

Government Policy and Probabilistic Equilibrium Selection^α

Huberto Ennis^γ Todd Keister^z

January 27, 2000

Abstract

We study an economy where search frictions create a coordination problem among agents and thereby generate multiple equilibria. Our interest is in how likely it is that the economy will find its way to each of these equilibria when agents learn as Bayesians. We show how using learning as an equilibrium selection device generates a probability distribution over the set of equilibria, and we study the effect of different government policies on this distribution. We show how in our model a tradeoff arises – policies that increase the value of being in a particular equilibrium tend to reduce the probability of reaching that equilibrium.

^αVERY PRELIMINARY. Please do not cite without permission. We thank Marcelo Veracierto for sharing his expertise in Fortran with us.

^γDepartment of Economics, Cornell University, 402 Uris Hall, Ithaca, NY 14850, U.S.A. E-mail: hme1@cornell.edu.

^zCentro de Investigación Económica, ITAM, Av. Camino Santa Teresa 930, México, D.F. 10700, Mexico. E-mail: keister@master.ster.itam.mx

1. Introduction

Models with multiple equilibria are often criticized as not being useful for policy analysis because they fail to assign a unique prediction to each possible policy choice. Rather, each choice leads to a set of possible outcomes, and this set may be quite large. In some cases it is possible to make meaningful comparisons of different sets of equilibria and thereby determine an “optimal” policy.¹ However, such cases are clearly the exception and not the rule; set-wise comparisons do not generally yield unambiguous results. Various criteria have been proposed for selecting one (or a small set) of the equilibria as “relevant” for the particular issue at hand so that the optimal policy can be determined relative to that equilibrium.² Unfortunately, none of these criteria have proven completely satisfactory. One such criterion is stability with respect to Bayesian learning about some aspect(s) of the economic environment.³ In this approach, agents have some initial beliefs about the data-generating processes in the economy and update these beliefs based on what they observe according to Bayes’ rule. As they act on these updated beliefs, more data is generated and they update their beliefs again. Under some conditions, the economy will converge to one of the rational expectations equilibria of the model. One perceived problem with this approach is that it often does not select a unique equilibrium; the economy may asymptotically approach several equilibria with positive probability. Our goal in this paper is to demonstrate that this is *not* a shortcoming of the learning approach. We show how Bayesian learning generates a probability distribution over the set of equilibria in a simple example, and we focus on the properties of this distribution. This distribution provides a different type of answer to the question of what will happen when a policy changes. Models with a unique equilibrium give a point estimate – if the policy parameter changes to x ; the new equilibrium will be y : Our model will make a probabilistic statement: if the policy parameter changes to x ; the possible outcomes are $y_1; \dots; y_n$; and equilibrium y_i will occur with probability $\frac{1}{2}$: We argue that this is a precise and valuable prediction. We show that the optimal policy choice resulting from such a view of the model is, in some cases, different from that which would be derived using any deterministic selection criterion.

Bayesian learning is one element of a broad class of dynamic processes that have been proposed

¹ See, for example, Grandmont [5], Woodford [13], Smith [12], and Keister [8], among others.

² See Guesnerie and Woodford [6], section 7, for an introduction to this topic and an extensive list of references.

³ See Blume and Easley [2] for an excellent survey of the literature on Bayesian learning.

as equilibrium selection mechanisms. Our analysis focuses entirely on this type of learning because it allows us to present our results in a concise and relatively transparent way, but the basic ideas we present are in no way tied to this choice. Similarly, the model we employ is highly stylized. Our goal is to demonstrate the value of probabilistic equilibrium selection in this particular setting, and in future work to show how the ideas presented here extend to other environments.

In our model, individual agents must decide whether or not to engage in production. Because of an output-market externality, the value of producing depends on how many other agents produce. The model is similar in many respects to that of Howitt and McAfee [7] , but our focus is on search and matching in the output market rather than in the labor market. Our externality generates a coordination problem in that an individual agent wants to produce if and only if enough other agents are producing. This, in turn, generates a pair of Pareto-ranked equilibria, one where everyone produces (the “good” equilibrium) and the other where no one produces (the “bad” equilibrium). We use learning to determine the likelihood that the economy will find its way to each of these equilibria. To address the issue of optimal policy determination, we introduce a government which is composed of a fraction of the agents in the economy. By determining the production and search decisions of these agents, the government can affect the value to a private agent of producing. This, in turn, affects not only the existence of each of the two equilibria, but also the likelihood of the economy converging to each of them. We show how the government may face a tradeoff: policies that increase the value of the good equilibrium to private agents tend to make that equilibrium less likely. We compute this tradeoff for a numerical example.

The next section contains a detailed description of the model and the policy tools available to the government. Section 3 presents some properties of equilibrium both with and without active government policy. It also provides a characterization of optimal allocations in this setting (which typically differ from equilibrium allocations because of the output-market externality). Section 4 then demonstrates how learning generates a probability distribution over the equilibrium set and presents some results characterizing this distribution, including the tradeoff that may arise. Finally, section 5 contains some concluding remarks.

2. The Model

Time is discrete and the horizon is infinite. There is a $[0; 1]$ continuum of identical agents and a single commodity in each period which is produced using only labor. Agents can consume their own output, but they prefer the output of others.⁴ Finding a match in the output market requires costly effort. Agents who have produced choose a level of search effort, with higher effort leading to a higher probability of finding a buyer (match). If a buyer is found, the two agents trade output and consume; if not, the agent consumes her own output. There is no credit and hence all exchanges are quid-pro-quo. These assumptions and the perishability of goods imply that the actions of an agent do not have intertemporal effects and therefore that the only dynamics in our model are with respect to expectations.

We allow for government policy in this model by assuming that the government is composed of fixed fraction \bar{A} of the agents in the economy. The actions taken by these agents need not be utility maximizing; they instead represent the policy of the government.⁵ By having all of these agents produce and then engage in output-market search, the government can increase demand in the output market, possibly making it more attractive for private agents to produce. We will consider two types of government policy, which we term *passive* and *active*. This is identical to a model with no government. In a passive policy, government agents simply act as private agents and maximize their own utility. In an active policy, government agents always produce, and then search with a common effort level. This effort level is the government's policy parameter.

2.1 Production

At the beginning of each period, each private agent must decide whether or not to produce. Production is a binary decision; it must be operated at a fixed scale or not operated at all. The utility value of the output when consumed by another agent is given by F : The utility value of consuming one's own output is $\frac{3}{4}F$; where $\frac{3}{4}$ is strictly less than unity. Producing requires labor input that gives disutility c : The cost c is an i.i.d. random variable each period with support $c^L; c^H$; where $c^L < c^H$. The probability that $c = c^L$ is given by β ; while the expected value of c is given by \bar{c} :

⁴ We think of this as a form of specialization in production, as in Diamond [4].

⁵ Aiyagari and Wallace [1], also within a search theoretic environment, model the government participation in the economy following the same principles used here (however, they are mainly concerned with monetary issues).

The realization of c is not yet known when the production decision is made. If an agent chooses not to produce, she does not consume and has zero utility that period (a normalization). We use r to denote the fraction of all agents that choose to produce.

Before we can study the problem of an agent deciding whether or not to produce, we need to calculate the benefit of producing. For this we need to look at the workings of the output market.

2.2 Output Market

As mentioned above, agents prefer goods produced by others to those that they produce themselves. Agents who have produced have an opportunity to search for a buyer, and they must choose a search intensity $\theta \in [0, 1]$: The probability of each agent finding a match, and hence the fraction of agents finding a match in equilibrium, is given by an aggregate matching function. We first describe the properties of this function and then analyze the problem faced by an agent in choosing her search intensity.

The Matching Function: The fraction of all agents in the economy that find a match is given by an aggregate matching function m .⁶ Let $\bar{\theta}$ denote the average level of search intensity in the economy, so that we have

$$\bar{\theta} = \int_0^1 \theta_i di$$

(we will obviously have $\theta_i = 0$ for agents that have not produced). Then the value of m depends on $\bar{\theta}$ and on the number of agents that are searching r : We assume that m takes the form

$$m(\bar{\theta}; r) = \psi(\bar{\theta})r$$

The function ψ then gives the number of matches per searching agent as a function of the total level of search intensity in the market. We assume that $\psi : \mathbb{R}_+ \rightarrow [0, 1]$ is continuous on \mathbb{R}_+ ; is C^3 on \mathbb{R}_{++} ; and satisfies the following conditions

⁶ Assuming a matching function simplifies matters considerably, but it clearly does not come without cost. In particular, policy changes may affect the nature of the matching process, and this effect will be absent in our analysis. We leave the study of this effect for future work. See Lagos [9] for an interesting analysis in this direction.

$$(A1) \quad \phi(0) = 0$$

$$(A2) \quad \phi'(x) > 0 \text{ for all } x > 0$$

$$(A3) \quad \lim_{x \rightarrow 1} \phi'(x) = 1$$

$$(A4) \quad \lim_{x \rightarrow 0} \phi''(x) = 0$$

$$(A5) \quad \phi'(x) < x \text{ for all } x > 0$$

$$(A6) \quad \text{There exists an } \bar{x} > 0 \text{ such that}$$

$$\phi''(x) > 0 \text{ for } 0 < x < \bar{x}$$

$$\phi''(x) < 0 \text{ for } \bar{x} < x < 1$$

$$(A7) \quad 1 < \lim_{x \rightarrow 0} \phi''(x) < \frac{4a}{(1-\frac{3}{4})F}$$

$$(A8) \quad \phi'''(x) < 0 \text{ for } 0 < x < \bar{x}:$$

The first assumption is simply that when no agents search, no one gets matched, while the second is that search effort is productive as long as the average intensity in the economy is positive. The third assumption says that for high enough levels of effort, nearly all searching agents will be matched.

The reasons for the fourth and fifth assumptions follow from the interpretation of ϕ' as a probability facing individual agents. Note that ϕ' gives the fraction of *searching* agents that find a match. If an individual agent chooses search intensity θ_i , taking the search intensity of other agents (θ^*) as given, then the resulting probability of getting matched with a buyer is

$$\frac{\phi'(\theta^*) \theta_i}{\theta^*}$$

Define

$$\frac{1}{2}(\theta^*; r) = \frac{\phi'(\theta^*) \theta_i}{\theta^*};$$

so that $\frac{1}{2}$ is the probability of finding a match per unit of search intensity. The fourth assumption guarantees that we have

$$\lim_{\theta^* \rightarrow 0} \frac{1}{2}(\theta^*; r) = 0;$$

so that if no other agents are searching, there is no reason to search for a buyer. The fifth assumption guarantees that the probability $\frac{1}{2}$ is less than one. The final three assumptions are regularity

conditions that generate a unique interior perfect-foresight equilibrium in the output market.⁷ The previous assumptions imply that ϕ is initially convex and eventually concave. Condition (A6) is simply that there is an unique point where the second derivative is zero. Condition (A7) states that the degree of convexity is initially large enough, which guarantees that the average product of search effort initially increases quickly enough to make positive search effort worthwhile. The final assumption states that the degree of convexity of ϕ is strictly decreasing up to the changeover point of the second derivative, or that the marginal product of search effort is initially concave. The importance of this condition for generating uniqueness of output-market equilibrium is discussed in the results below.

Throughout the paper we will present examples and computations that use the function

$$\phi(x) = 1 - \frac{1}{1+x^2} \quad (1)$$




Figure 1: A Matching Function

This function is depicted in figure 1. It satisfies all of our assumptions as long as we have

$$\frac{2a}{(1 - \frac{3}{4})F} < 1;$$

since $\phi''(0)$ for this function is equal to 2:

The Choice of Search Intensity: Agent i chooses a nonnegative search intensity σ_i : The cost of

⁷ There will still often be multiple equilibria in the economy as a whole because of the coordination problem in deciding whether or not to recruit.

this intensity is convex and given by a_i^2 : If an agent meets a buyer, they exchange output. If an agent does not meet a buyer, she can consume her own output, but the value to her is only $\frac{3}{4}F$. The search cost is measured in units of utility. The agent's output-market problem is to choose θ_i to maximize expected utility, that is, to solve

$$\begin{aligned} \max_{\theta_i} & \frac{1}{2}\theta_i F - a_i^2 \theta_i^2 + (1 - \frac{1}{2}\theta_i) \frac{3}{4}F - a_i^2 \theta_i^2 \\ \text{subject to} & \quad \frac{1}{2}\theta_i \leq 1; \end{aligned} \tag{2}$$

where the agent takes the value of $\frac{1}{2}$ as given (in equilibrium this will be $\frac{1}{2}(\frac{c}{r}; r)$). The constraint says that θ_i must be chosen so that the implied probability of finding a match is no greater than one. Equivalently, we could change the objective function so that if θ_i is chosen to make $\frac{1}{2}\theta_i$ greater than one, then the probability of a match would still be one. Since this would entail a higher search cost and no benefit, no agent would ever pick $\frac{1}{2}\theta_i > 1$; and the solution to that problem is the same as the solution to the constrained problem we have written here. The objective function can be reduced to

$$\frac{3}{4}F + (1 - \frac{3}{4}) \frac{1}{2}F \theta_i - a_i^2 \theta_i^2$$

The first-order condition for an interior solution is

$$(1 - \frac{3}{4}) \frac{1}{2}F = 2a_i^2 \theta_i$$

which generates the general solution

$$\theta_i^* = \min \left\{ \frac{(1 - \frac{3}{4}) \frac{1}{2}F}{2a_i^2}; \frac{1}{\frac{1}{2}} \right\} \tag{3}$$

Note that we have $\theta_i^* \geq 0$; with $\theta_i^* > 0$ as long as $\frac{1}{2}$ is greater than zero. Since the expected benefit of search effort is linear in θ and the marginal cost of effort is zero when θ is zero, it is optimal to engage in a positive level of search whenever $\frac{1}{2}$ is positive.

2.3 The Production Decision

With this information about the value of producing, we are now ready to examine the agent's production decision. The value of producing clearly depends on how many other agents will be par-

icipating in the output market. Let v_i denote the (utility) value to an agent of producing. The agent faces a binary decision problem, and will produce if and only if this value is greater than the expected cost, that is, if and only if

$$v_i > \tau$$

holds. The value v_i is given by

$$v_i = E \max_i \left\{ \frac{1}{2} \left(\frac{1}{i} F - a_i^{\circ} \right)^2 \mathbb{1} + \left(1 - \frac{1}{2} \right) \left(\frac{1}{i} F - a_i^{\circ} \right)^2 \mathbb{1}^{\circ} \right\};$$

where E denotes the expectation with respect to the value of $\frac{1}{2}$; which is determined by the decisions of the other agents. We assume that an individual agent believes that all other private agents will act identically (that is, we rule out asymmetric and mixed-strategy equilibria). Hence, in the absence of an active government policy, aggregate employment will be either zero or unity. When r is expected to be unity, the value of producing can be written as

$$v_i^H = \frac{1}{2} \left(\frac{1}{i} F - a_i^{\circ} \right)^2 \mathbb{1} + \left(1 - \frac{1}{2} \right) \left(\frac{1}{i} F - a_i^{\circ} \right)^2 \mathbb{1}^{\circ}$$

or

$$v_i^H = \frac{1}{3} F + \left(1 - \frac{1}{3} \right) \frac{1}{2} \left(\frac{1}{i} F - a_i^{\circ} \right)^2 \mathbb{1}^{\circ};$$

When r is expected to be zero, the value is simply given by

$$v_i^L = \frac{1}{3} F;$$

2.4 Government Policy

Specifying a government policy in this environment amounts to specifying a decision rule for public agents. A passive policy is the same as a model with no government: all agents maximize their own utility. We also consider an active policy where all public agents produce and then search with a common intensity level θ_G : In choosing the search intensity of public agents, the government is affecting the probability that searching private agents will find a buyer. We therefore view this policy as a form of aggregate demand management. Our interest is in (i) comparing the set of equilibria under the two policies and (ii) analyzing the optimal-policy question of how the government

should set ϕ_G :

3. Equilibrium

The equilibrium conditions for this economy are simply that the decisions made by agents generate the market conditions that each private agent takes as given. Formally, we have the following definition.

Definition: An *equilibrium with a passive government* is a function $\phi^p : [0; 1] \rightarrow \mathbb{R}$ giving the search intensity of every agent and a scalar $r \in [0; 1]$ such that when $\phi^p = \int_0^1 \phi_i^p di$; we have

- (i) each agent chooses optimally whether or not to produce and sets ϕ_i^p by (3),
- (ii) r is equal to the fraction of all agents that choose to produce, and
- (iii) $\frac{1}{2} = \frac{1}{2}(\phi^p; r) = \int_0^1 \phi_i^p r = \phi^p$:

An *equilibrium with an active government* is $(\phi^p; r)$ such that we have

- (i) $\phi_i^p = \phi_G$ for each public agent,
- (ii) each private agent chooses optimally whether or not to produce and sets ϕ_i^p by (3),
- (iii) $r \in \bar{A}$ is equal to the fraction of all agents that choose to produce, and
- (iv) $\frac{1}{2} = \frac{1}{2}(\phi^p; r) = \int_0^1 \phi_i^p r = \phi^p$:

We first examine the passive regime, and then turn to the analysis of equilibrium with active government policy.

3.1 Equilibrium with a Passive Government

We begin by showing that the optimal choice of ϕ_i is always interior, that is, that the constraint to problem (2) is never binding in equilibrium.

Lemma 1 The solution to problem (2) at the equilibrium value of $\frac{1}{2}$ is given by

$$\phi_i^p = \frac{(1 - \frac{3}{4}) \frac{1}{2} F}{2a} \tag{4}$$

Proof: See appendix. ¥

Since the agent's problem has a unique solution and all agents are identical, all agents in the

output market choose the same level of θ : This implies that we have

$$\theta = \int_0^1 \theta_i di = r^\theta:$$

Substituting this information and the definition of θ into equation (4), we have that the equilibrium intensity value θ^* is implicitly defined by

$$\theta^* = \frac{(1 - \frac{3}{4}) \frac{1}{\theta^*} F}{2a}: \quad (5)$$

We now show that our assumptions on θ guarantee that there exists a unique nonzero solution to this equation.

Lemma 2 There exists a unique interior solution for θ^* to equation (5).

Proof: See appendix . ¥

The uniqueness property allows us to fully characterize the comparative-static properties of the equilibrium search intensity θ^* :

Lemma 3 Let θ^* be the solution to (5). Then we have

$$\frac{\partial \theta^*}{\partial F} > 0; \frac{\partial \theta^*}{\partial a} < 0; \text{ and } \frac{\partial \theta^*}{\partial \frac{3}{4}} < 0:$$

Proof: Define $f(\theta; F; a; \frac{3}{4})$ to be the right-hand side of (5). Then the uniqueness of the solution implies that f crosses the forty-five degree line from above, or that

$$\frac{\partial f}{\partial \theta} < 1$$

must hold. This implies that the sign of the derivative of f with respect to each of the remaining variables directly determines the sign of the derivative of θ^* with respect to that variable. ¥

Example: Equation (5) can be illustrated using the matching function (1) and a “baseline” set of parameter values

$$a = \frac{1}{2}; F = 5; \frac{3}{4} = \frac{1}{2}; \tau = 2:8:$$

Plotting the left-hand and right-hand sides of (5) separately yields figure 2. The two curves cross at $\theta = 0$ and $\theta = \frac{p_2}{2}$: ¤




Figure 2: The equilibrium level of ϕ

With a passive government policy, there exist at most two (pure-strategy) equilibria under perfect foresight. We follow Howitt and McAfee [7] in calling these equilibria “optimistic” and “pessimistic.” First suppose that each agent believes that all other agents will produce. Given this optimistic belief, if $\phi^H > \bar{c}$ holds, then it is indeed optimal for every agent to produce, and hence there is an equilibrium with r equal to unity. Note that we have $\phi_i = \phi^H$ for all agents i that engage in production, and hence $\phi = \phi^H$. The equilibrium value of ϕ^H can therefore be written as

$$\phi^H = \frac{1}{2} (\phi^H)^2 F + (1 - \frac{1}{2} (\phi^H)^2) \frac{1}{3} F + a (\phi^H)^2 \bar{c}.$$

From equation (5) we have that

$$a (\phi^H)^2 = \frac{1}{2} (1 - \frac{1}{3}) (\phi^H)^2 F;$$

holds, so that the expression for ϕ^H simplifies to

$$\phi^H = \frac{1}{3} F + \frac{1}{2} (1 - \frac{1}{3}) (\phi^H)^2 F$$

or

$$\phi^H = \frac{1}{3} F + a (\phi^H)^2 \bar{c}.$$

The optimistic equilibrium exists whenever this expression is greater than the expected cost of producing. Now suppose that each agent believes that no other agent will produce. Given this pessimistic belief, not producing is optimal if $\phi^L < \bar{c}$ holds, in which case there is an equilibrium

with r equal to zero. The equilibrium value of ρ^L does not depend on ρ^H (since if no other agent produces there is no reason to search for a buyer), and is given by $\frac{1}{3}F$: Therefore the pessimistic equilibrium exists if $\frac{1}{3}F < \tau$ holds.

Example: For our example (with the parameter values given above), we have

$$\rho^H = 3:25; \tau = 2:8; \rho^L = 2:5:$$

Clearly both equilibria exist. □

3.2 Equilibrium with an Active Government

If the government decides to have all public agents produce, then total government employment will be \tilde{A} and aggregate employment can never fall below this. A private agent who decides to produce still faces the maximization problem given in (2), and hence will choose ρ_i according to (4). The government policy affects this choice through its effect on $\frac{1}{2}$: We assume that the government policy is fixed and known, but as before an agent's beliefs about the actions of other private agents are critical. Again we will consider only symmetric outcomes, so that either no private agents produce or all private agents produce. We begin with the case of pessimistic beliefs.

Pessimistic Beliefs: If no private agents produce, total employment is \tilde{A} and total search intensity is $\rho = \tilde{A} \rho^G$: In deciding whether to produce or not, a private agent looks at the probability of getting matched with a public agent. The probability per unit of search effort is given by

$$\frac{1}{2}(\rho; r) = \frac{\tilde{A} \rho^G}{\rho^G}:$$

The next lemma shows that as long as the public sector is not too big, the optimal search effort of a private agent (if she produces) is again given by the interior solution.⁸

Lemma 4 If

$$\tilde{A} \cdot \frac{s}{(1 - \frac{1}{3})F} < \frac{2a}{2a}$$

⁸ Note that the public-employment policy guarantees that $\frac{1}{2}$ will be positive, even if no private agents produce. Then, since there is no fixed cost of searching, an agent who produces will necessarily choose a positive level of ρ :

holds, then the optimal level of search intensity for a private agent is given by

$$i_i^{\circ} = \frac{(1 - \frac{3}{4}) \frac{1}{\alpha^G} F}{2a}; \quad (6)$$

Proof: See appendix. ¥

The value of producing in this case is a function of the policy parameter α^G and given by

$$L_i^{\circ} = \frac{1}{2} i_i^{\circ} F + (1 - \frac{1}{2} i_i^{\circ}) \frac{3}{4} F - a (i_i^{\circ})^2;$$

where the values of $\frac{1}{2}$ and i_i° are given above. Substituting in the optimal value i_i° , this expression reduces to

$$L_i^{\circ} = \frac{3}{4} F + \frac{(1 - \frac{3}{4})^2 \frac{1}{2} F^2}{4a};$$

which shows that the value of producing is strictly increasing in $\frac{1}{2}$: If the government chooses α^G so that

$$L_i^{\circ} > \bar{c}$$

holds, then the policy has eliminated the bad equilibrium and all agents will engage in production. Note, however, that the government can only increase $\frac{1}{2}$ up to a point. Once α^G is large enough that

$$i_{\bar{A}}^{\circ} < \frac{i_{\bar{A}}^{\circ}}{\bar{A}^{\circ}}$$

holds, further increases in effort decrease $\frac{1}{2}$ (due to the crowding-out effect). Hence it may be the case the no level of government effort can eliminate the bad equilibrium. However, as we show in the next section, by changing L_i° the government may be able to decrease the likelihood of the bad equilibrium even if it cannot eliminate it.

Example: For our chosen parameter values, the constraint in lemma 4 reduces to $\bar{A} \cdot \frac{2}{5} \frac{1}{4}$
 0:6325: We choose $\bar{A} = \frac{1}{2}$ as a base. Figure 3 plots i_i° as a function of α^G . Initially i_i° is increasing in α^G : As α^G increases further, the economy moves into the negative externality region and i_i° starts to decrease. The value of producing under pessimistic beliefs L_i° is given by

$$L_i^{\circ} = \frac{3}{4} F + a \alpha^2;$$




Figure 3: \bar{i}^G under pessimistic beliefs

Figure 4 plots \bar{s}^L as θ^G varies. The value of \bar{s}^L is initially increasing in θ^G ; but becomes decreasing in the negative externality region. Notice that for this example, no value of θ^G raises \bar{s}^L above \bar{c} ; so that it is not possible for this policy to eliminate the bad equilibrium. If instead we increase F to 5:59; the picture becomes that in figure 5. Here the policy can eliminate the bad equilibrium if θ^G is chosen in the appropriate range. Notice, however, the small range of the horizontal axis. Only for carefully chosen parameter values will the policy be able to affect the existence of the bad equilibrium. □

Optimistic Beliefs: Even if the government knew that all private agents held optimistic beliefs and would produce, it might still want to use θ^G as a policy tool to offset the externalities involved in output market search. Private agents still take the value of $\frac{1}{2}$ as given and choose θ^G according to (4). In this case, it is not clear if the solution will necessarily be interior. A sufficient condition for interiority would be

$$s > \frac{2a}{(1 - \frac{3}{4})F} - 1;$$

but this may be very strong. The proof would proceed exactly as in the previous lemma, but with \tilde{A} set equal to one.

Define \mathbf{b} to be the average level of search intensity

$$\mathbf{b} = (1 - \tilde{A}) \theta^G + \tilde{A} \theta^G;$$




Figure 4: $\frac{1}{2} \circ^G$

In this case the probability of finding a match per unit of search effort is given by

$$\frac{1}{2} \circ^G(r) = \frac{1(b)}{b}$$

The value of producing is then given by

$$\frac{1}{2} \circ^H = \frac{1}{2} \circ^G F + (1 - \frac{1}{2} \circ^G) \frac{3}{4} F - a (\circ^G)^2$$

If the solution to the private agents' problem is interior, then the equilibrium intensity is implicitly defined by

$$\circ^G = \frac{(1 - \frac{3}{4}) \frac{1(b)}{b} F}{2a}$$

and the payoff to producing simplifies, as before, to

$$\frac{1}{2} \circ^H = \frac{1}{3} F + a (\circ^G)^2 \tag{7}$$

This expression tells us that public employment only changes $\frac{1}{2} \circ^H$ through its effect on the equilibrium value of \circ^G . Since the government has some control over b , it can use \circ^G to influence $\frac{1}{2}$ and hence the value to private agents of producing.

Example: Plotting \circ^G as a function of \circ^G under optimistic beliefs yields figure 6. Note the scale of the horizontal axis. High values of \circ^G discourage search effort under optimistic beliefs much sooner than under pessimistic beliefs. This is because when private agents are producing and searching, the




Figure 5: ψ^L vs θ^G for $F = 5:59$

economy enters the negative externality region of the matching function more quickly. Hence the value of θ^G that maximizes ψ^H will be much lower than the value that maximizes ψ^L : This can be seen in figure 7, which graphs ψ^H as a function of θ^G : The value of ψ^H is initially increasing in θ^G ; an indication that when public agents are not searching the economy is in the positive externality region of 1: The value of θ^G that maximizes ψ^H is approximately $0.75 < \mathbf{b}^G$: Note that this is below the equilibrium search level with a passive government; this indicates that the laissez-faire economy is in the negative externality region. As θ^G increases past \mathbf{b}^G , the economy enters the negative externality region and ψ^H starts to fall. If the government agents search enough, the optimistic equilibrium can be eliminated. □

In summary, this type of demand management policy changes the value of producing under both pessimistic and optimistic beliefs. It may be possible for the government to use this policy to eliminate the bad equilibrium. For our chosen parameter values, however, it is not. What value of θ^G should the government set? The model does not give a clear answer, but there is one compelling candidate. Since θ^G does not affect the value of the bad equilibrium (which is always zero), it might seem reasonable to set $\theta^G = \mathbf{b}^G$ and maximize the value to private agents of being in the good equilibrium. The problem with this reasoning is that, as we shall see below, θ^G affects the likelihood that the economy will reach the good equilibrium through learning and that from this point of view \mathbf{b}^G is a bad choice. Before we present our analysis of learning, however, we briefly discuss optimal allocations in this model.




Figure 6: α_i under optimistic beliefs

3.3 Optimal Allocations

Because of the externalities in the output market, decentralized equilibria in the model without government intervention are unlikely to be optimal. We now attempt to characterize optimal allocations in this economy. The social planner first chooses r ; the fraction of agents that engage in production. To minimize on notation, we will arrange the names of agents so that those in $[0; r)$ have produced and those in $[r; 1)$ have not. The planner then picks an intensity level α_i for each agent (although it is clear that α_i will be set to zero for any agent that did not produce). The entire planning problem can be written as

$$\max_{\alpha_i; r} \int_0^r \mu \cdot \int_0^r \alpha_j dj \cdot F_i a_i^2 + \int_r^1 \mu \cdot \int_0^r \alpha_j dj \cdot \frac{1}{4} F_i a_i^2 \quad (8)$$

subject to

$$\alpha_i \geq 0 \text{ for all } i$$

$$0 \leq r \leq 1$$

Note that the formulation of this problem allows the planner to choose different values of α_i for different producing agents. However, since the value of λ depends on the average α_i while the search cost is convex in α_i ; it is never optimal to do this. This is proven in the following lemma.

Lemma 5 The solution to (8) has $\alpha_i = \alpha$ for all $i \in [0; r)$:

Proof: See appendix. ¥




Figure 7: $H_i \circ G^C$

Using this result, the planning problem can be rewritten as

$$\begin{aligned} \max_{\theta, r} & r^i (1 - \frac{3}{4})^{-1} (r^\theta) F_i a^{\theta^2} + r (\frac{3}{4} F_i \tau) \\ \text{subject to} & \end{aligned} \tag{9}$$

$$\theta \geq 0; 0 \leq r \leq 1;$$

Note that for any (positive) value of r ; θ will be chosen to maximize

$$(1 - \frac{3}{4})^{-1} (r^\theta) F_i a^{\theta^2};$$

Since this expression is zero when θ is zero, the optimal value of this expression will be nonnegative. Hence, when θ is chosen optimally the first term in the objective function is nondecreasing in r : For a fixed value of r ; the first and second derivatives of the objective function with respect to θ are

$$\begin{aligned} \frac{\partial}{\partial \theta} &= r ((1 - \frac{3}{4})^{-1} (r^\theta) r F_i 2a^\theta) \\ \frac{\partial^2}{\partial \theta^2} &= r^i (1 - \frac{3}{4})^{-1} (r^\theta) r^2 F_i 2a^\theta : \end{aligned}$$

The first derivative is always zero at $\theta = 0$: For small enough values of r ; the second derivative is negative at zero. Furthermore, as θ increases, $r^{2\theta}$ decreases until it is negative (and then remains negative), so that the second derivative is negative for all values of θ : Hence $\theta = 0$ would be the

optimal choice for any value of r satisfying

$$r \cdot \frac{S}{(1 - \frac{3}{4}) F^{-1}(0)} = \frac{2a}{F^{-1}(0)} \quad (10)$$

The expected cost of production is high enough that it cannot be optimal to have any agent produce and then not engage in search. Hence any value of r satisfying (10) *cannot* be an optimal choice.

For higher values of r ; the next lemma shows that there is a unique interior solution to the first-order condition, and that it gives the optimal value of θ :

Lemma 6 There is a unique interior solution to the equation

$$(1 - \frac{3}{4})^{-1} (r^\circ) r F = 2a^\circ;$$

and this value of θ solves problem (9).

Proof: See appendix. ¥

It is helpful to rewrite the equation defining the optimal value of θ (which we denote by θ°) as

$$\theta^\circ = \frac{(1 - \frac{3}{4})^{-1} (r^\circ) r F}{2a} \quad (11)$$

It is interesting to compare this to (5), which defines the (interior) equilibrium value of θ . They differ only in that the optimal θ depends on the marginal product of θ in increasing θ whereas the equilibrium θ depends on the average product. Hence the equilibrium is unlikely to be optimal, and may involve either too little or too much search effort. For low values of θ the marginal product is above the average and hence the equilibrium amount of search effort will tend to be too low. This is the case of a positive trade externality: additional search effort by some agents will make the efforts of other agents more productive. For high values of θ ; however, the marginal product falls below the average. In this case there is a negative trade externality, or a congestion (crowding-out) effect. The equilibrium level of search intensity will be too high in this case.

Example: Figure 8 demonstrates the relationship between the optimal and equilibrium levels of search effort for our chosen parameter values. The intersection of the 45-degree line with the solid curve gives the optimal θ ; while its intersection with the dashed curve gives the equilibrium θ with a passive government. The graph confirms that this particular economy is in the negative externality

region.

α




Figure 8: Equilibrium and optimal levels of α

In principle, the function $\alpha(r)$ implicitly defined in (11) should be substituted into the objective function, which would then be maximized with respect to r : This seems to be analytically intractable, so we turn to our example.

Example: For the specified matching function, the social planner's objective function is given by

$$r \left(1 - \frac{3}{4}\right) - \frac{1}{1 + (r\alpha)^2} \left(F - a\alpha^2 + r\left(\frac{3}{4}F - c\right)\right)$$

The optimal value of gamma is then given by

$$\alpha(r) = \frac{\sqrt{\frac{1}{2a} \left(2a + 2 \sqrt{(1 - \frac{3}{4}ar^2F + ar^2F)}\right)}}{r^2} \quad (12)$$

for high enough values of r : Let $V(r)$ denote the value of the objective function when α is given by (12). We use the parameter values from above. In this case, $\alpha(r)$ is zero for r less than $\frac{1}{5} \approx 0.447$: Figure 9 plots the objective as a function of r above this value. It is clear that in this case $r = 1$ is the optimal choice. Note that having a relatively low number of agents producing is worse than having no agents producing, that is, it yields a value below zero. The optimal value of α from equation (12) is approximately 1.1118; which is below the equilibrium value of approximately 1.2248. This again verifies that the economy with passive policy is in the negative externality region. If we reduce F to 4, the picture changes to that in figure 10. In this case $r = 1$ is still better than any other positive value of r ; but now the value of the objective function is negative at this point, so




Figure 9: $V(r)$ when $F = 5$

that $r = 0$ is the optimal choice.

This example has the property that the optimal plan has either all agents producing or no agents producing. We do not know if it is possible to construct examples where the optimal choice of r is interior. □

4. Learning

The existence of multiple equilibria brings up the difficult issue of equilibrium selection. Several different criteria for selecting among multiple rational expectations equilibria have been proposed in the literature. Some of these criteria are axiomatic, in the sense that they select equilibria that possess certain properties. Perhaps the best known of these is the “minimum state variable” criterion of McCallum [11], which involves simply ignoring equilibria in which extrinsic uncertainty has real effects. Another possibility is to select Pareto optimal equilibria. This latter approach is clearly inappropriate for our analysis, since it rules out by assumption equilibrium explanations of coordination failures, the very phenomenon we study. In addition, Guesnerie and Woodford [6] show that such purely formal criteria often have undesirable properties, such as failing to exist or lacking upper hemi-continuity in model parameters.

The other main class of selection criteria involves examining the stability of each equilibrium with respect to some dynamic process. This process is typically interpreted as one of agents learning about some parameters of the economy. This is the approach we take. The literature on learning has




Figure 10: $V(r)$ when $F = 4$

focused primarily on (i) whether or not learning converges to a rational expectations equilibrium and (ii) whether particular equilibria are stable or unstable with respect to a given learning process. It is easy to show that in our model (i) is not a concern; learning always converges to one of the equilibria studied above. Which equilibrium the economy converges to depends on the actual sequence of realizations of uncertainty, and hence is random. We take issue (ii) a step further than the previous literature by studying the properties of the probability distribution over the set of equilibria induced by Bayesian learning.

Lucas [10], in a discussion on the use of “rationality” in economic theory, suggests that rational expectations equilibria can often be “interpreted as a description of a kind of stationary ‘point’ of [a] dynamic, adaptive process.” We follow Lucas’ approach in the sense that we use the learning scheme to investigate the plausibility of the different possible equilibria.⁹ But we go a step further in that we study the probability distribution over the set of equilibria that is induced by an specific adaptive behavior mechanism.

There is always a degree of arbitrariness in the selection of a learning rule, and our specific results clearly depend on the rule that we employ. Nevertheless, we think that our adaptive mechanism is not completely arbitrary. First, it is well known that expected utility maximizers use Bayes rule to update beliefs (see Blume and Easley [2]). Second, we assume that agents begin with diffuse priors, which seems to be a reasonable assumption in large economies (“on average”). There is still

⁹ To use the adaptive behavior to predict macroeconomic performance along the learning transition does not constitute, for Lucas, “a serious hypothesis.” For example, the initial date of any data set being studied is not the $t = 0$ of the theoretical model in any behavioral sense.

some nonrationality in our assumptions about agent's beliefs over the probability distribution of the random endogenous variables *during the learning transition*. This, however, could be justified with a bounded rationality argument (see Guesnerie and Woodford [6]). Finally, this approach represents a starting point for our analysis of probabilistic equilibrium selection, and we plan to address other approaches in future work.

Let us consider the problem faced by an agent at the outset of the economy (this could be interpreted as the problem faced by agents after a radical structural change in the economy). Assume that, without any additional inference about the economy, the agents are able to compute the value of producing when every other agent produces, v^H , and the value when no other private agents produce, v^L , and that they know the two possible values that c can take: c^H and c^L .

However, let us also assume that agents do not know the true probability distribution associated to the two relevant random variables in the environment: the cost c and the other agents' production decisions r . We want to consider the implications of "rational" learning for the plausibility of the different possible final equilibria (after the learning process has converged). We assume that agents update beliefs using Bayes rule. As is usual in the learning literature, though, some extra assumptions about the prior distributions are necessary to be able to proceed. These assumptions are mainly directed at having the distribution of prior beliefs be a member of the same conjugate family of distributions as the final (after convergence) probability distributions underlying the economy.¹⁰

We introduce the following three assumptions. First, we assume that every agent believes that all other private agents in the economy will act identically, either producing or not producing. As a result of this, the support of the random variables subject to learning are fixed and known. They both have only two elements: $c \in \{c^H, c^L\}$; and $r \in \{1, \bar{A}\}$. The learning problem consists of figuring out the following two probabilities

$$p \equiv \Pr(c_t = c^L)$$

$$q \equiv \Pr(r_t = 1):$$

Second, we assume that agents believe that these probabilities are constant. This is not true during the transition, while learning is taking place, for the case of q . The variable r is an endoge-

¹⁰ Note that since r is an endogenous variable, the true final distribution depends also on what happens during the transitional period.

nous variable in the model that shows, for example, a strong autocorrelation during the transitional period. However, as the economy converges to a rational expectations equilibrium, this belief becomes eventually correct. We can interpret this assumption as part of a bounded rationality limitation on the part of the agents. As will become clear, the *realization* of the stochastic process for r during the transition can actually show a very complicated pattern, not necessarily appearing incompatible with the simple (although incorrect) belief held by agents.

Third, we assume that each agent has identical beliefs and that they start the learning process with independent diffuse priors over the values of p and q . That is, their beliefs are initially represented by a uniform distribution on $[0; 1]$ for each of the probabilities.

Let $(p_t; q_t)$ denote the expected value of $(p; q)$ according to the beliefs held by the agents at date t . In particular, we have $p_0 = q_0 = \frac{1}{2}$. In each iteration of the learning algorithm, the observation of $(c_t; r_t)$ provides useful information for updating $(p_t; q_t)$. Bayesian updating of beliefs allows us to write the expected value of the parameter of the distributions after t iterations as

$$p_{t+1} = \hat{p}_t \quad \text{if } c_t = c^H$$

$$p_{t+1} = \hat{p}_t + (1 - \hat{p}_t) \lambda \quad \text{if } c_t = c^L$$

and

$$q_{t+1} = \hat{q}_t + (1 - \hat{q}_t) \lambda \quad \text{if } r_t = 1$$

$$q_{t+1} = \hat{q}_t \quad \text{if } r_t = \tilde{A}$$

where $\lambda = \frac{1}{(t+2)(t+3)}$.¹¹ The agent's expected utility based on the beliefs held at the beginning of time t are given by

$$u(p_t; q_t) = q_t \lambda^H + (1 - q_t) \lambda^L + p_t c^L + (1 - p_t) c^H$$

and the agent will decide to produce if and only if $u(p_t; q_t) > 0$.

As in Howitt and McAfee [7], we can represent the learning process in $(p; q)$ space as in figure

¹¹ This updating rule is a special case of the following recursive algorithm for estimating the mean of the distribution of x : $x_{t+1}^e = x_t^e + \theta_t (x_t - x_t^e)$; where x_t^e is the estimation, x_t is the period- t realization and θ_t is a decreasing sequence of positive numbers that approaches zero as t goes to infinity (see Guesnerie and Woodford [6]). In our case, x_t has a Bernoulli distribution and $\theta_t = \frac{1}{(t+3)}$. It is not hard to show that when instead $\theta_t = \frac{1}{t}$ the recursive algorithm uses the sample mean as the relevant statistic.

11. The box in the figure represents the set of possible beliefs. A point in the box, such as point




Figure 11: Dynamics of beliefs

a, corresponds to a particular value of $(p_t; q_t)$: If the point falls in the region P; where the expected profit from producing is positive, all agents will produce and hence the value of q will increase. Howitt and McAfee [7] show that posterior beliefs always lie on the line segment connecting the prior beliefs with one of the corners of the box. From point a we would move to b if c^H is observed and to c if c^L is observed. Similarly, if the original point is x; then expected profits are negative, no agents will produce, and q will decrease. We move to point y if c^H is observed and to z if c^L is observed.

Bayesian updating consistently estimates the value of \bar{p} ; i.e. $p_t \rightarrow \bar{p}$ as $t \rightarrow \infty$; and hence the learning process must converge to $(p; q)$ equal to either $(\bar{p}; 0)$ or $(\bar{p}; 1)$: Define $\frac{1}{4}$ to be the probability that $q_t \rightarrow 1$ as $t \rightarrow \infty$. It is easy to see that $c^L < \bar{c}^L < \bar{c} < \bar{c}^H < c^H$ implies that we have $0 < \frac{1}{4} < 1$. Hence the adaptive learning process induces a probability distribution over the two equilibria to which it can converge. We now examine how the government policies in the previous section influence $\frac{1}{4}$ and hence the likelihood of each equilibrium. The analytics of this problem seem intractable, and therefore we turn to our numerical example.




Figure 12: $i_s^H; \frac{1}{4}$ generated by some θ^G

Example: We have shown above that for our particular parameter values, the bad equilibrium exists for all values of θ^G and hence $\frac{1}{4}$ is always less than unity. We construct a grid of values for θ^G , and for each of these values we compute the resulting $\frac{1}{4}$ through simulation.¹² It is straightforward to calculate i_s^H ; the value of the good equilibrium to private agents, for each θ^G in the grid using the equations in the previous section. The result is a collection of points $i_s^H; \frac{1}{4}$ that are generated by different possible values of the policy parameter. Figure 12 presents a plot of one collection of such points.

Increases in θ^G correspond to counterclockwise movements along the arc traced out by these points. Very low values of θ^G result in a relatively low probability of attaining the good equilibrium. As θ^G increases, i_s^H also increases and attains its maximum value at \mathbf{b}^G ; while $\frac{1}{4}$ is simultaneously increasing. This movement to the northeast on the graph is unambiguously good - the good equilibrium is becoming both better and more likely. Afterwards, however, a tradeoff sets in. Further increases in θ^G continue to increase $\frac{1}{4}$; making the good equilibrium more and more likely. However, at the same time the value of the good equilibrium is *decreasing*. There is a value of θ^G that maximizes the probability of reaching the good equilibrium (call it \mathbf{e}^G), but it yields a substantially lower value of being in this equilibrium than does \mathbf{b}^G : Beyond this point, further increases in θ^G

¹² The computations presented here are preliminary and more robustness checks need to be performed. All simulations were done in FORTRAN; the code is available upon request.

are unambiguously bad - they decrease both β^H and β^G :

The values of β^G in the range $\beta^G; \beta^G$ trace out the frontier of the set of possibilities in $\beta^H; \beta^G$ space. One can then ask which of these policies is optimal. The answer, of course, depends on the preferences of the policymaker in this space. Since β^H is measured in utility terms, one candidate objective would be to maximize expected utility across equilibria, that is, to maximize

$$\begin{aligned} & \beta^H + (1 - \beta) 0 \\ = & \beta^H: \end{aligned}$$

Since private agents do not produce in the bad equilibrium, the utility value of that equilibrium is zero. Preferences such as these would generate indifference curves that are convex to the origin and hence would select a policy strictly between β^G and β^G :

Recall that before we considered learning, β^G seemed like an obvious candidate for the optimal policy. There is no way to eliminate the bad equilibrium, and at leave this policy maximizes the value of being in the good equilibrium. Using learning to do probabilistic equilibrium selection, however, changes the way we think about optimal policy in this example, and leads to the conclusion that we should select a higher value of β^G : 2

5. Concluding Remarks

Our goal in this paper has been to show that using adaptive learning to generate a probability distribution over the set of equilibria can lead to useful and interesting results. In particular, this type of analysis provides a unique answer to optimal policy questions, and for our example this answer differs from that which would be generated by *any* deterministic selection criterion. We have framed our analysis in terms of a particular model and a particular learning rule, but we believe that these ideas apply much more broadly and we plan to extend this analysis in future work. We conclude here by offering a brief discussion of three issues that we plan to address.

The type of model we study here also possesses sunspot equilibria, as shown by Howitt and McAfee [7]. Suppose that at the beginning of each period, one of two extrinsic states of nature is revealed; either there is sunspot activity that period or there is not. If agents are allowed to condition their beliefs about the actions of others on this random variable, then there exist equilibria where

all agents' production decisions depend on the realization of the sunspot variable. These equilibria are necessarily randomizations over the certainty equilibria of the model,¹³ and hence our analysis is easily extended to this case. If we assume that agent's prior beliefs are independent across states, then the probability $\frac{1}{4}$ is computed exactly as above. The probability of getting the good equilibrium in both sunspot states is then $\frac{1}{4}^2$; while the bad equilibrium in both states comes with probability $(1 - \frac{1}{4})^2$: The economy converges to a sunspot equilibrium with probability $2\frac{1}{4}(1 - \frac{1}{4})$: One interesting thing that comes from this analysis is that the probability of reaching a sunspot equilibrium is maximized when $\frac{1}{4}$ is equal to one-half, or when there is the greatest uncertainty about the eventual outcome. In other words, when the economy can go either way, it may go both ways.

Another obvious extension of the above analysis is to other types of government policies in this same model. Instead of using the demand-management policy studied here, the government could affect the value of producing by taxing or subsidizing search effort in order to correct for the market externality. Such a policy will have no effect on ψ^L ; but it will affect ψ^H and through this the probability $\frac{1}{4}$: In our example, it is possible to show that there is a unique tax rate that maximizes the value ψ^H : Whether or not the same type of tradeoff between maximizing ψ^H and maximizing $\frac{1}{4}$ arises under this policy as does under the demand-management policy discussed above remains to be seen. A final extension of the analysis that we find promising is to labor market issues. One could imagine changing the model so that workers and firms need to be matched in order for production to occur,¹⁴ and that the government has a policy that affects the efficiency of this matching process. Optimal policy analysis using probabilistic equilibrium selection could then be done with respect to a broad range of labor market policies.

¹³ In a richer model where there is a way for at least some agents to transfer wealth across sunspot states, sunspot equilibria are typically *not* randomizations over the certainty equilibria of the model. See Cass and Shell [3].

¹⁴ This is in the spirit of Howitt and McAfee [7].

Appendix A. Proofs of Lemmas

Lemma 1: The solution to problem (2) at the equilibrium value of $\frac{1}{2}$ is given by

$$\phi_i^* = \frac{(1 - \frac{3}{4}) \frac{1}{2} F}{2a}.$$

Proof: This value of ϕ_i^* is that given by the first-order condition for an interior solution. Hence it is the solution to the problem as long as it satisfies the inequality constraint. In equilibrium we have $\phi_i = \phi$ for all i , $\phi = r^*$; and $\frac{1}{2}(\phi; r) = 1(r^*) = \phi$: Our assumptions on 1 include $1(r^*) < 1$; which implies that we have

$$\phi < \frac{\phi}{1(r^*)}$$

or

$$\phi < \frac{1}{\frac{1}{2}}.$$

Hence the constraint is not binding at the equilibrium value of $\frac{1}{2}$: ✖

Lemma 2: There exists a unique interior solution for ϕ^* to equation (5).

Proof: Define

$$D = \frac{2a}{(1 - \frac{3}{4}) F}.$$

Then the right-hand side of (5) is given by

$$g(\phi) = \frac{1 - 1(r^*)}{D \phi}$$

and we have

$$g'(\phi) = \frac{1}{D \phi^2} \left[\mu \phi^0(r^*) - \frac{1(r^*)}{\phi} \right].$$

Evaluating this at zero (which requires applying L'Hopital's rule twice) yields

$$g'(0) = \frac{4a}{(1 - \frac{3}{4}) r^2 F} \phi^{00}(0) > 1;$$

where the inequality follows from (A7): Hence the right-hand side starts out above the left-hand side. The limit as ϕ goes to infinity of the right-hand side is zero (since 1 is bounded), so by continuity there is at least one interior solution.

For uniqueness, first note that our assumptions on f^1 imply that there exists an $\bar{x} > 0$ such that

$$f^1(x) > \frac{f^1(x)}{x} \text{ for } x < \bar{x}$$

and

$$f^1(x) < \frac{f^1(x)}{x} \text{ for } x > \bar{x}.$$

This implies that the right-hand side is increasing up to $r^\circ = \bar{x}$ and then decreasing afterwards. The second derivative of the right-hand side with respect to r° is given by

$$g^{00}(r^\circ) = \frac{1}{D} \frac{r^\mu}{r^\circ} f^{100}(r^\circ) r^\circ - \frac{2}{r^\circ} f^{10}(r^\circ) - \frac{f^1(r^\circ)}{r^\circ} \quad (A-1)$$

We first need to evaluate this at $r^\circ = 0$:

$$\lim_{r^\circ \rightarrow 0} g^{00}(r^\circ) = \lim_{r^\circ \rightarrow 0} \frac{1}{D} \left[\frac{f^{100}(r^\circ) r^2}{r^\circ} - \frac{2 f^{10}(r^\circ) r}{r^\circ} - \frac{f^1(r^\circ)}{r^\circ} \right]$$

Applying L'Hopital's rule (several times to the latter terms) yields

$$\begin{aligned} &= \lim_{r^\circ \rightarrow 0} \frac{1}{D} \left[f^{1000}(r^\circ) r^3 - f^{10}(r^\circ) r^3 - \frac{f^1(r^\circ) r^3}{3} \right] \\ &= \lim_{r^\circ \rightarrow 0} \frac{1}{D} m^{000}(r^\circ) r^3; \end{aligned}$$

which is negative by (A8): Hence the right-hand side is initially concave. As r° increases, the f^{100} term in A-1 falls, while (A6) implies that the difference between the marginal and the average product of f^1 increases until $r^\circ = \bar{x}$: This keeps the expression for g^{00} negative until at least \bar{x} : After \bar{x} , f^{100} becomes negative, so that g^{00} is negative until at least \bar{x} ; where the marginal and average products of f^1 cross.

This establishes that g is a concave function on $(0; \bar{x})$: Recall that g is decreasing after \bar{x} : Together these facts imply that g can cross the 45-degree line only once, or that the solution to (5) must be unique. ✎

Lemma 5: The solution to (8) has $r_i^\circ = r^\circ$ for all $i \in [0; r]$:

Proof: Fix an arbitrary value of r : Then the problem is to choose the function r° to solve

$$\max_{r^\circ(i)} r \left(1 - \frac{3}{4} \right) F^1 \left(\int_0^r r^\circ_i di \right) - \int_0^r r^\circ_i^2 di:$$

Define

$$\bar{\omega} = \frac{\int_0^r \omega_i di}{r}.$$

Suppose that it is not the case that $\omega_i = \bar{\omega}$ for almost all $i \in [0; r)$ and consider the alternative plan \mathbf{b} where $b_i = \bar{\omega}$ for all $i \in [0; r)$: Clearly we have

$$\int_0^r b_i di = \int_0^r \omega_i di;$$

so that the value of the first term in the objective function is the same under the two policies. To evaluate the second term, consider the problem of minimizing the integral of ω_i^2 subject to the constraint that the integral of ω_i is equal to some constant (in this case $\bar{\omega}r$). The first-order necessary condition for this problem entails setting ω_i equal to $\bar{\omega}$ for almost all values of i : This implies that we must have

$$\int_0^r \omega_i^2 di > \int_0^r b_i^2 di$$

and hence the value of the second term is lower under the new policy, contradicting the optimality of the original policy. ✖

Lemma 6: There is a unique interior solution to the equation

$$(1 - \frac{3}{4})^{1/\theta} (r^\theta) rF = 2a^\theta;$$

and this value of θ solves problem (9).

Proof: Both sides of the equation start at the origin, and condition (A7) guarantees that the right-hand side initially increases faster than the left-hand side. Since r^θ goes to zero as θ goes to infinity, the left-hand side is eventually larger and by continuity there is at least one solution.

Let θ^* be the lowest (positive) value of θ that satisfies the equation. Since the right-hand side is crossing (or at least touching) the left-hand side from above, its slope must be no greater than that of the left-hand side, i.e., we must have

$$(1 - \frac{3}{4})^{1/\theta^*} (r^{\theta^*}) F \leq 2a:$$

Conditions (A6) and (A8) then imply that we have

$$(1 - \frac{3}{4})^{100} (r^\circ) F < 2a$$

for all $\circ > \bar{\circ}$; since 100 is decreasing up to $\bar{\circ}$ and negative thereafter. Hence the slope of the right-hand side is less than that of the left-hand side for all $\circ > \bar{\circ}$: This implies that (i) the two curves cross at $\bar{\circ}$ (and are not just tangent) and (ii) the two curves do not cross again. This demonstrates that there is a unique interior solution to the first-order condition.

Since the first-order condition is a necessary condition for an interior solution, the solution to the maximization problem must either be zero or $\bar{\circ}$: Since the objective function is increasing in \circ for \circ close to zero, the solution must be $\bar{\circ}$: \yen

Lemma 4: If

$$\tilde{A} \cdot \frac{2a}{(1 - \frac{3}{4}) F}$$

holds, then the optimal level of search intensity for a private agent is given by

$$\circ_i^* = \frac{(1 - \frac{3}{4}) \frac{1(\tilde{A}^\circ G)}{\circ G} F}{2a}:$$

Proof: Condition (A4) implies that we have

$$\frac{1 \tilde{A}^\circ G^\dagger}{\tilde{A}^\circ G} \cdot 1$$

and hence

$$\frac{1}{2} = \frac{1 \tilde{A}^\circ G^\dagger}{\circ G} \cdot \tilde{A} \cdot \frac{2a}{(1 - \frac{3}{4}) F}$$

must hold. This implies that

$$\frac{1}{2}^2 \cdot \frac{2a}{(1 - \frac{3}{4}) F}$$

holds, as does

$$\frac{(1 - \frac{3}{4}) \frac{1}{2} F}{2a} \cdot \frac{1}{\frac{1}{2}}:$$

This demonstrates that the interior solution satisfies the constraint. \yen

References

- [1] R. Aiyagari and N. Wallace, Government transaction policy, the medium of exchange, and welfare, *Journal of Economic Theory*, **74** (1997), 1-18.
- [2] L. Blume and D. Easley, Rational expectations and rational learning, in “Organizations with Incomplete Information” (M. Majumdar, Ed.), pp. ??-??, Cambridge University Press, Cambridge, U.K., 1998.
- [3] D. Cass and K. Shell, Do sunspots matter? *Journal of Political Economy* **91** (1983), 193-227
- [4] P. Diamond, Aggregate demand management in search equilibrium, *Journal of Political Economy* **90** (1982), 881-894.
- [5] J. M. Grandmont, Stabilizing competitive business cycles, , *Journal of Economic Theory* **40** (1986) 57-76.
- [6] R. Guesnerie and M. Woodford, Endogenous fluctuations, in “Advances in Economic Theory, Sixth World Congress,” Vol II (J.-J. Laffont, Ed.), pp. 289-412, Cambridge University Press, Cambridge, U.K., 1992.
- [7] P. Howitt and R. P. McAfee, Animal spirits, *American Economic Review* **82** (1992), 493-507.
- [8] T. Keister, Money taxes and efficiency when sunspots matter, *Journal of Economic Theory* **83** (1998) 43-68.
- [9] R. Lagos, An alternative approach to market frictions: An application to the market for taxicab rides, manuscript, London School of Economics (1997)
- [10] R. E. Lucas, Jr., Adaptive behavior and economic theory, *Journal of Business* **59** (1986), S401-S426.
- [11] B. T. McCallum, On non-uniqueness in linear rational expectations models: An attempt at perspective, *Journal of Monetary Economics* **11** (1983), 139-168.
- [12] B. Smith, Efficiency and determinacy of equilibrium under inflation targeting, *Economic Theory* **4** (1994), 327-344.
- [13] M. Woodford, Stationary sunspot equilibria in a finance constrained economy, *Journal of Economic Theory* **40** (1986) 128-137.
- [14] M. Woodford, Learning to believe in sunspots, *Econometrica* **58** (1990), 277-307.