Aggregating Heterogeneous-Agent Models with Permanent Income Shocks

Harmenberg, Karl

Link to publication in CBS Research Portal

Department of Economics

Copenhagen Business School

Working paper 13-2020

Aggregating Heterogeneous-Agent
Models with Permanent Income
Shocks

Karl Harmenberg

Aggregating Heterogeneous-Agent Models

with Permanent Income Shocks[*]

Karl Harmenberg[†]

September 21, 2020

Abstract

I introduce a method for simulating aggregate dynamics of heterogeneous-agent models where log permanent income follows a random walk. The idea is to simulate the model using a counterfactual *permanent-income-neutral measure* which incorporates the effect that permanent income shocks have on macroeconomic aggregates. With the permanent-income-neutral measure, one does not need to keep track of the permanent-income distribution. The permanent-income-neutral measure is both useful for the analytical characterization of aggregate consumption-savings behavior and for simulating numerical models. Furthermore, it is trivial to implement with a few lines of code.

1 Introduction

The heterogeneous-agent macroeconomic paradigm emphasizes the importance of rich heterogeneity at the micro level for macroeconomic aggregates. Although conceptually appealing, realistic heterogeneous-agent models are computationally challenging to solve, requiring the modeler to weigh computational tractability against a realistic description of the economic environment. It is therefore important to find model descriptions which are simultaneously realistic and computationally tractable.

A strand of the heterogeneous-agent literature, following Zeldes (1989), Deaton (1991), and Carroll (1997), models income as subject to fully permanent income shocks and transitory income shocks.[1] Not only does a unit-root process for log permanent income match the data well, but such an income process, combined with CRRA preferences, permits a simple description of household behavior (e.g., consumption) in terms of permanent income P_t and normalized cash-on-hand $m_t := M_t/P_t$ (where M_t is cash on hand) on the convenient functional form $C_t = P_t c(m_t)$.

Being able to describe micro-level consumption behavior in terms of normalized cash on hand, without explicit reference to permanent income, helps keep the computational problem tractable since a state variable in the household problem can be eliminated. Furthermore, the elimination of permanent income as a state variable has permitted a relatively precise theoretical description of household behavior. (Carroll, 2020) However, the computational tractability and theoretical clarity is lost when describing macro-level consumption behavior. For aggregate variables such as aggregate consumption, it is not sufficient to keep track of the distribution of households along the normalized cash-on-hand dimension since one needs to weigh the households by their permanent income. It is therefore seemingly necessary to keep track of the distribution of households both with respect to normalized cash on hand and with respect to permanent income, and the tractability gained at the micro level is lost when aggregating over all households.

This paper shows a way to recover both computational tractability and theoretical clarity for the macro-level behavior of models with fully permanent income shocks. I introduce a sufficient statistic for aggregate variables such as aggregate consumption, the *permanent-income-weighted distribution*, and show that there is a simple way to characterize the law of motion for the permanent-income-weighted distribution. The law of motion for the permanent-income-weighted distribution is equivalent to the law of motion for the distribution of normalized cash-on-hand with the adjustment that permanent income shocks are drawn using a counterfactual *permanent-income-neutral measure* which oversamples the positive permanent income shocks and undersamples the negative permanent income shocks. To be precise, if the objective distribution of permanent income shocks is given by a density function $f_\eta(\eta)$, then the permanent-income-neutral measure

---

[1] Zeldes (1989), Deaton (1991), and Carroll (1997), as well as, e.g., Gourinchas and Parker (2002) and Campbell and Cocco (2007) use an income process with permanent income shocks to study micro-level household behavior. More recently, the setup has been used to study aggregate macroeconomic behavior by McKay (2017), Carroll et al. (2017), Carroll et al. (2019) and Harmenberg and Öberg (2019).

is given by $\tilde{f}_\eta(\eta) := \eta f_\eta(\eta)$.

Using the permanent-income-neutral measure together with Szeidl (2013)'s characterization of when a stable invariant distribution of cash on hand exists, I characterize when a stable invariant permanent-income-weighted distribution exists. I then prove the conjecture from Carroll (2020) that, in the long run, aggregate consumption grows at the same rate as aggregate income.

The permanent-income-neutral measure also yields computational improvements. As an example, I show how using the permanent-income-neutral measure improves the precision of Monte Carlo simulations of the household distribution for an Aiyagari model. The implementation of the permanent-income-neutral measure is trivial, simply replace $f_\eta(\eta)$ by $\tilde{f}_\eta(\eta)$ in the code when aggregating the model.

The disposition of the paper is as follows. In Section 2, I introduce notation and prove the main theorem of the paper, a characterization of the law of motion for the permanent-income-weighted distribution. In Section 3, I use the main theorem to theoretically characterize the aggregate behavior of buffer-stock savings models. In Section 4, I show how the permanent-income-neutral measure can be used to improve computations. Section 5 comments on the similarities with the risk-neutral measure used in asset pricing and concludes.

2   Main Theorem

In this section, I consider the problem of aggregating models with fully permanent income shocks. It is assumed that behavior at the micro level has been solved for through, e.g., value-function iteration or other methods.

In preparation for the main theorem, I introduce the class of models under consideration. The models are ones such that the state space $\mathbf{m} \times \mathbf{P}$ consists of a normalized-state dimension $\mathbf{m}$ (e.g., normalized cash on hand) and a permanent-income dimension $\mathbf{P}$. The law of motion for permanent income only depends on previous-period permanent income and the permanent-income shock while the law of motion for the normalized-state dimension depends on the previous-period normalized-state dimension, possibly other shocks, and the permanent income shock but *not* the previous-period permanent income. This relatively abstract formulation nests, e.g., the canonical buffer-stock savings model (Carroll, 1997), models with risky assets (Haliassos and Michaelides, 2003), and models with durable goods subject to non-convex adjustment costs (Harmenberg and Öberg, 2019) (in this case, the normalized state is two dimensional). For simplicity, the exposition presumes no aggregate shocks but aggregate dynamics can easily be accommodated by introducing time subscripts for the law of motion.

Formally, a model with permanent-income shocks that allows normalization is described by

1. a multiplicative law of motion for permanent income $P_t \in \mathbf{P} = \mathbb{R}_+$ given by $P_{t+1} = G\eta_{t+1}P_t$ where G is the growth rate of permanent income and the shock η_{t+1} is drawn from the probability density

function $f_\eta(\eta_{t+1})$ such that $E\eta_{t+1} = 1$

2. and a law of motion for the normalized state as described by a density kernel $\phi(m_{t+1}, m_t, \eta_{t+1})$, i.e., the probability density function for the normalized state $m_{t+1} \in \mathbf{m} \subseteq \mathbb{R}^k$ given the previous state $m_t \in \mathbf{m}$ and the permanent-income shock $\eta_{t+1} \in \mathbb{R}_+$.

With these primitives, the Markov operator that maps a distribution $\psi_t \in D(\mathbf{m} \times \mathbf{P})$ to the next-period distribution $\psi_{t+1} \in D(\mathbf{m} \times \mathbf{P})$ is explicitly described by

$$\psi_{t+1}(m_{t+1}, P_{t+1}) = \int \phi\left(m_{t+1}, m_t, \frac{P_{t+1}}{GP_t}\right) f_\eta\left(\frac{P_{t+1}}{GP_t}\right) \frac{1}{GP_t}\psi_t(m_t, P_t)dm_t dP_t. \tag{1}$$

Often, we are interested in the distribution of households along the normalized-state dimension, "forgetting" the permanent-income dimension. The following definition introduces the distribution of households along the normalized-state dimension.

Definition 2.1. *The* marginal distribution *(along the normalized-state dimension) is defined as* $\psi_t^m(m_t) := \int \psi_t(m_t, P_t)dP_t$.

The evolution of the marginal distribution is easy to simulate, just simulate the evolution of many households and do not bother keeping track of the permanent-income dimension. However, for computing aggregates such as aggregate consumption, a different distribution along the normalized-state dimension is needed. Intuitively, we need to weigh households by their permanent income. Since consumption scales with permanent income, the consumption of a household i is given by the consumption function $C_{it} = c(m_{it})P_{it}$. Therefore, aggregate consumption is given by

$$C_t = \int c(m_t)P_t\psi_t(m_t, P_t)dm_t dP_t$$
$$= \int c(m_t)\left(\int P_t\psi_t(m_t, P_t)dP_t\right) dm_t.$$

Capturing the integral inside the parenthesis, we introduce the following notation:

Definition 2.2. *The* permanent-income-weighted distribution *is defined as* $\tilde{\psi}_t^m(m_t) := G^{-t}\int P_t\psi_t(m_t, P_t)dP_t$.

Total consumption is then given as a growth factor G^t times the integral of normalized consumption over the permanent-income-weighted distribution,

$$C_t = G^t \int c(m_t)\tilde{\psi}_t^m(m_t)dm_t. \tag{2}$$

Note that the permanent-income-weighted distribution is a sufficient statistic for computing aggregate consumption, aggregate savings and similar aggregate variables where household behavior is weighted by

permanent income. The main result of this paper is that there exists an explicit characterization of the law of motion of the permanent-income-weighted distribution $\tilde{\psi}_t^m$ without reference to the full state ψ_t. Furthermore, the law of motion for the permanent-income-weighted distribution $\tilde{\psi}_t^m$ is similar to the law of motion for the marginal distribution ψ_t^m.

Theorem 2.1. *The law of motion for ψ_t^m is given by*

$$\psi_{t+1}^m(m_{t+1}) = \int \phi\left(m_{t+1}, m_t, \eta_{t+1}\right) f_\eta\left(\eta_{t+1}\right) \psi_t^m(m_t) dm_t d\eta_{t+1} \tag{3}$$

and the law of motion for $\tilde{\psi}_t^m$ is given by

$$\tilde{\psi}_{t+1}^m(m_{t+1}) = \int \phi\left(m_{t+1}, m_t, \eta_{t+1}\right) \tilde{f}_\eta\left(\eta_{t+1}\right) \tilde{\psi}_t^m(m_t) dm_t d\eta_{t+1} \tag{4}$$

where $\tilde{f}_\eta(\eta_{t+1}) := \eta_{t+1} f_\eta(\eta_{t+1})$. We call the distribution $\tilde{f}_\eta$ the permanent-income-neutral measure.

Proof. Write $\tilde{\psi}_t^{m,\alpha} := G^{-\alpha t} \int P_t^\alpha \psi_t(m_t, P_t) dP_t$. The two distributions ψ_t^m and $\tilde{\psi}_t^m$ are the special cases with $\alpha = 0$ and $\alpha = 1$ respectively. The law of motion for $\tilde{\psi}_t^{m,\alpha}$ is given by

$$
\begin{aligned}
\tilde{\psi}_{t+1}^{m,\alpha}(m_{t+1}) &= G^{-\alpha(t+1)} \int P_{t+1}^\alpha \psi_{t+1}(m_{t+1}, P_{t+1}) dP_{t+1} \\
&= G^{-\alpha(t+1)} \int P_{t+1}^\alpha \phi\left(m_{t+1}, m_t, \frac{P_{t+1}}{GP_t}\right) f_\eta\left(\frac{P_{t+1}}{GP_t}\right) \frac{1}{GP_t} \psi_t(m_t, P_t) dm_t dP_t dP_{t+1} \\
&= G^{-\alpha(t+1)} \int \eta_{t+1}^\alpha G^\alpha P_t^\alpha \phi\left(m_{t+1}, m_t, \eta_{t+1}\right) f_\eta\left(\eta_{t+1}\right) \psi_t(m_t, P_t) dm_t dP_t d\eta_{t+1} \\
&= \int \phi\left(m_{t+1}, m_t, \eta_{t+1}\right) \eta_{t+1}^\alpha f_\eta\left(\eta_{t+1}\right) \left(G^{-\alpha t} \int P_t^\alpha \psi_t(m_t, P_t) dP_t\right) dm_t d\eta_{t+1} \\
&= \int \phi\left(m_{t+1}, m_t, \eta_{t+1}\right) \eta_{t+1}^\alpha f_\eta\left(\eta_{t+1}\right) \tilde{\psi}_t^{m,\alpha}(m_t) dm_t d\eta_{t+1},
\end{aligned}
$$

where the first equality is true by definition, the second equality by Equation 1, the third equality by the change of variable $d\eta_{t+1} = \frac{dP_{t+1}}{GP_t}$, the fourth by Fubini's theorem, and the fifth by definition. □

Remark 2.1. *With CRRA utility, household period utility is given by $C^{1-\gamma}/(1-\gamma)$. By setting $\alpha = 1 - \gamma$, the proof of Theorem 2.1 also describes the law of motion for $\tilde{\psi}_t^{m,1-\gamma}$ which is the sufficient statistic for aggregate welfare. Similarly, the proof explicitly characterizes the law of motion for $\tilde{\psi}_t^{m,2}$, the sufficient statistic for computing consumption squared. It is therefore possible to compute cross-sectional consumption variance by keeping track of $\tilde{\psi}_t^m$ and $\tilde{\psi}_t^{m,2}$.*

The law of motion for the permanent-income-weighted distribution is obtained by formally replacing the permanent-income shock distribution $f_\eta(\eta)$ with the permanent-income-neutral shock distribution $\tilde{f}_\eta(\eta) := \eta f_\eta(\eta)$. What is the intuition behind this result? The marginal distribution ψ_t^m describes how the households

are distributed along the normalized-state dimension while the permanent-income-weighted distribution describes how the units of permanent income are distributed along the normalized-state dimension. A particular realization of the permanent-income shock η affects $f(\eta)$ households, but in terms of permanent income these households are now given weight $\eta \times f(\eta)$. Therefore, if we want to keep track of the distribution of permanent income along the normalized-state dimension, we have to use the permanent-income-neutral measure.

Theorem 2.1 suggests an immediate way to simulate the permanent-income-weighted distribution, and thereby the evolution of aggregate variables. Take any method that simulates the evolution of households along the normalized-state dimension. It can straightforwardly be adapted to simulate the evolution of the permanent-income-weighted distribution: just change the probability distribution for permanent income shocks from f_η to $\tilde{f}_\eta$ when simulating the distribution.

3 Theoretical characterization of aggregate behavior in buffer-stock saving models

Theorem 2.1 is the main result of the paper. In this section, I show how the theorem allows a characterization of aggregate behavior in buffer-stock savings models, extending the work of Szeidl (2013) and Carroll (2020).

Consider the following model from Szeidl (2013), similar to Carroll (1997) and Haliassos and Michaelides (2003). There is a continuum of infinitely-lived households with stochastic labor income facing a menu of two financial securities: A risky investment S and a safe bond B. The household faces borrowing and short-sales constraints so that the portfolio shares are between zero and one for both assets. The agent solves the problem

$$\max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta_t \frac{C_t^{1-\gamma}}{1-\gamma} \qquad \text{s.t.} \qquad C_t + B_t + S_t = M_t,$$

$$M_{t+1} = S_t R_{t+1} + B_t R_f + Y_{t+1},$$

$$Y_{t+1} = P_{t+1}\epsilon_{t+1},$$

$$P_{t+1} = P_t G \eta_{t+1},$$

$$S_t \geqslant 0,$$

$$B_t \geqslant 0,$$

where M_t is the cash on hand of the household. The household's permanent income is denoted P_t and grows at rate G, subject to a permanent income shock η_t. The labor income is also subject to a transitory income shock ϵ_t. Both η_t and ϵ_t are non-negative, independent and i.i.d. with mean 1. The gross return on the

risky investment, R_{t+1}, is non-negative, i.i.d., independent of labor income, with $ER_{t+1} > R_f$ finite and the excess return $R_{t+1} - R_f$ assumes both non-negative and non-positive values with positive probability.

The consumer problem can be reformulated in terms of normalized variables $m_t = M_t/P_t, c_t = C_t/P_t$ and consumer behavior scales linearly with permanent income. Szeidl (2013) provides a characterization of when a stable invariant distribution of m_t, ψ^m, exists. The proof, mutatis mutandi, directly translates to a characterization of when a stable invariant permanent-income-weighted distribution, $\tilde{\psi}^m$, exists for this environment.

Consider the version of the above model without labor income (i.e., $G = 0$). It is well known that, without labor income, consumption is proportional to cash on hand and the portfolio share of the risk investment is constant. In this auxilliary model, let b^* denote the optimal share of consumption out of cash on hand, $C_t = b^* M_t$, let a^* denote the optimal share of risky investment, and write $R^*_{p,t+1} = a^* R_{t+1} + (1 - a^*)R_f$.

With these preliminaries, Szeidl (2013) proves the following result:

Proposition 3.1 (from Szeidl (2013))**.** *There exists a stable invariant marginal distribution ψ^m if*

$$E \log[R^*_{p,t+1}(1 - b^*)] < E \log[G\eta_{t+1}]. \tag{5}$$

Furthermore, if

$$E \log[R^*_{p,t+1}(1 - b^*)] > E \log[G\eta_{t+1}] \tag{6}$$

then there does not exist an invariant marginal distribution.

Using Theorem 2.1, the following proposition which is a simple extension of Szeidl (2013)'s result, provides a characterization of when a stable invariant permanent-income weighted distribution exists:

Proposition 3.2 (adapted from Szeidl (2013))**.** *There exists a stable invariant permanent-income weighted distribution $\tilde{\psi}^m$ if*

$$E \log[R^*_{p,t+1}(1 - b^*)] < \tilde{E} \log[G\eta_{t+1}] \tag{7}$$

where the expectation $\tilde{E}$ is taken with respect to the permanent-income-neutral measure given by the density function $\tilde{f}_\eta(\eta_{t+1}) = \eta_{t+1}f(\eta_{t+1})$. Furthermore, if

$$E \log[R^*_{p,t+1}(1 - b^*)] > \tilde{E} \log[G\eta_{t+1}] \tag{8}$$

then there does not exist an invariant permanent-income-weighted distribution.

Proof. By Theorem 2.1, the law of motion for ψ^m_t is the same as the law of motion for $\tilde{\psi}^m_t$, except f_η is

replaced by $\tilde{f}_\eta$. Therefore, the condition of Szeidl (2013) translates except the expectation is taken with respect to the permanent-income-neutral measure. □

Note that $\mathsf{E}\log[\mathsf{G}\eta_{t+1}] < \tilde{\mathsf{E}}\log[\mathsf{G}\eta_{t+1}]$ so the existence of a stable invariant marginal distribution is sufficient but not necessary to ensure the existence of a stable invariant permanent-income-weighted distribution. For some parameter values, there does not exist a stable invariant marginal distribution but there exists a stable invariant permanent-income-weighted distribution. Intuitively, a stable invariant marginal distribution does not exist if sufficiently many households have their normalized cash on hand $\mathfrak{m}_t$ increasing in an unbounded fashion. This happens for households who see their permanent income falling many times in a row. However, for the permanent-income-weighted distribution, these households contribute much less to the aggregate (since their permanent income fell) and therefore the condition for the existence of a stable invariant permanent-income-weighted distribution is less restrictive.

Armed with Proposition 3.2, we can now prove a conjecture from Carroll (2020).

Proposition 3.3 (conjecture from Carroll (2020))**.** *Under the conditions of Proposition 3.2, aggregate consumption grows at the same rate as permanent income in the long run.*

Proof. Recall that $C_t = \mathsf{G}^t \int c(\mathfrak{m}_t)\tilde{\psi}_t^m(\mathfrak{m}_t)d\mathfrak{m}_t$. Given the existence of a stable invariant permanent-income-weighted distribution $\tilde{\psi}^m$, in the long run we have

$$C_{t+1} = \mathsf{G}^{t+1}\int c(\mathfrak{m})\tilde{\psi}^m(\mathfrak{m})d\mathfrak{m} = \mathsf{G}\left(\mathsf{G}^t\int c(\mathfrak{m})\tilde{\psi}^m(\mathfrak{m})d\mathfrak{m}\right) = \mathsf{G}C_t \qquad (9)$$

so $\frac{C_{t+1}}{C_t} = \mathsf{G}$. □

Similarly, aggregate savings grows at the same rate as permanent income.

Following Carroll (2020), denote by $\mathbb{M}$ the cross-sectional average. We can compute aggregate consumption in two ways. First, we can use the objective measure,

$$C_t = \int c(\mathfrak{m}_t)P_t\psi_t(\mathfrak{m}_t, P_t)d\mathfrak{m}_t P_t = \mathbb{M}[c(\mathfrak{m}_t)P_t] = \mathsf{G}^t\mathbb{M}[c(\mathfrak{m}_t)] + \mathrm{Cov}_t(c(\mathfrak{m}_t), P_t). \qquad (10)$$

Second, we also have

$$C_t = \mathsf{G}^t\int c(\mathfrak{m}_t)\tilde{\psi}_t^m(\mathfrak{m}_t)d\mathfrak{m}_t = \mathsf{G}^t\tilde{\mathbb{M}}[c(\mathfrak{m}_t)] \qquad (11)$$

where $\tilde{\mathbb{M}}[\cdot]$ is taken with respect to the permanent-income-weighted distribution. Therefore, we get the result that

$$\mathrm{Cov}_t(c(\mathfrak{m}_t), P_t/\mathsf{G}^t) = -\left(\mathbb{M}[c(\mathfrak{m}_t)] - \tilde{\mathbb{M}}[c(\mathfrak{m}_t)]\right). \qquad (12)$$
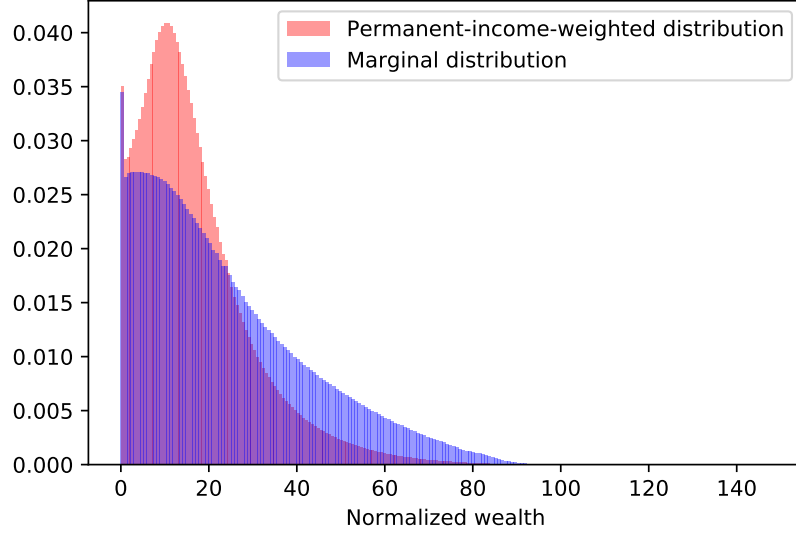
8

Figure 3.1: The steady-state permanent-income-weighted distribution of normalized wealth plotted together with the marginal distribution.

In other words, the covariance between permanent income and consumption is the gap in average normalized consumption between the marginal distribution and the permanent-income-weighted distribution.

In terms of the distribution of permanent income shocks, the permanent-income-neutral measure stochastically dominates the objective measure since it overweights the positive permanent income shocks and underweights the negative permanent income shocks. Therefore, the resulting dynamics in cash on hand under the permanent-income-neutral measure is stochastically dominated by the dynamics under the objective measure. Figure 3.1 shows the permanent-income-weighted distribution and the marginal distribution of normalized wealth from the Aiyagari model of Section 4, note that the permanent-income-weighted distribution has substantially less normalized wealth. Since the permanent-income-weighted distribution is stochastically dominated by the marginal distribution along the permanent-income dimension, the aggregate economy is more financially constrained than the average household in the economy.

4 Using the permanent-income-neutral measure for computations

To compute aggregate behavior of heterogeneous-agent models, an integral part is to compute model aggregates such as aggregate consumption and aggregate investment. By simulating the law of motion for the permanent-income-weighted distribution $\tilde{\psi}_t^m$ instead of the law of motion for the (full) distribution ψ_t, we reduce the dimensionality of the relevant state space needed to compute model aggregates. Furthermore,

9

since the permanent-income dimension of the state space is unbounded with the permanent-income distribution featuring a fat tail, eliminating this dimension is associated with sizeable computational improvements.[2] In this section, I solve an Aiyagari model with and without using the permanent-income-neutral measure.

The Aiyagari model is a minimal example where computation of cross-sectional aggregates is needed to solve for the equilibrium, and it therefore serves as an introduction to the application of the permanent-income-neutral measure. Note however that the method is applicable in all settings where cross-sectional aggregates are needed for the computation of equilibria, for example when computing aggregate dynamics in the presence of aggregate shocks using, e.g., the Krusell and Smith (1998) algorithm. I describe the economic environment of the Aiyagari model including parameter values, in Appendix A. The model is similar to Carroll et al. (2017) and in comparison with Aiyagari (1994), there are two differences. First, the income process features fully permanent income shocks. Second, to maintain a stationary income distribution, I introduce a perpetual-youth structure as in Blanchard (1985).

In the literature, e.g., McKay (2017), Carroll et al. (2017) and Carroll et al. (2019), it is necessary to introduce a perpetual-youth structure to maintain a stationary income distribution and compute aggregate variables. However, using the permanent-income-neutral measure, it is not *necessary* to have a well-defined income distribution in order to compute model aggregates.[3] I nonetheless introduce the perpetual-youth structure in order to compare simulating the model with and without the permanent-income-neutral measure.

To compute the steady-state equilibrium, we employ a simple bisection algorithm. Guess that the equilibrium value of aggregate capital K lies in $[K_{low}, K_{high}]$.

1. Set $K = 0.5K_{low} + 0.5K_{high}$.

2. Compute return on capital R and wage w given K.

3. Solve for the normalized consumption function $c(\cdot)$ using the endogenous-grid method.

4. Simulate a panel of households and compute the average value of end-of-period assets $P(m - c(m))$.

5. If $E[P(m - c(m))] > K$, set $K_{low} = K$. Else, set $K_{high} = K$.

6. Start over, and iterate until $K_{high} - K_{low} < \epsilon$.

The permanent-income-neutral measure helps with Step 4. The laws of motion for the full distribution ψ_t and the permanent-income-weighted distribution are given by Table 4.1. To quantify the computational improvement from using the permanent-income-weighted distribution, I simulate 1,000 households for N periods where $N = 10^3, 10^4, 10^5, 10^6$ using both the standard law of motion for ψ_t and the law of motion for the permanent-income-weighted distribution $\tilde{\psi}_t^m$. In Figure 4.1, the standard error of total savings is

---

[2]The permanent-income distribution features a fat tail in McKay (2017), Carroll et al. (2017) and Carroll et al. (2019).

[3]This is reminiscent of Constantinides and Duffie (1996)'s model of asset prices. In their model, the income distribution is degenerate but asset prices are well defined.

Distribution	Law of motion	Shock distributions
ψ_t	$m' = \epsilon' + R\frac{(m-c(m))}{\eta'}$	$\epsilon' \sim f_\epsilon$
	$P' = P\eta'$	$\eta' \sim f_\eta$
ψ_t^m	$m' = \epsilon' + R\frac{(m-c(m))}{\eta'}$	$\epsilon' \sim f_\epsilon$
		$\eta' \sim f_\eta$
$\tilde{\psi}_t^m$	$m' = \epsilon' + R\frac{(m-c(m))}{\eta'}$	$\epsilon' \sim f_\epsilon$
		$\eta' \sim \tilde{f}_\eta$

Table 4.1: The laws of motion for the joint distribution ψ_t of both normalized cash on hand and permanent income, the marginal distribution of normalized cash on hand ψ_t^m, and the permanent-income-weighted distribution of normalized cash on hand $\tilde{\psi}_t^m$. The permanent-income-neutral measure is given by $\tilde{f}(\eta) = \eta f(\eta)$.

plotted against the number of periods.[4] The standard error is decreasing in the number of periods for both the permanent-income-neutral simulation and the standard simulation, with the standard error proportional to the inverse square root of the number of periods. There is however a large gap between the standard errors of the two simulation methods, the standard error of the estimate of total savings using the permanent-income-neutral measure is only 17% of the standard error using the standard simulation method. Therefore, to achieve a given precision using the permanent-income-neutral measure one only needs a sample of size $0.17^2 = 3\%$ of the size required using the standard simulation method. In other words, the speedup for the simulation is more than thirtyfold.

This speedup is substantial since most of the computation time is spent simulating the model, as shown in Table 4.2.[5] I solve the model using the algorithm outlined above with (i) the endogenous-grid method with 300 grid points, and (ii) Monte-Carlo simulation using 1000 households for 10,000 periods using the permanent-income-neutral measure. The standard error of the estimate of aggregate saving is 0.05 and the bisection algorithm takes two minutes on a MacBook Pro. For comparison, I then solve the model without the permanent-income-neutral measure, with 1000 households for 339,000 periods. The standard error of the estimate of aggregate saving is still 0.05 and the bisection algorithm takes an hour on a MacBook Pro.

Compatibility with other simulation methods   In the computational exercise, the simulation was conducted using stochastic Monte-Carlo simulation to highlight that the permanent-income-neutral measure is trivial to use. An alternative approach to simulation is non-stochastic simulation (Young, 2010), where

---

[4]View the 1,000 households as random draws. The standard error is then the cross-sectional standard deviation of average capital divided by $\sqrt{1000}$.

[5]That it takes a long time to simulate the model rather than solving for optimal behavior is consistent with what Carroll et al. (2019) write in the documentation of their code, which uses an improved variant of Monte-Carlo simulation: "The HA-DSGE mode takes roughly 16 hours to solve the sticky and frictionless versions, and will display progress toward finding the equilibrium aggregate saving rule after each iteration. The very long run time for HA-DSGE is largely due to the extremely large number of periods."

Figure 4.1: The standard error of the estimate of aggregate savings from the simulation of 1,000 households for N periods where $N = 10^3, 10^4, 10^5, 10^6$, for both the standard simulation method and using the permanent-income-neutral measure.

	P-I-N simulation	Standard simulation
Computation time endogenous-grid method	1.8s	1.6s
Computation time simulation	119.0s	3610.3s
Total computation time	120.8s	3612.0s

Table 4.2: Computation times for solving the Aiyagari model with and without the permanent-income-neutral measure. The number of periods for the standard simulation is 33.9 times the number of periods for the permanent-income-neutral simulation, to keep the standard error of the estimate of aggregate savings constant.

the transition dynamics are summarized by a discretized transition matrix. One of the advantages of non-stochastic simulation is that the stationary distribution can be computed as the eigenvector associated with the largest eigenvalue of the transition matrix. The major potential drawback with non-stochastic simulation is that it suffers from a curse of dimensionality. As the dimensionality of the state space N grows, the size of the transition matrix N^2 grows even faster. Although there are ways around this curse of dimensionality, see in particular Tan (2020), using the permanent-income-neutral measure is a trivial way to reduce the dimensionality of the state space.

5 Discussion

The role of the permanent-income-neutral measure is analogous to the role of the risk-neutral measure in asset pricing. In asset pricing, the price of an asset, e.g. a stock, S depends on the payoff d and the stochastic discount factor Λ,

$$\underbrace{S}_{\text{Price}} = E[\Lambda d] = \underbrace{R^{-1}E[d]}_{\text{Discounted expected return}} + \underbrace{\text{cov}(\Lambda, d)}_{\text{Risk premium}} \tag{13}$$

where $R = 1/E[\Lambda]$. Notice the structural similarity with Equation 10. In asset pricing, the main challenge is the covariance between the stochastic discount factor and the payoff, i.e. pricing risk. In heterogeneous-agent macroeconomics, the main difficulty is the covariance between permanent income and the normalized state. In both cases, it helps to perform a change of measure to the risk-neutral measure/permanent-income-neutral measure.

Lately, following Achdou et al. (2014), there has been an explosion of work with heterogeneous-agent models in continuous time. How can we use the permanent-income-neutral measure in this setting? The mathematical machinery necessary, Girsanov's theorem, is well known in mathematical finance and directly applicable (for a textbook treatment of Girsanov's theorem for economists, see Björk (2009)). Let permanent income follow a geometric Brownian motion, $dP_t = gP_t dt + \sigma P_t dW_t$. Aggregate consumption is given by $E[P_t c(m_t)]$. Girsanov's theorem states that $E[P_t c(m_t)] = e^{gt} E^Q[c(m_t)]$ where the dynamics under the equivalent-martingale measure Q is given by replacing $dW_t = \sigma dt + dW^Q$. Therefore, to simulate the model under the permanent-income-neutral measure, formally replace dW_t by $\sigma dt + dW^Q$ in the stochastic differential equation for the evolution of m_t.

To conclude, the permanent-income-neutral measure both provides a simple computational improvement for simulating heterogeneous-agent models with permanent income shocks and helps clarify the theoretical properties of these models. The improvement in computational performance is as close to a free lunch as possible, since it only requires replacing f_η with $\tilde{f}_\eta$ in pre-existing code for simulating the model.

References

Achdou, Y., Han, J., Lasry, J.M., Lions, P.L., Moll, B., 2014. Heterogeneous Agent Models in Continuous Time.

Aiyagari, S.R., 1994. Uninsured Idiosyncratic Risk and Aggregate Saving. The Quarterly Journal of Economics 109, 659–684. URL: http://qje.oxfordjournals.org/content/109/3/659.abstract.

Björk, T., 2009. Arbitrage Theory in Continuous Time. 3rd ed., Oxford University Press.

Blanchard, O.J., 1985. Debt, Deficits, and Finite Horizons. Journal of Political Economy 93, 223–247.

Campbell, J.Y., Cocco, J.F., 2007. How do house prices affect consumption? Evidence from micro data. Journal of Monetary Economics 54, 591–621. doi:10.1016/j.jmoneco.2005.10.016.

Carroll, C., Crawley, E., Slacalek, J., Tokuoka, K., White, M., 2019. Sticky Expectations and Consumption Dynamics. Mimeo doi:10.5281/zenodo.1001068.

Carroll, C.D., 1997. Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis. Quarterly Journal of Economics 112, 1–55.

Carroll, C.D., 2020. Theoretical Foundations of Buffer Stock Saving. Quantitative Economics doi:10.3386/w10867.

Carroll, C.D., Slacalek, J., Tokuoka, K., White, M.N., 2017. The distribution of wealth and the marginal propensity to consume. Quantitative Economics 8, 977–1020. doi:10.3982/QE694.

Constantinides, G.M., Duffie, D., 1996. Asset pricing with heterogeneous consumers. Journal of Political Economy 104, 219–240. doi:10.1086/262023.

Deaton, A., 1991. Saving and Liquidity Constraints. Econometrica 59, 1221–1248.

Gourinchas, P.O., Parker, J.A., 2002. Consumption over the Life Cycle. Econometrica 70, 47–89.

Haliassos, M., Michaelides, A., 2003. Portfolio Choice and Liquidity Constraints. International Economic Review 44, 143–177.

Harmenberg, K., Öberg, E., 2019. Consumption Dynamics under Time-Varying Unemployment Risk.

Krusell, P., Smith, A.A., 1998. Income and Wealth Heterogeneity in the Macroeconomy. Journal of Political Economy 106, 867–896.

McKay, A., 2017. Time-varying idiosyncratic risk and aggregate consumption dynamics. Journal of Monetary Economics 88, 1–14. doi:10.1016/j.jmoneco.2017.05.002.

Szeidl, A., 2013. Stable Invariant Distribution in Buffer-Stock Saving and Stochastic Growth Models.

Tan, E., 2020. A fast and low computational memory algorithm for non-stochastic simulations in heterogeneous agent models. Economics Letters 193, 109285. URL: https://doi.org/10.1016/j.econlet.2020.109285, doi:10.1016/j.econlet.2020.109285.

Young, E.R., 2010. Solving the incomplete markets model with aggregate uncertainty using the Krusell-Smith algorithm and non-stochastic simulations. Journal of Economic Dynamics and Control 34, 36–41. URL: http://dx.doi.org/10.1016/j.jedc.2008.11.010, doi:10.1016/j.jedc.2008.11.010.

Zeldes, S.P., 1989. Optimal consumption with stochastic income: Deviations from certainty equivalence. Quarterly Journal of Economics 104, 275–297. doi:10.2307/2937848.

A   Appendix

The household consumption-saving problem is

$$\max_{\{B_t, C_t\}_{t=0}^{\infty}} E_0 \sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\gamma}}{1-\gamma} \qquad \text{s.t.} \qquad B_t + C_t = Y_t + RB_{t-1}, \qquad (14)$$

$$Y_t = P_t \epsilon_t, \qquad (15)$$

$$P_t = P_{t-1} \eta_t, \qquad (16)$$

with $B_{-1}, P_{-1}, \epsilon_0, \eta_0$ given at time 0 and $\epsilon_t \sim F_\epsilon$, $\eta_t \sim F_\eta$ i.i.d with $E\eta = E\epsilon = 1$.

There is a continuum of households solving the household problem described by Equations 14-16. To maintain a stationary income distribution, we employ a perpetual-youth structure. All households face a probability ω of dying each period. When a household dies, it is replaced by a newborn with no initial assets and permanent income $P = 1$. Total capital K is given by the cross-sectional average bond holdings $E[B_t]$, or equivalently $E[P_t(a_t - c(a_t))]$. The wage w and gross return on capital R are determined by a Cobb-Douglas production function $Y = K^\alpha$.

A steady state equilibrium is given by

i A consumption function $c(\cdot)$ that solves Equations 14-16 given R and w,

ii A stationary distribution of households ψ over the state space $\mathbf{m} \times \mathbf{P}$ generated by the shock distributions

f_η, f_ϵ, the probability ω of death $\chi = 1$, and the micro dynamics

$$m' = \begin{cases} w\epsilon' + R\frac{(m-c(m))}{\eta'} & \chi' = 0, \\ w\epsilon' & \chi' = 1, \end{cases}$$

$$P' = \begin{cases} P\eta' & \chi' = 0, \\ 1 & \chi' = 1. \end{cases}$$

iii Market clearing for capital and factor-price determination,

$$K = \int_{\mathbf{m}\times\mathbf{P}} (m - c(m))P \, dm \, dP,$$

$$R = \frac{\alpha K^{\alpha-1} + (1-\delta)}{1-\omega},$$

$$w = (1-\alpha)K^\alpha.$$

The return R is multiplied by a factor $1/(1-\omega)$, with the interpretation that the assets of a dead household is distributed among the living households.

A.1   Parameters

The subjective discount factor is set to $\beta = 0.97$ and risk aversion is set to $\gamma = 1$. The transitory shock is log-normally distributed with $\sigma_\epsilon = 0.158$ and the permanent shock is log-normally distributed with $\sigma_\eta = 0.073$. Capital depreciation is set to $\delta = 0.025$ and the capital share is set to $\alpha = 0.33$. The death rate is set to $\omega = .005$.