Available online at www.sciencedirect.com
ScienceDirect
Journal of Economic Theory 144 (2009) 280–295
JOURNAL OF Economic Theory
www.elsevier.com/locate/jet

Optimal growth and uncertainty: Learning

Christos Koulovatianos [a], Leonard J. Mirman [b], Marc Santugini [c,*]

[a] *Department of Economics, University of Vienna, Austria*
[b] *Department of Economics, University of Virginia, USA*
[c] *Institute of Applied Economics and CIRPÉE, HEC Montréal, Montréal, Canada*

Received 8 June 2007; final version received 4 April 2008; accepted 10 April 2008

Available online 23 May 2008

Abstract

We introduce learning in a Brock–Mirman environment and study the effect of risk generated by the planner's econometric activity on optimal consumption and investment. Here, learning introduces two sources of risk about future payoffs: structural uncertainty and uncertainty due to the anticipation of learning. The latter renders control and learning nonseparable. We present two sets of results in a learning environment. First, conditions under which the introduction of learning increases or decreases optimal consumption are provided. The effect depends on the strengths and directions of the two sources of risk, which may pull in opposite directions. Second, the effects of the mean and riskiness of the distribution of the signal and initial beliefs on optimal consumption are studied.

JEL classification: D8; D9; E2

Keywords: Dynamic programming; Learning; Growth; Risk; Uncertainty

1. Introduction

In the early literature on optimal growth, the evolution of output was deterministic, see Cass [10] and Koopmans [30]. This was a natural place to begin the study of optimal growth since growth had already been studied in a deterministic environment by Ramsey [40]. Brock and Mirman [9] introduced uncertainty in outcomes in an optimal growth model, which built

* Corresponding author.
E-mail address: marc.santugini@hec.ca (M. Santugini).

0022-0531/$ – see front matter
doi:10.1016/j.jet.2008.04.002

on earlier studies of stochastic positive growth, see Mirman [35,36]. Uncertainty in outcomes is modeled by introducing a random shock in the production function.

There is, however, another aspect of uncertainty that has yet to be studied in optimal growth: uncertainty about the structure of the economy. Unlike uncertainty in outcomes, structural uncertainty evolves through learning. Indeed, the planner becomes an econometrician by gathering and analyzing data in order to learn about an unknown parameter, and, thus, reducing structural uncertainty. At the same time, the planner also makes consumption and investment decisions. These two functions of the planner are interrelated.

We introduce learning in a Brock–Mirman environment with structural uncertainty. Here, the planner's econometric activity adds an element of risk to the decision making process, just as introducing uncertainty in outcomes to the deterministic model adds an element of risk in Brock–Mirman. We study the effect that learning has on optimal policies and extend the literature on learning and risk to a model of economic growth.

Indeed, there is a two-way interaction between decision making and learning. On the one hand, decision making may have an effect on learning, which is referred as experimentation.[1] On the other hand, learning may have an effect on decision making, which is due to an increase in risk of future payoffs. While the literature on learning has focused on experimentation, we concentrate on the effect of learning on decision making, extending the literature on risk to a dynamic learning environment. Previous literature on both risk and growth has focused only on models in which the planner knows the distributions of stochastic variables.[2]

In order to study the effect of learning, we specialize the Brock–Mirman model to the Mirman and Zilcha [38] model, with specific utility and production functions, but general distributions of the production shock and beliefs. We provide two sets of results in this class of optimal stochastic growth models.

First, conditions under which the introduction of learning increases or decreases optimal consumption are provided. Here, learning introduces two sources of risk about future payoffs: structural uncertainty and uncertainty due to the anticipation of learning.

On the one hand, the risk generated by structural uncertainty is determined by the curvature of the mean of the production shock with respect to the unknown parameter. Specifically, as structural uncertainty is introduced, the expected marginal utility of investment decreases with a concave mean, inducing more consumption, while the expected marginal utility of investment increases with convexity, inducing less consumption. On the other hand, the risk due to the anticipation of learning always increases the expected marginal utility of investment, leading to a decrease in consumption.

The total effect of learning depends on the strengths and directions of the two sources of risks, which may pull in opposite directions. If the mean of the production shock is convex with respect to the unknown parameter, both types of risk work in the same direction and consumption decreases. If the mean of the production shock is concave, both types of risk pull in opposite directions and the effect of learning depends on the strength of each risk.

[1] Experimentation was initially studied in models in which the only link between periods is beliefs. See [1,3,12,16,17, 20,22,23,28,29,37,39,43,44]. Experimentation in a model with capital accumulation has also been studied. See [7,8,13, 14,18,21,25].

[2] See [15,24,31,41,42] for a finite-period analysis. For an infinite-horizon setup, see [34] in a model with a single agent, and [2] for the case of a game. Finally, see [26] for a detailed review of other issues studied in this literature.

Second, we perform a comparative analysis of distributions on the learning planner's optimal consumption. Specifically, the effects of the mean and riskiness of the distributions of the production shock and beliefs on optimal consumption are studied.

We show that, while a higher mean of the production shock decreases consumption, a riskier distribution of the production shock has no effect on consumption. The first result is due to the structure of the Mirman–Zilcha model, i.e., a higher mean of the production shock makes investment more profitable. The second result follows from the fact that the uncertainty in outcomes due to the random production shock is determined solely through its mean in a Mirman–Zilcha model, so that the variability of the production shock does not affect behavior. Hence, in this class of models, the learning planner reacts to the anticipation of learning, independently of the amount of learning that takes place.

We also show that more optimistic beliefs decrease consumption if the mean of the production shock is positively related to the unknown parameter. Indeed, more optimistic beliefs increase the expected marginal utility of investment, inducing less consumption. Finally, unlike riskier distributions of the production shock, riskier beliefs affect consumption. A riskier distribution of beliefs leads to an increase in uncertainty through both sources of risk. The total effect of riskier beliefs depends on the strengths and directions of these two sources of risk.

The paper is organized as follows. In Section 2, we introduce learning in a general Brock–Mirman environment. In Section 3, optimal policies are characterized in the class of optimal stochastic growth models studied by Mirman–Zilcha. In Section 4, we study the effect of introducing learning on optimal policies. In Section 5, we perform a comparative analysis of distributions on the learning planner's optimal policies. Section 6 presents some final remarks for future research. All proofs are relegated to Appendix A.

2. Model

2.1. Brock–Mirman environment

Consider an economy in which output is determined by the production function $f(k,\eta)$, $f_1 > 0$, $f_{11} < 0$, as introduced in Mirman [33]. Here, k is capital and η is a realization of the random production shock $\tilde{\eta}$. The p.d.f. of $\tilde{\eta}$ is $\phi(\eta|\theta^*)$ for $\eta \in H \subset \mathbb{R}$, which depends on a parameter $\theta^* \in \Theta \subset \mathbb{R}^N$ for $N \in \mathbb{N}$. The relationship between the distribution of $\tilde{\eta}$ and the parameter θ^* is strictly monotonic.

Each period, a planner divides output y between consumption c and investment $k = y - c$. Capital k is used for the production of output $\hat{y}$ in the subsequent period, i.e.,

$$\hat{y} = f(y - c, \eta). \tag{1}$$

The objective is to maximize the expected sum of discounted utilities, where the discount factor is $\delta \in (0,1)$ and the utility function is $u(c)$, $u' > 0$, $u'' < 0$. Expectations are taken with respect to the sequence of future production shocks.

2.2. Full information planner

We first recall the full information growth model of Brock–Mirman, where the planner faces no structural uncertainty, i.e., the planner is informed because θ^* is known. Given θ^*, the informed planner anticipates the effect of the production shock on future output. The value function is

$$V_I\big(y;\theta^*\big) = \max_{c\in[0,y]} \left\{ u(c) + \delta \int_H V_I\big(f(y-c,\eta);\theta^*\big)\phi\big(\eta\big|\theta^*\big)\,\mathrm{d}\eta \right\}, \tag{2}$$

yielding optimal consumption $g_I(y;\theta^*)$.

2.3. Learning planner

We now relax the assumption of no structural uncertainty. Here, the planner faces structural uncertainty because θ^* is not known. Structural uncertainty is characterized by a priori beliefs about θ^*, expressed as a prior p.d.f. ξ on Θ. That is, the probability that $\theta^* \in S$ is $\int_S \xi(\theta)\,\mathrm{d}\theta$ for any $S \subset \Theta$.

Structural uncertainty leads to learning and, thus, evolves over time. Indeed, the planner observes η, which yields information, and uses Bayesian methods to learn about θ^*. Formally, given ξ and η, the posterior $\hat{\xi}(\cdot|\eta)$ is

$$\hat{\xi}(\theta|\eta) = \frac{\phi(\eta|\theta)\xi(\theta)}{\int_\Theta \phi(\eta|x)\xi(x)\,\mathrm{d}x} \tag{3}$$

for $\theta \in \Theta$, by Bayes' Theorem. Bayes' rule (3) characterizes the learning process through the updating of beliefs in light of the information gleaned from observing η. Observing η directly, allows us to focus on an environment with learning but no experimentation. Indeed, (3) is independent of consumption.

The learning planner makes consumption and investment decisions, while learning about θ^*. That is, endowed with initial output and beliefs, consumption and investment are chosen. The production shock η is then realized and the output, in the subsequent period, is determined from (1). Information is gleaned from observing η, which, from (3), affects beliefs about θ^*.

A learning planner's decisions are subject to both (1) and (3). Indeed, the learning planner anticipates the effect of the production shock on both future output and posterior beliefs. The value function of the learning planner is

$$V_L(y,\xi) = \max_{c\in[0,y]} \left\{ u(c) + \delta \int_H V_L\big(f(y-c,\eta),\hat{\xi}(\cdot|\eta)\big)\left[\int_\Theta \phi(\eta|\theta)\xi(\theta)\,\mathrm{d}\theta\right]\mathrm{d}\eta \right\} \tag{4}$$

subject to (3). The optimization in (4) yields the optimal consumption $g_L(y,\xi)$.

Learning increases the uncertainty of future payoffs by introducing two sources of risk: structural uncertainty and uncertainty due to the anticipation of learning. In other words, there are two distinct components of learning.

The first is about *beliefs*. While the informed planner's beliefs about θ^* are degenerate, the learning planner's are nondegenerate. There is an increase in uncertainty of future payoffs when knowledge of the distribution of the production shock, $\phi(\eta|\theta^*)$ in (2), is replaced by the expected p.d.f. of $\tilde{\eta}$ with respect to beliefs ξ, $\int_\Theta \phi(\eta|\theta)\xi(\theta)\,\mathrm{d}\theta$ in (4).

The second component concerns *anticipation*, i.e., learning is anticipated using Bayesian updating. In a dynamic context, rational expectations imply that the information contained in the future production shock is anticipated. The anticipation of learning is integrated into (4) by anticipating the updated beliefs from ξ to $\hat{\xi}(\cdot|\eta)$ using (3).

The anticipation of learning is related to the nonseparability of control and learning since the dynamics given in (1) and (3) are entwined through the production shock. If the only link between periods were beliefs, i.e., no capital accumulation, then the anticipation of learning would have no effect on optimization.

2.4. Adaptive learning planner

In order to study the effect of introducing learning, we introduce the intermediate case of an adaptive learner.[3] As with the learning planner, the adaptive learning planner does not know θ^*, and has beliefs about it expressed as a p.d.f. ξ on Θ. However, unlike the learning planner, the adaptive learning planner does not anticipate learning.

Given beliefs, the adaptive learning planner anticipates the effect of the production shock solely on future output, while beliefs are assumed to remain constant in his objective function. Therefore, the value function of the adaptive learning planner is

$$V_{AL}(y;\xi) = \max_{c\in[0,y]}\left\{u(c) + \delta \int_H V_{AL}\big(f(y-c,\eta);\xi\big)\left[\int_\Theta \phi(\eta|\theta)\xi(\theta)\,\mathrm{d}\theta\right]\mathrm{d}\eta\right\}, \tag{5}$$

yielding optimal consumption $g_{AL}(y;\xi)$. The adaptive learning planner does, however, update beliefs in each period. Once information arrives, the adaptive learning planner adapts and updates beliefs, subject to (3). Therefore, the adaptive learning planner reacts to, but does not anticipate new information.

Note that the informed and adaptive learning planners differ solely in the distribution of the production shock. Indeed, knowledge of the distribution of the production shock, $\phi(\eta|\theta^*)$ in (2), is replaced by the expected p.d.f. of $\tilde{\eta}$ with respect to beliefs ξ, $\int_\Theta \phi(\eta|\theta)\xi(\theta)\,\mathrm{d}\theta$ in (5). Thus, the adaptive learning planner faces a more variable distribution of the production shock than the informed planner.

2.5. Comparisons

While comparing informed and learning planners captures the overall effect of introducing learning in growth, the introduction of the intermediate case of an adaptive learning planner allows us to study the beliefs and anticipation components independently. First, comparing (2) and (5) captures the beliefs component, i.e., the risk generated from not knowing θ^*. Second, comparing (4) and (5) captures the anticipation component, i.e., the risk generated from uncertain posterior beliefs.

2.6. Remarks

In general, dynamic programs with learning such as (4) are intractable, i.e., they are not solvable either analytically or numerically, when there is no separability of control and learning.[4] The problem is not only whether a solution exists, but if a solution can be characterized and its properties studied. Two aspects of dynamic programming with learning should be noted.

First, (4) depends on the variable y and the prior p.d.f. ξ on Θ. Unless the space Θ contains a finite number of elements, the state space (y,ξ) is infinite-dimensional, leading to the curse of dimensionality.

Second, the evolution of beliefs, according to Bayes' law, does not prevent the prior and posterior p.d.f.'s ξ and $\hat{\xi}(\cdot|\eta)$ from belonging to different families. This makes the solution of an

[3] See [19] for a detailed exposition of adaptive learning. See also [32].

[4] When there is separability, the dynamic program becomes a standard growth problem, so that the learning planner is identical to the adaptive learning planner.

infinite-horizon dynamic programming problem with Bayesian dynamics generally intractable. Indeed, the learning planner makes consumption and investment decisions, anticipating updating beliefs every period. In other words, $V(f(y-c,\eta), \hat{\xi}(\cdot|\eta))$ in (4) encompasses beliefs that have been updated infinitely many times.

3. Optimal policies

In order to deal with the complexities of learning in growth, we focus on the class of optimal stochastic growth models studied in Mirman–Zilcha with the following assumptions.

Assumption 3.1. The utility function is $u(c) = \ln c$.

Assumption 3.2. The production function is Cobb–Douglas, $f(k,\eta) = k^{\eta}$.

Assumption 3.3. The support of $\tilde{\eta}$ is $H = [0, 1]$ and η is observable.

Assumptions 3.1, 3.2, and 3.3 hold for the remainder of the paper. The model with log utility, Cobb–Douglas production, and general distributions of the production shock and beliefs about θ^* yields closed-form solutions for optimal consumption in the cases of full information, adaptive learning, and learning.

The combination of log utility and Cobb–Douglas production is needed to obtain a tractable characterization of optimal consumption and investment in a learning context. The Mirman–Zilcha class of models has three features that makes the analysis possible.

First, Assumptions 3.1 and 3.2 imply that optimal consumption and investment are linear in output in the full information case. The linearity property remains under learning, although the fraction of output consumed now depends on beliefs and evolves with new information.

Second, from Assumptions 3.1 and 3.2, the uncertainty in outcomes, i.e., the random production shock, enters the optimization problem through its mean. In other words, the Mirman–Zilcha class of models displays certainty equivalence. This feature is exploited in the learning case since the uncertainty in outcomes is mapped to its mean, implying that the unknown parameter affects optimal consumption solely through $\mu(\theta) = \int_0^1 \eta\phi(\eta|\theta)\,\mathrm{d}\eta$, the mean of $\tilde{\eta}$ given $\theta \in \Theta$. The relationship between the mean of the production shock and the unknown parameter is the key in determining the effect of learning on the optimal consumption function and comparative analysis.

Third, Assumptions 3.1 and 3.2 imply that no assumption is needed on the production shock, as well as on the distribution of prior beliefs. The Mirman–Zilcha class of models does away with all the difficulties inherent in Bayesian analysis. In particular, the prior need not belong to the conjugate family of the distribution of the production shock. In other words, solutions for optimal consumption and investment are valid for a wide range of priors, even those that are outside of families of distributions that are closed under sampling.

To illustrate the importance of our assumptions, consider an alternative, in which the shock affects the level of production, i.e., $f(k,\eta) = \eta k$. If the utility is $u(c) = (1-1/\gamma)^{-1} c^{1-\frac{1}{\gamma}}$, $\gamma \neq 1$, then, formally, $g_L(y,\xi) = \kappa_1(\xi)^{-\gamma} y$, where κ_1 is implicitly defined by

$$\kappa_1(\xi) = \left(1 + \delta^{\gamma}\left(\int_0^1 \eta^{1-\frac{1}{\gamma}} \kappa_1\big(\hat{\xi}(\cdot|\eta)\big)\left[\int_{\Theta} \phi(\eta|\theta)\xi(\theta)\,\mathrm{d}\theta\right]\mathrm{d}\eta\right)^{\gamma}\right)^{\frac{1}{\gamma}}. \tag{6}$$

However, κ_1 might not exist. If it does exist, its properties cannot generally be studied. Here, although optimal consumption is linear in output, the distribution of the shock is no longer characterized by a statistic, such as the mean as in Mirman–Zilcha. Therefore, results with general distribution functions cannot be obtained.

We first state the optimal consumption of both the informed and adaptive learning planners. We then present and illustrate the optimal consumption of the learning planner.

3.1. Benchmark models

From Mirman–Zilcha, the optimal consumption of the informed planner, corresponding to (2), is

$$g_I(y;\theta^*) = \left(1-\delta\mu(\theta^*)\right)y, \tag{7}$$

while the optimal consumption of the adaptive learning planner, corresponding to (5), is

$$g_{AL}(y,\xi) = \left(1-\delta\int_{\Theta}\mu(\theta)\xi(\theta)\,\mathrm{d}\theta\right)y. \tag{8}$$

The presence of structural uncertainty does not affect the optimal consumption function since the true expectation of $\tilde{\eta}$, $E[\tilde{\eta}|\theta^*] = \mu(\theta^*)$ in (7), is replaced by the expectation of $\tilde{\eta}$ given beliefs, $E[\tilde{\eta}|\xi] = \int_{\Theta}\mu(\theta)\xi(\theta)\,\mathrm{d}\theta$ in (8).

3.2. Learning planner

In Appendix A, we show that the value function of the learning planner is of the form,

$$V_L(y,\xi) = \kappa_1(\xi)\ln y + \kappa_2(\xi), \tag{9}$$

where $\kappa_1(\xi) = \int_{\Theta}(1-\delta\mu(\theta))^{-1}\xi(\theta)\,\mathrm{d}\theta$ and $\kappa_2(\xi)$ depends on ξ.

Proposition 3.4. *The optimal consumption of the learning planner is*

$$g_L(y,\xi) = \left(\int_{\Theta}\frac{\xi(\theta)\,\mathrm{d}\theta}{1-\delta\mu(\theta)}\right)^{-1}y. \tag{10}$$

Despite the fact that this class of growth models displays certainty equivalence, certainty equivalence does not imply the separation of control and learning. Indeed, the anticipation of learning changes the optimal consumption function for the learning planner.

The dynamics of optimal consumption is changed as well. Indeed, a realization of a production shock has two effects on optimal consumption in the next period. First, the realization affects future output through the production function. Second, it affects the fraction of output consumed through Bayesian updating.

We present three examples that show the applicability of our model, not only in terms of distributions, but also in terms of general unknown structures. For instance, normal distributions are not needed to get analytic results. In Example 3.5, the case of learning about two unknown parameters is presented. Example 3.6 deals with a uniform distribution for $\tilde{\eta}$ with unknown support. Finally, Example 3.7 illustrates the case in which the learning planner does not know to which family $\tilde{\eta}$ belongs, as well as not knowing the parameters characterizing each family.

Example 3.5. Let $\tilde{\eta}$ have a beta distribution with unknown parameters $\theta = (\alpha, \beta)$, and beliefs $\xi(\alpha, \beta)$, $\alpha, \beta > 0$. Then, $\mu(\theta) = \alpha/(\alpha + \beta)$ and

$$g_L(y, \xi) = \left(\int_{\mathbb{R}^2_{++}} \frac{\xi(\alpha, \beta)\, \mathrm{d}\alpha\, \mathrm{d}\beta}{1 - \delta\alpha/(\alpha + \beta)} \right)^{-1} y. \tag{11}$$

Example 3.6. Let $\tilde{\eta}$ have a uniform distribution with unknown support $[0, \theta]$, and beliefs $\xi(\theta)$, $\theta \in [0, 1]$. Then, $\mu(\theta) = \theta/2$ and

$$g_L(y, \xi) = \left(\int_0^1 \frac{\xi(\theta)\, \mathrm{d}\theta}{1 - \delta\theta/2} \right)^{-1} y. \tag{12}$$

Example 3.7. Let $\Theta = \{\theta_1, \theta_2\}$, where θ_1 represents a beta distribution with unknown parameters (α, β), and beliefs $\xi_B(\alpha, \beta)$, $\alpha, \beta > 0$, while θ_2 represents a truncated normal distribution with support $[0, 1]$, unknown parameters (m, σ^2), and beliefs $\xi_N(m, \sigma^2)$, $m > 0, \sigma^2 \in \mathbb{R}_{++}$. If $0 \leqslant \rho \leqslant 1$ is the prior probability that the production shock is beta distributed, then

$$\begin{aligned} &g_L(y, \rho, \xi_B, \xi_N) \\ &\quad = \left(\rho \int_{\mathbb{R}^2_{++}} \frac{\xi_B(\alpha, \beta)\, \mathrm{d}\alpha\, \mathrm{d}\beta}{1 - \delta\mu_1(\alpha, \beta)} + (1 - \rho) \int_{\mathbb{R}_{++}} \int_{\mathbb{R}} \frac{\xi_N(m, \sigma^2)\, \mathrm{d}m\, \mathrm{d}\sigma^2}{1 - \delta\mu_2(m, \sigma^2)} \right)^{-1} y, \end{aligned} \tag{13}$$

where $\mu_1(\alpha, \beta)$ is the mean of a beta random variable with parameters (α, β) and $\mu_2(m, \sigma^2)$ is the mean of a truncated normal random variable with parameters (m, σ^2).

4. The effect of learning on optimal policies

In this section, conditions under which the introduction of learning increases or decreases optimal consumption are provided. Here, the effect depends on the strengths and directions of the beliefs and the anticipation components of learning. Formally, the effect depends upon the curvature of the marginal utility of investment with respect to the random variable.[5]

To see this, consider the first-order conditions of the informed planner,

$$\frac{1}{c} = \frac{\delta R(\mu(\theta^*))}{y - c}, \tag{14}$$

the adaptive learning planner,

$$\frac{1}{c} = \frac{\delta R(\int_\Theta \mu(\theta)\xi(\theta)\, \mathrm{d}\theta)}{y - c}, \tag{15}$$

and the learning planner,

$$\frac{1}{c} = \frac{\delta \int_\Theta R(\mu(\theta))\xi(\theta)\, \mathrm{d}\theta}{y - c}. \tag{16}$$

[5] Consider a two-period model in which the planner maximizes $u(c) + \delta E[u(f(y - c, \tilde{\eta}))]$ over c. If $\hat{y} = f(y - c) + \eta$, then the convexity of the marginal utility of consumption leads to less consumption as risk increases. If $\hat{y} = \eta f(y - c)$, then the convexity of $\eta f'(y - c)u'(\eta f(y - c))$ with respect to η leads to less consumption as risk increases.

Here, $R(x) = x(1-\delta x)^{-1}$, $R', R'' > 0$, for $x \in [0,1]$ characterizes the effect of uncertainty in outcomes due to the random production shock $\tilde{\eta}$ on the expected marginal utility of investment.

From (14) and (15), structural uncertainty affects the expected marginal utility of investment. Here, it is the second derivative of the mean of the production shock, with respect to the unknown parameter, that determines the effect of an increase in risk due to structural uncertainty.

Moreover, from (15) and (16), the anticipation of learning affects the expected marginal utility of investment. Here, it is the convexity of R with respect to μ that determines the effect of an increase in risk due to the anticipation of learning.

Finally, from (14) and (16), the overall effect of learning on optimal consumption is characterized by the expectation of R with respect to beliefs ξ. Here, it is the second derivative of R with respect to the unknown parameter θ that determines the overall effect of introducing learning through both its beliefs and anticipation components.

The impact of the anticipation component of learning is revealed by comparing (15) and (16). Proposition 4.1 states that the anticipation of learning always decreases optimal consumption. Formally,

Proposition 4.1. $g_{AL}(y;\xi) > g_L(y,\xi)$.

Proposition 4.1 is due to the convexity of R with respect to μ, and the use of Jensen's inequality on the right-hand sides of (15) and (16). The risk generated from the anticipation of learning increases the expected marginal utility of investment, leading to a decrease in consumption.

Next, the effect of introducing learning in an optimal growth model, when beliefs are unbiased, is studied. First, we consider beliefs about the mean of the production shock that are unbiased, i.e., $\mu(\theta^*) = \int_\Theta \mu(\theta)\xi(\theta)\,\mathrm{d}\theta$. Second, we focus on beliefs about the parameter θ^* that are unbiased, i.e., $\theta^* = \int_\Theta \theta\xi(\theta)\,\mathrm{d}\theta$. In both cases, conditions are established under which the introduction of learning, overall and through each of its components, increases or decreases optimal consumption using (14), (15) and (16). In other words, $g_I(y;\theta^*)$, $g_{AL}(y;\xi)$, and $g_L(y,\xi)$ are ordered.

In Proposition 4.2, the effect of learning when beliefs are unbiased about the mean of the production shock is studied. From (14) and (15), risk from structural uncertainty does not change the expected marginal utility of investment, since the uncertainty in outcomes is characterized only through its mean. Since the true mean of the production shock and unbiased beliefs about the true mean of the production shock have the same effect on behavior, there is certainty equivalence. Therefore, the total effect of learning is due to the anticipation of learning. As established in Proposition 4.1, the risk generated from the anticipation component increases the expected marginal utility of investment, leading to less consumption. Formally,

Proposition 4.2. *Suppose beliefs are unbiased about the mean of the production shock,* $\mu(\theta^*) = \int_\Theta \mu(\theta)\xi(\theta)\,\mathrm{d}\theta$. *Then, learning decreases optimal consumption, and* $g_I(y;\theta^*) = g_{AL}(y;\xi) > g_L(y,\xi)$.

In Proposition 4.3, the effect of learning when beliefs are unbiased about the unknown parameter, $\theta^* = \int_\Theta \theta\xi(\theta)\,\mathrm{d}\theta$, is studied. The effect of learning in this case is not as simple as in Proposition 4.2. The reason is that both sources of risk due to learning are at work here. Indeed, the effect of structural uncertainty depends on the second derivative of the mean of the production shock with respect to θ. If the mean of the production shock is concave with respect to θ, then structural uncertainty increases consumption. In other words, as structural uncertainty is introduced, with θ^* replaced by unbiased beliefs about θ^*, the expected marginal utility of investment decreases, inducing more consumption. On the other hand, the expected marginal utility of investment increases if the mean of the production shock is convex, inducing less consumption.

This point is illustrated in Example 3.5 in which $\tilde{\eta}$ has a beta distribution with parameters $\alpha, \beta > 0$. If $\alpha \equiv \theta$ is unknown and β is known, then $\mu''(\theta) = -2\beta/(\theta+\beta)^3 < 0$, and structural uncertainty increases consumption. However, if α is known and $\beta \equiv \theta$ is unknown, then $\mu''(\theta) = 2\alpha/(\alpha+\theta)^3 > 0$, and structural uncertainty decreases consumption.

The total effect depends on the strengths and directions of the beliefs and anticipation components. If the mean of the production shock with respect to the parameter is convex, then the two types of risk work in the same direction and consumption decreases. However, if the mean of the production shock is concave, then both types of risk pull in opposite directions and the effect of learning depends on the strength of each risk. Mathematically, it is the second derivative of R with respect to θ that determines the strength of the overall effect, i.e.,

$$\frac{\mathrm{d}^2 R}{\mathrm{d}\theta^2} = \frac{\mu''(\theta)(1-\delta\mu(\theta)) + 2\delta\mu'(\theta)^2}{(1-\delta\mu(\theta))^3}, \tag{17}$$

for $\theta \in \Theta$. The sign of (17) is determined by the sign of μ'' and the relationship $\mu'' \gtreqless -2\delta\mu'^2/(1-\delta\mu)$. Formally,

Proposition 4.3. *Suppose that beliefs are unbiased about the parameter, $\theta^* = \int_\Theta \theta\xi(\theta)\,\mathrm{d}\theta$.*

(1) *If $\mu'' > 0$, then $g_I(y;\theta^*) > g_{AL}(y;\xi) > g_L(y,\xi)$.*
(2) *If $\mu'' = 0$, then $g_I(y;\theta^*) = g_{AL}(y;\xi) > g_L(y,\xi)$.*
(3) *If $-2\delta\mu'^2/(1-\delta\mu) < \mu'' < 0$, then $g_L(y,\xi) < g_I(y;\theta^*) < g_{AL}(y;\xi)$.*
(4) *If $\mu'' = -2\delta\mu'^2/(1-\delta\mu)$, then $g_L(y,\xi) = g_I(y;\theta^*) < g_{AL}(y;\xi)$.*
(5) *If $\mu'' < -2\delta\mu'^2/(1-\delta\mu)$, then $g_I(y;\theta^*) < g_L(y,\xi) < g_{AL}(y;\xi)$.*

In case (1), the convexity of the mean of the production shock implies that structural uncertainty decreases consumption, as does the anticipation component. In other words, the two types of risk work in the same direction. In case (2), the mean of the production shock is linear in θ, so that structural uncertainty has no effect on the expected marginal utility of investment. Here, consumption decreases solely due to the anticipation component. In case (3), the mean of the production shock is concave. Here, the beliefs and anticipation components pull in opposite directions. The beliefs component increases, while the anticipation component, as is always the case, decreases consumption. But the mean of the production shock is not concave enough for the beliefs component to be dominant, and the overall effect of learning is to decrease consumption. Case (4) is a knife-edge case in which beliefs and anticipation components pull in opposite directions in equal strength. In case (5), the beliefs and anticipation components pull in opposite directions, but the mean of the production shock is concave enough to overwhelm the anticipation component. Thus, consumption increases.

5. Comparative analysis

In this section, the effect of different properties of the signal and initial beliefs on the learning planner's optimal consumption is studied. Specifically, we study the effects of the mean and riskiness of the distribution ϕ of the production shock $\tilde{\eta}$ as well as beliefs ξ about θ^*. The effect of riskier distributions on optimal consumption has been studied only in stochastic dynamic models in which the planner knows the distributions of stochastic variables. This analysis is extended to the learning case.

To facilitate the discussion, let $g_L^j(y, \xi)$ denote optimal consumption and $\mu^j(\theta) = \int_0^1 \eta \phi^j(\eta|\theta)\,\mathrm{d}\eta$, for the distribution ϕ^j. Moreover, let $g_L(y, \xi^j)$ denote optimal consumption with respect to ξ^j, $j = 1, 2$. Finally, we use the following definitions.

Definition 5.1. The p.d.f. φ^1 first-order stochastically dominates the p.d.f. φ^2, $\varphi^1 \succ_1 \varphi^2$, i.e., φ^1 has a higher mean than φ^2, if, for every nondecreasing function $\lambda : \mathbb{R} \to \mathbb{R}$, $\int_{\mathbb{R}} \lambda(x)\varphi^1(x)\,\mathrm{d}x \geqslant \int_{\mathbb{R}} \lambda(x)\varphi^2(x)\,\mathrm{d}x$.

Definition 5.2. For any two p.d.f.'s φ^1 and φ^2, φ^1 second-order stochastically dominates the p.d.f. φ^2, $\varphi^1 \succ_2 \varphi^2$, i.e., φ_1 is less risky than φ_2, if, for every concave function $\lambda : \mathbb{R} \to \mathbb{R}$, $\int_{\mathbb{R}} \lambda(x)\varphi^1(x)\,\mathrm{d}x \geqslant \int_{\mathbb{R}} \lambda(x)\varphi^2(x)\,\mathrm{d}x$.

5.1. Properties of the signal

Proposition 5.3 shows that a higher mean of the production shock $\tilde{\eta}$ decreases consumption.

Proposition 5.3. *If $\phi^1 \succ_1 \phi^2$, then $g_L^1(y, \xi) \leqslant g_L^2(y, \xi)$.*

From (16), the expected marginal utility of investment is greater under ϕ^1 than under ϕ^2 for $\phi^1 \succ_1 \phi^2$, inducing more investment and less consumption. Here, the planner faces a higher expected production shock for the next period, hence current investment is more profitable, and, thus, increases.[6]

Proposition 5.4 shows that an increase in the riskiness of the distribution of the production shock $\tilde{\eta}$ has no effect on consumption.

Proposition 5.4. *If $\phi^1 \succ_2 \phi^2$, then $g_L^1(y, \xi) = g_L^2(y, \xi)$.*

From (16), only $\mu(\theta)$ affects the expected marginal utility of investment. Hence, in the Mirman–Zilcha class of models, the learning planner's reaction to the anticipation of learning is independent of the amount of learning that takes place. Specifically, the informativeness of the signal has no effect on decisions. In other words, certainty equivalence regarding the random production shock continues to hold in this model with the introduction of learning.

Finally, changes in the mean and riskiness of the distribution of the production shock do have a dynamic effect on consumption in the subsequent period through updating beliefs. Indeed, if $\phi^1 \succ_1 \phi^2$ or $\phi^1 \succ_2 \phi^2$, then $\hat{\xi}^1 \neq \hat{\xi}^2$, for the same η. Hence, $g_L^1(\hat{y}, \hat{\xi}^1) \neq g_L^2(\hat{y}, \hat{\xi}^2)$.

5.2. Properties of prior beliefs

Proposition 5.5 shows the effect of more optimistic beliefs about θ^* on consumption. The effect of more optimistic beliefs depends on the first derivative of $\mu(\theta)$.

6 For studies on the effect of news about future productivity shocks on decision making in dynamic models that do not embed learning in the maximization problem, see [4–6,11,27] for the literature on expectation driven business cycles.

Proposition 5.5. *Suppose that* $\xi^1 \succ_1 \xi^2$.

(1) *If* $\mu' > 0$*, then* $g_L(y, \xi^1) \leqslant g_L(y, \xi^2)$.
(2) *If* $\mu' < 0$*, then* $g_L(y, \xi^1) \geqslant g_L(y, \xi^2)$.
(3) *If* $\mu' = 0$*, then* $g_L(y, \xi^1) = g_L(y, \xi^2)$.

From (16), if $\mu' > 0$, then $\xi^1 \succ_1 \xi^2$ implies that the expected marginal utility of investment is greater under ξ^1 than under ξ^2. Here, more optimistic beliefs about the production shock induces more investment and less consumption.

While, as stated in Proposition 5.4, a riskier distribution of $\tilde{\eta}$ does not affect consumption, Proposition 5.6 shows that riskier beliefs about θ^* do affect consumption. Proposition 5.6 generalizes Proposition 4.3. Recall that in Proposition 4.3, informed and learning planners are compared by increasing risk around θ^*. Here, two learning planners, one with riskier beliefs about θ^* than the other are compared.

As in Proposition 4.3, the effect of a riskier prior on consumption is determined by the sign of (17). Formally,

Proposition 5.6. *Suppose that* $\xi^1 \succ_2 \xi^2$.

(1) *If* $\mu'' < -2\delta\mu'^2/(1-\delta\mu)$*, then* $g_L(y, \xi^1) \leqslant g_L(y, \xi^2)$.
(2) *If* $\mu'' > -2\delta\mu'^2/(1-\delta\mu)$*, then* $g_L(y, \xi^1) \geqslant g_L(y, \xi^2)$.
(3) *If* $\mu'' = -2\delta\mu'^2/(1-\delta\mu)$*, then* $g_L(y, \xi^1) = g_L(y, \xi^2)$.

The discussion is similar to the one for Proposition 4.3. In case (1), the beliefs and anticipation components pull in opposite directions, but the mean of the production shock is concave enough to overwhelm the anticipation component. Thus, as beliefs become riskier, consumption increases. In case (2), the anticipation component is dominant, implying that the marginal utility of investment is convex in θ, leading to less consumption as beliefs become riskier. In case (3), the beliefs and anticipation components pull in opposite directions in equal strength, implying that the marginal utility of investment is linear in θ. There is no reaction to riskier beliefs.[7]

Finally, changes in the mean and riskiness of beliefs have a dynamic effect on consumption in the subsequent period through updating beliefs. Indeed, if $\xi^1 \succ_1 \xi^2$ or $\xi^1 \succ_2 \xi^2$, then $\hat{\xi}^1 \neq \hat{\xi}^2$, for the same η. Hence, $g_L(\hat{y}, \hat{\xi}^1) \neq g_L(\hat{y}, \hat{\xi}^2)$.

6. Final remarks

We have introduced learning in an optimal growth model in which the planner learns about the structure of the economy while making consumption and investment decisions. Since the planner observes the production shock directly there is no experimentation in this model, i.e., consumption and investment decisions cannot affect the flow of information.

There are situations in which it is reasonable to assume that the production shock is not observable. In this case, output is the signal used to update beliefs about the unknown parameter, which, in general, leads to experimentation. Previous work on experimentation in growth has

[7] It is possible to extend Proposition 4.2 by comparing two learning planners, one with riskier beliefs about $\mu(\theta^*)$ than the other one. As in Proposition 4.2, consumption always decreases as beliefs about $\mu(\theta^*)$ become riskier.

only characterized optimal policies in models with three periods and a two-value support of the unknown parameter. See Bertocchi and Spagat [8] and Datta et al. [13]. Future research should focus on characterizing optimal policies under experimentation in a more general setting.

Acknowledgments

We are grateful to the associate editor and three referees for their very helpful comments. We thank Hafedh Bouakez, Marco Cagetti, Manjira Datta, David K. Levine, Toshihiko Mukoyama, Kevin Reffett, John Stachurski, as well as seminar participants at the Federal Reserve Bank in Richmond, the 2006 Spring CAM conference in stochastic modeling at the University of Notre Dame, and the Fall 2006 Midwest Economic Theory Meetings at Purdue University.

Appendix A. Proofs

Proof of Proposition 3.4. We conjecture that the value function of the learning planner is of the form $V_L(y,\xi) = \kappa_1(\xi)\ln y + \kappa_2(\xi)$, where κ_1 and κ_2 depend on ξ. From (4),

$$V_L(y,\xi) = \max_{c\in(0,y)} \left\{ \ln c + \delta \ln(y-c) \int_0^1 \kappa_1\big(\hat{\xi}(\cdot|\eta)\big)\eta \left[\int_\Theta \phi(\eta|\theta)\xi(\theta)\,\mathrm{d}\theta \right] \mathrm{d}\eta \right.$$
$$\left. + \delta \int_0^1 \kappa_2\big(\hat{\xi}(\cdot|\eta)\big) \left[\int_\Theta \phi(\eta|\theta)\xi(\theta)\,\mathrm{d}\theta \right] \mathrm{d}\eta \right\}. \tag{A.1}$$

The first-order condition is $c^{-1} = \delta(y-c)^{-1}\int_0^1 \kappa_1(\hat{\xi}(\cdot|\eta))\eta[\int_\Theta \phi(\eta|\theta)\xi(\theta)\,\mathrm{d}\theta]\,\mathrm{d}\eta$, evaluated at $c = g_L(y,\xi)$, so that

$$g_L(y,\xi) = \left(1 + \delta \int_0^1 \kappa_1\big(\hat{\xi}(\cdot|\eta)\big)\eta \left[\int_\Theta \phi(\eta|\theta)\xi(\theta)\,\mathrm{d}\theta \right] \mathrm{d}\eta \right)^{-1} y. \tag{A.2}$$

Plugging (A.2) into (A.1) yields

$$V_L(y,\xi) = \left(1 + \delta \int_0^1 \kappa_1\big(\hat{\xi}(\cdot|\eta)\big)\eta \left[\int_\Theta \phi(\eta|\theta)\xi(\theta)\,\mathrm{d}\theta \right] \mathrm{d}\eta \right) \ln y + \kappa_3(\xi), \tag{A.3}$$
$$\equiv \kappa_1(\xi)\ln y + \kappa_2(\xi), \tag{A.4}$$

where κ_3 depends on ξ as well. Therefore,

$$\kappa_1(\xi) \equiv 1 + \delta \int_0^1 \kappa_1\big(\hat{\xi}(\cdot|\eta)\big)\eta \left[\int_\Theta \phi(\eta|\theta)\xi(\theta)\,\mathrm{d}\theta \right] \mathrm{d}\eta. \tag{A.5}$$

The solution to (A.5) is

$$\kappa_1(\xi) = \int_\Theta \frac{\xi(\theta)\,\mathrm{d}\theta}{1-\delta\mu(\theta)}, \tag{A.6}$$

where $\mu(\theta)=\int_0^1 \eta\phi(\eta|\theta)\,\mathrm{d}\eta$. To verify that (A.6) is the solution to (A.5), updating (A.6) to the next period yields

$$\kappa_1\big(\hat{\xi}(\cdot|\eta)\big)=\int\limits_{\Theta}\frac{\hat{\xi}(\theta|\eta)\,\mathrm{d}\theta}{1-\delta\mu(\theta)}, \tag{A.7}$$

$$=\int\limits_{\Theta}\frac{1}{1-\delta\mu(\theta)}\frac{\phi(\eta|\theta)\xi(\theta)}{\int_{\Theta}\phi(\eta|x)\xi(x)\,\mathrm{d}x}\,\mathrm{d}\theta. \tag{A.8}$$

Then, plugging (A.8) into (A.5) yields (A.6), verifying the conjecture of the value function. Combining (A.5) and (A.6) yields

$$g_L(y,\xi)=\left(1+\delta\int\limits_0^1\kappa_1\big(\hat{\xi}(\cdot|\eta)\big)\eta\left[\int\limits_{\Theta}\phi(\eta|\theta)\xi(\theta)\,\mathrm{d}\theta\right]\mathrm{d}\eta\right)^{-1}y, \tag{A.9}$$

$$=y/\kappa_1(\xi), \tag{A.10}$$

$$=\left(\int\limits_{\Theta}\frac{\xi(\theta)\,\mathrm{d}\theta}{1-\delta\mu(\theta)}\right)^{-1}y. \tag{A.11}$$

Since both the utility and production functions are strictly concave in c, (A.11) is the unique maximizer corresponding to (A.1). □

Proof of Proposition 4.1. Since $R''(x)=2\delta/(1-\delta x)^3>0$, the right-hand side of (15) is less than the right-hand side of (16) for any c, by Jensen's inequality. Therefore, $g_{AL}(y;\xi)>g_L(y,\xi)$. □

Proof of Proposition 4.2. Suppose $\mu(\theta^*)=\int_{\Theta}\mu(\theta)\xi(\theta)\,\mathrm{d}\theta$. First, $g_I(y;\theta^*)=g_{AL}(y;\xi)$ from (14) and (15). Second, $g_{AL}(y;\xi)>g_L(y,\xi)$ from Proposition 4.1. Therefore, $g_I(y;\theta^*)=g_{AL}(y;\xi)>g_L(y,\xi)$. □

Proof of Proposition 4.3. Suppose $\theta^*=\int_{\Theta}\theta\xi(\theta)\,\mathrm{d}\theta$. First, if $\mu''<-2\delta\mu'^2/(1-\delta\mu)$, then, for any c, the right-hand side of (14) is greater than the right-hand side of (16) by Jensen's inequality. Therefore, $g_I(y;\theta^*)<g_L(y,\xi)$. The proofs for $\mu''>-2\delta\mu'^2/(1-\delta\mu)$ and $\mu''=-2\delta\mu'^2/(1-\delta\mu)$ are identical. Second, if $\mu''<0$, then, for any c, the right-hand side of (14) is greater than the right-hand side of (15), since $\mu(\theta^*)>\int_{\Theta}\mu(\theta)\xi(\theta)\,\mathrm{d}\theta$ by Jensen's inequality. Therefore, $g_I(y;\theta^*)<g_{AL}(y;\xi)$. The proofs for $\mu''>0$ and $\mu''=0$ are identical. Third, $g_{AL}(y;\xi)>g_L(y,\xi)$ from Proposition 4.1. Combining these three points yields Proposition 4.2. □

Proof of Proposition 5.3. From (10), if $\phi^1\succ_1\phi^2$, then $\mu^1(\theta)=\int_0^1\eta\phi^1(\eta|\theta)\,\mathrm{d}\eta\geqslant\int_0^1\eta\phi^2(\eta|\theta)\,\mathrm{d}\eta=\mu^2(\theta)$ implying that $g_L^1(y,\xi)\leqslant g_L^2(y,\xi)$. □

Proof of Proposition 5.4. From (10), if $\phi^1\succ_2\phi^2$, then $\mu^1(\theta)=\int_0^1\eta\phi^1(\eta|\theta)\,\mathrm{d}\eta=\int_0^1\eta\phi^2(\eta|\theta)\,\mathrm{d}\eta=\mu^2(\theta)$ implying that $g_L^1(y,\xi)=g_L^2(y,\xi)$. □

Proof of Proposition 5.5. Suppose that $\xi^1\succ_1\xi^2$. If $\mu'>0$, then, for every c, the expected marginal return on investment in (16) is greater under ξ^1 than under ξ^2. Therefore, $g_L(y,\xi^1)\leqslant g_L(y,\xi^2)$. The proofs for $\mu'<0$ and $\mu'=0$ are identical. □

Proof of Proposition 5.6. Suppose that $\xi^1 \succ_2 \xi^2$. If $\mu'' < -2\delta\mu'^2/(1-\delta\mu)$, then, for every c, the expected marginal return on investment in (16) is greater under ξ^1 than under ξ^2, by Jensen's inequality. Therefore, $g_L(y,\xi^1) \leqslant g_L(y,\xi^2)$. The proofs for $\mu'' > -2\delta\mu'^2/(1-\delta\mu)$ and $\mu'' = -2\delta\mu'^2/(1-\delta\mu)$ are identical. □

References

[1] P. Aghion, P. Bolton, C. Harris, B. Jullien, Optimal learning by experimentation, Rev. Econ. Stud. 58 (4) (1991) 621–654.
[2] E. Antoniadou, C. Koulovatianos, L. Mirman, Strategic exploitation of a common property resource under uncertainty, Vienna Economic Papers 0703, University of Vienna, Department of Economics, 2007.
[3] R. Balvers, T. Cosimano, Actively learning about demand and the dynamics of price adjustment, Econ. J. 100 (402) (1990) 882–898.
[4] P. Beaudry, F. Portier, An exploration into Pigou's theory of cycles, J. Monet. Econ. 51 (6) (2004) 1183–1216.
[5] P. Beaudry, F. Portier, Stock prices, news and economic fluctuations, Amer. Econ. Rev. 96 (4) (2006) 1293–1307.
[6] P. Beaudry, F. Portier, When can changes in expectations cause business cycle fluctuations, J. Econ. Theory 135 (1) (2007) 458–477.
[7] G. Beck, V. Wieland, Learning and control in a changing economic environment, J. Econ. Dynam. Control 26 (9) (2002) 1359–1377.
[8] G. Bertocchi, M. Spagat, Growth under uncertainty with experimentation, J. Econ. Dynam. Control 23 (2) (1998) 209–231.
[9] W. Brock, L. Mirman, Optimal economic growth and uncertainty: The discounted case, J. Econ. Theory 4 (3) (1972) 479–513.
[10] D. Cass, Optimum growth in an aggregative model of capital accumulation, Rev. Econ. Stud. 32 (3) (1965) 233–240.
[11] L. Christiano, R. Motto, M. Rostagno, Monetary policy and a stock market boom-bust cycle, Mimeo, University of Northwestern, 2006.
[12] A. Creane, Experimentation with heteroskedastic noise, Econ. Theory 4 (2) (1994) 275–286.
[13] M. Datta, L. Mirman, E. Schlee, Optimal experimentation in signal-dependent decision problems, Int. Econ. Rev. 43 (2) (2002) 577–607.
[14] W. Dechert, S. O'Donnell, W. Brock, Bayes' learning of unknown parameters, J. Difference Equ. Appl. 13 (2–3) (2007) 121–133.
[15] J. Drèze, F. Modigliani, Consumption decisions under uncertainty, J. Econ. Theory 5 (3) (1972) 308–335.
[16] D. Easley, N. Kiefer, Controlling a stochastic process with unknown parameters, Econometrica 56 (5) (1988) 1045–1064.
[17] D. Easley, N. Kiefer, Optimal learning with endogenous data, Int. Econ. Rev. 30 (4) (1989) 963–978.
[18] M. El-Gamal, R. Sundaram, Bayesian economists … Bayesian agents: An alternative approach to optimal learning, J. Econ. Dynam. Control 17 (3) (1993) 355–383.
[19] G. Evans, S. Honkapohja, Learning and Expectations in Macroeconomics, Princeton University Press, Princeton, NJ, 2001.
[20] A. Fishman, N. Gandal, Experimentation and learning with network effects, Econ. Letters 44 (1–2) (1994) 103–108.
[21] X. Freixas, Optimal growth with experimentation, J. Econ. Theory 24 (2) (1981) 296–309.
[22] J. Fusselman, L. Mirman, Experimental consumption for a general class of disturbance densities, in: R. Becker, M. Boldrin, R. Jones, M. Thompson (Eds.), General Equilibrium, Growth, and Trade. The Legacy of Lionel McKenzie. Economic Theory, Econometrics, and Mathematical Economics Series, vol. 2, Academy Press, London, 1993, pp. 367–392.
[23] S. Grossman, R. Kihlstrom, L. Mirman, A Bayesian approach to the production of information and learning by doing, Rev. Econ. Stud. 44 (3) (1977) 533–547.
[24] F. Hahn, Savings and uncertainty, Rev. Econ. Stud. 37 (1) (1970) 21–24.
[25] G. Huffman, N. Kiefer, Optimal learning in a one-sector growth model, Working Papers 90-24, Cornell University, Center for Analytic Economics, 1994.
[26] M. Huggett, Precautionary wealth accumulation, Rev. Econ. Stud. 71 (3) (2004) 769–781.
[27] N. Jaimovich, S. Rebelo, Can news about the future drive the business cycle?, NBER Working Papers 12537, National Bureau of Economic Research, Inc., 2006.
[28] G. Keller, S. Rady, Optimal experimentation in a changing world, Rev. Econ. Stud. 66 (3) (1999) 475–507.

[29] N. Kiefer, Y. Nyarko, Optimal control of an unknown linear process with learning, Int. Econ. Rev. 30 (3) (1989) 571–586.
[30] T. Koopmans, On the concept of economic growth, in: Semaine d'Étude sur le Rôle de l'Analyse Économétrique dans la Formulation de Plans de Développement, vol. 28, Pontificau Academiae Scientiarum Scripta Varia, Vatican, 1965, pp. 225–300.
[31] H. Leland, Saving and uncertainty: The precautionary demand for saving, Quart. J. Econ. 82 (3) (1968) 465–473.
[32] F. Milani, Expectations, learning and macroeconomic persistence, J. Monet. Econ. 54 (7) (2007) 2065–2082.
[33] L. Mirman, Two essays on uncertainty and economics, PhD thesis, University of Rochester, 1970.
[34] L. Mirman, Uncertainty and optimal consumption decisions, Econometrica 39 (1) (1971) 179–185.
[35] L. Mirman, On the existence of steady state measures for one sector growth models with uncertain technology, Int. Econ. Rev. 13 (2) (1972) 271–286.
[36] L. Mirman, The steady state behavior of a class of one sector growth models with uncertain technology, J. Econ. Theory 6 (3) (1973) 219–242.
[37] L. Mirman, L. Samuelson, A. Urbano, Monopoly experimentation, Int. Econ. Rev. 34 (3) (1993) 549–563.
[38] L. Mirman, I. Zilcha, On optimal growth under uncertainty, J. Econ. Theory 11 (3) (1975) 329–339.
[39] E. Prescott, The multi-period control problem under uncertainty, Econometrica 40 (6) (1972) 1043–1058.
[40] F. Ramsey, A mathematical theory of saving, Econ. J. 38 (152) (1928) 543–559.
[41] M. Rothschild, J. Stiglitz, Increasing risk II: Its economic consequences, J. Econ. Theory 3 (1) (1971) 66–84.
[42] A. Sandmo, The effect of uncertainty on saving decisions, Rev. Econ. Stud. 37 (3) (1970) 353–360.
[43] D. Trefler, The ignorant monopolist: Optimal learning with endogenous information, Int. Econ. Rev. 34 (3) (1993) 565–581.
[44] V. Wieland, Learning by doing and the value of optimal experimentation, J. Econ. Dynam. Control 24 (4) (2000) 501–534.