Shock Tree versus Markov Chain in the Life-cycle Income Fluctuation Model*

Chintal A. Desai† Gautham Udupa‡ Sripad Atri§

June 24, 2026

Abstract

We study how discretizing a unit-root income process in a life-cycle income fluctuation model affects consumption-saving behavior and wealth outcomes. The macroeconomics literature typically discretizes the income state, whereas the asset pricing literature emphasizes discretization of income innovations. We begin by comparing two discrete approximations—(i) a Rouwenhorst-type Markov chain and (ii) a recombining binomial shock tree—against a continuous-normal benchmark. We find that the binomial shock-tree discretization closely matches benchmark moments, while increasing the number of states in the Rouwenhorst approach improves its accuracy and leads to convergence toward the benchmark. In addition, we show that a pentanomial shock-tree discretization closely matches the higher-order moments of a mixture-of-normals benchmark. These results suggest that discretizing innovations—an idea borrowed from finance—can provide a useful alternative to the state-based discretization methods commonly used in macroeconomics.

JEL Codes: G5; C63.
Keywords: Household finance; Discretization; Binomial Tree; Pentanomial Tree.

*We benefited from the replication toolkit of Fella et al. (2019). Microsoft 365 Copilot also contributed on this research. The standard disclaimer applies.
†Virginia Commonwealth University. Email: cdesai2@vcu.edu Corresponding Author
‡CAFRAL, Mumbai, India. Email: gautham.udupa@cafral.org.in
§CAFRAL, Mumbai, India. Email: sripad.atri@cafral.org.in

1 Introduction

Quantitative life-cycle models of household consumption, saving, and wealth accumulation typically require discretizing the stochastic income process. A common macroeconomic approach replaces the continuous income state with a Markov chain of finite income states. Standard methods include Tauchen-type approximations and their refinements, as well as the Rouwenhorst construction based on moment matching (Tauchen, 1986; Adda and Cooper, 2003; Tauchen and Hussey, 1991; Rouwenhorst, 1995). These numerical choices can affect computed consumption-saving policies and, in turn, the model's implications for wealth accumulation and distributional outcomes. The issue is particularly relevant for nonstationary income risk, such as a unit-root (random walk) process, where dispersion within a cohort grows with age. At the same time, the asset pricing literature often discretizes innovations directly using recombining shock-tree methods (Cox et al., 1979). This innovation-based discretization approach offers a natural alternative to the state-based methods typically used in macroeconomics.

In this paper, we study how the choice of income discretization affects quantitative outcomes in a life-cycle income fluctuation model. We assume that log income follows a unit-root process. We begin by comparing two approaches: a Markov-chain discretization of the income state using the Rouwenhorst method and a discretization of income innovations using a recombining binomial shock tree. We evaluate these two approaches against a continuous-normal benchmark. We then extend the innovation-based approach by considering a pentanomial shock tree, which allows us to match not only the mean and variance but also the skewness and kurtosis of the innovation distribution (Primbs et al., 2007). Finally, we assess the performance of this pentanomial discretization against a mixture-of-normals benchmark (Fella et al., 2019).

We find that the choice of discretization affects not only income and consumption dispersion but also wealth accumulation and wealth concentration. With a small number of income states, the Rouwenhorst approximation understates dispersion in income and consumption and attenuates wealth concentration. In contrast, the binomial shock-tree discretization closely matches the moments of the continuous-normal benchmark. Increasing the number of income states improves the accuracy of the Rouwenhorst approximation and leads to convergence toward the continuous-normal benchmark. We also find that extending the innovation-based approach to a pentanomial shock tree allows the model to capture higher-order moments of the income process. In particular, the skewness and kurtosis of income, consumption, and assets generated by the pentanomial discretization closely match those obtained from a mixture-of-normals benchmark. These findings suggest that innovation-based discretizations provide a useful alternative to conventional state-based methods in life-cycle models.

Our paper lies at the intersection of macroeconomics and finance. In quantitative life-cycle macro models, idiosyncratic income risk plays a central role, and discretization choices can affect model predictions and policy conclusions. Kopecky and Suen (2010) show that the Rouwenhorst method provides an accurate approximation for highly persistent income processes, while Fella et al. (2019) find a similar result for nonstationary income processes. More recently, Guvenen et al. (2021) and De Nardi et al. (2020) document non-Gaussian features of income risk and highlight the importance of capturing higher-order moments in quantitative models. To address this problem, we draw on the finance literature, where uncertainty is often modeled using recombining trees (Cox et al., 1979). To capture higher-order moments, this literature also suggests recombining pentanomial trees (Primbs et al., 2007). Our paper brings recombining shock-tree methods and Markov-chain discretizations into a common life-cycle framework and compares their implications for consumption-saving decisions and wealth outcomes.

2 Model

2.1 Economic Environment

We consider a standard finite-horizon life-cycle income fluctuation model. Time is discrete, and the household lives for T periods, indexed by $t = 1, \dots, T$. In each period, the household chooses consumption and savings to maximize lifetime utility subject to a borrowing constraint.

Preferences are given by time-separable log utility over consumption. The household discounts the future at a constant rate $\beta \in (0, 1)$. The lifetime objective function is

$$\mathbb{E}_0 \left[\sum_{t=1}^{T} \beta^{t-1} \log c_t \right]. \tag{1}$$

At each age t, the household enters the period with assets a_t and receives labor income y_t. Assets evolve according to the budget constraint

$$c_t + a_{t+1} = (1 + r)a_t + y_t, \tag{2}$$

where r is a constant interest rate. The household faces a borrowing constraint $a_{t+1} \geq 0$.

Income is exogenous and follows a unit-root (random walk) process in logs. Specifically, log income evolves as

$$\log y_{t+1} = \log y_t + \varepsilon_{t+1}, \tag{3}$$

where the innovation ε_{t+1} is independently and identically distributed (i.i.d.) with mean zero and constant variance σ_ε^2. There is no transitory component, hence income risk is entirely driven by the persistent innovation. Because shocks accumulate under the unit-root process, cross-sectional income dispersion increases with age.

The household's problem can be written recursively as a standard dynamic programming problem with state variables (a_t, y_t). The value function satisfies

$$V_t(a_t, y_t) = \max_{c_t,\, a_{t+1} \geq 0} \left\{ \log c_t + \beta \, \mathbb{E}_t \big[V_{t+1}(a_{t+1}, y_{t+1}) \big] \right\}, \tag{4}$$

subject to the budget constraint and the income process.

2.2 Solution Method (EGM)

We follow Fella et al. (2019) and solve the household's dynamic programming problem using the endogenous grid method (EGM) (Carroll, 2006). The EGM is well suited to our setting because it avoids explicit root-finding and is computationally efficient. The method uses the Euler equation to recover current consumption from a grid of next-period assets. Specifically, for each value of next-period assets, the Euler equation implies current consumption and, hence, an endogenous value of cash-on-hand.[1] Because these endogenous cash-on-hand values do not generally coincide with the exogenous grid, the consumption policy is interpolated back onto the exogenous grid. This solution method allows us to solve the model repeatedly and consistently across discretization schemes. Appendix A provides a detailed two-period illustration of the EGM.

3 Markov Chain versus Binomial

This section describes the two approaches used to discretize the income process in the life-cycle model. In both cases, the model environment and solution method remain the same as in Section 2; the only difference lies in how income risk is approximated.

The first approach discretizes the income state using a finite-state Markov chain constructed via the Rouwenhorst method. The second approach discretizes the income innovations using a recombining (binomial) shock-tree. While both methods provide tractable

[1] Cash-on-hand is defined as $z_t \equiv (1 + r)a_t + y_t$, that is, the resources available before the consumption-saving decision.

approximations of income risk, they differ in how the distribution of income evolves over the life-cycle.

The following subsections describe each discretization method in detail and explain how they are implemented within the EGM framework.

3.1 Markov-Chain Discretization (Rouwenhorst)

In quantitative macro models of the household consumption-saving problem, the persistent component of the income process is typically assumed to follow a first-order autoregressive process. Tauchen (1986), Adda and Cooper (2003), and Rouwenhorst (1995) provide popular Markov chain methods for approximating this process. In all three methods, for a stationary income process, the number and values of income states, as well as the transition matrix, remain constant across ages.

Kopecky and Suen (2010) compare the performance of these methods in an infinite-horizon income fluctuation model with a stationary persistent income component. They document that the Rouwenhorst method performs particularly well, especially for highly persistent processes. Fella et al. (2019) extend these methods to nonstationary income processes by allowing the income grid and transition matrix to vary with age while keeping the number of states fixed. They show that the Rouwenhorst construction can perform well in such environments even with a relatively small number of states.

This result is particularly relevant in our setting, where the income process follows a unit-root (random walk) specification. An important property of the Rouwenhorst method is that it preserves the conditional mean of the income process through the transition probabilities. At the same time, the income grid is constructed so that its spacing expands with age to match the growth in unconditional variance implied by the unit-root process. Appendix B provides a simple illustration of the Rouwenhorst method for the case of two income states ($N = 2$) and autoregressive coefficient $\rho = 1$.

In applying the Rouwenhorst method, we follow Fella et al. (2019). Log income at each age t takes values on a grid with N points, and transitions across these values are governed by an $N \times N$ transition matrix. The income grid and transition matrix are allowed to vary across age while keeping the number of states fixed. In the implementation of the EGM, expectations over future income are computed using the transition matrix.

From a conceptual perspective, the Rouwenhorst method is based on a moment matching approach that is similar in spirit to binomial approximations of geometric Brownian motion of stock prices used in finance (Cox et al., 1979). In such approximations, the jump size and branch (transition) probabilities are chosen to match key moments—mean and variance—of the underlying continuous process.

3.2 Binomial Shock-Tree Discretization

The second approach discretizes the income innovations rather than the income state. In our setting, as shown in equation (3), the innovation ε_t is i.i.d. with mean zero and variance σ_ε^2. To approximate this continuous shock, we replace it with a finite set of discrete outcomes.

A natural starting point is a two-point approximation to the innovation. To match the zero mean of the continuous distribution, these states must be symmetric around zero, so that

$$\varepsilon_t \in \{-a,\ +a\}. \tag{5}$$

Symmetry and zero mean further imply that the two outcomes occur with equal probability.

To determine the magnitude of the shock, we match the variance of the continuous distribution. For the two-point approximation, the variance is given by a^2, so matching $\mathrm{Var}(\varepsilon_t) = \sigma_\varepsilon^2$ implies

$$a = \sigma_\varepsilon. \tag{6}$$

Thus, the discrete approximation to the innovation is

$$\varepsilon_t \in \{-\sigma_\varepsilon,\ +\sigma_\varepsilon\}, \quad \text{each with probability } \tfrac{1}{2}. \tag{7}$$

Given this discrete representation of the innovation, log income evolves through the accumulation of shocks over time. Starting from an initial value, income at age t can be written as

$$\log y_t = \log y_0 + \sum_{s=1}^{t} \varepsilon_s. \tag{8}$$

Since each innovation takes one of two values, each realization of income corresponds to a sequence of positive and negative shocks. This generates a recombining tree structure, in which nodes represent distinct income realizations.

An important feature of this construction is that the number of distinct income states increases linearly with age. At age t, the set of possible values of $\log y_t$ is given by all combinations of $\pm\sigma_\varepsilon$, resulting in $N_t = t + 1$ distinct income states.

The implementation of the EGM under the binomial shock-tree discretization proceeds as follows. At a given age t, a household is located at a particular income node and enters the period with a given level of assets. In the next period, age $t+1$, income innovation takes one of two possible values—equal in magnitude and opposite in sign—with equal probability. This allows us to compute the expected marginal utility of next period consumption. Using the Euler equation, we then recover the corresponding endogenous current consumption and cash-on-hand levels. Finally, interpolation is used to obtain policy functions at that income node.

In Appendix C, we illustrate the EGM under binomial shock-tree discretization using a two-period example. Appendix D provides a corresponding two-period illustration under the Rouwenhorst (Markov-chain) discretization, allowing for a direct comparison of the resulting consumption and saving policies.

3.3 Continuous-Normal Benchmark

To evaluate the accuracy of the discrete approximations, we follow Fella et al. (2019) and use a benchmark in which the income process remains continuous. In this continuous-normal benchmark, the household's problem is identical to that described in Section 2.1, except that neither the income state nor the income innovations are discretized.

Expectations over future income are computed using Gauss–Hermite quadrature. This method approximates conditional expectations by integrating over the distribution of income innovations using a set of quadrature nodes and associated weights. It provides an accurate and computationally efficient way to evaluate expectations when shocks are normally distributed.

To implement this approach numerically, we construct a dense grid for the income state. Given current income and a realization of the Gaussian innovation, next-period income generally does not lie exactly on this grid. We therefore use interpolation to evaluate the continuation value at off-grid income realizations.

The model is solved using the EGM, as described in Section 2.2. The continuous-income solution serves as a reference point for assessing discretization error. By comparing the distributions of income, consumption, and wealth implied by alternative discrete approximations with those generated under the continuous specification, we can isolate the effect of discretizing the income process.

3.4 Results

This subsection reports the results of our comparison of two income discretization methods—the Rouwenhorst (Markov-chain) approximation and the binomial shock tree—against a continuous-normal benchmark. We consider a finite-horizon income fluctuation model, and for a meaningful comparison, all model components other than the income discretization are held fixed.

Following Fella et al. (2019), preferences are logarithmic, the discount factor is set to $\beta = 0.96$, and the gross interest rate is $R = 1 + r = 1.04$. The length of the life-cycle is $T = 40$, and the variance of the income innovation is $\sigma_\varepsilon^2 = 0.0161$.

We use an equally spaced asset grid with 1,000 points on $[0, 50]$. For the continuous-income benchmark, we approximate the income state using an equally spaced and symmetric grid of 10,000 points; the upper bound of this grid is set to twice the cumulative standard deviation of income innovations over the life-cycle. Expectations in the benchmark are evaluated using Gauss–Hermite quadrature with five nodes, combined with interpolation on the income grid.

Given these specifications, we solve the household's problem using the EGM in conjunction with backward induction. For each age, income realization, and asset level, we obtain consumption and asset policy functions. We then simulate the model for 2,000,000 households and compute pooled moments across all ages.[2] In particular, we report the mean and standard deviation of income (in levels), consumption, and assets, as well as the top 5% wealth share.

Table 1 reports the results. Mean levels of income and consumption are close across the three approaches. The main quantitative differences arise in dispersion and concentration statistics. Relative to the continuous-normal benchmark, the Rouwenhorst approximation with $N = 2$ substantially understates dispersion in both income and consumption, whereas the binomial shock tree produces dispersion measures that are much closer to the benchmark. For example, the benchmark standard deviation of income is 0.8164, compared with 0.6251 under Rouwenhorst and 0.8073 under the binomial shock tree.

These differences in dispersion are accompanied by differences in asset accumulation and wealth concentration. Mean assets are higher under Rouwenhorst (0.7216) than in the benchmark (0.6845), while the binomial shock tree closely matches the benchmark mean assets (0.6842). Regarding wealth concentration, the benchmark top 5% wealth share is 0.1306, the binomial shock tree generates 0.1297, and the Rouwenhorst discretization yields 0.1165.

Taken together, these results suggest that, in the unit-root (random walk) setting, discretizing income innovations via a recombining binomial shock tree tends to deliver distri-

[2] In the simulations, households enter age 1 with initial assets equal to zero.

butional statistics close to the continuous-normal benchmark, whereas a coarse income-state discretization can distort dispersion and inequality measures.

A useful way to interpret these results is through how each discretization captures the growth of income dispersion in a unit-root process. The shock-tree discretization builds income from accumulated innovations and generates an expanding set of income realizations over the life-cycle. It closely tracks the increase in dispersion implied by the underlying process and, as a result, provides a good approximation to the continuous-normal benchmark. In contrast, a coarse Markov-chain discretization with a fixed number of states ($N = 2$) compresses the income distribution and tends to understate dispersion, which is reflected in consumption and wealth outcomes.

We next analyze how the Rouwenhorst approach performs when the number of income states is increased. Specifically, following Fella et al. (2019), we consider $N = 5$, $N = 10$, and $N = 25$. Implementing the Rouwenhorst method with a larger number of income states requires constructing a transition matrix of size $N \times N$. This matrix can be obtained recursively, starting from the 2×2 transition matrix. In addition, the symmetric and evenly spaced income grid is adjusted while keeping the same upper and lower bounds as in the $N = 2$ case.[3]

Table 2 reports the same moments for the Rouwenhorst discretization with a larger number of income states. As N increases from 2 to 5, 10, and 25, the approximation improves and the results move closer to the continuous-normal benchmark values reported in Table 1. In particular, dispersion in income and consumption increases and approaches the benchmark levels, while wealth concentration also becomes closer to the benchmark. This indicates that the distortions observed with $N = 2$ are driven by the coarse discretization of the income process.

[3] For details on the recursive construction of the Rouwenhorst transition matrix and the income grid for larger values of N, see Kopecky and Suen (2010).

4 Shock-tree for Higher-order Moments

4.1 Pentanomial Shock-tree Discretization

Section 3 showed that lower-order moments—in particular, the mean and variance of income, consumption, and assets—generated by the recombining binomial shock tree closely match those obtained from the continuous-normal benchmark. Recent studies, however, document non-Gaussian features of income dynamics over the life-cycle, including asymmetry and fat tails (De Nardi et al., 2020; Guvenen et al., 2021). When the innovation process exhibits non-zero skewness and excess kurtosis, a richer innovation-based discretization is needed.

To address this issue, we extend the shock-tree approach by adopting a recombining pentanomial lattice. Our construction follows Primbs et al. (2007), who develop a pentanomial tree in the context of option pricing. Their approach uses moment matching to incorporate skewness and kurtosis. We apply this idea to income-shock discretization in the life-cycle model. In this sense, the pentanomial tree is a natural extension of the binomial shock tree used in Section 3.

A quadrinomial tree is not especially useful for this purpose. Since the probabilities must sum to one, a four-point discretization leaves limited flexibility to match the mean, variance, skewness, and kurtosis of the innovation distribution. At the same time, it becomes harder to preserve a convenient recombining structure and non-negative branch probabilities. By contrast, a five-point support provides the additional flexibility needed to approximate higher-order moments directly.

Let ε_t denote the innovation to log income in period t, as in equation (3). In the pentanomial construction, we approximate the innovation distribution using the five-point support

$$\varepsilon_t \in \{-4\alpha, -2\alpha, 0, 2\alpha, 4\alpha\},$$

with associated branch probabilities

$$(p_{-2},\, p_{-1},\, p_0,\, p_{+1},\, p_{+2}).$$

The step size α and the branch probabilities are chosen so that the implied discrete distribution matches the target moments of the innovation process. In our application, the mean of the innovation is normalized to zero, so the pentanomial discretization is constructed to match the variance, skewness, and kurtosis of the innovation distribution, together with the probability adding-up restriction. This is the key distinction relative to the binomial case, which is designed to approximate the lower-order moments of the innovation process. Following Primbs et al. (2007), the pentanomial lattice is admissible only if the implied branch probabilities are non-negative and sum to one.

As in the binomial case, the discretization is applied to the innovation term rather than to the income state itself. Since log income evolves according to $y_{t+1} = y_t + \varepsilon_{t+1}$, the pentanomial discretization generates a recombining lattice for log income over the life-cycle. The support is centered at zero and the innovation is discretized directly, so the resulting tree remains computationally tractable while allowing for richer distributional features.[4]

4.2 Mixture-of-Normal Benchmark

Section 4.1 introduced a pentanomial shock tree to discretize the income process while matching higher-order moments of its innovation term. To evaluate the accuracy of the pentanomial discretization, a continuous benchmark with the same non-Gaussian features is needed.

We again follow Fella et al. (2019) and use a benchmark in which the innovation to log

[4]For derivations of the branch probabilities and the admissible range of values of α that ensure non-negative branch probabilities, see Proposition 1 and equation (5) of Primbs et al. (2007).

income is modeled as a mixture-of-normal distributions. Specifically,

$$\varepsilon_t \sim \begin{cases} \mathcal{N}(\mu_1, \sigma_1^2) & \text{with probability } \pi_1, \\ \mathcal{N}(\mu_2, \sigma_2^2) & \text{with probability } 1 - \pi_1. \end{cases}$$

Because the innovation is drawn from two distinct normal distributions, the resulting distribution is non-normal and can exhibit non-zero skewness and excess kurtosis.[5]

Expectations over future income are computed by integrating over the mixture distribution of innovations. Numerically, we implement this integration using Gauss–Hermite quadrature. For each normal component, we evaluate the relevant conditional expectation at a set of Gauss-Hermite quadrature nodes and weights, and then combine the component-specific expectations using the corresponding mixture probabilities.

To implement the benchmark numerically, we again construct a dense grid for the income state. Given current income and a realization of the mixture-of-normal innovation, next-period income generally does not lie exactly on this grid. We therefore use interpolation to evaluate the continuation value at off-grid income realizations. The model is solved using the EGM described in Section 2.2.

4.3 Results

Following Fella et al. (2019), we set the variance of the innovation term to $\sigma_\varepsilon^2 = 0.0161$, the skewness to -1.36, and the kurtosis to 17.95. In the mixture-of-normal benchmark, the innovation to log income is distributed as $\mathcal{N}(0.0089, 0.0635^2)$ with probability 0.9000 and as $\mathcal{N}(-0.0800, 0.3430^2)$ with probability 0.1000.

Table 3 compares the moment-based statistics—mean, standard deviation, skewness, and kurtosis—generated by the pentanomial shock tree with those of the continuous mixture-of-normal benchmark. The pentanomial discretization closely matches the moments of income

[5] Timmermann (2000) and Civale et al. (2016) derive moments for stochastic processes in which the innovation term follows a mixture-of-normals distribution.

and consumption. In particular, the mean and standard deviation of income and consumption are nearly identical under the pentanomial discretization and the benchmark. The pentanomial tree also provides a close approximation to the skewness and kurtosis of these variables.

The approximation is somewhat less precise for assets, but the pentanomial discretization still captures the main features of the benchmark asset distribution. In particular, the skewness and kurtosis of assets remain close to their benchmark counterparts, and the top 5% wealth share differs only slightly across the two specifications. Taken together, these results suggest that the pentanomial shock tree provides a useful approximation to the continuous benchmark when higher-order moments of the innovation process matter.

Table 4 reports a robustness check based on quantile-based shape statistics. The same pattern emerges. Kelly skewness and Crow–Siddiqui (CS) kurtosis of income, consumption, and assets under the pentanomial discretization are close to those obtained from the mixture-of-normal benchmark.[6] The similarity of these quantile-based measures reinforces the findings in Table 3. Overall, in a quantitative life-cycle model, the pentanomial shock-tree discretization provides a useful way to approximate an income process with higher-order moments.

5 Conclusion

Our main finding is that innovation-based discretizations provide a useful alternative to conventional state-based methods in life-cycle models. The binomial shock tree provides a close approximation to the continuous-normal benchmark, while the pentanomial shock-tree discretization extends this approach to settings in which higher-order moments of the income process matter. Taken together, these results show that ideas borrowed from the finance literature can be fruitfully adapted to macroeconomic models of household income risk.

[6] For the computation of Kelly skewness and Crow–Siddiqui kurtosis, see Fella et al. (2019).

References

J. Adda and R. W. Cooper. *Dynamic economics: quantitative methods and applications*. MIT press, 2003.

C. D. Carroll. The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91(3):312–320, 2006.

S. Civale, L. Díez-Catalán, and F. Fazilet. Discretizing a process with non-zero skewness and high kurtosis. *Available at SSRN 2636485*, 2016.

J. C. Cox, S. A. Ross, and M. Rubinstein. Option pricing: A simplified approach. *Journal of Financial Economics*, 7(3):229–263, 1979. doi: 10.1016/0304-405X(79)90015-1.

M. De Nardi, G. Fella, and G. Paz-Pardo. Nonlinear household earnings dynamics, self-insurance, and welfare. *Journal of the European Economic Association*, 18(2):890–926, 2020. doi: 10.1093/jeea/jvz010.

G. Fella, G. Gallipoli, and J. Pan. Markov-chain approximations for life-cycle models. *Review of Economic Dynamics*, 34:183–201, 2019. doi: 10.1016/j.red.2019.03.013.

F. Guvenen, F. Karahan, S. Ozkan, and J. Song. What do data on millions of u.s. workers reveal about lifecycle earnings dynamics? *Econometrica*, 89(5):2303–2339, 2021. doi: 10.3982/ECTA14603.

K. A. Kopecky and R. M. H. Suen. Finite state markov-chain approximations to highly persistent processes. *Review of Economic Dynamics*, 13(3):701–714, 2010. doi: 10.1016/j.red.2010.02.002.

J. A. Primbs, M. Rathinam, and Y. Yamada. Option pricing with a pentanomial lattice model that incorporates skewness and kurtosis. *Applied Mathematical Finance*, 14(1): 1–17, 2007.

K. G. Rouwenhorst. Asset pricing implications of equilibrium business cycle models. In T. F. Cooley, editor, *Frontiers of Business Cycle Research*, pages 294–330. Princeton University Press, 1995.

G. Tauchen. Finite state markov-chain approximations to univariate and vector autoregressions. *Economics Letters*, 20(2):177–181, 1986. doi: 10.1016/0165-1765(86)90168-0.

G. Tauchen and R. Hussey. Quadrature-based methods for obtaining approximate solutions to nonlinear asset pricing models. *Econometrica*, 59(2):371–396, 1991.

A. Timmermann. Moments of markov switching models. *Journal of econometrics*, 96(1): 75–111, 2000.

Table 1: Pooled Moments: Normal Benchmark and Discrete Approximations

Moment	Normal Benchmark	Rouwenhorst (Markov chain, $N = 2$)	Binomial Shock Tree ($N_\varepsilon = 2$)
Mean(y)	1.1839	1.1711	1.1840
SD(y)	0.8164	0.6251	0.8073
Mean(c)	1.2113	1.2000	1.2114
SD(c)	0.8386	0.6525	0.8294
Mean(a)	0.6845	0.7216	0.6842
SD(a)	0.4258	0.4984	0.4237
Top 5% wealth share	0.1306	0.1165	0.1297

Notes: Income is reported in levels ($y = \exp(\log y)$), and moments are pooled across all ages. The top 5% wealth share is the share of total assets held by the richest 5% of households in the simulated cross-section.

Table 2: Rouwenhorst Discretization with Larger Numbers of States

Moment	$N = 5$	$N = 10$	$N = 25$
Mean(y)	1.1807	1.1831	1.1840
SD(y)	0.7525	0.7870	0.8059
Mean(c)	1.2087	1.2107	1.2115
SD(c)	0.7766	0.8102	0.8285
Mean(a)	0.6981	0.6912	0.6873
SD(a)	0.4525	0.4392	0.4313
Top 5% wealth share	0.1354	0.1335	0.1318

Notes: The table reports the same moments as in Table 1 for the Rouwenhorst discretization with larger numbers of income states. The continuous-normal benchmark and the $N = 2$ case are reported in Table 1.

Table 3: Pooled Moments: Mixture-of-Normal Benchmark and Pentanomial Discretization

Moment	Mixture-of-Normal Benchmark	Pentanomial Shock Tree
Mean(y)	1.1762	1.1763
SD(y)	0.7866	0.7869
Skew(y)	4.4540	4.4155
Kurt(y)	64.6433	63.7549
Mean(c)	1.2024	1.2048
SD(c)	0.8074	0.8098
Skew(c)	4.5212	4.4851
Kurt(c)	64.6633	63.6049
Mean(a)	0.6542	0.7124
SD(a)	0.4020	0.4384
Skew(a)	1.2312	1.2258
Kurt(a)	9.0346	8.4913
Top 5% wealth share	0.1282	0.1293

Notes: Income is reported in levels ($y = \exp(\log y)$), and moments are pooled across all ages. Skewness and kurtosis are ordinary moment-based measures (not Kelly skewness or Crow–Siddiqui kurtosis). The top 5% wealth share is the fraction of total assets held by the richest 5% of households.

Table 4: Robustness Check: Quantile-Based Shape Statistics

Statistic	Mixture-of-Normal Benchmark	Pentanomial Shock Tree
Kelly skewness of y	0.2711	0.3008
CS kurtosis of y	4.5858	5.0286
Kelly skewness of c	0.3434	0.3579
CS kurtosis of c	4.6749	5.1210
Kelly skewness of a	0.0415	0.0311
CS kurtosis of a	3.0100	3.0424
Top 5% wealth share	0.1282	0.1293

Notes: This table reports a robustness check based on quantile-based shape statistics. The pentanomial shock tree uses the Primbs et al. (2007) construction. For all variables, shape statistics are computed using levels and are pooled across all ages. Kelly skewness is defined as $(p_{90}+p_{10}-2p_{50})/(p_{90}-p_{10})$, and Crow–Siddiqui (CS) kurtosis is defined as $(p_{97.5}-p_{2.5})/(p_{75}-p_{25})$. The top 5% wealth share is the fraction of total assets held by the richest 5% of households.

Appendix A: Endogenous Grid Method: A Two-Period Illustration

A.1 Setup of the Two-Period Problem

To illustrate the intuition behind the EGM, we consider a simple two-period version of the consumption–savings problem. Time is indexed by $t = 1, 2$. The household chooses consumption and savings in period 1, knowing that it consumes all remaining resources in period 2.

Preferences are given by

$$\log c_1 + \beta\, \mathbb{E}[\log c_2], \tag{9}$$

where $\beta \in (0, 1)$ is the discount factor.

In period 1, the household enters with assets a_1 and receives income y_1. The budget constraint is

$$c_1 + a_2 = (1 + r)a_1 + y_1, \tag{10}$$

where a_2 denotes savings carried into period 2. The household faces a borrowing constraint $a_2 \geq 0$. In this setting, consumption c_1 and next-period assets a_2 are choice (policy) variables, while current assets a_1 and income y_1 summarize the state of the household at the beginning of the period.

Let income in period 2 be uncertain and given by y_2. Since the household consumes all available resource in period 2, the consumption in period 2 is

$$c_2 = (1 + r)a_2 + y_2. \tag{11}$$

It is convenient to define cash-on-hand as

$$z_t = (1 + r)a_t + y_t, \tag{12}$$

which represents total resources available before consumption and saving decisions. Using this notation, the period 1 budget constraint becomes $c_1 + a_2 = z_1$.

A.2 Optimality Condition and Construction of Endogenous Grid

The household chooses consumption in period 1 and savings for period 2 to maximize lifetime utility subject to the budget constraint. The first order condition for this problem is given by the Euler equation, which equates the marginal utility of consumption across periods.

For log utility, the Euler equation takes the form

$$\frac{1}{c_1} = \beta(1+r)\,\mathbb{E}\left[\frac{1}{c_2}\right]. \tag{13}$$

This condition has the following economic interpretation: Saving one additional unit in period 1 reduces current consumption by one unit, which increases marginal utility by $1/c_1$. The saved unit yields $(1+r)$ units of resources in period 2, which can be used for consumption. The expected marginal utility of this additional consumption in period 2 is therefore $\mathbb{E}_{y}[1/c_2]$, which reflects uncertainty in future income y_2. The Euler equation requires that the marginal cost of saving today equals the expected marginal benefit tomorrow, discounted by β.

The budget constraint of period 2, $c_2 = (1+r)a_2 + y_2$, allows the Euler equation to link the choice of next-period assets a_2 to current consumption c_1. This relationship is the key ingredient of the EGM: given a level of savings a_2, the Euler equation determines the corresponding optimal consumption c_1.

For each value of a_2 on an asset grid, we compute the corresponding optimal level of current consumption using period 2 budget constraint (11) and Euler equation (13). We denote this consumption by c_1^{endo} to emphasize that it is implied by the optimality condition. Given c_1^{endo}, the definition of cash-on-hand implies $z_1^{\text{endo}} = c_1^{\text{endo}} + a_2$.

Thus, each grid point for a_2 generates a pair $(z_1^{\text{endo}}, c_1^{\text{endo}})$, where z_1^{endo} and c_1^{endo} denote the level of cash-on-hand and consumption consistent with optimal behavior.

A.3 Why Interpolation Is Needed

The construction in the preceding subsection produces a set of $(z_1^{\text{endo}}, c_1^{\text{endo}})$ pairs by starting from a grid over next-period assets a_2 and recovering the corresponding levels of cash-on-hand z_1^{endo} and optimal consumption c_1^{endo}. Because z_1^{endo} is determined by the Euler equation, these values do not, in general, lie on a fixed and evenly spaced grid.

However, in solving dynamic models, policy functions must be evaluated on an exogenous grid for the state variables. Let z_1^{grid} denote such a fixed grid for cash-on-hand. This grid is chosen in advance and is independent of the optimality condition.

The EGM therefore requires mapping the policy function from the irregular (endogenous) grid z_1^{endo} onto the fixed (exogenous) grid z_1^{grid}. This is accomplished using interpolation, typically linear interpolation. In particular, given the values of c_1^{endo} defined at z_1^{endo}, interpolation is used to construct the consumption policy function $c_1(z_1^{\text{grid}})$.

Once consumption is obtained on the fixed grid, the asset policy is recovered from the budget constraint as

$$a_2(z_1^{\text{grid}}) = z_1^{\text{grid}} - c_1(z_1^{\text{grid}}). \tag{14}$$

A.4 Numerical Illustration

The preceding subsections explain the logic of the endogenous grid method. For numerical two-period illustrations in the context of the discretizations used in this paper, see Appendix C (binomial shock-tree discretization) and Appendix D (Rouwenhorst discretization).

Appendix B: Rouwenhorst Approximation in a Nonstationary Setting

B.1 Overview of the Rouwenhorst Method

The Rouwenhorst method provides a discrete Markov-chain approximation to a continuous autoregressive process for log income. Let $x_t = \log y_t$ and consider

$$x_{t+1} = \rho x_t + \varepsilon_{t+1}, \tag{15}$$

where ε_{t+1} has mean zero and variance σ_ε^2. The method represents this process using an income grid with N states and an associated $N \times N$ transition matrix.

The key idea is a moment matching construction. Transition probabilities preserve the conditional mean (and hence persistence), while the spacing of the income grid is chosen to match the relevant variance behavior implied by the continuous process. In stationary settings, the grid and transition matrix are time-invariant. In the nonstationary setting considered in this paper, these objects are allowed to vary with age while keeping the number of states fixed. The next subsection specializes this construction to the unit-root case, and Section B.3 provides a worked example for $N = 2$.

B.2 Nonstationary Extension

We consider a nonstationary income process in which log income follows a unit-root (random walk) specification, $\rho = 1$ in (15).

This implies that the conditional moments of the process are given by

$$\mathbb{E}[x_{t+1} \mid x_t] = x_t, \qquad \mathrm{Var}(x_{t+1} \mid x_t) = \sigma_\varepsilon^2,$$

while the unconditional variance grows over time:

$$\mathrm{Var}(x_t) = t\sigma_\varepsilon^2.$$

As a result, the cross-sectional dispersion of income increases with age.

To approximate this process using the Rouwenhorst method, we retain a fixed number of income states but allow the values of the income grid to vary across age. In particular, the income grid expands over time in order to match the increasing unconditional variance of the continuous process. At the same time, the transition probabilities are chosen to preserve the conditional mean of the process at each age. The next subsection illustrates this construction using an example with $N = 2$.

B.3 Worked Example: $N = 2$ and $\rho = 1$

Building on the nonstationary specification described in the previous subsection, we illustrate the Rouwenhorst approximation in the case of two income states ($N = 2$).

At each age t, the distribution of log income x_t is approximated by two symmetric points,

$$x_t \in \{-\sigma_{x,t},\ +\sigma_{x,t}\}, \tag{16}$$

where $\sigma_{x,t}$ is chosen to match the unconditional variance of the continuous process. Since $\mathrm{Var}(x_t) = t\sigma_\varepsilon^2$, this implies

$$\sigma_{x,t} = \sqrt{t}\,\sigma_\varepsilon. \tag{17}$$

Thus, the grid expands with age, reflecting the growth in income dispersion.

At age $t + 1$, the grid is given by

$$x_{t+1} \in \{-\sigma_{x,t+1},\ +\sigma_{x,t+1}\}, \quad \sigma_{x,t+1} = \sqrt{t+1}\,\sigma_\varepsilon. \tag{18}$$

Transitions across these states, from age t to age $t+1$, are governed by a symmetric matrix of the form

$$\Pi_t = \begin{pmatrix} p_t & 1-p_t \\ 1-p_t & p_t \end{pmatrix}. \tag{19}$$

Here, p_t is the probability of remaining in the same state from age t to age $t+1$, while $1-p_t$ is the probability of moving to the other state. The value of p_t is determined by matching the conditional mean of the continuous process.

To derive p_t, consider the state $x_t = \sigma_{x,t}$. Under the discrete approximation, the conditional expectation of next-period income is

$$\mathbb{E}[x_{t+1} \mid x_t = \sigma_{x,t}] = p_t \sigma_{x,t+1} + (1-p_t)(-\sigma_{x,t+1}) = (2p_t - 1)\sigma_{x,t+1}. \tag{20}$$

To match the conditional mean of the continuous process, we require

$$\mathbb{E}[x_{t+1} \mid x_t] = x_t = \sigma_{x,t}. \tag{21}$$

Equating these expressions yields

$$(2p_t - 1)\sigma_{x,t+1} = \sigma_{x,t}. \tag{22}$$

Substituting $\sigma_{x,t} = \sqrt{t}\sigma_\varepsilon$ and $\sigma_{x,t+1} = \sqrt{t+1}\sigma_\varepsilon$, we obtain

$$2p_t - 1 = \frac{\sqrt{t}}{\sqrt{t+1}},$$

which implies

$$p_t = \frac{1}{2}\left(1 + \frac{\sqrt{t}}{\sqrt{t+1}}\right). \tag{23}$$

This construction highlights the key mechanism underlying the Rouwenhorst approximation in the nonstationary setting. The spacing of the income grid is determined by matching the growth in unconditional variance, while the transition probabilities are chosen to preserve the conditional mean of the process. As age increases, the grid expands and the probability p_t approaches one, reflecting an increase in income persistence. The number of discrete income states, however, remains fixed ($N = 2$).

Appendix C: EGM Under Binomial Shock-Tree Discretization (Two-Period Illustration)

C.1 Setup

This appendix illustrates the implementation of the EGM under binomial shock-tree discretization for a single decision step from $t = 1$ to $t = 2$. At $t = 1$, the household has assets a_1 and is located at one of two income nodes. It chooses current consumption c_1 and next-period assets a_2. Income at $t = 2$ is uncertain. Finally, the household consumes all remaining resources at $t = 2$, since it is a terminal period.

Following Fella et al. (2019), preferences are logarithmic, the interest rate is $r = 0.04$ (so $R = 1 + r = 1.04$), and the discount factor is $\beta = 0.96$ (so $\beta R = 0.9984$). The borrowing constraint is $a_2 \geq 0$.

Income follows the unit-root process in logs,

$$\log y_2 = \log y_1 + \varepsilon_2, \qquad \varepsilon_2 \in \{-\sigma_\varepsilon, +\sigma_\varepsilon\}, \quad \mathbb{P} = \tfrac{1}{2}, \tfrac{1}{2},$$

with $\sigma_\varepsilon^2 = 0.0161$. We use the asset grid $\mathcal{A} = \{0,\ 0.5,\ 1\}$ both as the EGM grid for a_2 and as the feasible grid for a_1.

Cash-on-hand is defined as $z_t = Ra_t + y_t$, and the budget constraint at $t = 1$ is $c_1 + a_2 = z_1$.

C.2 Upper Income Node at $t = 1$

For illustration, we focus on the upper income node at $t = 1$:

$$\log y_1 = +\sigma_\varepsilon \qquad \Rightarrow \qquad y_1 = e^{\sigma_\varepsilon} = 1.1353.$$

Under shock-tree discretization, next-period ($t = 2$) income is generated by adding the income innovation to current log income:

$$\log y_2 \in \{\log y_1 - \sigma_\varepsilon,\ \log y_1 + \sigma_\varepsilon\} = \{0,\ 2\sigma_\varepsilon\}.$$

Thus,

$$y_2 \in \{e^0,\ e^{2\sigma_\varepsilon}\} = \{1.0000,\ 1.2889\}, \qquad \mathbb{P} = \tfrac{1}{2}, \tfrac{1}{2}.$$

Finally, the feasible cash-on-hand grid at this upper income node at $t = 1$, using $a_1 \in \mathcal{A}$, is

$$z_1^{\text{grid}} = Ra_1 + y_1 \in \{1.1353,\ 1.6553,\ 2.1753\}.$$

C.3 EGM Calculations at the Upper Income Node at $t = 1$

Since $t = 2$ is terminal, consumption equals cash-on-hand:

$$c_2 = z_2 = Ra_2 + y_2.$$

The EGM fixes $a_2 \in \mathcal{A}$ and computes

$$\frac{1}{c_1} = \beta R\,\mathbb{E}\left[\frac{1}{c_2}\right], \qquad c_1^{\text{endo}} = \frac{1}{\beta R\,\mathbb{E}[1/c_2]}, \qquad z_1^{\text{endo}} = c_1^{\text{endo}} + a_2.$$

Step 1: Compute c_2 for each a_2 and each shock. With $y_2 \in \{1.0000,\ 1.2889\}$ and $R = 1.04$:

$$\begin{aligned} a_2 = 0: &\quad c_2 \in \{1.0000,\ 1.2889\}, \\ a_2 = 0.5: &\quad c_2 \in \{1.5200,\ 1.8089\}, \\ a_2 = 1: &\quad c_2 \in \{2.0400,\ 2.3289\}. \end{aligned}$$

Step 2: Compute $\mathbb{E}[1/c_2]$ and invert the Euler equation. For each a_2,

$$\mathbb{E}\left[\frac{1}{c_2}\right] = \tfrac{1}{2}\frac{1}{c_2(y_2 = 1.0000)} + \tfrac{1}{2}\frac{1}{c_2(y_2 = 1.2889)}, \qquad c_1^{\text{endo}} = \frac{1}{0.9984\,\mathbb{E}[1/c_2]}.$$

This yields the endogenous objects (rounded to four decimals):

a_2	0	0.5	1
c_1^{endo}	1.1280	1.6546	2.1784
$z_1^{\text{endo}} = c_1^{\text{endo}} + a_2$	1.1280	2.1546	3.1784

Step 3: Interpolate consumption onto the feasible grid and recover next-period assets. Because z_1^{grid} does not coincide with z_1^{endo}, we linearly interpolate c_1 over the points $(z_1^{\text{endo}}, c_1^{\text{endo}})$ to obtain $c_1(z_1^{\text{grid}})$, and then compute

$$a_2(z_1^{\text{grid}}) = z_1^{\text{grid}} - c_1(z_1^{\text{grid}}).$$

For example, the middle grid point $z_1^{\text{grid}} = 1.6553$ lies between

$$z_1^{\text{endo}} = 1.1280 \qquad \text{and} \qquad z_1^{\text{endo}} = 2.1546.$$

Using linear interpolation between the two adjacent endogenous points,

$$c_1(1.6553) = 1.1280 + \frac{1.6553 - 1.1280}{2.1546 - 1.1280}(1.6546 - 1.1280) = 1.3985.$$

Applying the same procedure to all grid points yields:

z_1^{grid}	1.1353	1.6553	2.1753
$c_1(z_1^{\text{grid}})$	1.1317	1.3985	1.6652
$a_2(z_1^{\text{grid}})$	0.0035	0.2568	0.5101

The same procedure applies to the lower income node at $t = 1$, with the corresponding values of y_1 and y_2.

Appendix D: EGM Under Rouwenhorst Discretization (Two-Period Illustration)

D.1 Setup (upper node at $t = 1$)

This appendix provides a contrast to Appendix C by illustrating the EGM under the Rouwenhorst (Markov-chain) discretization. We focus on the upper income node at $t = 1$, using the same parameters and asset grids as in Appendix C: $r = 0.04$ (so $R = 1.04$), $\beta = 0.96$ (so $\beta R = 0.9984$), and $\mathcal{A} = \{0, 0.5, 1\}$.

At $t = 1$, the log income state is

$$\log y_1 = +\sigma_\varepsilon, \quad \sigma_\varepsilon = \sqrt{0.0161} = 0.1269,$$

so that $y_1 = e^{\sigma_\varepsilon} = 1.1353$.

Under the Rouwenhorst discretization for a unit-root process with $N = 2$, the support of next-period log income expands with age (see Appendix B). At $t = 2$, we have

$$\log y_2 \in \{-\sqrt{2}\sigma_\varepsilon,\ +\sqrt{2}\sigma_\varepsilon\},$$

which implies

$$y_2 \in \{0.8358,\ 1.1966\}.$$

Conditional on being in the upper node at $t = 1$, the transition probabilities are asymmetric:

$$\mathbb{P}(y_2 = e^{+\sqrt{2}\sigma_\varepsilon}) = p, \quad \mathbb{P}(y_2 = e^{-\sqrt{2}\sigma_\varepsilon}) = 1 - p,$$

where

$$p = \frac{1}{2}\left(1 + \frac{\sqrt{1}}{\sqrt{2}}\right) = 0.8536.$$

D.2 EGM at the Upper Node

At $t = 2$, consumption equals cash-on-hand:

$$c_2 = z_2 = Ra_2 + y_2.$$

Fixing $a_2 \in \{0, 0.5, 1\}$, EGM computes

$$\mathbb{E}\left[\frac{1}{c_2}\right] = p\frac{1}{Ra_2 + e^{+\sqrt{2}\sigma_\varepsilon}} + (1-p)\frac{1}{Ra_2 + e^{-\sqrt{2}\sigma_\varepsilon}},$$

followed by

$$c_1^{endo} = \frac{1}{\beta R\,\mathbb{E}[1/c_2]}, \qquad z_1^{endo} = c_1^{endo} + a_2.$$

This yields the endogenous pairs:

a_2	0	0.5	1
c_1^{endo}	1.0714	1.5963	2.1190
z_1^{endo}	1.0714	2.0963	3.1190

The feasible cash-on-hand grid is

$$z_1^{grid} = Ra_1 + y_1 \in \{1.1353,\ 1.6553,\ 2.1753\}.$$

Interpolation is performed exactly as in Appendix C and is therefore not repeated here. Interpolating c_1 over (z_1^{endo}, c_1^{endo}) and recovering next-period assets yields:

z_1^{grid}	1.1353	1.6553	2.1753
$c_1(z_1^{grid})$	1.1042	1.3698	1.6353
$a_2(z_1^{grid})$	0.0311	0.2854	0.5400

D.3 Comparison with Appendix C

Appendices C and D use the same asset grids and the same EGM procedure, starting from the same income node at $t = 1$. The difference arises entirely from how future income is approximated.

In Appendix C (binomial shock-tree discretization), next-period income is generated by adding symmetric innovations to the current state. In contrast, the Rouwenhorst method constructs a Markov chain in which next-period income lies on a time-varying grid and is reached through state-dependent transition probabilities.

These differences lead to distinct expectations of future marginal utility and, therefore, to different consumption and saving policies even in this simple two-period example. At the upper income node at $t = 1$, consumption under the Rouwenhorst method is lower than under the shock discretization method for the values of cash-on-hand considered here.