

# Buffer stock saving

## Model description

*Society for Computational Economics Working Group No. 1 on Language and*

*Semantics*

October 20, 2026

### 1. Model theory

We consider a consumer who saves in one risk-free asset and receives labour income with permanent and transitory shocks, following Carroll’s buffer-stock models (1992, 1997).<sup>1</sup> We use the model and notation of Carroll and Shanker, *Theoretical Foundations of Buffer Stock Saving*, version of 5 December 2025. The stationary solution is the limit of finite-horizon solutions as the horizon lengthens; there is no mortality and no bequest motive. Implementations in the QuantEcon Community Library keep this model and calibration fixed.

Bold Latin letters denote levels and upright letters denote functions. The consumer has CRRA utility  $u(c) = c^{1-\gamma}/(1-\gamma)$ , with relative risk aversion  $\gamma > 1$  and discount factor  $\beta > 0$ . Permanent income  $\mathbf{p}_t$  grows by the realised growth factor  $\tilde{\mathcal{G}}_{t+1} = \mathcal{G}\psi_{t+1}$ , where  $\mathcal{G} > 0$  is the deterministic growth factor and  $\psi_{t+1}$  the permanent shock. Labour income is  $y_t = \mathbf{p}_t \boldsymbol{\xi}_t$ , with transitory shock

$$\boldsymbol{\xi} = \begin{cases} 0 & \text{with probability } \wp, \\ \theta/(1-\wp) & \text{with probability } 1-\wp, \end{cases} \quad 0 < \wp < 1, \quad (1)$$

where  $\theta$  is the positive part of the transitory shock and bold  $\boldsymbol{\xi}$ , following the paper, is the full shock including the zero-income event. The positive shocks  $\psi$  and  $\theta$  have mean one and are independent of each other and over time, as is the zero-income event, so  $\mathbb{E}[\boldsymbol{\xi}] = 1$ .

Within a period the shocks are realised; the bank balance  $\mathbf{b}_t = R\mathbf{a}_{t-1}$ , with gross interest factor  $R > 0$ , and labour income from market resources  $\mathbf{m}_t = \mathbf{b}_t + y_t$  (hereafter resources); the consumer then chooses consumption  $\mathbf{c}_t$  and carries end-of-period assets  $\mathbf{a}_t = \mathbf{m}_t - \mathbf{c}_t$  into the next period. Divide resources, consumption and end-of-period assets by permanent income:  $m_t = \mathbf{m}_t/\mathbf{p}_t$ ,  $c_t = \mathbf{c}_t/\mathbf{p}_t$  and  $a_t = \mathbf{a}_t/\mathbf{p}_t$ . The normalised budget equations are

<sup>1</sup>Carroll, C. D. (1992), “The Buffer-Stock Theory of Saving: Some Macroeconomic Evidence”, *Brookings Papers on Economic Activity*, 1992(2), 61–156; Carroll, C. D. (1997), “Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis”, *Quarterly Journal of Economics*, 112(1), 1–55.

$$a_t = m_t - c_t, \quad b_{t+1} = \tilde{\mathcal{R}}_{t+1} a_t, \quad m_{t+1} = b_{t+1} + \xi_{t+1}, \quad \tilde{\mathcal{R}}_{t+1} = \frac{R}{\tilde{\mathcal{G}}_{t+1}}. \quad (2)$$

Here  $b_{t+1}$  is the bank balance relative to next period's permanent income and  $\tilde{\mathcal{R}}_{t+1}$  the growth-normalised return, the interest factor divided by the realised growth factor.

Because income may be zero in any period, debt could leave resources negative and no positive consumption feasible; zero assets is therefore the natural debt limit, and admissible plans satisfy  $0 < c_t \leq m_t$  after every shock history. At the terminal date  $T$ ,  $c_T(m) = m$  and  $v_T(m) = u(m)$ . Before  $T$ , saving nothing risks having nothing to consume after a zero-income draw, and utility tends to minus infinity at zero consumption, so saving nothing is suboptimal and  $0 < c_t(m) < m$ .

Let  $v$  and  $c$  denote the stationary normalised value and consumption functions, with value in levels  $\mathbf{v}(m, \mathbf{p}) = \mathbf{p}^{1-\gamma} v(m/\mathbf{p})$ . The normalised Bellman equation is

$$v(m) = \max_{0 < c < m} \left\{ u(c) + \beta \mathbb{E} \left[ (\tilde{\mathcal{G}}')^{1-\gamma} v \left( \tilde{\mathcal{R}}'(m - c) + \xi' \right) \right] \right\}, \quad m > 0. \quad (3)$$

Primes denote next-period quantities; expectations are over both income shocks. The factor  $(\tilde{\mathcal{G}}')^{1-\gamma}$  converts next period's value, which scales with  $(\mathbf{p}')^{1-\gamma}$ , into units of this period's permanent income. For the reasons just given, the objective tends to minus infinity at both endpoints of the choice set, so the maximum is attained in the interior. The first-order and envelope conditions (the paper establishes that  $v$  is differentiable) give the normalised Euler equation

$$c(m)^{-\gamma} = \beta R \mathbb{E} \left[ (\tilde{\mathcal{G}}')^{-\gamma} c(m')^{-\gamma} \right], \quad m' = \tilde{\mathcal{R}}'(m - c(m)) + \xi'. \quad (4)$$

Precautionary saving becomes more attractive as resources fall, while impatience encourages consumers to spend down high resources. Under the calibration of section 2 these incentives produce a target level of resources  $\hat{m}$ , below which expected resources rise and above which they fall. By the budget equations and  $\mathbb{E}[\xi] = 1$ , expected next-period resources are

$$\mathbb{E}[m_{t+1} \mid m_t] = (m_t - c(m_t)) \frac{R}{\mathcal{G}} \mathbb{E}[\psi^{-1}] + 1, \quad (5)$$

and the target is the fixed point  $\mathbb{E}[m_{t+1} \mid m_t = \hat{m}] = \hat{m}$ ; it is not the mean of the stationary distribution of  $m_t$ . The paper's pseudo-target, at which the level of resources is expected to grow at the rate of permanent income, solves the same equation with  $\mathbb{E}[\psi^{-1}]$  replaced by one and lies below  $\hat{m}$ ; it is not reported here.

## 2. Calibration

We use the paper’s annual calibration. Income growth and risk draw on the Panel Study of Income Dynamics (Carroll, 1992); preferences and the interest rate are conventional.

| Parameter                           | Symbol          | Value |
|-------------------------------------|-----------------|-------|
| Income growth factor                | $\mathcal{G}$   | 1.03  |
| Gross interest factor               | R               | 1.04  |
| Discount factor                     | $\beta$         | 0.96  |
| Relative risk aversion              | $\gamma$        | 2     |
| Probability of zero income          | $\wp$           | 0.005 |
| Standard deviation of $\log \psi$   | $\sigma_\psi$   | 0.1   |
| Standard deviation of $\log \theta$ | $\sigma_\theta$ | 0.1   |

The positive shocks are log-normal with means normalised to one,  $\log \psi \sim N(-\sigma_\psi^2/2, \sigma_\psi^2)$  and  $\log \theta \sim N(-\sigma_\theta^2/2, \sigma_\theta^2)$ .

## 3. Numerical solution

Every implementation retains the calibration but approximates the two shock distributions with finitely many values of its own choosing that preserve their means and the zero-income probability. The paper’s existence results require positive shocks bounded above and away from zero, which finite approximations with strictly positive  $\psi$  and  $\theta$  satisfy and the log-normal laws do not. Under this assumption, weak return impatience and finite value of autarky give a finite value function and a positive limiting consumption function,

$$\wp^{1/\gamma} \frac{\mathbf{P}}{R} < 1, \quad \beta \mathcal{G}^{1-\gamma} \mathbb{E}[\psi^{1-\gamma}] < 1, \quad \mathbf{P} = (R\beta)^{1/\gamma}, \quad (6)$$

where  $\mathbf{P}$  is the absolute patience factor, the growth factor of consumption under perfect foresight. Adding strong growth impatience,  $(\mathbf{P}/\mathcal{G})\mathbb{E}[\psi^{-1}] < 1$ , gives a unique target  $\hat{m}$ . The table lists these and the paper’s four other conditions, each holding when its factor is below one, with values at the calibration from the continuous laws; all seven also hold under a seven-point equiprobable approximation. Each implementation solves the target equation with its own shock approximation and should confirm that the conditions hold under it.

| Condition                | Factor                                                     | Value at the calibration |
|--------------------------|------------------------------------------------------------|--------------------------|
| Finite value of autarky  | $\beta \mathcal{G}^{1-\gamma} \mathbb{E}[\psi^{1-\gamma}]$ | 0.941                    |
| Absolute impatience      | $\mathbf{P}$                                               | 0.999                    |
| Return impatience        | $\mathbf{P}/R$                                             | 0.961                    |
| Weak return impatience   | $\wp^{1/\gamma} \mathbf{P}/R$                              | 0.068                    |
| Finite human wealth      | $\mathcal{G}/R$                                            | 0.990                    |
| Growth impatience        | $\mathbf{P}/\mathcal{G}$                                   | 0.970                    |
| Strong growth impatience | $(\mathbf{P}/\mathcal{G}) \mathbb{E}[\psi^{-1}]$           | 0.980                    |

For weak return impatience the paper’s calibration table prints  $\wp^{1/\gamma} \mathbf{P} \approx 0.071$ ; the factor above also divides by  $R$ .

The consumption function is increasing and concave. Under weak return impatience its marginal propensity to consume (MPC) tends to  $\bar{\kappa}$  as resources tend to zero and to  $\underline{\kappa}$  as resources tend to infinity, where

$$\underline{\kappa} = \max\left\{0, 1 - \frac{\mathbf{P}}{R}\right\}, \quad \bar{\kappa} = 1 - \wp^{1/\gamma} \frac{\mathbf{P}}{R}, \quad \underline{\kappa}m \leq c(m) \leq \bar{\kappa}m. \quad (7)$$

The lower limit is the perfect-foresight MPC; the upper limit falls short of one because a consumer with few resources still saves against the zero-income event.

The solution method is free; value function iteration is one option. The [paper’s own notebook](#) uses the endogenous grid method, linear interpolation and 960 asset grid points, iterating from the terminal consumption rule to convergence. Several figures in that notebook use other parameter values; implementations keep the calibration of section 2.

The stopping tolerance alone does not measure accuracy. Useful checks are a finer shock approximation, the extrapolated consumption function above the top of the asset grid, Euler residuals, monotonicity, concavity and the bounds on  $c(m)$ ; the paper’s notebook re-solves under greater permanent-income risk with a wider asset range, a denser grid and a tighter tolerance. Every implementation reports, in a format of its choosing, consumption  $c(m)$  at  $m \in \{0.25, 0.5, 0.75, 1, 1.25, 1.5, 2, 2.5, 3, 4, 5, 6, 8, 10, 15, 20\}$ , the target  $\hat{m}$  and the MPC at the target. Because a piecewise-linear consumption function has a kink at every grid point, the MPC at the target is the central difference  $[c(\hat{m} + h) - c(\hat{m} - h)]/(2h)$  with  $h = 0.01$ .