

Precautionary Saving and the Marginal Propensity to Consume Out of Permanent Income

August 15, 2009

Christopher D. Carroll¹

Johns Hopkins University

Abstract

The budget constraint requires that, eventually, consumption must adjust fully to any permanent shock to income. Intuition suggests that, knowing this, optimizing agents will fully adjust their spending immediately upon experiencing a permanent shock. However, this paper shows that if consumers are impatient and are subject to transitory as well as permanent shocks, the optimal marginal propensity to consume out of permanent shocks (the MPCP) is strictly less than 1, because buffer stock savers have a target wealth-to-permanent-income *ratio*; a positive shock to permanent income moves the ratio below its target, temporarily boosting saving.

Keywords risk, uncertainty, consumption, precautionary saving, buffer stock saving, permanent income hypothesis

JEL codes D81, D91, E21

As Published in the Journal of Monetary Economics, Volume 56, Issue 6, September 2009, pp 780-790
<http://dx.doi.org/10.1016/j.jmoneco.2009.06.016>

PDF: <http://econ.jhu.edu/people/ccarroll/papers/MPCPerm.pdf>

Web: <http://econ.jhu.edu/people/ccarroll/papers/MPCPerm/>

Archive: <http://econ.jhu.edu/people/ccarroll/papers/MPCPerm.zip>
(Contains the LaTeX source and Mathematica code solving the model)

¹ccarroll@jhu.edu, Department of Economics, 440 Mergenthaler Hall, Johns Hopkins University, Baltimore, MD 21218, and National Bureau of Economic Research, <http://econ.jhu.edu/people/ccarroll>.

Thanks to an anonymous referee whose comments substantially improved the paper. The paper's title is intended as homage to Miles Kimball's paper "Precautionary Saving and the Marginal Propensity to Consume."

1 Introduction

Arguably the core idea of Friedman (1957)’s Permanent Income Hypothesis is that an optimizing consumer’s response to an income shock should be much larger if that shock is permanent than if it is transitory.

A large empirical literature has shown that household income dynamics are reasonably well characterized by the Friedman (1957)-Muth (1960) dichotomy between permanent and transitory shocks.¹ And much of the subsequent theoretical literature can be interpreted as construction of the theoretical foundations for evaluating Friedman’s proposition under plausible assumptions about income dynamics, utility functions, and expectations.

The hardest part of the theoretical enterprise has been incorporation of a rigorous treatment of labor income uncertainty. Indeed, full understanding of the theoretical effects of such uncertainty on the marginal propensity to consume (MPC) out of transitory shocks is relatively recent: Kimball (1990a,b) showed that under standard assumptions about utility and expectations, the introduction of uncertainty in noncapital income increases the MPC at a given level of consumption, but not necessarily at a given level of wealth; and Carroll and Kimball (1996) show that the introduction of uncertainty causes the MPC to rise at any given level of wealth, and to increase more for consumers at lower levels of wealth.²

Surprisingly, no previous paper has systematically analyzed the complementary question of how uncertainty affects the marginal propensity to consume out of permanent shocks (the ‘MPCP’),³ though the question is important not only as a loose end in consumption theory, but also for microeconomic analysis of inequality (in both consumption and income) and for both micro- and macroeconomic analysis of tax policies and business cycles. Indeed, the topic can occasionally become headline news: The 2001 U.S. income tax cut was promoted by some economists as providing economic ‘stimulus’ on the explicit grounds that it was a permanent tax cut and therefore would have an immediate one-for-one effect on consumption.⁴

The lack of a formal treatment probably reflects a sense among researchers that they already know the answer: The MPCP should equal one. Because it is impossible to permanently insulate consumption from a permanent shock, if consumption does not adjust immediately and fully to such a shock, it will eventually need to adjust *more than* one-for-one to make up for any initial period of less-than-full adjustment. Consumption-smoothers, the thinking goes, will prefer to adjust fully now rather than less-than-fully now and more-than-fully later.

¹See, e.g., MaCurdy (1982); Abowd and Card (1989); Carroll and Samwick (1997); Jappelli and Pistaferri (2000); Storesletten, Telmer, and Yaron (2004); Blundell, Low, and Preston (2008).

²This result is a direct implication of the concavity of the consumption function that Carroll and Kimball (1996) prove.

³By ‘permanent shocks’ here I mean shocks to noncapital income; the terms permanent income and permanent noncapital income are used interchangeably in this paper, except where doing so might cause confusion because of the ambiguity the term ‘permanent income’ can have when consumers receive both capital and noncapital income.

⁴Evidence on the actual outcome is difficult to interpret; see Johnson, Parker, and Souleles (2006) for the best attempt.

But the only rigorous theoretical underpinning for this view is provided by Deaton (1991), who examines the problem of a liquidity-constrained consumer whose *only* uncertainty comes in the form of permanent shocks to income; Deaton shows that, under a particular ‘impatience’ condition, such a consumer with zero wealth will exhibit an MPCP of 1 (because under these assumptions it is always optimal to consume all current income).

After deriving some new results that bolster Deaton’s conjecture that, in his model, wealth tends to fall toward the absorbing state of zero where the MPCP is indeed one, this paper shows that if there are transitory as well as permanent shocks, under realistic calibrations the optimal MPCP can be substantially (though not enormously) less than one. The alteration is a consequence of the target-saving behavior that emerges when consumers are both prudent (Kimball (1990b)) and impatient. For a consumer starting at the target ratio of assets to permanent income, a positive shock to permanent income leaves the *target* unchanged. But for a given level of initial assets, a positive shock to the level of permanent income reduces the ratio of those assets to permanent income. For a consumer starting at the target, consumption therefore does not move up by the full amount of the income shock; the reciprocal logic holds for negative shocks.

The paper is organized as follows. The first section sets up the model and notation, and shows how the requirement of intertemporal budget balance is reflected in the consumption function. The second section derives an expression for the MPCP and explains qualitatively why it can be different from one; it then shows the relationship between that expression and Deaton’s results, and derives a formula that applies to the more general model with both transitory and permanent shocks. Because the exact value of the MPCP cannot be determined except by numerical methods, the fourth section numerically solves and simulates and finds that the marginal propensity to consume out of permanent shocks tends to fall between 0.75 and 0.92 for a wide range of plausible parameter settings. This section concludes by showing that behavior of the ergodic population of consumers that arises in the model is very close to behavior of a single consumer with assets equal to the target value, suggesting that the inconvenient step of simulation may be unnecessary for many kinds of analysis.

2 The Model

The consumer is assumed to behave according to the limiting solution to the problem

$$\begin{aligned}
 \mathbf{v}_t(\mathbf{m}_t, \mathbf{p}_t) &= \max_{\mathbf{c}_t} \mathbb{E}_t \left[\sum_{n=0}^{T-t} \beta^n u(\mathbf{c}_{t+n}) \right] \\
 &\text{s.t.} \\
 \mathbf{a}_t &= \mathbf{m}_t - \mathbf{c}_t, \\
 \mathbf{p}_{t+1} &= \mathbf{p}_t \Gamma \psi_{t+1}, \\
 \mathbf{m}_{t+1} &= R \mathbf{a}_t + \mathbf{p}_{t+1} \xi_{t+1},
 \end{aligned} \tag{1}$$

as the horizon T approaches infinity, where for clarity we have separately specified the various transitional steps that are often combined when the problem is written in its most compact (Bellman equation) form: $\mathbf{a}_t$ indicates assets after all actions at the end of period t ; $R = (1+r)$ is the interest factor for assets held between periods; permanent noncapital income $\mathbf{p}_{t+1}$ is equal to its previous value, multiplied by a growth factor Γ , and modified by a mean-one shock ψ_{t+1} , $\mathbb{E}_t[\psi_{t+1}] = 1$ (and we henceforth denote the combination $\Gamma\psi_{t+1}$ compactly by Γ_{t+1});⁵ $\mathbf{m}_{t+1}$ indicates the level of the consumer's ‘cash-on-hand’ at the time the consumption decision is made (the sum of beginning-of-period assets plus current-period noncapital income, where noncapital income equals permanent noncapital income $\mathbf{p}_{t+1}$ multiplied by a mean-one transitory shock ξ_{t+1} , $\mathbb{E}_t[\xi_{t+n}] = 1 \ \forall \ n > 0$, and we henceforth designate total noncapital income by $\mathbf{y}_{t+1} \equiv \mathbf{p}_{t+1}\xi_{t+1}$).⁶ As usual, the recursive nature of the problem allows us to express it more compactly as:

$$\mathbf{v}_t(\mathbf{m}_t, \mathbf{p}_t) = \max_{\mathbf{c}_t} u(\mathbf{c}_t) + \beta \mathbb{E}_t [\mathbf{v}_{t+1}(\mathbf{m}_{t+1}, \mathbf{p}_{t+1})].$$

As written, the problem has two state variables, the level of permanent income $\mathbf{p}_t$ and the level of cash-on-hand $\mathbf{m}_t$. Carroll (2009) shows that if utility is of the Constant Relative Risk Aversion (CRRA) form $u(c) = c^{1-\rho}/(1-\rho)$, it is possible to normalize the problem by the level of permanent income $\mathbf{p}_t$, thereby reducing the effective number of state variables to one. Specifically, defining nonbold variables as the bold equivalent divided by permanent noncapital income,⁷ $m_t = \mathbf{m}_t/\mathbf{p}_t$, $c_t = \mathbf{c}_t/\mathbf{p}_t$, and so on, and defining $\mathcal{R}_{t+1} \equiv R/\Gamma_{t+1}$, if we solve the problem

$$\begin{aligned} v_t(m_t) &= \max_{c_t} u(c_t) + \beta \mathbb{E}_t[\Gamma_{t+1}^{1-\rho} v_{t+1}(m_{t+1})] \\ \text{s.t.} \\ a_t &= m_t - c_t, \\ m_{t+1} &= a_t \mathcal{R}_{t+1} + \xi_{t+1} \end{aligned} \tag{2}$$

backwards from a final period of life in which $v_T(m) = u(m)$, the full value function $\mathbf{v}_t(\mathbf{m}_t, \mathbf{p}_t)$ at any prior period $t < T$ can be recovered from $\mathbf{v}_t(\mathbf{m}_t, \mathbf{p}_t) = \mathbf{p}_t^{1-\rho} v_t(m_t/\mathbf{p}_t)$ and the consumption function from $\mathbf{c}_t(\mathbf{m}_t, \mathbf{p}_t) = c_t(m_t/\mathbf{p}_t)\mathbf{p}_t$.

Carroll (2009) proves that the problem defines a contraction mapping with a limiting consumption function $c(m)$, under certain conditions including a requirement that the limiting discounted value of optimal behavior is finite and well-defined, which is guaranteed by the ‘finite value condition’ (FVC)

$$\beta \mathbb{E}_t[\Gamma_{t+1}^{1-\rho}] < 1. \tag{3}$$

The most interesting class of solutions is those that obtain when, in addition to the FVC, a ‘growth impatience condition’ (GIC) also holds. Defining the GIC requires

⁵Note that the definition of permanent income here differs from Deaton (1992)’s definition (which is often used in the macro literature), in which permanent income is the amount that a perfect foresight consumer could spend while leaving total (human and nonhuman) wealth constant.

⁶This problem is identical to problems that have been analyzed in a number of papers on ‘buffer-stock saving’ beginning with Carroll (1992); it differs from the problem analyzed by Deaton (1991) primarily because liquidity constraints are absent.

⁷See the appendix for an atlas of the variable names and notational conventions in this paper.

construction of an uncertainty-adjusted permanent income growth factor

$$\dot{\Gamma} = \Gamma(\mathbb{E}_t[\psi_{t+1}^{-1}])^{-1} \quad (4)$$

and the specification of an ‘absolute patience factor’

$$\mathbf{P} = (R\beta)^{1/\rho} \quad (5)$$

which measures the growth factor for consumption that would be chosen by the unconstrained perfect foresight consumer (the symbol $\mathbf{P}$ is the Old English letter ‘thorn’).

The GIC can be stated as a requirement that

$$\mathbf{P}_{\dot{\Gamma}} \equiv \left(\frac{\mathbf{P}}{\dot{\Gamma}} \right) < 1 \quad (6)$$

where we call the scaled version of $\mathbf{P}$ in (6) the ‘growth patience factor.’⁸

Some important conclusions can be drawn simply from the fact that the model can be rewritten in ratio form. The first is that because the level of consumption can be rewritten as $\mathbf{c}_t = c(m_t)\mathbf{p}_t$ for some invariant $c(m)$, the only way the elasticity of consumption with respect to permanent income $\mathbf{p}_t$ can be different from one is if there is a correlation between $\mathbf{p}_t$ and m_t . Of course, such a correlation does exist: Both $\mathbf{p}_t$ and m_t are influenced by the realization of the stochastic shock to permanent income ψ_t . Furthermore, both will reflect residual effects of the previous shocks to permanent income, $\psi_{t-1}, \psi_{t-2}, \dots$. It is these effects of the permanent shocks on the cash-on-hand to permanent-income *ratio* that will be the key to understanding the results below.

Another important insight comes from the fact, recently proven by Szeidl (2006), that the distribution of m_t is ergodic in models in this class. This implies that *eventually* the *infinite-horizon* MPCP must be one because ergodicity of m_t means that the expectation as of time t of m_{t+n} as $n \rightarrow \infty$ is the same for any particular realizations of $\psi_t, \psi_{t-1}, \psi_{t-2}, \dots$, implying that as $n \rightarrow \infty$ the time- t expectation of $c(m_{t+n})\mathbf{p}_{t+n}$ depends only on the level of $\mathbf{p}_t$.

But the ‘marginal propensity to consume’ out of a shock has traditionally been defined as the immediate effect, not the total eventual effect, and so we now ask how consumption is affected in period t by the contemporaneous realization of the shock to permanent income ψ_t .

3 The Marginal Propensity to Consume Out Of Permanent Income

As a benchmark, it is useful to begin by deriving the relationship between consumption and permanent income in the perfect foresight framework.⁹

⁸In fact, the paper shows that the problem defines a contraction mapping even under a somewhat weaker condition than the GIC (though the FVC is always required); however, for present purposes the only interesting solutions are those for which the GIC condition holds, so we impose it as sufficient even if not necessary.

⁹Results would be very similar for the model analyzed by Caballero (1990), with labor income uncertainty but constant absolute risk aversion utility, or for the certainty equivalent model.

3.1 The Perfect Foresight Case

A standard result in consumption theory¹⁰ is that for the infinite horizon perfect foresight version of the model above (i.e. a version in which $\xi_t = \psi_t = 1 \forall t$), the level of consumption is given by

$$\mathbf{c}_t = (1 - R^{-1}(R\beta)^{1/\rho}) \left[R\mathbf{a}_{t-1} + \left(\frac{\mathbf{p}_t}{1 - \Gamma/R} \right) \right]. \quad (7)$$

While strictly speaking there is no such thing as a ‘shock’ to permanent income in the perfect foresight model, it is possible to calculate how consumption would be different if permanent income were different. The answer is given by

$$\left(\frac{d\mathbf{c}_t}{d\mathbf{p}_t} \right) = \left(\frac{1 - R^{-1}(R\beta)^{1/\rho}}{1 - \Gamma/R} \right), \quad (8)$$

which we will refer to henceforth as the MPCP for the perfect foresight model. This quantity is less than one if

$$\begin{aligned} \Gamma/R &< R^{-1}(R\beta)^{1/\rho} \\ \Gamma &< (R\beta)^{1/\rho} \\ 1 &< R\beta\Gamma^{-\rho}. \end{aligned}$$

Notice that for $\rho > 1$ this can hold only if the GIC condition (6) *fails* (if we capture the perfect foresight version of the GIC by setting $\Gamma_{t+1} = \Gamma \forall t$). The interpretation is that in the perfect foresight framework, only the *patient* consumers have an MPCP of less than one. This makes intuitive sense: Patient consumers prefer to consume more in the future than in the present, so they do not spend all of the increase in income today.

Although this perfect foresight framework is often presented as the formalization of Friedman (1957)’s Permanent Income Hypothesis, the model implies that consumption responds one-for-one to a change in permanent noncapital income $\mathbf{p}_t$ only if $(R\beta)^{1/\rho} = \Gamma$. For plausible parameter values the model can easily predict an MPCP of anywhere between 0 and 6 (see table 1 for a parameterization that implies an MPCP of 6). This observation casts doubt upon the proposition that it is appropriate to treat the perfect foresight model as a formalization of Friedman (1957). For an argument that the buffer-stock model (that is, the solution to the model described above with impatient but prudent consumers) is a much better match than the perfect foresight model to Friedman’s original description of the PIH, see Carroll (2001).

3.2 The Response to Permanent Income Shocks

The natural definition of the MPCP in a model with shocks is the derivative of $\mathbf{c}_{t+1}$ with respect to ψ_{t+1} , given an initial level of assets $\mathbf{a}_t = a_t \mathbf{p}_t$,

$$\frac{d\mathbf{c}_{t+1}}{d\psi_{t+1}} = \frac{d\mathbf{p}_{t+1}c(m_{t+1})}{d\psi_{t+1}}$$

¹⁰For a derivation, see, e.g., the graduate lecture notes on the author’s home page, <http://econ.jhu.edu/people/ccarroll/public/lecturenotes>.

$$= \mathbf{p}_t \Gamma \left(\frac{d(\psi_{t+1} c(\mathcal{R}_{t+1} a_t + \xi_{t+1}))}{d\psi_{t+1}} \right).$$

This equation reveals a minor conceptual difficulty: The effect of ψ_{t+1} on c_{t+1} depends not only on the value of a_t but also on the realization of ξ_{t+1} , and so in principle there are two ‘state variables’ (other than the scaling variable $\hat{\mathbf{p}}_{t+1} \equiv \Gamma \mathbf{p}_t$) that determine the *ex post* MPCP. However, since ξ_{t+1} is an i.i.d. random variable, it is easy and intuitive to calculate the expectation of the derivative as

$$\begin{aligned} \mathbb{E}_t \left[\frac{d}{d\psi_{t+1}} \hat{\mathbf{p}}_{t+1} \psi_{t+1} c_{t+1} \right] &= \hat{\mathbf{p}}_{t+1} \mathbb{E}_t \left[\psi_{t+1} \frac{dc(m_{t+1})}{d\psi_{t+1}} + c(m_{t+1}) \right] \\ &= \hat{\mathbf{p}}_{t+1} \mathbb{E}_t \left[\psi_{t+1} c'(m_{t+1}) \frac{dm_{t+1}}{d\psi_{t+1}} + c(m_{t+1}) \right] \\ &= \hat{\mathbf{p}}_{t+1} \mathbb{E}_t \left[\psi_{t+1} c'(m_{t+1}) \frac{d}{d\psi_{t+1}} (\mathcal{R}_{t+1} a_t + \xi_{t+1}) + c(m_{t+1}) \right] \\ &= \hat{\mathbf{p}}_{t+1} \mathbb{E}_t [c(m_{t+1}) - c'(m_{t+1}) \mathcal{R}_{t+1} a_t], \end{aligned} \quad (9)$$

where the last line follows because $\mathcal{R}_{t+1} = (R/\Gamma) \psi_{t+1}^{-1}$ and $\psi(d/d\psi) \psi^{-1} = -\psi \psi^{-2} = -\psi^{-1}$. This expression leads to the natural definition of the MPCP, $\pi(a_t)$, as the expression multiplying the expected level of permanent income $\hat{\mathbf{p}}_{t+1}$,

$$\pi(a_t) \equiv \mathbb{E}_t [c(m_{t+1}) - c'(m_{t+1}) \mathcal{R}_{t+1} a_t]. \quad (10)$$

3.3 The Deaton Case (Permanent Shocks Only)

This expression maps nicely into Deaton (1991)’s finding that for consumers who begin with zero market resources the marginal propensity to consume out of $\mathbf{p}_{t+1}$ is one. Such consumers have $a_t = 0$ and therefore the second term on the RHS in equation (9) drops out. Deaton also assumed that there were no transitory shocks to income, so that $\xi_{t+1} = 1$. Finally, his consumers were sufficiently impatient so that their consumption at $c(m)$ was equal to one at $m = 1$. Hence the MPCP was given by $\pi(0) = \mathbb{E}_t [c(1)] = 1$.

To really understand Deaton’s result, it is necessary to recall why it must be that $c(1) = 1$.¹¹ Consider the first order condition for the unconstrained optimization problem,

$$c(m_t)^{-\rho} = R\beta \mathbb{E}_t [\Gamma_{t+1}^{-\rho} c(m_{t+1})^{-\rho}].$$

The consumer will be constrained at $m_t = 1$ iff the marginal utility of consuming 1 (which is $1^{-\rho} = 1$) is greater than the marginal utility of saving $a_t = 0$, i.e. if

$$\begin{aligned} 1 &> R\beta \mathbb{E}_t [\Gamma_{t+1}^{-\rho} c((\mathcal{R}_{t+1}) \times 0 + 1)^{-\rho}] \\ 1 &> R\beta \mathbb{E}_t [\Gamma_{t+1}^{-\rho}] \end{aligned} \quad (11)$$

where the second line follows from the first because with $a_t = 0$, $m_{t+1} = \xi_{t+1} = 1 = m_t$. Deaton directly imposes condition (11), thus guaranteeing his result that a consumer

¹¹ The following is intended as a loose intuitive argument rather than a rigorous derivation; in particular it mixes logic from the constrained and unconstrained optimization problems. See Deaton (1991) for the rigorous version.

with zero a who experiences only permanent shocks will remain at zero a forever. Zero a is an absorbing state.¹²

What Deaton was unable to prove, but conjectured must be true, was that a liquidity-constrained consumer who starts with positive a will always eventually run down that a to reach the absorbing state of $a = 0$. Consider the accumulation equation for m ,

$$a_{t+1} = \mathcal{R}_{t+1}a_t + 1 - c(\mathcal{R}_{t+1}a_t + 1). \quad (12)$$

Carroll and Kimball (2005) show that the marginal propensity to consume out of transitory income in a problem with liquidity constraints is always greater than the MPC in the unconstrained case. We also know, from combining Kimball (1990a) and Carroll and Kimball (1996), that the MPC in the unconstrained case with noncapital income risk is greater than the MPC without noncapital income risk. But from (7) we know that the MPC in the unconstrained case with no uncertainty is

$$\kappa \equiv (1 - R^{-1}(R\beta)^{1/\rho}), \quad (13)$$

and so the Carroll and Kimball (1996) results tell us that

$$c(1 + \mathcal{R}_{t+1}a_t) > c(1) + \mathcal{R}_{t+1}a_t\kappa = 1 + \mathcal{R}_{t+1}a_t\kappa$$

where the equality holds because $c(1) = 1$. Substituting in equation (12),

$$\begin{aligned} a_{t+1} &< \mathcal{R}_{t+1}a_t - \kappa\mathcal{R}_{t+1}a_t \\ &< \mathcal{R}_{t+1}a_t(1 - \kappa). \end{aligned} \quad (14)$$

From this we have (substituting (13) into (14))

$$a_{t+1} < \mathcal{R}_{t+1}a_t R^{-1}(R\beta)^{1/\rho} = a_t(R\beta)^{1/\rho}/\Gamma_{t+1}. \quad (15)$$

But for $\rho > 1$, Jensen's inequality implies that Deaton's impatience condition (11) is stronger than the GIC imposed in (6),¹³ which is that the expectation of the expression multiplying a_t on the RHS of equation (15) is less than one, so that

$$\mathbb{E}_t[a_{t+1}] < a_t.$$

Thus, at any positive level of assets $a_t > 0$, assets are expected to fall toward zero. Note that this condition does not guarantee that assets ever reach zero in finite time, because in principle it is possible (though arbitrarily improbable) to draw an arbitrarily long sequence of low draws of ψ_t . On the other hand, equation (15) does rule out the possibility that Deaton raised (but doubted) that some positive level of assets $\underline{a}$ could exist such that if $a_t > \underline{a}$ the consumption rule might never allow assets to fall below $\underline{a}$, thus preventing the consumer from ever reaching the absorbing state of $a_t = 0$. Hence, in Deaton's model, a falls unboundedly toward zero, and if it ever reaches zero, the MPCP equals one ever after.

¹²Note that Deaton's condition is stronger than the one required for the problem to define a contraction mapping. This reflects a subtle distinction: If the weaker condition (6) is imposed, but Deaton's stronger condition (11) is not satisfied, then a consumer who begins the period with zero resources will choose to save some strictly positive amount. In this case, zero wealth is NOT an absorbing state, and the target asset-to-permanent-income ratio is actually positive.

¹³Because $\mathbb{E}_t[\Gamma_{t+1}^{-\rho}]^{1/\rho} \geq \mathbb{E}_t[\Gamma_{t+1}^{-1}]^{-1}$.

3.4 The General Case (Transitory and Permanent Shocks)

Carroll (2009) proves that a ‘target’ value of $\tilde{a}$ will exist, where the target is defined as the level of assets such that $\mathbb{E}_t[a_{t+1}] = a_t$. Consider the behavior of consumption around the target,

$$\begin{aligned} a_{t+1} &= \mathcal{R}_{t+1}a_t + \xi_{t+1} - c(\mathcal{R}_{t+1}a_t + \xi_{t+1}) \\ \mathbb{E}_t[a_{t+1}] &= \mathbb{E}_t[\mathcal{R}_{t+1}]a_t + 1 - \mathbb{E}_t[c(\mathcal{R}_{t+1}a_t + \xi_{t+1})] \\ \tilde{a} &= \tilde{a} \mathbb{E}_t[\mathcal{R}_{t+1}] + 1 - \mathbb{E}_t[c(\mathcal{R}_{t+1}\tilde{a} + \xi_{t+1})] \\ \mathbb{E}_t[c(\mathcal{R}_{t+1}\tilde{a} + \xi_{t+1})] &= 1 + (\mathbb{E}_t[\mathcal{R}_{t+1}] - 1)\tilde{a}. \end{aligned} \tag{16}$$

With this observation about the nature of the target, we are now in position to walk through the key result of the paper. At $a_t = \tilde{a}$, from (10) the definition of the MPCP is

$$\begin{aligned} \pi(\tilde{a}) &= \mathbb{E}_t[c(\mathcal{R}_{t+1}\tilde{a} + \xi_{t+1})] - \mathbb{E}_t[c'(\mathcal{R}_{t+1}\tilde{a} + \xi_{t+1})\mathcal{R}_{t+1}]\tilde{a} \\ &= 1 + ((\mathbb{E}_t[\mathcal{R}_{t+1}] - 1) - \mathbb{E}_t[c'(\mathcal{R}_{t+1}\tilde{a} + \xi_{t+1})\mathcal{R}_{t+1}])\tilde{a} \end{aligned}$$

so that if $\tilde{a} > 0$ (which will be shown below), it is clear that the MPCP will be less than one if

$$0 > ((\mathbb{E}_t[\mathcal{R}_{t+1}] - 1) - \mathbb{E}_t[c'(\mathcal{R}_{t+1}\tilde{a} + \xi_{t+1})\mathcal{R}_{t+1}]). \tag{17}$$

Before we prove that this condition holds, consider what it means in intuitive terms. Since R , Γ and ψ are all numbers close to one, the latter term will be very close to the expected marginal propensity to consume $\mathbb{E}_t[c'_{t+1}(m_{t+1})]$. The former term is the intrinsic geometric growth rate of the assets/permanent-labor-income ratio (intrinsic, in the sense that it reflects both the return on assets R and the dilution of assets by permanent income growth and shocks, $1/\Gamma_{t+1}$). So this condition boils down to whether the MPC out of transitory income is greater than the intrinsic growth of a . But that is fundamentally what the impatience condition is about: If consumers are impatient, they will want to spend more than the amount justified by intrinsic growth of a . Thus, the assumption of impatience ensures an MPC out of transitory income that is large enough to overcome the intrinsic growth of a .

The key question therefore is whether we know the MPC out of transitory income is large enough. But recall that Carroll and Kimball (1996) have shown that the marginal propensity to consume under uncertainty is strictly greater than the MPC in the corresponding perfect certainty model, which turns out to be precisely the lower bound we need. That is, we know that $c'(m_{t+1}) > \kappa$ where as above $\kappa = 1 - R^{-1}(R\beta)^{1/\rho}$ is the MPC in the perfect foresight infinite horizon case. Using this fact gives

$$\mathbb{E}_t[c'(m_{t+1})\mathcal{R}_{t+1}] > \kappa \mathbb{E}_t[\mathcal{R}_{t+1}]$$

so (17) will certainly hold if the weaker condition

$$\begin{aligned} 0 &> (\mathbb{E}_t[\mathcal{R}_{t+1}] - 1) - \mathbb{E}_t[\mathcal{R}_{t+1}]\kappa \\ 1 &> \mathbb{E}_t[(R\beta)^{1/\rho}/\Gamma_{t+1}] \end{aligned}$$

holds. But this is just the GIC imposed above. Hence, at the target a the MPCP is strictly less than one.

Thus, the bottom line is that the growth impatience condition (6) guarantees a marginal propensity to consume out of transitory income that is large enough that, at the target $\check{a}$, the reduction in a induced by the permanent income shock cuts consumption by more than the amount that consumption increases as a result of the higher permanent income.

The final loose end is to show that $\check{a} > 0$. However, a result long-established in this literature is that with a CRRA utility function and no liquidity constraints, the lower bound on assets is the present discounted value of the minimum possible realization of future labor income. With lognormal permanent income shocks with no lower bound (as assumed here), the lower bound on future labor income is zero, so assets will always be strictly greater than zero. With actual assets always strictly positive, the target a must be positive if it exists.¹⁴

A brief discussion of how the results would be modified in the presence of liquidity constraints is in order. The first point to note is that for the model exactly as presented above, the addition of constraints would have no effect on behavior, because the consumer voluntarily chooses never to borrow even if constraints are not present. However, if lower bounds are placed on the transitory and permanent shocks, then consumers will wish to borrow in some circumstances. In this case constraints can make a difference. Carroll and Kimball (2005) provide a rigorous analysis of the effects of constraints on the decision rule, and it is clear from that analysis that a comprehensive and rigorous analysis of the effects of constraints here would be very complex. But intuition provides a clear bottom line. In the case with constraints, the minimum value of a is zero. It is also possible that the target $\check{a}$ is zero. But there will generally be some consumers who in some circumstances will hold positive assets. For these consumers, the logic above should hold, so that the MPCP is less than one. Simulation analysis of the model with constraints presented in Carroll (2001) confirms these intuitions.

We can also say something about how $\pi(a_t)$ varies with the level of assets. Its derivative with respect to assets is given by

$$\begin{aligned} \left(\frac{d}{da_t} \right) \pi(a_t) &= \mathbb{E}_t [c'(m_{t+1})\mathcal{R}_{t+1} - c'(m_{t+1})\mathcal{R}_{t+1} - c''(m_{t+1})\mathcal{R}_{t+1}^2] \\ &= \mathbb{E}_t [-c''(m_{t+1})\mathcal{R}_{t+1}^2]. \end{aligned} \tag{18}$$

But Carroll and Kimball (1996) prove that for problems in the class considered here the consumption function is strictly concave, $c''(m) < 0$, and since $\mathcal{R}_{t+1}^2$ is certainly positive, equation (18) implies that the marginal propensity to consume out of permanent shocks is increasing in the level of assets.

Indeed, we can even show that for a large enough level of actual assets, the MPCP will rise above one. This is because as the ratio of actual assets to permanent income approaches infinity, behavior in the model becomes arbitrarily close to behavior in the perfect foresight model. (For a proof, see Carroll and Kimball (2005)). Equation (8) implies that if the impatience condition is satisfied, the MPCP for the perfect foresight

¹⁴For a proof that a target ratio exists, see Carroll (2009); a positive value of the target is also an implication of the results in Szeidl (2006), who shows that the support of the distribution of a is strictly positive, and the target must be inside the support of the distribution.

model is greater than one, so the limit of the MPCP for the buffer-stock model as assets approaches infinity must exceed one. Note, however, the peculiar nature of the thought experiment here: The impatience condition is precisely what *prevents* assets from rising to infinity, so the question of what happens to the MPCP as we mechanically move assets toward infinity despite the fact that they are predicted to fall, is very much a curiosum.

These results appear to be the most that can be said analytically about the characteristics of $\pi(a_t)$. To obtain quantitative results for the average behavior of a population of consumers it is necessary to simulate.

4 Simulation Results

Table 1 presents simulation results for the average value of π (labelled “Mean π ”) that arises in steady-state among a population of consumers all behaving according to the model outlined above, under a baseline set of parameter values and a variety of alternatives.

The baseline calibration of the income process is taken from [Carroll \(1992\)](#), who finds that household-level data from the *Panel Study of Income Dynamics* are reasonably well characterized by the assumption that ψ_t is lognormally distributed with standard deviation $\sigma_\psi = 0.10$, while the process for transitory income has two parts: With probability $\wp$, income is zero, and with probability $(1 - \wp)$ the transitory shock θ_t is equal to $1/(1 - \wp)$ times the value of a shock drawn from a lognormal distribution with standard deviation $\sigma_\theta = 0.10$ and mean value one, so that $\mathbb{E}_t[\xi_{t+1}] = 1$ as assumed above. Permanent noncapital income growth at the household level is assumed to be $\Gamma = 1.03$. The baseline calibration for the interest rate and time preference rate are commonly-used values in macroeconomics, $R = 1.04$, $\beta = 0.96$. The baseline coefficient of relative risk aversion is $\rho = 3$, in the middle of the range from 1 to 5 generally considered plausible.

The first row of the table presents results for the baseline parameter values. The main result is found in the column labelled “Mean π .” To be perfectly clear about what this object is, assume a population of mass 1 is distributed uniformly on the unit interval, and define the operator $\mathbb{M}$ which calculates the mean value of variables in a population whose members are indexed by i ; thus, “Mean π ” is

$$\mathbb{M}[\pi(a_{t,i})] = \int_0^1 \mathbb{E}_{t,i}[c(m_{t+1,i}) - c'(m_{t+1,i})\mathcal{R}_{t+1,i}a_{t,i}]di, \quad (19)$$

where the mean is calculated in a period t in which the distribution has converged to the invariant distribution whose existence is proven by [Szeidl](#).

For comparison, the table also presents, where applicable,¹⁵ the MPCP implied by the perfect foresight infinite horizon version of the model (labelled “ Π_∞ ”), and from a

¹⁵For the infinite horizon MPCP to exist, the condition $R > \Gamma$ must hold, but this condition is not required to solve the stochastic model.

Deviations from Baseline [‡]	Patience $\mathbf{P}_{\hat{\Gamma}}$	Mean c	Mean a	Mean κ	Mean π	Perfect Foresight	
						Π_{∞}	Π_{T-40}
None (baseline)	0.980	1.012	0.619	0.236	0.783	4.053	1.632
$\beta = 0.98$	0.987	1.016	0.800	0.163	0.787	3.364	1.474
$\beta = 0.90$	0.959	1.009	0.440	0.370	0.792	6.181	2.168
$R = 1.02$	0.974	1.000	0.545	0.277	0.781	- $\diamond$	1.958
$R = 1.06$	0.986	1.030	0.769	0.181	0.790	1.806	1.405
$\Gamma = 1.02$	0.990	1.030	1.002	0.125	0.803	2.027	1.375
$\Gamma = 1.04$	0.971	1.005	0.514	0.302	0.784	- $\diamond$	1.958
$\rho = 1$	0.979	1.005	0.230	0.320	0.913	4.160	1.657
$\rho = 4$	0.980	1.017	0.837	0.201	0.731	4.040	1.629
$\sigma_{\psi} = 0.05$	0.973	1.006	0.496	0.316	0.785	4.053	1.632
$\sigma_{\psi} = 0.12$	0.984	1.020	0.832	0.160	0.791	4.053	1.632
$\wp = 0.0005$	0.980	1.006	0.329	0.288	0.871	4.053	1.632
$\wp = 0.05$	0.980	1.028	1.388	0.177	0.658	4.053	1.632
$\sigma_{\theta} = 0.05$	0.980	1.012	0.581	0.240	0.787	4.053	1.632
$\sigma_{\theta} = 0.15$	0.980	1.014	0.680	0.229	0.777	4.053	1.632
$\sigma_{\psi}=0, \Gamma = 1.00$	0.999	1.069	1.729	0.062	0.882	1.013	1.008
$\sigma_{\psi}=0, \Gamma = 1.01$	0.990	1.020	0.667	0.211	0.783	1.351	1.171
$\sigma_{\psi}=0, \Gamma = 1.02$	0.980	1.010	0.538	0.283	0.784	2.027	1.375

[‡] The first column indicates parameters that differ from the baseline. The baseline values are $R = 1.04, \beta = 0.96, \Gamma = 1.03, \rho = 3, \sigma_{\psi} = 0.1, \sigma_{\theta} = 0.1, \wp = 0.005$. The first row presents results when all parameters are at their baseline values.

$\diamond$ The infinite horizon perfect foresight solution is not well defined for these configurations of parameter values because $R \leq \Gamma$.

$\mathbf{P}_{\hat{\Gamma}}$ is the value of the growth patience factor defined in equation (6).

Table 1 Simulated Population-Mean MPCP π and Other Statistics

perfect foresight model for a consumer of average age (45) who has a horizon of 40 years (twenty years of work and twenty years of retirement), labelled “ Π_{T-40} .”¹⁶

Under the baseline parameter values, the population-average value of π is about 0.79. As the remainder of the table shows, the population-average value of π is between about 0.75 and 0.92 for most parametric configurations.

In addition to π , the table presents population-average values of each of the terms that made up π from (10).

Recall that at the target level of $\check{a}$ equation (16) tells us that

$$\mathbb{E}_t[c(m_{t+1})] = 1 + (\mathbb{E}_t[\mathcal{R}_{t+1}] - 1)\check{a}.$$

Since $\mathbb{E}_t[\mathcal{R}_{t+1}]$ will generally be a number close to one, this first term in the $\pi(a)$ expression could be substantially different from one only if consumers ended up holding large values of a . But since they are impatient by assumption, they are not likely to end up with large values of a . This reasoning is confirmed by the column of the table labelled “Mean c ,” which finds values very close to 1 for all parametric combinations.

Thus, most of the variation in the average value of π across parametric choices is attributable to differences in the $-a_t \mathbb{E}_t[c'(m_{t+1})\mathcal{R}_{t+1}]$ term. Making consumers more patient has two effects on this term. On the one hand, it increases target assets $\check{a}$ and therefore average assets, which makes the term more negative, reducing π ; on the other hand, the MPC c' declines with the level of assets, which would tend to shrink the term and therefore increase π . The near-constancy of population-mean π indicates that these two effects are roughly offsetting across different parametric choices.

The relative stability of π for the buffer-stock model contrasts sharply with the MPCP for the infinite horizon perfect foresight model, for which the MPCP is always greater than 1.8 in the first panel of the table, and rises as high as 6.2. The reason the MPCP in the PF model is always greater than one is that our consumers all satisfy the impatience criterion; inspection of (8) will verify that the MPCP must be greater than one if the impatience criterion is satisfied. This makes sense; impatient perfect-foresight consumers, upon learning that their income will be higher forever, will tend to increase their consumption by more than the increase in current income. However, what may not have been obvious *ex ante* is how *much* greater than 1 the MPCP typically is in the PF model. Results for the finite-horizon perfect foresight model are less extreme than for the infinite horizon version, but even in the finite-horizon model the MPCP is always at least 1.2 in the upper panel of the table.

The last three rows of the table present results when the permanent shocks are shut down and income growth is reduced; the most important result is for the case where there is no growth at all in income, so that $(R\beta)^{1/\rho} \approx \Gamma$, which, as noted earlier, is the condition that guarantees an MPCP of 1 in the perfect foresight model (the actual perfect foresight MPCP reported in the table is slightly above 1 because $R\beta$ is slightly below 1 for the baseline values $(R, \beta) = (1.04, 0.96)$). In the absence of permanent shocks, the

¹⁶The assumption is that upon retirement, total noncapital income (including Social Security and pension income) drops permanently to about 70 percent of preretirement salary, a calibration that roughly matches empirical evidence for the U.S.

impatience condition is (barely) satisfied and the stochastic version of the model can be solved with transitory shocks, generating an average π of about 0.88.

The remaining two rows show the consequences when the expected growth rate of income rises to 1 percent and 2 percent: The PF MPCP increases sharply, to slightly over 2 when $\Gamma = 1.02$; in the finite-horizon PF model, the MPCP rises to slightly over 1.2. In contrast, π falls to about 0.79 in the stochastic version of the model. This experiment highlights the interesting point that the relationship between impatience and the MPCP is of opposite *sign* in the stochastic and nonstochastic versions of the model.

The principal message from the table is that if consumers are impatient but prudent, optimal behavior implies an immediate MPC out of permanent shocks that is somewhat less than one (but not enormously less) for a wide variety of parameter values. More broadly, the value of the MPCP is much less sensitive to parameter values in the stochastic version of the model than in the perfect foresight version. And of course, the MPCP would be even lower for a finite horizon version of the stochastic model (just as in the perfect foresight model), because over a finite horizon a “permanent” shock has less effect on future resources than in an infinite horizon model.

A final point deserves elaboration. The theoretical results derived in section 3 applied only at the target level of assets. Yet table 1 shows that the conclusions reached for the target level of assets hold for populations distributed according to the invariant distributions. Since constructing the invariant distributions requires considerable extra work, it would be worthwhile to see whether results at exactly the target levels of assets are a good proxy for results from the invariant populations.

Table 2 presents the main statistics of interest, calculated both as an average across consumers distributed according to the invariant distribution, and for a consumer exactly at the target value of m or a (depending on the argument of the function). The message is simple: The target values are always very close to the population-average values. This suggests that theoretical work along the lines of that conducted in section 3 is likely to be both qualitatively and quantitatively a good guide to the behavior of an entire population. Since more propositions can be proven for the target level of assets than for the behavior of the ergodic population, and since it is possible to obtain quantitative results for the target values of a model without simulating, this suggests that future theoretical and quantitative work with this model may be able to dispense with simulation altogether, considerably reducing the computational demands of working with this class of models.

5 Conclusion

Intuition suggests that rational forward-looking consumers should have a marginal propensity to consume of one out of permanent shocks. This paper shows that while this intuition is not correct, or even close to correct, for the canonical infinite horizon perfect-foresight version of the CRRA-utility optimization model, it is approximately right for the ‘buffer-stock’ version of the model that arises when consumers are impatient and have a standard precautionary saving motive. The reason the MPCP is somewhat

Deviations [‡]	$\mathbf{p}_{\Gamma}$	Assets a		MPC κ		MPCP π	
		Mean	Target	Mean	Target	Mean	Target
None (baseline)	0.980	0.619	0.600	0.236	0.230	0.783	0.782
$\beta = 0.98$	0.987	0.800	0.761	0.163	0.157	0.787	0.781
$\beta = 0.90$	0.959	0.440	0.433	0.370	0.368	0.792	0.792
$R = 1.02$	0.974	0.545	0.533	0.277	0.272	0.781	0.780
$R = 1.06$	0.986	0.769	0.730	0.181	0.175	0.790	0.785
$\Gamma = 1.02$	0.990	1.002	0.921	0.125	0.116	0.803	0.789
$\Gamma = 1.04$	0.971	0.514	0.504	0.302	0.298	0.784	0.784
$\rho = 1$	0.979	0.230	0.219	0.320	0.308	0.913	0.913
$\rho = 4$	0.980	0.837	0.808	0.201	0.195	0.731	0.727
$\sigma_{\psi} = 0.05$	0.973	0.496	0.487	0.316	0.312	0.785	0.785
$\sigma_{\psi} = 0.12$	0.984	0.832	0.781	0.160	0.151	0.791	0.783
$\wp = 0.0005$	0.980	0.329	0.303	0.288	0.278	0.871	0.870
$\wp = 0.05$	0.980	1.388	1.382	0.177	0.174	0.658	0.654
$\sigma_{\theta} = 0.05$	0.980	0.581	0.571	0.240	0.238	0.787	0.787
$\sigma_{\theta} = 0.15$	0.980	0.680	0.647	0.229	0.219	0.777	0.774
$\sigma_{\psi}=0, \Gamma = 1.00$	0.999	1.729	1.456	0.062	0.060	0.882	0.854
$\sigma_{\psi}=0, \Gamma = 1.01$	0.990	0.667	0.650	0.211	0.206	0.783	0.782
$\sigma_{\psi}=0, \Gamma = 1.02$	0.980	0.538	0.528	0.283	0.279	0.784	0.784

[‡] This column indicates parameters that differ from the baseline. The baseline values are $R = 1.04, \beta = 0.96, \Gamma = 1.03, \rho = 3, \sigma_{\psi} = 0.1, \sigma_{\theta} = 0.1, \wp = 0.005$. The first row presents results when all parameters are at their baseline values.

[◇] The infinite horizon perfect foresight solution is not well defined for these configurations of parameter values because $R \leq \Gamma$.

$\mathbf{p}_{\Gamma}$ is the value of the growth patience factor defined in equation (6).

Table 2 Population-Mean Values Versus Values At Target

less than one in the buffer-stock model is that an increase in permanent income reduces the ratio of assets to permanent income, thus (temporarily) increasing the amount of precautionary saving. Simulations show that across a wide range of assumptions about the degree of impatience, the marginal propensity to consume out of permanent shocks is generally in or near the range from 0.75 to about 0.92.

The results in this paper are important for three reasons. First, empirical evidence from household surveys indicates that households experience large permanent shocks to their incomes of precisely the kind studied here, and no existing paper has provided a general theoretical analysis of the effects of these kinds of shocks on consumption. Second, the sharp contrast between the results for the stochastic and nonstochastic models, and the fact that the results for the stochastic model are much more plausible, provides another reason (if any were needed) that economists should avoid using the perfect foresight model for quantitative analysis. Finally, the paper provides a formal justification (that many economists probably did not know was lacking in the perfect foresight framework) for the assertion that permanent increases or decreases in taxes should result in consumption responses of roughly the same size, though the scrupulous economic advisor should warn that the response should be slightly less than one-for-one in the short run.

References

- ABOWD, JOHN M., AND DAVID CARD (1989): “On the Covariance Structure of Earnings and Hours Changes,” *Econometrica*, 57, 411–445.
- BLUNDELL, RICHARD, HAMISH LOW, AND IAN PRESTON (2008): “Decomposing Changes in Income Risk Using Consumption Data,” *Manuscript, University College London*.
- CABALLERO, RICARDO J. (1990): “Consumption Puzzles and Precautionary Savings,” *Journal of Monetary Economics*, 25, 113–136, Available at http://ideas.repec.org/p/clu/wpaper/1988_05.html.
- CARROLL, CHRISTOPHER D. (1992): “The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence,” *Brookings Papers on Economic Activity*, 1992(2), 61–156, Available at <http://econ.jhu.edu/people/ccarroll/BufferStockBPEA.pdf>.
- (2001): “A Theory of the Consumption Function, With and Without Liquidity Constraints,” *Journal of Economic Perspectives*, 15(3), 23–46, <http://econ.jhu.edu/people/ccarroll/ATheoryv3JEP.pdf> (as published) <http://econ.jhu.edu/people/ccarroll/ATheoryv3NBER.pdf> (more rigorous), <http://econ.jhu.edu/people/ccarroll/ATheoryMath.zip> (software archive).
- (2009): “Theoretical Foundations of Buffer Stock Saving,” *Manuscript, Department of Economics, Johns Hopkins University*, Previous version was published as NBER Working Paper No. 10867. Latest version available at <http://econ.jhu.edu/people/ccarroll/BufferStockTheory.pdf>.
- CARROLL, CHRISTOPHER D., AND MILES S. KIMBALL (1996): “On the Concavity of the Consumption Function,” *Econometrica*, 64(4), 981–992, Available at <http://econ.jhu.edu/people/ccarroll/concavity.pdf>.
- (2005): “Liquidity Constraints and Precautionary Saving,” *Manuscript, Johns Hopkins University*, Available at <http://econ.jhu.edu/people/ccarroll/papers/liquidRevised.pdf>.
- CARROLL, CHRISTOPHER D., AND ANDREW A. SAMWICK (1997): “The Nature of Precautionary Wealth,” *Journal of Monetary Economics*, 40(1), 41–71, Available at <http://econ.jhu.edu/people/ccarroll/nature.pdf>.
- DEATON, ANGUS S. (1991): “Saving and Liquidity Constraints,” *Econometrica*, 59, 1221–1248, Available at <http://ideas.repec.org/a/ecm/emetrp/v59y1991i5p1221-48.html>.
- (1992): *Understanding Consumption*. Oxford University Press, New York.
- FRIEDMAN, MILTON A. (1957): *A Theory of the Consumption Function*. Princeton University Press.

- JAPPELLI, TULLIO, AND LUIGI PISTAFERRI (2000): “Intertemporal Choice and Consumption Mobility,” *Econometric Society World Congress 2000 Contributed Paper Number 0118*, <http://fmwww.bc.edu/RePEc/es2000/0118.pdf>.
- JOHNSON, DAVID S., JONATHAN A. PARKER, AND NICHOLAS S. SOULELES (2006): “Household Expenditure and the Income Tax Rebates of 2001,” *American Economic Review*, 96(5), 1589–1610.
- KIMBALL, MILES S. (1990a): “Precautionary Saving and the Marginal Propensity to Consume,” *NBER Working Paper Number 3403*.
- (1990b): “Precautionary Saving in the Small and in the Large,” *Econometrica*, 58, 53–73.
- MACURDY, THOMAS (1982): “The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis,” *Journal of Econometrics*, 18(1), 83–114.
- MUTH, JOHN F. (1960): “Optimal Properties of Exponentially Weighted Forecasts,” *Journal of the American Statistical Association*, 55(290), 299–306.
- STORESLETTEN, KJETIL, CHRIS I. TELMER, AND AMIR YARON (2004): “Consumption and Risk Sharing Over the Life Cycle,” *Journal of Monetary Economics*, 51(3), 609–633.
- SZEIDL, ADAM (2006): “Invariant Distribution in Buffer-Stock Saving and Stochastic Growth Models,” *Manuscript, University of California at Berkeley*.

A Appendix

The following tables are provided to aid the reader in keeping track of nonstandard elements of the paper's notation.

Parameter	Definition
R	Riskfree Interest Factor
Γ	Nonstochastic Part of Permanent Noncapital Income Growth
ψ	Permanent Shock to Income
ξ	Transitory Shock to Income
θ	Transitory Shock to Income, Conditional On Nonzero Income
ρ	Coefficient of Relative Risk Aversion

Some combinations of the parameters above are used as convenient shorthand:

Symbol	Definition
$\Gamma_{t+1} \equiv \Gamma\psi_{t+1}$	Growth Factor for Permanent Noncapital Income
$\hat{\Gamma} \equiv \Gamma(\mathbb{E}_t[\psi_{t+1}^{-1}])^{-1}$	Uncertainty-adjusted Growth Factor
$\mathcal{R}_{t+1} \equiv R/\Gamma_{t+1}$	Return Normalized By Growth
$\mathbf{P}_{\hat{\Gamma}} \equiv \frac{(\beta R)^{1/\rho}}{\hat{\Gamma}}$	Growth Patience Factor

Endogenous variables:

Variable	Definition
a	Assets After All Actions (at end of period)
b	Beginning Bank Balances Before Consumption Choice
c	Consumption
m	Market Resources (Sum of Bank Balances and Income)

The meaning of typographical accents:

Accent	Meaning
Bold	Level of a Variable
Plain	Ratio of The Variable To Permanent Labor Income
$\vee$	$\bullet$ is the target value of variable $\bullet$

Abbreviations:

Acronym	Meaning	Definition
FVC	Finite Value Condition	$\beta \mathbb{E}_t[\Gamma_{t+1}^{1-\rho}] < 1$
GIC	Growth Impatience Condition	$\mathbf{P}_{\hat{\Gamma}} < 1$