inflation targeting and low taxation of capital are cast. It promises more.

I devote sections 5 and 6 to a framework for thinking about learning, then use it in section 7 and appendix A to describe some undirected evolutionary processes of special relevance to macroeconomists. These historical processes show how ideas that were once prevalent, but have now been discarded, shaped policies and generated experiments that brought us to where we are now. The adaptive models of sections 5 and 6 indicate why we should not always expect these historical processes to have limiting outcomes that are rational expectations equilibria or that solve dynamic mechanism design problems. The unmanaged historical episodes that I describe in section 7 exhibit decision making by trial and error, conflicting theories, purposeful and inadvertent experiments, unintended consequences, and discoveries that formed modern doctrines about central banking.

I title this essay "Evolution *and* Intelligent Design" rather than "Evolution *versus* Intelligent Design" because both ways of thinking are useful. By characterizing how self-confirming equilibria and the transition paths to and from them can differ from rational expectations equilibria, models of evolution and learning in macroeconomics can teach us about some leaps of faith that we make as we use the rational expectations hypothesis to advance the intelligent design of macroeconomic policies.

2 Ramifications of the rational expectations equilibrium concept

The rational expectations version of a common beliefs assumption goes beyond the agreement that emerges from the analysis of Aumann (1976) because the rational expectations hypothesis assumes that nature (a.k.a., the data generating mechanism) shares the model too. Furthermore, the econometricians know the form of the model, though they are uncertain about some parameters. This gives rational expectations much of its empirical power.[3]

[3]Muth (1960) provided an unusual perspective on the disappearance of free parameters in his first paper on rational expectations. Muth reverse engineered a univariate stochastic process for income that renders the adaptive expectations scheme used by Friedman (1956) (i.e., a one-parameter geometric distributed lag with weights adding up to one) consistent with optimal forecasting. In Muth's example, one free parameter determines the consumer's forecasting function and what disappears are any additional free parameters governing a univariate income process. Sargent (1977) reverse engineered a bivariate money growth and inflation process that renders the one-parameter adaptive expectation scheme used by Cagan (1956) consistent with

Important theoretical findings and empirical strategies exploit the communism of beliefs in a rational expectations equilibrium:

1. *Neutrality theorems.* Early applications of rational expectations communism led to the monetary neutrality theorems of Lucas (1972b), Sargent (1973), and Sargent and Wallace (1975) that asserted that one deterministic money supply rule yields the same probability distribution of output and unemployment outcomes as any other, so that Friedman's k percent rule for money growth cannot be dominated. In related contexts, these authors derived that result by taking the model of the natural unemployment rate hypothesis created by Phelps (1967) and replacing the diversity of beliefs that Phelps had assumed (in Phelps's model, the modeler can forecast better than the agents in the model) with rational expectations. This converted the nontrivial dynamic programming problem that Phelps studied for managing dynamic inflation-unemployment trade-offs into a trivial problem of choosing a monetary policy that keeps inflation low. This work provides an important part of the intellectual foundations for inflation targeting, a main purpose of which is to convince the monetary authorities to abstain from exploiting the Phillips curve either when they prefer to push unemployment below the natural rate or when they under estimate the natural rate (see Orphanides (2003) and Primiceri (2006)).

2. *Time consistency.* The availability of the rational expectations equilibrium concept enabled Kydland and Prescott (1977) and Calvo (1978) to explain how alternative timing protocols affect a benevolent government's incentives to manipulate and then confirm prior expectations about its actions.[4] The time consistency 'problem' is the observation that equilibrium outcomes in a representative agent economy depend on the timing protocol for decision making that the modeler imposes on a benevolent government. Better outcomes emerge if the government chooses a history-contingent plan once-and-for-all at time 0 than if it chooses sequentially. By choosing future actions at time 0, the government can take into account how expectations about its actions at times $t > 0$ influence private agents' actions at all dates between 0 and t. A government must ignore those beneficial expectations effects if it is forced to choose sequentially.[5]

   By choosing at time 0, the government uses a 'commitment technology' that binds it to confirm those private sector expectations when time t comes to pass. Optimal government plans under commitment can be formulated recursively in terms of a vector of nonnegative Lagrange multipliers on the private sector Euler equations. These serve as implementability conditions that require time t government actions to confirm prior expectations.

rational expectations.

[4]While technical treatments of the time consistency problem rely heavily on the rational expectations equilibrium concept, what is really needed to spot the problem is that private agents care about future government actions. At the U.S. Constitutional Convention on August 16, 1787, Gouverneur Morris, Oliver Ellsworth, James Madison, and George Mason were evidently aware of the time consistency problem, while Edmund Randolph and George Mason raised doubts about tying the hands of the government because they could not foresee all contingencies. See Madison (1987, pp. 470-471).

[5]A time t government takes actions of time $\tau > t$ governments fixed in a no-commitment or Markov perfect political-economic equilibrium. For example, see Krusell et al. (1997).

3. *Can reputation substitute for commitment?* A credible public policy is an equilibrium system of expectations that gives a benevolent government incentives to confirm prior expectations about its future actions, actions to which to it cannot commit because it chooses sequentially. When applied to the credible public plans models of Stokey (1989, 1991) and Chari and Kehoe (1993b,a), the powerful recursive formulation of Abreu et al. (1986, 1990) indexes equilibria by the government's continuation value.[6] By making an intrinsically 'forward-looking' variable, a promised discounted value for the representative household, also be a 'backward-looking' state variable that encodes history, Abreu et al. (1986, 1990) tie past and future together in a subtle way that exploits the rational expectations equilibrium concept. There are multiple reputational equilibria, i.e., multiple systems of common expectations that a benevolent government would want to confirm, with good and bad equilibria being tied together via incentive constraints. Continuation values associated with future paths that will not be taken in equilibrium induce the government to take actions that confirm what people expect along an equilibrium path.

   The key object in this literature is a history-dependent government strategy, that is, a sequence of functions mapping histories into time t government actions. A government strategy plays two roles, first, as a *decision rule* for the government and, second, as a *system of private sector expectations* about government actions that the government always wants to confirm. The theory is silent about who chooses an equilibrium system of beliefs, the government (after all, it *is* the *government's* decision rule) or the public (but then again, they *are* the private sector's expectations). This makes it difficult to use this theory to formulate advice to policy makers about actions that can help it to earn a good reputation. Instead, the theory is about how a government comes into a period confronting a set of private sector expectations about its actions that it will want to confirm.[7] The next type of rational expectations theory that I will mention seems to be a more promising framework for thinking about how to acquire a good reputation. But the theory of credible public policy seems to explain why some policy makers who surely knew about better decision rules chose instead to administer ones supporting bad outcomes.[8]

4. *Reputational models with two types of policy maker.* These models (for example Ball (1995)) also exploit the assumption of common expectations. Two types of policy makers are distinguished by their preferences and therefore how prone they are to opportunistic behavior. The policy makers know their own preferences but the pub-

[6]For some applications, see Ljungqvist and Sargent (2004, ch. 22), Chang (1998), and Phelan and Stacchetti (2001).

[7]Blinder (1998, pp. 60-62) struggles with this issue when he describes the pressures he perceived as Fed Vice Chairman not to disappoint the market. While Blinder's discussion can be phrased almost entirely within the rational expectations, the account by Bernanke (2007) of the problems the Fed experiences in anchoring private sector expectations cannot. Bernanke explicitly acknowledges that he is thinking in terms of objects outside a rational expectations equilibrium.

[8]Chari et al. (1998) and Albanesi et al. (2002) interpret the big inflation of the 1970s and its stabilization in the 1980s in terms of the actions of benevolent and knowledgeable policy makers who became trapped within but, thanks to a sunspot, eventually managed to escape expectations traps within subgame perfect or Markov perfect equilibria.

lic doesn't.[9] An equilibrium system of expectations imparts incentives to the more opportunistic type to masquerade as a less opportunistic preference type.[10]

This theory seems to be on the minds of monetary policy committee members and critics who talk about taking actions to earn reputations, i.e., to convince the market that the committee is dominated by policy makers of a particular type.

5. *No-trade theorems.* A rational expectations communism of beliefs underlies a benchmark model of trading volumes in financial markets. The no-trade theorems of Milgrom and Stokey (1982) and Tirole (1982) rely heavily on the assumption that all agents share a common probability model and come from applying conditional expectations operators taken with respect to that common distribution. Harrison and Kreps (1978) show that the no-trade outcome vanishes when only small perturbations are made away from the common beliefs assumption, while Morris (1996) shows how outcomes like Harrison and Kreps's can long endure even as a Bayesian consistency theorem eventually drives disparate beliefs toward common ones.

6. *Forecasting the forecasts of others.* Allen et al. (2006) and Amato and Shin (2006) apply the single-probability model discipline that comes from rational expectations to produce a model that, by formalizing the beauty contest metaphor of Keynes (1936, p. 156), significantly extends its implications. They assume a set of agents, each of whom has access to a common public signal and also to his own private signal about a fundamental variable of interest. They also assume that agents care not about the fundamental variable itself but rather about the average over all agents' beliefs about that fundamental. They construct an operator that averages beliefs across agents. Because it does not obey something analogous to a law of iterated expectations, each successive iteration of this operator raises attention paid to the public signal. They use that property to explain why the existence of the public signal and the assumption that agents care not about the fundamental *per se* but about the average expectation of the fundamental cause a degradation of private information. As the number of agent grows large, and therefore as the aggregate of private information becomes very informative about the fundamental, instead of pooling private information, agents in the Allen et al. (2006) and Amato and Shin (2006) settings discard it.[11] This work has enlivened and enlightened controversies about central bank transparency.[12]

7. *Empirical asset pricing.* While the pure theory of asset pricing does not impose the rational expectations assumption, many empirical researchers have. Hansen and

[9]Strategic delegation, that is, assigning monetary policy to someone whose preferences give him no temptation to depart from a Ramsey plan as time unfolds, is a special case of this structure. See Rogoff (1985).

[10]Some statements by Bernanke (2007) about the public not knowing the monetary authorities' objective functions can be interpreted in terms of these models.

[11]Lucas (1975) assumed that private agents could pool their private information. Kasa (2000) and Pearlman and Sargent (2005) established that pooling is an outcome in the setting of Townsend (1983)), a difference in outcomes relative to those of Allen et al. (2006) and Amato and Shin (2006) that can be traced to partly to agents' objective functions.

[12]It is testimony to the ambiguity of words that Keynes's beauty contest metaphor propelled Harrison and Kreps (1978) to move away from rational expectations while it prompted the authors mentioned in the text to use it more.

Singleton (1983) and many others have deduced restrictions on time $t+1$ returns $R_{j,t+1}(x_{t+1})$ for asset j by using the Euler equation for consumer i

$$1 = \beta \int_{x_{t+1}} \frac{u_i'(c_{t+1}^i(x^{t+1}))}{u_i'(c_t^i(x^t))} R_{j,t+1}(x_{t+1}) f_i(x_{t+1}|x^t) dx_{t+1} \tag{1}$$

where $f_i(x_{t+1}|x^t)$ is consumer i's subjective one-step-ahead transition density for a state variable x_{t+1} that determines both returns and time $t+1$ consumption c_{t+1}^i. Throughout this paper, I use f to denote a probability density and x^t to denote a history $x_t, x_{t-1}, \ldots, x_0$. To make (1) operational, empirical workers in macroeconomics and finance have supplemented it with a rational expectations assumption that eliminates the substantial heterogeneity of beliefs that is permitted in asset pricing theory. In a finite-horizon setting, Harrison and Kreps (1979) showed that, when there are complete markets, the *stochastic discount factor*

$$m_{t+1} = \beta \frac{u_i'(c_{t+1}^i(x^{t+1}))}{u_i'(c_t^i(x^t))} \frac{f_i(x_{t+1}|x^t)}{f(x_{t+1}|x^t)} \tag{2}$$

is unique, meaning that it does not depend on i. Here $f(x_{t+1}|x^t)$ is a common physical conditional measure that does not depend on i. Because offsetting differences in utility functions and probabilities leave the left side of (2) fixed, the uniqueness of the stochastic discount factor allows different densities $f_i(x_{t+1}|x^t)$'s for different is. Suppose that f is the true measure that actually governs outcomes. Then Blume and Easley (2006) showed that in complete markets economies with Pareto optimal allocations and an *infinite* horizon, for agents i who survive in the limit, the measures f_i and f merge, meaning that they agree about probabilities of tail events.[13] For a complete markets economy with a Pareto optimal allocation, presumably we could use the Blume and Easley result to defend a rational expectations assumption by assuming that at time 0 all agents have access to an infinite history of observations.

Grossman and Shiller (1981), Hansen and Singleton (1983), Hansen and Richard (1987) sought an econometric framework to apply when markets are incomplete, in which case it is not enough to appeal Blume and Easley's market survival justification for assuming beliefs that are common or eventually common. So Hansen and Singleton (1983) and Hansen and Richard (1987) simply imposed rational expectations and made enough stationarity assumptions to validate the Law of Large Numbers that gives GMM or maximum likelihood estimation good asymptotic properties. Under these assumptions, (1) imposes testable restrictions on the empirical joint distribution of returns and either individual or aggregate consumption. Many economists have specified theories of the stochastic discount factor defined in terms of aggregate or individual consumption, for example, by letting $u(c)$ be a constant relative risk aversion utility function $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$ and defining the stochastic discount factor as the intertemporal marginal rate of substitution

$$m_{t+1} = \frac{\beta u'(c_{t+1})}{u'(c_t)}. \tag{3}$$

[13] In the context of a complete markets economy with a Lucas tree, Sandroni (2000) had argued that a disagreement about tail events would present some consumers with arbitrage opportunities that cannot exist in equilibrium.

The data have reflected poorly on (1) and (3) under the rational expectations assumption that $f = f_i$.[14]

One reaction has been to add backward-looking or forward-looking contributions to time t felicity while retaining rational expectations (see Campbell and Cochrane (1999) and Epstein and Zin (1989), respectively). Another reaction by Hansen and Jagannathan (1991) was to treat the stochastic discount factor m_{t+1} as an unknown nonnegative random variable and to find what observed returns $R_{j,t+1}$ and the restriction

$$1 = \int_{x_{t+1}} m_{t+1}(x^{t+1}) R_{j,t+1}(x_{t+1}) f(x_{t+1}|x^t) dx_{t+1} \tag{4}$$

imply about the first and second moments of admissible stochastic discount factors (with incomplete markets, there exist multiple stochastic discount factors). This research strategy aims to characterize the mean and standard deviation that an empirically successful m must have before specifying a particular theory about the utility function and beliefs that underly that m and that link m to real variables like consumption. Notice how this approach potentially relaxes rational expectations by leaving open the possibility that a successful theory of a stochastic discount factor will assign a nontrivial role to the probability ratio $\frac{f_i(x_{t+1}|x^t)}{f(x_{t+1}|x^t)}$. This likelihood ratio creates a wedge relative to the Euler equation that has usually been fit in the rational expectations macroeconomic tradition originating in Hansen and Singleton (1983) and Mehra and Prescott (1985). Likelihood ratio wedge approaches have been investigated by Bossaerts (2002, 2004) and Hansen (2007), who for the case of logarithmic preferences points out an observational equivalence between belief distortions and a rational expectations implementation of the recursive utility specification of Kreps and Porteus (1978) and Epstein and Zin (1989).

3 Intelligent design in macroeconomics

What I call the intelligent design approach in macroeconomics realizes much of the promise of rational expectations macroeconomics. By solving Pareto problems in which a planner and all agents optimize in light of information and incentive constraints and a common probability model, it is a coherent response to the Lucas (1976) indictment of pre-rational expectations macroeconomic policy design procedures. Lucas alleged that those procedures took private agents' decision rules as invariant with respect to hypothetical government interventions that altered the laws of motion for government policy instruments that impinge on private agents' constraint sets. Via their cross-equation restrictions, rational expectations models automatically make private agents' decision rules be functions of a government policy. At its most ambitious, the intelligent design approach in macroeconomics uses rational expectations

[14] After making a pair of substitutions from a production function and national income identity, macroeconomists have relabeled equation (3) as an IS curve. The poor statistical fit of this keystone equation in the best contemporary macroeconometric models has led macroeconomists to append what are usually less than less than fully interpreted shocks or 'wedges' to (3). This is one reason that macroeconomists have enthusiastically welcomed some of the efforts to explain those wedges mentioned in the text and in section 3.

communism both to process historical data and to design a new and better equilibrium. Thus, a complete implementation of intelligent design involves these steps:

1. Apply rational expectations econometrics to historical data to estimate parameters that describe private agents' preferences, technology, endowments, and information sets.[15]

2. Posit a timing protocol and an objective function for a government, typically a Pareto criterion.

3. Find a new rational expectations equilibrium that maximizes the government's objective.

4. Proclaim as advice the government policy that implements that rational expectations equilibrium.

The intelligent design tradition dates back at least to David Hume:

> Political writers have established it as a maxim, that, in contriving any system of government, and fixing the several checks and controuls of the constitution, every man ought to be supposed a knave, and to have no other end, in all his actions, than private interest. By this interest we must govern him, and by means of it, make him, notwithstanding his insatiable avarice and ambition, co-operate to public good. (Hume 1985, p. 43)

The intelligent design approach in macroeconomics supplements Hume's knave assumption with an assumption that everybody in the model shares common beliefs about probability distributions.

1. *Optimal fiscal and monetary policy cast as dynamic Ramsey problems.* Rational expectations models with exogenously specified taxes and government expenditure processes have been and continue to be fruitful tools for analyzing intertemporal and cross-country differences in flat rate taxes on various types of income and wealth (see Hall (1971)). Such models ask how alternative exogenous dynamic paths for distorting taxes affect the dynamics of prices and quantities in competitive equilibria. Because they associate a distinct competitive equilibrium with each budget-feasible tax policy, these models have been useful tools in the hands of researchers who want to study the consequences of observed intertemporal or international patterns of taxes that they take as given (e.g., see Prescott (2002)).

   Given an available set of distorting taxes, it is natural to study what a good history-contingent tax policy would be. Papers in the literature on Ramsey taxation choose a tax policy that yields a best competitive equilibrium as measured by a Pareto criterion. Judd (1985) and Chamley (1986) showed that a flat rate tax on capital should converge to zero. Lucas and Stokey (1983) and Aiyagari et al. (2002) focused on how a government should cope with uncertain future expenditures drawn from a known distribution

[15] There are differences of opinion about how to model the government in the historical data set.

by exchanging more or less state-contingent loans with a representative private agent and how access or lack of access to complete insurance markets should impinge on the stochastic processes of labor taxes and government indebtedness to the public. Zhu (1992) and Chari et al. (1994) added capital to the Lucas and Stokey (1983) model and showed that the expected value of capital taxation should converge to zero. Lucas (1986) used insights from the Ramsey approach to recommend a coherent monetary and fiscal policy.

2. *Dynamic mechanism design.* Papers in the Ramsey taxation literature exploit the assumption that the government and the public share a common probability measure about the aggregate states. Almost all applications of dynamic mechanism design assume even more, namely, that agents share a joint probability distribution over processes for aggregate and individual-specific shocks that are hidden to the planner.

   Dynamic mechanism design is the principal tool that macroeconomists have used to build macroeconomic models that incorporate enforcement and information imperfections. A principal or planner and an agent share knowledge of a complete set of conditional distributions $f(x_t|x^t, a^t, b^t)$ of time t outcomes x_t where a^t are histories of actions by an agent and principal, respectively. For example, in a model of optimal unemployment compensation by Shavell and Weiss (1979) and Hopenhayn and Nicolini (1997), a_t is the time t search effort of an unemployed worker and b_t is a time t transfer from an insurance agency to the worker. While the principal and the agent both know these distributions, they observe different information. The insurance agency does not observe the worker's search effort and compensates for its information disadvantage by balancing its desire to offer insurance and its need to provide incentives that will solicit information. That has the consequence that it is optimal for the agency to make payments to unemployed workers that diminish as a worker's duration of unemployment grows.[16]

   It is important that the principal and the agents share an accurate description of the distribution of outcomes for choices that will not be observed along the equilibrium path. The principal offers payoffs off an equilibrium path that induce agents to choose to stay along an equilibrium path.

   A key step in shaping compelling applied work along these lines involves specifying an empirically plausible probability distribution with which the theory will confront the principal and the agents. For examples of ambitious applications, I recommend the analysis of European programs for unemployed workers by Pavoni and Violante (2007) and the analysis of complicated mortgage contracts by Piskorksi and Tchistyi (2006).[17]

[16] In light of the outcomes of the Shavell and Weiss (1979) and Hopenhayn and Nicolini (1997) analysis, it is interesting that Sawhill (1995, p. 6) used the principle that "welfare should be time limited and conditional on behavior" as an example of her thesis that politics and not economics sets the agenda for economic discussions and the activities for advisors. She saw the 1990s pressure for welfare reform as emerging not from economics but from political vote-gathering calculations, leading her to ask "Why, then, has policy veered off in directions not supported by the available research?" (Sawhill 1995, p. 8).

[17] Also see Fuchs and Lippi (2006) for an interesting theoretical analysis of when currency unions break up. Their model has the interesting feature that states can be reached in which participation constraints of both countries bind, a feature that could not happen in the earlier models with lack of commitment that

3. *Dynamic Mirlees taxation.* A new dynamic public finance literature applies the mechanism design approach partly in an effort to refine results that had been attained in the Ramsey literature mentioned above.[18] This literature gets interesting results from the assumption that the planner and agents share a common joint distribution over aggregate variables and a complete panel of individuals' skills, but that each individual alone observes the history of evolution of his own skills. The literature focuses on a tradeoff between efficiency and incentives. The planner needs to provide agents (who are better informed than the planner) with incentives to reveal information about their skills or effort. The agents' informational advantage imposes constraints on the allocations that the planner can implement relative to ones he could achieve if he had more information. As a way to highlight the distortions in marginal conditions attributable to the planner's limited information, various papers in this literature focus on what they call "wedges" that they construct by comparing marginal conditions for planning problems with and without an information disadvantage for the planner. For example, a common condition that emerges without an information advantage is the consumption Euler equation for person i:

$$1 = E\Big[R_{t+1}\frac{\beta u'(c_{i,t+1})}{u'(c_{i,t})}\Big|I_t\Big]. \tag{5}$$

With private information about skills or effort, the corresponding marginal condition becomes (see Rogerson (1985) and Kocherlakota (2005))[19]

$$1 = E\Big[\frac{u'(c_{it})}{\beta R_{t+1}u'(c_{i,t+1})}\Big|I_t\Big], \tag{6}$$

which in general differs from equation (5). One reason that this difference is interesting is that (5) is a key ingredient of the Judd-Chamley argument that the tax rate on capital should converge to zero. The difference between equations (6) and (5) reflects the planner's incentives to deter deviations by exploiting the fact that savings affect incentives to work. Another reason it is potentially interesting is that equation (5), which is the foundation of the new Keynesian IS curve, has empirical inadequacies, so macroeconomists want workable models of "wedges" (a.k.a. Euler equation errors) (see Kocherlakota and Pistaferri (2004, 2007)).

Models that impute rich cross dynamics among aggregate variables and individual skills that are hidden from the planner allow the tradeoffs between efficiency and incentives to evolve over time.[20] This feature has interesting implications. For example, Werning (2005) shows how the planner's knowing the correlation between aggregate outcomes and the distribution of skills causes him to move allocations over time that

they cite.

[18] See for example Golosov et al. (2003), Kocherlakota (2005), and for a useful introduction Golosov et al. (2007).

[19] This formula is correct when there is no aggregate risk, but must be modified when there is aggregate risk. See Kocherlakota (2005). Ljungqvist and Sargent (2004, ch. 1) is a history of post WWII macroeconomics in terms of equations (5) and (6).

[20] For this reason, empirical research along the lines described by Cunha and Heckman (2007) should assist builders of the new dynamic fiscal policy.

make marginal wedges behave differently than would be dictated by the tax-smoothing conditions found by Lucas and Stokey (1983).

4. *Design of monetary policy in sticky price models.* A vast literature summarized and extended by Clarida et al. (1999) and Woodford (2003) uses dynamic macroeconomic models with sticky prices to design monetary policy rules by solving Ramsey plans and finding attractive ways to represent and implement them. This work has caught on at central banks and forms the technical language in which arguments about procedures like inflation targeting are framed and discussed.

4 Rational expectations econometrics

A rational expectations equilibrium is a joint probability distribution $f(x^t|\theta)$ over histories x^t indexed by free parameters $\theta \in \Theta$ that describe preferences, technologies, endowments, and information. Rational expectations econometrics tells an econometrician who is outside the model how to learn the parameter θ that pins down that joint distribution. The econometrician knows only a parametric form for the model and therefore initially knows less than the agents about the equilibrium joint probability distribution that nature and the agents inside the model share. The econometrician's tools for learning the parameter vector θ are (1) knowledge of a likelihood function that expresses the equilibrium distribution as a function of a set of unknown parameters, (2) a time series or panel of observations drawn from the equilibrium distribution, and (3) a Law of Large Numbers, a Central Limit Theorem, and some large deviations theorems that allow him to characterize the convergence, rate of convergence, and tail behavior of his estimates. Thus, though the econometrician starts outside the commune sharing the same model, with enough data and a correct likelihood function, he can approach it.

Another name for a rational expectations equilibrium evaluated at a particular history is a *likelihood function*

$$L(\theta|x^t) = f(x^t|\theta) = f(x_t|x^{t-1};\theta)f(x_{t-1}|x^{t-2};\theta)\cdots f(x_1|x_0;\theta)f(x_0|\theta). \tag{7}$$

The factorization of a likelihood on the right side is convenient for estimation. The factorization displays the restrictions that a rational expectations model imposes on a vector autoregression, the conditional density $f(x_t|x^{t-1};\theta)$ being a (possibly nonlinear) vector autoregression.

The most ambitious branch of rational expectations econometrics recommends maximizing this likelihood function or combining it with a Bayesian prior $p(\theta)$ to construct a posterior $p(\theta|x^t)$.[21] In choosing θ to maximize a likelihood function, a rational expectations econometrician searches for a system of expectations that prompts the forward-looking artificial agents inside the model to make decisions that best fit the data.[22] Factor the log likelihood function for a rational expectations model as

$$\log L(x^t|\theta) = \ell(x^t|x^{t-1};\theta) + \ell(x_{t-1}|x^{t-2};\theta) + \cdots \ell(x_1|x_0;\theta) + \ell(x_0|\theta) \tag{8}$$

[21] For early applications of this empirical approach, see Sargent (1977), Sargent (1979), Hansen and Sargent (1980), Taylor (1980), and Dagli and Taylor (1984).

[22] As the econometrician searches over probability measures indexed by θ, he imputes to the agents inside the system of expectations implied by the θ under consideration.

where $\ell(x_t|x^{t-1};\theta) = \log f(x_t|x^{t-1};\theta)$ and define the score function as $s_t(\theta) = \frac{\partial \ell(x_t|x^{t-1};\theta)}{\partial \theta}$. In population, the first-order conditions for maximum likelihood estimation are the conditional moment conditions

$$E[s_t|x^{t-1}] = 0, \tag{9}$$

which imply that the score is a martingale difference sequence, a fact that is the starting point for a useful theory of statistical inference. By replacing the mathematical expectation E in equation (9) and the related moment equation (10) below with a sample average $T^{-1}\sum_{t=1}^{T}$, the econometrician finds a θ that allows him to approximate the equilibrium distribution better as $T \to +\infty$. The absence of free parameters that characterize decision makers' beliefs underlies the cross-equation restrictions that are the econometric hall marks and the source of information for identifying parameters in rational expectations models.[23]

Rational expectations econometrics has solid accomplishments to its credit. By reorienting econometric attention away from parameterized decision rules to deeper parameters characterizing preferences, technology, endowments, and information, it offered a compelling response to the Lucas (1976) critique. It provided ways of thinking about identification that refined our understanding of the exclusion restrictions that were the mainstay of the Cowles commission methods that underlay Keynesian macroeconometric models. It contributed a way to interpret error terms in econometric equations as discrepancies that reflect an information advantage of the agents inside a model (e.g., see Hansen and Sargent (1980)). Thinking about error terms in this way naturally led researchers to treat the serial correlation properties of those errors as nuisance parameters, then pursue a semi-parametric estimation strategy.[24] These and other limited information estimation strategies back off from the communism stressed above because the econometrician is not presumed fully to share with agents the entire joint probability distribution over histories, including variables not observed by the econometrician.

Generalized Methods of Moments (GMM) estimation also retreats marginally from communism about beliefs (see Hansen (1982)). It estimates a subset of the model parameters, say, $\theta_1 \in \Theta_1 \subset \Theta$ by using conditional moment conditions

$$E[h(x_t|z^t, \theta_1)] = 0 \tag{10}$$

where $z^t \subset x^t$ is a time t information set observed by the econometrician, which is possibly a subset of the information observed by agents in the model. Euler equations for some of the agents in the model are popular candidates for constructing a function h. GMM estimation is less ambitious, or depending on your point of view, less pretentious, than maximum likelihood estimation because it does not aspire to estimate a complete model. Knowing a subset of parameters θ_1 is typically insufficient to determine a rational expectations equilibrium. But in using GMM, the econometrician still exploits a milder brand of communism of beliefs because he heavily exploits the assumption that the conditional distribution that he uses to estimate (10) equals the conditional distribution that the agents inside the model believe.

[23]See Imrohoroglu (1993) for a model that is an exception to the letter but not the spirit of the statement in the text. Cross-equation restrictions allow Imrohoroglu to use maximum likelihood estimation to pin down parameters including one that indexes a continuum of sunspot equilibria. Imrohoroglu usefully distinguishes econometric identification, which prevails, from uniqueness of equilibrium, which does not.

[24]See Hansen and Sargent (2009, ch. 10).

4.1 Indirect inference

The Lucas (1976) critique advocated that researchers not use good fitting non-structural models for policy analysis. But lately we have learned that purely statistical models can be useful tools for extracting good estimates of the parameters of some intractable structural rational expectations models. Indirect inference methods tell a researcher how to use the first-order conditions for estimating the parameters of a good fitting but purely statistical model to make good inferences about parameters of a structural economic model.

A rational expectations modeler juggles two statistical models when he uses indirect inference methods. Indirect inference assumes that a researcher wants to estimate parameters of a rational expectations model of interest for which (1) analytical difficulties prevent him from being able directly to evaluate a likelihood function for a model of interest $f(x^t|\rho)$, and (2) computational methods allow him to simulate time series from $f(x^t|\rho)$ at given parameter values ρ. See Gourieroux et al. (1993), Smith (1993), and Gallant and Tauchen (1996). An indirect inference estimator carries along two models, the model of interest with the untractable likelihood function, and an auxiliary model with a tractable likelihood function that approximates the historical data well. The two models have different sets of parameters, the parameters of the economist's model ρ being interpretable in terms of preferences, technologies, information sets, and government policies, the parameters θ of the auxiliary model $f(x^t|\theta)$ being data fitting devices. The idea of Gallant and Tauchen (1996) is first to estimate the auxiliary model by maximum likelihood, then to use the score functions for the auxiliary model and the first-order conditions in equation (9) to define a criterion for a GMM estimator that can be used in conjunction with simulations of the economic model to estimate the parameters ρ of the economist's model. Thus, let the auxiliary model have log likelihood function given by equation (8) and for the data sample in hand, compute the maximum likelihood estimate $\hat{\theta}$ for the auxiliary model. Now for a given artificial data set $\{x_t(\rho)\}$ from the *economic* model, think of forming the score function for the *auxiliary* model $s_t(x_t(\rho)|x^{t-1}(\rho), \hat{\theta})$ for each t for a given sample, evaluated at the maximum likelihood estimate $\hat{\theta}$ of the parameters of the auxiliary model. Simulate paths of $x_\tau(\rho)$ for $\tau = 1, \ldots, N$ from the economic model. Gallant and Tauchen estimate ρ by setting the average score[25]

$$\frac{1}{N} \sum_{\tau=1}^{N} s_\tau(x_\tau(\rho)|x^{\tau-1}(\rho), \hat{\theta}) \tag{11}$$

as close to zero as possible when measured by a quadratic form of the type used in GMM. If the auxiliary model fits well, this method gives good estimates of the parameters ρ of the economic model. Technically, the indirect estimator is as efficient as maximum likelihood when the economic and auxiliary models are observationally equivalent.

This ideal case raises the following question: what would happen if macroeconomic policy makers were incorrectly to use what from nature's point of view is actually an auxiliary model? Historical data can give the government no indication that it should abandon its model. Nevertheless, it could be making major policy design mistakes because its misunderstands the consequences of policies that it has not chosen.[26] The possibility that the

[25]This description fits their Case 2.

[26]See Lucas (1976), Sargent (1999, ch. 7), and Fudenberg and Levine (2007).

government uses what, unbeknownst to it, is just an auxiliary model, not a structural one, sets the stage for the adaptive models and the self-confirming equilibria to which they converge to be described in section 6 and 7.

5 Issues that rational expectations equilibria sidestep

As with all good theoretical tools, the rational expectations assumption sharpens the mind by narrowing it. It puts some potentially interesting issues aside in order to focus on others. The communism of beliefs that gives rational expectations models the discipline to make the sharp and useful statements mentioned above also limits their use as tools for helping us to understand problems that require that multiple models be on the table simultaneously, either because different agents have different models or because individual agents express their doubts about model specification by contemplating multiple models.

5.1 Misspecification

Macroeconomists routinely describe their models as approximations, but this rarely affects their formal analyses of them.[27] The coherence attained by equating objective and subjective probability distributions means that it is impossible to carry out a self-contained analysis of model misspecification within a rational expectations model. This observation has ramifications for empirical work that are yet to be fully understood. It often occurs that subjecting a rational expectations model to a statistical specification test makes an author doubt his model. How should the modeler react to adverse evidence about a most important aspect of the communism of beliefs he hoped for in constructing the model, namely, that the model's implied joint probability over histories actually generates the data?

An influential group of rational expectations modelers in macroeconomics refuses to use estimators based on a likelihood function. Another way of expressing this is to say that they refuse to take all of the probability implications of their model seriously. A frequently pronounced excuse for "calibrating" parameters instead of using the likelihood function is that because the model is an approximation (i.e., it is misspecified), justifications for applying likelihood methods to all of the data available do not apply.[28] But if the modeler believes this, might it be appropriate to share his specification doubts with the agents inside his model?[29]

The intricate dependence of outcomes of intelligent design analyses on the details of the stochastic specification has led some researchers to modify the rational expectations assumption by adding some doubts about stochastic specification to at least some of the agents. See Woodford (2005), Bergemann and Morris (2005), Kocherlakota and Phelan (2007), and Hansen and Sargent (2008, chs. 15 and 16) for some examples in which principals

[27]Hansen and Jagannathan (1997) is an interesting exception that constructs a measure of misspecification.

[28]The most formal justifications for this stance are by Sims (1993) and Hansen and Sargent (1993), who show how to filter data to manipulate an appropriate Kullback-Leibler statistical approximation criterion with the aim of improving estimates of preference and technology parameters for misspecified models. With correctly specified models, using filtered data only degrades estimates.

[29]See Hansen (2007) and Hansen and Sargent (2008) for a research agenda that suggests that the answer to this question is yes.

or agents or both express those doubts by simultaneously entertaining multiple probability models and seeking allocations that are robust to those doubts.

5.2 The law of unintended consequences

In discussing the consequences of the Roosevelt administration's silver purchase program, Friedman (1991) alluded to a "law of unintended consequences." Friedman told how a policy originally designed to reward senators from western silver-producing states inadvertently transmitted deflation to China, drove it off the silver standard that had previously insulated it from the world wide deflation and depression, and unleashed a fiat money hyperinflation that facilitated an ultimate Communist takeover.[30] Notice that the common beliefs assumption used in the intelligent design approach excludes a law of unintended consequences.

5.3 Learning about an equilibrium

If they did not know them in the beginning, it could take a long time for people to learn about some features of the distributions that they are presumed to know in a rational expectations equilibrium. The research program described by Hansen (2007) studies how difficult-to-learn-about but easy-to-care about long-run risks impinge on asset prices. Those long-run risks are easy to care about when consumers have Kreps and Porteus (1978) preferences. They are difficult to learn about because they are subtle low-frequency features.

5.4 Disparate macroeconomic theories

Rational expectations models' communism of beliefs prevents the models from containing macroeconomic policy controversies inspired by competing models.[31] During my life in economics, distinguished macroeconomists have passionately disagreed about the macroeconomic consequences of hypothetical policy experiments, for example, whether unemployment could be reduced and output growth raised by exploiting a 'cruel choice' between inflation and unemployment, how costly in terms of output and inflation it would be to suppress inflation, and what gains supply side economics truly offered. These controversies pitted different models against each other. Of course, you can pit *alternative* rational expectations models against one another, but within each model, all agents share beliefs with each other and with nature, and they do not comprehend the multiplicity of models present in controversies among macroeconomists.[32]

As we shall see in the next section, theories of learning in games and macroeconomics tell us that an adaptive process that gropes for better models and better forecasts can stall

[30]Friedman is thus reaffirming the earlier account by Friedman and Schwartz (1963).

[31]In principle, a thoroughgoing incorporation of Bayesian model averaging within a rational expectations equilibrium could capture such policy disputes. The imperialistic Bayesian model averaging procedure would produce a unique model that the agents inside the model use to construct forecasts. This approach has rarely been used so far as I know. An incomplete approach to such a model is contained in Cogley and Sargent (2005).

[32]This observation indicates a sense in which the eclectic model of Cogley and Sargent (2005) to be discussed in subsection 7.5 lacks coherence and why for a rational expectations theorist, 'eclectic' is a code word for 'incoherent'.

short of a rational expectations equilibrium and can leave economic advisors with incorrect models.

6 Learning in games and macroeconomics

One main point of this section is that adding even small doses of disagreement and adaptation can sometimes substantially improve fits relative to rational expectations models. Another point is to indicate how macroeconomic models with adaptive agents converge to a self-confirming equilibrium, a situation that can leave a government with mistaken views about the consequences of policy experiments that could improve outcomes. This is interesting because it warns us against automatically expecting outcomes from a highly evolved system to be optimal.

Rational expectations models are not theories about how agents inside a model *form* beliefs that equal an objective distribution. At most, they describe the *outcome* of an unspecified learning process that has settled down and endowed everyone with beliefs that agree with nature's.[33] The learning literature aspires to construct what Bray and Kreps (1987) call theories of learning *about* (as opposed to *within*) a rational expectations (or Nash) equilibrium. By saying *about* and not *within*, Bray and Kreps emphasize that the challenge is to model how a system of agents can come to learn the objective distribution by using adaptive algorithms that do not simply apply Bayes' law to a correct probability model.[34] Why can't we just appeal to the same Law of Large Numbers that fulfills the purposes of a rational expectations econometrician who is learning *about* an equilibrium? The reason is that a rational expectations econometrician is *outside* the model and his learning is a side-show that does not affect the data generating mechanism. It is different when the people learning about an equilibrium are inside the model. Their learning affects decisions and alters the distribution of endogenous variables over time, making other adaptive learners aim at moving targets. This feature of the learning problem makes it substantially more difficult than the problem of proving consistency of parameter estimates in stationary statistical settings.

6.1 Learning in games

In a game, a Nash equilibrium is the natural counterpart of a rational expectations equilibrium or a recursive competitive equilibrium. An extensive literature studies whether a system of adaptive players converges to a Nash equilibrium. A range of plausible adaptive

[33] A difficult challenge in the machine learning literature is to construct an adaptive algorithm that learns dynamic programming. For a recent significant advance based on the application of the adjoint of a resolvent operator and a law of large numbers, see Meyn (2007, ch. 11).

[34] I see Bray and Kreps's 'about' versus 'within' tension running through the literature that attempts to set down Bayesian theories of convergence to Nash equilibria. Marimon (1997) said that a Bayesian knows the truth from the beginning. Young (2004) pointed out that the absolute continuity assumption underlying the beautiful convergence result of Kalai and Lehrer (1993, 1994) requires that players have substantial prior knowledge of their opponents' strategies. Young is skeptical that Kalai and Lehrer have provided a compelling answer to the key question of whether "... can one identify priors [over opponents strategies] whose support is wide enough to capture the strategies that one's (rational) opponents are actually using, without assuming away the uncertainty inherent in the situation?" Young (2004, p. 95)

algorithms have been proposed that are differentiated by how much foresight and theorizing they attribute to the players.[35] At one extreme are adaptive models that have naive players who ignore strategic interactions and either play against histograms of their opponents past actions (this is called fictitious play) or alter their moves in directions that *ex post* reduce their *regret* at not having taken other actions in the past, given their opponents' histories of actions. At the other extreme are models in which players construct statistical theories about their opponents' behavior, use them for a while to make forward-looking decisions, occasionally subject their theories to hypothesis tests, discard rejected ones and choose new specifications.

This literature has sought plausible and robust algorithms that converge to a Nash equilibrium. Hart and Mas-Colell tell us that this is a tall order:

> It is notoriously difficult to formulate sensible adaptive dynamics that guarantee convergence to Nash equilibrium. In fact, short of variants of exhaustive search (deterministic or stochastic), there are no general results. Hart and Mas-Colell (2003, p. 1830)

Hart and Mas-Colell and Foster and Vohra (1999) show that the source of the difficulty is that most adaptive schemes specify that adjustments in a player's strategy do not depend on the payoff functions of other players, an uncoupling of the dynamics that in general dooms the system not to converge to a Nash equilibrium. Many examples of the adaptive schemes in the literature are uncoupled. Because many game theorists find uncoupled schemes desirable, parts of the literature have lowered the bar by looking for convergence to something weaker than Nash equilibria, namely, correlated equilibria or coarse correlated equilibria. Hart and Mas-Colell (2003, p. 1834) make the telling remark that "It is thus interesting that Nash equilibrium, a notion that does not predicate coordinated behavior, cannot be guaranteed to be reached in an uncoupled way, while correlated equilibrium, a notion based on coordination, can."[36]

Hart and Mas-Colell (2000, 2001, 2003) study adaptive schemes that are backward looking. For example, some of the most interesting ones have a player construct counterfactual historical payoffs that he would have received had he played other strategies, then compute a measure of regret, then adjust his future play in directions that would have minimized his regret. These schemes impute little or no theorizing and foresight to the players.

For my present purposes, one of the most interesting contributions comes from part of the literature that attributes more sophistication to players, in particular, the work of Foster and Young (2003), which is also summarized in Young (2004, ch. 8).[37] Their model has

[35]For a critical survey of this literature, see Young (2004).

[36]Experimental economics has supplied data sets designed to check ideas from the literature on adaptive learning in games. It is remarkable that laboratory experiments using macroeconomics are rarer than those using microeconomics. See Duffy (2006) for an account of the existing experiments. I suspect that the main reason for fewer experiments in macro than in micro is that the choices confronting artificial agents within even one of the simpler recursive competitive equilibria used in macroeconomics are very complicated relative to the settings that experimentalists usually confront their subjects with.

[37]For a distinct but related approach, see Jehiel (1995, 1998). The Foster and Young (2003) model seems to me to capture some of the flavor of the anticipated utility framework advocated by Kreps (1998). The classifier models in Marimon et al. (1990) have a similar flavor.

the following components: (1) each player has a large set of potential models that describe his opponents' strategies; (2) players use a random device to select a particular model; (3) after that model is selected, there is an 'act and collect data' period during which a player (incorrectly) assumes that he will believe his current model forever; during this period, each player chooses his actions via a smoothed best response to what his model tells him about opponents' actions (e.g., a quantal response function); (4) after a data collection period, a player compares the empirical pattern of his opponents' play with that predicted by his model. He performs an hypothesis test that compares the theoretical and empirical distributions. If he rejects his current model, he randomly draws a new model from his set of models, then returns to step 2. If he accepts the model, he returns to step 3, waits a random number of periods, and then begins another data collection period.

With suitable assumptions about the lengths of testing periods and the tolerances of the hypothesis tests, Foster and Young (2003) show that behaviors eventually emerge that are often close to Nash equilibria. Their notion of hypothesis tests is sufficiently broad to include many plausible procedures. Their convergence result seems to be an uncoupled multi-agent learning scheme that actually approaches Nash equilibria, not something weaker like the coarse correlated equilibrium that the entirely backward-looking schemes mentioned above can approach. They avoid the conundrum of Hart and Mas-Colell partly by weakening the notion of convergence.

Besides admiring their convergence result, as a macroeconomist I am attracted to the Foster and Young (2003) setup because I like the parable about conventional wisdom followed by hypothesis testing that their model captures. I will use that parable in section 7 when I describe some macroeconomic history. Before I do that, I shall describe a macroeconomics literature about adaptive learning that adopts a vision like Foster and Young's.

6.2 Learning rational expectations equilibria

A literature on least squares learning in self-referential systems studies whether a system of agents who use recursive least squares algorithms to update their temporary models and forward looking decision algorithms based on those temporary models will converge to a rational expectations equilibrium (see Marcet and Sargent (1989a), Evans and Honkapohja (1999, 2001), Woodford (1990), and Fudenberg and Levine (1998)). These models have the following structure: (1) one or more decision makers take actions at time t by solving a dynamic programming problem based on time t econometric estimates of a possibly misspecified time t model, under the false assumption used to formulate the dynamic programming problem that decision makers will retain their currently estimated models forever; (2) the actions of some of those decision makers influence the data-generating process; and (3) the decision makers update estimates of their models each period using some version of least squares. The literature studies limiting properties of this system. A main finding is that convergence to a rational expectations equilibrium does not occur in general but that convergence to a *self-confirming* equilibrium often does. In a self-confirming equilibrium, the agents inside the model share common beliefs about probabilities of events that occur infinitely often but can have differing beliefs about events that occur less often. There are insufficient observations off the equilibrium path to allow a Law of Large Numbers to eradicate disagreements.[38] Stochastic approximation and large deviations theory have been used to characterize how the system produces probability distributions that approach and occasionally deviate from self-confirming equilibria (see Fudenberg and Levine (1993, 2007), Kreps (1998), Sargent (1999), Cho et al. (2002), and Williams (2004)).

6.2.1 Self-confirming equilibrium

It is convenient to partition the observables $x_t = y_t \quad v_t$ where v_t are some decisions taken by a government and y_t are all other observables. There are a true data generating process and an approximating model, respectively,

$$f(y^\infty, v^\infty|\rho) \text{ and } f(y^\infty, v^\infty|\theta). \tag{12}$$

An agent has preferences ordered by

$$\int U(y^\infty, v^\infty) f(y^\infty, v^\infty|\theta) d(y^\infty, v^\infty) \tag{13}$$

and chooses a history-dependent plan

$$v_t = h(y^t|\theta) \tag{14}$$

that maximizes (13) gives rise to the sequence of decisions $v(h|\theta)^\infty$. We call maximizing (13) a Phelps problem in honor of a particular version of a government control problem of this type that was solved by Phelps (1967) and that will play an important role in the models to be discussed in subsection 7.2.

Definition 6.1. *A self-confirming equilibrium (SCE) is a parameter vector θ_o for the approximating model that satisfies the data-matching conditions*

$$f(y^\infty, v(h|\theta_o)^\infty|\theta_o) = f(y^\infty, v(h|\theta_o)^\infty|\rho). \tag{15}$$

An SCE builds in (1) optimization of (13) given beliefs indexed by θ_o, and (2) a choice of $\theta = \theta_o$ that satisfies the data matching conditions (15). Data matching prevails for events that occur under the equilibrium policy $v(h|\theta)^\infty$, but it is possible that

$$f(y^\infty, v^\infty|\theta_o) \neq f(y^\infty, v^\infty) \tag{16}$$

for $v^\infty \neq v(h|\theta)^\infty$. In an SCE, the approximating model is observationally equivalent with the true model for events that occur under the policy implied by equilibrium decisions, but not necessarily under other policies.

[38] Many examples exist in which a system of least squares learners does converge to a rational expectations equilibrium. However, the model builder in those examples prompts the agents in directions that facilitate complete learning, typically by endowing him with parameterized functional forms that can support a rational expectations equilibrium as a limit point and by restricting what has to be learned to the parameters of those functions. Marcet and Sargent (1989a) and Evans and Honkapohja (1999, 2001) contain such examples.

6.2.2 Learning converges to an SCE

What makes an SCE especially interesting is its role as a limit point of an adaptive system. We suppose that an adaptive learner begins with an initial estimate $\hat{\theta}_0$ at time 0 and uses a recursive least squares learning algorithm

$$\hat{\theta}_{t+1} - \hat{\theta}_t = a(\hat{\theta}_t, y^t, v^t, t). \tag{17}$$

As in the models of learning in games of Foster and Young (2003) and Young (2004, ch. 8), we assume that time t decision makers mistakenly regard their time t model indexed by $\hat{\theta}_t$ as permanent and form the sequence of decisions[39]

$$\hat{v}(h)_t = h(y|\hat{\theta}_t) \tag{18}$$

where $h(\cdot|\theta)$ is the same function (14) that solves the original Phelps problem (13) under the model $f(y^\infty, v^\infty|\theta)$. Under this scheme for making decisions, the joint density of $(y^\infty, v^\infty, \hat{\theta}^\infty)$ is

$$f(y^\infty, \hat{v}(h)^\infty, \hat{\theta}^\infty|\rho). \tag{19}$$

The learning literature studies the limiting behavior of (19) and imposes restrictions on the estimator a and the densities $f(\cdot|\theta)$ and $f(\cdot|\rho)$ that imply that

$$\hat{\theta}_t \to \theta_o, \tag{20}$$

where the sense of convergence can be either almost surely or in distribution, depending on details of the estimator a in (17).[40]

6.2.3 REE or SCE?

Sometimes researchers have specified the approximating model to equal the true one, meaning that there exists a value θ for which $f(y^\infty, v^\infty|\rho) = f(y^\infty, v^\infty|\theta_o)$ for *all* plans v^∞, not just equilibrium ones. This assumption underlies findings like those of Woodford (1990) and Marcet and Sargent (1989b) in which least squares learning schemes converge to rational expectations equilibria. When $f(y^\infty, v^\infty|\rho) \neq f(y^\infty, v^\infty|\theta_o)$ for some choices of v, the best that can be hoped for is convergence to an SCE.

[39] Cho and Kasa (2006) create a model structure closer to the vision of Foster and Young (2003). In particular, their model has the following structure: (1) one or more decision makers take actions at time t by solving a dynamic programming problem based on a possibly misspecified time t model, (2) the actions of some of those decision makers influence the data-generating process; (3) the decision maker shows that he is aware of the possible misspecification of his model by trying to detect misspecifications with an econometric specification test, (4) if the specification test rejects the model, the decision maker selects an improved model, while (5) if the current model is not rejected, the decision maker formulates policy using the model under the assumption (used to formulate the dynamic programming problem) that he will retain this model forever. Cho and Kasa define useful mathematical senses in which the same stochastic approximation and large deviations results that pertain to a least-squares learning setup also describe the outcomes of their model-validation setup.

[40] For example, so-called 'constant gain' algorithms give rise to convergence in distribution, while estimators whose gains diminish at the proper rates converge almost surely. For example, see Williams (2004). A few papers have studied rates of convergence. There are examples in which convergence occurs at a $\sqrt{T}$ rate, but also examples where convergence occurs markedly more slowly.

Sargent (1999, ch. 6) works with a weaker notion of an SCE thatBranch and Evans (2005, 2006) call a misspecification equilibrium. Branch and Evans construct misspecification equilibria in which agents i and j have different models parameterized, say, by θ_i and θ_j, and in which $f(y^t|\theta_i) \neq f(y^t|\theta_j) \neq f(y^t|\rho)$, where again ρ parameterizes the data generating mechanism. A misspecification equilibrium imposes moment conditions on the approximating models that imply parameters θ_i that give equal minimum mean square error forecast errors $E_{\theta_j}(y_{t+1} - E_{\theta_j}(y_{t+1}|y^t)^2)$ for all surviving models. Branch and Evans use this setup to model equilibria in which beliefs and forecasts are heterogeneous across agents. They provide conditions under which recursive least squares learning algorithms converge to a subset of the possible misspecification equilibria.[41]

6.3 Lessons for macroeconomic policy analysis

6.3.1 Cross-equation restrictions and descriptions of policy variables

Rational expectations econometrics identifies parameters partly from the cross-equation restrictions between outcomes and government policy rules embedded in the joint theoretical density $f(y^\infty, v(h|\theta_o)^\infty|\theta_o)$ that appears in (15). Having a good statistical description of the policy process v^t is important for acquiring good estimates of θ_o or ρ. When the historical government policies are better described by an adaptive process like (19), it affects cross-equation restrictions on θ_o or ρ.

6.3.2 SCE-REE gaps and policy analysis

Why is a gap between a rational expectations equilibrium and a self-confirming equilibrium important for a macroeconomist? Macroeconomists build models with many small players and a small number (often one) of large players called governments. The small players are the private agents who can take aggregate laws of motion as given within a recursive competitive equilibrium. It is sufficient for them that their views are correct along the equilibrium path. If a small agent has access to a long enough history of observations drawn from a self-confirming equilibrium, he can form unimprovable forecasts by simply taking appropriate (conditional) averages of past outcomes. It doesn't matter to a small agent that his views may be incorrect views off the equilibrium path.

But it can matter very much when a government, a large agent, has incorrect views off the equilibrium path because in designing its policy, we suppose that a government solves a Ramsey problem in which it contemplates the consequences of off-equilibrium path experiments. Wrong views about off-equilibrium path events shape government policy and the equilibrium path. Self-confirming equilibria leave ample room for mistaken policies, unintended consequences, disagreements about macroeconomic theories, and issues about the value of social experimentation.

[41] I view the models of Brock and Hommes (1997) and Brock and de Fontnouvelle (2000) as early versions of misspecification equilibria.

7 Learning good monetary policy

The non mathematical appendix A describes a 700 year process of theorizing and experimenting that transformed the prevailing commodity money system from one with many nominal anchors – mint-point, melt-point pairs for coins of all denominations – to a system that retained gold points for only one standard coin and used government-issued convertible token coins and notes for other denominations. After another 100 years, governments abolished the gold points for the standard coin too, leaving monetary authorities to search for a nominal anchor based on their good intentions and their knowledge of the quantity theory of money. The appendix notes how a commodity money system concealed the quantity theory of money. The purpose of the gold and silver points was to make the price level a low variance, small trend exogenous variable and the money supply into a low variance, small trend endogenous variable. Eventually, some policy mistakes revealed the quantity theory to empiricists.[42]

Friedman (1991, pp. 249-252) noted how our present fiat money system is historically unprecedented and cited with approval the observation of Fisher (1926, p.131) that "Irredeemable paper money has almost invariably proved a curse to the country employing it". Friedman highlighted two major obstacles that obstruct the path to the a well managed fiat currency, namely, political pressures to use fiat money to finance the government and the temptation to exploit a Phillips curve (Friedman (1991, p. 207)). Empirical learning models have been used to interpret outcomes after monetary authorities have yielded to one of these pressures.

7.1 The temptation to finance deficits with the printing press

Empirical studies of high inflation episodes fueled by large money-financed deficits have formed a useful laboratory for both rational expectations models and models of adaptation. A model that combines a rational expectations version of a Cagan (1956) money demand function with a government budget constraint with an exogenous monetized government deficit as a driving force has been a workhorse for understanding money and prices of governments during big inflations. That model has a continuum of equilibria, something that must be taken into account in empirical applications (see Imrohoroglu (1993)). The rational expectations dynamics are perverse in the sense that stable equilibria are on the wrong side of the Laffer curve and have comparative dynamics that associate higher deficits with lower inflation. In response to that difficulty, Bruno and Fischer (1990) and Marcet and Sargent (1989b) showed how imputing learning dynamics to holders of money leads to dynamic systems that converge to rational expectations equilibria that are on the good side of the Laffer curve. These stable-under-learning equilibria support the old time religion that associates larger deficits with *higher* inflation. Marcet and Nicolini (2003) and Sargent et al. (2006a) show how adaptive dynamics that allow temporary escapes from the domain of attraction of those good-side-of-the-Laffer curve equilibria can improve the fit of these models relative to those attained by purely rational expectations versions. They do so by modestly modifying rational expectations equilibria to allow escape dynamics occasionally to liberate agents'

[42]Fetter (1978, p. 16) and Friedman (1991, pp. 150-151) discuss how concerns about small denomination coins shaped the gold standard.

expectations from the monetary and fiscal determinants of inflation that would anchor expectations under rational expectations. In loosening those connections, these papers make contact with some of issues about anchoring inflation raised by Bernanke (2007).

These models are set up so that all of the learning is being done by private agents (i.e., the holders of fiat money) and so that SCE's are also rational expectations equilibria. In the next subsection, we focus on models in which a government is learning.

7.2 Learning inflation-unemployment dynamics

In the remainder of this paper, I reconstruct how the U.S. monetary authorities struggled to learn to manage inflation-unemployment dynamics after the collapse of Bretton Woods. To introduce these stories, I note that today it is widely accepted that a monetary authority can control inflation if it wants. Then why did the U.S. monetary authority allow inflation to rise in the late 1960s and 1970s and why did it choose to bring inflation down in the 1980s and 1990s? If we assume that the monetary authority's purposes did not change, and that it always disliked inflation and unemployment, then it seems natural to suspect alterations over time in the monetary authority's model of inflation-unemployment dynamics. Can a theory of learning (and maybe forgetting) about inflation-unemployment dynamics rationalize the inflation outcomes over which the US monetary authority chose to preside?[43]

Sims (1988) advocated applying the real-time dynamics from an adaptive model to explain the conquest of U.S. inflation in the 1980s. We want also to explain how the Fed allowed inflation to ignite. I'll describe empirical models that highlight temporary discrepancies between the government's model and the true data generating mechanism, a government that struggles to optimize and learn by solving Phelps problems and revising its parameter estimates to bring them into line with the data, and a force driving the government's parameter estimates toward a point where they can't be improved. These models are designed to capture the idea that "... if the central bank and the public learn from experience that high inflation imposes greater costs and fewer benefits than previously thought, then the equilibrium will adjust toward one with lower inflation and lower inflation expectations." Bernanke (2007).

In constructing these models, one cannot avoid playing god any more than one can when constructing a rational expectations model. We must set down a true data generating model, parameterized models for the government and other decision makers inside the model, particular recursive learning algorithms, and initial conditions for agents' beliefs.[44] In selecting the models to impute to nature and the agents, it is natural to use models that had serious adherents at the time. The models to be described below have made different reasonable selections from the specifications described in the 'revisionist history' of the U.S. Phillips curve by King and Watson (1994). King and Watson studied the consequences of specification decisions about econometric directions of fit (i.e., should you regress inflation

[43]For testimony that policy authorities in the U.S. are concerned about related issues, see Bernanke (2007) and Mishkin (2007).

[44]See Evans and Honkapohja (2003), Orphanides and Williams (2005, 2007), and Bullard and Mitra (2007) for applications of models of this type to evaluating the stability and performance of alternative monetary policy rules. See Cogley (2005) and Piazzesi and Schneider (2007) for applications to the yield curve.

on unemployment or unemployment on inflation?[45]) and how they might affect government decisions. Sargent (1999, ch. 7) described how those specification decisions can affect self-confirming equilibrium outcomes.

In interpreting empirical findings that emerge from these adaptive models, it is useful to watch how authors have chosen the specifications of the true data generating model $f(y^\infty, v^\infty|\rho)$ in (12), the government's model $f(y^\infty, v^\infty|\theta)$ in (12), the government's adaptive estimator a in (17), the return function U in the government's Phelps problem (13), and, finally, the initial conditions for the government's beliefs $\hat{\theta}_0$. Primiceri (2006) posits that the government's model equals the true one, but that it does not know the true parameter values, while Sims (1988), Chung (1990), and Sargent (1999) posit that the government's model is distinct from nature's even at the population self-confirming equilibrium parameter values.[46] Thus, beliefs in Primiceri's model converge to a SCE in which the government has correct beliefs about on-equilibrium and off-equilibrium path events. In contrast, the government beliefs in other models converge to a SCE in which the government has the wrong model off equilibrium paths and which it engages in suboptimal policy because it is caught in a lack-of-experimentation trap.

7.3 A classical account

The model used by Sims (1988), Chung (1990), Sargent (1999), Cho et al. (2002), and Sargent et al. (2006b) features a structure in which nature's model is a version of the rational expectations-natural rate model of Lucas (1972b), while the government's model is a non-expectational Phillips curve that, depending on the coefficients in distributed lags, potentially asserts an exploitable Phillips curve. The model has the following components: (1) a time-invariant true data generating model consisting of a Lucas style expectational Phillips curve with rational expectations that implies no exploitable tradeoff between inflation and unemployment; (2) a government model that consists of a Samuelson-Solow type exploitable Phillips curve in which the public's expectations of inflation are not identified as a variable positioning the Phillips curve (which, unbeknownst to the government, it truly does); (3) an adaptive algorithm that the government uses to update its model each period; (4) a government that can set the systematic part of inflation and a Phelps problem that the government solves each period to determine the current period's setting of the systematic part of inflation in light of its latest estimates. (5)a population self-confirming equilibrium that the mean dynamics of the system heads toward and whose outcomes equal those of the time-consistent equilibrium of Kydland and Prescott (1977).[47] The timing protocol in the model is such that if the government had the correct model, it would attain Ramsey outcomes, in the language of Stokey (1989); (6) some escape dynamics that can allow the adaptive system recurrently to visit the Ramsey outcome. These escapes occur when a sequence of unlikely shocks teaches the government a too-strong version of the natural rate hypothesis, despite

[45] To align with various studies in the 1970s, King and Watson call inflation on unemployment the Keynesian direction and unemployment on inflation the classical direction.

[46] It is also useful to note the free parameters describing the government's beliefs that these structure introduce relative to a rational expectations model. For example, in Primiceri's model what are added are the initial conditions for θ, R.

[47] Kydland and Prescott (1977, p. 481) sketched an adaptive algorithm with a curve-fitting government that they said would converge to their time-consistent equilibrium.

its erroneous model.[48] The visits are necessarily temporary because the learning dynamics are bound to reassert themselves and pull outcomes toward the self-confirming equilibrium.

The first implementations of this type of model imputed constant gain algorithms to the government. Simulations of Sims (1988) generated sample paths that seemed promising for explaining a Volcker-like stabilization prompted by the government's being able to learn a good enough version of the natural rate hypothesis. However, formal econometric attempts to implement the model by Chung (1990) and Sargent (1999) failed to fit the U.S. data well, mainly because the government's adaptive algorithm catches on to the adverse shifts in Phillips curve so quickly in the early 1970s that it tells the Phelps problem to tell the government to stabilize inflation much earlier than actually occurred. Sargent et al. (2006b) replaced the constant gain algorithm used in the earlier models with the Bayesian updating procedure implied by a drifting coefficients model with a covariance matrix V for the innovations in the drifts to the coefficients. When they estimated V along with the parameters of nature's model by maximum likelihood, they found that could reverse engineer a drifting set of government beliefs that when put into the Phelps problem each period produces a sequence of first period Phelps policy recommendations that do a good job of matching the actual inflation data. The estimated V makes the intercept in the Fed's quite volatile and thus makes contact with the account of Arthur Burns's Fed, which according to Hetzel (1998), attributed much of the inflation of the 1970s to special factors that are akin to adding dummy variables to regression that capture intercept drift. It should be noted that the maximum likelihood estimate of V is so large that it conveys the image of a government that expects coefficients to drift so much that it is very open to discounting past data heavily in order to fit new observations. The model's conjuring up a Fed that over fits its models to recent data is food for thought for Fed watchers. The synthesized government beliefs succeed in rationalizing inflation *ex post* as a response to these government beliefs, and the beliefs themselves do a good job of forecasting inflation, thus capturing what seems to have been a remarkably good record of inflation forecasting by the Fed (see Bernanke (2007)). But relative to available alternatives, the imputed beliefs do a poor job of forecasting unemployment, a deficiency of the model that hints that the reverse-engineering exercise may be imputing unrealistic views about *joint* inflation-unemployment dynamics to the Phelps problem in order to rationalize observed inflation outcomes.

7.4 A Keynesian account

By making other choices about the models to impute to nature and the government, Primiceri (2006) modifies the above structure in a way that enables him to tell a very different story about the post WWII U.S. inflation history. In his story, there is only a temporary gap between the government's model and nature's because he assumes that the government has a correct model but initially does not know its parameter values. Primiceri (2006) captures the evolution of U.S. inflation after 1960 with a policy authority's learning dynamics along a path that converges to a self-confirming equilibrium in which the government's

[48]This allows these models to make contact with Albert Ando's remarks to me that by the mid 70s even those observers who were using the econometric specifications that Lucas (1972a) and Sargent (1971) had criticized had enabled researchers to detect that the tradeoff between inflation and unemployment was much less exploitable than many observers had thought earlier.

model equals nature's. The success of this paper comes from Primiceri having been able to calibrate a plausible set of initial government beliefs that allows the dynamics of the authority's beliefs that follow a path that captures the 1960s-1980s rise and fall of U.S. inflation and the associated broad movements in unemployment, then converges to a self-confirming equilibrium that describes the data after the mid 1990s.

Primiceri's model has these components: (1) a time invariant true data generating model featuring (i) an expectations augmented Phillips curve, (ii) an aggregate demand equation that describes how the time t value of an uninterpreted government policy instrument v_t affects current and future gaps between the unemployment rate u_t and a natural rate of unemployment u_t^N,[49] and (iii) a Cagan (1956)-Friedman (1956) adaptive expectations scheme that describes how the public forms the expectations of inflation that appear in (i)[50]; (2) initial beliefs for the government about the value of the natural rate of unemployment and the coefficients in a reduced-form Phillips curve[51]; (3) constant-gain recursive least squares algorithms for updating the government's beliefs; (4) a Phelps problem for choosing v_t on the basis of the government's time t beliefs.

The model neatly allows the government's misperception of the natural rate to influence policy, as advocated by Orphanides (2002, 2003), while adding other potentially important government misperceptions that impinge on the first-period outcome v_t of the Phelps problem. Primiceri shows that the lower is the sum of the weights on lagged inflation in the expectational Phillips curve, and therefore the less persistent is inflation under a passive government policy, the less counterinflationary is the policy that emerges from the Phelps problem. A lower estimated persistence of inflation indicates to the government that mean reverting inflation will evaporate soon enough on its own. The coefficients in the government's time t estimated model measure the strength of the feedback from unemployment to inflation and therefore how costly is the inflation-unemployment perceived by the government at time t. They therefore also influence how actively counterinflationary is the policy emerging from the time t Phelps problem.

Primiceri uses the following disciplined estimation strategy. He calibrates initial government beliefs by using data between 1948 and 1960 and to estimate them by least squares. An important feature of these calibrated beliefs is that they feature a level of persistence of inflation in the Phillips curve that is much lower than the high persistence that prevails in the estimated model's self-confirming equilibrium. In addition to these initial conditions, Primiceri sets two constant gain parameters, a separate one for the natural rate, another for all other coefficients in the government's beliefs. These calibrated objects, the data, and

[49]Feature (ii) of Primiceri's model embraces a Keynesian spirit of assuming that the authority influences output directly through the aggregate demand function, then inflation indirectly through the expectations-augmented Phillips curve. Contrast this with the classical specification adopted by Sims (1988), Chung (1990), Sargent (1999), Cho et al. (2002), and Sargent et al. (2006b). Primiceri uses the notation V_t rather than v_t.

[50]Primiceri assumes that a fraction of agents form expectations this way, while the remainder have rational expectations. Primiceri's specification imposes that the sum of weights on lagged inflation equals unity. Lucas (1972a) and Sargent (1971) argued that, except in a special case, the sum of the weights on lagged inflation being one is not a valid characterization of the natural rate hypothesis. Despite those papers, this characterization continues to be adopted. See King and Watson (1994) and Sargent (1999).

[51]The reduced form is derived by substituting the adaptive expectations scheme into the expectations augmented Phillips curve.

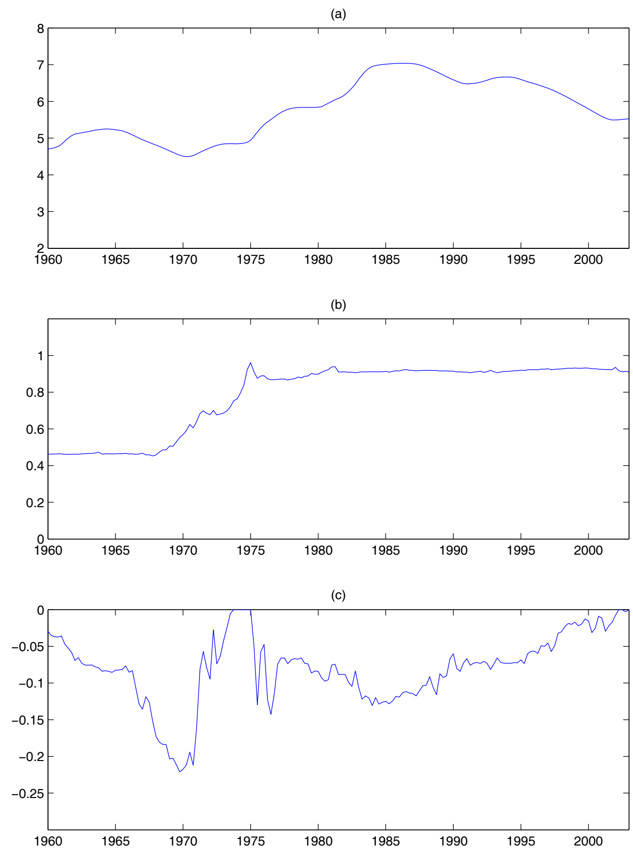

Figure 1: Evolution of policy-maker's beliefs about: (a) the natural rate of unemployment; (b) the persistence of inflation in the Phillips curve; and (c) the slope of the Phillips curve. (Primiceri 2006, p. 882)

the parameters of the structural relations pin down the government's beliefs. There are no additional free parameters in the estimation describing the government's beliefs. Primiceri uses maximum likelihood to estimate parameters appearing in the government's objective function and the time-invariant structural equations.

Primiceri explains the acceleration of inflation in the 1960s and 1970s, then the fall in the 1980s in terms of the government's initial underestimates of the natural rate hypothesis that were followed along a learning path by a temporal pattern of underestimates of the persistence of inflation and overestimates of the costs of disinflation coming from its estimated inflation-unemployment tradeoff. Figure 1 reproduces Primiceri's figure II, which shows his estimates of the evolution of the Fed's estimates of the natural rate of unemployment, the persistence inflation, and the slope of the Phillips curve. The Phelps problem attributes the acceleration of inflation to the authority's initial underestimates of the natural rate and the persistence of inflation. It attributes the reluctance to deflate to its overestimation of the costs of disinflation as captured by the slope of the Phillips curve. We will return to this point in subsection 7.5, where we link it to some econometric issues about direction of fit raised by King and Watson (1994).[52]

There is a link between under-estimates of the natural rate and over-estimates of the sacrifice ratio. When the Fed under-estimates the natural rate and over-estimates the unemployment gap, it over-predicts the amount of disinflation. When that disinflation fails to materialize, it revises its estimate of the slope of the Phillips curve towards zero. Thus, Orphanides's story complements stories like Primiceri's about sacrifice ratio pessimism.

7.5 An eclectic account

Cogley and Sargent (2005) performed an exercise that did not require them to specify a true data generating mechanism, it being enough for their purposes to consult the empirical distribution. But the government's perceptions about policy paths not taken play a leading role in their story. A government entertains three models that Cogley and Sargent chose to represent specifications that had at one time or another received prominent support in the literature about U.S. unemployment-inflation dynamics described by King and Watson (1994). The models are (1) a Samuelson-Solow Phillips curve with King and Watson's Keynesian direction of fit, a model that implies a long-run exploitable trade-off between inflation and unemployment, (2) a Solow-Tobin model with a Keynesian direction of fit that features a short-run but no long-run trade-off (according to what Lucas (1972a) and Sargent (1971) claimed was a dodgy notion of long-run) between inflation and unemployment; and (3) a Lucas specification with a classical direction of fit that implies no exploitable trade-off between inflation and unemployment. If probability one is put on the Lucas model, the Phelps problem gives the trivial solution that the government should set the systematic part of inflation equal to zero. If probability one is put on either of the other models, inflation is a linear function of the state variables appearing in those exploitable dynamic Phillips curves. The government puts positive probability on all three models, so the Phelps problem brokers

[52] Among many interesting features of Primiceri's results are his estimate of k, a parameter in the government objective function that allows Primiceri to evaluate the government's temptation to deviate from the natural rate (he finds that the temptation is small) and the time series that he extracts for v_t, which tracks a real interest rate very well after 1980.

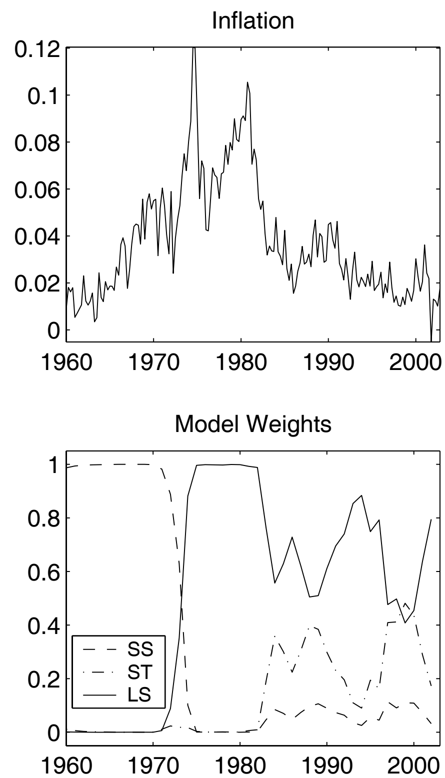

Figure 2: CPI inflation and Bayesian posterior model weights on the Samuelson-Solow (SS), Solow-Tobin (ST), and Lucas (LS) models.

some kind of compromise among the recommendations of the three models, but what kind of compromise?

The government starts with a prior with non-zero weights on all three models in 1960, estimates each sub model using Bayesian methods, and updates its prior over the three sub models. In each period, the government solves a Phelps problem that penalizes inflation and unemployment and that uses its time t prior to average over its time t estimates of its three submodels. Cogley and Sargent put prior probabilities in 1960 of .98 on the Samuelson-Solow model and .01 each on the Solow-Tobin and the Lucas model. Those prior probabilities on the Lucas and Solow-Tobin models were intended to respect the spirit of the above quote from McDonald (1985) because only the Samuelson-Solow model had been invented in 1960. Putting U.S. inflation-unemployment data into this machine, Cogley and Sargent compute time series of (1) the posterior model weights $\alpha_{i,t}$, and (2) the systematic part of the inflation rate set by the government in the Phelps problem.

Figures 2 and 3 taken from Cogley and Sargent (2005) frame the following puzzles. By the early 1970s, the data had moved the government's prior to put probability approaching 1 on the Lucas model (see figure 2). Since that model recommends zero inflation, why was actual inflation so high and variable in the 1970s? And why was the systematic part of inflation that emerges from the Phelps problem (see figure 3) even higher and more variable? Why does the Phelps planner seem to disregard the recommendations of the Lucas model

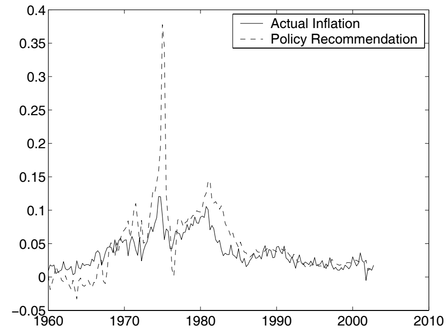

Figure 3: CPI inflation and recommendation from Phelps problem.

and crank out high target inflation throughout the 1970s?

As Cogley and Sargent (2005)) explain, the answer is to be found in what the Samuelson-Solow and Solow-Tobin models say would happen if the Lucas zero-target-inflation policy were to be adopted. In the early 1970s, the coefficients in those submodels, with their Keynesian direction of fit,[53] moved in ways that pointed to very high sacrifice ratios. Despite their low posterior probabilities, those models implied very high expected discounted losses were the Lucas program to be implemented immediately. In contrast, the high-probability Lucas model implied less adverse consequences if the recommendations of the Samuelson-Solow or Solow-Tobin models were allowed to prevail. The Phelps problem weights the submodel posterior probabilities against losses associated with various off-taken-path recommendations. The Lucas models policy recommendation did not prevail in the 1970s because there was a low probability that it would be disastrous. In order for a low-inflation recommendation to emerge from the Phelps problem, it was necessary that the estimated coefficients in the Samuelson-Solow and Solow-Tobin models adjust in ways that would render less adverse the consequences of a low-inflation policy. The data indicate that happened by the early 1980s.[54]

The direction-of-fit issue discussed by King and Watson (1994) is important in understanding how some of Primiceri's results relate to Cogley and Sargent's. Both models emphasize how monetary policy changed as the authorities updated their estimates, and Primiceri also attributes the inflation of the 1970s to the high perceived sacrifice ratio that Keynesian Phillips curve models presented to policy makers. But Primiceri assumes that the Fed relied exclusively on a version of the SolowTobin model and does not address why the Fed disregarded the recommendations of the Lucas model. The central element of his story – the high perceived cost of disinflation or sacrifice ratio – is not a robust prediction across

[53] Again in the language of King and Watson.

[54] The data also indicate that Bayes' law sponsors comebacks for the Samuelson-Solow and Solow-Tobin models in the 1980s and 1990s. One reaction that a true believer in the Lucas model might have is that Bayes' law is just too forgiving in still putting positive probability on those other models after the early 1970s data had come in, and that the inflation problem of the 1970s would have been solved by driving a stake through those other models. But no one has the authority to drive stakes, and models with operating characteristics much like those two survive today. The dispute between the fallacious (according to Friedman and Schwartz (1963, p. 191)) real bills doctrine and the quantity theory of money is mottled with repeated episodes having one of these doctrines being disposed of in favor of the other, then the other making a comeback. The real bills doctrine rides high in times like these when the Fed pegs the Federal Funds rate.

the three submodels used by Cogley and Sargent because it depends critically on the direction of fit, as documented by Cogley and Sargent (2005, p. 546-547). The reason that the sacrifice ratios differ so much across submodels comes from how the submodels interpret the diminished, near-zero contemporaneous covariance between inflation and unemployment that had emerged by the early 1970s. In a Keynesian Phillips curve, this diminished covariance flattens the short-term tradeoff, making the authorities believe that a long spell of high unemployment would be needed to bring inflation down, prompting Keynesian modelers to be less inclined to disinflate. But for a classical Phillips curve, the shift toward a zero covariance steepens the short-term tradeoff, making the authorities believe that inflation could be reduced at less cost in terms of higher unemployment. Thus, a classically-oriented policy maker would be more inclined to disinflate.

7.6 A monetary policy rules literature

The adaptive models described in the preceding three subsections all explain the rise and fall of post WWII U.S. inflation in terms of monetary policy rules that drifted over time in response to drifts over time in the monetary authorities' models of the economy. All three models embed very crude descriptions of the monetary policy rules and completely sidestep interesting questions about monetary policy transmission mechanisms. It is appropriate to say a few words about a related literature that uses time series data to infer the structure of post WWII U.S. monetary policy rules and how they have changed over time. The bottom line is that this literature has mixed evidence about whether monetary policy rules shifted enough to validate stories along the lines of our three adaptive models.[55]

Bernanke and Mihov (1998) developed an SVAR methodology for measuring innovations in monetary policy and their macroeconomic effects. They compared alternative ways of measuring monetary policy shocks and derived a new measure of policy innovations based on possibly time-varying estimates of the Fed's operating procedures. They presented a measure of the overall stance of policy (see Bernanke and Mihov (1998, Fig. III, p. 899)) that is striking in how the distribution of tight and loose policies seems not to have changed much in the periods before and after 1980.

But Clarida et al. (2000) estimated a forward-looking monetary policy reaction function for the postwar United States economy before and after Volcker's appointment as Fed Chairman in 1979 and found substantial differences in the estimated rules across periods. They found that interest rate policy in the Volcker-Greenspan period has been much more sensitive to changes in expected inflation than in the pre-Volcker period. They then extracted implications of the estimated rules for the equilibrium properties of inflation and output in a new Keynesian DSGE model and found that the Volcker-Greenspan rule is stabilizing, but that the earlier rule was not. Lubik and Schorfheide (2004) estimated a new Keynesian model like Clarida et al.'s in which the equilibrium is undetermined if monetary policy is passive and constructed posterior weights for the determinacy and indeterminacy region of the parameter space as well as estimates for the propagation of fundamental and sunspot shocks. They found that U.S. monetary policy post-1982 was consistent with determinacy but that

[55] This mixed news partly reflects the theoretical property of time series models that it is statistically difficult to detect drifts or shifts in the systematic part of a vector autoregression and much easier to detect changes in volatilities.

the pre-Volcker policy was not, and also that before 1979 indeterminacy substantially altered the propagation of shocks.

In contrast, working in terms of less fully interpreted models, Sims and Zha (2006) estimated a multivariate regime-switching model for monetary policy and found that the best fit allows time variation in disturbance variances only. When they permitted the systematic VAR coefficients to change, the best fit was with change only in the monetary policy rule. They estimated three regimes that correspond to periods across which the folk-wisdom states that monetary policy differed. But they found that those differences among regimes were not large enough to account for the rise and decline of inflation of the 1970s and 1980s. Likewise, by estimating a time-varying VAR with stochastic volatility, Primiceri (2005) found that both the systematic and non-systematic components of monetary policy had changed. In particular, he found that the systematic responses of the interest rate to inflation and unemployment exhibited a trend toward a more aggressive behavior, while also having sizeable high frequency oscillations. But Primiceri concluded that those had small effects on the rest of the economy and that exogenous non-policy shocks were more important than interest rate policy in explaining the high inflation and unemployment episodes described above, thus coming down more on the 'bad luck' than the 'bad policies' side of the argument. One can only hope that conclusion is too pessimistic because we have learned to do better.

8 Concluding remarks

It remains true that we are far "... from being able to solve, with full knowledge of the case, a multitude of questions which are boldly decided every day." (Cournot 1838, p. 5) It is unreasonable to criticize rational expectations models for not shedding light on problems that they were not designed to study, things like model specification doubts and policy makers with divergent and shifting models that can't be expressed in terms of a single model shared by nature and humankind. It is useful to recall how Lucas (1976) motivated his critique of econometric policy evaluation procedures with a summary of evidence for drifting coefficients and the lack of a self-contained theory of those drifts in pre-rational expectations theoretical structures. But the theoretical examples in that paper, and in much of the subsequent rational expectations literatures in macroeconomics, imply time-invariant VARS that fail to shed light on those drifting coefficients.[56] Models of adaptation can capture those drifts while retaining much of the structure of the cross-equations restrictions brought by rational expectations.

Rational expectations models are good tools for answering some questions. Learning models are good tools for answering others.

> The traditional rational-expectations model of inflation and inflation expectations has been a useful workhorse for thinking about issues of credibility and institutional design, but, to my mind, it is less helpful for thinking about economies in which (1) the structure of the economy is constantly evolving in ways that are imperfectly understood by both the public and policymakers and (2) the policymakers' objective function is not fully known by private agents. In particular,

[56] See Sargent (1999, ch. 2) for a discussion of this loose end in responses to the Lucas critique.

> together with the assumption that the central bank's objective function is fixed and known to the public, the traditional rational-expectations approach implies that the public has firm knowledge of the long-run equilibrium inflation rate; consequently, their long-run inflation expectations do not vary over time in response to new information.
>
> Although variations in the extent to which inflation expectations are anchored are not easily handled in a traditional rational-expectations framework, they seem to fit quite naturally into the burgeoning literature on learning in macroeconomics. The premise of this literature is that people do not have full information about the economy or about the objectives of the central bank, but they instead must make statistical inferences about the unknown parameters governing the evolution of the economy. Bernanke (2007)

By stressing the possibility that learning has propelled us to a self-confirming equilibrium in which the government chooses an optimal policy based on a wrong model, the learning literature changes how we should think about promoting the novel policies that will allow misguided governments to break out of the lack-of-experimentation traps to which self-confirming equilibria confine them.

Non mathematical appendix

A From commodity to fiat money

To introduce my theme, I cite David Ricardo:

> The introduction of the precious metals for the purposes of money may with truth be considered as one of the most important steps towards the improvement of commerce, and the arts of civilised life; but it is no less true that, with the advancement of knowledge and science, we discover that it would be another improvement to banish them again from the employment to which, during a less enlightened period, they had been so advantageously applied. Ricardo (1816, p. 65)

A long and disorderly process with "much backing and filling and confusion about purpose and power" led to Ricardo's idea.[57] Keynes and others made that idea the foundation of their proposals for a well managed fiat currency.

A.1 Learning to manage a commodity currency by supplementing it with tokens

Redish (1990, 2000) and Sargent and Velde (2002) described how it took 800 years to understand and cope with two imperfections that marred an ideal self-regulating commodity money system in which coins of all denominations were meant to exchange at values proportional to silver (or gold) content. In that ideal system, a government instructed a mint to offer to sell coins of different denominations for silver at prices proportional to their weights in silver. The mint did not buy coins for silver, but citizens were free to melt silver coins to recover silver. If minting and melting were costless, this self-regulating system would automatically adjust the denomination structure of coins to suit coin holder's preferences by letting them melting coins of a denomination they wanted less of, then taking the silver to the mint to buy coins of the denomination they wanted to acquire.[58] In the ideal system, a silver melt point equaled a silver mint point, denomination by denomination.

In practice, two imperfections hampered this system: (1) it was costly to produce coins; and (2) coins depreciated through wear and tear and sweating and clipping. The first imperfection gave rise to nonempty intervals between melt and mint points for gold or silver coins of each denomination – an upper point that indicated a melting point for that coin and a lower one that prompted minting. The proportionate spreads between minting and melting points differed because as a fraction of the value of the coin, it is cheaper to produce a large denomination coin than a small denomination coin. Unless the government were to subsidize the mint for producing low denomination coins, the spread between minting

[57] I borrowed the words in quotes from Friedman and Schwartz (1963, p.193), who used them to describe the evolution of the beliefs and policies of the Federal Reserve.

[58] Sargent and Velde (2002, p. 95) cited Bernando Davanzati, who in 1588 wrote that "metal should be worth as much in bullion as in coin, and be able to change from metal to money and money to metal without loss, like an amphibious animal."

and melting points would be proportionately wider for low denomination coins. The second imperfection allowed underweight coins to circulate along side full weight coins.

A nonempty interval between melting and minting points allowed coins to circulate by *tale* (i.e., by what is written on the coin rather than by weight) at an exchange value that exceeded their value by weight. Indeed, as Adam Smith pointed out, in the presence of costs of coinage, the money supply mechanism provided incentives for people to mint coins only when their value in exchange exceeded their value by weight by enough to cover the brassage and seigniorage fees that the mint charged to cover costs of production and taxes (Smith 1789, Book I, ch. 5).

Nonempty intervals with proportionately wider widths for lower denomination coins and a consequent exchange rate indeterminacy allowed the intervals to shift over time and eventually to become so misaligned that they recurrently provided incentive for coins of lower denominations to be melted. Sargent and Velde (2002) described why small denomination coins depreciated relative to large ones during periods of especially high demand for small change, and how that perversely pushed the mint-melt point intervals for small coins toward positions at which the small coins would be *melted*, thereby creating the recurring shortages of small coins documented by Cipolla (1956, 1982).[59]

Cipolla (1956) described a temporary practical remedy for these shortages. To cure a shortage of small denomination coins, the authorities debased them, thereby shifting the mint-melt intervals for small denomination coins in a direction that motivated citizens to take silver to the mint to purchase new coins. Monetary authorities throughout Europe used this method for hundreds of years. There were repeated debasements in small denomination silver coins and secular declines in rates of exchange of small denomination for large denomination coins.

Many experiments, some inadvertent, others purposeful, were performed, and numerous theoretical tracts were written and disputed before what Cipolla (1956) called the 'standard formula' for issuing token small denomination coins was put into practice in the mid 19th century.[60] It solved the problem of misaligned mint-melt intervals for coins of different denominations. by, first, having only one large denomination full weight coin that the mint would stand ready to sell for a precious metal, and, second, having the government issue difficult-to-counterfeit small denomination token coins that it would promise to convert on demand into the large denomination coin. This required a technology for manufacturing coins that are difficult to counterfeit.[61]

As examples of inadvertent experiments, token monies were occasionally issued inside besieged cities and sometimes they worked well enough. A purposeful experiment was sparked by a document that prefigured later arguments of John Law, Adam Smith, and David Ricardo. Advisors to King Ferdinand II of Spain told him that he could issue token copper coins that Spanish residents would voluntarily accept from the government in exchange for full bodied silver coins. They described how this could be done in a noninflationary way

[59]This multi-interval commodity money system in which coins circulate by tale is taken for granted by Smith (1789, book I, ch. 5).

[60]This process of shuttling through experiments, reformulations of theories, and further experiments reminds me of the hypothesis-testing learning models of Foster and Young (2003) and Cho and Kasa (2006), but I might be imagining things.

[61]See Redish (1990, 2000) and Selgin (2003).

and how it would provide a fiscal boon to the Spanish treasury.[62] Three successive Spanish Kings tried this experiment, which had all of the ingredients of the 19th century standard formula except convertibility. For 25 years, the experiment worked well, yielding the government substantial revenues without inflation. But eventually excessive issues of copper coins caused inflation, in the aftermath of which the Spanish monetary authorities pursued a fascinating sequence of experiments that manipulated the exchange rate of copper coins to adjust the price level and/or raise revenues for the Spanish government by restamping copper coins or manipulating the unit of account. In a commodity money system, the quantity theory is muted because it can operate only in the limited interval allowed between the mint and melt points for the precious metal. When the Spanish experiment broke through those restrictions, it gave the British statistician Sir William Petty the data that he used to discover a quantity theory of money (see Hull (1899)). Other episodes created more data to substantiate the quantity theory of money, for example, the construction and collapse of John Law's system (see Velde (2007)) and the overissuing of French assignats after the sales of the church lands that had initially backed them were suspended after war broke out in 1792 (see Sargent and Velde (1995)). But while those episodes lent vivid empirical support to a quantity theory, they also brought evidence that government monetary authorities could not be trusted to administer a pure fiat standard in ways that stabilized prices.[63]

In 1660, the master of the British mint, Henry Slingsby, added an element missing from the Spanish experiment, namely, convertibility of token coins, and went on to recommend what in the 19th century became the standard formula.[64] But perhaps because the inflation accompanying the Spanish and some other similar experiments had given token coins such a bad name, the British government ignored Slingsby's recommendations. Many experts, including Locke (1691), continued to insist that token coins of any denomination were dangerous and that a good faith commodity money system required that coins of all denominations be full bodied. For a long time, that sentiment convinced national governments not to issues tokens, but that did not stop other entities from creating them. In seventeenth and eighteenth century Britain, hundreds of private firms and municipalities issued small denomination tokens that formed a substantial part of the country's coinage. Between 1816 and 1836, the British government implemented the standard formula by nationalizing a token coin industry that had long existed.

A.2 Ricardo's proposal

It required 156 years to take the short logical step from Slingsby's 1660 standard formula for issuing convertible token subsidiary coins to David Ricardo's 1816 recommendation. Ricardo suggested that a country's domestic money supply should ideally consist of paper notes that the government promises to exchange at a pegged price for gold bullion bars, but that no gold coins should actually be minted. A variant of Ricardo's scheme in which a government

[62]See the document cited in Sargent and Velde (2002, pp. 231-232).

[63]I suspect that is why later advocates for replacing the gold standard with 'more scientific' systems of managed currencies including Adam Smith and Ricardo to Keynes purposefully omitted references to some of the historical experiments that generated the data that were sources for the quantity theory of money. For example, Smith (1789) did not cite John Law's theoretical writings as among the sources for his monetary recommendations.

[64]See Sargent and Velde (2002, pp. 268-269).

promises to redeem domestic notes for gold, but only for foreign residents, came to be practiced around 1900. This arrangement, by which "a cheap local currency [is] artificially maintained at par with the international standard of value," (Keynes 1913, p. 25) was called the "gold exchange standard." Keynes described how by 1913 this system had come to prevail in India through a sequence of haphazard administrative decisions that eventually produced a coherent system that no one had planned but that Keynes applauded. Keynes (1913, p. 25) predicted that Ricardo's scheme would be an essential part of "the ideal currency system of the future."[65]

The standard formula eliminates the gold or silver points for all coins except one standard coin and uses the mint and melt points for that coin to regulate the total quantity of money, while it uses its promise freely to convert tokens into that standard coin to produce the correct denomination composition. It was one more step from the standard formula or Ricardo's proposal to the idea of Fisher (1920), Keynes, and others that well intentioned government officials should administer a fiat currency in ways that stabilize the price level. Doing that would allow them to to remove the mint and melt points for the one standard coin too. Discovering the quantity theory of money was an essential step in learning the conditions under which a fiat money system could be managed to provide greater price level stability than could be achieved with a gold standard. Under a gold standard, the domain over which the quantity theory could act was restricted to the interval between the gold points[66], a fact that qualifies Friedman and Schwartz's interpretation of U.S. price movements in terms of a quantity theory at least until 1933.

Despite the discovery of the quantity theory of money, Keynes acknowledged that many people doubted that a well-managed fiat currency would be incentive-feasible:

> The advocates of gold, as against a more scientific standard, base their cause on the double contention, that in practice gold has provided and will provide a reasonably stable standard of value, and that in practice, since governing authorities lack wisdom as often as not, a managed currency will, sooner or later, come to grief. Conservatism and scepticism join arms – as they often do. Perhaps superstition comes in too; for gold still enjoys its smell and colour. Keynes (1924, p. 132)

But sceptics were swimming against the tide:

> "... in the modern world of paper money and bank credit there is no escape from a 'managed' currency, whether we wish it or not (p. 136) ... In truth, the gold standard is already a barbarous relic. ... Advocates of the ancient standard do

[65] Speaking of how a change in Indians' preferences for holding gold could cause world-wide inflation in prices:

> The time may not be far distant when Europe, having perfected her mechanism of exchange on the basis of a gold standard, will find it possible to regulate her standard of value on a more rational and stable basis. It is not likely that we shall leave permanently the most intimate adjustments of our economic organism at the mercy of a lucky prospector, a new chemical process, or a change of ideas [preferences for holding gold] in Asia. (Keynes 1913, p. 71)

[66] See Sargent and Smith (1997). Because the money supply is purely endogenous at both the gold mint point and the gold melt point, exogenous movements in the money supply can occur only within those points.

> not observe how remote it now is from the spirit and requirements of the age. A regulated non-metallic standard has slipped in unnoticed. *It exists.* Whilst the economists dozed, the academic dream of a hundred years, doffing its cap and gown, clad in paper rags, has crept into the real world by means of the bad fairies – always so much more potent than the good – the wicked ministers of finance. Keynes (1924, p. 138)

As Keynes wanted, in the twentieth century governments throughout the world carried out the historically unprecedented experiment of managing currencies completely cut off from gold backing (see Friedman (1991, p. 245)). Figure 4 documents that, at least until very recently, the monetary authorities in four hard-currency countries failed to live up to Keynes's high expectations for them and to deliver the kind of price stability that their predecessors had attained when they were restrained by that barbarous relic. Figures 5 and 6 show price indexes for Istanbul and Argentina, places with softer currencies (compare the vertical scales).

I let Milton Friedman have the last word.

> ... the world is now engaged in a great experiment to see whether it can fashion a different anchor, one that depends on government restraint rather than on the cost of acquiring a physical commodity ... The verdict is far from in on whether fiat money will involve a lower cost than commodity money ... Friedman (1991, p. 42).
>
> Nonetheless, it remains an open question whether the temptation to use fiat money as a source of revenue will lead to a situation that will ultimately force a return to a commodity standard ... The final answer will come only as history unfolds over the next decades. What that answer will be depends critically on our success in learning from historical episodes such as those that have been examined in this book. Such a learning process has been under way for centuries, ever since the first appearance of systematic analyses of money and monetary institutions. It has entered a new and urgent stage as the world ventures into hitherto unexplored terrain. Friedman (1991, pp. 259-260).

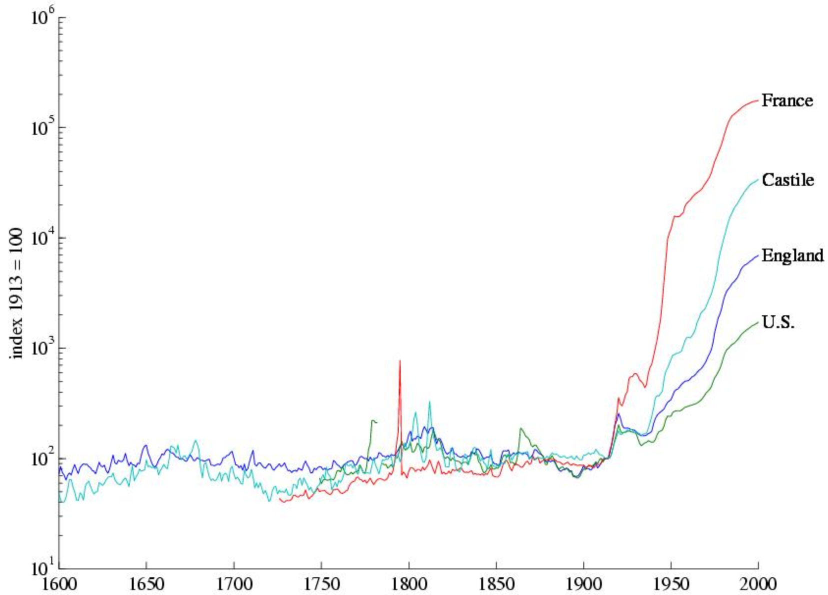

Figure 4: Indices of prices in terms of unit of account in England, the United States, France, and Spain. Sargent and Velde (2002, p. 35)

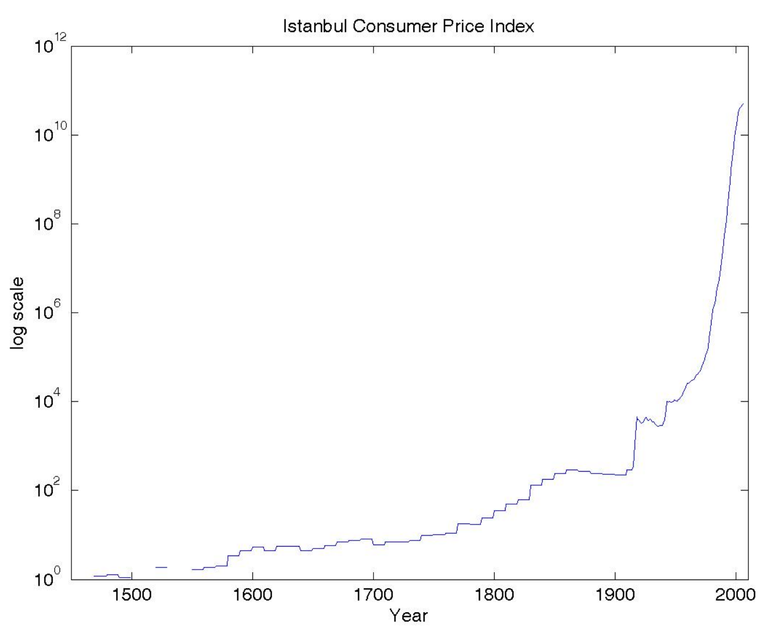

Figure 5: Indices of prices in Istanbul.

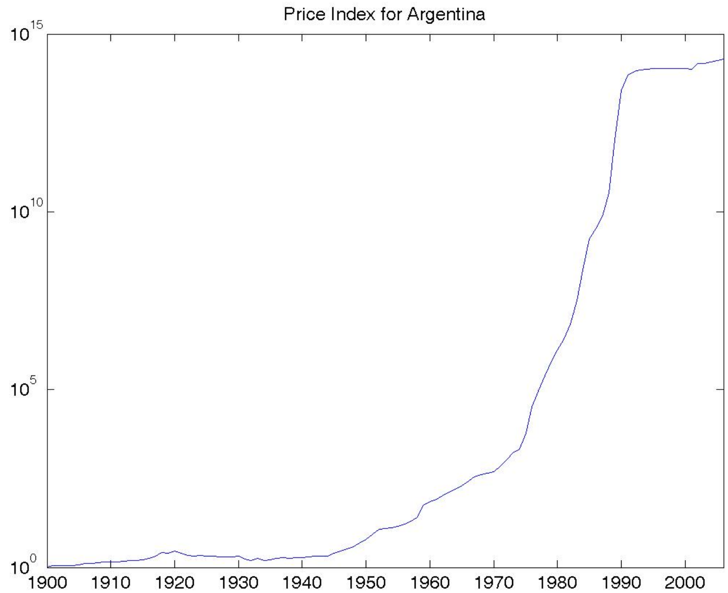

Figure 6: Price index for Argentina.

References

Abreu, Dilip, David Pearce, and Ennio Stacchetti. 1986. Optimal Cartel Equilibria with Imperfect Monitoring. *Journal of Economic Theory* 39:251–269.

———. 1990. Toward a Theory of Discounted Repeated Games with Imperfect Monitoring. *Econometrica* 58 (5):1041–1063.

Aiyagari, S. Rao, Albert Marcet, Thomas J. Sargent, and Juha Seppala. 2002. Optimal Taxation without State-Contingent Debt. *Journal of Political Economy* 110 (6):1220–1254.

Albanesi, Stefania, V.V. Chari, and Lawrence J. Christiano. 2002. Expectation Traps and Monetary Policy. Nber working papers, National Bureau of Economic Research, Inc.

Allen, Franklin, Stephen Morris, and Hyun Song Shin. 2006. Beauty Contests and Iterated Expectations in Asset Markets. *Review of Financial Studies* 19 (3):719–752.

Amato, Jeffery D. and Hyun Song Shin. 2006. Imperfect common knowledge and the information value of prices. *Economic Theory* 27:213–241.

Aumann, Robert J. 1976. Agreeing to disagree. *The Annals of Statistics* 4:1236–1239.

Ball, Laurence M. 1995. Time-Consistent Policy and Persistent Changes in Inflation. *Journal of Monetary Economics* 329–350.

Bergemann, Dirk and Stephen Morris. 2005. Robust Mechanism Design. *Econometrica* 73 (6):1771–1813.

Bernanke, Ben S. 2007. Inflation Expectations and Inflation Forecasting. Remarks at the Monetary Economics Workshop of the National Bureau of Economic Research Summer Institute, Cambridge, Massachusetts, available on the Federal Reserve Board web site.

Bernanke, Ben S. and Ilian Mihov. 1998. Measuring Monetary Policy. *The Quarterly Journal of Economics* 113 (3):869–902.

Blinder, Alan S. 1998. *Central Banking in Theory and Practice (Lionel Robbins Lectures)*. Cambridge, Massachusetts: MIT Press.

Blume, Lawrence and David Easley. 2006. If You're so Smart, why Aren't You Rich? Belief Selection in Complete and Incomplete Markets. *Econometrica* 74 (4):929–966.

Bossaerts, Peter. 2002. *The Paradox of Asset Pricing*. Princeton, New Jersey: Princeton University Press.

———. 2004. Filtering Returns for Unspecified Biases in Priors when Testing Asset Pricing Theory. *Review of Economic Studies* 71:6386.

Branch, William A. and George W. Evans. 2005. Model Uncertainty and Endogenous Volatility. Computing in economics and finance 2005, Society for Computational Economics.

———. 2006. Intrinsic heterogeneity in expectation formation. *Journal of Economic Theory* 127 (1):264–295.

Bray, Margaret M. and David M. Kreps. 1987. Rational Learning and Rational Expectations. In *Arrow and the Ascent of Modern Economic Theory*, edited by George R. Feiwel, 597–625. New York University Press.

Brock, William A. and Patrick de Fontnouvelle. 2000. Expectational diversity in monetary economies1. *Journal of Economic Dynamics and Control* 24 (5-7):725–759.

Brock, William A. and Cars H. Hommes. 1997. A Rational Route to Randomness. *Econometrica* 65 (5):1059–1096.

Brunnermeier, Markus K. and Jonathan A. Parker. 2005. Optimal Expectations. *American Economic Review* 95 (4):1092–1118.

Brunnermeier, Markus K, Christian Gollier, and Jonathan A Parker. 2007. Optimal Beliefs, Asset Prices and the Preference for Skewed Returns. Cepr discussion papers, C.E.P.R. Discussion Papers.

Bruno, Michael and Stanley Fischer. 1990. Seigniorage, Operating Rules, and the High Inflation Trap. *Quarterly Journal of Economics* May:353–374.

Bullard, James and Kaushik Mitra. 2007. Determinacy, Learnability, and Monetary Policy Inertia. *Journal of Money, Credit and Banking* 39 (5):1177–1212.

Cagan, Phillip. 1956. *The Monetary Dynamics of Hyperinflation*, 25–117. Chicago: University of Chicago Press.

Calvo, Guillermo A. 1978. On the Time Consistency of Optimal Policy in a Monetary Economy. *Econometrica* 46:1411–1428.

Campbell, John Y. and John H. Cochrane. 1999. By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior. *Journal of Political Economy* 107:205–251.

Chamley, Christophe. 1986. Optimal Taxation of Capital Income in General Equilibrium with Infinite Lives. *Econometrica* 54 (3):607–22.

Chang, Roberto. 1998. Credible monetary policy with long-lived agents: recursive approaches. *Journal of Economic Theory* xxx–xxx.

Chari, V. V. and Patrick J. Kehoe. 1993a. Sustainable Plans and Debt. *Journal of Economic Theory* 61 (2):230–261.

Chari, V V and Patrick J Kehoe. 1993b. Sustainable Plans and Mutual Default. *Review of Economic Studies* 60 (1):175–95.

Chari, V V, Lawrence J Christiano, and Patrick J Kehoe. 1994. Optimal Fiscal Policy in a Business Cycle Model. *Journal of Political Economy* 102 (4):617–52.

Chari, V. V., Lawrence J. Christiano, and Martin Eichenbaum. 1998. Expectation Traps and Discretion. *Journal of Economic Theory* 81 (2):462–492.

Cho, In-Koo and Kenneth Kasa. 2006. Learning and Model Validation. University of Illinois.

Cho, In-Koo, Noah Williams, and Thomas J. Sargent. 2002. Escaping Nash Inflation. *Review of Economic Studies* 69:1–40.

Chung, Heetaik. 1990. *Did Policy Makers Really Believe in the Phillips Curve?: An Econometric Test*. Ph.D. thesis, University of Minnesota.

Cipolla, Carlo. 1956. *Money, Prices, and Civilization in the Mediterranean World, Fifth to Seventeenth Century*. New York: Gordian Press.

———. 1982. *The Monetary Policy of Fourteenth-Century Florence*. Berkeley: University of California Press.

Clarida, Richard, Jordi Gali, and Mark Gertler. 1999. The Science of Monetary Policy: A New Keynesian Perspective. *Journal of Economic Literature* 37 (4):1661–1707.

Clarida, Richard, Jordi Gal, and Mark Gertler. 2000. Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory. *The Quarterly Journal of Economics* 115 (1):147–180.

Cogley, Timothy. 2005. Changing Beliefs and the Term Structure of Interest Rates: Cross-Equation Restrictions with Drifting Parameters. *Review of Economic Dynamics* 8 (2):420–451.

Cogley, Timothy and Thomas J. Sargent. 2005. The Conquest of U.S. Inflation: Learning and Robutsness to Model Uncertainty. *Review of Economic Dynamics* 8:528–563.

Cournot, Augustin. 1838. *Researches into the Mathematical Principles of the Theory of Wealth*. New York: The Macmillan Company. Translated by Nathanial T. Bacon with a bibliography of mathematical economics by Irving Fisher.

Cunha, Flavior and James Heckman. 2007. The Technology of Skill Formation. *American Economic Review, Papers and Proceedings* 31–47.

Dagli, Ates C. and John B. Taylor. 1984. Estimation and solution of linear rational expectations models using a polynomial matrix factorization. *Journal of Economic Dynamics and Control* 8 (3):341–348.

Duffy, John. 2006. Experimental Macroeconomics. Prepared for the New Palgrave Dictionary of Economics.

Epstein, L. and S. Zin. 1989. Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework. *Econometrica* 57:937–969.

Evans, George W. and Seppo Honkapohja. 1999. Learning Dynamics. In *Handbook of Macroeconomics*, edited by John B. Taylor and Michael Woodford. Elsevier.

———. 2001. *Learning and Expectations in Macroeconomics*. Princeton University Press.

———. 2003. Expectations and the Stability Problem for Optimal Monetary Policies. *Review of Economic Studies* 70 (4):807–824.

Fetter, Frank Whitson. 1978. *Development of British Monetary Orthodoxy, 1797-1875*. Fairfield, New Jersey: Augustus M. Kelley. First published 1965 by Harvard University Press, Cambridge, Ma.

Fisher, Irving. 1920. *Stabilizing the Dollar: A Plan to Stabilize the General Price Level without Fixing Individual Prices*. New York: Macmillan.

———. 1926. *The Purchasing Power of Money: Its Determination and Relation to Credit, Interest and Crises*. New York: Macmillan. New and revised edition, originally published in 1911.

Foster, Dean P. and Rakesh Vohra. 1999. Regret in the On-Line Decision Problem. *Games and Economic Behavior* 29 (1-2):7–35.

Foster, Dean P. and H. Peyton Young. 2003. Learning, hypothesis testing, and Nash equilibrium. *Games and Economic Behavior* 45 (1):73–96.

Friedman, M. and A.J. Schwartz. 1963. *A Monetary History of the United States, 1857-1960*. Princeton, New Jersey: Princeton University Press.

Friedman, Milton. 1956. *A Theory of the Consumption Function*. Princeton, N.J.: Princeton University Press.

———. 1991. *Money Mischief: Episodes in Monetary History*. New York, San Dieto, and London: Harcourt Brace Jovanovich.

Fuchs, William and Francesco Lippi. 2006. Monetary Union with Voluntary Participation. *Review of Economic Studies* 73 (2):437–457.

Fudenberg, Drew and David K. Levine. 1993. Self-Confirming Equilibrium. *Econometrica* 61:523–545.

———. 1998. *Learning in Games*. Cambridge, MA.: M.I.T. Press.

———. 2007. Self-Confirming and the Lucas Critique. Unpublished manuscript, Harvard University and Washington University.

Gallant, A Ronald and George Tauchen. 1996. Which Moments to Match? *Econometric Theory* 12 (4):657–81.

Golosov, Mikhail, Narayana Kocherlakota, and Aleh Tsyvinski. 2003. Optimal Indirect and Capital Taxation. *Review of Economic Studies* 70 (3):569–587.

Golosov, Mikhail, Aleh Tsyvinski, and Ivan Werning. 2007. New Dynamic Public Finance: A User's Guide. *NBER Macroeconomics Annual 2006* 317–364.

Gourieroux, C., A. Monfort, and E. Renault. 1993. Indirect Inference. *Journal of Applied Econometrics* 8:S85S118.

Grossman, Sanford J and Robert J Shiller. 1981. The Determinants of the Variability of Stock Market Prices. *American Economic Review* 71 (2):222–27.

Hall, Robert E. 1971. The Dynamic Effects of Fiscal Policy in an Economy With Foresight. *Review of Economic Studies* 38 (114):229–44.

Hansen, Lars Peter. 1982. Large Sample Properties of Generalized Method of Moments Estimators. *Econometrica* 50 (4):1029–54.

———. 2007. Beliefs, Doubts and Learning: Valuing Economic Risk. *American Economic Review, Papers and Proceedings* 97 (2):1–30.

Hansen, Lars Peter and Ravi Jagannathan. 1991. Implications of Security Market Data for Models of Dynamic Economies. *Journal of Political Economy* 99 (2):225–62.

———. 1997. Assessing Specification Errors in Stochastic Discount Factor Models. *Journal of Finance* 52 (2):557–90.

Hansen, Lars Peter and Scott F Richard. 1987. The Role of Conditioning Information in Deducing Testable Restrictions. *Econometrica* 55 (3):587–613.

Hansen, Lars Peter and Thomas J. Sargent. 1980. Formulating and estimating dynamic linear rational expectations models. *Journal of Economic Dynamics and Control* 2 (2):7–46.

———. 1993. Seasonality and approximation errors in rational expectations models. *Journal of Econometrics* 55 (1-2):21–55.

———. 2008. Robustness.

———. 2009. *Recursive Models of Dynamic Linear Economies*. Princeton, New Jersey: Princeton University Press.

Hansen, Lars Peter and Kenneth J Singleton. 1983. Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns. *Journal of Political Economy* 91 (2):249–65.

Harrison, J. Michael and David M. Kreps. 1978. Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations. *Quarterly Journal of Economics* 92 (2):323–336.

———. 1979. Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory* 20 (3):381–408.

Hart, Sergiu and Andreu Mas-Colell. 2000. A Simple Adaptive Algorithm Leading to Correlated Equilibrium. *Econometrica* 68:1127–50.

———. 2001. A General Class of Adaptive Strategies. *Journal of Economic Theory* 98 (1):26–54.

———. 2003. Uncoupled Dynamics Do Not Lead to Nash Equilibrium. *The American Economic Review* 93 (5):1830–1836.

Hetzel, Robert. 1998. Arthur Burns. *Federal Reserve Bank of Richmond Economic Quarterly* 84 (1):21–44.

Hopenhayn, Hugo A and Juan Pablo Nicolini. 1997. Optimal Unemployment Insurance. *Journal of Political Economy* 105 (2):412–38.

Hull, Charles Henry, ed. 1899. *The Economic Writings of Sir William Petty*. Cambridge: Cambridge University Press.

Hume, David. 1985. Of the Independency of Parliament. In *Essays: Moral, Political, and Literary*. Liberty Classics. Edited and with a Forward, Notes, and Glossary by Eugene F. Miller.

Imrohoroglu, Selahattin. 1993. Testing for sunspot equilibria in the German hyperinflation. *Journal of Economic Dynamics and Control* 17 (1-2):289–317.

Jehiel, Philippe. 1995. Limited Horizon Forecast in Repeated Alternate Games. *Journal of Economic Theory* 67 (2):497–519.

———. 1998. Learning to Play Limited Forecast Equilibria. *Games and Economic Behavior* 22 (2):274–298.

Judd, Kenneth L. 1985. Redistributive taxation in a simple perfect foresight model. *Journal of Public Economics* 28 (1):59–83.

Kalai, Ehud and Ehud Lehrer. 1993. Rational Learning Leads to Nash Equilibrium. *Econometrica* 61 (5):1019–45.

———. 1994. Weak and strong merging of opinions. *Journal of Mathematical Economics* 23 (1):73–86.

Kasa, Kenneth. 2000. Forecasting the Forecasts of Others in the Frequency Domain. *Review of Economic Dynamics* 3:726–56.

Keynes, John Maynard. 1913. *Indian Currency and Finance*. Cambridge, UK: University Printing House. The Royal Economic Society edition, 1971.

———. 1924. *A Tract on Monetary Reform*. Macmillan St Martin's Press. Volume IV of the Collected Writings of John Maynard Keynes, published for the Royal Economic Society in 1971.

———. 1936. *The General Theory of Employment Interest and Money*. New York: Harcourt, Brace and Company.

King, Robert G. and Mark W. Watson. 1994. The post-war U.S. phillips curve: a revisionist econometric history. *Carnegie-Rochester Conference Series on Public Policy* 41:157–219.

Kocherlakota, Narayana R. 2005. Zero Expected Wealth Taxes: A Mirrlees Approach to Dynamic Optimal Taxation. *Econometrica* 73 (5):1587–1621.

Kocherlakota, Narayana R. and Christopher Phelan. 2007. On the Robustness of Laissez-Faire. Levine's bibliography, UCLA Department of Economics.

Kocherlakota, Narayana R. and Luigi Pistaferri. 2004. Asset Pricing Implications of Pareto Optimality with Private Information. Tech. rep.

———. 2007. Household Heterogeneity and Real Exchange Rates. *Economic Journal* 117 (519):C1–C25.

Kreps, D. M. and E. L. Porteus. 1978. Temporal Resolution of Uncertainty and Dynamic Choice. *Econometrica* 46:185–200.

Kreps, David M. 1998. Anticipated Utility and Dynamic Choice. In *Frontiers of Research in Economic Theory: The Nancy L. Schwartz Memorial Lectures, 1983-1997*. Cambridge University Press.

Krusell, Per, Vincenzo Quadrini, and Jose-Victor Rios-Rull. 1997. Politico-economic equilibrium and economic growth. *Journal of Economic Dynamics and Control* 21 (1):243–272.

Kydland, Finn E and Edward C Prescott. 1977. Rules Rather Than Discretion: The Inconsistency of Optimal Plans. *Journal of Political Economy* 85 (3):473–91.

Ljungqvist, Lars and Thomas J. Sargent. 2004. *Recursive Macroeconomic Theory, Second Edition*. Cambridge, Mass.: MIT Press.

Locke, John. 1691. *Some Considerations of the Consequences of the Lowering of Interest and the Raising of the Value of Money*. London: Awnsham and John Churchill.

Lubik, Thomas A. and Frank Schorfheide. 2004. Testing for Indeterminacy: An Application to U.S. Monetary Policy. *American Economic Review* 94 (1):190–217.

Lucas, Robert E., Jr. 1972a. Econometric Testing of the Natural Rate Hypothesis. In *The Econometrics of Price Determination: Conference, October 30-31*, edited by Otto Eckstein. Washington, D.C.: Board of Governors, Federal Reserve System.

———. 1972b. Expectations and the neutrality of money. *Journal of Economic Theory* 4 (2):103–124.

———. 1975. An Equilibrium Model of the Business Cycle. *Journal of Political Economy* 83 (6):1113–44.

———. 1976. Econometric Policy Evaluation: A Critique. In *The Phillips Curve and Labor Markets*, edited by K. Brunner and A. H. Meltzer, 19–46. Amsterdam: North-Holland.

———. 1986. Principles of fiscal and monetary policy. *Journal of Monetary Economics* 17 (1):117–134.

Lucas, Robert E., Jr. and Nancy L. Stokey. 1983. Optimal fiscal and monetary policy in an economy without capital. *Journal of Monetary Economics* 12 (1):55–93.

Madison, James. 1987. *Notes of Debates in the Federal Convention of 1787 Reported by James Madison*. London: W.W. Norton & Company. Bicentennial Edition.

Marcet, Albert and Juan P. Nicolini. 2003. Recurrent Hyperinflations and Learning. *American Economic Review* 93 (5):1476–1498.

Marcet, Albert and Thomas J. Sargent. 1989a. Convergence of Least Squares Learning Mechanisms in Self Referential Linear Stochastic Models. *Journal of Economic Theory* 48:337–368.

———. 1989b. Least-Squares Learning and the Dynamics of Hyperinflation. In *International Symposia in Economic Theory and Econometrics*, edited by William Barnett, John Geweke, and Karl Shell, 119–137. Cambridge, England: Cambridge University Press.

Marimon, Ramon. 1997. Learning from Learning in Economics. In *Advances in Economics and Econometrics: Theory and Applications*, edited by David Kreps and Kenneth Wallis. Cambridge University Press.

Marimon, Ramon, Ellen McGrattan, and Thomas J. Sargent. 1990. Money as a medium of exchange in an economy with artificially intelligent agents. *Journal of Economic Dynamics and Control* 14 (2):329–373.

McDonald, Forrest. 1985. *Novus Ordo Seclorum: The Intellectual Origins of the Constitution*. Lawrence, Kansas: University of Kansas Press.

Mehra, Rajnish and Edward C. Prescott. 1985. The equity premium: A puzzle. *Journal of Monetary Economics* 15 (2):145–161.

Meyn, S. P. 2007. Control Techniques for Complex Networks. To appear, Cambridge University Press.

Milgrom, Paul and Nancy Stokey. 1982. Information, trade and common knowledge. *Journal of Economic Theory* 26 (1):17–27.

Mishkin, Frederic S. 2007. Inflation Dynamics. Remarks at the Annual Macro Conference, Federal Reserve Bank of San Francisco, available on the Federal Reserve Board web site.

Morris, Stephen. 1996. Speculative Investor Behavior and Learning. *The Quarterly Journal of Economics* 111 (4):1111–33.

Muth, John F. 1960. Optimal Properties of Exponentially Weighted Forecasts. *Journal of the American Statistical Association* 55:299–306.

Orphanides, Athanasios. 2002. Monetary-Policy Rules and the Great Inflation. *American Economic Review* 92 (2):115–120.

———. 2003. The quest for prosperity without inflation. *Journal of Monetary Economics* 50 (3):633–663.

Orphanides, Athanasios and John C. Williams. 2005. The decline of activist stabilization policy: Natural rate misperceptions, learning, and expectations. *Journal of Economic Dynamics and Control* 29 (11):1927–1950.

———. 2007. Inflation targeting under imperfect knowledge. *Economic Review* 1–23.

Pavoni, Nicola and Gianluca Violante. 2007. Optimal Welfare-to-Work Programs. *Review of Economic Studies* 74 (1):283–318.

Pearlman, Joseph G. and Thomas J. Sargent. 2005. Knowing the Forecasts of Others. *Review of Economic Dynamics* 8 (2):480–497.

Phelan, Christopher and Ennio Stacchetti. 2001. Sequential Equilibria in a Ramsey Tax Model. *Econometrica* 69 (6):1491–1518.

Phelps, Edmund S. 1967. Phillips Curves, Expectations of Inflation and Optimal Unemployment over Time. *Economica* 2 (3):22–44.

Piazzesi, Monika and Martin Schneider. 2007. Equilibrium Yield Curves. *NBER Macroeconomics Annual 2006* 389–442.

Piskorksi, Tomasz and Alexei Tchistyi. 2006. Optimal Mortgage Design. Stern School, New York University.

Prescott, Edward C. 2002. Prosperity and Depression. *American Economic Review* 92:1–15.

Primiceri, Giorgio. 2005. Time Varying Structural Vector Autoregressions and Monetary Policy. *Review of Economic Studies* 72:82152.

———. 2006. Why Inflation Rose and Fell: Policymakers' Beliefs and US Postwar Stabilization Policy. *Quarterly Journal of Economics* 121 (3):867–901.

Redish, Angela. 1990. The Evolution of the Gold Standard in England. *Journal of Economic History* 50 (4):789–805.

———. 2000. *Bimetallism: An Economic and Historical Analysis*. Cambridge: Cambridge University Press.

Ricardo, David. 1816. *Proposals for an Economical and Secure Currency*. London: John Murray, Albermarle-Street. Reprinted in Vol. IV, of *The Works and Correspondence of David Ricardo*, edited by Pierro Sraffa, with the collaboration of M.H. Dobb, Liberty Fund, Indianapolis, 2004.

Rogerson, William P. 1985. Repeated Moral Hazard. *Econometrica* 53 (1):69–76.

Rogoff, Kenneth. 1985. The Optimal Degree of Commitment to an Intermediate Monetary Target. *The Quarterly Journal of Economics* 100 (4):1169–89.

Sandroni, Alvaro. 2000. Do Markets Favor Agents Able to Make Accurate Predictions? *Econometrica* 68:1303–1342.

Sargent, Thomas, Noah Williams, and Tao Zha. 2006a. The Conquest of South American Inflation. Nber working papers, National Bureau of Economic Research, Inc.

———. 2006b. Shocks and Government Beliefs: The Rise and Fall of American Inflation. *American Economic Review* 96 (4):1193–1224.

Sargent, Thomas J. 1971. A Note on the 'Accelerationist' Controversy. *Journal of Money, Credit and Banking* 3 (3):721–25.

Sargent, Thomas J. 1973. Rational Expectations, the Real Rate of Interest, and the Natural Rate of Unemployment. *Brookings Papers on Economic Activity* 1973 (2):429–480.

Sargent, Thomas J. 1977. The Demand for Money During Hyperinflations under Rational Expectations: I. *International Economic Review* 18 (1):59–82.

Sargent, Thomas J. 1979. A note on maximum likelihood estimation of the rational expectations model of the term structure. *Journal of Monetary Economics* 5 (1):133–143.

Sargent, Thomas J. 1984. Autoregressions, Expectations, and Advice. *American Economic Review* 74 (2):408–15.

Sargent, Thomas J. 1999. *The Conquest of American Inflation*. Princeton University Press.

Sargent, Thomas J. and Bruce D. Smith. 1997. Coinage, debasements, and Gresham's laws. *Economic Theory* 10 (2):197–226.

Sargent, Thomas J. and François Velde. 2002. *The Big Problem of Small Change*. Princeton, New Jersey: Princeton University Press.

Sargent, Thomas J. and François R. Velde. 1995. Macroeconomic Features of the French Revolution. *Journal of Political Economy* 103 (3):474–518.

Sargent, Thomas J and Neil Wallace. 1975. ”Rational” Expectations, the Optimal Monetary Instrument, and the Optimal Money Supply Rule. *Journal of Political Economy* 83 (2):241–54.

Sawhill, Isabel V. 1995. Distinguished Lecture on Economics in Government: The Economist vs. Madmen in Authority. *Journal of Economic Perspectives* 9 (3):3–13.

Selgin, George. 2003. Steam, Hot Air, and Small Change: Matthew Boulton and the Reform of Britain's Coinage. *Economic History Review* 56:478–509.

Shavell, Steven and Laurence Weiss. 1979. The Optimal Payment of Unemployment Insurance Benefits over Time. *Journal of Political Economy* 87 (6):1347–62.

Sims, Christopher A. 1982. Policy Analysis with Econometric Models. *Brookings Papers on Economic Activity* (1):107–152.

———. 1988. Projecting Policy Effects with Statistical Models. *Revista de Analsysis Economico* 3:3–20.

———. 1993. Rational expectations modeling with seasonally adjusted data. *Journal of Econometrics* 55 (1-2):9–19.

Sims, Christopher A. and Tao Zha. 2006. Were There Regime Switches in U.S. Monetary Policy? *American Economic Review* 96 (1):54–81.

Smith, Adam. 1789. *An Inquiry into the Nature and Causes of the Wealth of Nations*. New York: The Modern Library. The Modern library edition is copied from the fifth edition published in 1789, the last life-time edition.

Smith, Jr, A A. 1993. Estimating Nonlinear Time-Series Models Using Simulated Vector Autoregressions. *Journal of Applied Econometrics* 8 (S):S63–84.

Stokey, Nancy L. 1989. Reputation and Time Consistency. *American Economic Review* 79:134–139.

———. 1991. Credible Public Policy. *Journal of Economic Dynamics and Control* 15.

Taylor, John B. 1980. Aggregate Dynamics and Staggered Contracts. *Journal of Political Economy* 88 (1):1–23.

Tirole, Jean. 1982. On the Possibility of Speculation under Rational Expectations. *Econometrica* 50 (5):1163–81.

Townsend, Robert M. 1983. Forecasting the Forecasts of Others. *Journal of Political Economy* 91 (4):546–88.

Velde, François. 2007. John Law's System. *American Economic Review Papers and Proceedings* 97 (2):276–270.

Werning, Ivn. 2005. Tax smoothing with redistribution. Staff Report 365, Federal Reserve Bank of Minneapolis. Available at http://ideas.repec.org/p/fip/fedmsr/365.html.

Williams, Noah. 2004. Escape Dynamics in Learning Models Princeton University manuscript.

Woodford, Michael D. 1990. Learning to Believe in Sunspots. *Econometrica* 58:277–307.

———. 2003. *Interest and Prices: Foundations of a Theory of Monetary Policy*. Princeton, New Jersey: Princeton University Press.

———. 2005. Robustly Optimal Monetary Policy with Near Rational Expectations. Nber working papers, National Bureau of Economic Research, Inc.

Young, H. Peyton. 2004. *Strategic Learning and its Limits*. Oxford and New York: Oxford University Press. The Arne Ryde Memorial Lectures Series.

Zhu, Xiaodong. 1992. Optimal Fiscal Policy in a Stochastic Growth Model. *Journal of Economic Theory* 58:250–289.