# Aggregating Heterogeneous-Agent Models with Permanent Income Shocks*

Karl Harmenberg†

April 8, 2020

**Abstract**

I propose 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 counterfactual *permanent-income-neutral* probabilities which incorporate the effect that permanent-income shocks have on macroeconomic aggregates. With the permanent-income-neutral probabilities, one does not need to keep track of the permanent-income distribution. The method can implemented with a few lines of code and greatly improves the speed and accuracy of the simulation. Furthermore, the permanent-income-neutral probability distribution is useful for the analytical characterization of aggregate consumption-savings behavior in this class of models.

## 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, requiring the modeler to weigh computational tractability against a realistic description of the economic environment. To numerically solve a heterogeneous-agent model requires three distinct steps, all of which can be a computational obstacle. First, postulate and compute behavior at the micro level using, e.g., value-function iteration. Second, aggregate the micro-level behavior into macroeconomic aggregates, either by stochastic Monte-Carlo simulation or non-stochastic simulation. Third, specify general equilibrium and compute equilibrium dynamics using e.g., the Krusell-Smith algorithm or linearization.

This paper is about the second step, aggregation. When solving for general equilibrium using, e.g., the Krusell and Smith (1998) algorithm, it is required to repeatedly aggregate and simulate the micro behavior

*I am grateful for helpful comments from Jeppe Druedahl and Erik Öberg.

†kha.eco@cbs.dk, Copenhagen Business School

into macro aggregates and in applications this step can be a major computational bottleneck. I provide a way to easily perform this aggregation for heterogeneous-agent models where log permanent income follows a random walk. The method uses the *permanent-income-neutral* measure, distorted probabilities which reduce the dimensionality of the state space. Implementing the method is trivial and amounts to changing a few lines of code when simulating the economy. As a computational example, I solve for equilibrium in a simple Aiyagari environment and show that the permanent-income-neutral measure speeds up the computation by an order of magnitude.

For the method to be applicable, it is important that household behavior scales with permanent income. With standard CRRA preferences and log permanent income following a random walk, households' decisions only depend on their asset positions as a share of the household's permanent income. Since the seminal work by Deaton (1991) and Carroll (1992, 1997), the combination of CRRA preferences and a random-walk structure for log permanent income has been used in many partial-equilibrium life-cycle models, including Gourinchas and Parker (2002) and Cocco, Gomes, and Maenhout (2005). Recently, McKay (2017), Carroll, Slacalek, Tokuoka, and White (2017) and Carroll, Crawley, Slacalek, Tokuoka, and White (2019) study aggregate consumption dynamics of heterogeneous-agent models using income processes with a random-walk structure for permanent income. For these models, the simulation method using the permanent-income-neutral measure is directly applicable, with large potential gains in speed.

Most of the heterogeneous-agent macroeconomic literature postulate a persistent but not fully permanent component of the households' income process. For example, in their handbook chapter on macroeconomics and household heterogeneity, Krueger, Mitman, and Perri (2016) model income as subject to (i) unemployment shocks, (ii) transitory shocks, and (iii) persistent shocks (with autocorrelation $\rho = 0.99$). In the recent HANK literature, McKay et al. (2016), Bayer et al. (2019), Auclert et al. (2020), and Kaplan et al. (2020) all model household income as an $AR(1)$ in logs with high persistence and the income process of Kaplan, Moll, and Violante (2018) features a rare income shock with a half-life of 18 years. Although the method of the permanent-income-neutral measure is not directly applicable to these models, an economic modeler may want to trade off the simplifying assumption of fully permanent income shocks (i.e., $\rho = 1$) against the improvement in computational tractability associated with the permanent-income-neutral measure.[1]

For the household problem, the key computational advantage of models with permanent income shocks is that household behavior can be normalized with respect to permanent income. The relevant state for households is "cash on hand" as a multiple of their permanent income, and their consumption takes the form $C = c(m)P$ where $m = M/P$ is their cash on hand $M$ divided by permanent income $P$. When computing optimal behavior, we can eliminate permanent income as a state variable, with both conceptual clarity and computational gains as a result. Although the class of models where permanent income follows a random

[1]Druedahl and Jørgensen (2017) show that economic behavior, such as the marginal propensity to consume, is largely unaffected by a jump from persistent income shocks to permanent income shocks.

walk have this tractable property at the micro level, they have not been tractable with respect to aggregation at the macro level. Normalized household consumption $c(m)$ needs to be weighted by permanent income $P$ when computing total consumption, $\bar{C} = \mathbb{E}[c(m)P]$, where $\mathbb{E}[\cdot]$ denotes the household cross-sectional average. We therefore seemingly need to keep track of the joint distribution of both $m$ and $P$ when simulating the model, despite being able to eliminate permanent income $P$ as a state variable when computing optimal behavior at the micro level. Furthermore, the random walk in permanent income implies that there is no stationary permanent-income distribution and even in finite-horizon settings such as life-cycle models, the permanent-income distribution becomes very dispersed. Therefore, it can be challenging to keep track of the joint distribution of permanent income and normalized cash on hand.

The contribution of this paper is to show that there is a way to aggregate and simulate this class of models without needing to keep track of the joint distribution of both $m$ and $P$. The key insight is that, rather than describing the law of motion of households, we can describe the law of motion for "units of permanent income". The method amounts to simulating the model with distorted *permanent-income-neutral* probabilities instead of the objective probabilities. Implementation is trivial in discrete-time and continuous-time applications, both for stochastic Monte-Carlo simulation and non-stochastic simulation à la Young (2010), yielding order-of-magnitude improvements.

The fundamental idea of the permanent-income-neutral probabilities is a *change of probability measure*. There is a close analogy between the permanent-income-neutral measure in our setting and the risk-neutral measure in asset pricing. In both cases, a change of measure eliminates a state variable with computational gains as a result. Beyond the computational advantages, the shift to the permanent-income-neutral measure is also theoretically clarifying. In Section 3, I adopt the theoretical results of Carroll (2020) and Szeidl (2013) to analytically characterize aggregate consumption behavior.

## 2 The permanent-income-neutral measure

### 2.1 Setting

**Household problem** To fix ideas, we introduce a minimal model featuring a random walk in permanent income following Deaton (1991), Carroll (1992, 1997). The household consumption-saving problem is

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

$$Y_t = P_t \epsilon_t, \tag{2}$$

$$P_t = P_{t-1} \eta_t, \tag{3}$$

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 $E\eta = E\epsilon = 1$.

The assumptions of CRRA utility together with a log random walk in permanent income yield that optimal consumption in this model depends on "normalized cash-on-hand" $m_t = \frac{Y_t + RB_{t-1}}{P_t}$ and is linear in permanent income $P_t$, $C_t = c(m_t)P_t$. It is straightforward to compute the normalized consumption function $c(\cdot)$ using value-function iteration or the endogenous-grid method. The existence of a normalized consumption function $c(\cdot)$ gives computational tractability since permanent income can be eliminated as a state variable when computing the household decision functions.

**Simulating a household at the micro level** At the micro level, the consumption function $c(\cdot)$ together with the shock distributions are sufficient to simulate an individual household.

To concretely describe the law of motion for an individual household, consider a household at state $(m, P)$. The household's normalized consumption and savings are given by $c(m)$ and $b = m - c(m)$ respectively. Therefore, consumption and savings are given by $C = Pc(m)$ and $B = P(m - c(m))$ respectively.

In the next period, the household receives a transitory income shock $\epsilon' \sim F_\epsilon$ and a permanent income shock $\eta' \sim F_\eta$. Therefore, the next-period permanent income is $P' = \eta' P$ and next-period income is $Y' = \epsilon' P'$ so next-period cash on hand is $M' = Y' + RB = \epsilon' P' + R(m - c(m))P$ and normalized cash on hand is $m' = M'/P' = \epsilon' + R\frac{m - c(m)}{\eta'}$. The resulting next-period state is $(m', P') = (\epsilon' + R\frac{m-c(m)}{\eta'}, \eta' P)$.

To summarize, the law of motion is given by

$$m' = \epsilon' + R\frac{(m - c(m))}{\eta'}, \tag{4}$$

$$P' = P\eta', \tag{5}$$

together with a specification of the probability distributions,

$$\mathbb{P}(\epsilon' = \epsilon_i) = p_i^\epsilon, \tag{6}$$

$$\mathbb{P}(\eta' = \eta_j) = p_j^\eta. \tag{7}$$

## 2.2 Aggregation of micro behavior

The focus of this paper is how to aggregate the micro behavior described by Equations 4-7 into aggregate behavior. In this section, I describe two relevant macro aggregates, average marginal propensity to consume, which aggregates trivially, and aggregate consumption, which does not. The main result of the paper is that, using permanent-income-neutral probabilities, aggregate consumption, aggregate savings, and many other moments can be aggregated as easily as marginal propensity to consume.

At time $t$, there is a continuum $\mu_t$ of households in the state space $\mathbf{m} \times \mathbf{P}$, where $\mathbf{m}$ is the normalized cash-on-hand dimension and **P** is the permanent-income dimension. We now set out to compute average marginal propensity to consume and total consumption for this distribution of households.

**Statically aggregating** We first set out to compute average MPC for a continuum of households in a given period. Marginal propensity to consume for a given household is $\mathsf{MPC} := \frac{d(c(M/P)P)}{dM} = c'(m)$. Average marginal propensity to consume is given by the integral of $\mathsf{MPC}$ over the state space,

$$\overline{\mathsf{MPC}}_t = \int_{\mathbf{m}\times\mathbf{P}} c'(m)\mu_t(m, P)dmdP.$$

We can rewrite this integral further as

$$\overline{\mathsf{MPC}}_t = \int_{\mathbf{m}} c'(m)\left(\int_{\mathbf{P}} \mu_t(m, P)dP\right) dm = \int_{\mathbf{m}} c'(m)\mu_t^{\mathbf{m}}(m)dm.$$

where $\mu_t^{\mathbf{m}}$ is the marginal distribution of households along the normalized cash-on-hand dimension. We see that average marginal propensity to consume only depends on the normalized-asset dimension.

In contrast with average marginal propensity to consume, aggregate consumption depends on the distribution of households both along the **m** dimension and **P** dimension. Consumption for a given household is $C = c(a)P$. Total consumption is therefore given by

$$\overline{C}_t = \int_{\mathbf{m}\times\mathbf{P}} c(m)P\mu(m, P)dmdP.$$

We can rewrite this integral further as

$$\overline{C}_t = \int_{\mathbf{m}} c(m)\left(\int_{\mathbf{P}} \mu_t(m, P)PdP\right) dm = \int_{\mathbf{m}} c(m)\tilde{\mu}_t^{\mathbf{m}}dm$$

where $\tilde{\mu}_t^{\mathbf{m}}$ is the permanent-income weighted marginal distribution of households along the normalized cash-on-hand dimension.

For both the average marginal propensity to consume and total consumption, there is a one-dimensional sufficient statistic of the two-dimensional distribution of households. For average marginal propensity to consume, the sufficient statistic is the distribution of households along the normalized-asset dimension, $\mu_t^{\mathbf{m}}$. For total consumption, the sufficient statistic is the distribution of households, weighted by their permanent income, along the normalized-asset dimension, $\tilde{\mu}_t^{\mathbf{m}}$. The unweighted distribution $\mu_t^{\mathbf{m}}$ is a democratic measure where all households are equally important. For average marginal propensity to consume, the democratic weighting of households is the right one, since average marginal propensity to consume tells us how much consumption would increase (in partial equilibrium) if all households were given 100 dollars. By contrast, the weighted distribution $\tilde{\mu}^{\mathbf{m}}$ is a plutocratic measure where households are weighted according to their permanent income. For the classical macroeconomic aggregates such as aggregate saving and aggregate consumption, households contribute to the aggregate in proportion to their permanent income.

**Law of motion for the sufficient statistics** One way of computing the dynamics of aggregate variables is to compute the law of motion for the distribution of households $\mu_t$ in the state space $\mathbf{m} \times \mathbf{P}$, and in every period compute the aggregate variables from the distribution $\mu_t$. The law of motion for $\mu_t$ is implied by Equations 4-5 together with the probability distributions of Equations 6-7. In terms of computation, the law of motion can be approximated by (stochastic) Monte-Carlo simulation.

Although this approach is straightforward to implement, it is computationally costly since the evolution both in the permanent-income dimension and the normalized cash-on-hand dimension needs to be tracked. The distribution of permanent income does in general have unbounded support with a fat tail, which means that outliers in the permanent-income dimension can disproportionately affect the macroeconomic aggregates. For example, both Carroll (2020) and McKay (2017) consider infinite-horizon models with perpetual-youth households, which generate a Pareto distribution of permanent income in steady state. With a fat tail in the permanent-income dimension, the variance of permanent income may be infinite and even if it is finite, there is in practice slow convergence to the true mean.[2]

It is advantageous to describe a law of motion directly for the sufficient statistics $\mu_t^{\mathbf{m}}$ and $\tilde{\mu}_t^{\mathbf{m}}$ since we can then reduce the dimensionality of the problem. There is an obvious alternative to tracking $\mu_t$ if we want to track the evolution of average marginal propensity to consume. Since the marginal distribution $\mu_t^{\mathbf{m}}$ is a sufficient statistic for average marginal propensity to consume, we only need to keep track of the normalized-asset dimension. The law of motion for $\mu_t^{\mathbf{m}}$ is described by Equation 4 together with the probability distributions of Equations 6-7. We can therefore compute the evolution of the marginal distribution $\mu_t^{\mathbf{m}}$ directly and forget about the full distribution $\mu_t$.

For describing the law of motion for $\tilde{\mu}_t^{\mathbf{m}}$, the sufficient statistic for computing aggregates such as total consumption and total savings, we cannot ignore the permanent-income dimension since we need to weigh households by their permanent income. The main result of the paper is that, despite this obstacle, there is a straightforward way of describing the law of motion for $\tilde{\mu}_t^{\mathbf{m}}$.

**Theorem 1.** *The law of motion for the permanent-income weighted distribution $\tilde{\mu}_m$ is given by*

$$m' = \epsilon' + R\frac{(m - c(m))}{\eta'}$$

[2] The granularity literature, e.g., Gabaix (2011) and Carvalho and Grassi (2019), explore whether fat tails together with finite populations and multiplicative idiosyncratic shocks can generate aggregate business cycle dynamics.

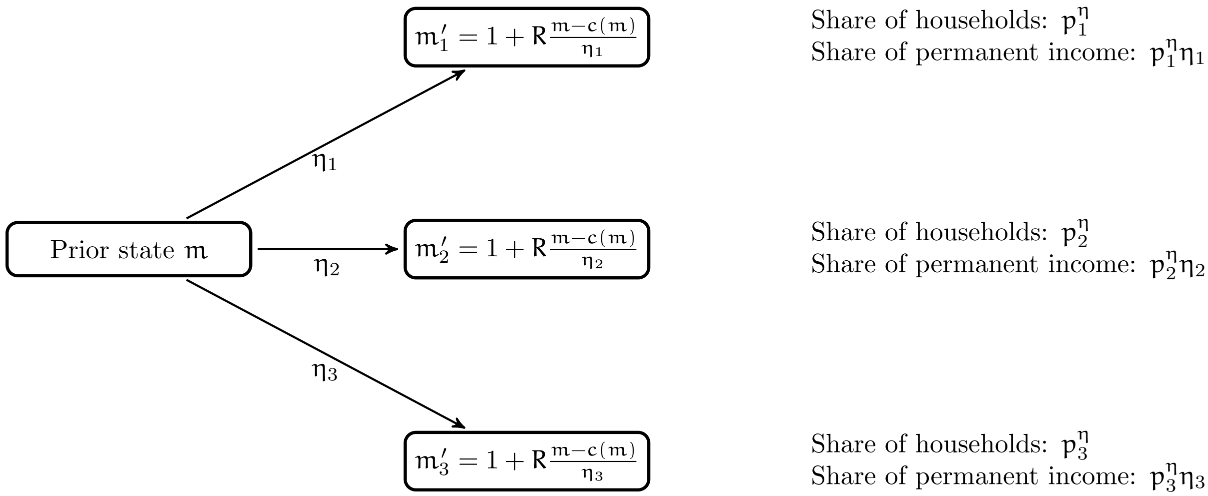


Figure 2.1: Intuition behind the permanent-income-neutral measure. The permanent-income-neutral measure weighs a shock not by the share of households hit by the shock, but by the share of permanent income that they account for.

*together with the probability weights*

$$\mathbb{P}(\epsilon' = \epsilon_i) = p_i^\epsilon, \tag{8}$$

$$\mathbb{P}(\eta' = \eta_j) = \eta_j p_j^\eta. \tag{9}$$

*We call the counterfactual probability weights* $\eta_j p_j^\eta$ *the* permanent-income-neutral *probabilities.*

*Proof.* The proof, which is relegated to the appendix, amounts to little more than accounting using the language of measure theory. □

**Remark 1.** *Note that the permanent-income-neutral probabilities sum to* 1 *since* $\mathbb{E}\eta = 1$.

For the intuition behind Theorem 1, consider for simplicity an environment without transitory income shocks. For such an environment, Figure 1 depicts how households transition from normalized cash on hand $m$ to their future normalized cash on hand $m'$. Of the households with normalized cash on hand $m$, a share $p_j^\eta$ receive the permanent-income shock $\eta_j$ and transition to $m_j'$. For macroeconomic aggregates such as aggregate consumption, we also need to weigh the households by their importance for the aggregates, therefore we need to keep track of how "units of permanent income" transition. Of the units of permanent income at state $m$, a share $p_j^\eta \eta_j$ ends up at $m_j'$. This share $p_j^\eta \eta_j$ is a composite of the share $p_j^\eta$ of the households, and these households' new weight in the aggregate which is adjusted by $\eta_j$.

Using Theorem 1, we can simulate the evolution of the permanent-income-weighted distribution $\tilde{\mu}_t^m$ by Monte-Carlo simulation. The only adjustment, compared to simulating the evolution of the (unweighted)

| Distribution | Law of motion | Probabilities/Weights |
|---|---|---|
| $\mu_t$ | $m' = \epsilon' + R\frac{(m-c(m))}{\eta'}$<br>$P' = P\eta'$ | $\mathbf{P}(\epsilon' = \epsilon_i) = p_i^{\epsilon}$,<br>$\mathbf{P}(\eta' = \eta_j) = p_j^{\eta}$ |
| $\mu_t^{\mathbf{m}}$ | $m' = \epsilon' + R\frac{(m-c(m))}{\eta'}$ | $\mathbf{P}(\epsilon' = \epsilon_i) = p_i^{\epsilon}$,<br>$\mathbf{P}(\eta' = \eta_j) = p_j^{\eta}$ |
| $\tilde{\mu}_t^{\mathbf{m}}$ | $m' = \epsilon' + R\frac{(m-c(m))}{\eta'}$ | $\mathbf{P}(\epsilon' = \epsilon_i) = p_i^{\epsilon}$,<br>$\mathbf{P}(\eta' = \eta_j) = \eta_j p_j^{\eta}$ |

Table 2.1: The laws of motion for the joint distribution $\mu_t$ of both normalized cash on hand and permanent income, the (unweighted) distribution of normalized cash on hand $\mu_t^{\mathbf{m}}$, and the permanent-income-weighted distribution of normalized cash on hand $\tilde{\mu}_t^{\mathbf{m}}$.

marginal distribution $\mu_t^{\mathbf{m}}$, is that the permanent-income shocks are drawn using the permanent-income-neutral probabilities $\{\tilde{p}_j^{\eta}\} = \{\eta_j p_j^{\eta}\}$. We summarize the laws of motion for $\mu_t$, $\mu_t^{\mathbf{m}}$, and $\mu_t^{\mathbf{m}}$ in Table 2.1.

## 2.3 Analogy to the risk-neutral measure

The mathematical idea behind Theorem 1, a *change of measure*, has been used in asset pricing where it is often useful to recast a problem using the risk-neutral measure (see, e.g., Cochrane (2005) and Björk (2009)). In our setting, the permanent-income-neutral measure combines the permanent income shock $\eta_j$ and the objective probability $p_j$. In asset pricing, the risk-neutral measure combines the stochastic discount factor $\Lambda_j$ and the objective probability $p_j$. To see the correspondence, note that in asset pricing, the price of an asset $P$ is given by

$$P = \frac{1}{R}\mathbb{E}[\Pi] + \mathrm{Cov}(\Pi, \Lambda)$$

where $\Pi$ is the payoff of the asset and $\Lambda$ is the stochastic discount factor. Correspondingly, aggregate consumption in our setting is given by

$$\overline{C} = \mathbb{E}[c(m)] + \mathrm{Cov}(c(m), P).$$

In both settings, the covariance term is the main obstacle for aggregation, and a change of measure helps with describing the dynamics of asset prices/aggregate consumption.

## 2.4 Some generality

Consider the following more general model with possibly time-varying and stochastic interest rate $R_t$ and wage $w_t$, possibly time-varying and stochastic growth $G_t$ in permanent income as well as possibly time-

| Distribution | Law of motion | Probabilities/Weights |
|---|---|---|
| $\mu_t$ | $m_{t+1} = w_{t+1}\epsilon_{t+1} + \frac{R_t}{G_{t+1}}\frac{(m-c_t(m))}{\eta_{t+1}}$<br>$P_{t+1} = G_{t+1}P\eta_{t+1}$ | $\mathbb{P}(\epsilon' = \epsilon_{i,t}) = p^{\epsilon}_{i,t}$,<br>$\mathbb{P}(\eta' = \eta_{j,t}) = p^{\eta}_{j,t}$ |
| $\mu^{\mathbf{m}}_t$ | $m_{t+1} = w_{t+1}\epsilon_{t+1} + \frac{R_t}{G_{t+1}}\frac{(m-c_t(m))}{\eta_{t+1}}$ | $\mathbb{P}(\epsilon' = \epsilon_{i,t}) = p^{\epsilon}_{i,t}$,<br>$\mathbb{P}(\eta' = \eta_{j,t}) = p^{\eta}_{j,t}$ |
| $\tilde{\mu}^{\mathbf{m}}_t$ | $m_{t+1} = w_{t+1}\epsilon_{t+1} + \frac{R_t}{G_{t+1}}\frac{(m-c_t(m))}{\eta_{t+1}}$ | $\mathbb{P}(\epsilon' = \epsilon_{i,t}) = p^{\epsilon}_{i,t}$,<br>$\mathbb{P}(\eta' = \eta_{j,t}) = \eta_{j,t}p^{\eta}_{j,t}$. |

Table 2.2: The laws of motion for the joint distribution $\mu_t$ of both normalized cash on hand and permanent income, the (unweighted) distribution of normalized cash on hand $\mu^{\mathbf{m}}_t$, and the permanent-income-weighted distribution of normalized cash on hand $\tilde{\mu}^{\mathbf{m}}_t$, for the general model of Subsection 2.4.

varying and stochastic shock distributions $F_{\epsilon,t}, F_{\eta,t}$.

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

$$Y_t = w_t P_t \epsilon_t,$$

$$P_t = G_t P_{t-1} \eta_t,$$

with $B_{-1}, P_{-1}, \epsilon_0, \eta_0$ given at time 0 and $\epsilon_t \sim F_{\epsilon,t}$, $\eta_t \sim F_{\eta,t}$ i.i.d with $E\eta_t = E\epsilon_t = 1$. For this generalized model, dynamics are simply adjusted as indicated by Table 2.2.

For a model with multiple assets, for example a durable good subject to non-convex adjustmet costs as in Harmenberg and Öberg (2019), we can apply the permanent-income-neutral measure by interpreting $m$ as a state vector and $c(m)$ as a decision vector. As long as the micro-level behavior scales linearly with permanent income, we can readily apply the permanent-income-neutral measure.

The argument also extends to continuous time using Girsanov's theorem, which describes how to shift to an equivalent-martingale measure under Brownian motion in continuous time. Appendix B describes how to shift to the permanent-income neutral measure in a continuous-time setting.

The method of the permanent-income-neutral measure can be adopted to compute higher order moments such as consumption variance. To compute the variance of consumption, first compute $\mathbb{E}[C]$ using the permanent-income-neutral measure. Define $P_{sq} = P^2$. Note that $P_{sq}$ grows at rate $G_{sq} = \mathbb{E}[\eta^2]$. Define $\eta_{sq} = \eta^2/G_{sq}$. Then, $\mathbb{E}[C^2] = \mathbb{E}[P_{sq}c(a)^2]$ and we can apply the permanent-income-neutral method with the shocks $\eta_{sq}$ and growth rate $G_{sq}$ to compute $\mathbb{E}[C^2]$. The consumption variance is then computed using the formula $\sigma_C^2 = \mathbb{E}[C^2] - \mathbb{E}[C]^2$. Higher order moments such as skewness and kurtosis are computed in an analogous fashion.

# 3 Existence of a stable invariant permanent-income-weighted distribution

A shift to the permanent-income-neutral measure allows us to translate the analytical characterization of the micro behavior in buffer-stock savings models from Szeidl (2013) and Carroll (2020) into an analytical characterization of the macro behavior of buffer-stock savings models.

Szeidl (2013) showed that the household problem described by Equations 1-3, which implies a law of motion for the marginal distribution $\mu^{\mathbf{m}}$ given by

$$m' = \epsilon' + \frac{R}{G}\frac{(m - c(m))}{\eta'}$$

together with the probability weights

$$\mathbb{P}(\epsilon' = \epsilon_i) = p_i^{\epsilon},$$
$$\mathbb{P}(\eta' = \eta_j) = p_j^{\eta},$$

generates a stable invariant marginal distribution $\mu^{\mathbf{m}}$ if

$$\log\left[(R\beta)^{1/\gamma}\right] < \log G + \mathbb{E}\log\eta.$$

If the inequality is reversed, then a stable invariant marginal distribution does not exist.

Adopting his argument, the household problem described by Equations 1-3, which implies a law of motion for the permanent-income weighted distribution $\tilde{\mu}^{\mathbf{m}}$ given by

$$m' = \epsilon' + \frac{R}{G}\frac{(m - c(m))}{\eta'}$$

together with the probability weights

$$\mathbb{P}(\epsilon' = \epsilon_i) = p_i^{\epsilon},$$
$$\mathbb{P}(\eta' = \eta_j) = \eta_j p_j^{\eta},$$

generates a stable invariant permanent-income weighted distribution $\tilde{\mu}^{\mathbf{m}}$ if

$$\log\left[(R\beta)^{1/\gamma}\right] < \log G + \tilde{\mathbb{E}}\log\eta.$$

where the expectation $\tilde{\mathbb{E}}$ is taken with respect to the permanent-income-neutral probabilities $\tilde{p}_i = p_i\eta_i$. If the inequality is reversed, then a stable invariant permanent-income-weighted distribution does not exist.

Since $\tilde{\mathbb{E}} \log \eta > \mathbb{E} \log \eta$, the existence of a stable invariant permanent-income-weighted distribution $\tilde{\mu}^{\mathbf{m}}$ is ensured by the existence of a stable invariant marginal distribution $\mu^{\mathbf{m}}$. Loosely speaking, a stable invariant marginal distribution does not exist if too many households accumulate normalized cash on hand in an unbounded fashion. However, the households that do accumulate normalized cash on hand in an unbounded fashion are the ones that receive many consecutive bad permanent income shocks. For the permanent-income-weighted distribution, these households are given less weight and therefore the condition for the existence of a stable invariant permanent-income-weighted distribution is weaker than the condition for the existence of a stable invariant marginal distribution. Both the existence of a stable invariant marginal distribution and the existence of a stable invariant permanent-income-weighted distribution are ensured by the standard impatience condition $\beta R \mathbb{E}\left[(G\eta)^{-\gamma}\right] < 1$.

**Degenerate income distribution, well-defined aggregates** The evolution of the permanent-income-weighted distribution $\tilde{\mu}_t^{\mathbf{m}}$ can be described without reference to the underlying distribution $\mu_t$. In particular, there is a stable invariant permanent-income-weighted distribution $\tilde{\mu}^{\mathbf{m}}$ even though the long-run distribution of permanent income is degenerate. As in Constantinides and Duffie (1996), aggregate behavior is well-defined although the income distribution is not defined. Under the permanent-income-neutral measure, it is therefore not necessary to adopt a perpetual-youth structure in order to be able to compute model aggregates.

Completing the theoretical characterization initiated by Carroll (2020), under the stable invariant permanent-income-weighted distribution $\tilde{\mu}^{\mathbf{m}}$, aggregate consumption and aggregate savings grow at the same rate as permanent income. Normalizing $\mathbb{E}[P_0] = 1$, we have

$$\bar{C}_t = G^t \tilde{\mathbb{E}} c(m) = G^t \bar{C}_0.$$

where the cross-sectional expectation $\tilde{\mathbb{E}}$ is taken with respect to the permanent-income-weighted distribution $\tilde{\mu}^{\mathbf{m}}$.

If we compute aggregate consumption using the objective probability measure instead, we have

$$\bar{C}_t = \mathbb{E}\left[P_t c(m_t)\right] = G^t \mathbb{E}\left[c(m_t)\right] + \mathrm{Cov}\left(P_t, c(m_t)\right).$$

We therefore arrive at the formula

$$\mathrm{Cov}(P_t, c(m_t)) = -G^t \left(\mathbb{E} c(m) - \tilde{\mathbb{E}} c(m)\right).$$

In terms of the distribution of permanent income shocks, the permanent-income-weighted measure

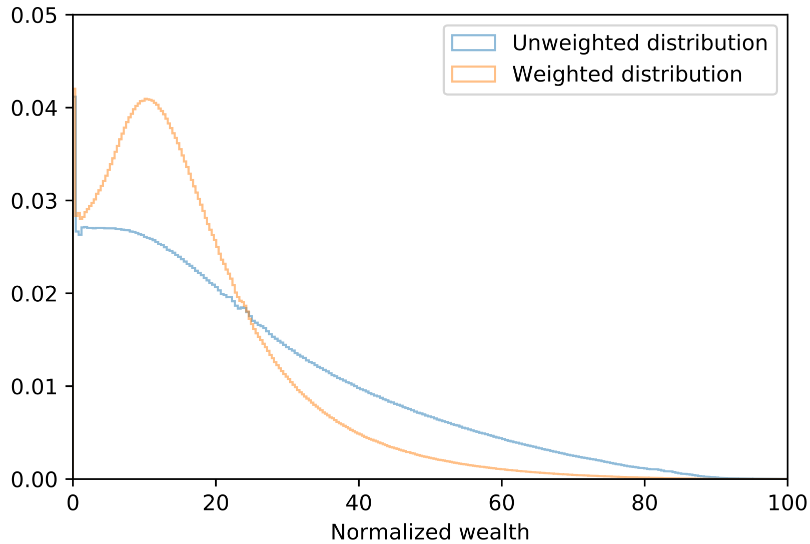


Figure 3.1: The steady-state permanent-income-weighted distribution of normalized wealth plotted together with the unweighted distribution.

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 unweighted distribution of normalized wealth from the numerical exercise in Section 4, the permanent-income-weighted distribution has substantially less normalized wealth. Since the permanent-income-weighted distribution is stochastically dominated, we have $\mathbb{E}\mathfrak{c}(\mathfrak{m}) > \tilde{\mathbb{E}}\mathfrak{c}(\mathfrak{m})$ and the covariance between normalized consumption and permanent income is negative. The argument is analogous for aggregate savings.

Since the permanent-income-weighted distribution is stochastically dominated by the marginal distribution, the aggregate economy behaves as if it holds less assets than the average household. For example, the partial-equilibrium aggregate consumption response to a small proportional change in both the wage and the interest rate is given by $\tilde{\mathbb{E}}\mathfrak{c}'(\mathfrak{m})$ which is larger than the average household response $\mathbb{E}\mathfrak{c}'(\mathfrak{m})$.

# 4 Applying the permanent-income-neutral measure: computations

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. The permanent-income-neutral measure yields a thirtyfold improvement in computational speed/precision.

There is a continuum of households solving the household problem described by Equations 1-3. To maintain a stationary income distribution, we employ a perpetual-youth structure.[3] All households face a probability $\omega$ 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 $K^{\alpha}$.

A steady state equilibrium is given by

i A consumption function $c(\cdot)$ that solves Equations 1-3 given $R$ and $w$,

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

$$\mathbf{P}(\epsilon' = \epsilon_i) = p_i^{\epsilon},$$
$$\mathbf{P}(\eta' = \eta_j) = p_j^{\eta},$$
$$\mathbf{P}(\chi' = 1) = \omega,$$

(where $\chi'$ is a death shock) together with 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,

$$K = \int_{\mathbf{m} \times \mathbf{P}} (m - c(m)) P dm dP,$$
$$R = \frac{\alpha K^{\alpha-1} + (1-\delta)}{1-\omega},$$
$$w = (1-\alpha)K^{\alpha}.$$

[3]Note that, using the permanent-income-neutral measure, it is not necessary to introduce death for computational reasons. I introduce the perpetual-youth structure to be able to compute model aggregates using the objective measure.

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.

### 4.1 Solving the model

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 $\alpha = 0.33$.

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

1. Set $K = 0.5K_{low} + 0.5K_{high}$.
2. Compute $R$ and $w$ given $K$.
3. Solve for $c(\cdot)$ using the endogenous-grid method.
4. Simulate a panel of households and compute the average value of $P(m - c(m))$.
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. To see this, we simulate 1,000 households for $N$ periods where $N = 10^3, 10^4, 10^5, 10^6$ using both the standard law of motion (Equations 4-7) and the permanent-income-neutral measure (Theorem 1). In Figure 4.1, the standard error of total savings is 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 4.1. 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

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

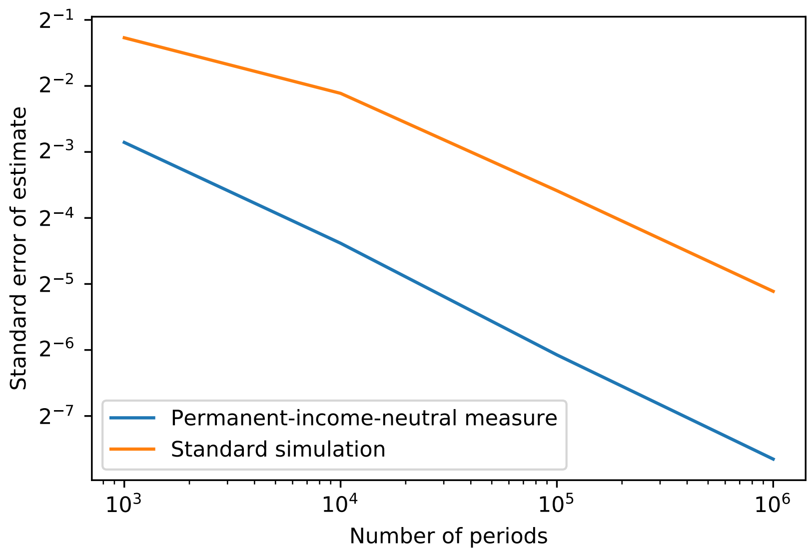


Figure 4.1: The standard error of the estimate of aggregate savings from the simulation of 1,000 households for N periods where $\mathsf{N} = 10^3, 10^4, 10^5, 10^6$, for both the standard simulation method and using the permanent-income-neutral measure.

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 one hour on a MacBook Pro.

**Compatibility with other simulation methods** In this 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 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

| | 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 4.1: 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.

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 $\mathsf{N}$ grows, the size of the transition matrix $\mathsf{N}^2$ grows even faster. The permanent-income-neutral measure helps with non-stochastic simulation since it reduces the dimensionality of the state space.

In Harmenberg and Öberg (2019), we use the permanent-income-neutral measure together with non-stochastic simulation for a model where households can save both in a liquid asset and durable goods, with the durable good subject to a nonconvex adjustment cost. With 200 grid points in the cash-on-hand dimension, 200 grid points in the durable-goods dimension, 2 employment states and a total of 4 aggregate states, we have a state space of size $200 \times 200 \times 2 \times 4 = 320{,}000$ and the transition matrix is of size $320{,}000^2 = 1.024 \times 10^{11}$. If stored naively as a dense matrix, the matrix would be of size 820 gigabytes. In practice, the matrix is sparse so it fits on the RAM memory of, e.g., a MacBook Pro. Computing the transition dynamics for Harmenberg and Öberg (2019) using non-stochastic simulation would simply be unfeasible without the permanent-income-neutral measure. With 30 grid points in the permanent-income dimension, the size of the transition matrix would be 900 times as large and the dense matrix would be 740 terabytes.

## 5 Conclusion

The permanent-income-neutral measure improves the computation of model aggregates such as aggregate consumption and aggregate saving in a broad class of heterogeneous-agent models. The implementation of the method is trivial and yields order-of-magnitude computational improvements. Furthermore, the permanent-income-neutral measure allows us to prove the existence of a stable invariant permanent-income-weighted distribution, and therefore analytically characterize the macro behavior of buffer-stock savings models.

## References


Auclert, A., Rognlie, M., Straub, L., 2020. The Intertemporal Keynesian Cross. Journal of Political Economy Forthcomin.

Bayer, C., Lütticke, R., Pham-Dao, L., Tjaden, V., 2019. Precautionary Savings, Illiquid Assets, and the Aggregate Consequences of Shocks to Household Income Risk. Econometrica 87 (1), 255–290.

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

Carroll, C., Crawley, E., Slacalek, J., Tokuoka, K., White, M., 2019. Sticky Expectations and Consumption Dynamics. Mimeo.

Carroll, C. D., 1992. The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence. Brookings Papers on Economic Activity 1992 (2), 61.

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

Carroll, C. D., 2020. Theoretical Foundations of Buffer Stock Saving. Quantitative Economics.

Carroll, C. D., Slacalek, J., Tokuoka, K., White, M. N., 2017. The distribution of wealth and the marginal propensity to consume. Quantitative Economics 8 (3), 977–1020.

Carvalho, V. M., Grassi, B., 2019. Large Firm Dynamics and the Business Cycle. American Economic Review Forth.

Cocco, J. F., Gomes, F. J., Maenhout, P. J., 2005. Consumption and portfolio choice over the life cycle. Review of Financial Studies 18 (2), 491–533.

Cochrane, J. H., 2005. Asset Pricing, 2nd Edition. Princeton University Press, Princeton, New Jersey.

Constantinides, G. M., Duffie, D., 1996. Asset pricing with heterogeneous consumers. Journal of Political Economy 104 (2), 219–240.

Deaton, A., 1991. Saving and Liquidity Constraints. Econometrica 59 (5), 1221–1248.

Druedahl, J., Jørgensen, T. H., 2017. Persistent vs. Permanent Income Shocks in the Buffer-Stock Model. B.E. Journal of Macroeconomics 17 (1), 1–16.

Gabaix, X., 2011. The Granular Origins of Aggregate Fluctuations. Econometrica 79 (3), 733–772. URL `http://doi.wiley.com/10.3982/ECTA8769`

Gourinchas, P.-O., Parker, J. A., 2002. Consumption over the Life Cycle. Econometrica 70 (1), 47–89.

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

Kaplan, G., Mitman, K., Violante, G. L., 2020. The Housing Boom and Bust: Model Meets Evidence. Journal of Political Economy Forthcomin.

Kaplan, G., Moll, B., Violante, G. L., 2018. Monetary Policy According to HANK. American Economic Review 108 (3), 697–743.

Krueger, D., Mitman, K., Perri, F., 2016. Macroeconomics and Household Heterogeneity. In: Taylor, J., Uhlig, H. (Eds.), Handbook of Macroeconomics, Vol. 2.

Krusell, P., Smith, A. A., 1998. Income and Wealth Heterogeneity in the Macroeconomy. Journal of Political Economy 106 (5), 867–896.

McKay, A., 2017. Time-varying idiosyncratic risk and aggregate consumption dynamics. Journal of Monetary Economics 88, 1–14.

McKay, A., Nakamura, E., Steinsson, J., 2016. The Power of Forward Guidance Revisited. American Economic Review 106 (10), 3133–3158.

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

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 (1), 36–41. URL `http://dx.doi.org/10.1016/j.jedc.2008.11.010`

# A Proof of Theorem 1

In this appendix, I prove the equivalence between the evolution of aggregates for the original model and the evolution of aggregates for the auxilliary model with permanent-income neutral probabilities.

**Environment** For expositional purposes, assume that shocks are discretely distributed (which allows us to easily separate sums belonging to shocks from integrals belonging to the distribution).

- An individual's state is given by cash on hand $m$ and permanent income $P$.
- An individual makes decisions which lead to bond holdings $b = g(am$. We do not need to specify the micro foundations of $g$, as long as the individual's decision only depends on $m$ and not on $P$. We do not assume that $g$ is continuous or invertible, merely that $g$ is measurable.
- The stochastic environment maps $b, P$ to a new state, $(b, P) \mapsto (b/\eta + \epsilon, P\eta)$, where $\eta$ is a shock to permanent income and $\epsilon$ is a transitory shock to income.
- The distribution of individuals over states $(m, P) \in \mathbf{m} \times \mathbf{P}$ is given by the measure $\mu^{\mathbf{m}\times\mathbf{P}}$ (where both $\mathbf{m} \subset \mathbb{R}$ and $\mathbf{P} \subset \mathbb{R}$ are endowed with the Borel $\sigma$-algebra).
- The distribution of households over normalized cash on hand is given by $\mu^{\mathbf{m}}(\mathcal{M}) = \int_{\mathcal{M}\times\mathbf{P}} d\mu_{\mathbf{m}\times\mathbf{P}}$, where $\mathcal{M}$ is a measurable subset of $\mathbf{m}$. This distribution is not the right distribution for macroeconomic aggregates. For macroeconomic aggregates, the contribution of a household needs to be weighted by its permanent income.

- For aggregate variables, $\tilde{\mu}^{\mathbf{m}}(\mathcal{M}) = \int_{\mathcal{M}\times\mathbf{P}} P d\mu^{\mathbf{m}\times\mathbf{P}}$ is a sufficient statistic for the distribution. For example, total consumption is given by $C = \int c(m) d\tilde{\mu}^{\mathbf{m}}$.

**Proof** We now proceed to show how to get a law of motion for $\tilde{\mu}^{\mathbf{m}}$ by using the permanent-income-neutral measure. For a given pair of shocks $\eta', \epsilon'$ we have $(m, P) \mapsto (m', P') = (g(m)/\eta' + \epsilon', \eta' P)$. Write $p_i^{\eta} = \mathbb{P}(\eta' = \eta_i)$ and $p_j^{\epsilon} = \mathbb{P}(\epsilon' = \epsilon_i)$. The measure of households with state $(m', P')$ in the set $\mathcal{M}' \times \mathcal{P}'$ is then given by

$$\mu^{\mathbf{m}'\times\mathbf{P}'}(\mathcal{M}' \times \mathcal{P}') = \sum_i \sum_j p_i^{\eta} p_j^{\epsilon} \mu_{\mathbf{B}\times\mathbf{Z}}((\eta_i(\mathcal{M}' - \epsilon_i)) \times (\mathcal{P}'/\eta_i)) = \sum_i \sum_j p_i^{\eta} p_j^{\epsilon} \mu^{\mathbf{m}\times\mathbf{P}}(g^{-1}(\eta_i(\mathcal{M}' - \epsilon_i)) \times (\mathcal{P}'/\eta_i))$$

and therefore the weighted measure $\tilde{\mu}^{\mathbf{m}'}$ is given by

$$\begin{aligned}
\tilde{\mu}^{\mathbf{m}'}(\mathcal{M}') &= \int_{\mathcal{M}'\times\mathbf{P}'} P' d\mu_{\mathbf{m}'\times\mathbf{P}'} = \sum_i \sum_j p_i^{\eta} p_j^{\epsilon} \int_{(g^{-1}(\eta_i(\mathcal{M}'-\epsilon_i))\times\mathbf{P}} (\eta_i P) d\mu^{\mathbf{m}\times\mathbf{P}} \\
&= \sum_i \sum_j (p_i^{\eta}\eta_i) p_j^{\epsilon} \tilde{\mu}_{\mathbf{m}}(g^{-1}(\eta_i(\mathcal{M}' - \epsilon_i))) = \\
&= \tilde{E}_{\eta}\left[E_{\epsilon}[\tilde{\mu}^{\mathbf{m}}(g^{-1}(\eta_i(\mathcal{M}' - \epsilon_i)))]\right]
\end{aligned}$$

where $\tilde{E}_{\eta}$ is computed using the permanent-income-neutral probabilities $\{\tilde{p}_i\} = \{p_i \eta_i\}$.

Similarly, the unweighted measure is given by

$$\begin{aligned}
\mu^{\mathbf{m}'}(\mathcal{M}') &= \int_{\mathcal{M}'\times\mathbf{P}'} d\mu_{\mathbf{m}'\times\mathbf{P}'} = \sum_i \sum_j p_i^{\eta} p_j^{\epsilon} \int_{(g^{-1}(\eta_i(\mathcal{M}'-\epsilon_i))\times\mathbf{P}} d\mu^{\mathbf{m}\times\mathbf{P}} \\
&= \sum_i \sum_j p_i^{\eta} p_j^{\epsilon} \mu^{\mathbf{m}}(g^{-1}(\eta_i(\mathcal{M}' - \epsilon_i))) = \\
&= E_{\eta}\left[E_{\epsilon}[\mu^{\mathbf{m}}(g^{-1}(\eta_i(\mathcal{M}' - \epsilon_i)))]\right]
\end{aligned}$$

We see that the law of motion for the weighted measure is the same as the law of motion for the unweighted measure, except that the probability distribution for the permanent shock is replaced by the permanent-income-neutral distribution $\{\tilde{p}\} = \{p_i \eta_i\}$.

# B The permanent-income-neutral measure in continuous time

For a textbook treatment of Girsanov's theorem for economists, see Björk (2009).

Consider the following continuous-time model where income follows a geometric Brownian motion,

$$\max \int_0^\infty e^{-\rho t} \frac{C_t^{1-\gamma}}{1-\gamma} \qquad \text{s.t.} \qquad \dot{M}_t = rM_t + Y_t - C_t,$$

$$dY_t = \sigma Y_t dW_t.$$

Like its discrete-time counterpart, optimal behavior in continuous time is characterized by a consumption function $C_t = Y_t c(m_t)$ where $m_t = M_t/Y_t$.

To simulate the model, we need to describe the dynamics of $m$ and $Y$. The dynamics of $Y$ are exogenously given by

$$dY_t = \sigma_Y Y_t dW_t.$$

To retrieve the dynamics of $m$, we recall that $m = M/Y$ and apply Itô's lemma.

$$\begin{aligned} dm_t &= \frac{dM_t}{Y_t} - \frac{M_t}{Y_t^2} dY_t - \frac{dM_t}{Y_t^2} dY_t + \frac{1}{2} \frac{2M_t}{Y_t^3} (dY_t)^2 = \\ &= \frac{rM_t + Y_t - C_t}{Y_t} dt - \frac{M_t}{Y_t^2} (\sigma_Y Y_t dW_t) + \frac{M_t}{Y_t^3} \sigma_Y^2 Y_t^2 dt = \\ &= ((r + \sigma^2) m_t + 1 - c(m_t)) dt - \sigma m_t dW_t. \end{aligned}$$

To summarize, the consumption model is summarized by

$$C_t = c(m_t) Y_t, \tag{10}$$

$$dY_t = \sigma Y_t dW_t^P, \tag{11}$$

$$dm_t = ((r + \sigma^2) m_t + 1 - c(m_t)) dt - \sigma m_t dW_t, \tag{12}$$

where the consumption function $c(\cdot)$ is determined by the household problem.

Aggregate consumption is given by $\mathbb{E}[Y_t c(m_t)]$. Since $Y_t$ is a martingale, we can shift to the equivalent-martingale measure using Girsanov's theorem. Girsanov's theorem states that $\mathbb{E}[Y_t c(m_t)] = \mathbb{E}^Q[c(m_t)]$ where the dynamics under the equivalent-martingale measure $Q$ is given by replacing $dW_t = \sigma_Y dt + dW_t^Q$. The dynamics of the permanent-income-weighted distribution is therefore given by the dynamics,

$$dm_t = (rm_t + 1 - c(m_t)) - \sigma m_t dW_t^Q. \tag{13}$$

With the continuous-time formulation, it is analytically clear what the permanent-income-neutral measure does. It induces a negative drift $-\sigma^2 dt$ which captures that households who received positive permanent income shocks both carry a larger weight in the aggregate and have lower normalized assets. Therefore, it appears as if the aggregate economy faces a negative drift in assets compared with the average household.

The continuous-time dynamics of the permanent-income-weighted distribution $\tilde{\mu}_t^{\mathbf{m}}$ is given by the Kolmogorov forward equation associated with Equation 13,

$$\frac{\partial \tilde{\mu}_t^{\mathbf{m}}(m)}{\partial t} = -\frac{\partial}{\partial m}\left((rm + 1 - c(m))\tilde{\mu}_t^{\mathbf{m}}(m)\right) + \frac{1}{2}\frac{\partial^2}{\partial m^2}\left(\sigma^2 m^2 \tilde{\mu}_t^{\mathbf{m}}(m)\right)$$

which can be contrasted with the Kolmogorov forward equation for the unweighted distribution,

$$\frac{\partial \mu_t^{\mathbf{m}}(m)}{\partial t} = -\frac{\partial}{\partial m}\left(((r + \sigma^2)m + 1 - c(m))\mu_t^{\mathbf{m}}(m)\right) + \frac{1}{2}\frac{\partial^2}{\partial m^2}\left(\sigma^2 m^2 \mu_t^{\mathbf{m}}(m)\right).$$

With time-varying interest rate and shock variance,, we replace $r$ with $r_t - g_t$ and $\sigma$ with $\sigma_t$. If permanent income is a jump process (with arrival rate $\lambda_j$ associated to shock $\eta_j$), then under the permanent-income-neutral measure the arrival rates are $\lambda_j \eta_j$.