

Practical Policy Evaluation*

Narayana R. Kocherlakota

April 11, 2018

Abstract

In the wake of the Lucas Critique, the study of appropriate macroeconomic policy has largely focused on the comparison of different regimes/rules. In practice, few policymakers are faced with making those kinds of choices. In this paper, I examine the problem of a policymaker making but one in a long sequence of similar decisions (like to raise or cut interest rates by a quarter percentage point). I model the policymaker as playing a dynamic game against a forward-looking private sector. My main result is that, under relatively weak conditions, the policymaker's optimal response to the current state can be found by applying statistical regression methods to past macroeconomic data. Based on this result, I argue that macroeconomic policy evaluation intended to be of practical value should rely considerably less on putatively structural macroeconomic models and considerably more on regression-based approaches.

1 Introduction

There is an enormous academic literature about macroeconomic policy evaluation under the assumption that private agents are forward-looking. Most of this literature is about the following kinds of questions:

*University of Rochester and NBER. This paper was prepared for the April 2018 Carnegie-Rochester-NYU Conference on Public Policy honoring the contributions of Charles Plosser to economics. It contains some minor corrections relative to the March 2018 version.

- What is the optimal specification of date- and state-contingent monetary/fiscal policies?
- What is the optimal specification of a monetary-fiscal policy rule that maps endogenous and exogenous variables into policy choices?
- Out of the many possible equilibria in an infinite horizon dynamic policy game, which one is optimal?

Here, the word “optimal” typically refers to “ex-ante” or “steady-state” welfare. The research addressing these kinds of questions is careful to respond to the Lucas Critique (1976) by using dynamic stochastic general equilibrium (DSGE) models that are intended to be invariant to different choices among possible policy regimes.

These broad questions are indeed important ones for those policymakers who have been asked to consider a once-and-for-all change in an overarching policy regime. But these situations are rare ones. In practice, policymakers are usually faced with a much more limited problem: what should they do today, given that their current action is only one of a long temporal chain of such choices? Thus, at fixed time intervals, central banks choose the level of short-term interest rates or long-term asset purchases. In the US, the Federal Reserve decides on an annual basis whether large banks should be allowed to pay out dividends. Democratically elected governments periodically formulate budget plans.

In this paper, I ask the following question: how should a policymaker, who is making but one in a long sequence of choices over time, use the available data to arrive at a decision? I address the question by embedding the policymaker within a standard dynamic policy game in which he or she makes choices at a point in time, taking into account the response of a forward-looking private sector to those choices. I then analyze *within the model* how the policymaker can make use of available data to formulate a best response. My main result is that, under relatively weak conditions, the policymaker’s optimal response to the current state can be estimated using theory-free statistical methods applied to past data on the state of the economy, policymaker choices, and economic outcomes (as summarized through

the policymaker's payoff). In other words, I use (game) theory to demonstrate that the policymaker doesn't need to use (macroeconomic) theory.

What do I mean by "relatively weak conditions"? I make two main assumptions. First, the policymaker's objective varies over time due to non-economic factors that, conditional on his/her available information, only affect economic outcomes through the policy choice itself. This *conditional independence restriction* serves to create a sequence of quasi-experiments that allows the policymaker to identify the impact of policy choices on his/her payoffs. Second, the variation in the objective induces the policymaker to make all possible policy choices, conditional on any observable economic states. This *full support restriction* ensures that no policy change leads the private sector to update its beliefs about the policymaker's future strategy. It also ensures that (a long string of) past data is sufficient to identify the impact of any possible current policy choice on the economy.

This argument does rely on the policymaker having access to a large number of observations of different policy choices in a similar economic state. In the paper, I show how theory can be a useful supplement to purely statistical methods when data is scarce (or, relatedly, when the conditional independence restriction is only occasionally satisfied). Even under these circumstances, though, theory is only being used as a way to impose functional form restrictions on the economy's response to policymaker choices. There is no clear sense in which the relevant theory or parameters of interest is required to be structural.

Following the seminal work of Barro and Gordon (1983), there has been a large literature on dynamic games in macroeconomics. This literature has two distinct goals. Positively, it seeks to understand properties of the data as properties of equilibria of the dynamic game. Normatively, it uses the equilibria of the dynamic game as a way to characterize the range of what can be accomplished in the absence of commitment.

My aim is different: I use dynamic game theory as a tool to help policymakers and advisors to figure out how to make choices within the context of a given equilibrium. The basic intuition of my main result is simple: To solve his/her decision problem, a player within

a game only needs to know how other players will respond to his/her choices. The decision-maker doesn't need to know *why* those other players will respond the way that they do. The point of this paper is that the government can glean the limited information that it needs about the private sector's response function using a simple (possibly nonlinear) regression.

There has been much criticism of modern macroeconomic models in the wake of the financial crisis, in terms of their treatment of financial markets, their treatment of labor markets, their treatment of heterogeneity, and their treatment of expectations. All of these criticisms are largely orthogonal to my argument. In the wake of the Lucas Critique, macroeconomists have designed their models so as to allow policymakers and their advisors to analyze the impact of regime changes. I offer no assessment of the models' value in that regard. My main message is that economic theory teaches us that most practical policy questions can be addressed with only limited reference to these models (or others). Macroeconomists - both within and without policy institutions - who want to contribute to the making of policy should be doing much more research on how best to develop the statistical methods necessary to answer these "practical policy" questions, instead of continuing to focus almost exclusively on developing models suitable for addressing "regime change" questions.

My distinction between "practical policy" questions and "regime change" questions echoes Sims' (1982) views on the limitations of the relevance of the Lucas Critique for actual policymaking. Indeed, I view one contribution of my paper as formalizing several of his core ideas.¹ Since Sims wrote, there has been a revolution in applied microeconomics in the use of atheoretical statistical methods to obtain answers to important policy questions. This paper argues that a similar change could be of value in applied macroeconomics

Plosser (2014) argues that the Federal Reserve would be well-served if it were to pre-commit to an interest rate rule of the kind described by Taylor (1993).² This paper doesn't address this argument. Instead, I simply take it as a feature of the institutional environment

¹I discuss the connection of my results to the structural vector autoregression literature initiated by Sims (1980) in Section 6.

²See Kocherlakota (2016) for a different perspective from Plosser's.

that the central bank has (transitory) discretion at any date.³ The question then is how, given that discretion, the central bank can best use the available data to achieve its objective.

2 Differences From the Lucas Critique

In this section, I describe the two main differences between the analysis of Lucas (1976) and mine. The first difference has to do with the differing nature of the policy questions. The second difference has to do with my use of more modern game theory than was available to Lucas to model private agent beliefs about public sector behavior.

2.1 Differing Policy Questions

To illustrate how the policy questions in this paper differ from those considered by Lucas, it is useful to consider a version of the permanent income hypothesis that is essentially equivalent to that in Lucas (1976). I define permanent income by:

$$y_{pt} = (1 - \beta) \sum_{s=0}^{\infty} \beta^s E_t y_{t+s},$$

where y_t represents after-tax income at date t and the information available at date t embedded in the conditional expectation consists of $\{y_{t-s}\}_{s=0}^{\infty}$. Consumption at date t is given by:

$$c_t = ky_{pt} + u_t$$

where u_t can be viewed as a mean zero, constant variance shock that is i.i.d. over time.

The consumer's pre-tax income is constant at $2a$. At each date t , the government chooses a lump-sum tax τ_t so as to maximize the objective function:

$$-(2w_t + \tau_t)^2 - (y_t - 2v_t)^2$$

³Plosser (2013) seems to do the same when discussing the challenges associated with forward guidance.

where the objective shifter w_t is a random walk, so that $(w_t - w_{t-1})$ is mean zero, constant variance, and i.i.d. over time and the objective shifter v_t is mean zero, constant variance, and i.i.d. over time.

The solution to the government's problem is to set τ_t equal to $a - v_t - w_t$. After-tax income is then given by:

$$y_t = a + v_t + w_t \tag{1}$$

In this model, as Lucas (1976, p. 26) argues, we can derive a reduced form relating consumption to current and past after-tax incomes:

$$c_t = k(1 - \beta)y_t + k\beta(1 - \lambda) \sum_{j=0}^{\infty} \lambda^j y_{t-j} + u_t \tag{2}$$

where λ is determined by the relative variances of the transitory income shock v_t and the permanent tax shock w_t .

With one key difference, to which I'll turn in a moment, the above description follows Lucas' (1976) exposition of the permanent income hypothesis. Lucas then turns to a consideration of a wide variety of different tax policies, like a permanent increase in the level of taxes or a permanent increase in their volatility. He shows that the impact of such policy changes cannot be accurately predicted via the reduced form relationship (2).

The key difference lies in how I motivated the government's choices that give rise to (2). Lucas posits that the government randomizes taxes in order to generate the law of motion (1). But he does not explain *why* the government is randomizing taxes in this fashion. In contrast, I posit a particular (random) objective for the government. The variation in taxes that leads to (2) is a consequence of the government's making its choices according to this objective.

My formalization of the policy problem imposes a great deal of discipline on the government. The government is required to make its tax choices periodically. It makes those choices according to an exogenous objective that it cannot alter. The only policy question

is: how can the government choose τ_t at each date so as to maximize its objective? None of the Lucas tax proposals are possible.

I admit that, in my desire to follow Lucas' lead, I've assumed away any dependence of the government's objective on the private sector's response to the policy choice. These responses do play a key role in my more general analysis. Later, I'll return to a modified version of the consumption tax problem that includes this effect.

2.2 Rules and Strategies

Lucas (1980b) argues that: “our ability as economists to predict the responses of agents rests, in situations where expectations about the future matter, on our understanding of the stochastic environment agents believe themselves to be operating in. In practice, this limits the class of policies the consequences of which we can hope to assess in advance to *policies generated by fixed, well understood, relatively permanent rules* (or functions relating policy actions taken to the state of the economy).” (emphasis mine). Similar considerations led him to state in the second part of the Critique that “the only *scientific* quantitative evaluations available to us are comparisons of the consequences of alternative policy rules” (emphasis his). These views have been highly influential in circumscribing the scope of macroeconomic policy evaluation among academics over the past forty plus years.

In this paper, I diverge considerably from Lucas' perspective. I don't require my hypothetical policymaker to commit in advance to a policy rule of the kind that Lucas describes. Rather, the policymaker is allowed to choose freely at each date. Lucas' words seem to suggest that there is no systematic way to analyze such a situation. But dynamic game theory - developed largely in the 1980s after Lucas wrote - provides exactly that kind of systematic approach. We shall see that in a (time-independent recursive) equilibrium the policymaker finds it optimal to follow a stable strategy (as a function of the state of the economy). This strategy is known to the agents in the economy and they make their choices, based on that knowledge. Within the equilibrium, the policymaker is seen as following a “rule” by the pri-

vate sector. But that “rule” emerges endogenously from periodic policymaker optimization, not from exogenously imposed restrictions on policymaker behavior.

How can the policymaker ever contemplate deviations from this strategy that is seen as a rule by private agents? I offer two answers to this query. First, within an equilibrium to the game, the private sector’s (commonly known) strategy describes its response to any such deviations. Second, and I believe more usefully, I allow for shocks to the policymaker’s objective that are unobservable to the private sector. These shocks are sufficiently rich that agents never see any reason to change their (concentrated) beliefs about the policymaker’s strategy.

3 A Dynamic Game and Its Equilibrium

In this section, I set forth a baseline dynamic model of periodic policy choice when the private sector is forward-looking, and describe a natural notion of equilibrium for this game.

3.1 Timing of Play

In this subsection, I describe the timing of play in a given period in a dynamic game between two players: the government and the private sector.⁴ The description is abstract, in an attempt to capture many policy choices. However, I frequently connect the high-level description back to the specific problem of a central bank choosing the short-term interest rate.

In what follows, I refer to “sets” and to “densities” over those sets. I intend for the former term to include both discrete sets and compact intervals in Euclidean spaces. I intend for the latter term to include, correspondingly, both vectors of probability weights and probability density functions.

⁴I treat the private sector as a group of agents that are required to co-ordinate their actions. The analysis can readily be extended to a model in which the private sector is a group of atomless agents who freely choose to co-ordinate on the same choice (as in Chari and Kehoe (1990)).

Stage 1: Common knowledge of persistent factors.

At the beginning of the period, a state variable $x \in X$ is common knowledge between the government and the private sector.

The state variable x captures all relevant information from prior periods that helps predict the course of the economy in the current period, conditional on the choices of the government.⁵ The information x could be about endogenous variables (such as asset prices or lagged data on output in the context of monetary policy) or about exogenous variables. The dimension of X is potentially large (so it could include microeconomic information like (approximate) distributions of net worth across households, firms, and/or financial institutions.).

Stage 2: Government learns its preferences.

The government privately observes a payoff shifter θ drawn from the set Θ according to the density $p_\theta(\cdot|x)$. This density has full support for all x .

The government privately learns the realization of a shock θ that influences its payoff function. In the monetary policy context, we could view θ as a gauge for how dovish or hawkish the central bank is. However, θ could also measure the central bank's desire for interest rate smoothing that is independent of their macroeconomic objectives. (In this case, lagged interest rates could be incorporated into the public state variable x .)

Stage 3: Government chooses an action.

The government publicly chooses g from the set G .

In this stage, the government chooses a policy action. In the context of monetary policy, the action would be the choice of a short-term interest rate target, in combination with some form of communication (like a policy statement or press release).

⁵The state variable x necessarily includes all payoff-relevant state variables. But it may also include payoff-irrelevant state variables that inform the players' strategies, like past communication or promised utility.

Stage 4: Private sector chooses an action.

The private sector publicly chooses an action a from the set A .

The private sector responds to the government's policy action. The set A could be very high-dimensional. As we'll see, the private sector's action choice also shapes the evolution of the state variable x .

I treat the private sector as a single agent. But this is without loss of generality in terms of my main results - I could readily incorporate arbitrary degrees of observable or unobservable heterogeneity into the description of the private sector. As we will see, from the point of view of the policymaker, the private sector is simply a multi-armed bandit.

Stage 5: The economic outcome is determined.

The economic outcome y is publicly drawn from the set Y according to the density $p_y(\cdot|x, g, a)$.

The (within-period) economic outcome y is determined. It is potentially affected by the current state and the actions of the government and private sector.

The taste shifter θ is only allowed to affect p_y indirectly through the government's choice of g . This conditional (on x) independence restriction is key in what follows: it means that the variation in g due to θ can be seen as a form of randomized experiment.⁶ Note that even more restrictive independence assumptions are common in the macroeconomics literature. Thus, most of the New Keynesian literature adds an error term to the monetary policy rule that is *unconditionally* independent of all other disturbances in the model.

⁶As in Imbens and Wooldridge (IW) (2007), there are actually two distinct restrictions being imposed on the taste shifter θ . The first is that, conditional on x , the random variable θ is stochastically independent of all other disturbances that influence the economic outcome y . This is what IW call a random assignment restriction. The second is that changes in the taste shifter θ have no direct effect on the economic outcome y . This is what IW call an exclusion restriction. In the settings of interest to IW, they view the exclusion restriction as much more substantive. In the context under study in this paper, I see the random assignment restriction as likely to be the one with more content.

Stage 6: Period payoffs are realized.

The government realizes its within-period payoff $w(\theta, g, y)$. The private sector realizes its within-period payoff $u(x, a, y)$.

The private sector payoff is a function of its action and the economic outcome. The government payoff is a function of the economic outcome, its choice, and its taste shifter. The exclusion of the state variable x from the policymaker's payoff function is only for simplicity, and has no material effect on the analysis.

It is often argued that governments need economic models to know what they should maximize. The government's payoff here should be understood as being determined by statutory requirements and political factors, not by economic reasoning. Thus, the Federal Reserve seeks to stabilize inflation and unemployment at their long-run levels because of Congressional mandates.

Stage 7: Next period's state is realized.

Next period's state x' is publicly drawn from the set X , according to the density $p_{x'}(\cdot|x, a, y)$. This density has full support over X for all (x, a, y) .

The state evolves in response to the private sector choices and outcomes. The exclusion of g as an influence on $p_{x'}$ is only for simplicity and doesn't affect the results.

3.2 Building an Equilibrium

Given the above within-period timing, I now turn to the construction of an equilibrium.⁷

⁷The equilibrium concept is Bayesian-Nash. (Perfection has no additional bite because all information sets of the two players have positive probability.) I focus further on equilibria that are pure strategy, time-independent, and that depend on the past only through the state variable x . This formulation of equilibrium is a natural analog of the "recursive equilibrium" concept used extensively by macroeconomists (see Ljungqvist and Sargent (2018) for many applications). It is distinct from (Bayesian-)Markov equilibrium because the state x may include payoff-irrelevant variables.

3.2.1 Private Sector Strategy

I start with the private sector. It is *forward-looking*, in the sense that its overall payoff in a given period is given by the expectation of:

$$u(.) + \delta U', 0 < \delta < 1$$

where the within-period payoff u was defined above and U' is the discounted sum of future within-period payoffs. Recall that the private sector's action choice affects the future state x' of the economy. This ability to influence the future, combined with the forward-looking nature of its objective, means that the private sector's action in a given state depends on its beliefs about how future governments will respond to the evolution of the state. In equilibrium, the government uses a stable (pure) strategy that maps the public state x and its private information θ into an element of G . Consistent with this equilibrium restriction, the private sector believes that the government uses a strategy $\hat{\gamma}$ of the form $\hat{\gamma} : X \times \Theta \rightarrow G$.

Given those beliefs, and given its observations of the state x and a government action choice g , the private sector's problem in any period is given by:

$$\max_a \int_Y \{u(x, a, y) + \delta \int_X V(x'; \hat{\gamma}) p_{x'}(x'|x, a, y) dx'\} p_y(y|x, g, a) dy$$

Here, $V(x'; \hat{\gamma})$ represents the private sector's continuation value, given that it observes the public state x' and believes that the government has a time-invariant strategy $\hat{\gamma}$. Within a period, the government chooses its action before the private sector does. As a result, the private sector beliefs about the government's strategy only matter because the private sector is forward-looking, which means that $\hat{\gamma}$ affects the private sector's problem through the continuation value function.

The solution to this maximization problem depends on the private sector's information (x, g) and on the private sector's beliefs $\hat{\gamma}$ about the government's strategy. Hence, the

solution can be denoted $\alpha^*(x, g; \hat{\gamma})$.

The value function V satisfies a natural recursive relationship. First, we define $\bar{u}(x; \hat{\gamma})$ to be the expectation of the private sector's within-period payoff:

$$u(x, \alpha^*(x, \hat{\gamma}(x, \theta); \hat{\gamma}), y)$$

over the randomness generated by (θ, y) . (Note this expectation is based on the private sector's belief that, in the future, it will use the strategy α^* .) We calculate that expectation by using the product of the conditional densities:

$$p_y(y|x, \gamma(x, \theta), \alpha^*(x, \hat{\gamma}(x, \theta); \hat{\gamma})) \\ p_\theta(\theta|x)$$

Then, the value function $V(x; \hat{\gamma})$ is equal to the expectation of:

$$\bar{u}(x; \hat{\gamma}) + \delta V(x'; \hat{\gamma}).$$

over the randomness generated by x' , conditional on x . We calculate the expectation using the product of the conditional densities:

$$p_{x'}(x'|x, \alpha^*(x, \hat{\gamma}(x, \theta); \hat{\gamma}), y) \\ p_y(y|x, \hat{\gamma}(x, \theta), \alpha^*(x, \hat{\gamma}(x, \theta); \hat{\gamma})) \\ p_\theta(\cdot|x)$$

3.2.2 Full Support

The private sector's beliefs about the government's strategy are concentrated on the single function $\hat{\gamma}$. This is how Nash equilibrium and its refinements work: every player's strategy is common knowledge. It follows that the private sector will never change its beliefs about the

government's strategy *as long as the government's (public) actions have positive probability given the information known to the private sector*. That kind of consistency is ensured if the private sector's belief $\hat{\gamma}$ has full support in the sense that:

$$\forall x \in X \text{ and } \forall g \in G, \exists \theta \in \Theta \text{ s.t. } \hat{\gamma}(x, \theta) = g$$

In words, the strategy $\hat{\gamma}$ has full support if the private sector can always rationalize any observed action choice g as being due to a particular realization of the government's private information about its payoffs.⁸

The full support requirement is important because it serves to limit the history dependence of equilibrium strategies. Suppose $\hat{\gamma}$ didn't satisfy the full support requirement. Then, in any period, the private sector might see a history of government choices that is inconsistent with the government's having always used the strategy $\hat{\gamma}$. In any such history (and in any continuation of any such history), the private sector's beliefs about future government choices may differ from $\hat{\gamma}$. The private sector's strategy could then depend on the full history of government choices, not just on the current choice. In what follows, I simply assume that the nature of the government's objective, and the shocks to that objective, imply that it is a best response for the government to use a strategy that has full support.

3.2.3 The Government's Strategy

I turn next to the government. For now, I'll assume that the government is myopic in each period; I'll discuss the (minor) consequences of relaxing that assumption later.

Because the private sector makes its choices after the government in each period, the government must form beliefs about the private sector's strategy (that is, how it will respond to various choices of the government). I assume the government believes that the private

⁸In the next section, I introduce the possibility that the government has information about the economy that is not known to the private sector. In this context, the full support restriction is weaker: it only requires that the private sector can always rationalize any choice by the government as being due to some realization of the government's private information about its payoffs *or the economy*.

sector's strategy is $\hat{\alpha}$, where:

$$\hat{\alpha} : X \times G \rightarrow A$$

Thus, the government's problem in any period, given its observations of (x, θ) , is to choose g to solve:

$$\max_{g \in G} \int_Y w(\theta, g, y) p_y(y|x, g, \hat{\alpha}(x, g)) dy$$

The resulting strategy can be written as a function $\gamma^*(x, \theta; \hat{\alpha})$ of the government's information (x, θ) , given the government's beliefs $\hat{\alpha}$ about the private sector's strategy.

3.2.4 Equilibrium

An equilibrium is a government strategy γ^* and a private sector strategy α^* such that:

- $\gamma^*(x, \theta; \alpha^*)$ solves the government's problem for all (x, θ) , given that the government believes that the private sector uses strategy α^* .
- $\alpha^*(x, g; \gamma^*)$ solves the private sector's problem for all (x, g) , given that the private sector believes that the government uses strategy γ^* and the private sector's continuation value function $V(\cdot; \gamma^*)$ is defined as above.
- γ^* has full support (so that, for any x in X and any g in G , there exists θ in Θ such that $\gamma^*(x, \theta) = g$).

3.3 Discussion

The above describes a dynamic game between the government and the private sector, in which both make periodic choices. Within the equilibrium to the game, the government and the private sector follow time-invariant strategies (that map their payoff-relevant information into actions). These strategies are commonly known.

Within this equilibrium, given its observations (x, θ) , the government contemplates choosing among all possible g in G . In his Critique, Lucas expresses concern that deviations from

past behavior on the part of the government will lead to confusion in the private sector about the government's future strategy. But within an equilibrium to the above game, this potential confusion simply can't happen: it is common knowledge between the players how the private sector updates its beliefs in response to any government action choice. This statement is true even if the government's strategy doesn't have full support. With full support, the relevant updating is considerably simplified: no matter what the government does, the private sector *never* changes its beliefs about the government's future strategy.

All persistent factors in the economy are embedded in the publicly observable state variable x . Later, I'll discuss the consequences of relaxing that assumption, so that, conditional on the public state, the government's (privately known) payoff shifter θ is persistent. As we shall see, that persistence makes the private sector's beliefs about θ into a key state variable in the problems of both the private sector and the government.

4 Econometrics of the Government's Problem

The government enters period t . It observes (x_t, θ_t) . Recall that, in equilibrium, the government's problem is to solve:

$$\max_{g \in G} \Pi_G(g; x_t, \theta_t, \hat{\alpha})$$

where:

$$\Pi_G(g; x_t, \theta_t, \hat{\alpha}) = \int_Y w(\theta_t, y, g) p_y(y|x_t, g, \hat{\alpha}(x_t, g)) dy$$

Here, $\hat{\alpha}$ represents the private sector's response to government choices and p_y is the (exogenous) density of the economic outcome, conditional on government/private sector choices.

In this section, I now suppose that the government in period t doesn't know p_y or $\hat{\alpha}$. Rather, it has an infinite data set:

$$(x_{t-s}, y_{t-s-1}, g_{t-s-1})_{s=0}^{\infty}$$

of current and past realizations of the public state x , past realizations of the economic outcome y , and the government's past actions. (In a later section, I consider some of the important and interesting complications associated with having a finite sample.) The data set doesn't (need to) include the past realizations of the government's payoff shifter θ . However, I do suppose that the government knows θ_t - that is, its current within-period payoff as a function of its choice and the economic outcome (g, y) .

In this section, I address the question of how the government can use these data to solve its problem. I demonstrate the main result in the paper: the decision problem can be solved by the government's:

- using observations in which x is close to x_t
- running a nonlinear regression of its past payoffs (evaluated using the current realization of θ_t) on its past choices of g
- and then choosing the value of g that maximizes this regression function.

Thus, the government can solve its problem through purely statistical methods without any knowledge of structural elements like p_y or u .

4.1 Conditionally Random Government Choices

The government's objective is the expectation of its payoff, given (x_t, θ_t) , for each possible choice of g . The goal is to use past data to estimate this conditional expectation. In order to do so, I assume that the private sector's past strategy α^* is the same as the private sector's current strategy $\hat{\alpha}$. (This stationarity assumption is non-testable using the available data, but is a property of equilibrium as defined above.)

The bigger issue concerns a potential endogeneity problem. The government's various past choices of g are associated with the variables (x, θ) differing from the current observation (x_t, θ_t) . These different realizations of (x, θ) are, in turn, associated with different conditional

densities for y . The government needs to distinguish the possible effects of (x, θ) on y from the effects of g itself.

However, in the description of stage 5 of the game, we imposed the restriction that the density of the economic outcome y , conditional on the state x_t , only depends on the payoff shifter θ_t through the government's choice of g . This conditional independence restriction implies we can rewrite the government's period t objective for any g as:

$$\Pi_G(g; x_t, \theta_t, \hat{\alpha}) = E(w(\theta_t, \gamma^*(x, \theta), y) | x = x_t, \gamma^*(x, \theta) = g; \hat{\alpha})$$

where the expectation is over the realizations of (θ, y) and γ^* is the past strategy of the government.⁹ This converts the government's policy projection problem (what should be projected to happen if it chooses g ?) into a conditional expectation (what should be expected to happen in the event that its strategy implies a choice of g ?). In the remainder of this section, I show how this result implies that the government can use regression methods to find the solution to its problem, even though it doesn't know p_y or $\hat{\alpha}$.

4.1.1 Finite Sets

In this subsection, I assume that the sets (X, G) are both finite.¹⁰

Let $S_g = \{s_n\}_{n=1}^\infty$ be a subset of the natural numbers such that:

$$(x_{t-s_n}, g_{t-s_n}) = (x_t, g).$$

The set S_g is the subset of past data in which the realization of x was the same as in the

⁹The “for any g ” part of this statement relies on the government's past strategy γ^* having full support. This restriction is satisfied within an equilibrium and is verifiable using the infinite sample.

¹⁰In the case that G has two elements, this subsection becomes a special case of the vast treatment effects literature. The policy decision $g = 1$ can be viewed as the economy being assigned to “treatment” while the policy decision $g = 0$ can be viewed as an assignment to “no treatment”. The conditional independence assumption on θ is then that the decision to “treat”, conditional on x , is independent of all other factors that affect the economic outcome. As is well-known, under this assumption, the expected impact of treatment relative to non-treatment is equal to the difference in the average outcomes between treated and non-treatment.

current period and the government's action choice was g . The full support assumptions made about $p_{x'}$ (in stage 7) of the dynamic game, combined with the full support assumption about the government's past strategy γ^* , ensure that S_g has an infinite number of observations.

The sequence $\{y_{t-s_n}\}_{n=1}^{\infty}$ is a countably infinite sequence of i.i.d. draws from the density:

$$p_y(y|x_t, g, \hat{\alpha}(x_t, g))$$

Consider the sample average:

$$\bar{w}_g = \lim_{N \rightarrow \infty} \sum_{n=1}^{\infty} w(\theta_t, g, y_{t-s_n})/N$$

This infinite-sample average of i.i.d. draws is equal to the expectation of the government's within-period payoff, conditional on x_t and the government's choice being g . Hence, we can conclude that:

$$\begin{aligned} \bar{w}_g &= E(w(\theta_t, \gamma^*(x_t, \theta), y) | x = x_t, \gamma^*(x_t, \theta) = g; \hat{\alpha}) \\ &= \Pi_G(g; x_t, \theta_t, \hat{\alpha}) \end{aligned}$$

The government can then choose g by looking for the value of g that delivers the highest \bar{w}_g .

To summarize: to evaluate a particular option g , the government:

1. discards all past observations except those with the same x as it observes today and with the same g as it is considering. These past observations likely are based on a different θ than the current θ_t . But that different θ doesn't affect the empirical density of the y 's.
2. transforms the economic outcome y for each observation in the retained data into a payoff w , using the current realization of θ and the g under consideration.
3. average the resulting w 's.

The government then picks the option g that delivers the highest average in step 3.

4.1.2 Interval Case

I now turn to the case in which the sets (X, G) are compact intervals in Euclidean spaces.

Given the current state x_t , let $\Lambda(\epsilon)$ be a neighborhood of size ϵ around (x_t) . Let $S_\epsilon(x_t) = \{s_n\}_{n=1}^\infty$ be the subset of the natural numbers such that $(x_{t-s_n}) \in \Lambda(\epsilon)$. As argued above, $S_\epsilon(x_t)$ has an infinite number of observations. To simplify the argument, I assume that ϵ is sufficiently small that we can view the density of (θ, y) , conditional on x , as being the same for all x in $\Lambda(\epsilon)$.

Consider the data set:

$$(W_n, g_{t-s_n})_{n=1}^\infty$$

where:

$$W_n \equiv w(\theta_t, g_{t-s_n}, y_{t-s_n})$$

Note that the payoff W_n is defined using the current realization of θ . Given the assumption that $(\theta_{t-s_n}, y_{t-s_n})$ is independent of (x_{t-s_n}) in $\Lambda(\epsilon)$, it follows that $\{W_n, g_{t-s_n}\}_{n=1}^\infty$ is a sequence of i.i.d. vectors. We can use this data set to estimate a nonlinear regression of W on g :

$$W = h(g; x_t) + u$$

where u is mean-independent of g . Note that, given x_t , the regression function h is identified because of the full support restriction on the government's strategy.

The regression function $h(g; x_t)$ measures the expectation of the government's payoff,

conditional on x equalling x_t and the government's choice $\gamma^*(x_t, \theta_t)$ equalling g . Hence:

$$\begin{aligned} h(g; x_t) &= E(w(\theta_t, \gamma^*(x_t, \theta), y) | x = x_t, \gamma^*(x_t, \theta) = g; \hat{\alpha}) \\ &= \Pi_G(g; x_t, \theta_t, \hat{\alpha}) \end{aligned}$$

It follows that the government can solve its problem by choosing the value of g that optimizes the regression function $h(g; x_t)$.

4.1.3 A Potential Role for Theory

In the prior two subsections, I described how the government can estimate its objective function Π_G using a nonparametric regression. These nonparametric methods essentially threw out all past data in which the realization of x differed (too much) from the current observation x_t . I now turn to how the restrictions implied by theory or other auxiliary information can potentially allow the government to use past data when the state x differs markedly from its current value. These kinds of restrictions will turn out to be helpful when the government's data is more limited or when the conditional independence restriction is only satisfied for some values of x .

To recall (yet again), the government's objective function takes the form:

$$\Pi_G(g; x_t, \theta_t, \hat{\alpha}) = \int_Y w(\theta_t, y, g) p_y(y | x_t, g, \hat{\alpha}(x_t, g; \hat{\gamma})) dy$$

when its information set is (x_t, θ_t) . Here, $\hat{\alpha}$ is the private sector's strategy determined by its optimization problem. In many situations, theory (combined with other data/information) could be used to restrict the form of the private sector's objective. There may also be theoretical or data-based restrictions on the conditional density p_y that describes the response of the economic outcome to the actions of the government and public sector. Collectively, these various restrictions serve to connect the government's objective function at different realizations of (x, θ) . At an abstract level, we can capture these connections by requiring Π_G

to be indexed by a low-dimensional vector $\beta \in B$, so that, given θ_t , there exists a function $\Upsilon : B \times X \times G \rightarrow \mathbb{R}$ such that:

$$\Upsilon(\beta, \bar{x}, g) = \int_Y w(\theta_t, y, g) p_y(y|\bar{x}, g, \hat{\alpha}(\bar{x}, g; \hat{\gamma})) dy$$

for all (\bar{x}, g) . As above, the conditional independence restriction implies that:

$$\Upsilon(\beta, \bar{x}, g) = E(w(\theta_t, \gamma^*(x, \theta), y) | x = \bar{x}, \gamma^*(x, \theta) = g; \hat{\alpha})$$

for all (\bar{x}, g) , so that $\Upsilon(\beta, \cdot)$ is the nonlinear regression function of w on g , given that $x = \bar{x}$.

There are many ways to estimate β given the large data set available to the government. Here's one that is based on method of moments. Form the data set $(W_n, g_{t-n}, x_{t-n})_{n=1}^{\infty}$, where:

$$W_n = w(\theta_t, g_{t-n}, y_{t-n})$$

The government can then look for the value of $\hat{\beta}$ in B that satisfies moment conditions of the kind:

$$N^{-1} \sum_{n=1}^N [W_n - \Upsilon(\hat{\beta}, x_{t-n}, g_{t-n})] 1_{\{(x_{t-n}, g_{t-n}) \in \Phi\}} = 0$$

for various Φ that are (Borel) subsets of $(X \times G)$. With that estimate in hand, the government can solve its problem by finding the value of g that maximizes $\Upsilon(\hat{\beta}, g, x_t)$ (that is, by setting $x = x_t$).

In this fashion, theory can play a role in helping the government solve its problem. To be clear, that role is certainly not a star turn. The government is not looking to build an elaborate theoretical edifice that can be described as structural in some sense. In particular, the government is completely uninterested in the economic mechanisms behind the past data on (g, x) . Rather, it is using theory as a source of plausible functional form restrictions.

These kinds of auxiliary restrictions can also make it possible for the government to proceed when the conditional independence restriction only holds in some states. In particular,

suppose that:

$$p_y(\cdot|x, \theta, g, a)$$

is independent of θ when x is in E , but may depend on θ if x is not in E . In this case, the variation in θ is only generating a policy “experiment” for those x in E . The key question is whether it is possible to identify β using only the data from these “experiment” states - that is, is there a unique β^* in B such that:

$$\Upsilon(\beta, \bar{x}, g) = \int_Y w(\theta_t, y, g) p_y(y|\bar{x}, g, \hat{\alpha}(\bar{x}, g; \hat{\gamma})) dy$$

for all (\bar{x}, g) such that \bar{x} is in E . If so, the government can discard the past observations in which the economic state x_{t-n} is not in E , and then use the remaining data, as above, to estimate β^* . It can then find the optimal policy choice by maximizing $\Upsilon(\beta^*, x_t, g)$, even if x_t is not in E . The government is using the additional restrictions provided by theory (or other) information to substitute for the lack of experimental variation in g .

4.1.4 Linear-Quadratic Case

In this subsection, I illustrate the value of the theoretical restrictions via a linear-quadratic example. It serves as the basis of the two specific economic cases that I analyze later.

Suppose it is known that the public state $\{x_t\}_{t=1}^{\infty}$ follows an exogenous stochastic process. As well, it is known that the economic outcome y is an affine function of (x, g, a) and a disturbance term:

$$y = \beta_0 + \beta_x x + \beta_g g + \beta_a a + \varepsilon$$

where ε has mean zero and constant variance. The government’s payoff is given by:

$$-(g - \theta)^2 - (y - \bar{y})^2$$

where θ is a non-degenerate shock with mean bx that is conditionally (on x) independent of

ε . Finally, the government believes that the endogenous private sector strategy is linear:

$$\hat{\alpha}(x, g; \hat{\gamma}) = \alpha_0 + \alpha_x x + \alpha_g g$$

This belief is correct if the private sector's objective is quadratic in (x, a, y) and it believes that the government's strategy is using a strategy that is linear in (x, θ) .

Given the linear strategy of the private sector, the economic outcome is affine in the current state x and the government's choice g :

$$y = \beta'_0 + \beta'_x x + \beta'_g g + \varepsilon$$

Suppose that, in the past, the government's strategy has put positive weight on θ . Then, x and g are imperfectly correlated. The government can consistently estimate β' by regressing y on (x, g) . It can then choose g so as to maximize its objective, taking this estimated response as given:

$$\max_g - (g - \theta)^2 - (\beta'_0 + \beta'_x x + \beta'_g g - \bar{y})^2$$

Note the solution to this maximization problem does in fact put nonzero weights on both (x, θ) . So, assuming that the government has always behaved optimally in the past, the government can identify the effect of g on y by regressing y on (g, x) .

4.1.5 Summary

In the context of the dynamic policy game described in the prior section, I contemplated a government with an infinite amount of data about its past choices (g), past states (x) and past economic outcomes (y). I showed that the government can solve for its best response to the private sector's strategy using a *purely statistical* two-step procedure. First, it regresses past economic outcomes (transformed through its current payoff function $w(\theta_t, \cdot)$) on its past choices. Second, it finds the choice of g that maximizes this estimated regression function.

The procedure works (in the sense of solving the government’s decision problem in any period) in the context of the equilibrium defined in the prior section. But it’s more generally applicable: the government can use this procedure to find its best response as long as:

- the private sector’s current strategy is the same as its past strategy
- and the government’s past strategy has full support conditional on publicly observable information

The key requirement is that, in the past data, the government’s action varies over its action set G because of a taste shifter that, conditional on the publicly observable state x_t , only affects the economic outcome through the government’s policy choice. This conditional independence restriction allows the government to use the estimated regression function as a way to do conditional policy projections.

4.2 Additional Government Information

The above analysis assumed that the government’s information about the economy was limited to what is publicly observable (x). Suppose instead that the government also privately observes a factor ξ such that the conditional density p_y depends in some fashion on ξ :

$$p_y(y|x, \xi, g, a)$$

and that ξ is (imperfectly) correlated with θ . This interaction between the government’s taste shifter and the conditional density of y would eliminate the (conditional on x) “experimental” nature of the government’s past policy choices.

In this case, we need to expand the definition of the state to include ξ . In particular, suppose that, conditional on (x, ξ) , θ affects y only through g and that, conditional on (x, ξ) , the government’s strategy $\gamma^*(x, \theta, \xi)$ has full support over G . Then, the government can, as above, solve its problem by using data in which (x, ξ) is close to (x_t, ξ_t) and regressing its past

payoffs on its past choices.¹¹ The conditioning in the conditional independence restriction is with respect to the government’s information set.

4.3 Why Not “Structural”?

The government can identify its best response using a simple regression approach. Why doesn’t it need or want to use a more structural approach?

As section 1.2.1 describes, the private sector’s strategy depends on its beliefs about the government’s strategy. In much of academic macroeconomics, it is presumed that the government can influence those private sector beliefs through credible announcements or overt commitments. Most academic analyses of policy hinge then on how the private sector’s strategy would change if its beliefs about the government’s strategy were to change. But gauging that strategic response requires the government to know many details about the private sector, like its within-period payoff function, its discount factor, and the nature of state-to-state transitions. More broadly, the government needs to know elements of the private sector’s problem that are *structural*, in the sense that they are invariant to its beliefs about the government’s strategy.

In contrast, I consider the government’s decision problem within an equilibrium (as defined in the prior section). In this decision problem, the government has no ability to influence the beliefs or strategy of the private sector. Instead, the government’s problem is to choose its best action taking those beliefs/strategy as given. As long as those beliefs/strategy don’t change over time, this problem can be effectively addressed using purely statistical methods.

4.4 Why the Rules Approach is Flawed

It is common to recommend that the government should use the past data in a different fashion: estimate its past strategy/rule and then play according to that rule. But there are

¹¹In this case, the government’s choice of g typically conveys information to the private sector about ξ . This information about y influences the private sector’s decision problem. However, in equilibrium, the private sector finds it optimal to follow a strategy that is a time-invariant function of (x, g) .

two problems with this approach.

First, what does it mean to “estimate” the past rule? Presumably, it means using a (possibly nonlinear) regression to decompose the government’s past policy choices:

$$g_{t-n} = \pi(x_{t-n}) + \epsilon_{t-n}$$

into a systematic part (due to the publicly observable information x) and a non-systematic part ϵ , which is mean-independent of x . The recommendation would then be to play according to the rule $\pi(x_t)$ today. But this recommendation ignores the government’s taste shifter θ and the government’s private information ξ about the economy. There is a reason why the error term ϵ exists in the past data, and the government needs to take account of that when making its current decision.

Second, and perhaps more importantly, there is no reason why the current government should assume that its past behavior was in fact optimal. The rules-based approach locks the government into repeating any past mistakes. The statistical approach described above instead allows the government to improve on any past errors.

5 Examples

In this section, I provide two linear-quadratic examples of the more general analysis in the prior section. The first is a simple permanent income hypothesis example of the kind discussed in Section 2. The second is a New Keynesian monetary example. The main conclusion in both examples is that the government can use a reduced-form regression to solve its within-equilibrium decision problem without knowing the “structural” or “deep” parameters.

I don’t require the private sector’s beliefs $\hat{\gamma}$ about the government’s strategy to be the same as the government’s actual strategy γ in either example. I could add this (equilibrium) restriction without affecting the argument. What matters is that the private sector believes that the government’s strategy has full support. This aspect of the private sector’s beliefs

ensures that they are independent of the government's choices.

5.1 Consumption Tax Example

I posit that pre-tax income (relative to its mean) follows:

$$e_t = e_{t-1} + \eta_t$$

where $\{\eta_t\}_{t=1}^{\infty}$ is i.i.d.. Current consumption (relative to its mean) is given by:

$$c_t = ay_t^p + by_t$$

where y_t is after-tax income and:

$$y_t^p = (1 - \beta)E_t \sum_{s=0}^{\infty} \beta^s y_{t+s}$$

is permanent after-tax income. The government's objective is $-c_t^2 - (\tau_t - u_t)^2$, where u_t is i.i.d. over time.

The private sector assumes that the government will use a strategy of the form:

$$\tau_t = \hat{A}e_t + \hat{B}u_t$$

in the future. Given those beliefs, the private sector's consumption in period t is given by a linear function¹²:

$$c_t = A_{PS}e_t + B_{PS}\tau_t$$

The key to the government's problem is that this function is invariant to its choice of τ_t . This is because the private sector ascribes different τ_t 's, conditional on e_t , to the government's having seen different realizations of u_t .

¹²Specifically, $A_{PS} = a(1 - \beta\hat{A}) + b$ and $B_{PS} = -(a(1 - \beta) + b)$.

The government regresses past c_t on past (e_t, τ_t) to uncover the coefficients (A_{PS}, B_{PS}) in the private sector's linear strategy. It then optimizes:

$$-(B_{PS}\tau_t + A_{PS}e_t)^2 - (\tau_t - u_t)^2$$

with respect to τ_t . This gives rise to a linear function¹³:

$$\tau_t = Ae_t + Bu_t$$

where $B \neq 0$. Thus, as surmised by the private sector, the government uses a linear strategy that puts non-zero weight on both (e_t, u_t) .

Note that, in solving its decision problem, the government only needs to know the reduced-form coefficients (A_{PS}, B_{PS}) . It doesn't need to know how the private sector's strategy depends on its beliefs \hat{A} .

5.2 Monetary Policy Example

In this example, I consider a simple log-linearized New Keynesian model in which:

$$y_t = E_t y_{t+1} - \sigma(R_t^{CB} - r_t^{nat} - E_t \pi_{t+1})$$

$$\pi_t = \kappa y_t + \mu_t$$

$$E_t r_{t+1}^{nat} = (1 - \rho_{nat})\bar{r} + \rho_{nat} r_t^{nat}$$

$$E_t \mu_{t+1} = \rho_\mu \mu_t$$

Here, the variable r_t^{nat} is the natural real rate of interest and μ_t is the firms' mark-up shock; they're exogenous variables that are commonly known to evolve according to first-order autoregressive processes. (However, the coefficients of these processes are not known to the government.) The variables (y_t, π_t) represent log-linearized output gaps and inflation; they're

¹³More specifically, $A = -A_{PS}B_{PS}/(1 + B_{PS}^2)$ and $B = 1/(B_{PS}^2 + 1)$.

endogenous choices of the private sector in period t . The variable R_t^{CB} is the nominal interest rate to be chosen by the central bank in period t . (I've left out the forward-looking part of the New Keynesian Phillips curve. This can be rationalized through the inflexible firms' setting their prices equal to last period's price level, as opposed to their own prices.)

I assume that the central bank wants to maximize:

$$-y_t^2 - \pi_t^2 - (R_t^{CB} - \bar{r} - \epsilon_t)^2$$

Here, besides eliminating output gaps and non-zero inflation, the central bank wants to keep nominal interest rates stable around a target level that is subject to mean zero shocks in the form of ϵ_t . These shocks are i.i.d. over time and have support equal to the real line.

I first solve for the representative household's strategy, as a linear function of the state variable (μ_t, r_t^{nat}) and the central bank's choice R_t^{CB} . To that end, re-write the representative household's Euler equation in terms of y as:

$$y_t = -\sigma R_t^{CB} + \sigma r_t^{nat} + \sigma \rho_\mu \mu_t + (1 + \sigma \kappa) E_t y_{t+1}$$

Suppose that the household believes that the central bank's period $(t + 1)$ strategy is linear:

$$R_{t+1}^{CB} = \hat{A}_0 + \hat{A}_1^{CB} r_{t+1}^{nat} + \hat{A}_2^{CB} \mu_{t+1} + \hat{A}_3^{CB} \epsilon_{t+1}$$

where $\hat{A}_3^{CB} > 0$. Given these beliefs, there is a time-invariant linear output strategy:

$$y_t = A_0^H + A_1^H r_t^{nat} + A_2^H \mu_t + A_3^H R_t^{CB}$$

that satisfies the Euler equation.¹⁴

Next, I assume that the central bank has an infinite data set of current and past observations of the exogenous variables $(r_{t-s}^{nat}, \mu_{t-s})_{s=0}^{\infty}$ and past observations of the endogenous variables $(y_{t-s}, R_{t-s}^{CB})_{s=1}^{\infty}$. Given these data, the central bank regresses y_{t-s} on contemporaneous $(r_{t-s}^{nat}, \mu_{t-s}, R_{t-s}^{CB})$. This regression identifies the household's strategic parameters $(A_0^H, A_1^H, A_2^H, A_3^H)$ as long as the central bank's past strategy puts non-zero weight on ϵ_{t-s} , so that there is variation in R_{t-s}^{CB} that is independent of the variation in the state variables $(r_{t-s}^{nat}, \mu_{t-s})$. The central bank can identify κ by regressing π_{t-s} on y_{t-s} and μ_{t-s} .

Given that information from its past data, the central bank can find its optimal R_t^{CB} in the current period by solving the maximization problem:

$$\max -(R_t^{CB} - \bar{r} - \epsilon_t)^2 - (\kappa y_t + \mu_t)^2 - y_t^2$$

subject to the constraint that $y_t = A_0^H + A_1^H r_t^{nat} + A_2^H \mu_t + A_3^H R_t^{CB}$. The solution to this optimization problem is:

$$R_t^{CB} = A_0^{CB} + A_1^{CB} r_t^{nat} + A_2^{CB} \mu_t + A_3^{CB} \epsilon_t$$

where A_3^{CB} is non-zero. Hence, the central bank does find it optimal to use a strategy with

¹⁴In particular, the parameters $(A_0^H, A_1^H, A_2^H, A_3^H)$ satisfy:

$$\begin{aligned} & A_0^H + A_1^H r_t^{nat} + A_2^H \mu_t + A_3^H R_t^{CB} \\ = & -\sigma R_t^{CB} + \sigma r_t^{nat} + \sigma \rho_\mu \mu_t + (1 + \sigma \kappa)[A_0^H + A_1^H((1 - \rho_{nat})\bar{r} + \rho_{nat} r_t^{nat}) + A_2^H \rho_\mu \mu_t + A_3^H E_t R_{t+1}^{CB}] \\ = & -\sigma R_t^{CB} + \sigma r_t^{nat} + \sigma \rho_\mu \mu_t + (1 + \sigma \kappa)[A_0^H + A_1^H((1 - \rho_{nat})\bar{r} + \rho_{nat} r_t^{nat}) + A_2^H \rho_\mu \mu_t \\ & + (1 + \sigma \kappa)A_3^H(\hat{A}_0^{CB} + \hat{A}_1^{CB}((1 - \rho_{nat})\bar{r} + \rho_{nat} r_t^{nat}) + \hat{A}_2^{CB} \rho_\mu \mu_t)] \end{aligned}$$

This equation implies, via the method of undetermined coefficients, that:

$$\begin{aligned} A_0^H &= (1 + \sigma \kappa)[A_0^H + A_1^H(1 - \rho_{nat})\bar{r} + A_3^H \hat{A}_0^{CB} + A_3^H \hat{A}_1^{CB}(1 - \rho_{nat})\bar{r}] \\ A_1^H &= \sigma + (1 + \sigma \kappa)[A_1^H \rho_{nat} + A_3^H \hat{A}_1^{CB} \rho_{nat}] \\ A_2^H &= \sigma \rho_\mu + (1 + \sigma \kappa)[A_2^H \rho_\mu + A_3^H \hat{A}_2^{CB} \rho_\mu] \\ A_3^H &= -\sigma \end{aligned}$$

full support.

As in the prior example, the government can solve its problems using only the reduced-form coefficients $(A_0^H, A_1^H, A_2^H, A_3^H)$. It doesn't need to know those coefficients relate to the private sector's beliefs $(\hat{A}_0^{CB}, \hat{A}_1^{CB}, \hat{A}_2^{CB}, \hat{A}_3^{CB})$ or other more structural aspects of the environment.

6 Discussion

In this section, I discuss some extensions and relate the above discussion to prior literature.

6.1 Relationship to Structural Vector Autoregressions

The statistical approach that I discuss above is reminiscent of what has been done in the structural vector autoregression (SVAR) literature initiated by Sims (1980). But I see at least three important differences between that literature¹⁵ and my own discussion.

First, I provide a game-theoretic justification for why a government faced with a periodic policy choice can use a purely statistical approach to estimate its best response at each date.¹⁶ I have not seen that done for the SVAR approach.

Second, the SVAR literature typically is based on the perspective of an economist who does not know the information set of the government. Thus, from the SVAR perspective, the relevant notion of a shock in this model is the unanticipated component of government policy:

$$\epsilon_t \equiv g_t - E(g_t|x_t)$$

given the information x_t available to the public sector. But, as discussed in Section 4.2, the government might have more information (ξ) about the evolution of the economy than what is in x . This kind of information would, from the SVAR perspective, make ϵ a “non-structural”

¹⁵See Stock and Watson (2016) for an extensive recent summary.

¹⁶In this sense, my analysis provides a formal conceptualization of what Leeper and Zha (2003) term “modest” policy interventions.

innovation.

In contrast, in this paper the relevant perspective is that of the policymaker and his/her advisers. It seems plausible that they will know their own information (that is, what's included in the public shock x_t and any private shock ξ_t) about the economy, and so be able to tell what past policy variation can be ascribed to that information. To be clear, I don't mean to suggest that the conditional independence restriction imposed on θ in this paper is non-substantive - but it does seem considerably more plausible than the typical identification restrictions in the SVAR literature. (This policymaker-oriented attack on the identification problem is similar to the narrative approach to identification used by Romer and Romer (1989).)

Finally, the SVAR literature is overly ambitious in terms of its goals relative to what is needed here. The SVAR literature typically attempts to identify the response of a large number of economic variables to an identified policy shock. Here, because I am focused on the policymaker's problem, there is only one variable of interest: the policymaker's payoff. As we saw when we discussed the role of theory in Section 4.1.3, the policymaker has no interest in the statistical or economic details of how past data on its past choices or the economic state were generated. This wedding between estimation and decision-making will play an even larger role in the analysis when I discuss finite sample issues in the next section.

6.2 Hidden Persistence in Government Tastes

In the description of the game, the density of the policy shifter θ_t depends only on the publicly observable variable x_t . This doesn't mean that θ lacks temporal dependence. But it does mean that any temporal dependence in θ_t is encoded in the publicly observable x_t . As a result, at the beginning of each period, all agents - public and private - agree on the likelihoods of different realizations of θ .

Suppose alternatively that, conditional on x , the taste shifter θ has persistence in the form of a Markov structure. The government's knowledge of θ_t then provides it with superior

knowledge about the future evolution of the policy taster shifter. The private sector’s choice of its action a_t would depend on its possibly imperfect beliefs about θ_t (updated after it observes the government’s choice of g_t). Thus, the translation of any government policy choice into the density of y now depends on the private sector’s beliefs about θ_t .

We can only treat this case if the government, at the time of making its choice, knows the private sector’s beliefs about θ_t . Without this knowledge, the government’s strategy would depend on the full history of past private sector actions, because that full history is informative about the private sector’s beliefs. We could no longer view the government’s strategy as being Markov in any simply described state.

So, to extend the results to allow for hidden persistence, we need to augment the public state variable x_t with a verifiable measure of the private sector’s beliefs about θ_t . Once we do so, the main result (that the government can identify its optimal choice using a simple regression of its payoff on g , conditional on x) can be extended to this case. In the monetary policy context, I view asset prices and/or surveys about future monetary policy as containing the relevant kind of information.¹⁷

6.3 Private Sector Private Information

In the above discussion. the private sector’s equilibrium strategy α^* is a deterministic function of its information x_t and g_t . It is easy to extend the baseline model to relax this (possibly overly strong) implication. Suppose that, at the same time as the government learns its payoff type, the private sector privately observes a factor φ drawn from a conditional density $p_\varphi(\cdot|x)$. Suppose too that this factor influences the private sector’s within-period payoff function and/or the conditional density of the economic outcome y . Then, the private sector’s equilibrium strategy α^* is a function of (x, g, φ) .

The government can still solve its policy problem using the same regression-based ap-

¹⁷The discussion in this subsection likely seems somewhat arcane. But it is highly pertinent to actual policymaking. The whole issue of whether “expectations are anchored” is tightly connected to what agents believe θ_t is.

proach as long as this additional factor φ satisfies two restrictions. First, it is important that, conditional on the government's information (x, ξ) , θ (the government's taste shifter) and φ are independent. Otherwise, the realization of θ affects the density of the economic outcome y through its interaction with ϕ , and the past realizations of g can no longer be viewed as a conditional experiment. Second, it is important that φ only depends on the past through x . Otherwise, the private sector and the government will enter the period with different beliefs about the evolution of the economy.

6.4 Non-Myopic Government

Throughout, I have imposed the restriction that the government only cares about its within-period payoff. Suppose that its objective in period t is instead:

$$w_t + \delta_g E_t w_{t+1} + \delta_g^2 E_t w_{t+2} + \dots + \delta_g^T E_t w_{t+T}$$

where E_t represents the government's conditional expectation¹⁸ as of date t . How would the above analysis change?

The answer is: not much. The current period t government can't control what the future governments will do. Rather, it makes its choices taking those future governments' Markov strategy γ^* as given. From the point of view of the current government, the future governments are really no different from the private sector.¹⁹

All that changes is the government's measure of its payoff. Suppose as before the government has a data set:

$$\{y_{t-n}, g_{t-n}, x_{t-n}\}_{n=1}^{\infty}$$

where (X, G) are compact intervals in Euclidean spaces. Suppose in addition that the government has data on both current and past realizations of its taste shifter $\{\theta_{t-n}\}_{n=0}^{\infty}$. In light

¹⁸In contrast, Feldman, et. al. (2016) argue that a non-myopic government should use market-based expectations rather than statistical expectations.

¹⁹Note that this analysis of a non-myopic government presumes that, conditional on x_t , θ_t is independent of past θ 's.

of its multi-period payoff, the government now defines:

$$W'_n = w(\theta_t, g_{t-n}, y_{t-n}) + \sum_{s=1}^T \delta_g^s w(\theta_{t-n+s}, g_{t-n+s}, y_{t-n+s})$$

to be its (realized) payoff in period $(t - n)$, where $n > T$. Without loss of generality, assume that the observations of $\{x_{t-n(T+1)}\}_{n=1}^\infty$ are all sufficiently close to x_t that we can treat:

$$(W'_{n(T+1)}, g_{t-n(T+1)})_{n=1}^\infty$$

as a sequence of i.i.d. random vectors. (We can always discard data so that this is true.) The government can now run a nonlinear regression of W' on g using this data set. Under the same conditional independence restriction on θ_t that we made earlier (it doesn't affect the economic outcome y_t or future state x_{t+1} except through g), this regression function is again equivalent to the government's current projection of its expected (multi-period) payoff, conditional on having chosen g .

7 Finite Samples

In this section, I return to the baseline setup described in Sections 3 and 4. There, I assumed that the government has an infinitely long data set of past observations on its choices and the outcomes of those choices. I now turn to the more realistic situation in which the government in a given period t only has a finite data set. I sketch two possible approaches to dealing with finite samples: Bayesian and classical. As will become clear, my discussion is mainly intended to be suggestive of possible routes for further research, as opposed to being definitive.²⁰

One unresolved issue is that there is a basic internal inconsistency in my treatment of the finite sample case. I consider the problem of a government at date t that only has a

²⁰I use a traditional probabilistic approach to thinking about the government's problem when it has a finite sample. However, it could be interesting to instead explore the pluses and minuses of a more algorithmic approach (that is, machine learning) as described in Breiman (2001) and as discussed in Mullainathan and Speiss (2017).

finite data set. As will become clear, that government's optimal choice necessarily depends on this entire sample and its strategy is no longer recursive in a state variable like x . But I nonetheless continue to assume that the government at date t treats the strategies used by the private sector and past governments as functions of the past only through x .

7.1 Bayesian

Suppose that, in period t , the government has a finite data set $\{x_{t-n}, g_{t-n}, y_{t-n}\}_{n=1}^N$ of past observations on the economic state, on its own choice, and on the economic outcome. It also has current data x_t on the economic state and on its taste shifter θ_t .

Define $W_n = w(\theta_t, g_{t-n}, y_{t-n})$ to be the government's payoff in period $(t - n)$ (evaluated using today's realized tastes). As noted above, suppose that the government's and the private sector's past strategies are time invariant functions of the current state. Then, conditional on the sequence $\{x_{t-n}, g_{t-n}\}_{n=1}^N$, $\{W_n\}_{n=1}^N$ is a realization of N independent random variables, where the density of W_n is determined by the pair (x_{t-n}, g_{t-n}) . (This means that the density of W_n and W_m is necessarily the same in any two dates in which $(x_{t-n}, g_{t-n}) = (x_{t-m}, g_{t-m})$.)

At this stage, I assume that using economic theory or other auxiliary sources of information²¹, the government knows that the density of W , conditional on (x, g) , can be written:

$$h(W, g, x, \beta)$$

for some β in a finite dimensional parameter space B . The government's information also gives it a prior belief about β representable by a density p over B .

Given the government's information about β , it can write the (conditional) likelihood of the data string $W^N \equiv (W_n)_{n=1}^N$ as:

$$L(W^N; \beta) = \prod_{n=1}^N h(W_n, g_{t-n}, x_{t-n}, \beta)$$

²¹By using a conditional approach, I ignore any potential information about β contained in the observations about (g, x) .

It can then update the prior p using this conditional likelihood to obtain a posterior belief over B , given the data string W^N :

$$p'(\beta) \propto p(\beta) \prod_{n=1}^N h(W_n, g_{t-n}, x_{t-n}, \beta)$$

With this posterior in hand, and given its current observations of (θ_t, x_t) , the government can find its optimal choice of g by solving the problem:

$$\max_{g \in G} \int_B p'(\beta) \int_{\mathbb{R}} Wh(W, g, x_t, \beta) dW d\beta$$

Here, the government is evaluating each g by calculating its implied within-period expected payoff, for each possible specification of β , and then integrating with respect to its (data-based) posterior p' .

Earlier, we saw how, with an infinite dataset, the government could use its past experience as its sole source of information about the consequences of different policy choices. But with a finite sample, the government may have few or no past observations with potentially desirable combinations of (x, g) . The Bayesian approach provides an automatic way to substitute theory or any other sources of information embedded in the prior to compensate for these data limitations. Note that, with an infinite number of observations, the posterior density $p'(\beta)$ would be concentrated on a single value.

7.2 Classical

I begin with a classical approach that is standard in the sense that it is typically used in practical applications. I then describe an alternative approach that is grounded in statistical decision theory.

The standard approach is to plug in the maximum likelihood estimate of the unknown parameters. Define the conditional likelihood $L(W^N; \beta)$ as above. We can then define the

maximum likelihood estimate $\hat{\beta}(W^N)$ as that value of β that solves:

$$\max_{\beta \in B} L(W^N; \beta).$$

Given this estimate, the government can find its optimal policy by solving the problem:

$$\max_{g \in G} \int_{\mathbb{R}} Wh(W, x_t, g, \hat{\beta}(W^N)) dW$$

However, there is a problem with this (standard) approach. If we look at the government's maximand in the above objective, we can see that its maximand depends on the sample W^N via $\hat{\beta}$. From an ex-ante (pre-sample) perspective, this variability influences the government's payoff. This effect is ignored by the standard approach. (Of course, if $N = \infty$, this variability disappears: $\hat{\beta}(W^N)$ is the same for all samples.)

We can take this sample variability into account by using basic statistical decision theory.²² Consider any decision rule d that maps datasets W^N into the government's action set G . (Again, as above, I treat the past realizations of $(g_{t-n}, x_{t-n})_{n=1}^N$ as given.) Then, if the true parameter is β , the government gets an ex-ante (pre-sample) payoff of:

$$\Psi_G(\beta; d) = \int_{\mathbb{R}^N} \int_{\mathbb{R}} Wh(W, x_t, d(W^N), \beta) L(W^N; \beta) dW dW^N$$

The government can then evaluate the ex-ante performance of different decision rules by comparing their payoff functions Ψ_G . For example, suppose there is a decision rule d^* such that its associated payoff function is larger than any other decision rule for all β :

$$\Psi_G(\beta; d^*) \geq \Psi_G(\beta; d)$$

Then it would clearly be appropriate for the government to use that d^* . Unfortunately,

²²Ferguson (1967) is a classic reference.

it seems highly unlikely in most settings that the government could find such a d^* . More generally, the government would have to choose a decision rule by trading off performance at some (true) values of β against performance at other values of β . The Bayesian approach accomplishes that trade-off via the prior belief on β .

8 The Right Questions to Ask

There may be no more important sentence in the intellectual history of macroeconomics than this one penned by Robert Lucas (1980a): “Our task ... is to write a FORTRAN program that will accept specific economic policy rules as ‘input’ and will generate as ‘output’ statistics describing the operating characteristics of time series we care about, which are predicted to result from these policies.” This research vision - radical at the time Lucas wrote - became foundational for the discipline in the 1980s and has remained so ever since. There is a sense that it has led to much progress. Thanks to advances in both technique and technology, the computer programs that Lucas foresaw are based on increasingly complex economic models of household and firm interactions. Thanks to the same advances in technique and technology, these complex models can now be parameterized with reference to amazingly rich micro-data.

The problem is that this research agenda is *by construction* largely irrelevant to the formation of most - not all, but most - actual economic policy. The agenda is based on the idea that policy evaluation is about the assessment of *rules* that specify how policy is to be formulated in all dates and states. The consideration of such sweeping counterfactuals is indeed a highly exciting intellectual activity. It should not be surprising that it has generated an enormous body of fascinating papers. But this literature about policy regimes is not of much *practical* interest: policymakers rarely have the power to engage in this kind of quasi-constitutional design. The typical policymaker is instead charged at a single point in time with making but one episodic decision in a long sequence of such decisions.

How should such a policymaker use the available data to make the best possible deci-

sion? In this paper, I analyze this problem using the tools of dynamic game theory. I show that the policymaker's problem can be addressed by first regressing policymaker payoffs on policymaker choices (and policymaker information), and then making the policy choice that maximizes that (possibly nonlinear) regression function.

Of course, I don't arrive at this conclusion without making assumptions. The most important ones are that:

- (Conditional Independence Restriction) Conditional on the policymaker's information about the economy, policy decisions vary because of some factor that is privately known to the policymaker and has no impact on the economy other than through the policy decision itself.
- (Full Support Restriction) The variation in policy decisions induced by this auxiliary factor is sufficiently large that no current policy decision will alter the private sector's beliefs about the strategy of future policymakers.

The full support restriction is testable (with a sufficiently large data set). As is often (typically) true in purely observational data, the conditional independence restriction is not.²³

But it would be a logical and substantive mistake to view the validity of these two sufficient restrictions as strictly necessary for the purely statistical approach recommended in this paper to be of use. Rather, I view the theory in this paper as pointing to the kinds of questions that policymakers and their advisors should be asking:

- Why have policy decisions varied in the past, given the available information about the economy?
- Conditional on that available information, can this source of variation be plausibly viewed as having had little impact on the economy except through the policy choice itself?

²³Cooley and Leroy (1984) criticize atheoretical macroeconometrics because of these nontestable restrictions. Despite their (reasonable) concerns, in the interim, these kinds of restrictions have become critical in much of applied microeconomics. Also, and as I noted earlier, much of theoretical macroeconomics employs even stronger independence restrictions.

- How will the policy choices under consideration affect the private sector's beliefs about the strategies of future policymakers?
- What kinds of variables should be seen as part of the economic state (x) that is pre-determined at the time policy is made?
- What kinds of plausible *a priori* parametric restrictions can be imposed on the form of the regression function of interest?

These kinds of questions lie at the heart of much modern applied microeconomics. They should become more central to macroeconomic policy evaluation that is intended to be of practical value.

References

- [1] Barro, R., and Gordon, D., “A Positive Theory of Monetary Policy in a Natural Rate Model,” *Journal of Political Economy* 91, 1983, 589-610.
- [2] Breiman, L., “Statistical Modeling: The Two Cultures,” *Statistical Science* 16, 2001, 199-231.
- [3] Chari, V. V., and Kehoe, P., “Sustainable Plans,” *Journal of Political Economy* 98, 1990, 783-802.
- [4] Cooley, T., and LeRoy, S., “Atheoretical Macroeconometrics: A Critique,” *Journal of Monetary Economics* 16, 1985, 283-308.
- [5] Ferguson, T., *Mathematical Statistics: A Decision-Theoretic Approach*, New York: Academic Press, 1967.
- [6] Feldman, R., Heinecke, K., Kocherlakota, N., Schulhofer-Wohl, S., and Tallarini, T., “Market-Based Expectations as a Tool for Policymakers,” University of Rochester working paper, 2016.
- [7] Imbens, G., and Wooldridge, J., “Instrumental Variables with Treatment Effect Heterogeneity: Local Average Treatment Effects,” NBER Lecture 5, Summer 2007.
- [8] Kocherlakota, N., “Rules versus Discretion: A Reconsideration,” *Brookings Papers on Economic Activity* 2016(2), 1-48.
- [9] Leeper, E., and Zha, T., “Modest Policy Interventions,” *Journal of Monetary Economics* 50, 1973, 1673-1700.
- [10] Ljungqvist, L. and Sargent, T., *Recursive Macroeconomic Theory, Third Edition*, Cambridge, MA: MIT Press, 2018.

- [11] Lucas, R. E., Jr., “Econometric Policy Evaluation: A Critique,” *Carnegie-Rochester Conference Series on Public Policy* 1, 1976, 19-46.
- [12] Lucas, R. E., Jr., “Methods and Problems in Business Cycle Theory,” *Journal of Money, Credit, and Banking* 12, 1980a, 696-715.
- [13] Lucas, R. E., Jr., “Rules, Discretion, and the Role of the Economic Advisor,” in *Rational Expectations and Economic Policy*, ed. S. Fischer, Chicago: The University of Chicago Press, 1980, 199-210.
- [14] Mullainathan, S., and Spiess, J., “Machine Learning: An Applied Econometric Approach,” *Journal of Economic Perspectives* 31, 2017, 87-106.
- [15] Plosser, C., “Forward Guidance,” speech given at Stanford Institute of Economic Policy Research, Stanford University, 2013.
- [16] Plosser, C., “Monetary Policy Rules: Theory and Practice,” speech given at “Frameworks for Central Banking in the Next Century Policy Conference,” Hoover Institution, Stanford University, 2014.
- [17] Romer, C., and Romer, D., “Does Monetary Policy Matter? A New Test in the Spirit of Friedman and Schwartz,” *NBER Macroeconomics Annual* 4, 1989, 121-170.
- [18] Sims, C., “Macroeconomics and Reality,” 1980, *Econometrica* 48, 1980, 1-48.
- [19] Sims, C., “Policy Analysis with Econometric Models,” *Brookings Papers on Economic Activity* 1982(1), 107-151.
- [20] Stock, J., and Watson, M., “Dynamic Factor Models, Factor-Augmented Vector Autoregressions, and Structural Vector Autoregressions in Macroeconomics,” *Handbook of Macroeconomics, vol 2.*, ed. J. Taylor and H. Uhlig, New York: Elsevier, 2016, 415-525.
- [21] Taylor, J., “Discretion Versus Policy Rules in Practice,” *Carnegie-Rochester Series on Public Policy* 39, 1993, 195–214.