

Optimal growth with unobservable resources and learning

Yaw Nyarko^a, Lars J. Olson^{b,*}

^a *Department of Economics, New York University, New York, NY 10003, USA*

^b *Department of Agricultural and Resource Economics, University of Maryland, 2200 Symons Hall,
College Park, MD 20742, USA*

Received 12 July 1994; revised 11 May 1995

Abstract

This paper examines the problem of choosing optimal resource consumption from an imperfectly observable aggregate capital, wealth or resource stock which the decision-maker learns about over time. Learning is complicated by the fact that the planner receives information about an object (the true resource stock) which is a moving target. Under the assumption that the utility of zero consumption is $-\infty$ we show that the optimal policy follows a completely deterministic 'cautious' or 'minmax' policy that assumes the worst in each period and optimizes against that. When this model is compared to a model with completely observable wealth levels the following insights are obtained: (1) there is 'over-saving,' (2) investment and output processes are more volatile than consumption, and (3) regressions underestimate the risk aversion of agents. Information about the wealth or stock level is only valuable if it alters the support of the agent's beliefs. Thus, information may be statistically informative yet economically valueless. If information changes the support of the agent's beliefs then the optimal solution features an endogenous resource 'discovery' process.

JEL classification: D90; D83

Keywords: Growth theory; Consumption; Learning; Unobservable stocks

* Corresponding author.

1. Introduction

In many dynamic economic models agents must allocate a resource between consumption for the current period and investment for the future. In analyzing such problems it is typical to assume that each agent has perfect knowledge of the resources available at the time the allocation decision is made. In reality, there are many intertemporal resource allocation problems for which the true resource stock is not perfectly observable.

For many individuals a major portion of their wealth is stored in their house. A precise value for this wealth can only be obtained when the house is sold. The prices of other homes on the market provide some information about a property's value; however, since no two homes are identical such information will be imperfect. Similar uncertainty also applies to other forms of wealth such as art or antiques.¹ More generally, an individual's total wealth is described by the present discounted value of all future income from physical and human capital. This human capital is inherently unobservable. Even if markets exist for capitalizing one's future stream of wages, such human wealth cannot be determined without actually going to the market.

The allocation of natural resources is another important problem where stock levels are unobservable. Oil and mineral resources are extracted from stocks about which our knowledge is uncertain and subject to change as new reserves are discovered. Fish and wildlife species are harvested from randomly fluctuating populations that can only be partially observed. Even though fisheries management agencies spend substantial sums of money to forecast resource stocks, confidence intervals of $\pm 50\%$ of the mean are prevalent (Clark and Kirkwood, 1986).

Finally, aggregate measures of resources such as the money supply or capital stock are subject to imperfect measurement, yet this is not captured in aggregate models of economic growth.

This paper examines the problem of choosing optimal consumption from an imperfectly observable resource stock which the decision-maker learns about over time. The problem of optimal consumption with an unobservable resource stock differs in a significant way from the recent literature on decision-making and learning with an unknown, but *fixed* parameter (e.g., Grossman et al., 1977; Freixas, 1981; McLennan, 1984; Walters, 1986; Easley and Kiefer, 1988; Kiefer and Nyarko, 1989; El-Gamal and Sundaram, 1993, and Huffman and Kiefer, 1990). In the literature on learning with an unknown parameter, beliefs of the agent evolve over time, but the quantity over which there is uncertainty is constant. Hence, uncertainty is always being reduced over time. In our framework the learning process is complicated by the fact that the decision-maker is learning

¹ Baumol (1986) provides an interesting discussion of the uncertainties associated with investment in art.

about an object (the true resource stock) which is itself a moving target. Since the resource stock may evolve according to a stochastic transition equation, perfect knowledge of the true resource stock at one point in time does not imply the true stock is known at all future points.

In this paper we focus on a one-good stochastic growth model. The model encompasses a number of important problems in the theory of economic growth and natural resource economics, and is widely used in the study of individual consumption/saving behavior. The agent's beliefs about the initial wealth or resource stock are represented by a prior probability distribution. At the end of each period the agent receives an information signal about the true wealth level. Learning by the agent is specified in the standard way using the operator defined by Bayes' Rule.

We begin by imposing the interiority condition that the utility of zero consumption equals $-\infty$. First, we examine the case where the agent receives no information signals about the unobservable stock. In Section 3 we show that the optimal policy for the problem with unobservable resources and learning is identical to the optimal policy under an associated problem that is completely deterministic and easily characterized. The decision-maker assumes the initial stock is the lowest among those entertained under the prior beliefs and that the worst shock to production occurs in each period. The agent then optimizes against this 'cautious' or 'minmax' deterministic problem. The optimality of minmax behavior holds under general conditions and does not require any assumptions of concavity on utility or the production function.

This rather striking result implies that well-known solutions to deterministic models can be used to characterize the behavior of optimal policies for the problem with unobservable wealth levels. The implications are examined in Section 4. For example, in the standard deterministic growth model it is well-known that optimal consumption and resource stocks converge to a golden rule steady state. With unobservable resources a form of over-saving occurs and with probability one the true resource stock eventually becomes bounded away from the golden rule stock that determines optimal consumption. This result may be of interest given empirical findings of old-age over-saving behavior. A second implication of minmax behavior is that optimal output and investment are much more volatile than consumption. Optimal consumption solves a deterministic 'cautious' or 'minmax' problem and, hence, consumption is not random. On the other hand, output and investment vary randomly due to the stochastic shocks to production. Finally, we note that an econometrician who does not account for the presence of incomplete information may underestimate the agent's coefficient of relative risk aversion. This may be germane to the controversy over the equity premium puzzle discussed by Mehra and Prescott (1985).

Sections 5 and 6 extend the analysis to situations where, in each period, agents receive information signals about the unobservable stock or wealth level. If the information signals do not change the support of the agent's beliefs (e.g., normally

distributed signals) then the results of Sections 3 and 4 generalize immediately. If information increases the minimum of the support of beliefs over stock levels, then the optimal solution is the same as that of an associated deterministic problem that features an endogenously determined ‘discovery’ of new resources on the part of the agent. This provides a more realistic representation of many problems where the known stock or wealth is augmented over time by discoveries that may depend on previous consumption levels.

The crucial assumption used in obtaining our results is that the utility of zero consumption equals minus infinity. In a representative agent model, the assumption may be based on an equivalence between zero consumption and non-survival. In other settings it implies an absolute aversion to bankruptcy. Such absolute aversion may occur due to social stigmas or an extreme fear of being left with no resources in old age. The assumption holds for many popular utility functions used in the literature including the log utility function and the class of all constant relative risk aversion utility functions with coefficient no less than one.² A natural question to ask is: what happens when this property does not hold? In Section 7 we show that an agent either behaves as if he/she satisfies this property or else the agent becomes bankrupt in finite time with strictly positive probability.

The final section of the paper extends the Hall (1978) consumption model to include unobservable wealth levels. Under somewhat more restrictive conditions agents exhibit minmax behavior in this model as well.

There is some previous work that examines problems similar to the one in this paper. Kemp (1976), Robson (1979), Bhattacharya (1982), and Quyen (1991) all focus on with the problem of allocating an *exhaustible* resource when there is learning about the stock. In fact, Kemp (1976, footnote 7), refers to the potential for min-max behavior, but he does not provide a formal analysis, nor does he examine the economic implications associated with such behavior. The other papers rely on frameworks that differ substantially from that used here. In other work, Majumdar (1982) proves the existence of stationary optimal policies under certain continuity and compactness conditions. Using specific functional forms on both technology and preferences, Clark and Kirkwood (1986) examine the value of

² Such utility functions are often used in applied work (and are frequently the only utility functions for which explicit solutions are available). Arrow (1971) explains theoretically that the coefficient of risk aversion, Γ , should be around unity. Kydland and Prescott (1982) show that values of Γ between one and two best explain aggregate investment and output data. Tobin and Dolde (1971) use $\Gamma = 1.5$ in a life-cycle savings model. The CRRA class of utility functions is the preeminent class used in theoretical and applied finance. In this regard, Rubinstein (1976) lists ten desirable properties of the (generalized) log utility function. The assumption is strictly a local condition and outside some arbitrarily small neighborhood of $c = 0$ the utility function may be any continuous function that is desired. For positive consumption the Euler conditions will be determined outside this small neighborhood of the origin.

stock surveys when the objective is to maximize the expected discounted sum of consumption.

2. The unobservable resources problem P1

Let $y_t \geq 0$ be the stock of resources available at date t ($t = 0, 1, \dots$). The agent does not know the exact value of y_t . The agent's beliefs about y_t are represented by the prior distribution μ_t , a probability measure on the non-negative real line. The agent receives utility in each period from consuming an amount, c_t . Since the stock of resources is not necessarily known by the agent it is possible in principle that the agent will attempt to consume more than the available resource stock. It is convenient to interpret c_t as 'attempted' consumption. If the agent attempts to consume an amount $c_t > y_t$ then the 'attainable' consumption of the agent is y_t . Later we shall characterize when attempted consumption is always less than the true stock level. In such cases, the distinction between attempted and attainable consumption becomes irrelevant. Formally, however, we shall always consider c_t to be 'attempted' consumption and $\min\{c_t, y_t\}$ to be 'attained' consumption.

A consumption of $c_t \leq y_t$ results in an investment of $x_t = y_t - c_t$. If $c_t > y_t$ then investment is $x_t = 0$. Given an investment x_t , the next period stock is $y_{t+1} = f(x_t, r_{t+1})$ where $f: \mathfrak{R}_+ \times [a, b] \rightarrow \mathfrak{R}_+$ is the production function and r_{t+1} is the realization at date $t + 1$ of an unobserved independent and identically distributed process $\{r_t\}_{t=1}^\infty$ which takes values in some compact interval $[a, b]$ and has known marginal distribution Γ . The return that the agent receives at date t is denoted by $R(c_t, y_t) = u(\min\{c_t, y_t\})$ where $u: \mathfrak{R}_+ \rightarrow \mathfrak{R}$ is a utility function of 'attainable' consumption. When $c_t \leq y_t$ this reduces to the standard optimal growth model where $R(c_t, y_t) = u(c_t)$.

Beliefs about the date $t + 1$ stock, y_{t+1} , are obtained from the prior μ_t by taking into account both the date t consumption level, c_t , and the production function f . Let $P(\mathfrak{R}_+)$ be the set of all probability distributions on the non-negative real line. The mapping from consumption and beliefs at date t into the beliefs about the date $t + 1$ stock is denoted by $F: \mathfrak{R}_+ \times P(\mathfrak{R}_+) \rightarrow P(\mathfrak{R}_+)$, where

$$F(c_t, \mu_t)(A) = \int_{\mathfrak{R}_+} \Gamma(\{r_{t+1} | f(y_t - \min\{c_t, y_t, r_{t+1}\}) \in A\}) \mu_t(dy_t) \tag{2.1}$$

for each (Borel measurable) set A in \mathfrak{R}_+ and for each $(c_t, \mu_t) \in \mathfrak{R}_+ \times P(\mathfrak{R}_+)$.

Before consumption is chosen at date t , the observable history consists of the vector of past consumptions and previous beliefs. Any such vector is referred to as a partial history at date t and is denoted by

$$h_t = \{(\mu_0, c_0), \dots, (\mu_{t-1}, c_{t-1}), \mu_t\} \in \left[\prod_{i=0}^{t-1} P(\mathfrak{R}_+) \times \mathfrak{R}_+ \right] \times P(\mathfrak{R}_+). \tag{2.2}$$

The elements of h_t obey the updating relationship (2.1). A *date t policy* is any (Borel measurable) function $c_t = \pi_t(h_t)$ which specifies the date t consumption as a function of the partial history h_t (where the set of measures on the non-negative real line, $P(\mathfrak{R}_+)$, is endowed with the weak topology of measures). A *policy* $\pi = \{\pi_t\}_{t=0}^\infty$ is any collection of date t policies for $t = 0, 1, \dots$. A policy π and initial prior μ_0 result in a sequence of partial histories $\{h_t\}_{t=0}^\infty$. Define P_π as the probability measure induced on the partial histories by the policy π and initial prior μ_0 .

A consumption process $\{c_t\}_{t=0}^\infty$ is called a *measurable process* if for each t , $c_t \geq 0$ and c_t depends (measurably) only upon the partial history at date t . It follows that any policy results in a measurable consumption process. Of course, if $\{y_t\}_{t=0}^\infty$ is the stock process corresponding to some policy, then in some periods an unattainable consumption level, $c_t > y_t$, may be attempted. A measurable consumption process which satisfies $c_t \leq y_t$ with P_π -probability one for all t is called an *attainable consumption process*. This is equivalent to $c_t \leq \min\{y \in \text{Supp } \mu_t\}$ for all t . The associated policy is called an *attainable policy*.

Given any date zero prior μ_0 , a policy π results in a sum of discounted returns given by

$$V_\pi(\mu_0) = E_0 \sum_{t=0}^\infty \delta^t u(\min\{c_t, y_t\}), \tag{2.3}$$

where $\delta \in (0, 1)$ is the discount factor. Since the $\{c_t, \mu_t\}_{t=0}^\infty$ process is stochastic, expression (2.3) takes expectations, E_0 , at date zero, where the expectations are those corresponding to P_π , the probability induced by initial prior μ_0 and policy π (consult Blackwell (1965), for details on this). A policy π^* is *optimal* if $V_{\pi^*}(\mu_0) \geq V_\pi(\mu_0)$ for all policies π and all date zero prior beliefs μ_0 . The value function is defined by $V(\mu_0) = \text{Supp}_\pi V_\pi(\mu_0)$ where the supremum is taken over all policies, π .

The agent’s problem of choosing a policy to maximize the expected sum of discounted returns is referred to as *Problem P1*. The optimal policy (when it exists), satisfies the functional equation:

$$V(\mu) = \max_{c \geq 0} E_\mu [u(\min\{c, y\}) + \delta V(F(c, \mu))]. \tag{2.4}$$

2.1. Assumptions

We impose the following assumptions on the model. Each of these assumptions, except possibly (R.1), are standard in the literature.

- R.1. $R(0, y) = -\infty$ for all $y \geq 0$, and given any sequence $\{c^n, y^n\}_{n \geq 1} \in \mathfrak{R}_+ \times \mathfrak{R}_+$ which converges to $(0, y)$, $R(c^n, y^n)$ converges to $-\infty$.
- R.2. $R(c, y)$ is continuous at each (c, y) such that $c > 0$.
 - f.1. $f(x, r)$ is continuous in x and r .
 - f.2. $f(0, r) = 0$ for all r and $f(x, r) > 0$ for all r and $x > 0$.

- f.3. $f(x, r)$ is strictly increasing in x for fixed r .
 f.4. There is an $\underline{x} > 0$ such that for all $x \in (0, \underline{x})$ and for all r , $f(x, r) > x$.
 f.5. There exists a $\bar{x} > 0$ such that for all $x \geq \bar{x}$ and for all r , $f(x, r) < x$.

Under (R.1) it may be possible for all policies starting from positive initial stock levels to lead to a sum of discounted returns equal to minus infinity. Assumption (f.4) is a ‘productivity’ assumption on the production function which requires that the production function lie above the 45 degree line in a small neighborhood of the origin. It implies that a fixed positive consumption level can be sustained from any arbitrarily small initial stock. This rules out the value function being unbounded below from strictly positive stock levels. Assumption (f.5) has the consequence of bounding the set of possible stock levels from above. In particular, it is easy to show that under (f.1) and (f.5) if y_0 is the initial stock and $\{y_t\}_{t=0}^{\infty}$ is the stock process resulting from any policy then for all $t \geq 0$, $y_t \leq \max\{y_0, \bar{x}\}$. Both (f.4) and (f.5) are standard in the optimal growth literature.

Since the quantity \bar{x} in (f.5) can be made arbitrarily large, we assume that the support of the agent’s beliefs about the date 0 stock level lies in the set $[0, \bar{x}]$. This imposes little loss of generality since the agent knows that the stock level will become less than \bar{x} in finite time, even if no consumption takes place at any date. We also assume that the initial beliefs of the agent are bounded away from zero. Without this assumption it is possible that, under (R.1), the value function equals minus infinity and any policy is optimal. Since \underline{x} in assumption (f.4) can be made arbitrarily small, we shall take this to be a lower bound of the support of the agent’s initial beliefs. To summarize, we assume

B.1. The support of the prior beliefs on the initial stock, μ_0 , is some subset of the set $[x, \bar{x}]$.

We always implicitly assume that the true initial stock lies in the support of the initial prior distribution of the agent. All results in the paper are stated ‘with P_{π} -probability one’. If the initial stock lies outside of the support of the prior distribution the agent may observe events which were previously assigned zero probability and P_{π} -probability one sets may have no relationship with probability one sets of an outside observer who knows the value of the initial stock. In this paper we do not provide a theory of how an agent might behave under such situations (but see Nyarko, 1991, for a discussion of this very issue).

3. Optimality of minmax behavior

We proceed to show that under (R.1) the agent behaves very ‘cautiously’ and that the optimal solution with unobservable resources also solves a ‘minmax’ deterministic optimal growth problem with:

1. an initial stock equal to the lower bound of the agent’s prior beliefs, $y_0 = \min\{y \geq 0 \mid y \in \text{Supp } \mu_0\}$; and
2. a deterministic production function equal to the original production function in the worst state, $\underline{f}(x) \equiv \min_r f(x, r)$.

In other words, the agent exhibits a sort of minmax behavior, assuming the worst possible stock level and the worst shock to production at each date and optimizing against that. This ‘minmax’ or ‘cautious’ deterministic problem is stated below as Problem P2. (Problem P1 is the general problem with unobservable stocks presented in Section 2.) Consider an initial stock level $\underline{y} > 0$ as given.

Problem P2. $\max_{\{c_t\}} \sum_{t=0}^{\infty} \delta^t R(c_t, y_t)$ s.t. $0 \leq c_t \leq y_t$ for all $t \geq 0$ where $y_0 = \underline{y}$, $y_{t+1} = \underline{f}(y_t - c_t)$ for all $t \geq 0$, and $\underline{f}(x) \equiv \min_r f(x, r)$. The significant characteristic of Problem P2 is the absence of any uncertainty. The agent simply assumes the worst with certainty.

That optimal consumption with unobservable resources is equivalent to optimal consumption in a deterministic problem may seem surprising; however, the rationale is relatively straightforward and follows immediately from two observations given below as Proposition 3.1 and Lemma 3.2. First, any policy that is not attainable leads to zero consumption in the next period with positive probability. Under (R.1) zero consumption in any period results in utility of minus infinity. Hence, any optimal policy for the unobservable resources problem must be attainable. Second, any consumption process that is attainable for the unobservable resources problem must be feasible for the ‘minmax’ deterministic problem. Together these two observations imply that the search for optimal policies for both problems P1 and P2 can be restricted to the same set of consumption policies. Since problems P1 and P2 have the same utility function they must have the same solutions. We now proceed with the formal analysis.

Proposition 3.1. *If π is an optimal policy from some initial prior μ_0 then π is attainable.*

Proof of Proposition 3.1. Let $\{c_t, \mu_t\}_{t=0}^{\infty}$ be the consumption and posterior beliefs processes generated by π . Define $\underline{y}_t \equiv \min\{y \in \text{Supp } \mu_t\}$. Suppose that at some date t , the ‘attempted’ consumption level c_t is not attainable. Then $c_t > \underline{y}_t$. From definition of \underline{y}_t , this implies that under the agent’s beliefs there is a strictly positive probability that the stock level will be depleted at date t . From (f.2) this implies that with strictly positive probability from periods $t + 1$ onwards the resource stock and consumption levels will be zero. Under (R.1) this will result in a utility of $-\infty$. Under assumptions (f.4) and (B.1) the agent can choose an attainable policy that yields finite utility at each date. Hence, an agent who satisfies (R.1) will not choose an unattainable policy, since such a policy gives an expected total return of $-\infty$. This implies that an *optimal* policy is attainable. \square

Lemma 3.2. *Let $\{c_t\}_{t=0}^{\infty}$ be an attainable consumption process for the unobservable resources Problem P1 with initial prior μ_0 . Then $\{c_t\}_{t=0}^{\infty}$ is feasible for the deterministic ‘minmax’ Problem P2 with initial stock $\underline{y}_0 = \min\{y \in \text{Supp } \mu_0\}$. In other words, for all t on each sample path,*

$$0 \leq c_t \leq \underline{y}_t, \text{ where } \underline{y}_{t+1} \equiv \underline{f}(\underline{y}_t - c_t). \tag{3.1}$$

Proof of Lemma 3.2. Let $\{c_t, \mu_t\}_{t=0}^\infty$ be the consumption and posterior beliefs processes generated by some attainable policy from initial prior μ_0 . An induction argument now proves both (3.1) and

$$\underline{y}_t = \min\{y \in \text{Supp } \mu_t\}. \tag{3.2}$$

Fix any sample path. From the definition of \underline{y}_0 , (3.2) holds for $t = 0$. Since the consumption process is attainable by assumption, this implies that $c_0 \leq \underline{y}_0$. Hence, (3.1) holds for $t = 0$.

Next, suppose that (3.1) and (3.2) hold for some date t . If y_t and y_{t+1} are the true stock levels at dates t and $t + 1$ respectively, then $y_{t+1} = f(y_t - c_t, r_{t+1})$ where r_{t+1} is the unobserved realization of the shock to the production function at date $t + 1$. Under the induction hypothesis, the lower bound of the support of the distribution of y_t is \underline{y}_t . Hence, the lower bound on the support of beliefs over y_{t+1} is $\min_{r_{t+1}} f(\underline{y}_t - c_t, r_{t+1}) \equiv f(\underline{y}_t - c_t) \equiv \underline{y}_{t+1}$ and (3.2) holds at date $t + 1$. Since the consumption process is attainable, $c_{t+1} \leq \underline{y}_{t+1}$ and (3.1) holds at date $t + 1$. By induction, (3.1) and (3.2) hold for all t . \square

Together, Proposition 3.1 and Lemma 3.2 lead to our main proposition on the behavior of optimal policies.

Proposition 3.3. Fix any initial prior μ_0 and define $\underline{y}_0 \equiv \min\{y \in \text{Supp } \mu_0\}$. Then the consumption process $\{c_t\}_{t=0}^\infty$ is optimal for the optimal growth problem with unobservable stocks and initial prior μ_0 (i.e., Problem P1), if, and only if, it is also optimal for the 'minmax' deterministic optimal growth problem with initial stock \underline{y}_0 and production function f (i.e., Problem P2).

Proof of Proposition 3.3. Fix any initial prior μ_0 and let $\underline{y}_0 \equiv \min\{y \in \text{Supp } \mu_0\}$. Define P_1^A to be the set of all consumption processes, $\{c_t\}_{t=0}^\infty$, that are attainable for Problem P1. Define P_2^F to be the set of all consumption processes that are feasible for Problem P2 from initial stock level \underline{y}_0 . Proposition 3.1 implies that any optimal consumption process for problem P1 lies in P_1^A . Hence, in obtaining optimal consumption processes for Problem P1 it suffices to optimize over P_1^A . From Lemma 3.2 it follows that $P_1^A \subseteq P_2^F$. It should be obvious that $P_2^F \subseteq P_1^A$. So $P_1^A = P_2^F$. Further, the objective functions in both Problems P1 and P2 are the same (i.e., defined via u , the utility function). Since the two problems have the same objective function and the same relevant constraint set, they must have the same optimal solutions. This proves Proposition 3.3. \square

4. Implications of minmax behavior

4.1. Dynamics and over-saving

We now review some well-known results for the optimal growth Problem P2 with deterministic production function, $f(x) = \min_r f(x, r)$, and observable stocks.

To aid in characterizing the behavior of optimal processes we impose standard concavity assumptions on the utility and production functions, and an Inada condition on the production function.

R.3 The utility function $u(c)$ is increasing and strictly concave in c .

f.6 The production function $f(x, r)$ is strictly concave in x for each r .

f.7 $\lim_{x \rightarrow 0} f'(x, r) = \infty$ and $\lim_{x \rightarrow \infty} f'(x, r) < 1$ for each r , where $f'(x, r)$ is the derivative of f with respect to x , which is assumed to exist at each $x > 0$ and to be continuous at each (x, r) with $x > 0$.

The golden rule input, x_δ , output, y_δ , and consumption, c_δ , for the minimum production function, \underline{f} , are defined by

$$\underline{f}(x_\delta) = 1/\delta, y_\delta = \underline{f}(x_\delta) \text{ and } c_\delta = \underline{f}(x_\delta) - x_\delta, \quad (4.1)$$

where δ is the discount factor. Under assumptions (R.1)–(R.3) and (f.1)–(f.7), it is well known that the optimal input, output and consumption processes for Problem P2 converge to their golden rule values. Proposition 3.3 then implies that the optimal consumption process for the model with unobservable stocks, Problem P1, converges to the golden rule consumption level, c_δ , of the production function \underline{f} .

These facts enable us to examine the evolution of the true resource stock level in Problem P1 with unobservable stocks. To rule out trivialities it is necessary to assume that the agent's initial prior is non-degenerate (i.e., its support contains at least two points) and that the true stock level is not equal to the minimum of the support of the agent's prior beliefs. Under these conditions, an optimizing agent exhibits a strong form of over-saving relative to the stochastic optimal growth problem with production function, $f(x, r)$, and observable stock levels.

This is most easily illustrated by examining the case where the production function is deterministic. Fig. 1 shows a typical production function $f(x)$ along with the mapping $y_{t+1} = f(y_t - c_\delta)$, where c_δ is the golden rule consumption level. At $y_t = y_\delta$, this mapping has slope $f'(y_\delta - c_\delta) = 1/\delta > 1$, cutting the 45° line from below. It also has an upper fixed point denoted by y^* . Suppose the agent solving Problem P1 with unobservable stocks has a prior whose support is any non-degenerate interval with lower bound y_δ . Because this is the golden rule stock for the deterministic Problem P2, the optimal policy for an agent solving Problem P1 is to consume c_δ in each period. Hence, the behavior of the actual stock is governed by the mapping $y_{t+1} = f(y_t - c_\delta)$. From Fig. 1 it is clear that if the true initial stock is $y_0 > y_\delta$, then the optimal resource stock converges to y^* , the greater fixed point of the mapping $y_{t+1} = f(y_t - c_\delta)$. The agent behaves as if the stock level is y_δ and chooses consumption equal to c_δ in each period, even though the true resource stock approaches y^* over time. While the agent knows that the true stock converges to y^* with probability one, the agent does not know exactly how close to y^* the true stock is at any given date. In each period the support of the agent's beliefs has a lower bound at y_δ , and over time the support of the agent's beliefs tends to the set $[y_\delta, y^*]$. Since there is a very large penalty




Fig. 1.

to making a mistake under (R.1), the agent behaves very cautiously and always chooses the minmax golden rule consumption.

Relative to the true resource stock the agent is consuming too little and saving too much in every period. This over-saving is a general phenomenon that also holds in models with stochastic production functions, $f(x, r)$, under assumptions (R.3), (f.6), and (f.7), where the initial prior is any non-degenerate prior. Define $\bar{f}(x) = \max_r f(x, r)$ and let \bar{y}^* and \bar{y}^{**} be the largest fixed points of the maps $y_{t+1} = \bar{f}(y_t - c_\delta)$ and $y_{t+1} = \underline{f}(y_t - c_\delta)$, respectively (see Fig. 2). Observe that y_δ is the smallest fixed point of the map $y_{t+1} = \bar{f}(y_t - c_\delta)$. Again, suppose that the true stock level is not equal to the minimum of the support of the prior beliefs.

For the cautious or minmax problem the optimal consumption converges to the golden rule level c_δ . By Proposition 3.3 this is also the limit of the optimal consumption process for the unobservable resources problem P1. For problem P1 the true stock eventually enters and remains in the set $[y^*, y^{**}]$. Once again there is over-saving relative to the optimal growth problem with observable stocks. This




Fig. 2.

discussion is formally summarized in the following proposition whose proof is given in the appendix.

Proposition 4.1. Assume (R.1-R.3), (f.1-f.7), and (B.1), and suppose that the initial prior μ_0 is non-degenerate and that the true stock is not equal to the minimum of the support of the prior beliefs. If $\{c_t, \mu_t, y_t\}_{t=0}^\infty$ are the optimal consumption process, the associated process of posterior beliefs, and true resource stocks, then $\lim_{t \rightarrow \infty} c_t = c_\delta$, $\lim_{t \rightarrow \infty} \min\{y \in \text{Supp } \mu_t\} = y_\delta$, and with probability one the true resource stock enters the set $[y^, y^{**}]$ and stays there forever.*

Remark 1. The limiting value of the difference between actual stocks and the golden rule stock is at least $K = (y^* - y_\delta)$. This provides a lower bound on the long-run ‘excess saving’ of the agent no matter how small the uncertainty about the initial stock. The agent may believe the initial (date 0) stock level lies in the

set $[y, y + \epsilon]$, with $\epsilon > 0$ arbitrarily small. Proposition 4.1 implies that the limiting value of the 'excess saving' is at least K , regardless of the value of $\epsilon > 0$.

Remark 2. Using the techniques of Majumdar et al. (1989) it is possible to show that the true stock process $\{y_t\}_{t=0}^{\infty}$ converges in distribution to a unique invariant distribution (or stochastic steady state) with support in $[y^*, y^{**}]$.

Remark 3. The assumption that the true stock is not equal to the minimum of the support of prior beliefs may be relaxed if the production function is stochastic (i.e., if there does not exist an $x > 0$ such that $f(x, r) = f(x, r')$ for all r and r').

Remark 4. It should be clear that similar types of over-saving behavior will be observed in a finite horizon version of the problem. Many empirical studies of agents' life cycle saving behavior have documented the phenomenon of over-saving of older agents (e.g., Mirer, 1979 or Danziger et al., 1983, or the literature surveys by Modigliani, 1988 or Kotlikoff, 1988). Further, some studies (e.g., Hurd, 1987) show the bequest motive is unimportant in explaining such behavior. While we do not claim that agents behave in strict accordance with the assumptions of this paper, the intuition offered by Proposition 4.1 is suggestive that some over-saving behavior may be associated with incomplete information over wealth levels.

4.2. *Output and investments are more volatile than consumption*

In the standard stochastic optimal growth model with observable resources (e.g., Brock and Mirman, 1972), optimal consumption is typically a monotone (and sometimes linear) function of the stock level. This implies that consumption closely follows movements in contemporaneous output and investment and has similar measures of variability. For example, if consumption, output and investment are linearly related, then their coefficients of variation are equal. On the other hand, empirical studies have typically found that the time series of consumption is relatively smooth compared to investment and output (e.g., Kydland and Prescott, 1982, Table 4, p. 1365).

Under Proposition 4.1, the limiting variance of the consumption process is zero. At the same time, the output process converges to a limiting distribution with non-trivial support and strictly positive variance. Thus, the optimal growth problem with unobservable stocks and learning yields a model where the time series behavior of consumption is much less volatile than investment and output.

4.3. *Underestimation of risk aversion*

Assume that the agent has a constant relative risk aversion (CRRA) utility with coefficient, Γ ; i.e., $u(c) = c^{1-\Gamma}/(1-\Gamma)$ for $\Gamma > 0$ but different from one, and $u(c) = \ln(c)$ for $\Gamma = 1$. To maintain (R.1), we assume that if $0 < \Gamma < 1$, utility is CRRA except for some arbitrarily small neighborhood of 0 in which the utility function satisfies (R.1). This assumption is unnecessary if $\Gamma \geq 1$ and otherwise

imposes little loss of generality since consumption is typically uniformly bounded away from zero. For simplicity, suppose the production function, $f(x)$, is some strictly concave deterministic function of investment.

Now, suppose the agent is solving a problem with incomplete observation of wealth levels, but that the econometrician assumes the agent is solving a standard optimal growth problem and that realized asset values enter the econometrician's information set. The latter may occur when there are lags between agents' decisions and the revelation of asset values. The Euler condition for the econometrician is:

$$u'(c_t) = \delta f'(x_t) u'(c_{t+1}) \text{ or } \Gamma = \ln(\delta f'(x_t)) / \ln(c_{t+1}/c_t), \quad (4.2)$$

while the agent's Euler equation is given in terms of the minmax consumption and investment processes generated according to Proposition 3.3.

Consider the situation where consumption is increasing over time, as occurs in most of the post-war U.S. aggregate consumption data. Under increasing consumption, $\ln(c_{t+1}/c_t) > 0$. Since $f(\cdot)$ is strictly concave and the true investment is larger than the minmax investment, an econometrician who estimates Γ will calculate a smaller value of $f'(x_t)$ than the agent. From (4.2) this will cause the econometrician's estimate of Γ to be smaller than the agent's true coefficient of risk aversion.

This suggests that unobservable wealth may offer some insight into the equity premium puzzle of Mehra and Prescott (1985). They show that a standard optimal growth model with CRRA utility model cannot explain the high return on equity relative to the return on risk free assets in post-war U.S. data under reasonable values of the model parameters. However, they mention that 'with large Γ virtually any pair of average equity and risk-free returns can be obtained' (p. 154).

5. Information signals

5.1. The information or signals process

In this section the model is expanded to allow for learning on the part of the agent. We now suppose that after the agent chooses a date t consumption level c_t , the agent receives a signal, s_t , that provides information about y_t , the stock of resources at the beginning of date t . For example, at the end of the period the agent may receive a noisy measurement of the resource stock, $s_t = y_t + \eta_t$, where η_t is an independent noise term with mean zero. We assume that the set of possible realizations of the signal lies in some subset of a finite dimensional Euclidean space, denoted by S . The conditional probability of the signal given y and c is $P(ds|y,c)$. This formulation allows there to be no signals in which case $P(ds|y,c)$ is independent of (y,c) . The objective of the agent is the same as described in Section 2 above; however, when making the date t consumption

decision the agent now has information on the history of all signals up to that date, $\{s_0, \dots, s_{t-1}\}$.

There are now two steps required to obtain the agent's beliefs about the date $t + 1$ resource stock, given the agent's date t beliefs, actions and signals. First, the agent revises his/her beliefs about the date t stock level, y_t , conditional upon the observed signal s_t . This step is the standard Bayesian updating. Denote the posterior distribution on y_t given the signal s_t , the action c_t , and prior distribution $\mu_t \in P(\mathfrak{R}_+)$ by

$$\mu'_t = B(s_t, c_t, \mu_t) \tag{5.1}$$

where $B: S \times \mathfrak{R}_+ \times P(\mathfrak{R}_+) \rightarrow P(\mathfrak{R}_+)$ is the Bayes' Rule updating formula. If the conditional probability distribution $P(ds_t | y_t, c_t)$ admits a density function $p(s_t | y_t, c_t)$ then Bayes' Rule is such that for each (Borel measurable) subset A of \mathfrak{R}_+ ,

$$\mu'_t(A) = B(s_t, c_t, \mu_t)(A) = \frac{\int_A p(s_t | y_t, c_t) \mu_t(dy_t)}{\int_{\mathfrak{R}_+} p(s_t | y_t, c_t) \mu_t(dy_t)}. \tag{5.2}$$

The second step of the updating process involves the mapping F defined in (2.1) where the consumption level, c_t , and the production function, $f(x, r)$, are used to transform μ'_t into updated beliefs about the period $t + 1$ resource stock. Together, these two steps define a mapping from the date t beliefs, consumption and information signals into beliefs about the resource stock at date $t + 1$ given by

$$\mu_{t+1} = F(c_t, B(s_t, c_t, \mu_t)). \tag{5.3}$$

We now illustrate the mappings F and B in an explicit example.

Example 5.1. (The mappings F and B). Let $N(m, v)$ denote the normal distribution with mean m and variance v . Suppose the agent has a prior over the date t stock, y_t , which is $N(\bar{y}_t, \sigma_t^2)$ and that the production function is linear with $f(x, r) = kx + r$ where k is a constant and the shock term, r , is $N(0, \sigma_r^2)$. (We ignore the fact that this means that stock levels may be negative; this example is to illustrate the updating only.) Let the information signal be given by $s_t = y_t + \eta_t$ where η_t is $N(0, 1)$. When the agent chooses action c_t and the signal s_t is observed the posterior distribution on y_t is given by $\mu'_t = B(s_t, c_t, \mu_t) = N(\bar{y}'_t, \sigma_t'^2)$, where

$$\bar{y}'_t = (\bar{y}_t + \sigma_t^2 s_t) / (1 + \sigma_t^2) \text{ and } \sigma_t'^2 = \sigma_t^2 / (1 + \sigma_t^2). \tag{5.4}$$

The posterior distribution of y_{t+1} is $\mu_{t+1} = F(c_t, \mu'_t) = N(\bar{y}_{t+1}, \sigma_{t+1}^2)$, where

$$\bar{y}_{t+1} = k(\bar{y}'_t - c_t) \text{ and } \sigma_{t+1}^2 = k^2 \sigma_t'^2 + \sigma_r^2 \tag{5.5}$$

with \bar{y}'_t and σ_t' defined as in (5.4).

5.2. Model where signals do not change the support of beliefs

We now assume that the information signal is such that the supports of the prior and posterior beliefs are the same:

S.1. For all $(s, c, \mu) \in S \times \mathfrak{R}_+ \times P(\mathfrak{R}_+)$, $\text{Supp } \mu = \text{Supp } B(s, c, \mu)$.

If there are no signals or, equivalently, if the signals are uninformative, then (S.1) holds trivially. Information signals can be informative and still satisfy (S.1). This is illustrated in the examples below where not only does (S.1) hold, but, if $y_t = \bar{y}$ at each date t (so that we have a standard statistical inference problem), then in the limit the signals lead to complete learning of \bar{y} .

First we have the following:

Lemma 5.2. Suppose $P(ds|y, c)$, the conditional distribution of the signals, admits a strictly positive Lebesgue density function $p(s|y, c)$ on the real line (or the non-negative half-line) so that for any set A in S , $P(s \in A|y, c) = \int_A p(s|y, c) ds$ with $p(s|y, c) > 0$ for all (s, c, y) . Then (S.1) above holds.

Proof of Lemma 5.2. Under the hypotheses of the Lemma, Bayes' Rule defines a posterior distribution over stocks conditional on the information signal such that for any set A in \mathfrak{R}_+ ,

$$B(s, c, \mu)(A) = \frac{\int_A p(s|y, c) \mu(dy)}{\int_{\mathfrak{R}_+} p(s|y, c) \mu(dy)}.$$

Since by assumption $p(s|y, c) > 0$ for all s and y , for any set A , $\mu(A) > 0$ implies $B(s, \mu)(A) > 0$ (and vice versa). Hence, $\text{Supp } \mu = \text{Supp } B(s, \mu)(A)$ for all s so (S.1) holds. \square

Example 5.3. (Signals with additive error). Suppose the signals are given by the relation $s_t = y_t + \eta_t$ where η_t is independent of y_t and has a density function $\psi(\eta)$ with $\psi(\eta) > 0$ for all η in $(-\infty, \infty)$. For example, η could be a normally distributed random variable. Then $P(ds|y, c)$ admits a density function $p(s|y, c) = \psi(s - y)$ which is positive for all s and y (and is independent of c). Hence, from Lemma 5.2, (S.1) holds.

Example 5.4. (Signals with multiplicative error) Suppose the signals are given by $s_t = y_t \cdot \eta_t$ where η_t is independent of y_t and has a density function $\psi(\eta)$ with $\psi(\eta) > 0$ for all η in $(0, \infty)$. For example, η could be a log-normally distributed random variable. Then $P(ds|y, c)$ admits a density function $p(s|y, c) = \psi(s/y)$ which is positive for all s and y . Again, from Lemma 5.2, (S.1) holds.

5.3. Under (S.1) the optimality of minmax behavior and its implications continue to hold

In this sub-section we note that all of the results of Sections 3 and 4 continue to hold even when we allow the agent to receive information signals at each date, so long as condition (S.1) holds. First consider the optimality of minmax behavior discussed in Section 3. For Problem P1 with information signals obeying (S.1) the precise arguments made in Proposition 3.1 imply that any optimal policy is attainable. Similarly, Lemma 3.2 continues to hold under (S.1). The arguments in Proposition 3.3 then imply that the solutions to Problems P1 and P2 are identical.³ Thus, the over-saving result of Section 4.1 carries over to the model with information signals that obey condition (S.1). In addition, consumption is more volatile than investment and output as in Section 4.2, and there is the possibility of underestimating the coefficient of risk aversion as in Section 4.3.

5.4. Under (S.1) the value of information is zero

We argued above that any solution to the optimal growth problem with unobservable resources and information signals obeying (S.1) is also a solution to the minmax problem P2 of Section 3. This result has the implication that the agent is not willing to pay any amount of money to receive additional observations of the signal. Under (S.1) the information signal does not change the behavior of the agent since it does not change the support of the prior distribution. Information is informative in a statistical sense but it is not economically valuable to the agent even though it may significantly reduce the posterior variance in the beliefs of the agent. To see this suppose the agent is offered the chance to observe a signal which is a noisy estimate of the stock level, $s_t = y_t + \eta_t$, where η_t is normally distributed with mean zero and variance of σ_η^2 . No matter how small the variance of η the agent does not gain by observing the signal.⁴

This result can be compared with Blackwell's Theorem (Blackwell, 1951) which states that a signal is statistically informative if, and only if, it is economically valuable. Blackwell's Theorem is not violated in our model for the following reasons. First, it is typically stated in terms of weak inequalities, i.e., a statistically

³ In Problem P1 the agent may observe information signals and in principle can use these signals as a conditioning device when the agent is indifferent between two consumption decisions (such indifference is possible in our model since we do not require strict concavity of the utility function in Section 3). To maintain the equivalence, an agent solving Problem P2 may similarly use a random number generator to obtain the same distribution of consumptions and utilities. For both Problems P1 and P2 such randomizations are sub-optimal when the utility function is strictly concave.

⁴ The agent will pay positive resources to observe a signal that is an observation of the true stock. Thus, in the example there is a discontinuity of the value of information in the variance parameter at $\sigma_\eta^2 = 0$.

informative signal cannot make the agent worse off. Second, it normally requires finitely many actions and signals and uniformly bounded utility functions which we do not have here.

6. Model with general informative signals

6.1. Endogenous shocks and the resource discovery process

We now examine the model with information signals without the restriction (S.1) that the signals have no effect on the support of the prior distribution. The main result is that the solution to the unobservable stocks Problem P1 is the same as the ‘minmax’ or ‘cautious’ optimal growth problem P2 with one modification; at the end of each period t in the ‘minmax’ optimal growth problem there is a *discovery* of resources, z_t . The stochastic process defining the discovery of resources at the end of each period is precisely the increase in the minimum of the support of prior beliefs over stocks that results from the information signal. In essence, the agent has ‘discovered’ that the certain resource stock is larger than previously believed. This discovery process has a natural interpretation in models of natural resource economics where new information increases the size of the known resource stock (either mineral or oil reserves or the population of a valuable fish or wildlife species). It is also relevant to situations where agents reassess the value of assets based on outside information.

We begin with some preliminary definitions. Recall that the production function is of the form $f(x, r)$ where x is the input level and r is the realization of an independently and identically distributed process $\{r_t\}_{t=1}^{\infty}$. To analyze the problem with resource discovery it is useful to consider production functions with a more general endogenous (real-valued) shock process $\{z_t\}_{t=1}^{\infty}$. The shock process is said to be *exogenous* if it is independent of the stock and consumption processes. The shock process is said to be *endogenous* if the date t shock, z_{t+1} , has a distribution that is a function of the history of consumptions, outputs and shocks at date t , $\{(c_i, y_i)\}_{i=0}^t, \{z_i\}_{i=1}^t\}$, but is independent of the future values $\{(c_i, y_i)\}_{i=t+1}^{\infty}, \{z_i\}_{i=t+2}^{\infty}$.

We now define a general optimal growth problem with endogenous shocks. Let $g: \mathfrak{R}_+ \times \mathfrak{R}_+ \rightarrow \mathfrak{R}_+$ denote a production function where $g(x, z)$ gives the output produced from the input x when z is the realization of an endogenous shock process at some given date. Fix an initial stock $y > 0$.

Problem P3. Optimal growth problem with endogenous shocks to production: $\max_{\{c_t\}} E\{\sum_{t=0}^{\infty} \delta^t R(c_t, y_t)\}$ s.t. $0 \leq c_t \leq y_t$ for all $t \geq 0$ where $y_0 = y$, $y_{t+1} = g(y_t - c_t, z_{t+1})$ for all $t \geq 0$ and the consumption level c_t is constrained to be a (Borel measurable) function of the date t partial history $\{(c_i, y_i)\}_{i=0}^{t-1}, \{z_i\}_{i=1}^t, y_t\}$ and where the expectations are those with respect to the *known* distribution of the endogenous process.

Problem P3 is a generalization of the optimal growth model with exogenous

shocks (e.g., Brock and Mirman, 1972) that allows the shocks to technology to depend upon past actions of the agent. In Problem P3 the beginning of period stocks are accurately observed and the agent knows the vector $\{\{c_t, y_t\}_{t=0}^{\infty}, \{z_t\}_{t=1}^{\infty}, y_t\}$ before the consumption at date t is chosen. The shock z_{t+1} is observed after c_t is chosen; however, the agent knows the probability distribution of the shock process conditional on the history of observations.

We now return to Problem P1, the model with unobservable resource stocks. Fix an initial prior, μ_0 , and let π be any policy. This generates a sequence of consumptions, signals and posterior beliefs $\{c_t, s_t, \mu_t\}_{t=0}^{\infty}$. At the beginning of date t the agent has beliefs μ_t about the beginning of date t resource level, y_t . At the end of date t , after observation of the signal s_t , the agent has beliefs $B(s_t, c_t, \mu_t)$ about the beginning of date t resource level, y_t . Define the date t discovery, z_{t+1} , to be the increase in the lower bound of the support of the posterior distribution at each date t associated with the signal observed at that date, i.e.,

$$z_{t+1} \equiv [\min\{y \in \text{Supp } B(s_t, c_t, \mu_t)\}] - [\min\{y \in \text{Supp } \mu_t\}], \tag{6.1}$$

where $B(s_t, c_t, \mu_t)$ is the posterior distribution on the resource stock (see (5.2)). Under (S.1), $z_t = 0$ for all t . Below we provide examples that illustrate possible non-trivial discovery processes generated by (6.1).

Example 6.1 (Signals with additive error). Let the signal received at date t be given by

$$s_t = y_t + \eta_t \text{ with } \text{Supp } \eta_t = [0, 1],$$

where $\{\eta_t\}_{t=0}^{\infty}$ is an i.i.d. process independent of the stock levels. By assumption, the noise in the signal, η_t , lies in $[0, 1]$ so observation of s_t indicates to the agent that $s_t - 1 \leq y_t \leq s_t$. Denote the lower bound of the support of μ_t by \underline{y}_t . After observing the signal s_t , the lower bound of the support of the updated beliefs about y_t is $\max\{\underline{y}_t, s_t - 1\}$ and the resource discovery is given by $z_{t+1} = [\max\{\underline{y}_t, s_t - 1\}] - \underline{y}_t$.

Example 6.2 (Signals with multiplicative error). Let the signal received at date t be given by

$$s_t = y_t \cdot \eta_t \text{ with } \text{supp } \eta_t = [0, 1],$$

where $\{\eta_t\}_{t=0}^{\infty}$ is an i.i.d process independent of the stock levels. By assumption, the noise in the signal, η_t , lies in $[0, 1]$ so observation of s_t indicates to the agent that $0 \leq s_t/y_t \leq 1$ or $y_t \geq s_t$. Denote the lower bound of the support of μ_t by \underline{y}_t . After observing the signal s_t , the lower bound of the support of the beliefs on y_t is $\max\{\underline{y}_t, s_t\}$ and the discovery of resource stocks is given by $z_t = [\max\{\underline{y}_t, s_t\}] - \underline{y}_t$.

6.2. Optimality of minmax behavior with resource discovery

We now show that for the unobservable stocks problem P1 the agent chooses a consumption process that is also optimal for Problem P3 with a discovery process

generated by the initial prior μ_0 . In particular, for the unobservable stocks problem an optimizing agent assumes that the initial resource level is the lower bound of the support of the prior distribution, μ_0 , that production corresponds to the lowest output at each date, and that resource stocks are augmented by discoveries as defined in (6.1). Recall that $\underline{f}(x) = \min_r f(x, r)$.

Proposition 6.1. Assume (R.1–R.3), (f.1–f.5), and (B.1). Fix an initial prior μ_0 and let π and $\{c_t\}_{t=0}^\infty$ be an optimal policy and associated optimal consumption process for Problem P1 with unobservable stocks and information signals. Let $\{z_t\}_{t=1}^\infty$ be the associated discovery process. Then $\{c_t\}_{t=0}^\infty$ is also optimal for problem P3 with observable stocks with initial stock $y_0 = \min\{y \in \text{Supp } \mu_0\}$ and production function $g(x, z) = \underline{f}(x + z)$ where $\underline{f} = \min_r f(x, r)$.

The proof of Proposition 6.1 follows similar arguments to the proof of Proposition 3.3, and is omitted. Note that Proposition 3.3 is a special case of Proposition 6.1, specializing to the situation where the discovery process is $z_t = 0$ for all t . It should also be clear that, with obvious modifications, the implications of Proposition 6.1 hold in a setting in which the agent receives information about the outcome of the random shock r_t before making the consumption decision c_t .

7. Agents either minmax or they fail to survive with positive probability

In this section we study the model where the assumption $U(0) = -\infty$ does not necessarily hold. We show that an optimal policy is one of two types. The optimal policy is either consistent with minmax behavior as in Problem P2, or the agent chooses actions that lead to extinction of the resource stock and bankruptcy of the agent with strictly positive probability. Hence, agents most likely to be ‘alive’ (i.e., with positive resources) in the distant future will be either those with utility functions obeying (R.1) or those who exhibit the same type of minmax behavior.

Proposition 7.1. Assume (R.2), (R.3), (f.1–f.5), (B.1) and (S.1). Let $\{c_t\}_{t=0}^\infty$ be any optimal consumption process generated by an optimal policy π . Then, either $\{c_t\}_{t=0}^\infty$ solves the minmax problem (P2) or it results in the stock becoming extinct with strictly positive P_π -probability in finite time.

Proof of Proposition 7.1. Fix any optimal policy, π . First, suppose that π satisfies the feasibility constraints of the associated minmax problem P2. The associated consumption process must solve Problem P2 (otherwise it could not be optimal for Problem P1, whose objective function is the same as that of Problem P2). Now, suppose that π does not obey the constraints of the associated minmax Problem P2. We argue that the stock eventually becomes extinct with strictly positive probability. To see this suppose the support of the prior, μ_0 , is the set $[\underline{y}, \bar{y}]$. If the

optimal policy violates the constraints of Problem P2 in the initial period ($t = 0$), this means that the optimal consumption c_0 lies in the set $[\underline{y}, \bar{y}]$. This implies extinction of the resource if the initial stock level lies in the set $[\underline{y}, c_0]$, an event which the agent's prior assigns strictly positive probability. More generally, one may show that under the optimal policy, if the constraints of the associated minmax Problem P2 are violated at some date with positive probability, then there will be extinction of the resource at some date with strictly positive probability. \square

It is a straightforward exercise to construct examples that illustrate the possibility of extinction of the resource when (R.1) fails to hold.

8. The Hall model of consumption

Thus far, the discussion has been in the context of a standard optimal growth model. Under somewhat stronger conditions the optimality of minmax behavior holds for the model of consumption and saving studied by Hall (1978) and many others. In that model, an agent has an initial wealth of Y_0 and receives labor income of w_t at the beginning of each date t . Wealth saved at the end of date t earns a gross interest rate or appreciates by a factor of ρ_t . Both $\{\rho_t\}_{t=1}^\infty$ and $\{w_t\}_{t=1}^\infty$ are exogenous i.i.d. stochastic processes, with $\rho_t \geq 1$. Let Y_t denote the wealth or asset holdings of the agent at the beginning of date t , after labor income has been received. Let c_t represent consumption at date t . If the wealth level at each date is observable so that there is no learning, then the agent solves the problem

$$\max_{\{c_t\}} E \sum_{t=0}^\infty \delta^t u(c_t), \text{ s.t. } Y_{t+1} = \rho_t(Y_t - c_t) + w_{t+1}, Y_0 \text{ given.} \quad (8.1)$$

Now, suppose that the asset levels, Y_t , are unobservable. The interest rate, ρ_t , and the labor income, w_{t+1} , are either observable or else have known marginal distributions. Let μ_0 denote the initial prior on Y_0 , and let μ_t denote the date t beliefs about Y_t . If the agent plans consumption of $c_t \geq 0$ the 'attained consumption' is $\min\{c_t, Y_t\}$. In each period the agent saves $x_t = Y_t - \min\{c_t, Y_t\}$ and the next period asset level is $y_{t+1} = \rho_t x_t + w_{t+1}$, where $x_t = Y_t - \min\{c_t, Y_t\}$. The agent's optimization problem is

$$\max_{\{c_t\}} E \sum_{t=0}^\infty \delta^t u(\min\{c_t, Y_t\}), \text{ s.t. } Y_{t+1} = \rho_t x_t + w_{t+1}, Y_0 \text{ given.} \quad (8.2)$$

If the production function is defined to be $f(x_t, r_{t+1}) = \rho_t x_t + w_{t+1}$, where $r_{t+1} = (\rho_t, w_{t+1})$, then this model is very similar to that discussed in previous sections of the paper. The shock term, r_{t+1} , is now a vector, with possibly one or both coordinates observable. Assumptions (f.1), (f.3) and (f.4) still hold. Assumptions (f.2) and (f.5) require modification. (f.5) was used to ensure that the value function is finite. To achieve the same effect we may replace (f.5) with (f.5') $\delta \rho_t$ is uniformly less than one with probability one.

Assumption (f.2) required that zero saving leads to zero output and consumption in the next period. (R.1) then implied that the agent would not choose to consume the entire stock in any given period. This reasoning does not carry over to the consumption/saving model if labor income is bounded away from zero since $f(0, r_t) = w_t > 0$ and (f.2) fails. It is therefore necessary replace assumptions (f.2) and (R.1) with a stronger joint restriction on the distribution of labor income and the utility function. Let $P(dw)$ be the distribution of labor income. Assume:

$$(R.1') \int u(w)P(dw) = -\infty.$$

Assumption (R.1') holds if w has positive mass at $w = 0$ and $u(0) = -\infty$; however, it is important to note that this is not required for (R.1') to hold. For example, suppose $u(c) = -1/c$ and w is uniformly distributed on the interval $[0, w']$ where w' is any arbitrary positive number. Then $\int_0^{w'} u(w)P(dw) = \int_0^{w'} (-1/w)(1/w')dw = (1/w')[-\ln(w)]_0^{w'} = -\infty$.

The 'minmax' problem is defined as follows. First, the agent supposes that the initial wealth level is the minimum of the support of beliefs over wealth, $Y_0 = \min\{Y \in \text{Supp } \mu_0\}$. Define $\underline{\rho}_t = \rho_t$ if ρ_t is observed and define $\underline{\rho}_t$ to be the infimum of ρ_t otherwise. Define \underline{w}_t analogously. With these definitions, the 'minmax' optimization problem with observable asset levels is:

Problem P2'. $\max_{\{c_t\}} \sum_{t=0}^{\infty} \delta^t u(c_t)$ s.t. $0 \leq c_t \leq y_t$ and $y_{t+1} = \underline{\rho}_t(y_t - c_t) + \underline{w}_{t+1}$ for all $t \geq 0$ with $y_0 = Y_0$.

When (f.5') and (R.1') replace (f.2), (f.5) and (R.1) and condition (S.1) is maintained, it is straight-forward to extend the arguments of previous sections to show that any solution to (8.2) is a solution to Problem P2' and vice versa.

In this version of the consumption model an agent with unobservable wealth exhibits 'over-saving' in the same sense as in Section 4.2, however, the amount of 'over-saving' may become arbitrarily large in the limit even when the initial ignorance about the asset level is small. To see this, let Y_t denote the true date t asset level and let y_t denote the 'minmax' asset level generated via Problem P2'. Denote the difference between the actual and 'minmax' asset level by $\Delta_t = Y_t - y_t$. Since $\rho_t \geq \underline{\rho}_t$, the definitions of Y_t and y_t imply $\Delta_{t+1} \geq \rho_t \Delta_t$, which in turn implies $\Delta_t \geq [\prod_{i=0}^{t-1} \rho_i] \Delta_0$. If $\rho_t > 1$ for all t and $\Delta_0 > 0$, then the difference between the actual and 'minmax' stocks tends to infinity, i.e., $\lim_{t \rightarrow \infty} \Delta_t = \infty$.

9. Concluding remarks

The results of this paper show that when there is infinite disutility from zero consumption, the optimal consumption policy for the problem with unobservable resources and learning is identical to the optimal policy under an associated deterministic problem in which the agent exhibits a form of minmax behavior. When the agent receives information signals that change the support of beliefs, the

optimal policy is equivalent to a stochastic growth model with an endogenous resource discovery process. The intuition offered by the results is intriguing and it would be of some interest to develop an empirical analysis of these issues. Further investigation of agents' behavior when (R.1) is relaxed also seems warranted.

Acknowledgements

We especially thank Professors Jess Benhabib, James P. Dow, Jr., Mark Gertler, Boyan Jovanovic, Robert Moffit, Simon Potter and an anonymous referee for their comments. We also appreciate remarks offered by participants at a number of seminars. The assistance of the C.V. Starr Center of Applied Economics at New York University to the first author is gratefully acknowledged. The second author was at the University of California, Riverside when this research was completed.

Appendix A. Proof of proposition 4.1

From Proposition 3.3, optimal consumptions for Problem P1 are also optimal for Problem P2. Using well known results from the optimal growth literature, it follows that $\lim_{t \rightarrow \infty} c_t = c_\delta$. By (3.2), if y_t is the date t stock level associated with the minmax Problem P2 then $y_t = \min\{y \in \text{supp } \mu_t\}$. This implies that $\lim_{t \rightarrow \infty} \min\{y \in \text{Supp } \mu_t\} = y_\delta$. We now show that with probability one the actual stock process enters the set $[y^*, y^{**}]$ and never leaves. First we need some definitions for which we refer to Fig. 3 for clarification. Since the map $y_{t+1} = f(y_t - c_\delta)$ has two fixed points, for any $\epsilon > 0$ sufficiently small the map $y_{t+1} = f(y_t - c_\delta - \epsilon)$ also has two fixed points. Recall that $\tilde{f}(x) = \max_r f(x, r)$. Let y' and y'' be the smaller fixed points of the maps $y_{t+1} = f(y_t - c_\delta - \epsilon)$ and $y_{t+1} = \tilde{f}(y_t - c_\delta + \epsilon)$, respectively, and let y_ϵ^* , y_ϵ^{**} be the larger fixed points of those two maps. Observe that $\tilde{f}(y_\delta - c_\delta + \epsilon) > f(y_\delta - c_\delta) = y_\delta$ so there exists a point, \tilde{y} , such that $\tilde{y} < y_\delta$ and $\tilde{f}(\tilde{y} - c_\delta + \epsilon) > y_\delta > \tilde{y}$.

The lower bound of the support of the posterior beliefs, y_t , converges to the golden rule stock, y_δ . By assumption, the true initial stock is strictly greater than the lower bound of the support of initial beliefs. Hence, the actual stock level at each date exceeds the lower bound of the support of beliefs at that date so that in the limit the actual stock level is no less than the golden rule stock. (Under the conditions of Remark 3 following the statement of Proposition 4.1 this will be true from date 2 onward.) In particular, from some finite date onwards the actual stock level always exceeds \tilde{y} . Further, c_t converges to c_δ . Without loss of generality we may suppose that for all t ,

$$y_t > \tilde{y} \text{ and } |c_t - c_\delta| < \epsilon. \tag{4.1.1}$$




Fig. 3.

Now, suppose that the shocks to the production function are finite and ordered so that there is a ‘best’ shock which results in the highest output from each input level. In Lemma 4.1.1 below we show that there exists a finite number N such that if the actual stock level exceeds \tilde{y} , and the best state occurs at least N consecutive times in a row, then the actual stock level enters the set (y', ∞) and never leaves it. In Lemma 4.1.2 we show that on almost every sample path one can find a time when the best state does occur for N consecutive periods. Combining Lemmas 4.1.1 and 4.1.2 implies that if the initial stock exceeds \tilde{y} , then with probability one the true stock enters the set (y', ∞) and never leaves. From Eq. (4.1.1) we may suppose the initial stock exceeds \tilde{y} . In the remark following Lemma 4.1.2, we indicate that this conclusion does not depend on the assumption that the shocks are finite and ordered.

In Lemma 4.1.3 we show that once the actual stock enters the set (y', ∞) it will enter the set $[y_\epsilon^*, y_\epsilon^{**}]$ in finite time and stay there forever. Since $\epsilon > 0$ is

arbitrary, this proves that the actual stock enters the set $[y^*, y^{**}]$ and stays there forever. Lemmas 4.1.1–4.1.3 therefore complete the proof of Proposition 4.1.

Lemma 4.1.1. Suppose there exists a ‘best’ shock r' such that for all $x \geq 0$, $f(x, r') \geq f(x, r)$ for all r , and Eq. (4.1.1) holds. Then there exists an $N < \infty$ such that if from date τ , the best state occurs for N consecutive periods, then $y_{\tau+N} > y'$.

Proof of Lemma 4.1.1. By definition of \bar{y} , $\bar{f}(\bar{y} - c_\delta + \epsilon) > \bar{y}$, so that $\bar{y} > y'$. With the aid of Fig. 3 it is easy to see that iterates of the map $y_{i+1} = \bar{f}(y_i - c_\delta + \epsilon)$ from \bar{y} converge to y_ϵ^{**} , the larger fixed point of that mapping. In particular, there exists an integer $N < \infty$ such that the N -th iterate of the mapping $y_{i+1} = \bar{f}(y_i - c_\delta + \epsilon)$ from $y_0 = \bar{y}$ exceeds y' . The lemma follows immediately. \square

Lemma 4.1.2. Suppose $\{r_i\}_{i=1}^\infty$ is an i.i.d. process with $\text{Prob}(\{r_1 = r'\}) > 0$ for some realization r' . For any $N < \infty$, if we define the set $B = \{\text{all sample paths such that there exists a } \tau < \infty \text{ with } r_{\tau+i} = r' \text{ for } i = 0, 1, \dots, N-1\}$, then $\text{Prob}(B) = 1$.

Proof of Lemma 4.1.2. For any $J = 0, 1, \dots$, define the set $B_J = \{r_{JN+i} = r' \text{ for } i = 1, 2, \dots, N\}$. Let $q = \text{Prob}(\{r_i = r'\}) > 0$. Then for all J , $\text{Prob}(B_J) = q^N > 0$, so using the independence of the sets B_J ,

$$\begin{aligned} \text{Prob}\left(\bigcup_{J=0}^\infty B_J\right) &= 1 - \text{Prob}\left(\bigcap_{J=0}^\infty B_J^c\right) = 1 - \lim_{K \rightarrow \infty} \text{Prob}\left(\bigcap_{J=0}^K B_J^c\right) \\ &= 1 - \lim_{K \rightarrow \infty} (1 - q^N)^K = 1, \end{aligned}$$

where B_J^c denotes the complement of B_J . The lemma follows from the fact that the set $\bigcup_{J=0}^\infty B_J$ is contained in the set B . \square

Remark. Lemmas 4.1.1 and 4.1.2 are needed to assert that from a stock level in excess of \bar{y} , on almost every sample path the stock level eventually enters the set (y', ∞) and stays there forever. Lemma 4.1.1 uses the fact that there is a ‘best’ shock that occurs with positive probability. This assumption is without loss of generality since the proof of Lemma 4.1.1 can be modified to show that there is an $N < \infty$ and a set D of N -tuples of realizations of the shock process with $\text{Prob}(D) > 0$ such that if the initial stock is \bar{y} , and the N -tuple of shocks (r_1, r_2, \dots, r_N) lies in the set D then y_N enters the set (y', ∞) and stays there forever. Lemma 4.1.2 can then be modified to show that on almost every sample path there exist N consecutive dates such that the N -tuple of shocks at those dates lies in the set D .

Lemma 4.1.3. If $y_0 > y'$ and Eq. (4.1.1) holds, then along almost every sample path the actual stock enters the set $[y_\epsilon^*, y_\epsilon^{**}]$ in finite time and stays there forever.

Proof of Lemma 4.1.3. From an initial stock $y_0 > y'$, it is easy to see (refer to Fig. 3) that iterates of the map $y_{t+1} = \tilde{f}(y_t - c_\delta + \epsilon)$ converge to the point y_ϵ^{**} , while iterates of map $y_{t+1} = \underline{f}(y_t - c_\delta - \epsilon)$ converge to the point y_ϵ^* . Note that the evolution of y_t is governed by the map $y_{t+1} = f(y_t - c_t, r_{t+1})$ and the definitions of \tilde{f} and \underline{f} and the assumption that c_t lies in $(c_\delta - \epsilon, c_\delta + \epsilon)$ for all t imply $\underline{f}(y_t - c_\delta - \epsilon) \leq f(y_t - c_t, r_{t+1}) \leq \tilde{f}(y_t - c_\delta + \epsilon)$; hence, y_t enters the set $[y_\epsilon^*, y_\epsilon^{**}]$ in finite time. It is clear that once the y_t process enters that set it stays there forever. \square

References

- Arrow, K.J., 1971, *Essays in the theory of risk-bearing* (North-Holland, Amsterdam).
- Baumol, W.J., 1986, Unnatural value: or art investment as a floating crap game, *American Economic Review Proceedings* 76, 10–14.
- Bhattacharya, G., 1982, The importance of learning in planning with exhaustible resources, in: W. Eichhorn et al., eds., *Economic theory of natural resources* (Physica-Verlag, Wurzburg) 13–24.
- Blackwell, D., 1951, The comparison of experiments, in: *Proceedings of the second Berkeley symposium on mathematical statistics and probability* (Berkeley, University of California Press) 93–102.
- Blackwell, D., 1965, Discounted dynamic programming, *Annals of Mathematical Statistics* 36, 226–235.
- Brock, W. and L.J. Mirman, 1972, Optimal growth and uncertainty: the discounted case, *Journal of Economic Theory* 4, 479–513.
- Clark, C. and G. Kirkwood, 1986, On uncertain renewable resource stock: optimal harvest policies and the value of stock surveys, *Journal of Environmental Economics and Management* 13, 235–244.
- Danziger, S., J. van der Gaag, E. Smolensky and M. Taussig, 1983, The life cycle hypothesis and the consumption behavior of the elderly, *Journal of Post-Keynesian Economics* 5, 208–227.
- Easley, D. and N. Kiefer, 1988, Controlling a stochastic process with unknown parameters, *Econometrica* 56, 1045–1064.
- El-Gamal, M.A. and R.K. Sundaram, 1993, Bayesian economists...Bayesian agents I: an alternative approach to optimal learning, *Journal of Economic Dynamics and Control* 173, 355–383.
- Freixas, X., 1981, Optimal growth with experimentation, *Journal of Economic Theory* 24, 296–309.
- Grossman, S.J., R.E. Kihlstrom and L.J. Mirman, 1977, A Bayesian approach to the production of information and learning by doing, *Review of Economic Studies* 44, 533–547.
- Hall, R.E., 1978, Stochastic implications of the life cycle-permanent income hypothesis: theory and evidence, *Journal of Political Economy* 866, 971–988.
- Huffman, G. and N.M. Kiefer, 1990, Optimal learning in a one-sector growth model, working paper (University of California, Riverside).
- Hurd, M., 1987, Savings of the elderly and desired bequests, *American Economic Review* 773, 298–312.
- Kemp, M.C., 1976, How to eat a cake of unknown size, in: *Three topics in the theory of international trade: distribution, welfare, and uncertainty* (North-Holland, Amsterdam) 297–308.
- Kiefer, N.M. and Y. Nyarko, 1989, Optimal control of an unknown linear process with learning, *International Economic Review* 30, 571–586.
- Kotlikoff, L.J., 1988, Intergenerational transfers and savings, *Journal of Economic Perspectives* 2, 41–58.
- Kydland, F.E. and E.C. Prescott, 1982, Time to build and aggregate fluctuations, *Econometrica* 506, 1345–1370.

- Majumdar, M., 1982, A note on learning and optimal decisions with a partially observable state space, in: L.J. Mirman and D. F. Spulber, eds., *Essays in the economics of renewable resources* (North-Holland, Amsterdam) 141–153.
- Majumdar, M., T. Mitra and Y. Nyarko, 1989, Dynamic optimization under uncertainty: non-convex feasible set, in: G. Feiwel, ed., *Joan Robinson and modern economic theory* (MacMillan Press, New York).
- McLennan, A., 1984, Price dispersion and incomplete learning in the long run, *Journal of Economic Dynamics and Control* 7, 331–347.
- Mehra, R. and E.C. Prescott, 1985, The equity premium: a puzzle, *Journal of Monetary Economics* 15, 145–161.
- Mirer, T.W., 1979, Wealth-age relation among the aged, *American Economic Review* 69, 435–443.
- Modigliani, F., 1988, The role of intergenerational transfers and life cycle saving in the accumulation of wealth, *Journal of Economic Perspectives* 2, 15–40.
- Nyarko, Y., 1991, Learning in mis-specified models and the possibility of cycles, *Journal of Economic Theory* 55, 416–427.
- Quyen, N.V., 1991, Exhaustible resources: a theory of exploration, *Review of Economic Studies* 58, 777–789.
- Robson, A.J., 1979, Sequential exploration of uncertain deposits of a depletable natural resource, *Journal of Economic Theory* 21, 88–110.
- Rubinstein, M., 1976, The strong case for the generalized logarithmic utility model as the premier model of financial markets, *Journal of Finance, Proceedings* 31, 551–571.
- Tobin, J. and W. Dolde, 1971, Wealth, liquidity and consumption, in: *Consumer spending and monetary policy: the linkage* (Federal Reserve Bank of Boston, Boston) 99–146.
- Walters, C.J., 1986, *Adaptive management of renewable resources* (Macmillan, New York).