Aggregating Heterogeneous-Agent Models with Permanent Income Shocks*

Karl Harmenberg†

June 15, 2021

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.

*I am grateful for helpful comments from editor Thomas Lubik, the associate editor, the referee, Jeppe Druedahl, Kasper Kragh-Sørensen, and Erik Öberg. I am also grateful to Tomas Björk for teaching me the mathematics of equivalent probability measures.

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

1 Introduction

The heterogeneous-agent macroeconomic paradigm emphasizes the importance of rich heterogeneity at the micro level for macroeconomic aggregates. One of the main challenges for the paradigm is that macroeconomic models with rich descriptions of the micro environment are computationally challenging to solve and it is therefore important to develop computational methods that reduce these challenges. In this paper, I provide a simple aggregation method which reduces the dimensionality of the state space for the class of macroeconomic models featuring permanent income shocks. The method, which only takes a few lines of code to implement, improves the computational performance by several orders of magnitude.

The class of models under consideration, following Zeldes (1989), Deaton (1991), and Carroll (1997), models income as subject to fully permanent income shocks and transitory income shocks.[1] An advantage of such an income process with fully permanent income shocks is that, combined with CRRA preferences, it 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 $C_t = P_t c(m_t)$. The unit-root process for log permanent income also closely approximates the benchmark income process for heterogeneous-agent models, a highly persistent $AR(1)$.[2]

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, 2021). However, the computational tractability and theoretical clarity seemed to be 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 was thought to be 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 was 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

[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. (2020) and Harmenberg and Öberg (2021).

[2] Although there is a growing literature emphasizing the importance of accurately capturing the full richness of income dynamics, see, e.g., Browning et al. (2010), De Nardi et al. (2020) and Guvenen et al. (2021), an income process with near-permanent income shocks remains the baseline in the heterogeneous-agent macroeconomic literature. For example, in their handbook chapter on macroeconomics and household heterogeneity, Krueger, Mitman and Perri (2016) model income as subject to persistent shocks with quarterly autocorrelation 0.99. In the recent HANK literature, McKay et al. (2016), Bayer et al. (2019), Auclert et al. (2018), and Kaplan et al. (2020) all model household income as an $AR(1)$ in logs with high persistence and the income process of Kaplan et al. (2018) features a rare income shock with a half-life of 18 years.

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 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 for buffer-stock savings model, I characterize when a stable invariant permanent-income-weighted distribution exists. I then prove the conjecture from Carroll (2021) that, in the long run, aggregate consumption grows at the same rate as aggregate income in buffer-stock savings models.

The permanent-income-neutral measure also yields computational improvements. I use the permanent-income-neutral measure to compute aggregate savings in an Aiyagari model both through Monte Carlo simulation and non-stochastic simulation. In both cases, the computation of aggregate savings is faster using the permanent-income-neutral measure. With non-stochastic simulation, the permanent-income-neutral measure yields a thousandfold improvement in computation time and the Aiyagari model can be solved in less than a third of a second. 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, comments on how to use the permanent-income-neutral in continuous time, 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, 2021) (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 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

$$\begin{aligned} C_t &= \int c(m_t)P_t\psi_t(m_t, P_t)dm_t dP_t \\ &= \int c(m_t)\left(\int P_t\psi_t(m_t, P_t)dP_t\right) dm_t. \end{aligned}$$

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

Definition 2. *The* permanent-income-weighted distribution *is defined as* $\tilde{\psi}_t^m(m_t) := G^{-t} \int P_t \psi_t(m_t, P_t) dP_t$.

The inclusion of the detrending factor G^{-t} in the definition of the permanent-income-weighted distribution is not essential but yields a cleaner statement of the main theorem. The detrending is analogous to, e.g., the detrending of the Solow model by productivity and population growth.

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

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

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

Remark 1. *With CRRA utility, household period utility is given by $C^{1-\gamma}/(1-\gamma)$. By setting $\rho = 1-\gamma$, the proof of Theorem 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? Consider a setup where all households have the same permanent income 1.0 at time $t = 0$. Of these households, half receive an increase in their permanent income by 50 percent and half receive a fall in their permanent income by 50 percent. In period $t = 1$, half of the households therefore have permanent income 1.5 and half have permanent income 0.5. Although only 50 percent of households received the positive shock to permanent income, 1.5×50 percent $= 75$ percent of permanent income resides with these households. For the purposes of determining aggregate consumption, tracking permanent income rather than the households is sufficient. The permanent-income-neutral shock distribution combines the objective probability of the shock ($p = 0.5$) with the ex-post weight assigned to the households that receive the shock ($\eta = 1.5$), describing the law of motion for "units of permanent income" (a share $\tilde{p} = p \times \eta = 0.75$ of permanent income resides with the households that received the positive shock). Therefore, if we want to keep track of the distribution of permanent income, rather than households, along the normalized-state dimension, we use the permanent-income-neutral measure.

Theorem 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 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 (2021).

Consider the following buffer-stock savings model from Carroll (1997). There is a continuum of infinitely-lived households with stochastic labor income who can consume and save in a risk-free bond subject to a no-borrowing constraint. The households solve the problem

$$\max \mathbb{E}_0 \sum_{t=0}^{\infty} \beta^t \frac{C_t^{1-\gamma}}{1-\gamma} \qquad \text{s.t.} \qquad \begin{aligned} C_t + B_{t+1} &= M_t, \\ M_{t+1} &= RB_t + Y_{t+1}, \\ Y_{t+1} &= wP_{t+1}\epsilon_{t+1}, \\ P_{t+1} &= P_t G \eta_{t+1}, \\ B_t &\geqslant 0, \end{aligned} \tag{5}$$

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. Labor income is also subject to a transitory income shock ϵ_t. Both $\eta_t \sim f_\eta$ and $\epsilon_t \sim f_\epsilon$ are non-negative, independent and i.i.d. with mean 1.

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. Reformulating the problem in its recursive formulation, the households are solving the following problem:[3]

$$v(m) = \max_{b,c} \frac{c^{1-\gamma}}{1-\gamma} + \beta \mathbb{E}\left[(G\eta')^{1-\gamma} v(m')\right] \qquad \text{s.t.} \qquad \begin{aligned} b + c &= m, \\ m' &= \frac{Rb}{G\eta'} + w\epsilon', \\ b &\geqslant 0, \end{aligned} \tag{6}$$

where primed variables (m', η', ϵ') denote next-period variables.

Szeidl (2013) provides a characterization of when a stable invariant distribution of normalized cash on

[3] Write the value function in terms of normalized cash on hand $m = M/P$ and permanent income P. It is then straightforward to guess and verify that the value function is on the form $V(m, P) = v(m)P^{1-\gamma}$, allowing us to eliminate P as a state variable in the household problem.

hand m_t, ψ^m, exists for this problem. The characterization is in terms of the asymptotic marginal propensity to consume as the household's assets tend to infinity, which we denote by mpc^*. Szeidl (2013) shows that this asymptotic marginal propensity to consume can be solved for analytically by considering an auxilliary model without labor income (i.e., a "cake eating problem"), in our setting $mpc^* = 1 - \frac{(\beta R)^{1/\gamma}}{R}$. With these preliminaries, Szeidl (2013) proves the following result:

Proposition 1 (from Szeidl (2013))**.** *There exists a stable invariant marginal distribution* ψ^m *if*

$$\log[R(1 - mpc^*)] < \mathbb{E}\log[G\eta']. \tag{7}$$

Furthermore, if

$$\log[R(1 - mpc^*)] > \mathbb{E}\log[G\eta'] \tag{8}$$

then there does not exist an invariant marginal distribution.

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. Using Theorem 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 2 (adapted from Szeidl (2013))**.** *There exists a stable invariant permanent-income weighted distribution* $\tilde{\psi}^m$ *if*

$$\log[R(1 - mpc^*)] < \tilde{\mathbb{E}}\log[G\eta'] \tag{9}$$

where the expectation $\tilde{\mathbb{E}}$ *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*

$$\log[R(1 - mpc^*)] > \tilde{\mathbb{E}}\log[G\eta_{t+1}] \tag{10}$$

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

Proof. By Theorem 1, the law of motion for ψ_t^m is the same as the law of motion for $\tilde{\psi}_t^m$, 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 $\mathbb{E}\log[G\eta'] < \tilde{\mathbb{E}}\log[G\eta']$ 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 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 2, we can now prove a conjecture from Carroll (2021).

Proposition 3 (conjecture from Carroll (2021))**.** *Under the conditions of Proposition 2, aggregate consumption grows at the same rate as permanent income in the long run.*

This conjecture may strike the reader as obvious. However, natural as it may look, it evaded being proven because the right way to approach the conjecture, with the permanent-income-neutral measure, was not available.

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

$$C_{t+1} = G^{t+1}\int c(m)\tilde{\psi}^m(m)dm = G\left(G^t\int c(m)\tilde{\psi}^m(m)dm\right) = GC_t \tag{11}$$

so $\frac{C_{t+1}}{C_t} = G$. □

Following Carroll (2021), denote by $\mathbb{M}_t[\cdot]$ the cross-sectional average operator at time t, that is, the expected value from drawing a household at random using the density $\psi_t(m_t, P_t)$. We can compute aggregate consumption in two ways. First, we can use the objective measure,

$$C_t = \int c(m_t)P_t\psi_t(m_t, P_t)dm_t dP_t = \mathbb{M}_t[c(m_t)P_t] = G^t\mathbb{M}_t[c(m_t)] + \mathrm{Cov}_t(c(m_t), P_t), \tag{12}$$

where the covariance Cov_t is also taken with respect to the cross-sectional distribution of households $\psi_t(m_t, P_t)$. The last equality in Equation (12) corresponds to the remark in Carroll (2021) where he writes that "[a] proof that the covariance shrinks fast enough would mean that the term could be neglected" in his discussion of a proof strategy for the conjecture proven in Proposition 3.

We can also compute aggregate consumption using the permanent-income-weighted distribution,

$$C_t = G^t\int c(m_t)\tilde{\psi}_t^m(m_t)dm_t = G^t\tilde{\mathbb{M}}[c(m_t)] \tag{13}$$

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

$$\mathrm{Cov}_t(c(m_t), P_t/G^t) = -\left(\mathbb{M}_t[c(m_t)] - \tilde{\mathbb{M}}_t[c(m_t)]\right). \tag{14}$$

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 particular, the covariance between normalized consumption and permanent income does not shrink asymptotically, and the proof strategy suggested by Carroll (2021) is bound to fail. Instead, a shift to the permanent-income-weighted measure provides an easy way to prove the conjecture.

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 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.

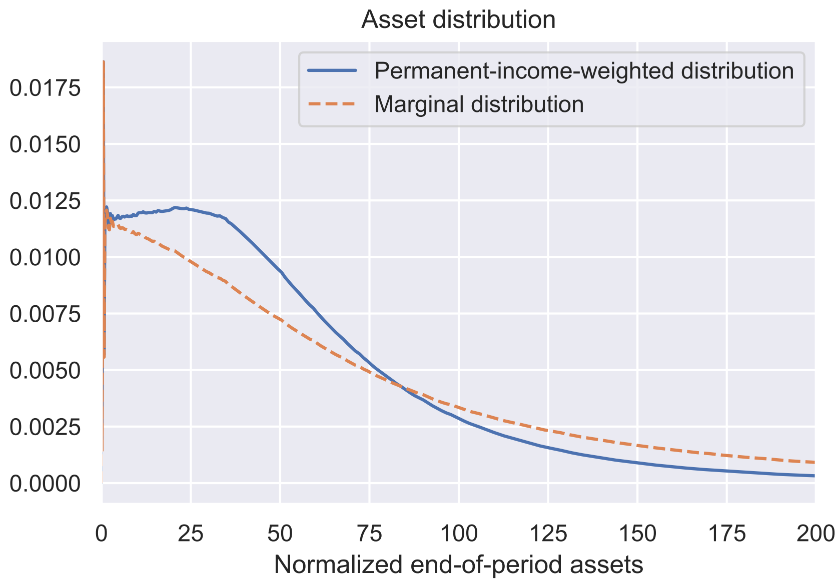

Figure 1: The steady-state permanent-income-weighted distribution of normalized wealth plotted together with the marginal distribution. The densities are jagged near the borrowing constraint because the shock distributions are discrete valued.

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, 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.[4] 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 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).

[4]The permanent-income distribution features a fat tail in McKay (2017), Carroll et al. (2017) and Carroll et al. (2020).

In the literature, e.g., McKay (2017), Carroll et al. (2017) and Carroll et al. (2020), 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.[5] I nonetheless introduce the perpetual-youth structure in order to compare simulating the model with and without the permanent-income-neutral measure.

Model environment We are looking for a stationary equilibrium to a stationary environment where households are facing the household problem described by Equation (5), or equivalently Equation (6). Since the environment is stationary, the expected growth rate of permanent income is set to zero, i.e., $G = 1$.

- *Household block/asset demand:* There is a continuum of perpetual-youth households solving the household problem described by Equation (5), or equivalently Equation (6). For this problem, household optimal behavior is summarized by a consumption function $c(\cdot)$ such that a household with normalized cash on hand m and permanent income P consumes $c(m)P$. 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 $P = 1$.[6]

- *Production block/asset supply:* Factor prices are determined by a Cobb-Douglas aggregate production function $Y = K^{\alpha}$. The interest rate also incorporates that the assets of the deceased households are distributed among the living households, e.g., through actuarially fair annuity markets.

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

- *Equilibrium/resource constraint:* Aggregate capital is equal to aggregate savings. For an individual household, savings are equal to cash on hand minus consumption, $b = m - c(m)$, so in aggregate we have

$$K = \int (m - c(m))P\psi(m, P)dmdP$$

  where ψ is the stationary distribution of households implied by the solution to the household problem.

Structure of the equilibrium An equilibrium is a tuple $(K, w, R, c(\cdot), \psi)$ such that:

[5] 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.

[6] Strictly speaking, the household problem described by Equation (6) was formulated for infinite-horizon households. However, if we interpret the discount rate β in Equation (6) as the combination of pure discounting $\widehat{\beta}$ and the mortality risk ω, $\beta = \widehat{\beta}(1-\omega)$, then Equation (6) describes the household problem for a perpetual-youth household.

1. The factor prices are $R = \frac{\alpha K^{\alpha-1}+(1-\delta)}{1-\omega}$ and $w = (1-\alpha)K^{\alpha}$.

2. The consumption function $c(\cdot)$ solves Equation (6), given R and w.

3. The stationary distribution ψ is given by the law of motion implied by the consumption function $c(\cdot)$ and the factor prices R and w.

4. The capital stock K equals aggregate savings $K = \int (m - c(m)) P \psi(m, P) dm dP$.

Solution algorithm Any solution algorithm for this problem involves computing a stationary distribution. The permanent-income-neutral measure and the results in this paper allow us to only compute the one-dimensional permanent-income-weighted distribution rather than the two-dimensional joint distribution of cash on hand and permanent income.

I implement a solution algorithm in Python (using NumPy, SciPy and Numba) on a MacBook Pro 2019 with details in Appendix A. The full code is available as an Online Appendix.[7] The algorithm consists of two layers. The inner layer is a function that returns the implied (partial equilibrium) aggregate savings, given a level of capital. The high-level outline of the code for this function is as follows:

1. Take K, the level of capital, as an input.

2. Compute factor prices $R = \frac{\alpha K^{\alpha-1}+(1-\delta)}{1-\omega}$ and $w = (1-\alpha)K^{\alpha}$ given capital K.

3. Compute the optimal consumption function $c(\cdot)$ given factor prices R and w using a combination of the endogenous-grid method (Carroll, 2006) and Howard's improvement algorithm.

4. Compute the stationary distribution ψ from the law of motion implied by the consumption function $c(\cdot)$ and the factor prices R and w, using non-stochastic simulation (Young, 2010).

5. Compute aggregate savings $\int (m - c(m)) P \psi(m, P) dm dP$ given the consumption function $c(\cdot)$ and the stationary distribution ψ.

6. Return aggregate savings.

In the outer layer, the algorithm solves for the level of capital for which implied savings equal the capital stock. I find the level of capital using Broyden's method.

Parameter values In what follows, I use parameter values from Carroll et al. (2017). The household discount factor is set to $\beta = 0.99$ and the risk aversion to $\gamma = 1.0$. The transitory shock is log-normally distributed with $\sigma_\epsilon = \sqrt{0.01 \times 4}$ and the permanent shock is log-normally distributed with $\sigma_\eta = \sqrt{0.01 \times 4/11}$. The mortality risk is set to $\omega = 0.00625$ yielding and average work life of 40 years. The capital share is set to $\alpha = 0.36$ and the depreciation rate to 0.025.

[7]The code is also available through the author's homepage, including an interactive Python notebook.

Discretization for computations Both the permanent and transitory income shocks are discretized using Gauss-Hermite quadrature with five nodes. The grid for the normalized cash-on-hand dimension has 300 grid points with more grid points for low levels of cash on hand.

Computing the optimal consumption function Optimal consumption behavior is solved for using the endogenous-grid method together with Howard's improvement algorithm, with an error tolerance of 10^{-10} using the L^1 norm between the value functions of two subsequent iterations of the combination of the endogenous-grid method and Howard's improvement algorithm. Details for the computation of optimal consumption is provided in Appendix A.

4.1 Computing the stationary distribution

The law of motion without the permanent-income-neutral measure The law of motion for the joint distribution ψ_t is implied by the household-level stochastic law of motion

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

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

where the death shock χ' is equal to 0 with probability $1-\omega$ and equal to 1 with probability ω, the transitory shock is drawn using its distribution $\epsilon' \sim f_\epsilon$ and the permanent income shock is drawn using its distribution $\eta' \sim f_\eta$.

The law of motion with the permanent-income-neutral measure The law of motion for the permanent-income-weighted distribution $\tilde{\psi}_t$ is implied by the household-level stochastic law of motion

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

where the death shock χ' is equal to 0 with probability $1-\omega$ and equal to 1 with probability ω, the transitory shock is drawn using its distribution $\epsilon' \sim f_\epsilon$ while the permanent income shock is drawn using the permanent-income-neutral measure $\eta' \sim \tilde{f}_\eta$.

For completeness, the law of motion (or lack thereof) for permanent income under the permanent-income-neutral measure is

$$P' = 1.$$

Computations The simplest strategy for computing either distribution, Monte Carlo simulation, is to simulate long time series of (m_t, P_t) using the above stochastic laws of motion and then view the long time series as representative of the stationary distribution. A better method, following Young (2010), is to discretize the state space and describe the transition probabilities between different grid points with a matrix M. The stationary distribution is then the unique eigenvector, normalized to sum to 1, associated with the eigenvalue 1.

I explore the computational benefits of using the permanent-income-neutral measure for both Monte Carlo simulation and non-stochastic simulation below. The factor prices are set to $R = 1.00965$ and $w = 2.67369$ (these are, as will later be shown, the equilibrium values for the factor prices). I compute implied aggregate savings by both Monte Carlo simulation and non-stochastic simulation.

4.1.1 Performance of the permanent-income-neutral measure under Monte Carlo simulation

With and without the permanent-income-neutral measure, I simulate histories of length 1 000 000 periods and compute the implied aggregate savings, viewing the 1 000 000 periods as draws from the steady state distribution. I repeat this exercise 100 times, reporting the mean aggregate capital and the standard deviation of aggregate savings divided by $\sqrt{100}$ (i.e., the standard error). The implementation of the permanent-income-neutral measure only requires changing a few lines of code in the simulation code. In the Appendix, I display the Python code for stochastic simulation (see Figure A.4). When the permanent-income-neutral measure is used, the permanent income shocks are drawn with the permanent-income-neutral probabilities.

Aggregate savings are computed as the average asset holdings over time, $\frac{1}{T}\sum_{t=1}^{T} b_t P_t$. The results of the simulation exercise are displayed in Table 1, the simulation of 100 draws of the 1 000 000 periods takes approximately 90 seconds. The simulated aggregate savings are not statistically different from each other but the standard error of the two estimates are different. The standard error without the permanent-income-neutral measure is 2.40 times larger than the standard error when using the permanent-income-neutral measure. Since the standard error falls of by $1/\sqrt{N}$, this means that the simulation with the permanent-income-neutral measure only requires $1/2.40 = 17.4\%$ as many draws for a given level of precision.

Why is precision increased? One of the nuisances with simulating the model without the permanent-income-neutral measure is that the permanent-income distribution features a fat tail. Some households have much greater permanent income than others and are therefore disproportionately important for aggregates. However, much simulation time is spent simulating the households with low permanent income that do not contribute much to aggregates. The permanent-income-neutral measure in effect oversamples the permanent-

	Monte Carlo w/o P-I-N measure	Monte Carlo w/ P-I-N measure	Ratio
Mean aggregate savings	52.94	52.85	
SE of aggregate savings	(0.15)	(0.06)	2.40

Table 1: The mean and standard error of aggregate savings from simulating 1 000 000 periods for 100 households. Without the permanent-income-neutral measure, the standard error of aggregate savings is 2.40 times as large. For a given precision, simulations with the permanent-income-neutral measure therefore only need $1/2.40^2 = 17.4\%$ as many Monte Carlo draws as simulations without the permanent-income-neutral measure.

income-rich households and thereby gains substantially better precision.[8]

4.1.2 Performance of the permanent-income-neutral measure under non-stochastic simulation

Next, I implement non-stochastic simulation of the model, following Young (2010). There are $\mathsf{N}^{\mathsf{m}} = 300$ grid points in the cash-on-hand dimension. For the permanent-income dimension, I use $\mathsf{N}^{\mathsf{P}} = 31$ grid points. The discretized state space is thus of size $\mathsf{N}^{\mathsf{m}} \times \mathsf{N}^{\mathsf{P}} = 300 \times 31 = 9300$.

Non-stochastic simulation for computing the steady state savings amounts to writing down the implied transition matrix M for the law of motion and then finding the eigenvector of M associated with eigenvalue 1. The transition matrix M has $(\mathsf{N}^{\mathsf{m}} \times \mathsf{N}^{\mathsf{P}})^2 = 9300^2 = 8.649 \times 10^7$ entries, and although the matrix is sparse and can be saved as a sparse matrix, it is a serious computational enterprise to find its eigenvector associated with eigenvalue 1 (I use Scipy's own sparse matrix eigenvalue routine scipy.sparse.linalg.eigs which calls the ARPACK routine written in Fortran 77).

By using the permanent-income-neutral measure, I only need to track the cash-on-hand dimension and the state space is thus of size $\mathsf{N}^{\mathsf{m}} = 300$. Implementing non-stochastic simulation with a one-dimensional state space is marginally easier than implementing non-stochastic simulation with a two-dimensional state space. It is therefore somewhat easier, in terms of programming, to implement non-stochastic simulation using the permanent-income-neutral measure than without. Further, non-stochastic simulation of the permanent-income-weighted distribution only requires a minimal modification of (potentially pre-existing) code for non-stochastic simulation of the marginal distribution for m, as shown in Figure 2. Finally, because of the lower dimensionality of the state space, the transition matrix $\tilde{\mathsf{M}}$ under the permanent-income-neutral measure only has $(\mathsf{N}^{\mathsf{m}})^2 = 300^2 = 90\,000$ entries.

Table 2 shows the computation time for non-stochastic simulation, with and without the permanent-income-neutral measure. Without the permanent-income-neutral measure (with 31 grid points), computing

[8] This is in close analogy with *importance sampling* used for Monte Carlo integration in Bayesian econometrics, see, e.g., Kloek and van Dijk (1978).

```python
if weighting_scheme == 'None':
        weight = prob
if weighting_scheme == 'Aggregate':
        weight = params.eta_val[eta_i]*prob
```

Figure 2: For non-stochastic simulation using the permanent-income-neutral measure, it is sufficient to change one line of code in pre-existing code for creating the transition matrix of non-stochastic simulation of the marginal distribution of cash on hand.

	Non-stochastic w/o P-I-N measure (31 grid points)	Non-stochastic w/o P-I-N measure (101 gridpoints)	Non-stochastic w/ P-I-N measure
Aggregate savings	50.86	53.11	53.12
Computation time	1.17s	12.58s	0.01s

Table 2: Aggregate savings using non-stochastic simulation with 31 grid points in the permanent-income dimension, 101 grid points in the permanent-income dimension, and with the permanent-income-neutral measure. The permanent-income-neutral measure yields a thousandfold speedup compared to 101 grid points and a hundredfold speedup compared to 31 grid points. While 101 grid points and the permanent-income-neutral measure yield very close values for aggregate savings, the discrepancy from 31 grid points is substantial.

aggregate savings takes approximately one second, while it only takes 0.01 seconds using the permanent-income-neutral measure. However, there is a noticeable discrepancy in aggregate savings between the two methods. Without the permanent-income-neutral measure, aggregate savings are computed to be 50.86 while aggregate savings are 53.12 when using the permanent-income-neutral measure. This discrepancy is due to 31 grid points providing an insufficient approximation of the permanent-income dimension. When we increase the number of grid points in the permanent-income dimension to 101, aggregate savings without the permanent-income-neutral measure are very close to aggregate savings with the permanent-income-neutral measure. The increase in grid points does however lead to an even larger state space, of size 30 000, yielding a transition matrix with 9×10^8 entries. The increase in the size of the state space leads to an increase in computational time to over ten seconds, three orders of magnitude slower than the computation time when using the permanent-income-neutral measure.

4.2 Computing the equilibrium with the permanent-income-neutral measure

In the previous section, we studied the performance of the permanent-income-neutral measure for computing a stationary distribution and the implied aggregate savings. In this subsection, I report the results from using Broyden's method together with non-stochastic simulation and the permanent-income-neutral measure to solve for the equilibrium level of capital.

The algorithm finds the equilibrium level of capital, $\mathsf{K} = 53.12$, in 0.25 seconds, with a difference between aggregate savings and the capital stock less than 10^{-12}. This involves 8 outer iterations with Broyden's method, for which less than 0.08 seconds are spent in total on aggregating the model. In other words, using the permanent-income-neutral measure, we solve for the equilibrium level of capital much faster than we compute aggregate savings for one single iteration of the outer algorithm if we do not use the permanent-income-neutral measure. The equilibrium permanent-income-weighted distribution and the (unweighted) marginal distribution of normalized assets are both shown in Figure 1.

The model considered in this paper is purposefully kept simple but the method is general. For a given level of computational complexity, eliminating permanent income as a state variable allows researchers to introduce an additional state variable. In Harmenberg and Öberg (2021), we use the permanent-income-neutral measure together with non-stochastic simulation to solve for aggregate dynamics of a model which has a three-dimensional state space (savings, durable goods, permanent income) and non-convex adjustment costs. By using the permanent-income-neutral measure, the state space is only two dimensional and computing aggregate dynamics is easy.

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,

$$\underbrace{S}_{\text{Price}} = E[\Lambda d] = \underbrace{R^{-1}E[d]}_{\text{Discounted expected return}} + \underbrace{\text{cov}\,(\Lambda, d)}_{\text{Risk premium}} \tag{15}$$

where $R = 1/E[\Lambda]$. Notice the structural similarity with Equation (12). 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. (2021), 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 (2019)). Let permanent income follow a geometric Brownian motion, $dP_t = gP_t dt + \sigma P_t dW_t$. Aggregate consumption is given by $\mathbb{E}[P_t c(m_t)]$. Girsanov's theorem states that $\mathbb{E}[P_t c(m_t)] = e^{gt}\mathbb{E}^Q[c(m_t)]$ 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., 2021. Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach. Review of Economic Studies doi:10.1093/restud/rdab002.

Aiyagari, S.R., 1994. Uninsured Idiosyncratic Risk and Aggregate Saving. The Quarterly Journal of Economics 109, 659–684. doi:10.2307/2118417.

Auclert, A., Rognlie, M., Straub, L., 2018. The Intertemporal Keynesian Cross.

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, 255–290. doi:10.3982/ECTA13601.

Björk, T., 2019. Arbitrage Theory in Continuous Time. 4th ed., Oxford University Press. doi:10.1093/oso/9780198851615.001.0001.

Blanchard, O.J., 1985. Debt, Deficits, and Finite Horizons. Journal of Political Economy 93, 223–247. doi:10.1086/261297.

Browning, M., Ejrnæs, M., Alvarez, J., 2010. Modelling Income Processes with Lots of Heterogeneity. Review of Economic Studies 77, 1353–1381. doi:10.1111/j.1467-937X.2010.00612.x.

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., 2020. Sticky Expectations and Consumption Dynamics. American Economic Journal: Macroeconomics 12, 40–76. doi:10.1257/mac.20180286.

Carroll, C.D., 1997. Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis. Quarterly Journal of Economics 112, 1–55. doi:10.1162/003355397555109.

Carroll, C.D., 2006. The method of endogenous gridpoints for solving dynamic stochastic optimization problems. Economics Letters 91, 312–320. doi:10.1016/j.econlet.2005.09.013.

Carroll, C.D., 2021. Theoretical Foundations of Buffer Stock Saving. Quantitative Economics URL: https://qeconomics.org/ojs/forth/354/354-1.pdf.

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.

De Nardi, M., Fella, G., Paz-Pardo, G., 2020. Nonlinear household earnings dynamics, self-insurance, and welfare. Journal of the European Economic Association 18, 890–926. doi:10.1093/jeea/jvz010.

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. doi:10.1111/1468-0262.00269.

Guvenen, F., Karahan, F., Ozkan, S., Song, J., 2021. What Do Data on Millions of U.S. Workers Reveal About Life-Cycle Earnings Risk?

Haliassos, M., Michaelides, A., 2003. Portfolio Choice and Liquidity Constraints. International Economic Review 44, 143–177. doi:10.1111/1468-2354.t01-1-00065.

Harmenberg, K., Öberg, E., 2021. Consumption dynamics under time-varying unemployment risk. Journal of Monetary Economics 118, 350–365. doi:10.1016/j.jmoneco.2020.10.004.

Kaplan, G., Mitman, K., Violante, G.L., 2020. The Housing Boom and Bust: Model Meets Evidence. Journal of Political Economy 128, 3285–3345. doi:10.1086/708816.

Kaplan, G., Moll, B., Violante, G.L., 2018. Monetary Policy According to HANK. American Economic Review 108, 697–743. doi:10.3386/w21897.

Kloek, T., van Dijk, H.K., 1978. Bayesian Estimates of Equation System Parameters: An Application of Integration by Monte Carlo. Econometrica 46, 1–19.

Krueger, D., Mitman, K., Perri, F., 2016. Macroeconomics and Household Heterogeneity, in: Taylor, J., Uhlig, H. (Eds.), Handbook of Macroeconomics, Vol. 2, pp. 843–921. doi:10.1016/bs.hesmac.2016.04.003.

Krusell, P., Smith, A.A., 1998. Income and Wealth Heterogeneity in the Macroeconomy. Journal of Political Economy 106, 867–896. doi:10.1086/250034.

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.

McKay, A., Nakamura, E., Steinsson, J., 2016. The Power of Forward Guidance Revisited. American Economic Review 106, 3133–3158. doi:10.1257/aer.20150063.

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, 36–41. 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: Details of computations

The full source code, including an interactive Python notebook, is available as an Online Appendix and through the author's homepage.

Grids The grid in normalized cash on hand m consists of 300 grid points from 0.1 to 400 spaced quadratically (the square of an array with equispaced grid points from $\sqrt{0.1}$ to $\sqrt{400}$). The grid in normalized end-of-period assets b consists of 300 grid points with 399 of the grid points from 0.1 to 400 spaced quadratically (the square of an array with equispaced grid points from $\sqrt{0.1}$ to $\sqrt{400}$), and an additional grid point at 0.0.

Computing the optimal consumption function The computation of the optimal consumption function uses a combination of the endogenous-grid method and Howard's improvement algorithm. The code, displayed in Figure A.1, alternates between the following two steps until convergence:

- Given the current guess for the value function, compute optimal consumption by the endogenous-grid method.
- Given the current guess for optimal consumption, use Howard's improvement algorithm to compute the implied value function from the consumption function. Use the implied value function as the new guess for the value function.

The application of the endogenous-grid method is completely standard. The application of Howard's improvment algorithm uses the transition matrix T, implied by the shocks and the consumption function, and matrix algebra. Letting u denote the vector of per-period utilities, the vector v of values is given by $v = u + u\beta T + u\beta^2 T^2 + \ldots = u(I - \beta T)^{-1}$. The code for the implementation of Howard's improvement algorithm is shown in Figure A.2.

The method only needs roughly ten iterations before convergence. The equilibrium consumption function is shown in Figure A.3.

Discretization of the permanent-income dimension Since the permanent-income dimension features a fat tail, the grid in permanent-income consists of 31 or 101 grid ponts exponentially distributed from $\exp(-10)$ to $\exp(10)$ (or equivalently, the grid for log permanent income is equispaced from -10 to 10). Since the stationary distribution for permanent income is double Pareto, the implied distributions for the permanent-income dimension captures the peak of the distribution better with an odd number of grid points.

Stochastic simulation The stochastic simulation generates a permanent-income-weighted distribution in accordance with the distribution obtained from non-stochastic simulation. The code for stochastically

```python
def compute_optimal_consumption_function(initial_v, R, wage, params):
   v = initial_v

   error = 1.0
   iterations = 0

   while error > 1e-10:
      consumption_function = egm_iteration(v, R, wage, params)
      transition_matrix_howard = \
         create_transition_matrix(consumption_function, 'Howard',
                                 params, R, wage)
      v_new = howard_improvement_algorithm(consumption_function,\
                                  transition_matrix_howard, params)

      error = np.sum(np.abs(v-v_new))
      v = v_new
      iterations += 1

   print("Iterations needed= ", iterations)

   return consumption_function, v
```

Figure A.1: The code for computing optimal consumption.

```python
def howard_improvement_algorithm(consumption_function, transition_matrix, params):
   period_utility = u(consumption_function, params.gamma)
   v = spsolve(eye(params.Nm)-params.beta*transition_matrix, period_utility)

   return v
```

Figure A.2: The implementation of Howard's improvement algorithm.

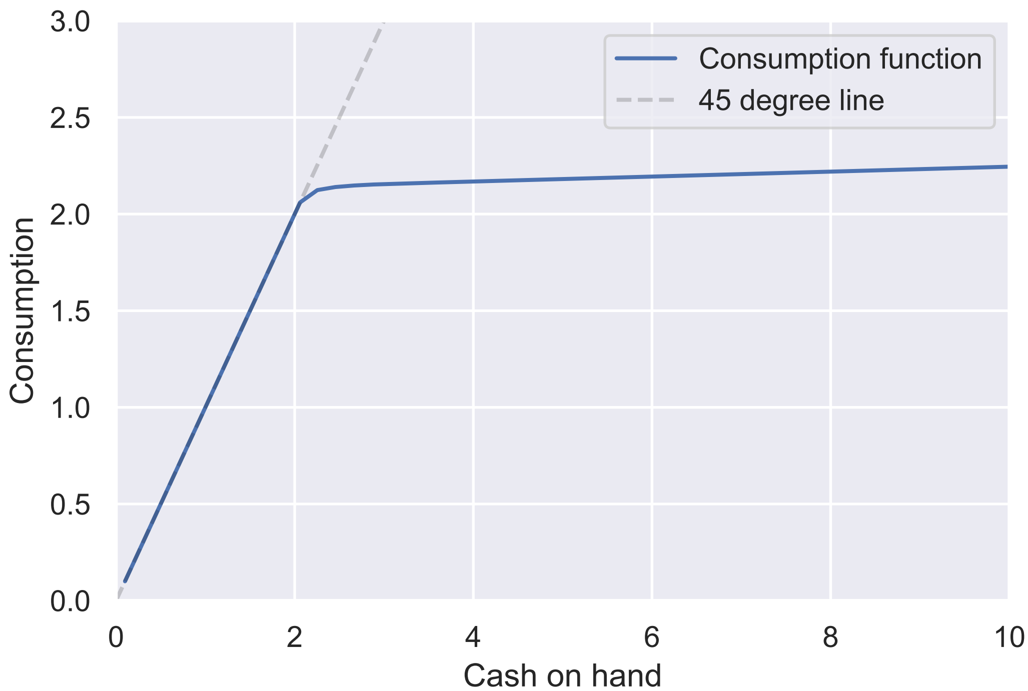

Figure A.3: The equilibrium consumption function.

simulating one period for one agent is shown in Figure A.4.

```python
def simulate_agent_one_period(m, P, consumption_function,
                              params, R, wage,
                              distorted_probabilities):

   #Transitory shock
   epsilon_shock = random_choice(params.epsilon_val, params.epsilon_prob)

   if distorted_probabilities == False:
      #Permanent-income shock if not using the
      #permanent-income-neutral measure
      eta_shock = random_choice(params.eta_val,
                                params.eta_prob)
   else:
      #If using the permanent-income-neutral measure, the shock
      #probability distribution is adjusted
      eta_shock = random_choice(params.eta_val,
                                params.eta_prob*params.eta_val)

   #Death shock
   death_shock = random_choice(np.array([0,1]),
   np.array([1-params.death_prob, params.death_prob]))

   if death_shock == 0:
      #Savings are cash on hand minus (linearly intepolated) consumption
      b = (m-linint(m, params.mgrid, consumption_function))
      m_new = wage*epsilon_shock + R*b/eta_shock

   if distorted_probabilities == False:
      P_new = eta_shock*P
   else:
      P_new = 1.0 #With permanent-income-neutral measure, the permanent
                  #income should not be updated.
   else:
      b = 0.0

   m_new = wage*epsilon_shock
   P_new = 1.0

   return m_new, P_new, b
```

Figure A.4: Python code for simulating the evolution of one agent for one period. This code is called repeatedly to generate a sample path of normalized cash on hand m and permanent income P. The permanent-income-neutral measure is used when distorted_probabilities is set to True. When the permanent-income-neutral measure is used, the permanent income shock (eta_shock) is drawn using the permanent-income-neutral measure $\tilde{f}(\eta) = \eta f(\eta)$ and permanent income (P and P_new) is kept at 1.0.