Optimal Consumption and Asset Allocation with Unknown Income Growth

Neng Wang[a,*]

[a] Columbia University and National Bureau of Economic Research

Received Date; Received in Revised Form Date; Accepted Date

Abstract

Recent empirical evidence supports the view that the income process has an individual-specific growth rate component (Baker (1997), Guvenen (2007b), and Huggett, Ventura, and Yaron (2007)). Moreover, the individual-specific growth component may be stochastic. Motivated by these empirical observations, I study an individual's optimal consumption-saving and portfolio choice problem when he does not observe his income growth. As in standard income fluctuation problems, the individual cannot fully insure himself against income shocks. In addition to the standard income-risk-induced precautionary saving demand, the individual also has learning-induced precautionary saving demand, which is greater when belief is more uncertain. With constant unobserved income growth, changes in belief are not predictable. However, with stationary stochastic income growth, belief is no longer a martingale. Mean reversion of belief reduces hedging demand on average and in turn mitigates the impact of estimation risk on consumption-saving and portfolio decisions.

Keywords: Incomplete markets, precautionary saving, learning, hedging, estimation risk.

JEL classification: E2, G11, G31

1. Introduction

Intertemporal consumption-saving and portfolio allocation is a fundamental topic in modern economics. Almost all existing research on this topic assumes that the individual has complete information about the parameters of his income process such as the growth rate

*I thank Patrick Bolton, Hui Chen, Darrell Duffie, Jan Eberly, Marc Giannoni, Bob Hall, Yingcong Lan, Tom Sargent, Mike Woodford, Steve Zeldes, and seminar participants at NYU and Columbia for comments. I am especially grateful to the anonymous referee and Bob King (Editor) for many insightful comments, which have substantially improved the paper. I thank Larry Shampine for his expert advice on numerical solutions and Jinqiang Yang for superb research assistance. Address: Columbia Business School, 3022 Broadway, Uris Hall 812, New York, NY 10027. Email: neng.wang@columbia.edu.

and volatility. The complete-information assumption may be a sensible starting point, if one believes that an individual's income process can be represented by a growth component common to all individuals plus an idiosyncratic shock process (MaCurdy (1982) and Abowd and Card (1989)). Intuitively, if all individuals share the same income growth rate (for example, among individuals within the same education group), each individual can then estimate the "common" income growth by using panel data of income within his associated group.

However, Baker (1997), Guvenen (2007a), and Huggett, Ventura, and Yaron (2007) provide convincing empirical evidence in support of a competing view that the income process has an important *individual-specific* growth component. If income growth is individual specific, it then becomes much more difficult for each individual to estimate his own income growth. An individual enters the labor market with a prior belief about his future income growth and updates his belief over time based on realized incomes. Learning affects the individual's consumption-saving and portfolio allocation decisions through two channels: the *expected* growth of income and precautionary saving demand induced by the individual's learning of his income growth.

Income volatility has two effects on consumption and saving. The first is the standard precautionary saving demand against income fluctuations. Naturally, hedging demand against income risk is higher when the income stream is more volatile. More interestingly, unknown income growth also induces hedging demand, which is stochastic and depends on the agent's time-varying belief. For a given fixed spread between the two possible levels of income growth rates, estimation risk (induced by learning about his income growth) decreases with income volatility. This seemingly counter-intuitive result may be explained as follows. Past incomes from a more volatile income process provide less information about the unknown true income growth rate. Hence, the agent updates his belief less in response to unanticipated income innovations. Therefore, estimation risk is smaller when the underlying income process is more volatile, *ceteris paribus*. The net impact of income volatility on hedging demand depends on the relative magnitude of these two opposing effects.

Recent work also provides strong empirical evidence consistent with the hypothesis that income growth is stochastic. Haider (2001) and Guvenen and Kuruscu (2008) document that cross-sectional dispersion in income growth has been rising since the 1970s. The stochastic feature of income growth further enriches the agent's learning problem, but makes his consumption-saving and portfolio decisions more complicated. The potential empirical importance of stochastic income growth on decision rules and utility costs calls for models incorporating *stochastic* income growth. Intuitively, when income growth is stochastic and unknown, belief change is locally predictable (due to the expected change in income growth) and hence belief is no longer a martingale as in the case of constant growth. For example, when the conditional probability of income growth being low is small, mean reversion (due to the stochastic transition of income growth from low to high) pulls the agent's belief upward in expectation. Intuitively, this mean reversion of belief makes shocks driving the change of belief no longer permanent, unlike in settings with unknown *constant* income growth. The stationary belief updating process in turn lowers the impact of estimation risk on consumption. Therefore, consumption responds less to change in belief.

This paper contributes to the literature on incomplete-markets consumption, saving, and portfolio choice with learning. Earlier papers that explore the role of partially observed income on consumption include Goodfriend (1992), Pischke (1995), and a collection of papers in Hansen and Sargent (1991). All these studies postulate that the agent's consumption is given by the certainty-equivalence-based permanent-income hypothesis (PIH) rule (Friedman (1957)), which precludes any possible effect of estimation risk on consumption. The most closely related papers are Guvenen (2007b) and Wang (2004). Guvenen (2007b) solves for the consumption rule numerically for agents with constant relative risk-averse utility. His work complements this one in terms of methodology and economic insights. Unlike Wang (2004), in the present paper, learning has implications not only on income volatility, but also on expected changes in income. More importantly, the conditional variance of belief updating is *stochastic*. As a result, learning induces *stochastic* belief-dependent precautionary saving demand in this paper. Finally, unlike Guvenen (2007b) and Wang (2004), this paper also studies the effect of estimation risk on hedging and portfolio allocation.

2. Model Setup

Consider a consumption-saving and portfolio allocation problem. An infinitely-lived agent receives an exogenous perpetual stream of stochastic income. Let $y(t)$ denote the *level* of the agent's time-t labor income. Assume that the dynamics of $\{y(t) : t \geq 0\}$ is given by

$$dy(t) = (\alpha(t) - \kappa y(t))\, dt + \sigma dZ(t)\,, \tag{1}$$

where Z is a standard Brownian motion. The parameter σ measures the conditional volatility of the income change over an incremental unit of time. The income growth parameter $\{\alpha(t) : t \geq 0\}$ may change stochastically. The detailed specification for α is deferred to the next section. For convergence, assume $r + \kappa > 0$, i.e. income cannot grow too fast. When $\kappa = 0$, the income process (1) has a unit root (non-stationary). When $\kappa > 0$, the process given in (1) is stationary and is known as an Ornstein-Uhlenbeck process.[1] In this case, the parameter κ measures the degree of mean reversion. The discrete-time counterpart of (1) when $\kappa > 0$ is represented by the following first-order autoregressive (AR1) process:

$$y(t+1) = a_0(t) + a_1 y(t) + \hat{\sigma}\epsilon(t+1), \tag{2}$$

where $a_1 = e^{-\kappa}$, $a_0(t) = \alpha(t)\left(1 - e^{-\kappa}\right)/\kappa$, $\hat{\sigma} = \sigma\sqrt{\left(1 - e^{-2\kappa}\right)/(2\kappa)}$, and $\epsilon(t+1)$ is a time-$(t+1)$ innovation drawn from the standard normal distribution. The above AR1 process has been widely used to model income (Deaton (1992) and Attanasio (1999)). In the precautionary saving literature, Caballero (1991) uses a discrete-time unit-root process ($\kappa = 0$), a special case of (2) to model labor income and derives a closed-form consumption rule. Wang (2006) obtains the closed-form consumption rule and characterizes the stochastic precautionary saving demand for a class of the income process known as "affine" models nesting (1) as a special case.

[1] I thank the anonymous referee and Bob King (Editor) for the suggestion to extend the model specification for the income process to allow for mean reversion.

The agent can invest in both a risk-free asset (with a constant rate of return r) and a risky financial asset (e.g. the market portfolio). Investing in the risky asset offers the agent both the opportunity to earn a higher expected return than the risk-free rate r and the benefit of hedging labor-income related risk. The instantaneous return $dR(t)$ of the market portfolio over time increment dt is given by:

$$dR(t) = (r + \zeta)\, dt + \nu dW(t), \tag{3}$$

where ζ is the market risk premium, ν is the volatility of the market return, and W is a standard Brownian motion. Equation (3) specifies that the market return is independently and identically distributed (iid). Let ρ be the (instantaneous) correlation between the labor income process (1) and the return of the risky asset, i.e. the correlation between Brownian motions Z and W is ρ. Let $\eta = \zeta/\nu$ denote the Sharpe ratio of the market portfolio. Let $\psi(t)$ denote the amount of wealth that the agent allocates to the market portfolio at time t, and hence $x(t) - \psi(t)$ corresponds to time-t wealth invested in the risk-free asset. The agent's financial wealth dynamics is then given by

$$dx(t) = (rx(t) + y(t) - c(t))\, dt + \psi(t)\, (\zeta dt + \nu dW(t))\ , \tag{4}$$

where the first term in (4) gives the sum of interest income rx (if all wealth is invested in the risk-free asset) and labor income y minus consumption c. That is, the first term gives the saving rate $s = rx + y - c$ in standard self-insurance models if the agent can only invest in the risk-free asset. The last term, $\psi(t)\, (\zeta dt + \nu dW(t))$, captures the "excess" return by borrowing at the risk-free rate and investing in the risky asset.

The agent has a time-additive separable utility function given by

$$U(c) = \mathbb{E}\left(\int_0^\infty e^{-\beta s}\, u(c(s))\, ds\right)\ , \tag{5}$$

where $\beta > 0$ is his subjective discount rate and $\gamma > 0$ is the coefficient of absolute risk aversion (CARA), i.e. $u(c) = -e^{-\gamma c}/\gamma$. It is well known that CARA utility gives much tractability in deriving the consumption rule because it ignores the wealth effect. Merton (1971), Kimball

(1989), Caballero (1990, 1991), Svensson and Werner (1993), Davis and Willen (2000), and Wang (2003, 2006) have all adopted CARA utility in analyzing the agent's consumption-saving decisions under incomplete markets with different income process specifications. The agent chooses his consumption c and wealth allocation to the risky asset ψ to maximize his utility given in (5) subject to his stochastic labor-income process (1), his wealth accumulation process (4), and the transversality condition specified in the online appendix.

3. Model Analysis

In standard consumption-saving models, the agent knows both his income process and the parameters governing his income process, such as the growth parameter α. However, much empirical evidence suggests that the agent's income growth may be individual specific and hence the agent does not necessarily know his income growth parameter. Learning about income growth could potentially have a significant impact on the agent's intertemporal consumption-saving and portfolio allocation rules. Moreover, income growth α may change stochastically over time, further complicating the agent's decision problem.

To understand the effects of learning and the stochastic feature of the income growth on consumption-saving and portfolio allocation in an intuitive and pedagogical way, we categorize our model into four special sub-models along two dimensions: whether the agent knows the value of $\{\alpha(t) : t \geq 0\}$, and whether the agent's income growth $\{\alpha(t) : t \geq 0\}$ is stochastic. Table 1 summarizes the structure of the model development in this paper.

[Insert Table 1 here.]

Each special case will provide new insights on the effect of learning on precautionary saving. I start with the models where the agent knows the value of his income growth parameter α.

3.1. Models I and III: Known (but possibly stochastic) income growth

First, I describe the dynamics for the income growth $\{\alpha(t) : t \geq 0\}$, and then analyze the agent's optimality. The agent's income growth is often subject to both aggregate and idiosyncratic risks. One parsimonious way to capture the stochastic nature of income growth is to postulate that income growth $\{\alpha(t) : t \geq 0\}$ varies stochastically over time between α_1 and $\alpha_2 < \alpha_1$, the two possible levels.[2] Let $N(t)$ denote the regime for the agent's income growth $\alpha(t)$ at time t. That is, income growth $\alpha(t)$ takes the value $\alpha_{N(t)}$ for $N(t) = 1, 2$. Fix a small time period Δt. If the time-t growth rate is high, i.e. $(\alpha(t) = \alpha_1)$, the growth rate remains high at time $(t + \Delta t)$ with probability $(1 - \lambda_1 \Delta t)$ and decreases to α_2 at time $(t + \Delta t)$ with the remaining probability $\lambda_1 \Delta t$. Similarly, if the time-t growth rate is low, i.e. $(\alpha(t) = \alpha_2)$, $\lambda_2 \Delta t$ is the transition probability from low growth rate α_2 to α_1, the high value.[3] Now, I use dynamic programming to characterize the agent's optimality.

Wealth x, income y, and income growth are the three state variables. Let $V(x, y, n)$ denote the value function when the income growth rate is α_n, where $n = 1, 2$. When the current income growth is high $(\alpha(t) = \alpha_1)$, the agent's Hamilton-Jacobi-Bellman (HJB) equation is given by

$$\begin{aligned}
\beta V(x, y, 1) = \max_{c,\,\psi}\; & u(c) + (rx + \psi\zeta + y - c)\, V_x(x, y, 1) + (\alpha_1 - \kappa y)\, V_y(x, y, 1) \\
& + \frac{\psi^2 \nu^2}{2} V_{xx}(x, y, 1) + \psi\rho\nu\sigma V_{xy}(x, y, 1) + \frac{\sigma^2}{2} V_{yy}(x, y, 1) \\
& + \lambda_1 \left(V(x, y, 2) - V(x, y, 1) \right). \qquad (6)
\end{aligned}$$

The left side of (6) is the "flow" value of the agent's value function. The right side of (6) is the sum of his utility flow $u(c)$ and the instantaneous expected changes in his value function. Optimality of consumption and portfolio rules implies that the two sides of (6) are equated. The V_x term describes the marginal increase of the agent's value function from saving. The V_y term captures the marginal increase of the agent's value function if income y increases by

[2] Kimball (1989) assume that the "level" of income, rather than the growth rate (drift) of the income process, stochastically switches between two states in their study of the impact of precautionary saving on Ricardian (tax) equivalence. Indeed, the income process in Kimball (1989) is a special case of (1).

[3] For the continuous-time regime switching model, the implied stationary probabilities for high-income-growth and low-income-growth states are $\lambda_2/(\lambda_1 + \lambda_2)$, and $\lambda_1/(\lambda_1 + \lambda_2)$, respectively. See Hamilton (1989) for an early and important application of regime-switching models to economics and econometric analysis in discrete time.

a unit. The V_{xx}, V_{xy}, and V_{yy} terms reflect the effects of stochastic return, income volatility and their correlation on the agent's value function. The last term captures the effect of stochastic transition of his income growth rate on the expected change in his value function. Note that the value function changes discretely from $V(x, y, 1)$ to $V(x, y, 2)$ when the growth rate changes.

The first-order condition (FOC) for consumption is $u'(c) = V_x(x, y, 1)$. That is, the marginal utility of consumption $u'(c)$ is equal to the marginal value of wealth V_x. The FOC with respect to the portfolio rule ψ gives

$$\psi = -\frac{\zeta V_x(x, y, 1)}{\nu^2 V_{xx}(x, y, 1)} - \frac{\rho\sigma V_{xy}(x, y, 1)}{\nu V_{xx}(x, y, 1)}, \tag{7}$$

where the first term captures the standard risk-return tradeoff from investing in the market portfolio, and the second term reflects the agent's motive to hedge against labor income shocks. Next, we incorporate the effect of learning.

3.2. Models II and IV: Unknown (but possibly stochastic) income growth

If the agent does not know his income growth, his time-t information set $\mathcal{F}_t$ then only contains the history of his past incomes $\{y(s) : s \leq t\}$, not the true (but possibly stochastic) value of $\alpha(t)$. Let $p(t)$ denote his time-t belief that the growth rate is high (i.e. $\alpha(t) = \alpha_1$), in that $p(t) = \text{Prob}\left(\alpha(t) = \alpha_1 | \mathcal{F}_t\right)$. Let μ denote the expected growth rate of the income process. By definition, the expected growth rate μ is a weighted average of the two possible income growth rates, in that

$$\mu(t) = \mathbb{E}_t\left(\alpha(t)\right) = p(t)\alpha_1 + \left(1 - p(t)\right)\alpha_2 = \alpha_2 + \delta p(t), \tag{8}$$

where

$$\delta = \alpha_1 - \alpha_2 \tag{9}$$

is the difference between the two possible values of α. For a given small time period $(t, t+\Delta t)$, the change in income is $(y(t+\Delta t) - y(t))$. Out of this total change, $\left(\mu(t) - \kappa y(t)\right)\Delta t$ is the expected change. The unanticipated change is given by $\left(y(t + \Delta t) - y(t) - \left(\mu(t) - \kappa y(t)\right)\Delta t\right)$.

Scaling by volatility $\sigma\sqrt{\Delta t}$ and taking the limit as $\Delta t \to 0$, we may construct a "new" Brownian motion process B as follows:

$$dB(t) = \left(dy(t) - \left(\mu(t) - \kappa y(t)\right) dt\right) / \sigma \,, \tag{10}$$

which will serve as the innovations process for belief updating.

Re-writing the innovations process (10), we have the following innovations-based representation of the income process (1):

$$dy(t) = \left(\alpha_2 + \delta p(t) - \kappa y(t)\right) dt + \sigma dB(t), \tag{11}$$

where we use (8) for $\mu(t)$. Since the agent only observes his past income, the innovations representation (11) will naturally be useful when we derive the agent's optimal consumption and portfolio rules. Using the results in Liptser and Shiryayev (1977), we can write the belief process as follows:

$$dp(t) = \left(\lambda_2 - (\lambda_1 + \lambda_2)\, p(t)\right) dt + \sigma^{-1}\delta p(t) \left(1 - p(t)\right) dB(t), \tag{12}$$

where B is given in (10). Note that belief p and income y are perfectly correlated (one shock model). We defer the economic interpretations of (12) to later sections when discussing model intuition.

When the income growth rate α is unknown, the optimization problem is not Markovian with respect to the original information set $\mathcal{F}_t$, which only contains the history of income y. The belief updating process (12) and the innovations-representation process (11) for income y jointly convey the same information as the agent's original income process (1) and his prior belief about his income growth do. We can transform the original non-Markovian optimization problem into a Markovian one. That is, the agent maximizes his utility function (5), subject to the innovations-based representation of his income process (11), his belief updating process (12), his wealth accumulation equation (4), and the standard transversality condition given in the online appendix.[4] Importantly, the agent's learning about his income

[4] Gennotte (1986) and Xia (2001) study the optimal asset allocation when the agent has incomplete information about his investment opportunities, such as the dividend growth rate or the expected stock return. See Detemple (1986), Wang (1993), and Veronesi (1999) for equilibrium asset pricing implications of learning.

growth implies that belief p is also a state variable in addition to wealth x and income y.

There are three state variables for the agent's optimization problem: wealth x, income y, and belief p that income growth α is high. The HJB equation for the agent's value function $J(x, y, p)$ is given as follows:

$$\begin{aligned} \beta J = \max_{c,\psi}\ & u(c) + (rx + \psi\zeta + y - c)\, J_x + \frac{\psi^2\nu^2}{2} J_{xx} + (\alpha_2 + \delta p - \kappa y)\, J_y \\ & + \psi\rho\nu\sigma J_{xy} + (\lambda_2 - (\lambda_1 + \lambda_2)\, p)\, J_p + \frac{\delta^2}{2\sigma^2} p^2 (1-p)^2 J_{pp} \\ & + \psi\rho\nu\sigma^{-1}\delta p\,(1-p)\, J_{xp} + \delta p\,(1-p)\, J_{yp} + \frac{1}{2}\sigma^2 J_{yy}. \end{aligned} \tag{13}$$

The left side of (13) is the annuity (flow measure) of his value function. As for the HJB equation (6) when the income growth α is known, the right side includes standard terms, such as J_x, J_y, J_{xx}, J_{yy}, and J_{xy}. Unlike the HJB equation (6) for the case with known income growth, the agent's learning has additional effects on decision making. For example, J_p and J_{pp} terms capture the effects of the agent's belief about income growth on his value function. Since belief updating is solely driven by realized incomes, the agent's income process is perfectly correlated with his belief updating, as reflected in the J_{yp} term in (13). Finally, the J_{xp} term captures the agent's hedging demand induced by his estimation risk (associated with belief updating).

4. Model I: Known and constant growth parameter α

I first solve the model and then discuss its economic implications.

Model Solution. First, consider the setting where $\alpha(t)$ is *known* and is constant over time. This is the standard consumption-saving and portfolio allocation problem for an agent endowed with uninsurable stochastic labor income. The transition probability out of the current income growth α is zero. That is, $\alpha(t) = \alpha$ for all t. The following proposition summarizes the main results on consumption and portfolio rules for this setting, dubbed as Model I:

Proposition 1 *If the agent knows his constant income growth α, his consumption c^* and wealth allocation to the risky asset ψ^* are given by*

$$c^*(t) = r\left(x(t) + g(y(t);\alpha)\right), \tag{14}$$

$$\psi^*(t) = \frac{\zeta}{\gamma r \nu^2} - \frac{\rho\sigma}{\nu}\frac{1}{r+\kappa}, \tag{15}$$

where the risk-adjusted certainty equivalent human wealth $g(y;\alpha)$ is given by

$$g(y;\alpha) = \frac{1}{r+\kappa}\left(y + \frac{\alpha - \rho\sigma\eta}{r}\right) - \frac{\gamma\left(1-\rho^2\right)\sigma^2}{2(r+\kappa)^2} + \frac{\beta - r}{\gamma r^2} + \frac{\eta^2}{2\gamma r^2}. \tag{16}$$

Model intuition and implications. The agent's investment opportunity in the risky asset has two effects. First, the risky asset offers a higher expected return and hence shall raise the forward-looking agent's current consumption (Merton (1971)). This effect is captured by the first term in the agent's portfolio rule (15), and also by the constant positive term $\eta^2/(2\gamma r^2)$ in the agent's risk-adjusted certainty equivalent wealth $g(y)$ given in (16). Second, investing in the risky asset allows the agent to partially hedge against his labor income risk (i.e. the second term in the portfolio allocation rule (15)). The agent has a higher hedging demand, if the systematic volatility $\rho\sigma$ is larger. Hedging changes the agent's labor income growth from α to $(\alpha - \rho\sigma\eta)$, and also reduces the non-diversifiable component of his labor income volatility from σ to $\sigma\sqrt{1-\rho^2}$. Since precautionary saving demand arises from the agent's non-diversifiable idiosyncratic risk, hedging lowers the agent's precautionary saving demand. We measure the agent's precautionary saving demand as the amount by which the certainty equivalent wealth $g(y)$ is lower than the corresponding certainty equivalent wealth under the PIH rule. It is immediate to see that the precautionary saving demand only depends on idiosyncratic volatility $\sigma\sqrt{1-\rho^2}$ and is given by

$$\pi(t) = \frac{\gamma\left(1-\rho^2\right)\sigma^2}{2\left(r+\kappa\right)^2}. \tag{17}$$

The more persistent income shocks are (lower κ), the greater the agent's precautionary saving demand π is, which is consistent with our analysis on hedging demand ψ. If labor income is perfectly correlated with the risky asset return, the agent can fully hedge his income risk. As a result, his precautionary saving demand is zero (i.e. complete markets setting). Finally, Model I nests Caballero (1991) and Wang (2006), settings where the agent cannot invest in the risky asset ($\psi(t) = 0$), as special cases.[5]

To sum up, Model I generates the following empirically testable predictions. First, precautionary saving demand is higher for the more persistent income process. Second, for incomes more correlated with the market portfolio, hedging demand is higher, and hence precautionary saving demand is lower.

Next, I analyze the case where income growth is constant but unknown.

5. Model II: Unknown and constant growth parameter α

First, I analyze the agent's Bayesian learning problem and then use dynamic programming to solve for his decision rules. Finally, I highlight the model implications on learning-induced precautionary saving.

Model Solution. When the agent does not know his income growth, he needs to use his past realized incomes to estimate the likelihood that his income growth α is high. Note that Model II is a special case of the general learning model of Section 3.2. with $\lambda_1 = \lambda_2 = 0$. We may write the belief updating process (12) as follows:

$$dp(t) = \sigma^{-1}\delta p(t)\,(1 - p(t))\,dB(t), \tag{18}$$

where B is the Brownian motion process under the *innovations* representation given in (10). Note that the belief updating process (18) is a martingale. This is because the unknown

[5] Wang (2006) extends the discrete-time CARA-Gaussian formulation of Caballero (1991) in a continuous-time setting to allow for conditionally heteroskedastic labor income process. The key advantage of introducing conditional heteroskedaticity of labor income process is that the agent's marginal propensity to consume (MPC) out of labor income may be less than the MPC out of financial wealth, a desirable feature argued in Friedman (1957) , Hall (1978) , and Zeldes (1989). Using isoelastic-utility-based buffer-stock-type saving models, Deaton (1991), Carroll (1997), and Gourinchas and Parker (2002) also generate this desirable feature on the MPC.

growth rate α is *constant*, and hence the change in the agent's belief must be unpredictable, in that $\mu(t) = \mathbb{E}_t\left(\alpha(t)\right) = \mathbb{E}_t\left(\mathbb{E}_s\left(\alpha(t)\right)\right) = \mathbb{E}_t\left(\mu(s)\right)$, for any $t < s$.

The instantaneous volatility of belief updating is symmetric in p and $(1-p)$ because the unobserved growth rate can only take two possible values: α_1 and α_2. The greater the wedge $\delta = \alpha_1 - \alpha_2 > 0$ is, the more volatile belief updating is. Moreover, a higher income volatility σ implies a less volatile belief updating. Intuitively, a higher realized value of income is more informative about the unknown income growth if the income process is less volatile (lower σ). The following proposition summarizes the results on consumption and portfolio rules when the agent learns about his (constant) income growth.

Proposition 2 *If income growth α is constant but unknown to the agent, his consumption c^* and wealth allocation to the risky asset ψ^* are given by*

$$c^*(t) = r\left(x(t) + g(y(t);\alpha_2) + f(p(t))\right), \tag{19}$$

$$\psi^*(t) = \frac{\zeta}{\gamma r \nu^2} - \frac{\rho}{\nu}\frac{\sigma}{r+\kappa} - \xi(t), \tag{20}$$

where $g(y;\alpha_2)$ is given by (16), learning-induced hedging demand ξ is

$$\xi(t) = \frac{\rho}{\nu}\frac{\delta p(t)\left(1-p(t)\right)f'(p(t))}{\sigma}, \tag{21}$$

and $\{f(p) : 0 \le p \le 1\}$ solves the following non-linear ODE:

$$\begin{aligned} rf(p) = {} & \frac{\delta p}{r+\kappa} - \frac{\rho\eta}{\sigma}\delta p\left(1-p\right)f'(p) + \frac{\delta^2}{2\sigma^2}p^2\left(1-p\right)^2 f''(p) \\ & - \gamma r\left(1-\rho^2\right)\left[\frac{1}{r+\kappa}\delta p\left(1-p\right)f'(p) + \frac{\delta^2}{2\sigma^2}p^2\left(1-p\right)^2 f'^2\right], \end{aligned} \tag{22}$$

subject to $f(0) = 0$ and $f(1) = \delta/(r(r+\kappa))$.

Model implications: Learning, precautionary saving, and hedging. When the agent learns about his income growth $\{\alpha(t) : t \ge 0\}$, his certainty equivalent wealth has an additional term $f(p)$, which depends on p. If income growth is low with probability one at all times (i.e. $\alpha(t) = \alpha_2$ for all t), we are back to Model I with constant known income growth α_2 with $f(0) = 0$. Similarly, if income growth α is always high (i.e. $\alpha(t) = \alpha_1$ for all t), we have $f(1) = \delta/(r(r+\kappa))$. Note that $p = 0$ and $p = 1$ are absorbing states here. For $0 < p < 1$, we need to solve for $f(p)$ given in (22).

Learning has two effects: the expected growth rate of income and the precautionary saving demand effect. In the following analysis, we separate these two effects on $f(p)$. Since the *expected* growth rate of income exists for any agent, we first solve for $f(p)$ for risk-neutral agents ($\gamma = 0$). Let $\bar{f}(p)$ denote the certainty equivalent wealth satisfying (22) for risk-neutral agents (i.e. $\gamma = 0$). Intuitively, $\bar{f}(p)$ captures the learning effect via the channel of the *expected* growth rate of income. Let $l(p) = \bar{f}(p) - f(p)$, where $f(p)$ solves (22) for a given $\gamma \geq 0$. The wedge $l(p)$ captures the second effect: learning-induced precautionary saving demand.

Consider the special setting with $\rho = 0$, which includes the standard self-insurance (income fluctuation) problem (with risk-free asset only) as a special case. When $\rho = 0$, the solution for $\bar{f}(p)$ is linear and is given by

$$\bar{f}(p) = \frac{\delta p}{r(r+\kappa)}. \tag{23}$$

Figure 1 plots the "risk-adjusted" certainty equivalent wealth $f(p)$ and the wedge $l(p) = \bar{f}(p) - f(p)$ for $\gamma = 0, 1, 2$. We use the following (annualized and continuously compounded) parameters for the remainder of the paper unless otherwise noted. The interest rate $r = 4\%$, the dispersion of income growth $\delta = \alpha_1 - \alpha_2 = 3\%$, income volatility $\sigma = 40\%$, and the degree of income mean reversion $\kappa = 5\%$.

[Insert Figure 1 here.]

The left panel of Figure 1 shows that $f(p)$ is increasing in p for a given γ. Moreover, $f(p)$ decreases with the coefficient of absolute risk aversion γ. The right panel plots the learning-induced precautionary saving $l(p) = \bar{f}(p) - f(p)$. Note that $l(0) = l(1) = 0$ because there is no more uncertainty if $p = 0, 1$ under constant but unknown growth (Model II).

Learning-induced precautionary saving $l(p)$ is concave in p. Intuitively, when the agent is more uncertain about his income growth (i.e. in the interior region of p), learning induced precautionary saving $l(p)$ is higher, *ceteris paribus*. However, note that $l(p)$ is not symmetric around $p = 1/2$, and is rather skewed. This is due to the fact that $f(p)$ is convex, (i.e. $f'(1-p) > f'(p)$ for $0 < p < 1/2$), and the fact that $l(p)$ depends both on $p(1-p)$ and $f'(p)$. The nonlinear term in ODE (22) and the right panel of Figure 1 capture this asymmetry.

Now I consider the effect of learning when the agent can invest in the risky asset. As in Model I and Merton (1971), the agent earns a higher expected return, (i.e. the first term $\zeta/(\gamma r \nu^2)$ in (20)), hedges the systematic component of his labor income risk (the second term in (20)), and also hedges the correlated component of his income growth risk (the $\xi(t)$ term given in (21)). The hedging demand with respect to the labor income risk is the same as in our Model I, Svensson and Werner (1993), and Davis and Willen (2000). The hedging demand with respect to the estimation risk is however, stochastic, and depends on the time-varying volatility $\sigma^{-1}\delta p(1-p)$ of the belief updating process (18) and $f'(p)$, which measures the sensitivity of $f(p)$ with respect to belief p.

Income volatility σ has two *opposite* effects on the total hedging demand. On the one hand, a higher income volatility σ increases the hedging demand of labor income risk. On the other hand, incomes from a more volatile income process provide less precise information about the unknown income growth α given a fixed dispersion $\delta = \alpha_1 - \alpha_2$. Hence, the agent updates his belief less in response to "unexpected" income news. Therefore, a higher income volatility σ maps to a lower estimation risk and a lower hedging demand against estimation risk, *ceteris paribus*.

[Insert Figure 2 here.]

The left panel of Figure 2 plots the certainty equivalent wealth $f(p)$ for three levels of income volatility $\sigma = 0.1, 0.2, 0.3$ with correlation $\rho = 0.5$ (the other parameters are the same as those for Figure 1). Note that $f(p)$ *increases* with income volatility σ for $\rho > 0$, whereas $g(y; \alpha_2)$ decreases with σ. If $\rho < 0$, the opposite holds: $f(p)$ increases with income volatility

σ and is concave in p. The right panel of Figure 2 plots learning-induced hedging demand $\xi(p)$. Note that hedging demand $\xi(p)$ *decreases* with income volatility σ. This seemingly counter-intuitive result is due to the assumption that income changes are less informative for a given $\delta = \alpha_1 - \alpha_2$ (recall that the conditional volatility of income changes (in terms of levels) is constant in our model.) However, in reality, income volatility may increase with the income level. If so, then estimation risk will also increase with income volatility. In that case, the estimation-risk-induced precautionary saving and the standard income risk effect on precautionary saving may potentially move in the same direction, as shown by Guvenen (2007b). In that paper, the logarithmic income process is conditionally homoskedastic, which implies that the volatility of income increases with the level of income and hence estimation risk may increase with the level of income.

Empirically, an interesting and testable prediction is the effect of income risks on the precautionary saving demand induced by estimating income growth. We have highlighted a potential mechanism which makes the estimation risk lower when income is riskier. Again, we need to carefully control for the level effect of income growth estimation risk on precautionary saving demand. Having analyzed the impact of learning when income growth is constant, we now turn to the more general setting where income growth is stochastic.

6. Stochastic income growth

I first solve the model with stochastic income growth but without learning as a benchmark. Then, I solve the model with learning and interpret the economics of learning about stochastic income growth.

Model III: Known and stochastic growth parameter α.

Recall that income growth is given by a regime-switching model. While stochastic, income growth is known to the agent. The agent's information set $\mathcal{F}_t$ includes $\{N(s) : s \leq t\}$, where $N(t) = 1, 2$ correspond to the high and the low income growth rates, respectively. The following proposition summarizes the main results of Model III.

Proposition 3 *When the agent knows income growth* $\{\alpha(t) : t \geq 0\}$*, his portfolio allocation is given by (15), and his consumption is given by*

$$c^*(t) = r\left(x(t) + g(y(t); \alpha_2) + \phi_{N(t)}\right), \tag{24}$$

where $g(y; \alpha_2)$ *is given in (16) and* $\{\phi_1, \phi_2\}$ *jointly solve*

$$r\phi_1 = -\frac{\lambda_1}{\gamma r}\left(e^{-\gamma r(\phi_2 - \phi_1)} - 1\right) + \frac{\delta}{r+\kappa}, \tag{25}$$

$$r\phi_2 = -\frac{\lambda_2}{\gamma r}\left(e^{-\gamma r(\phi_1 - \phi_2)} - 1\right). \tag{26}$$

The portfolio rule is also given by (15), the same as in Model I. Equations (25) and (26) jointly characterize the growth-dependent consumption profiles: ϕ_1 and ϕ_2. Compared with Model I, the *stochastic growth* feature of income induces additional precautionary saving demand.

Model IV: Unknown and stochastic growth parameter α.

Without observing his income growth α, the agent uses his past incomes to estimate the likelihood that his income growth is high. Equation (12) gives the belief updating process. The intuition for the volatility specification in (12) is the same as the one for the belief process (18) in Model II, a special case of Model IV. See discussions on volatility in Section 5 for Model II. The intuition for the drift specification in (12) is richer than the one for Model II. Because the underlying unknown income growth α is *stochastic*, the expected change in the agent's belief is no longer zero, unlike Model II in Section 5. Consider a small time period $(t, t + \Delta t)$. Suppose the current income growth is high (i.e. $\alpha(t) = \alpha_1$). The conditional probability that income growth changes from α_1 to α_2 is $\lambda_1 \Delta t$. The size of this change is $\alpha_2 - \alpha_1 = -\delta$. Therefore, the expected change in income growth (conditional on $\alpha(t) = \alpha_1$) is $-\delta\lambda_1\Delta t$. The time-t probability that $\alpha(t) = \alpha_1$ is $p(t) = \text{Prob}_t(\alpha(t) = \alpha_1)$.

The unconditional expected change in income growth is thus given by

$$\mathbb{E}_t\left(\mu(t+\Delta t)-\mu(t)\right) = -p(t)\delta\lambda_1\Delta t + (1-p(t))\,\delta\lambda_2\Delta t \tag{27}$$

$$= \delta\left[\lambda_2 - (\lambda_1+\lambda_2)\,p(t)\right]\Delta t. \tag{28}$$

Using $\mu(t) = \alpha_2 + \delta p(t)$, we can show that the drift of $p(t)$ is equal to $(\lambda_2 - (\lambda_1+\lambda_2)\,p(t))$. The above analysis on the drift and the analysis on volatility in the previous section jointly provide an economically intuitive explanation for the Bayesian updating rule (12). When the belief $p(t)$ is larger than the (unconditional) long-run probability $\lambda_2/(\lambda_1+\lambda_2)$, the belief $p(t)$ is expected to move downward on average. This reflects the mean reversion property of the belief process $\{p(t) : t \geq 0\}$. Using this belief updating rule, we solve the agent's decision problem and summarize the results in the following proposition.

Proposition 4 *When the agent does not know his stochastic income growth α, his consumption c^* and wealth allocation ψ^* are given by (19) and (20), respectively, where $g(y;\alpha_2)$ is given in (16), and $\{f(p) : 0 \leq p \leq 1\}$ solves*

$$\begin{aligned} rf(p) = {} & \frac{\delta p}{r+\kappa} + \left[(\lambda_2 - (\lambda_1+\lambda_2)\,p) - \left(\frac{\rho\eta}{\sigma} + \gamma\frac{r\left(1-\rho^2\right)}{r+\kappa}\right)\delta p\,(1-p)\right] f'(p) \\ & + \frac{\delta^2}{2\sigma^2}p^2\,(1-p)^2\,f''(p) - \frac{\gamma r\left(1-\rho^2\right)}{2\sigma^2}\delta^2 p^2\,(1-p)^2\,f'^2, \end{aligned} \tag{29}$$

subject to the following boundary conditions:

$$rf(0) = \lambda_2 f'(0), \tag{30}$$

$$rf(1) = \frac{\delta}{r+\kappa} - \lambda_1 f'(1). \tag{31}$$

Results, intuition, and implications.

Unlike Model II (with constant growth), the states $p = 0$ and $p = 1$ in Model IV are no longer absorbing. The intuition is as follows. With stochastic growth, income growth may become low with probability $\lambda_1 \Delta t$ over time increment $(t, t+\Delta t)$, even if the agent knows for sure his income growth is high at time t. Therefore, learning induces precautionary saving demand at all times for any levels of belief p, provided that income growth is stochastic.

I now turn to the impact of income growth persistence (λ_1, λ_2) on consumption-saving decisions. Recall that learning-induced precautionary saving $l(p)$ is given by the difference between the certainty equivalent wealth $f(p)$ and $\bar{f}(p)$, where $\bar{f}(p)$ solves (29) for $\gamma = 0$, i.e. $l(p) = \bar{f}(p) - f(p)$. The left and the right panels of Figure 3 plot the certainty equivalent wealth $f(p)$ and learning-induced precautionary saving $l(p)$, respectively. In addition to the parameters we have used in Figures 1 and 2, we set market (portfolio) risk premium $\zeta = 6\%$ and market (portfolio) return volatility $\nu = 20\%$.

[Insert Figure 3 here.]

On the left panel of Figure 3, we see that either a higher value of λ_1 or of λ_2 makes $f(p)$ flatter. Of course, increasing λ_2, the transition probability from low income growth to high income growth, on average makes income growth higher and hence $f(p)$ larger. For example, comparing the setting of $(\lambda_1, \lambda_2) = (0, 0)$ with that of $(\lambda_1, \lambda_2) = (0, 3\%)$, we see that $f(p)$ is higher when $\lambda_2 = 3\%$ than when $\lambda_2 = 0$. Moreover, with $\lambda_1 = 0$, the high-income-growth state is absorbing, and $f(1) = \delta/(r(r + \kappa))$ for both settings as we see from the figure.

The right panel of Figure 3 plots learning-induced precautionary saving demand $l(p)$. Intuitively, with stochastic growth, income growth is more transitory and hence precautionary saving demand is lower when belief in the interior region, *ceteris paribus*. Comparing the setting of $(\lambda_1, \lambda_2) = (0, 0)$ with the setting of $(\lambda_1, \lambda_2) = (3\%, 3\%)$, we see that learning-induced precautionary saving demand is higher with constant income growth (i.e. $(\lambda_1, \lambda_2) = (0, 0)$) other than near the boundaries $p = 0$ and $p = 1$.

Figure 4 plots learning-induced hedging demand ξ as a function of belief p. The left panel analyzes the impact of transition intensities (λ_1, λ_2) on $f(p)$. As in Model II, learning-induced hedging demand ξ is concave in p. More interestingly, hedging demand ξ decreases with λ. Intuitively, the more transitive the income growth is, the lower learning-induced hedging demand ξ is. The right panel of Figure 4 shows that learning-induced hedging demand ξ decreases with income volatility σ. This result is consistent with the one for Model II. Intuitively, the same realized changes in income from a less volatile stream are more informative and reflect more changes in income growth. Hence, learning-induced hedging demand is higher when income streams are less volatile.

[Insert Figure 4 here.]

When income growth is stochastic and unknown, the change in beliefs is locally predictable. Mean reverting beliefs imply that shocks are no longer permanent, unlike in settings with unknown *constant* income growth. The stationarity of belief implies that the agent's learning process is less volatile and the estimation risk is lower. Therefore, learning-induced precautionary saving demand is lower and consumption responds less to the change in belief, when unknown income growth is stochastic.

7. Conclusions

In this paper, I study the effect of learning about income growth on an individual's consumption-saving and portfolio choice decisions when he cannot fully diversify his labor-income risk. The individual uses the dynamic Bayesian rule to update his belief about his income growth. Estimation risk naturally arises from his learning process. Importantly, this estimation risk generates additional precautionary savings demand beyond the standard income-risk-induced precautionary savings. By investing in the risky asset, the agent partially hedges against both income risk and estimation risk. While higher income volatility induces greater hedging demand against income shocks, higher income volatility (for a fixed income growth wedge) also induces less volatile belief updating because realized income is less informative about unknown income growth. Hence, higher income volatility implies lower estimation risk, which in turn suggests a smaller learning-induced hedging demand.

When income growth is stochastic and unknown, the agent's learning about income growth becomes even less volatile. Intuitively, the change in beliefs is locally predictable due to the expected change in income growth and hence beliefs are no longer martingales. Mean reversion of beliefs makes shocks driving the change in beliefs more transitory, unlike in settings with unknown *constant* income growth. The stationary belief updating process in turn lowers the impact of estimation risk on consumption. Therefore, when beliefs are not extreme (i.e at corners $p = 0, 1$), learning-induced precautionary saving demand is lower and consumption responds less to belief change, when unknown income growth is stochastic rather than constant.

The main objective of this paper is to study the effects of incomplete information about the income growth rate on his consumption and portfolio allocations, when the agent's income shocks are not insurable. In order to deliver this intuition in a simplest possible way, I have intentionally chosen the CARA utility for technical convenience. While analytically convenient, this utility specification ignores the wealth effect on consumption and portfolio allocation rules. The natural next step is to extend the analysis to settings with iso-elastic utility, which capture the wealth effect and hence allow for making quantitative assessments on the role of learning about income growth.

References

Abowd, J., Card, D. A., 1989. On the covariance structure of earnings and hours changes. Econometrica 57, 411-445.

Attanasio, O. P., 1999. Consumption. In J. B. Taylor and M. Woodford (Eds.), Handbook of Macroeconomics, Volume 1B. Elsevier, Amsterdam, pp. 731-803.

Baker, M., 1997. Growth-rate heterogeneity and the covariance structure of life-cycle earnings. Journal of Labor Economics 15, 338-375.

Caballero, R. J., 1990. Consumption puzzles and precautionary savings. Journal of Monetary Economics 25 (1), 113-136.

Caballero, R. J., 1991. Earnings uncertainty and aggregate wealth accumulation. American Economic Review 81 (4), 859-871.

Carroll, C. D., 1997. Buffer-stock saving and the life cycle/permanent income hypothesis. Quarterly Journal of Economics 112 (1), 1-55.

Davis, S. J., Willen, P., 2000. Using financial assets to hedge labor income risks: Estimating the benefits. Manuscript, University of Chicago.

Deaton, A., 1991. Saving and liquidity constraints. Econometrica 59 (5), 1221-1248.

Deaton, A., 1992. Understanding Consumption. Oxford University Press, Oxford.

Detemple, J., 1986. Asset pricing in a production economy with incomplete information. Journal of Finance 41, 383-391.

Friedman, M., 1957. A Theory of the Consumption Function. Princeton University Press, Princeton.

Gennotte, G., 1986. Optimal portfolio choice under incomplete information. Journal of Finance 41, 733-749.

Goodfriend, M., 1992. Information-Aggregation bias. American Economic Review 82, 508-519.

Gourinchas, P.-O., Parker, J. A., 2002. Consumption over the life cycle. Econometrica 70, 47-89.

Guvenen, F., 2007a. An empirical investigation of labor income processes. University of Minnesota Working Paper.

Guvenen, F., 2007b. Learning your earning: Are labor income shocks really very persistent? American Economic Review 97, 687-712.

Guvenen, F., Kuruscu, B., 2008. A quantitative analysis of the evolution of the U.S. wage distribution: 1970-2000. University of Minnesota Working Paper.

Haider, S. J., 2001. Earnings instability and earnings inequality of males in the united states: 1967-1991. Journal of Labor Economics 19, 799-836.

Hall, R. E., 1978. Stochastic implications of the life cycle-permanent income hypothesis: Theory and evidence. Journal of Political Economy 91 (6), 249-265.

Hamilton, J. D., 1989. A new approach to the economic analysis of nonstationary time series and business cycle. Econometrica 57, 357-384.

Hansen, L. P., Sargent, T. J., 1991. Rational Expectations Econometrics. Westview Press, Boulder, Colorado.

Huggett, M., Ventura, G., Yaron, A., 2007. Sources of life-cycle inequality. University of

Pennsylvania Working Paper.
Kimball, M. S., Mankiw, N. G., 1989. Precautionary saving and the timing of taxes. Journal of Political Economy 97, 863-879.
Liptser, R. S., Shiryayev, A. N., 1977. Statistics of Random Processes, Volume I, II. Springer-Verlag, New York.
MaCurdy, T. E., 1982. The use of time series processes to model the error structure of earnings in a longitudinal data analysis. Journal of Econometrics 18, 83-114.
Merton, R. C., 1971. Optimal consumption and portfolio rules in a continuous-time model. Journal of Economic Theory 3, 373-413.
Pischke, J., 1995. Individual income, incomplete information, and aggregate consumption. Econometrica 63, 805-840.
Svensson, L., Werner, I., 1993. Nontraded assets in incomplete markets. European Economic Review 37, 1149-1168.
Veronesi, P., 1999. Stock market overreaction to bad news in good times: A rational expectations equilibrium model. Review of Financial Studies 12, 975-1007.
Wang, J., 1993. A model of intertemporal asset prices under asymmetric information. Review of Economic Studies 60, 249-282.
Wang, N., 2003. Caballero meets Bewley: The permanent-income hypothesis in general equilibrium. American Economic Review 93, 927-936.
Wang, N., 2004. Precautionary saving with partially observed income. Journal of Monetary Economics 51, 1645-1681.
Wang, N., 2006. Generalizing the permanent-income hypothesis: Revisiting Friedman's conjecture on consumption. Journal of Monetary Economics 53, 737-752.
Xia, Y., 2001. Learning about predictability: the effect of parameter uncertainty on dynamic portfolio allocation. Journal of Finance 56, 205-246.
Zeldes, S. P., 1989. Optimal consumption with stochastic income: Deviations from certainty equivalence. Quarterly Journal of Economics 104 (2), 275-298.

Table 1: Categorization of Models

income growth α	known	unknown
constant	Model I	Model II
stochastic	Model III	Model IV

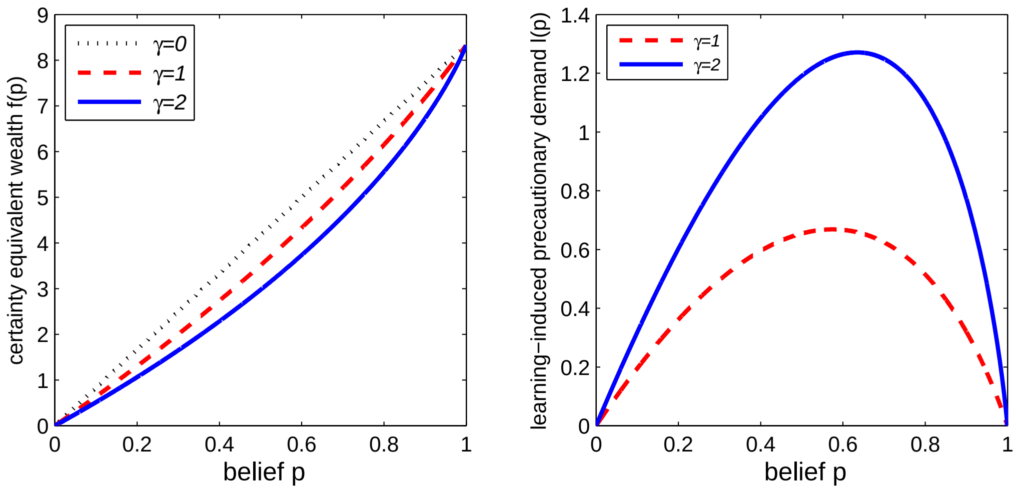

Figure 1: **The certainty equivalent wealth $f(p)$ and learning-induced precautionary saving $l(p) = \bar{f}(p) - f(p)$ in Model II (with $\rho = 0$): The effects of risk aversion.**

1

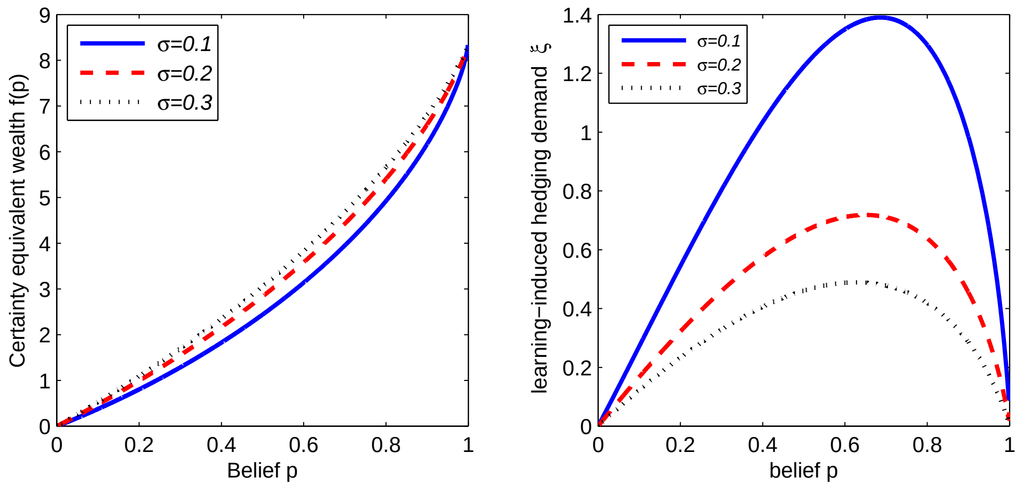

Figure 2: **The certainty equivalent wealth $f(p)$ and learning-induced hedging demand $\xi(p)$ in Model II: The effects of income volatility σ. The correlation coefficient: $\rho = 0.5$. Other parameter values are the same as those for Figure 1.**

1

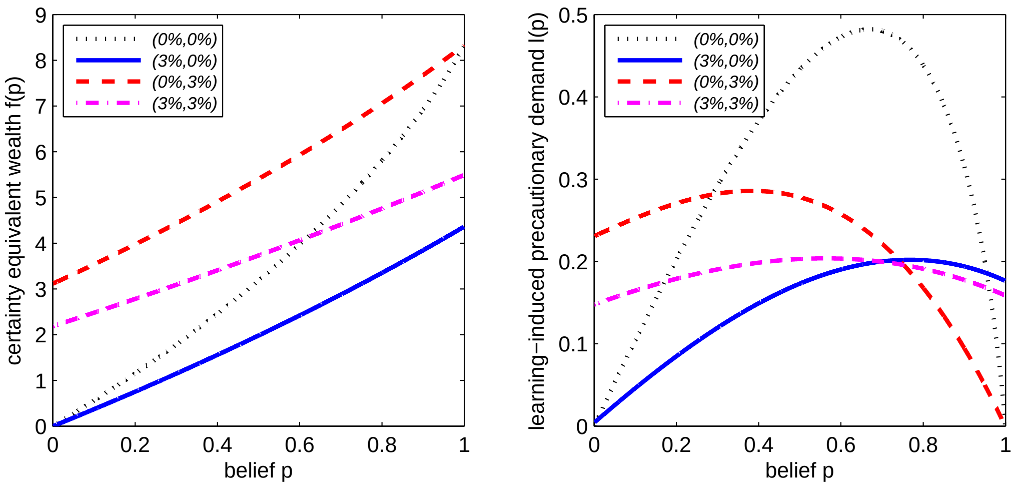

Figure 3: **The certainty equivalent wealth $f(p)$ and learning-induced precautionary demand $l(p)$ in a stochastic income growth model with learning in Model IV: The effects of transition intensities (λ_1, λ_2). Other parameter values are the same as those for Figure 2.**

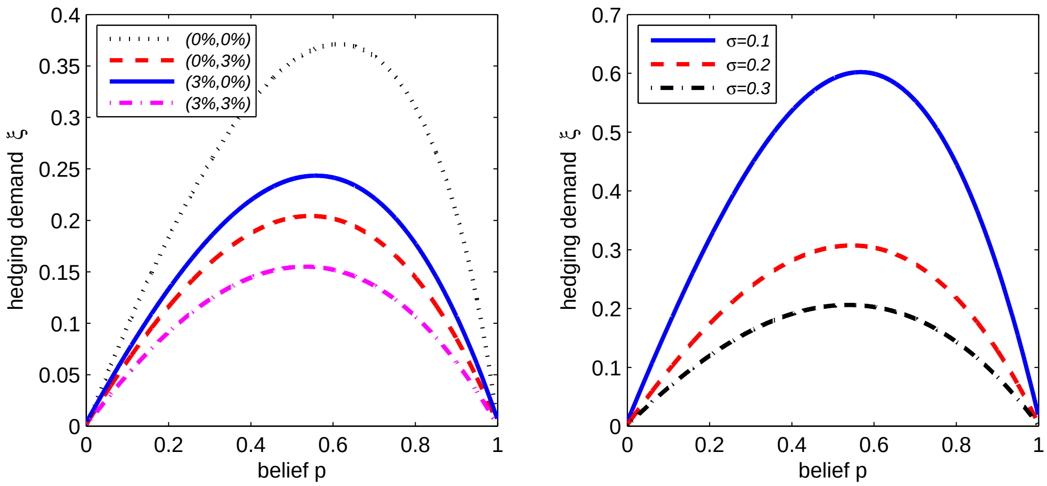

Figure 4: **Learning-induced hedging demand $\xi(p)$ in a stochastic income growth model with learning (Model IV). The left panel plots for various income growth transition parameters (λ_1, λ_2). The right panel plots for three levels of σ for the setting with $\lambda_1 = \lambda_2 = 3\%$. Other parameter values are the same as those for Figure 3.**

1