

Optimal Growth and Uncertainty: Learning*

Christos Koulovatianos[†] Leonard J. Mirman[‡]

Marc Santugini[§]

March 28, 2008

*We are grateful to the Associate Editor and three anonymous 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.

[†]University of Vienna, Department of Economics. Email: koulovc6@univie.ac.at.

[‡]University of Virginia, Department of Economics. Email: lm8h@virginia.edu.

[§]Corresponding author. HEC Montréal, Institute of Applied Economics, and CIRPÉE. Email: marc.santugini@hec.ca.

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 from 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 changes in the mean and riskiness of the distribution of the signal and initial beliefs on optimal consumption are studied.

1 Introduction

In the early literature on optimal growth, the evolution of output was deterministic, see Cass (1965) and Koopmans (1965). This was a natural place to begin the study of optimal growth since growth had already been studied in a deterministic environment by Ramsey (1928). Brock and Mirman (1972) introduced uncertainty in outcomes in an optimal growth model, which built on earlier studies of stochastic positive growth, see Mirman (1972, 1973). Uncertainty in outcomes is modeled by introducing a random shock in the production function. Hence, the future is riskier than in the deterministic case since future output is random, affecting optimal consumption and investment.

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, to reduce structural uncertainty. At the same time, the planner also makes consumption and investment decisions. These two functions of the planner are inter-related.

We introduce learning in a Brock-Mirman environment and study the effect of risk generated by the planner's econometric activity on optimal decisions. We extend and combine the literature on learning and risk in a growth context. 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

¹Experimentation was initially studied in models in which the only link between periods is beliefs. See Prescott (1972), Grossman et al. (1977), Easley and Kiefer (1988, 1989), Kiefer and Nyarko (1989), Balvers and Cosimano (1990), Aghion et al. (1991), Fusselman and Mirman (1993), Mirman et al. (1993), Treffer (1993), Creane (1994), Fishman and Gandal (1994), Keller and Rady (1999), and Wieland (2000). Experimentation in a model with capital accumulation has also been studied. See Freixas (1981), Bertocchi and Spagat (1998), Datta et al. (2002), El-Gamal and Sundaram (1993), Huffman and Kiefer (1994), Beck and Wieland (2002), and Dechert et al. (2007).

payoffs. While the literature on learning has focused on experimentation, we concentrate on the effect of learning on decision making. In that sense, we extend the literature on risk to a dynamic learning environment. Previous literature on risk has focused only on models in which the planner knows the distributions of stochastic variables.²

We provide two sets of results in the class of optimal stochastic growth models studied by Mirman and Zilcha (1975), with specific utility and production functions, but general distributions of the production shock and beliefs.

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 from the anticipation of learning. The effect depends on the strengths and directions of the two sources of risk, which may pull in opposite directions.

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 marginal utility of investment decreases with a concave mean, inducing more consumption, while the marginal utility of investment increases with convexity, inducing less consumption. On the other hand, the risk generated from the anticipation of learning always increases the 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. If the mean of the production shock with respect to the unknown parameter is convex, 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.

Second, we perform a comparative analysis of distributions on the learning

²See Leland (1968), Hahn (1970), Sandmo (1970), Rothschild and Stiglitz (1971), and Drèze and Modigliani (1972) for a finite-period analysis. For an infinite-horizon setup, see Mirman (1971) in a model with a single agent, and Antoniadou et al. (2007) for the case of a game. Finally, see Huggett (2004) for a detailed review of other issues studied in this literature.

planner's optimal consumption. Specifically, the effects of changes in 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 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 risks. 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 consumption and investment 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 consumption. Section 6 presents some final remarks for future research. All proofs are relegated to the appendix.

2 Model

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 (1970). 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.,

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

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.

Full Information Planner. We first recall the full information growth model of Brock and Mirman (1972), 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

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

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

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 $\int_S \xi(\theta) d\theta$ 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

$$\hat{\xi}(\theta|\eta) = \frac{\phi(\eta|\theta)\xi(\theta)}{\int_{\Theta} \phi(\eta|x)\xi(x)dx} \quad (3)$$

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

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

subject to (3), yielding optimal consumption $g_L(y, \xi)$.

Learning increases the uncertainty of future payoffs by introducing two sources of risk: structural uncertainty and uncertainty from 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 , $\int_{\Theta} \phi(\eta|\theta)\xi(\theta)d\theta$ 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. Anticipation of learning is integrated into (4) by anticipating the updated

beliefs from ξ to $\hat{\xi}(\cdot|\eta)$ using (3).

Nonseparability of Control and Learning. 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.

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

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

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 new information, but does not anticipate it.

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 , $\int_{\Theta} \phi(\eta|\theta) \xi(\theta) d\theta$ in (5). Thus, the adaptive learning planner faces a more variable distribution of the production shock than the informed planner.

Comparisons. While comparing informed and learning planners captures the overall effect of introducing learning in growth, the introduction

³See Evans and Honkapohja (2001) for a detailed exposition of adaptive learning. See also Milani (2007).

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.

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 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, the value function, $V\left(f(y - c, \eta), \hat{\xi}(\cdot|\eta)\right)$ in (4), encompasses beliefs that have been updated infinitely many times.

3 Optimal Consumption and Investment

In order to deal with the complexities of learning in growth, we focus on the class of optimal stochastic growth models studied by Mirman and Zilcha (1975) with the following assumptions.

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

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

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

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 consumptions 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 and Zilcha (1975) 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 $\mu(\theta) = \int_0^1 \eta \phi(\eta|\theta) d\eta$, 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.

We first state the optimal consumptions of both the informed and adaptive learning planners. We then present and illustrate the optimal consump-

tion of the learning agent.

Benchmark Models. From Mirman-Zilcha, the optimal consumption of the informed planner, corresponding to (2), is

$$g_I(y; \theta^*) = (1 - \delta\mu(\theta^*))y, \quad (6)$$

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

$$g_{AL}(y, \xi) = \left(1 - \delta \int_{\Theta} \mu(\theta)\xi(\theta)d\theta\right) y. \quad (7)$$

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 (6), is replaced by the expectation of $\tilde{\eta}$ given beliefs, $E[\tilde{\eta}|\xi] = \int_{\Theta} \mu(\theta)\xi(\theta)d\theta$ in (7).

Learning Planner. In the appendix, we show that the value function of the learning planner is of the form,

$$V_L(y, \xi) = \kappa_1(\xi) \ln y + \kappa_2(\xi), \quad (8)$$

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

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

$$g_L(y, \xi) = \left(\int_{\Theta} \frac{\xi(\theta)d\theta}{1 - \delta\mu(\theta)}\right)^{-1} y. \quad (9)$$

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 are 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 wide 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

$$g_L(y, \xi) = \left(\int_{\mathbb{R}_{++}^2} \frac{\xi(\alpha, \beta) d\alpha d\beta}{1 - \delta\alpha/(\alpha + \beta)} \right)^{-1} y. \quad (10)$$

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

$$g_L(y, \xi) = \left(\int_0^1 \frac{\xi(\theta) d\theta}{1 - \delta\theta/2} \right)^{-1} y. \quad (11)$$

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 $0 \leq \rho \leq 1$ is the prior probability that the production shock is beta distributed, then

$$g_L(y, \rho, \xi_B, \xi_N) = \left(\rho \int_{\mathbb{R}_{++}^2} \frac{\xi_B(\alpha, \beta) d\alpha d\beta}{1 - \delta\mu_1(\alpha, \beta)} + (1 - \rho) \int_{\mathbb{R}_{++}} \int_{\mathbb{R}} \frac{\xi_N(m, \sigma^2) dm d\sigma^2}{1 - \delta\mu_2(m, \sigma^2)} \right)^{-1} y, \quad (12)$$

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, we study the overall effect of introducing learning through both its beliefs and anticipation components. Our results extend the literature on the effect of an increase in risk on optimal consumption and investment to a learning environment. Previous literature has focused on models in which the planner knows the distributions of stochastic variables. An increase in risk of future payoffs results in substitution and precautionary behavior. First, increased risk in investment leads to a substitution from future consumption to present consumption. Second, increased risk raises the prospect of low consumption in future, leading to more investment for precautionary reasons. The effect of risk on consumption depends on the strengths of these two elements. Formally, the effect depends on the second derivative of some functions of the random variable being studied.⁵

In the learning case, the effect of an increase in risk also depends on the second derivatives of some functions of the appropriate random variable. To see this, consider the first-order conditions of the informed planner,

$$\frac{1}{c} = \frac{\delta R(\mu(\theta^*))}{y - c}, \quad (13)$$

the adaptive learning planner,

$$\frac{1}{c} = \frac{\delta R\left(\int_{\Theta} \mu(\theta)\xi(\theta)d\theta\right)}{y - c}, \quad (14)$$

and the learning planner,

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

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

⁵Consider 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 it is the convexity of the marginal utility of consumption that leads to precautionary investment. If $\hat{y} = \eta f(y - c)$, then it is the convexity of $\eta f'(y - c)u'(f(y - c))$ with respect to η that leads to precautionary investment.

expected marginal utility of investment.

From (13) and (14), 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 (14) and (15), the anticipation of learning affects the expected marginal utility of investment. Here, it is the convexity of R that determines the effect of an increase in risk due to the anticipation of learning.

Finally, from (13) and (15), 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 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 (14) and (15). 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 , and, thus, the expected marginal utility of investment, and the use of Jensen's inequality on the right-hand sides of (14) and (15). The risk generated from the anticipation of learning increases the expected marginal utility of investment, leading to a decrease in consumption or precautionary investment.

Beliefs and Anticipation Components. 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., $\mu(\theta^*) = \int_{\Theta} \mu(\theta)\xi(\theta)d\theta$. Second, we focus on beliefs about the parameter θ^* that are unbiased, i.e., $\theta^* = \int_{\Theta} \theta\xi(\theta)d\theta$. 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 (13), (14) and (15). 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 (13) and (14), 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 precautionary investment. Formally,

Proposition 4.2. *Suppose beliefs are unbiased about the mean of the production shock, $\mu(\theta^*) = \int_{\Theta} \mu(\theta)\xi(\theta)d\theta$. 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, $\theta^* = \int_{\Theta} \theta\xi(\theta)d\theta$, 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 with respect to θ is concave, then structural uncertainty increases optimal consumption. In other words, as structural uncertainty is introduced, with θ^* replaced by unbiased beliefs about θ^* , the marginal utility of investment decreases, inducing less investment. And, the marginal utility of investment increases with convexity, inducing more investment.

This point is illustrated in Example 3.5 in which $\tilde{\eta}$ has a beta distribution with parameters $\alpha, \beta > 0$ assuming that $\rho = 1$. 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 both types of risk work in the same direction and optimal 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.,

$$\frac{d^2 R}{d\theta^2} = \frac{\mu''(\theta)(1 - \delta\mu(\theta)) + 2\delta\mu'(\theta)^2}{(1 - \delta\mu(\theta))^3}, \quad (16)$$

for $\theta \in \Theta$. The sign of (16) is determined by the sign of μ'' and the relationship $\mu'' \begin{cases} \geq \\ < \end{cases} -2\delta\mu'^2/(1 - \delta\mu)$. Formally,

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

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 effect of changes in 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 $\mu^j(\theta) = \int_0^1 \eta \phi^j(\eta|\theta) d\eta$, 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 $\lambda : \mathbb{R} \rightarrow \mathbb{R}$, $\int_{\mathbb{R}} \lambda(x) \varphi^1(x) dx \geq \int_{\mathbb{R}} \lambda(x) \varphi^2(x) dx$.

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 $\lambda : \mathbb{R} \rightarrow \mathbb{R}$, $\int_{\mathbb{R}} \lambda(x) \varphi^1(x) dx \geq \int_{\mathbb{R}} \lambda(x) \varphi^2(x) dx$.

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 $g_L^1(y, \xi) \leq g_L^2(y, \xi)$.*

From (15), 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 immediately.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 (15), only $\mu(\theta)$ affects the expected marginal utility of investment. Hence, in this class of models, the learning agent reacts to the anticipation of learning, 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 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 $\mu(\theta)$.

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

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

⁶For 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 the literature on expectation driven business cycles beginning with Beaudry and Portier (2004, 2006, 2007), but also including Christiano et al. (2006), and Jaimovich and Rebelo (2006).

From (15), 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 (16). Formally,

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

1. *If $\mu'' < -2\delta\mu'^2/(1 - \delta\mu)$, then $g_L(y, \xi^1) \leq g_L(y, \xi^2)$.*
2. *If $\mu'' > -2\delta\mu'^2/(1 - \delta\mu)$, then $g_L(y, \xi^1) \geq g_L(y, \xi^2)$.*
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 decreases. In case 2, the anticipation component is dominant, implying that the expected marginal utility of investment is convex in θ , leading to precautionary investment as beliefs become riskier. In case 3, both beliefs and anticipation components pull in opposite directions in equal strength, implying that the expected 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.

⁷It 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.

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

In this paper, we have studied learning when the shock determines the curvature of the production. Consider an alternative production in which the shock affects the level, i.e., $f(k, \eta) = \eta k$. If the utility is $u(c) = (1 - 1/\gamma)^{-1} c^{1-1/\gamma}$, $\gamma \neq 1$, it can be shown that $g_L(y, \xi) = \kappa_1(\xi)^{-\gamma} y$, where κ_1 is implicitly defined by

$$\kappa_1(\xi) = \left(1 + \delta^\gamma \left(\int_0^1 \eta^{1-1/\gamma} \kappa_1(\hat{\xi}(\cdot|\eta)) \left[\int_{\Theta} \phi(\eta|\theta) \xi(\theta) d\theta \right] d\eta \right)^\gamma \right)^{1/\gamma}. \quad (17)$$

While optimal consumption is still linear in output, the distribution of the shock is no longer characterized by a statistic, such as the mean in Mirman-Zilcha. Therefore, results with general distribution functions cannot be obtained. Future numerical research on this alternative case will further our understanding of learning in different economic structures.⁸

⁸We thank the Associate Editor for suggesting this alternative.

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),

$$V_L(y, \xi) = \max_{c \in (0, y)} \left\{ \ln c + \delta \ln(y - c) \int_0^1 \kappa_1(\hat{\xi}(\cdot|\eta)) \eta \left[\int_{\Theta} \phi(\eta|\theta) \xi(\theta) d\theta \right] d\eta + \delta \int_0^1 \kappa_2(\hat{\xi}(\cdot|\eta)) \left[\int_{\Theta} \phi(\eta|\theta) \xi(\theta) d\theta \right] d\eta \right\}. \quad (18)$$

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

$$g_L(y, \xi) = \left(1 + \delta \int_0^1 \kappa_1(\hat{\xi}(\cdot|\eta)) \eta \left[\int_{\Theta} \phi(\eta|\theta) \xi(\theta) d\theta \right] d\eta \right)^{-1} y. \quad (19)$$

Plugging (19) into (18) yields

$$V_L(y, \xi) = \left(1 + \delta \int_0^1 \kappa_1(\hat{\xi}(\cdot|\eta)) \eta \left[\int_{\Theta} \phi(\eta|\theta) \xi(\theta) d\theta \right] d\eta \right) \ln y + \kappa_3(\xi), \quad (20)$$

$$\equiv \kappa_1(\xi) \ln y + \kappa_2(\xi), \quad (21)$$

where κ_3 depends on ξ as well. Therefore,

$$\kappa_1(\xi) \equiv 1 + \delta \int_0^1 \kappa_1(\hat{\xi}(\cdot|\eta)) \eta \left[\int_{\Theta} \phi(\eta|\theta) \xi(\theta) d\theta \right] d\eta. \quad (22)$$

The solution to (22) is

$$\kappa_1(\xi) = \int_{\Theta} \frac{\xi(\theta) d\theta}{1 - \delta \mu(\theta)}, \quad (23)$$

where $\mu(\theta) = \int_0^1 \eta \phi(\eta|\theta) d\eta$. To verify that (23) is the solution to (22), updating (23) to the next period yields

$$\kappa_1 \left(\hat{\xi}(\cdot|\eta) \right) = \int_{\Theta} \frac{\hat{\xi}(\theta|\eta) d\theta}{1 - \delta\mu(\theta)}, \quad (24)$$

$$= \int_{\Theta} \frac{1}{1 - \delta\mu(\theta)} \frac{\phi(\eta|\theta)\xi(\theta)}{\int_{\Theta} \phi(\eta|x)\xi(x) dx} d\theta. \quad (25)$$

Then, plugging (25) into (22) yields (23), verifying the conjecture of the value function. Combining (22) and (23) yields

$$g_L(y, \xi) = \left(1 + \delta \int_0^1 \kappa_1 \left(\hat{\xi}(\cdot|\eta) \right) \eta \left[\int_{\Theta} \phi(\eta|\theta)\xi(\theta) d\theta \right] d\eta \right)^{-1} y, \quad (26)$$

$$= y/\kappa_1(\xi), \quad (27)$$

$$= \left(\int_{\Theta} \frac{\xi(\theta) d\theta}{1 - \delta\mu(\theta)} \right)^{-1} y. \quad (28)$$

Since both the utility and production functions are strictly concave in c , (28) is the unique maximizer corresponding to (18).

Proof of Proposition 4.1. Since $R''(x) = 2\delta/(1 - \delta x)^3 > 0$, the right-hand side of (14) is less than the right-hand side of (15) for any c , by Jensen's inequality. Therefore, $g_{AL}(y; \xi) > g_L(y, \xi)$.

Proof of Proposition 4.2. Suppose $\mu(\theta^*) = \int_{\Theta} \mu(\theta)\xi(\theta) d\theta$. First, $g_I(y; \theta^*) = g_{AL}(y; \xi)$ from (13) and (14). 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 $\theta^* = \int_{\Theta} \theta\xi(\theta) d\theta$. First, if $\mu'' < -2\delta\mu^2/(1 - \delta\mu)$, then, for any c , the right-hand side of (13) is greater than the right-hand side of (15) 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 (13) is greater than the right-hand side of (14), since $\mu(\theta^*) > \int_{\Theta} \mu(\theta)\xi(\theta) d\theta$ 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 (9), if $\phi^1 \succ_1 \phi^2$, then $\mu^1(\theta) = \int_0^1 \eta \phi^1(\eta|\theta) d\eta \geq \int_0^1 \eta \phi^2(\eta|\theta) d\eta = \mu^2(\theta)$ implying that $g_L^1(y, \xi) \leq g_L^2(y, \xi)$.

Proof of Proposition 5.4. From (9), if $\phi^1 \succ_2 \phi^2$, then $\mu^1(\theta) = \int_0^1 \eta \phi^1(\eta|\theta) d\eta = \int_0^1 \eta \phi^2(\eta|\theta) d\eta = \mu^2(\theta)$ 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 (15) is greater under ξ^1 than under ξ^2 . Therefore, $g_L(y, \xi^1) \leq g_L(y, \xi^2)$. 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 (15) is greater under ξ^1 than under ξ^2 , by Jensen's inequality. Therefore, $g_L(y, \xi^1) \leq g_L(y, \xi^2)$. The proofs for $\mu'' > -2\delta\mu^2/(1 - \delta\mu)$ and $\mu'' = -2\delta\mu^2/(1 - \delta\mu)$ are identical.

References

- P. Aghion, P. Bolton, C. Harris, and B. Jullien. Optimal Learning by Experimentation. *Review of Economic Studies*, 58(4):621–654, 1991.
- E. Antoniadou, C. Koulovatianos, and L.J. Mirman. Strategic Exploitation of a Common Property Resource under Uncertainty. Vienna Economic Papers 0703, University of Vienna, Department of Economics, 2007.
- R.J. Balvers and T.F. Cosimano. Actively Learning About Demand and the Dynamics of Price Adjustment. *Economic Journal*, 100(402):882–898, 1990.
- P. Beaudry and F. Portier. An Exploration into Pigou’s Theory of Cycles. *Journal of Monetary Economics*, 51(6):1183–1216, 2004.
- P. Beaudry and F. Portier. Stock Prices, News and Economic Fluctuations. *American Economic Review*, 96(4):1293–1307, 2006.
- P. Beaudry and F. Portier. When can Changes in Expectations cause Business Cycle Fluctuations. *Journal of Economic Theory*, 135(1):458–477, 2007.
- G.W. Beck and V. Wieland. Learning and Control in a Changing Economic Environment. *Journal of Economic Dynamics and Control*, 26(9):1359–1377, 2002.
- G. Bertocchi and M. Spagat. Growth under Uncertainty with Experimentation. *Journal of Economic Dynamics and Control*, 23(2):209–231, 1998.
- W.A. Brock and L.J. Mirman. Optimal Economic Growth and Uncertainty: The Discounted Case. *Journal of Economic Theory*, 4(3):479–513, 1972.
- D. Cass. Optimum Growth in an Aggregative Model of Capital Accumulation. *Review of Economic Studies*, 32(3):233–240, 1965.
- L. Christiano, R. Motto, and M. Rostagno. Monetary Policy and a Stock Market Boom-Bust Cycle. Available at

<http://www.ecb.int/events/pdf/conferences/cbc4/ChristianoRostagnoMotto.pdf>, 2006.

- A. Creane. Experimentation with Heteroskedastic Noise. *Economic Theory*, 4(2):275–286, 1994.
- M. Datta, L.J. Mirman, and E.E. Schlee. Optimal Experimentation in Signal-Dependent Decision Problems. *International Economic Review*, 43(2):577–607, 2002.
- W.D. Dechert, S.I. O’Donnell, and W.A. Brock. Bayes’ Learning of Unknown Parameters. *Journal of Difference Equations and Applications*, 13(2–3):121–133, 2007.
- J.H. Drèze and F. Modigliani. Consumption Decisions under Uncertainty. *Journal of Economic Theory*, 5(3):308–335, 1972.
- D. Easley and N.M. Kiefer. Controlling a Stochastic Process with Unknown Parameters. *Econometrica*, 56(5):1045–1064, 1988.
- D. Easley and N.M. Kiefer. Optimal Learning with Endogenous Data. *International Economic Review*, 30(4):963–978, 1989.
- M.A. El-Gamal and R.K. Sundaram. Bayesian Economists ... Bayesian Agents: An Alternative Approach to Optimal Learning. *Journal of Economic Dynamics and Control*, 17(3):355–383, 1993.
- G.W. Evans and S. Honkapohja. *Learning and Expectations in Macroeconomics*. Princeton University Press, 2001.
- A. Fishman and N. Gandalf. Experimentation and Learning with Network Effects. *Economics Letters*, 44(1–2):103–108, 1994.
- X. Freixas. Optimal Growth With Experimentation. *Journal of Economic Theory*, 24(2):296–309, 1981.
- J.M. Fusselman and L.J. Mirman. Experimental Consumption for a General Class of Disturbance Densities. In *General Equilibrium, Growth, and Trade II*, pages 367–392. Academy Press, Inc., 1993.

- S.J. Grossman, R.E. Kihlstrom, and L.J. Mirman. A Bayesian Approach to the Production of Information and Learning by Doing. *Review of Economic Studies*, 44(3):533–547, 1977.
- F.H. Hahn. Savings and Uncertainty. *Review of Economic Studies*, 37(1):21–24, 1970.
- G.W. Huffman and N.M. Kiefer. Optimal Learning in a One-Sector Growth Model. Working Papers 90-24, Cornell University, Center for Analytic Economics, 1994.
- M. Huggett. Precautionary Wealth Accumulation. *Review of Economic Studies*, 71(3):769–781, 2004.
- N. Jaimovich and S. Rebelo. Can News About the Future Drive the Business Cycle? NBER Working Papers 12537, National Bureau of Economic Research, Inc, 2006.
- G. Keller and S. Rady. Optimal Experimentation in a Changing World. *Review of Economic Studies*, 66(3):475–507, 1999.
- N.M. Kiefer and Y. Nyarko. Optimal Control of an Unknown Linear Process with Learning. *International Economic Review*, 30(3):571–586, 1989.
- T.C. Koopmans. Optimum Growth in an Aggregative Model of Capital Accumulation. In *The Economic Approach to Development Planning*, pages 225–287. North-Holland, 1965.
- H.E. Leland. Saving and Uncertainty: The Precautionary Demand for Saving. *Quarterly Journal of Economics*, 82(3):465–473, 1968.
- F. Milani. Expectations, Learning and Macroeconomic Persistence. *Journal of Monetary Economics*, 54(7):2065–2082, 2007.
- L.J. Mirman. *Two Essays on Uncertainty and Economics*. PhD thesis, University of Rochester, 1970.

- L.J. Mirman. Uncertainty and Optimal Consumption Decisions. *Econometrica*, 39(1):179–185, 1971.
- L.J. Mirman. On the Existence of Steady State Measures for One Sector Growth Models with Uncertain Technology. *International Economic Review*, 13(2):271–286, 1972.
- L.J. Mirman. The Steady State Behavior of a Class of One Sector Growth Models with Uncertain Technology. *Journal of Economic Theory*, 6(3):219–242, 1973.
- L.J. Mirman and I. Zilcha. On Optimal Growth under Uncertainty. *Journal of Economic Theory*, 11(3):329–339, 1975.
- L.J. Mirman, L. Samuelson, and A. Urbano. Monopoly Experimentation. *International Economic Review*, 34(3):549–563, 1993.
- E.C. Prescott. The Multi-Period Control Problem Under Uncertainty. *Econometrica*, 40(6):1043–1058, 1972.
- F.P. Ramsey. A Mathematical Theory of Saving. *Economic Journal*, 38(152):543–559, 1928.
- M. Rothschild and J.E. Stiglitz. Increasing Risk II: Its Economic Consequences. *Journal of Economic Theory*, 3(1):66–84, 1971.
- A. Sandmo. The Effect of Uncertainty on Saving Decisions. *Review of Economic Studies*, 37(3):353–360, 1970.
- D. Trefler. The Ignorant Monopolist: Optimal Learning with Endogenous Information. *International Economic Review*, 34(3):565–581, 1993.
- V. Wieland. Learning by Doing and the Value of Optimal Experimentation. *Journal of Economic Dynamics and Control*, 24(4):501–534, 2000.