A Generalized Method of Moderation for Consumption under Uncertainty

Alan Lujan[1]
Johns Hopkins University
Econ-ARK

Christopher D. Carroll
Johns Hopkins University
Econ-ARK

Monday 16th March, 2026

Abstract

We embed behavioral bounds from buffer stock theory into a neural network architecture for consumption-saving problems, replacing the budget-constrained sigmoid output with a simple output-layer reparameterization: the network learns a ratio $\omega \in [0, 1]$ that a min formula maps into the interval between analytical pessimist and optimist consumption rules. Three constraints hold by construction for any network weights: a positive MPC floor from return impatience, the optimist upper bound from finite human wealth, and the borrowing constraint $c \leq m$. Applied to a buffer stock problem with six sources of uncertainty and a five-dimensional state space (permanent income normalizes out), the moderation architecture achieves Euler equation errors 0.7 orders of magnitude lower than the unconstrained benchmark (-2.75 vs. -2.05 mean $\log_{10}$ on the stylized calibration) while correctly tracking analytical bounds at high wealth and at initializations where the unconstrained architecture lacks structural guidance. A Rao-Blackwellized all-in-one expectation operator integrates out discrete unemployment analytically, eliminating the variance spikes that arise when rare zero-income events appear in both independent shock draws of the product estimator. Training completes in 60 seconds on a single GPU.

Keywords Deep Learning, Dynamic Programming, Consumption-Saving Models, Buffer Stock Theory, Neural Networks, Computational Economics, Precautionary Saving, Euler Equations, Economics Informed Neural Networks

1 Introduction

Heterogeneous-agent macroeconomic models [Aiyagari, 1994, Huggett, 1993, Krusell and Smith, 1998] require solving consumption-saving problems for agents across the entire wealth distribution, including the tails where computational grids become prohibitive. Maliar et al. [2021] showed that neural networks trained on Euler equation residuals can solve these problems in seconds even with five or more state variables, replacing grid-based methods that scale exponentially in dimension. Their architecture constrains consumption to lie within the budget set ($c \in [0, m]$) but imposes no further economic structure. Within the training domain, this suffices; the difficulty is extrapolation. When simulated agents visit wealth levels beyond the training range, the consumption function must still respect precautionary saving, MPC bounds, and the borrowing constraint; without structural guarantees, it may not.

We embed behavioral bounds from buffer stock theory into the network architecture through a simple reparameterization of the output layer. Under CRRA utility with prudence, an optimist who expects mean future

[1]Correspondence to: alujan@jhu.edu

shocks and a pessimist who anticipates permanent unemployment provide analytical upper and lower bounds on consumption [Kimball, 1990, Carroll and Shanker, 2024]. Our moderation approach has the network learn only a ratio $\omega \in [0, 1]$, which a min formula maps into the interval $[\underline{\kappa}m, \min(m, \underline{\kappa}(m + \bar{h}))]$, enforcing three constraints by construction: a positive MPC floor from return impatience, the optimist upper bound from finite human wealth, and the borrowing constraint $c \leq m$. No residual clipping, warmup schedule, or penalty term is needed. On the stylized calibration, the moderation architecture achieves mean Euler equation errors of -2.75 in $\log_{10}$, an improvement of 0.7 orders of magnitude over the unconstrained benchmark (-2.05).

The moderation approach also addresses two extrapolation concerns that arise in iterative heterogeneous-agent computation. First, because the MPC floor $\underline{\kappa}m > 0$ is embedded in the architecture, the degenerate $c \approx 0$ local minimum is unreachable regardless of initialization, so re-training after equilibrium updates cannot fall into the wrong basin. Second, at wealth levels beyond the training range, the min formula forces consumption toward the analytical optimist bound, providing bounded extrapolation error rather than unconstrained drift. Training completes in 60 seconds on a single GPU for a five-dimensional state space (permanent, persistent, and transitory income, stochastic returns, and time-varying preferences after normalization).

Buffer stock theory provides the analytical ingredients for this embedding. Carroll and Shanker [2024] derived MPC bounds and patience conditions for the two-shock model; we extend these to state-dependent AR(1) processes in returns, discounting, and preferences through multiplicative decompositions (Lemma 3.1) that yield closed-form patience factors. Setting persistence and volatility parameters to zero recovers the standard one-dimensional problem. The extension is necessary because the scalar condition $(\mathbf{BR})^{1/\gamma} < 1$ evaluated at mean parameters can hold even when the state-dependent factor $\Phi_R(s_t) \geq 1$ at extreme realizations of the AR(1) processes, giving a false guarantee of non-degenerate solutions. Our new results are the patience factor decomposition (Lemma 3.1), the asymptotic logit linearity of the moderation ratio (Proposition 4.1), the human wealth extension to AR(1) persistent income (Lemma 3.5), and the Rao-Blackwellized AIO estimator ((40)), which applies the classical Rao-Blackwell variance reduction to the specific structure of the AIO product estimator with discrete unemployment; the existence, contraction, and target theorems restate Carroll and Shanker [2024] with scalar conditions replaced by suprema over the bounded state support (see Remark 3.1).

Existing strategies for embedding economic structure into neural network solutions constrain training objectives or equilibrium conditions; we constrain the output space itself. Azinovic et al. [2022] and Han et al. [2021] enforce equilibrium conditions during training but place no restrictions on the policy's functional form. Azinovic et al. [2023] go further by hard-encoding market clearing conditions into the network architecture via Market Clearing Layers, but their constraints apply to equilibrium aggregates, not to individual policy bounds. Pascal [2024] develops bias-corrected Monte Carlo operators that reduce estimator bias but do not bound the policy itself; Ashwin et al. [2025] study whether neural network learning rules converge to rational expectations equilibria, a question orthogonal to the quality of the converged policy we address here; and Chen et al. [2024] use transfer learning from structural models to improve neural network solutions by warming up from a simpler calibration. Our min formula constrains the policy output space directly, ensuring consumption lies between analytical bounds for any network weights, not just at converged parameters.[1]

On the theoretical side, our analysis builds on the income fluctuation problem formulated by Schechtman [1976], the buffer stock interpretation of Carroll [1997], and the general existence results of Ma et al. [2020] for state-dependent returns and income with arbitrary serial correlation. Carroll and Shanker [2024] proved that weak return impatience and finite value of autarky suffice for non-degenerate solutions in the two-shock case, deriving the MPC bounds we generalize here; their construction defines the limiting solution through finite-horizon convergence, while our approach targets the stationary policy directly via Euler residual minimization.

Computationally, we combine stochastic grid sampling [Maliar et al., 2011] with an all-in-one expectation operator (Paragraph) and the Fischer-Burmeister function for the occasionally binding borrowing constraint. A Rao-Blackwellized variant ((40)) integrates out discrete unemployment analytically, eliminating the variance spikes that destabilize the stochastic product estimator. Grid-based methods [Judd, 1998, Carroll, 2006, Maliar and Maliar, 2013] become prohibitive beyond 4-5 dimensions.

The argument proceeds as follows. The Economic Model introduces the six shock sources and normalizes by permanent income to obtain a stationary five-dimensional problem. Because the moderation architecture needs analytical consumption bounds, A Generalized Method of Moderation derives state-dependent MPC

[1] For a survey of deep learning methods in quantitative economics, see Fernandez-Villaverde et al. [2024].

limits through a patience factor decomposition that extends scalar conditions from Carroll and Shanker [2024] to AR(1) processes; the pessimist-optimist interval that results is the constraint set for the neural network. Neural Network Training constructs the loss function and develops the Rao-Blackwellized AIO estimator that eliminates variance spikes from rare unemployment events. Results and Validation compares the two architectures on Euler equation accuracy, initialization robustness, and extrapolation beyond the training domain.

2 The Economic Model

We analyze a consumption-saving problem with uncertainty in preferences, discounting, returns, and income. Consumers accumulate a buffer of wealth relative to permanent income, balancing the precautionary motive against impatience [Ludvigson and Michaelides, 2001]. For the problem to be well-posed, we require a base discount factor $\mathbf{B} \in (0, 1)$ and a gross risk-free rate $\mathbf{R} \in (0, \frac{1}{\mathbf{B}})$, so that $\mathbf{BR} < 1$ and discounted utility streams converge.

The consumer chooses consumption $\mathbf{c}_t$ each period, carrying end-of-period savings into next-period market resources $\mathbf{m}_{t+1}$. Time-varying preference weights $\mu_t > 0$ scale period utility, and a stochastic discount factor $\mathcal{B}_\tau$ governs intertemporal substitution. The optimization problem in level variables is:[2]

$$\begin{aligned} \max_{\{\mathbf{c}_t, \mathbf{m}_{t+1}\}_{t=0}^{\infty}} \mathbb{E}_0 \left[\mu_0 u(\mathbf{c}_0) + \sum_{t=1}^{\infty} \left(\prod_{\tau=1}^{t} \mathcal{B}_\tau \right) \mu_t u\left(\mathbf{c}_t\right) \right] \\ \text{s.t. } \mathbf{m}_{t+1} = \left(\mathbf{m}_t - \mathbf{c}_t\right) R_{t+1} + \mathbf{y}_{t+1}, \\ \mathbf{c}_t \leq \mathbf{m}_t, \end{aligned}$$

where $\mathcal{B}_\tau = \mathbf{B}\beta_\tau$ is the one-period stochastic discount factor (combining base patience $\mathbf{B}$ with a stochastic component β_τ specified below) and the utility function u takes the constant relative risk aversion (CRRA) form:

$$u(c) \equiv \frac{c^{1-\gamma}}{1-\gamma}, \tag{1}$$

where γ denotes the coefficient of relative risk aversion, satisfying $\gamma > 0$ and $\gamma \neq 1$.[3]

In the objective, time-varying preference factors μ_t weight period utility $u(\mathbf{c}_t)$, while cumulative discount factors $\prod_{\tau=1}^{t} \mathcal{B}_\tau$ govern how future utility depreciates. Current utility enters with weight μ_0 and no discounting; each additional period compounds impatience through the product $\mathcal{B}_1\mathcal{B}_2 \cdots \mathcal{B}_t$.

Each period, the consumer splits market resources $\mathbf{m}_t$ between consumption $\mathbf{c}_t$ and saving $\mathbf{m}_t - \mathbf{c}_t$: saving earns stochastic return R_{t+1} and combines with labor income $\mathbf{y}_{t+1}$ to replenish next period's resources. The constraint $\mathbf{c}_t \leq \mathbf{m}_t$ rules out borrowing against future income.

2.1 Stochastic Processes

The model incorporates uncertainty in three fundamental dimensions: income (job loss and earnings fluctuations), discounting (time-varying patience from preference shocks), and returns (asset market risk). Each process follows a lognormal AR(1) specification, which admits closed-form moment generating functions for computing patience factors. The income process parameters are calibrated to empirical estimates (described in Results and Validation); the preference and discount factor processes are specified to demonstrate the method's capacity for high-dimensional state spaces rather than to match specific empirical targets.

[2] Bold font denotes level quantities not yet normalized by permanent income ($\mathbf{m}_t$, $\mathbf{c}_t$, $\mathbf{y}_t$, $\mathbf{p}_t$, $\mathbf{q}_t$); lowercase (m_t, c_t, y_t) denotes ratios to permanent income after normalization. The persistent component $\mathbf{q}_t$ remains bold in the normalized state vector because it is not divided out. The AR(1) states r_t, μ_t, β_t are dimensionless multipliers that never appear in level form.

[3] The CRRA specification exhibits homogeneity of degree $1-\gamma$: scaling consumption by α scales utility by $\alpha^{1-\gamma}$. This enables normalization by permanent income, making the problem stationary. We maintain $\gamma \neq 1$ because log utility ($\gamma = 1$) requires additive rather than multiplicative value function normalization.

2.1.1 The Income Process

The permanent income component $\mathbf{p}_t$ will be normalized out when we reformulate the problem below; the remaining components determine the normalized state. Level income is $\mathbf{y}_t = \mathbf{p}_t \mathbf{q}_t \theta_t$, decomposing into permanent, persistent, and transitory components following the Friedman-Muth permanent income framework [Friedman, 1957, Muth, 1960].[4] Permanent income $\mathbf{p}_t$ follows a random walk, capturing irreversible changes such as job switches or career shifts. Persistent income $\mathbf{q}_t$ follows an AR(1) process that generates slowly mean-reverting fluctuations (bonuses, commissions). Transitory income θ_t is white noise: one-time events (overtime, windfall) that vanish next period.

The permanent income component follows a log random walk with drift:

$$\log \mathbf{p}_{t+1} = \log \mathbf{p}_t + \log \mathbf{G} + \log \psi_{t+1}, \tag{2}$$

where $\mathbf{G} > 0$ is the deterministic growth rate and $\log \psi_{t+1} \sim \mathcal{N}(-\sigma_\psi^2/2, \sigma_\psi^2)$ is a mean-1 lognormal innovation. The permanent income growth factor is $\mathcal{G}_{t+1} \equiv \mathbf{p}_{t+1}/\mathbf{p}_t = \mathbf{G}\psi_{t+1}$ with $\mathbb{E}[\mathcal{G}_{t+1}] = \mathbf{G}$.

The persistent income component follows a log AR(1) process:

$$\log \mathbf{q}_{t+1} = \rho_q \log \mathbf{q}_t + \log \varrho_{t+1}, \tag{3}$$

where $\log \varrho_{t+1} \sim \mathcal{N}(-\sigma_q^2/2, \sigma_q^2)$ is a mean-1 lognormal innovation and $\rho_q \in [0, 1)$ governs persistence.[5]

The transitory shock θ_{t+1} incorporates unemployment risk following Zeldes [1989]:

$$\theta_{t+1} = \begin{cases} 0 & \text{with probability } \wp > 0 \\ \xi_{t+1}/(1-\wp) & \text{with probability } 1-\wp \end{cases} \tag{4}$$

where ξ_{t+1} is a mean-1 lognormal shock with $\log \xi_{t+1} \sim \mathcal{N}(-\sigma_\theta^2/2, \sigma_\theta^2)$. The probability $\wp$ of zero-income events generates a natural borrowing constraint: since consuming everything today risks arriving next period with zero resources and zero income (yielding utility $-\infty$ under CRRA), consumers always save a positive amount. Even a tiny probability of zero income generates a powerful precautionary motive: CRRA utility with $\gamma > 1$ assigns $-\infty$ to zero consumption, so the expected marginal utility of saving is unbounded near the constraint. This unemployment risk enters the weak return impatience condition Proposition 3.2 as the factor $\wp^{1/\gamma}$, ensuring a non-degenerate limiting consumption function even for highly patient consumers.

2.1.2 Discounting and Returns

Discounting, preferences, and returns follow the same lognormal AR(1) structure as persistent income. The time-varying discount factor $\mathcal{B}_{t+1} = \mathbf{B}\beta_{t+1}$ combines deterministic patience $\mathbf{B}$ with a stochastic component β_t (when $\beta_t > 1$ the consumer is temporarily more patient; when $\beta_t < 1$, impatience rises). The marginal utility shock μ_t scales period utility through $\mu_t u(c_t)$, playing a distinct role from β_t: the discount factor governs how the consumer values future utility relative to current utility, while μ_t scales the level of current utility. Stochastic returns $R_{t+1} = \mathbf{R} r_{t+1}$ add portfolio risk to the precautionary saving motive. All three processes follow log AR(1) dynamics:

$$\begin{aligned} \log \beta_{t+1} &= \rho_\beta \log \beta_t + \log \delta_{t+1}, & \log \delta_{t+1} &\sim \mathcal{N}(-\sigma_\beta^2/2, \sigma_\beta^2) \\ \log \mu_{t+1} &= \rho_\mu \log \mu_t + \log \eta_{t+1}, & \log \eta_{t+1} &\sim \mathcal{N}(-\sigma_\mu^2/2, \sigma_\mu^2) \\ \log r_{t+1} &= \rho_r \log r_t + \log \iota_{t+1}, & \log \iota_{t+1} &\sim \mathcal{N}(-\sigma_r^2/2, \sigma_r^2), \end{aligned} \tag{5}$$

where all persistence parameters satisfy $\rho \in [0, 1)$ and all innovations are mean-1 lognormal.[6] When the problem is normalized by permanent income below, returns are adjusted by the growth factor $\mathcal{G}_{t+1}$.

[4] This specification extends the income fluctuation problem originally formulated by Schechtman [1976] and further analyzed by Schechtman and Escudero [1977], building on the generalized treatment by Ma et al. [2020] that allows general serial correlation in income dynamics.

[5] This specification ensures $\mathbb{E}[\varrho_{t+1}] = 1$, so the one-step conditional expectation satisfies $\mathbb{E}_t[\mathbf{q}_{t+1}] = \mathbf{q}_t^{\rho_q}$. The unconditional (ergodic) mean of $\mathbf{q}$ equals $\exp(-\rho_q \sigma_q^2/(2(1-\rho_q^2)))$, which is less than 1 for $\rho_q > 0$. Since the model is solved in normalized variables where the relevant quantity is the conditional expectation of future income given current state, this unconditional mean does not affect the solution.

[6] The mean adjustments $-\sigma^2/2$ in each AR(1) specification ensure that all innovations (ι, η, δ, ϱ, ψ) are mean-1 lognormal variables, and the composite transitory shock θ_{t+1} is mean-1 (though not lognormal due to unemployment risk). This guarantees $\mathbb{E}[\mathcal{G}_{t+1}] = \mathbf{G}$ and simplifies moment calculations for the patience factors.

The vector of log-innovations $\varepsilon_{t+1} = (\log \iota_{t+1}, \log \eta_{t+1}, \log \delta_{t+1}, \log \varrho_{t+1}, \log \psi_{t+1}, \log \xi_{t+1})$ follows a multivariate normal distribution with diagonal covariance matrix (zero correlations), which makes the patience factor decompositions in Lemma 3.1 analytically tractable.[7] The transitory shock θ_{t+1} is constructed from ξ_{t+1} via equation (4): the continuous innovation ξ_{t+1} enters the vector, but the transformed shock θ_{t+1} (which equals zero with probability $\wp$) is excluded since its support includes a mass point at zero incompatible with lognormality. These six independent innovations drive all uncertainty in the model.

2.2 Normalized Bellman Equation

Homogeneity of the CRRA utility function allows normalizing by permanent income $\mathbf{p}_t$, reducing the state from $(\mathbf{m}_t, z_t)$ to (m_t, s_t) where $m_t \equiv \mathbf{m}_t/\mathbf{p}_t$. The normalized value function is:[8]

$$V(m_t, s_t) = \max_{c_t} \left\{ \mu_t u(c_t) + \mathbb{E}_t \left[\mathcal{B}_{t+1} \mathcal{G}_{t+1}^{1-\gamma} V(m_{t+1}, s_{t+1}) \right] \right\}, \tag{6}$$

where $s_t \equiv (r_t, \mu_t, \beta_t, \mathbf{q}_t)$ is the normalized exogenous state vector, obtained from z_t by dropping $\mathbf{p}_t$ (which scales out through normalization), and $\mathcal{R}_{t+1} \equiv R_{t+1}/\mathcal{G}_{t+1} = \mathbf{R} r_{t+1}/(\mathbf{G}\psi_{t+1})$ is the growth-adjusted return factor. The budget constraint, borrowing limit, and income definition are:

$$m_{t+1} = (m_t - c_t)\mathcal{R}_{t+1} + y_{t+1}, \tag{7}$$

$$c_t \leq m_t, \tag{8}$$

$$y_{t+1} = \mathbf{q}_{t+1}\theta_{t+1}, \tag{9}$$

where y_{t+1} is normalized income and $\mathcal{R}_{t+1}$ is the growth-adjusted return. The full normalized state (s_t, m_t) combines exogenous conditions with the endogenous wealth variable.

2.2.1 First-Order Conditions

Standard first-order and envelope conditions yield the Euler equation. With Lagrange multiplier $\Lambda_t \geq 0$ on the borrowing constraint $c_t \leq m_t$, the first-order condition is

$$\mu_t u'(c_t) = \mathbb{E}_t \left[\mathcal{B}_{t+1} \mathcal{R}_{t+1} \mathcal{G}_{t+1}^{1-\gamma} V'(m_{t+1}, s_{t+1}) \right] + \Lambda_t, \tag{10}$$

and the envelope theorem gives $V'(m_t, s_t) = \mu_t u'(c_t)$. Define the marginal utility growth factor $\mathcal{M}_{t+1} \equiv \mu_{t+1}/\mu_t$. Normalizing the Lagrange multiplier as $\lambda_t \equiv \Lambda_t/(\mu_t u'(c_t))$ and substituting, the Euler equation becomes

$$1 = \mathbb{E}_t \left[\mathcal{B}_{t+1} \mathcal{M}_{t+1} \mathcal{R}_{t+1} \mathcal{G}_{t+1}^{1-\gamma} \frac{u'(c_{t+1})}{u'(c_t)} \right] + \lambda_t. \tag{11}$$

When the constraint is slack ($c_t < m_t$), $\lambda_t = 0$ and the expected discounted marginal utility ratio equals unity. When the constraint binds ($c_t = m_t$), $\lambda_t > 0$ creates a wedge.

3 A Generalized Method of Moderation

We constrain the neural network's output space by embedding behavioral bounds from buffer stock theory directly into the architecture, guaranteeing that consumption lies between analytical pessimist and optimist rules for any network weights. The architectural formula appears in (32); the patience conditions derived here make that formula computable. We proceed in three steps: first establishing the analytical bounds from perfect-foresight analysis; then deriving patience conditions that ensure solution existence; and finally constructing the state-dependent linear rules between which the network interpolates.

[7] Independence enables the patience factor decompositions in Lemma 3.1 to factor into products of terms involving individual innovations. Extending to correlated shocks (e.g., income and returns falling together in recessions) would require numerical integration to compute patience factors, replacing the closed-form expressions with simulation-based estimates. The unconstrained (MMW) approach handles correlated shocks without modification since Monte Carlo integration draws from the joint distribution regardless of correlation structure. The independence assumption is standard in the calibration literature [Abowd and Card, 1989, MaCurdy, 1982].

[8] Level and normalized value functions are related by $\mathbf{V}(\mathbf{m}_t, z_t) = \mathbf{p}_t^{1-\gamma} V(m_t, s_t)$. Permanent income growth enters the continuation value through $\mathcal{G}_{t+1}^{1-\gamma}$.

3.1 Behavioral Bounds from Economic Theory

Without uncertainty, the optimal consumption function has a simple analytical solution. Buffer-stock theory establishes that this perfect-foresight solution provides bounds on behavior under uncertainty. A consumer who correctly perceives uninsurable income risk will consume less than a consumer with the same expected income but no perceived income uncertainty. This precautionary saving behavior means the realistic consumer's consumption must lie strictly below the perfect-foresight optimist's consumption. Conversely, the realistic consumer will consume strictly above the level chosen by a perfect-foresight pessimist who believes with certainty that future income will always be the worst possible outcome.

The optimist consumes at the asymptotic MPC that a wealthy buffer-stock agent would approach, using the limiting marginal propensity $\underline{\kappa}(s_t)$ (which accounts for innovation risk through Jensen's inequality) applied to market resources plus expected future income. The pessimist anticipates worst-case outcomes (permanent unemployment with transitory income at its minimum) and follows a more conservative but still linear rule. Both rules have closed-form linear expressions sharing a common slope (derived below as the limiting MPC), differing only in their intercepts reflecting different beliefs about future income. A consumer who correctly perceives uncertainty and exhibits prudence through positive precautionary saving selects consumption strictly between these bounds. The precise functional forms $\bar{c}(m_t, s_t)$ and $\underline{c}(m_t, s_t)$ are derived below after introducing the patience factors that ensure solution existence.

Proposition 3.1 (Consumption Bounds)**.** *Under CRRA utility with $\gamma > 1$ (implying prudence, $u''' > 0$), positive unemployment probability $\wp > 0$, and bounded lognormal innovations, consumption satisfies:*

$$\underline{c}(m_t, s_t) < c^*(m_t, s_t) < \bar{c}(m_t, s_t) \tag{12}$$

for all interior wealth levels $m_t > 0$, where $\underline{c}$ is pessimist consumption, c^* is optimal consumption, and $\bar{c}$ is optimist consumption.[9]

Given the bounds (12) exist, we can define the moderation ratio abstractly. The explicit functional forms for $\bar{c}$ and $\underline{c}$ are derived in Theorem 3.6 after introducing the patience factors that determine the limiting MPC.

Definition 3.1 (Moderation Ratio)**.** *Given optimist consumption $\bar{c}$ and pessimist consumption $\underline{c}$ satisfying the bounds in Proposition 3.1, the moderation ratio $\omega : \mathbb{R}_+ \times \mathcal{S} \to (0, 1)$, where $\mathcal{S} \equiv \{(r, \mu, \beta, \mathbf{q}) \in \mathbb{R}^4_{++}\}$ denotes the exogenous state space, is defined implicitly by:*

$$c^*(m_t, s_t) = \underline{c}(m_t, s_t) + \omega(m_t, s_t) \cdot [\bar{c}(m_t, s_t) - \underline{c}(m_t, s_t)] \tag{13}$$

where $s_t \in \mathcal{S}$ is the exogenous state vector. The ratio satisfies $0 < \omega(m_t, s_t) < 1$ strictly for all feasible wealth levels.

3.2 Patience Factors and Solution Existence

The conditions below generalize Carroll and Shanker [2024] from their scalar absolute patience factor $\Phi \equiv (\mathbf{RB})^{1/\gamma}$ to state-dependent functions $\Phi(s_t)$ that account for stochastic returns, discounting, and preferences. The logical structure of the existence and target-wealth results carries over unchanged: WRIC and FVAC suffice for a non-degenerate solution, GIC guarantees a target wealth ratio, and the implication chains link the conditions in the same hierarchy. The new element is the patience factor decomposition (Lemma 3.1), which factorizes each state-dependent patience factor into a product of current-state terms (from AR(1) persistence) and time-invariant innovation terms (from independent lognormal shocks). This multiplicative structure makes the supremum conditions computable in closed form, replacing the scalar comparisons of Carroll and Shanker [2024] with state-dependent conditions that can be verified analytically.[10]

[9] The upper bound $c^* < \bar{c}$ follows from Kimball [1990]: prudence implies that uncertainty raises the marginal value of saving, reducing consumption below the level chosen by an agent with the same expected income but no uncertainty. The lower bound $\underline{c} < c^*$ holds because the pessimist assumes permanent unemployment ($\wp = 1$) while the actual unemployment probability satisfies $\wp < 1$, so future expected income strictly exceeds zero. The strict inequalities hold at interior wealth levels; at $m = 0$, both c^* and $\underline{c}$ equal zero under the natural borrowing constraint. See Carroll and Shanker [2024] for the formal argument in the two-shock case; the restatement to state-dependent patience factors follows by replacing scalar conditions with their supremum counterparts (see Remark 3.1).

[10] Carroll and Shanker [2024] provide diagrams illustrating how parameter regions map to different combinations of binding conditions. Their theoretical construction defines the limiting solution as the pointwise limit of finite-horizon

Two economic forces must be properly balanced for the problem to be well-posed. If the consumer is too patient relative to income growth, wealth accumulates without bound (violating growth impatience). If the consumer is too patient relative to returns, consumption is deferred indefinitely (violating return impatience). Existence for progressively general settings has been established by Rincón-Zapatero and Rodrı

guez-Palmero [2003] (unbounded returns), Ma et al. [2022] (stochastic discounting), and Alvarez and Stokey [1998] (permanent income growth).[11]

When these conditions hold, consumers maintain a "buffer stock" of wealth near a target ratio where precautionary and impatience motives exactly balance. We consolidate all sources of uncertainty (including stochastic returns, discounting, and preferences absent from the canonical two-shock model) into a single stochastic kernel.

Definition 3.2 (Stochastic Kernel)**.** *The stochastic kernel consolidates all sources of uncertainty in the model:*

$$\mathcal{K}_{t+1} \equiv \mathcal{B}_{t+1}\mathcal{M}_{t+1}\mathcal{R}_{t+1}\mathcal{G}_{t+1}^{1-\gamma}, \tag{14}$$

where $\mathcal{B}_{t+1}$ is the stochastic discount factor, $\mathcal{M}_{t+1} \equiv \mu_{t+1}/\mu_t$ is the marginal utility growth factor, $\mathcal{R}_{t+1} \equiv R_{t+1}/\mathcal{G}_{t+1}$ is the growth-adjusted return, and $\mathcal{G}_{t+1} = \mathbf{G}\psi_{t+1}$ is the gross permanent-income growth factor.[12] *Using this notation, the Euler equation (11) from First-Order Conditions can be written compactly as*

$$1 = \mathbb{E}_t\left[\mathcal{K}_{t+1}\frac{u'(c_{t+1})}{u'(c_t)}\right] + \lambda_t, \tag{15}$$

where $\lambda_t \geq 0$ is the normalized Lagrange multiplier on the borrowing constraint.

Proposition 3.2 (Patience conditions)**.** *The following conditions generalize Carroll and Shanker [2024] from scalar patience factors to state-dependent functions $\Phi(s_t)$, with $\Phi^* \equiv \sup_{s_t}\Phi(s_t)$ denoting the supremum over the bounded state support. Solution existence requires WRIC + FVAC; the moderation architecture further requires RIC ($\underline{\kappa} > 0$) and FHWC ($\bar{h} < \infty$). Under FHWC, the conditions chain: strong GIC $\Rightarrow$ log GIC $\Rightarrow$ GIC $\Rightarrow$ FVAC $\Rightarrow$ RIC. RIC and GIC cannot fail simultaneously without violating FVAC [Carroll and Shanker, 2024].*

The return patience factor, which enters the moderation bounds through $\underline{\kappa}(s_t) = 1 - \Phi_R(s_t)$, is:

$$\Phi_R(s_t) \equiv \mathbb{E}_t\left[R_{t+1}^{(1-\gamma)/\gamma}(\mathcal{B}_{t+1}\mathcal{M}_{t+1})^{1/\gamma}\,\middle|\, s_t\right]. \tag{16}$$

Lemma 3.1 (Patience factor decomposition)**.** *Because AR(1) processes factor innovations from state values multiplicatively, each patience factor decomposes into a product of time-varying state terms and time-invariant innovation terms, making the supremum conditions computable in closed form.*

To derive the decomposition, consider the return patience factor (16). The AR(1) processes imply $R_{t+1} = \mathbf{R}r_t^{\rho_r}\iota_{t+1}$ (where ι_{t+1} is the return innovation), $\mathcal{B}_{t+1} = \mathbf{B}\beta_t^{\rho_\beta}\delta_{t+1}$, and $\mathcal{M}_{t+1} = \mu_t^{\rho_\mu - 1}\eta_{t+1}$. Substituting into the conditional expectation and using independence of innovations:

$$\begin{aligned}\Phi_R(s_t) &= \mathbb{E}_t\left[R_{t+1}^{(1-\gamma)/\gamma}(\mathcal{B}_{t+1}\mathcal{M}_{t+1})^{1/\gamma}\right] \\ &= (\mathbf{R}r_t^{\rho_r})^{(1-\gamma)/\gamma}(\mathbf{B}\beta_t^{\rho_\beta})^{1/\gamma}\mu_t^{(\rho_\mu-1)/\gamma}\cdot\mathbb{E}[\iota^{(1-\gamma)/\gamma}]\mathbb{E}[\delta^{1/\gamma}]\mathbb{E}[\eta^{1/\gamma}].\end{aligned} \tag{17}$$

problems; our neural network approach directly approximates the stationary infinite-horizon policy by minimizing Euler residuals. The patience conditions ensure both constructions yield the same non-degenerate limit.

[11] See also Martins-da Rocha and Vailakis [2010] (local contractions), Jaśkiewicz and Nowak [2011] (Markov decision processes), Li and Stachurski [2014] (CRRA computation), and Chamberlain and Wilson [2000] (stochastic interest rates).

[12] The factors $\mathcal{B}_{t+1}$ and $\mathcal{M}_{t+1}$ enter the Euler equation symmetrically as a product $\mathcal{B}_{t+1}\mathcal{M}_{t+1}$. When the persistence parameters coincide ($\rho_\mu = \rho_\beta$), these two AR(1) processes could be combined into a single "effective discount" process, reducing the state space from five to four dimensions. We maintain the richer specification to demonstrate neural network scalability and to preserve the distinct economic interpretations of time preference shocks (β_t) versus marginal utility shocks (μ_t, capturing taste or health fluctuations).

Table 1: *

Patience conditions [Carroll and Shanker, 2024]

Condition	Formula	Guarantees
FVAC	$\mathbf{BG}^{1-\gamma}\mathbb{E}[\beta_{t+1}(\mu_{t+1}/\mu_t)\psi_{t+1}^{1-\gamma}] < 1$	Bounded lifetime utility
AIC	$\Phi_A^* < 1$, where $\Phi_A(s_t) \equiv \mathbb{E}_t[(\mathcal{B}_{t+1}\mathcal{M}_{t+1}R_{t+1})^{1/\gamma}\|s_t]$	Consumption declines absent precaution
RIC	$\Phi_R^* < 1$ (see (16))	Positive MPC floor $\underline{\kappa} > 0$; moderation architecture
WRIC	$\wp^{1/\gamma}\Phi_R^* < 1$	Non-degenerate solution existence
FHWC	$\mathbf{G} < \mathbf{R}$	Finite human wealth $\bar{h} < \infty$; moderation architecture
GIC	$\Phi_G^* < 1$, where $\Phi_G(s_t) \equiv \mathbb{E}_t[\mathcal{K}_{t+1}^{1/\gamma}\|s_t]$	Pseudo-target $\hat{m}$; aggregate balanced growth
Log GIC	$\log\Phi_G^* < -\sigma_\psi^2/2$	Invariant distributions [Szeidl, 2013]
Strong GIC	$\mathbb{E}_{s,\psi}[\Phi_G(s)/\psi] < 1$	Individual target $\check{m}$

The first line substitutes the AR(1) transition equations; the second line separates known state variables from independent innovations and factors the expectation.

The return patience factor decomposes as:

$$\Phi_R(s_t) = \phi_R \cdot (\mathbf{R}r_t^{\rho_r})^{(1-\gamma)/\gamma} \cdot (\mathbf{B}\beta_t^{\rho_\beta})^{1/\gamma} \cdot \mu_t^{(\rho_\mu-1)/\gamma}, \tag{18}$$

where the innovation component consolidates the contribution from independent shocks:

$$\phi_R \equiv \mathbb{E}[\iota_{t+1}^{(1-\gamma)/\gamma}] \cdot \mathbb{E}[\delta_{t+1}^{1/\gamma}] \cdot \mathbb{E}[\eta_{t+1}^{1/\gamma}]. \tag{19}$$

Applying the lognormal moment generating function,[13] *the innovation component evaluates to:*

$$\phi_R = \exp\left[\frac{(1-\gamma)(1-2\gamma)}{2\gamma^2}\sigma_r^2 + \frac{1-\gamma}{2\gamma^2}\sigma_\beta^2 + \frac{1-\gamma}{2\gamma^2}\sigma_\mu^2\right]. \tag{20}$$

Permanent income innovation ψ_{t+1} does not appear in ϕ_R because the return patience factor governs the normalized problem where permanent income cancels: the Euler kernel $\mathcal{K}_{t+1}$ involves $R_{t+1}/\mathcal{G}_{t+1}$, and the ψ_{t+1} component of $\mathcal{G}_{t+1}$ enters only the growth patience factor ϕ_G through $\mathbb{E}[\psi^{-1}]$.

The growth patience factor decomposes as:

$$\Phi_G(s_t) = \phi_G \cdot (\mathbf{B}\beta_t^{\rho_\beta})^{1/\gamma} \cdot \mu_t^{(\rho_\mu-1)/\gamma} \cdot (\mathbf{R}r_t^{\rho_r})^{1/\gamma} \cdot \mathbf{G}^{-1}, \tag{21}$$

where the innovation component is:

$$\phi_G \equiv \mathbb{E}[\iota_{t+1}^{1/\gamma}] \cdot \mathbb{E}[\delta_{t+1}^{1/\gamma}] \cdot \mathbb{E}[\eta_{t+1}^{1/\gamma}] \cdot \mathbb{E}[\psi_{t+1}^{-1}]. \tag{22}$$

The lognormal MGF with $a = -1$ gives $\mathbb{E}[\psi^{-1}] = \exp[\sigma_\psi^2]$. Substituting into the product yields:

$$\phi_G = \exp\left[\frac{1-\gamma}{2\gamma^2}(\sigma_r^2 + \sigma_\beta^2 + \sigma_\mu^2) + \sigma_\psi^2\right]. \tag{23}$$

[13] For mean-1 lognormal variables with $\log x \sim \mathcal{N}(-\sigma^2/2, \sigma^2)$, we have $\mathbb{E}[x^a] = \exp(a \cdot (-\sigma^2/2) + a^2\sigma^2/2) = \exp[a(a-1)\sigma^2/2]$.

When all persistence parameters vanish ($\rho_r = \rho_\beta = \rho_\mu = 0$), the state-dependent factors collapse to constants and Φ_R reduces to $\mathbf{R}^{(1-\gamma)/\gamma}\mathbf{B}^{1/\gamma} \cdot \phi_R$, recovering the scalar absolute patience factor of Carroll and Shanker [2024].

This decomposition separates time-varying state dependence (captured by $r_t^{\rho_r}$, $\beta_t^{\rho_\beta}$, $\mu_t^{\rho_\mu}$) from time-invariant innovation effects (captured by ϕ_R and ϕ_G). Innovations ι, δ, η enter through R_{t+1}, $\mathcal{B}_{t+1}$, $\mathcal{M}_{t+1}$ respectively; permanent income innovation ψ_{t+1} enters through $\mathcal{G}_{t+1}^{-\gamma}$. To see this, substitute $\mathcal{R}_{t+1} = R_{t+1}/\mathcal{G}_{t+1}$ into the kernel definition (14): $\mathcal{K}_{t+1} = \mathcal{B}_{t+1}\mathcal{M}_{t+1}\mathcal{R}_{t+1}\mathcal{G}_{t+1}^{1-\gamma} = \mathcal{B}_{t+1}\mathcal{M}_{t+1}R_{t+1}\mathcal{G}_{t+1}^{-\gamma}$. The persistent income innovation ϱ_{t+1} and transitory shock θ_{t+1} affect future consumption through the budget constraint but do not appear directly in the stochastic kernel.

Computationally, the multiplicative structure is convenient: the innovation components ϕ_R and ϕ_G are computed once from parameters, while the state-dependent terms are evaluated at each sampled state during neural network training.

One complication arises from unbounded support: lognormal AR(1) processes have support on $(0, \infty)$, so the literal suprema $\Phi_R^* \equiv \sup_s \Phi_R(s)$ and $\Phi_G^* \equiv \sup_s \Phi_G(s)$ are infinite. The theoretical existence results (Theorem 3.2) therefore rely on the bounded-support assumption stated in the theorem: the continuous innovations have strictly positive, bounded lognormal distributions (equivalently, the log-innovations have bounded support).[14] In practice, truncating the lognormal distributions at $\pm 4\sigma$ of the log produces a compact state space over which the patience conditions hold uniformly; for the calibrations in Table 2, Φ_R and Φ_G remain well below unity throughout this truncated support.

Remark 3.1 (Carroll-Shanker restatements)**.** *The results below restate Carroll and Shanker [2024] in our state-dependent notation. Proofs carry over by replacing scalar patience factors (P/R, P/G) with suprema (Φ_R^*, Φ_G^*) over the bounded state support; the multiplicative AR(1) structure (Lemma 3.1) makes these computable. See Carroll and Shanker [2024] and Boyd [1990] for the weighted-norm contraction arguments; Durán [2003] for the stochastic discounting extension.*[15]

Theorem 3.2 (Existence and contraction)**.** *(CS Theorems 1–2.) Under WRIC Proposition 3.2 and FVAC Proposition 3.2, the normalized problem (6) admits a unique non-degenerate solution $c^*(m, s)$ that is C^2, strictly increasing, and strictly concave in m [Carroll and Kimball, 1996]. The MPC-bounded Bellman operator is*

$$\mathcal{T}V(m, s) = \max_{c \in [\underline{\kappa}(s)m,\, \bar{\kappa}(s)m]} \left\{ \mu\, u(c) + \mathbb{E}_s \left[\mathcal{B}' \mathcal{G}'^{1-\gamma} V(m', s') \right] \right\}, \tag{24}$$

where primes denote next-period values drawn from the transition kernel given (m, s), and μ is the current-state marginal utility weight (a component of s). The operator $\mathcal{T}$ is a contraction in the χ-weighted sup norm ($\chi(m) = \zeta + m^{1-\gamma}$, where $\zeta > 0$ is a regularization constant) with modulus $\bar{\delta} < 1$ from FVAC.

Definition 3.3 (State-Dependent)**.** ***MPC Bounds***
The MPC bounds generalize the scalar $\underline{\kappa} = 1 - \mathrm{P}/R$ of Carroll and Shanker [2024] to state-dependent functions of the AR(1) processes:

$$\begin{aligned} \underline{\kappa}(s_t) &\equiv 1 - \Phi_R(s_t), \\ \bar{\kappa}(s_t) &\equiv 1 - \wp^{1/\gamma}\Phi_R(s_t), \end{aligned} \tag{25}$$

where $\Phi_R(s_t)$ is the return patience factor (18). Under RIC, $\underline{\kappa}(s_t) \in (0, 1)$ throughout the ergodic distribution.

Lemma 3.3 (Limiting)**.** ***MPC bounds***
(CS Lemma 1.) Under WRIC, $\underline{\kappa}(s_t) < c'(m, s_t) < \bar{\kappa}(s_t)$, with $c' \to \underline{\kappa}$ as $m \to \infty$ and $c' \to \bar{\kappa}$ as $m \to 0^+$. These arise as infinite-horizon limits of finite-horizon MPC sequences converging geometrically at rates $\Phi_R(s_t)$ and $\wp^{1/\gamma}\Phi_R(s_t)$.

[14] Standard practice in computational economics truncates lognormal shocks at finite bounds for discretization [Tauchen, 1986]. The distinction between bounded and unbounded support matters for the existence theorems but not for the neural network training, which samples from the truncated distributions.

[15] The continuous innovations (ψ, ξ, ι, η, δ, ϱ) have strictly positive lognormal support. The discrete unemployment event $\theta = 0$ is handled separately: WRIC accounts for zero-income risk through $\wp^{1/\gamma}$, and the Rao-Blackwellized estimator integrates out the discrete component analytically.

For the neural network, $\Phi_R(s_t)$ is computed at each sampled state to obtain $\underline{\kappa}(s_t)$, which enters the moderation bounds. Both patience conditions are verified ex-ante from parameters; if either fails, the model may lack a unique stationary solution.

3.3 Target Wealth and Buffer Stock Stability

Theorem 3.4 (Buffer stock targets)**.** *(CS Theorems 3–4.) Under WRIC, FVAC, and the appropriate growth impatience condition:*

- ***Strong GIC*** $\Rightarrow$ *unique stable individual target* $\check{m}$*:* $\mathbb{E}_t[m_{t+1}] = m_t$ *at* $m_t = \check{m}$*.*
- ***GIC*** $\Rightarrow$ *unique pseudo-target* $\hat{m}$*:* $\mathbb{E}_t[\psi_{t+1} m_{t+1}] = m_t$ *at* $m_t = \hat{m}$*, enabling aggregate balanced growth.*
- ***Log GIC*** $\Rightarrow$ *unique log-target* $\tilde{m}$*:* $\mathbb{E}_t[\log \mathbf{m}_{t+1}] = \log(\mathbf{G}\mathbf{p}_t\tilde{m})$ *at* $m_t = \tilde{m}$*.*

When all three targets exist, the log-target lies below the pseudo-target, which lies below the individual target: $\tilde{m} < \hat{m} < \check{m}$*.*

The ordering follows from Jensen's inequality. At $\hat{m}$, the definition $\mathbb{E}[\psi_{t+1} m_{t+1}] = \hat{m}$ and $\mathbb{E}[\log X] < \log \mathbb{E}[X]$ give $\mathbb{E}[\log(\psi_{t+1} m_{t+1})] < \log \hat{m}$, so the log-level drift is negative at $\hat{m}$. Since this drift is strictly decreasing, its zero $\tilde{m}$ lies below $\hat{m}$. The inequality $\hat{m} < \check{m}$ follows because high permanent income growth ψ_{t+1} raises next-period normalized resources m_{t+1}, so $\mathbb{E}[\psi m_{t+1}] > \mathbb{E}[\psi]\mathbb{E}[m_{t+1}]$ and the permanent-income-weighted stationarity condition binds at lower wealth than the unweighted condition.

When GIC holds but strong GIC fails, $\hat{m}$ exists but $\check{m}$ does not. The moderation method applies without modification since solution existence requires only WRIC and FVAC (the moderation architecture also requires RIC and FHWC, which hold across all calibrations).

3.4 Aggregate Invariance

A continuum of ex ante identical buffer-stock households generates cross-sectional distributions whose stationarity depends on the growth impatience hierarchy. Szeidl [2013] proved that invariant distributions of m, c, and a exist under log growth impatience Proposition 3.2. Harmenberg [2021] showed that under the weaker growth impatience condition Proposition 3.2, permanent-income-weighted invariant distributions exist, and aggregate consumption grows at rate $\mathbf{G}$ along the balanced growth path. These results extend to our state-dependent setting: the joint ergodic distribution of (m, s) inherits stationarity from the marginal ergodic distributions of the AR(1) processes combined with the buffer stock stability of m.

Remark 3.2 (Moderation ratio at target wealth)**.** *At the individual target* $\check{m}$ *(when it exists), the moderation ratio satisfies* $\omega(\check{m}, s_t) \in (0, 1)$ *strictly interior, reflecting the balance between precautionary and impatience motives. As* $m \to \infty$*,* $\omega \to 1$ *(approaching optimist behavior) since precautionary motives vanish. As* $m \to 0$*,* ω *remains bounded away from both 0 and 1 due to weak return impatience: the fear of zero-income events prevents consumption from collapsing to the pessimist bound, while prudence prevents consumption from approaching the optimist bound.*

3.5 State-Dependent Linear Rules

Both perfect-foresight consumption functions are linear in market resources with state-dependent coefficients: consumption equals a marginal propensity to consume times total effective wealth (market resources plus human wealth). These closed-form linear rules provide the bounds that constrain the neural network's output space.

The perfect-foresight consumers differ in their beliefs about future income realizations.[16] The consumption bounds require computing the present value of future income that each behavioral type believes they will receive. This quantity, human wealth, determines the intercept of each linear consumption rule.

[16]The optimist behaves as if wealthy enough that uncertainty is irrelevant, consuming at the asymptotic MPC $\underline{\kappa}$ that accounts for innovation risk through Jensen's inequality. Human wealth for the optimist represents the present discounted value of expected future labor income. The pessimist believes future income will be perpetually zero (permanent unemployment with zero transitory income), yielding zero human wealth. Both consumption rules are linear in market resources with analytical solutions.

Definition 3.4 (Human Wealth)**.** *Human wealth is the present discounted value of expected future labor income. The optimist's human wealth $\bar{h}(\mathbf{q}_t)$ depends on the persistent income state and assumes future shocks equal their expected values. The pessimist's human wealth $\underline{h} = 0$ reflects the belief in permanent unemployment.*

The finite human wealth condition (FHWC) requires $\mathbf{G} < \mathbf{R}$, ensuring the optimist's human wealth series converges.

Lemma 3.5 (Human wealth with)**.** ***AR(1) persistent income***
Carroll and Shanker [2024] assume no persistent income component; the series below extends human wealth to AR(1) persistent income, introducing the Jensen correction Υ_n from the conditional lognormal moments. With persistent income following $\log \mathbf{q}_{t+1} = \rho_q \log \mathbf{q}_t + \log \varrho_{t+1}$ where ϱ_{t+1} is mean-1 lognormal with $\log \varrho_{t+1} \sim \mathcal{N}(-\sigma_q^2/2, \sigma_q^2)$, the optimist's human wealth is:

$$\bar{h}(\mathbf{q}_t) = \sum_{n=1}^{\infty} \left(\frac{\mathbf{G}}{\mathbf{R}}\right)^n \mathbb{E}_t[\mathbf{q}_{t+n}] = \sum_{n=1}^{\infty} \left(\frac{\mathbf{G}}{\mathbf{R}}\right)^n \mathbf{q}_t^{\rho_q^n} \cdot \Upsilon_n, \tag{26}$$

where the Jensen correction factor is:

$$\Upsilon_n \equiv \exp\left(\frac{\sigma_q^2}{2} \cdot \left[\frac{1-\rho_q^{2n}}{1-\rho_q^2} - \frac{1-\rho_q^n}{1-\rho_q}\right]\right). \tag{27}$$

The series converges when FHWC Proposition 3.2 holds ($\mathbf{G} < \mathbf{R}$). The ratio $(\mathbf{G}/\mathbf{R})^n$ appears because the level consumption rule is normalized by $\mathbf{p}_t$ while permanent income grows at rate $\mathbf{G}$, so the effective discount rate for future normalized income is $\mathbf{R}/\mathbf{G}$ rather than $\mathbf{R}$ alone. When persistence vanishes ($\rho_q = 0$, iid shocks), $\Upsilon_n = 1$ and $\mathbb{E}_t[\mathbf{q}_{t+n}] = 1$, yielding $\bar{h} = \mathbf{G}/(\mathbf{R}-\mathbf{G})$. The exponent in Υ_n equals $(\sigma_q^2/2)[(1+\rho_q^n)/(1+\rho_q) - 1] \cdot (1-\rho_q^n)/(1-\rho_q)$; since $(1+\rho_q^n)/(1+\rho_q) < 1$ for $n \geq 2$ when $\rho_q \in (0,1)$, the bracket is negative, confirming $\Upsilon_n < 1$.

The Jensen correction arises from the lognormal moment formula. Iterating the AR(1) process yields $\log \mathbf{q}_{t+n} | \mathbf{q}_t \sim \mathcal{N}(\ell_n, \tau_n^2)$, where $\ell_n \equiv \mathbb{E}[\log \mathbf{q}_{t+n} | \mathbf{q}_t]$ is the conditional log-mean at horizon n and $\tau_n^2 \equiv \mathrm{Var}[\log \mathbf{q}_{t+n} | \mathbf{q}_t]$ is the conditional log-variance, both given in closed form by the geometric series $(1-\rho_q^n)/(1-\rho_q) = 1 + \rho_q + \cdots + \rho_q^{n-1}$:

$$\begin{aligned} \ell_n &= \rho_q^n \log \mathbf{q}_t - \frac{\sigma_q^2}{2} \cdot \frac{1-\rho_q^n}{1-\rho_q}, \\ \tau_n^2 &= \sigma_q^2 \cdot \frac{1-\rho_q^{2n}}{1-\rho_q^2}. \end{aligned} \tag{28}$$

Applying the lognormal moment formula $\mathbb{E}[e^X] = e^{\ell + \tau^2/2}$ (for $X \sim \mathcal{N}(\ell, \tau^2)$) gives $\mathbb{E}_t[\mathbf{q}_{t+n}] = \exp(\ell_n + \tau_n^2/2) = \mathbf{q}_t^{\rho_q^n} \cdot \Upsilon_n$, where the factor Υ_n consolidates the variance terms.

The transitory shock θ_t does not appear in the human wealth formula because it is mean-1 white noise with no persistence: $\mathbb{E}_t[\mathbf{q}_{t+n}\theta_{t+n}] = \mathbb{E}_t[\mathbf{q}_{t+n}] \cdot \mathbb{E}[\theta_{t+n}] = \mathbb{E}_t[\mathbf{q}_{t+n}]$. The pessimist's human wealth $\underline{h} = 0$ reflects the belief that all future income will be zero (permanent unemployment), consistent with the worst-case scenario defining the lower consumption bound.

Theorem 3.6 (Perfect-Foresight Consumption Rules)**.** *The CRRA structure yields linear consumption rules as in Carroll and Shanker [2024], but with state-dependent slopes $\underline{\kappa}(s_t) = 1 - \Phi_R(s_t)$ that vary with the exogenous state through the decomposition Lemma 3.1. The optimist consumes according to*

$$\bar{c}(m_t, s_t) \equiv \underline{\kappa}(s_t) \cdot (m_t + \bar{h}(\mathbf{q}_t)) \tag{29}$$

while the pessimist follows

$$\underline{c}(m_t, s_t) \equiv \underline{\kappa}(s_t) \cdot m_t, \tag{30}$$

where $\underline{\kappa}(s_t) = 1 - \Phi_R(s_t)$ *is the state-dependent limiting MPC,* $\bar{h}(\mathbf{q}_t)$ *is optimist human wealth Lemma 3.5, and* $\underline{h} = 0$ *is pessimist human wealth under permanent unemployment. Since* $\underline{h} = 0$*, the pessimist rule simplifies to* $\underline{c} = \underline{\kappa}(s_t)m_t$*.*

Both rules share slope $\underline{\kappa}(s_t)$ but differ in their intercepts, reflecting different beliefs about future income. The unconstrained rule $\bar{c}$ is non-degenerate if and only if both FHWC and RIC hold; without FHWC, human wealth diverges; without RIC, the MPC vanishes [Carroll and Shanker, 2024]. In the constrained problem, the liquidity constraint binds at $m_t = 1$ if and only if growth impatience holds [Carroll and Shanker, 2024].

Under typical parameters ($\rho_q < 1$, $\mathbf{G} < \mathbf{R}$), the human wealth series converges geometrically and is computed numerically. The truncation error from omitting terms $n > N$ is bounded by $Q^* \cdot \Upsilon^*(\mathbf{G}/\mathbf{R})^{N+1}/(1 - \mathbf{G}/\mathbf{R})$, where $\Upsilon^* \equiv \sup_n \Upsilon_n \leq 1$ (since $\Upsilon_n < 1$ for $n \geq 2$ when $\rho_q > 0$) and $Q^* \equiv \sup_{n>N} \mathbf{q}_t^{\rho_q^n}$, which converges to 1 as $N \to \infty$ since $\rho_q^n \to 0$. At mean state values ($\mathbf{q}_t = 1$), $\bar{Q} = 1$ and the bound simplifies.

How many terms suffice depends on the ratio $\mathbf{G}/\mathbf{R}$: for the stylized calibration ($\mathbf{G}/\mathbf{R} = 1.00/1.04 \approx 0.96$), $N = 400$ terms yield a truncation bound below 10^{-5}; for the baseline calibration ($\mathbf{G}/\mathbf{R} = 1.03/1.04 \approx 0.99$), which sits near the FHWC boundary, $N = 1{,}000$ terms are needed to achieve a bound below 10^{-2}. The implementation adjusts N to the calibration.

Higher persistence ρ_q raises human wealth by increasing the weight on current income levels. Higher volatility σ_q lowers human wealth: the mean-adjustment in the AR(1) process ($-\sigma_q^2/2$ per period in log income) dominates the variance correction, so $\Upsilon_n < 1$ for $n \geq 2$ when $\rho_q > 0$, reducing $\mathbb{E}_t[\mathbf{q}_{t+n}]$ relative to the deterministic case.

With unemployment risk ($\wp > 0$), pessimist human wealth $\underline{h} = 0$ and the natural borrowing constraint $\underline{m} = 0$. The consumption function structure simplifies: the pessimist consumes $\underline{c} = \underline{\kappa}(s_t)m_t$ (30), the optimist consumes $\bar{c} = \underline{\kappa}(s_t)(m_t + \bar{h}(\mathbf{q}_t))$ (29), and the moderation formula (13) interpolates between them.

All behavioral bounds have closed-form expressions: $\underline{\kappa}(s_t) = 1 - \Phi_R(s_t)$ from the return patience factor and $\bar{h}(\mathbf{q}_t)$ from the geometric series for AR(1) income Lemma 3.5. Under return impatience, $\underline{\kappa} \in (0, 1)$, and $\bar{c} \geq \underline{c}$ follows from $\bar{h} \geq 0$.

4 Neural Network Training

With analytical bounds in hand, we turn to the neural network that learns ω within those bounds. Both the unconstrained and moderation architectures train in 60 seconds on a single GPU by minimizing Euler equation residuals on stochastically sampled states, avoiding grid-based discretization entirely. Each network outputs two values: a bounded ratio (consumption share $\varsigma_t \in [0, 1]$ for the unconstrained approach, moderation ratio $\omega_t \in [0, 1]$ for the structured approach) and a normalized Lagrange multiplier $\lambda_t \geq 0$ for the borrowing constraint.

4.1 Network Architecture and Sampling

Since the model is stationary, we seek a time-invariant decision rule $\phi_{\text{nn}}(m_t, s_t; \mathbf{W})$ mapping the full state (m_t, s_t) to two outputs. Following the architecture of Maliar et al. [2021], the network ϕ_{nn} is a multilayer perceptron with 3 hidden layers of 32 units each and ReLU activations, containing approximately 3,000 trainable parameters $\mathbf{W}$. The 5-dimensional input combines the endogenous wealth variable m_t with the exogenous state $s_t = (r_t, \mu_t, \beta_t, \mathbf{q}_t)$. Input normalization transforms exogenous states by $\log(x)/\bar{\sigma}_x$ where $\bar{\sigma}_x$ is the ergodic standard deviation, mapping $\pm 1\sigma$ deviations to ± 1; market resources m are normalized linearly to $[-1, 1]$ over the training domain. The 2-dimensional output (31) applies sigmoid activation for the first component (constraining it to $[0, 1]$) and exponential activation for the Lagrange multiplier $\lambda_t > 0$:

$$\phi_{\text{nn}}(m_t, s_t; \mathbf{W}) = \begin{pmatrix} \text{sigm}(f_1(m_t, s_t; \mathbf{W})) \\ \exp(f_2(m_t, s_t; \mathbf{W})) \end{pmatrix}, \tag{31}$$

where $\text{sigm}(x) = 1/(1 + e^{-x})$ is the logistic sigmoid function and f_1, f_2 are the two components of the final linear layer. The exponential activation ensures $\lambda_t > 0$, providing a smooth positive approximation to the Lagrange multiplier; as $f_2 \to -\infty$, $\lambda_t \to 0^+$, so the architecture can represent near-zero multipliers at interior solutions where the borrowing constraint is slack. The Fischer-Burmeister complementarity function accommodates this smooth relaxation without requiring exact $\lambda = 0$.

Training data is drawn from the 5-dimensional state space. The exogenous state variables $r, \mu, \beta, \mathbf{q}$ are sampled from their ergodic distributions, where each lognormal AR(1) process has ergodic variance $\sigma^2/(1-\rho^2)$ for its log. The endogenous state m is sampled uniformly from a finite interval $[\underline{m}, \bar{m}]$ containing the natural borrowing constraint and extending to wealth levels well above the ergodic mean.

The unconstrained approach sets the first output to the consumption share ς_t, yielding $c_t = \varsigma_t \cdot m_t$. This places no structural restrictions on the consumption function beyond $c_t \in [0, m_t]$ implied by the sigmoid activation. Figure 1 shows that the untrained network produces consumption near the 45-degree line, carrying no economic information about precautionary saving or MPC bounds.

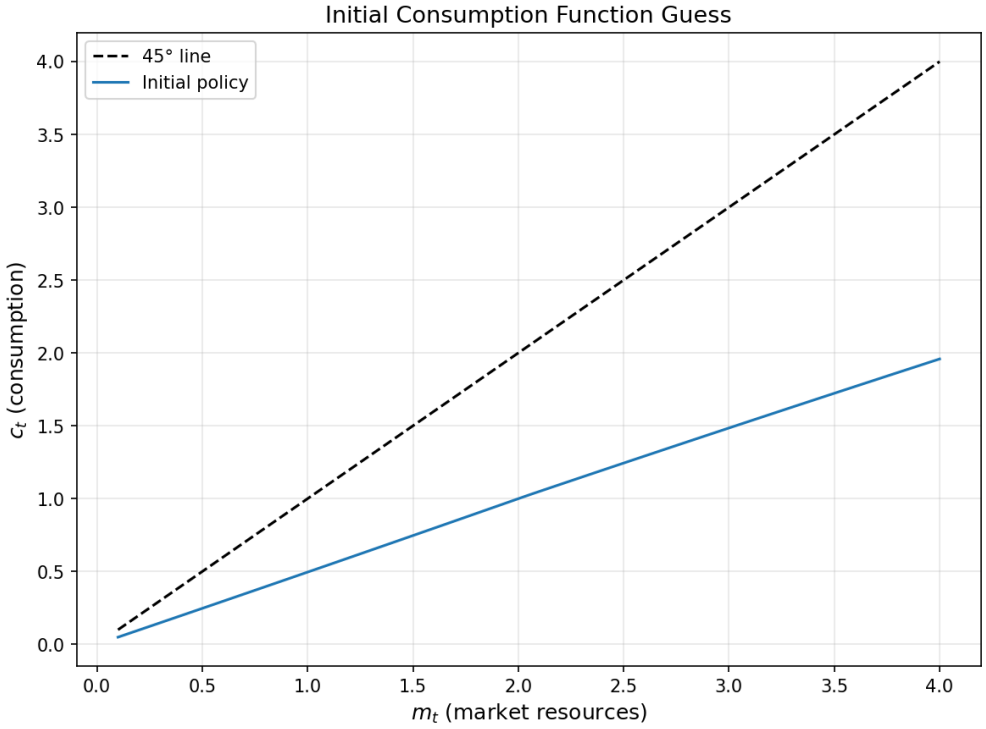


Figure 1: At initialization the network consumes half of market resources ($\varsigma \approx 0.5$) at every wealth level, a flat prior that carries no information about precautionary saving, target wealth, or MPC bounds. At three times the buffer-stock target, this implies saving exactly half of resources, far less than the optimist bound prescribes. The moderation architecture avoids this uninformative prior by construction. Policy evaluated at mean state values.

4.2 Moderation Implementation

We replace the budget-constrained sigmoid transformation $c = \text{sigm}(f) \cdot m$ of Maliar et al. [2021] with a theory-informed reparameterization that enforces all three theoretical bounds on consumption. The network still outputs a sigmoid-bounded ratio $\omega_t \in [0, 1]$, but we map it to the interval $[\underline{\kappa} m_t, \min(m_t, \underline{\kappa}(m_t + \bar{h}))]$ instead of $[0, m_t]$:

$$c_t = \underline{\kappa} m_t + \omega_t \cdot \min\big(m_t(1 - \underline{\kappa}),\ \underline{\kappa}\bar{h}\big) \tag{32}$$

where the limiting MPC $\underline{\kappa}(s_t) = 1 - \Phi_R(s_t)$ and optimist human wealth $\bar{h}(\mathbf{q}_t)$ are analytical functions of the exogenous state, defined in State-Dependent Linear Rules. We suppress state-dependence for readability. Three constraints follow immediately from $\omega \in [0, 1]$:

1. The pessimist floor holds: $c_t \geq \underline{\kappa} m_t$, because $\omega \geq 0$ and $\min(\cdot) \geq 0$.
2. The optimist ceiling holds: $c_t \leq \underline{\kappa}(m_t + \bar{h})$, because $\omega \leq 1$ and $\min(\cdot) \leq \underline{\kappa}\bar{h}$.
3. The borrowing constraint holds: $c_t \leq m_t$, because $\omega \leq 1$ and $\min(\cdot) \leq m_t(1 - \underline{\kappa})$.

The min selects which upper bound binds. Define the cusp point $m^* = \underline{\kappa}\bar{h}/(1-\underline{\kappa})$ where the optimist bound $\underline{\kappa}(m+\bar{h})$ crosses m. Below the cusp, the borrowing constraint is tighter and the feasible gap equals $m(1-\underline{\kappa})$; above the cusp, the optimist bound is tighter and the gap equals $\underline{\kappa}\bar{h}$, a constant independent of wealth. At mean state values in the stylized calibration ($\underline{\kappa} \approx 0.07$, $\bar{h} \approx 25$), the cusp falls at $m^* \approx 1.88$.

This formula has three consequences for training. First, because $c_t \leq m_t$ holds by construction, the savings rate $1 - c_t/m_t$ is non-negative for any ω. The Fischer-Burmeister complementarity condition (Paragraph) never penalizes overconsumption, so the loss reduces to a single-objective Euler equation optimization without competing penalty terms. Second, at the boundary $m = 0$ both terms in the min vanish, giving $c_t = 0$ regardless of ω: zero resources, zero consumption. Third, as $m \to \infty$ the gap converges to $\underline{\kappa}\bar{h}$, so consumption approaches $\underline{\kappa}(m + \bar{h})$ provided $\omega \to 1$. The min formula alone guarantees $c \in [\underline{\kappa}m, \underline{\kappa}(m + \bar{h})]$; the log-odds prior (36) drives $\omega \to 1$ at the rate theory predicts (Proposition 4.1), completing the asymptotic guarantee. The neural network learns only the interior profile; the boundary behavior follows from the architecture and the prior together.

While the min formula provides correct boundary behavior, the Euler equation gradient $\partial R_1/\partial\omega$ scales as $O(\bar{h}/m^2)$ at high wealth: the optimizer cannot distinguish $\omega = 0$ from $\omega = 1$ when m is large relative to $\bar{h}$. Without correction, training produces non-monotonic consumption profiles at high wealth (observed in preliminary experiments without the prior). The following result pins down the asymptotic behavior of ω and motivates a log-odds prior that resolves this gradient vanishing.

Proposition 4.1 (Asymptotic logit linearity)**.** *Under return impatience Proposition 3.2 and finite human wealth Proposition 3.2, the excess caution* $\Delta(m, s) \equiv \bar{c}(m, s) - c^*(m, s)$ *satisfies* $\Delta(m, s) = \Theta(m^{-\gamma})$ *as* $m \to \infty$ *for each state* s.[17]

$$\text{logit}(\omega(m, s)) = \gamma \log(m/m^*(s)) + O(1), \tag{33}$$

where $m^*(s) = \underline{\kappa}(s)\bar{h}/(1 - \underline{\kappa}(s))$ *is the cusp point.*

Proof. Fix the state s throughout; all quantities $\underline{\kappa} = \underline{\kappa}(s)$, $\bar{h} = \bar{h}(s)$, $m^* = m^*(s)$ are treated as constants.

Above the cusp ($m > m^*$), the pessimist-optimist gap $\bar{c} - \underline{c} = \underline{\kappa}\bar{h}$ is independent of m, so $1 - \omega = \Delta/(\underline{\kappa}\bar{h})$. It suffices to show $\Delta(m) = \Theta(m^{-\gamma})$.

Integration identity. The optimist rule gives $\bar{c}' = \underline{\kappa}$, and Lemma 3.3 gives $c'(m) > \underline{\kappa}$ for all finite m, so $\Delta' = \bar{c}' - c' = \underline{\kappa} - c' < 0$ and

$$\Delta(m) = \int_m^\infty [c'(x) - \underline{\kappa}]\, dx. \tag{34}$$

MPC decay rate. As the buffer stock grows, the probability of the borrowing constraint binding within any finite horizon shrinks, and the precautionary saving motive vanishes. Under CRRA utility, the excess MPC $c'(m) - \underline{\kappa}$ decays at rate $\Theta(m^{-\gamma-1})$: the formal argument follows from Carroll [2011] in the one-dimensional case. The extension to multiplicative AR(1) shocks is plausible because the CRRA Euler equation has the same asymptotic structure at high wealth, where all shocks become small relative to savings and the $m^{-\gamma-1}$ rate is determined by γ alone; a formal proof for state-dependent returns would require extending Carroll [2011] to the multidimensional setting. Integrating gives $\Delta(m) = \Theta(m^{-\gamma})$, since $\int_m^\infty x^{-\gamma-1}dx = m^{-\gamma}/\gamma$.

Logit consequence. From $1 - \omega = \Delta/(\underline{\kappa}\bar{h}) = \Theta(m^{-\gamma})$, we have $\omega \to 1$ as $m \to \infty$. Then $\log(1 - \omega) = -\gamma \log m + O(1)$ and $\log \omega = O(m^{-\gamma}) = O(1)$, so

$$\text{logit}(\omega) = \log \omega - \log(1 - \omega) = \gamma \log m + O(1) = \gamma \log(m/m^*) + O(1). \tag{35}$$

The logit of the moderation ratio therefore grows at rate γ in $\log m$, which is (33).

□

[17] The decay rate $\Theta(m^{-\gamma})$ is proved for the one-dimensional case by Carroll [2011]. The extension to state-dependent AR(1) shocks is conjectural: the CRRA Euler equation has the same asymptotic structure at high wealth, but a formal proof requires extending the one-dimensional argument to the multidimensional setting. The log-odds prior (36) is motivated by this rate and performs well empirically across all calibrations. The moderation ratio satisfies:

We encode this asymptotic by transforming the network's raw sigmoid output $\omega_{\text{nn}} \in [0, 1]$ through a log-odds prior with slope γ (the CRRA coefficient):

$$\omega_t = \text{sigm}\big(\text{logit}(\omega_{\text{nn}}) + \gamma \log(m_t/m^*)\big). \tag{36}$$

In odds space, this reads $\text{odds}(\omega_t) = \text{odds}(\omega_{\text{nn}}) \cdot (m_t/m^*)^\gamma$. The exponent γ matches the theoretical convergence rate (33): as $m_t \to \infty$, the prior's contribution $\gamma \log(m_t/m^*)$ dominates and $\omega_t \to 1$, so consumption approaches $\underline{\kappa}(m_t + \bar{h})$ from below at the rate theory predicts. When $m_t \ll m^*$, the prior pushes $\omega_t \to 0$, matching the pessimist rule at low wealth.

Because the prior encodes both the correct asymptotic direction and the correct convergence rate, the network receives $\log(m_t/m^*)$ as its wealth input and learns only an approximately constant residual correction at high wealth, replacing both the linear normalization and the need to discover the full sigmoid profile from scratch.

Both $\underline{\kappa}(s_t)$ and $\bar{h}(\mathbf{q}_t)$ depend on the exogenous state and model parameters, not on the neural network weights, so they need not be updated during training. Computing $\bar{h}$ requires a truncated geometric sum ((26)), with up to 1,000 terms depending on the calibration (see State-Dependent Linear Rules for truncation error analysis), but this is a single batched dot product across all n sample points, adding less than 1% to the per-iteration cost.

Figure 2 shows the untrained moderation consumption function alongside the analytical bounds. Even before training, the architecture places consumption between the pessimist and optimist rules by construction.

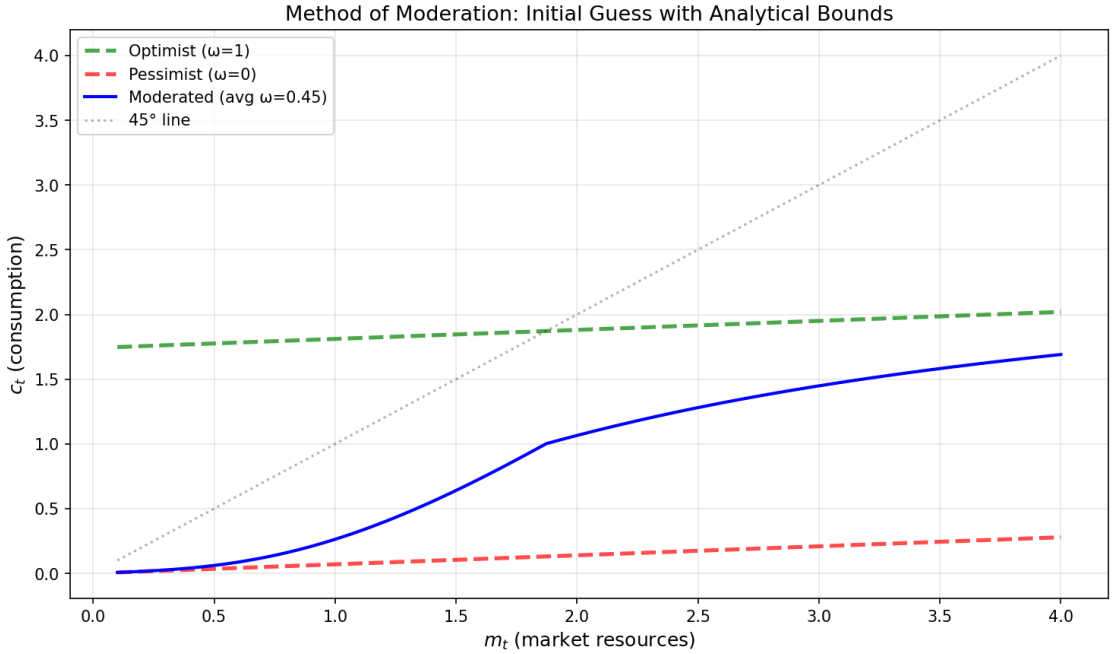


Figure 2: Untrained moderation consumption function (blue solid) with analytical bounds: the optimist (green dashed, $c = \underline{\kappa}(m + \bar{h})$) and pessimist (red dashed, $c = \underline{\kappa}m$). The 45-degree line (black dotted) shows the borrowing constraint $c = m$. At initialization ($\omega \approx 0.5$), the moderated consumption lies midway between bounds. Policy evaluated at mean exogenous state.

The change from Maliar et al. [2021] to (32) with the prior (36) requires only two modifications to the output layer: the log-odds transformation with exponent γ converts the sigmoid output, and the min formula replaces the linear scaling. Both architectures output a pair (ω, λ), share the same hidden layers, and minimize the identical loss function (Loss Function and Constraint Handling).

4.3 Loss Function and Constraint Handling

Following Maliar et al. [2021], we handle the borrowing constraint $c_t \leq m_t$ via the Fischer-Burmeister function [Fischer, 1992], $\Psi_{FB}(a,b) = a + b - \sqrt{a^2 + b^2}$, which satisfies $\Psi_{FB} = 0$ if and only if $a \geq 0$, $b \geq 0$, $ab = 0$. With the consumption share $\varsigma_t = c_t/m_t$ and the normalized Lagrange multiplier $\lambda_t \geq 0$ from First-Order Conditions, the two residuals are:

$$R_1(m_t, s_t, \varepsilon_{t+1}) = \mathcal{K}_{t+1} \frac{u'(c_{t+1})}{u'(c_t)} + \lambda_t - 1 \tag{37}$$

$$R_2(m_t, s_t) = \Psi_{FB}(1 - \varsigma_t, \lambda_t) \tag{38}$$

Residual R_1 is the Euler equation with the multiplier absorbed so that $\mathbb{E}[R_1] = 0$ at the optimum whether or not the borrowing constraint binds; R_2 enforces complementarity. The sigmoid and exponential activations ensure $(1 - \varsigma_t, \lambda_t)$ never reaches the origin, sidestepping the Fischer-Burmeister non-differentiability.

The loss follows the Euler-residual minimization and all-in-one (AIO) expectation operator of Maliar et al. [2021]. For i.i.d. draws $\varepsilon^{(1)}, \varepsilon^{(2)}$, the identity $(\mathbb{E}[R_1])^2 = \mathbb{E}[R_1(\varepsilon^{(1)})R_1(\varepsilon^{(2)})]$ eliminates nested expectations. The loss function becomes:

$$\Xi(\mathbf{W}) = \mathbb{E}_{(m_t, s_t), \varepsilon^{(1)}, \varepsilon^{(2)}} \left[R_1(m_t, s_t, \varepsilon^{(1)}) R_1(m_t, s_t, \varepsilon^{(2)}) + v \cdot R_2(m_t, s_t)^2 \right] \tag{39}$$

with weight $v = 1$. Each iteration draws one state and two independent shock vectors, reducing evaluations from $n \times J$ to $2n$.

The Rao-Blackwellized variant below doubles this to $4n$ by evaluating each residual at both employment states.

4.3.1 Rao-Blackwellized Estimation for Unemployment

Transitory income θ_t follows a mixture distribution: $\theta_t = 0$ with probability $\wp$ (unemployment), and $\theta_t = \xi_t/(1 - \wp)$ with probability $1 - \wp$ (employment), where the lognormal component ξ_t satisfies $\mathbb{E}[\xi] = 1$. The mixture preserves $\mathbb{E}[\theta] = 1$ while introducing the zero-income risk that motivates the pessimist bound $\underline{h} = 0$.

The standard AIO estimator Paragraph draws θ from the full mixture. When an unemployment event occurs ($\theta = 0$), next-period resources consist entirely of savings, the marginal utility ratio $(c_{t+1}/c_t)^{-\gamma}$ is large, and the Euler residual R_1 spikes. When both independent draws happen to include unemployment (probability $\wp^2 \approx 2.5 \times 10^{-5}$), the product $R_1^{(1)} \cdot R_1^{(2)}$ can exceed its typical magnitude by several orders, producing high estimator variance and unstable gradients.

The unemployment probability $\wp$ is a known model parameter, not a random variable to be sampled. We integrate out the discrete component of θ analytically and reserve Monte Carlo for the continuous lognormal component ξ alone. The conditional expectation of R_1 given the exogenous shocks $(r', \mu', \delta', q', \psi')$ decomposes as

$$\mathbb{E}_\theta[R_1 \,|\, r', \mu', \delta', q', \psi'] = \wp \cdot R_1\big|_{\theta=0} \;+\; (1 - \wp) \cdot \mathbb{E}_\xi \Big[R_1\big|_{\theta = \xi/(1-\wp)} \Big] \tag{40}$$

For each AIO shock draw we evaluate the Euler residual at $\theta = 0$ and at $\theta = \xi/(1 - \wp)$ (where ξ is a fresh lognormal draw), then combine with the known weights to form the Rao-Blackwellized residual R_1^{RB}:

$$R_1^{\mathrm{RB}} = \wp \cdot R_1\big|_{\theta=0} \;+\; (1 - \wp) \cdot R_1\big|_{\theta = \xi/(1-\wp)}. \tag{41}$$

The AIO loss (39) becomes $\Xi(\mathbf{W}) = \mathbb{E}[R_1^{\mathrm{RB}}(\varepsilon^{(1)}) \cdot R_1^{\mathrm{RB}}(\varepsilon^{(2)}) + v \cdot R_2^2]$. By the Rao-Blackwell theorem, conditioning on the discrete component θ reduces variance without introducing bias. The unemployment contribution, weighted by $\wp \approx 0.005$, now appears in every sample rather than in rare spikes, so the product $R_1^{\mathrm{RB},(1)} \cdot R_1^{\mathrm{RB},(2)}$ stays bounded throughout training.

This modification doubles the number of residual evaluations per AIO draw (from two to four: two values of θ for each of two independent shock draws). The variance reduction compensates: in our experiments,

Table 2: Calibration parameters

Parameter	Symbol	Baseline	Stylized	Quarterly
Risk aversion	γ	2	2	2
Discount factor	$\mathbf{B}$	0.96	0.90	0.99
Gross return	$\mathbf{R}$	1.04	1.04	1.01
Income growth	$\mathbf{G}$	1.03	1.00	1.0075
Unemployment prob.	$\wp$	0.005	0.005	0.00125
Perm. income s.d.	σ_ψ	0.06	0.02	0.03
Trans. income s.d.	σ_θ	0.20	0.02	0.10
Return s.d.	σ_r	0.08	0.02	0.04
Return persistence	ρ_r	0.20	0.20	0.20
Discount s.d.	σ_β	0.02	0.02	0.01
Discount persistence	ρ_β	0.20	0.20	0.20
Marg. util. s.d.	σ_μ	0.02	0.02	0.01
Marg. util. persistence	ρ_μ	0.20	0.20	0.20
Pers. income s.d.	σ_q	0.04	0.02	0.02
Pers. income persistence	ρ_q	0.90	0.90	0.90
Limiting MPC	$\underline{\kappa}$	0.039	0.070	0.010

the Rao-Blackwellized estimator converges to the same loss in fewer iterations than the stochastic mixture, producing a net speedup despite the per-iteration overhead.

4.4 Training Procedure

Following the training protocol of Maliar et al. [2021], both networks are trained for 50,000 iterations using the Adam optimizer [Kingma and Ba, 2015] with base learning rate $\alpha = 0.001$. Each iteration samples $n = 128$ state points from ergodic distributions and evaluates the all-in-one objective function. Automatic differentiation computes gradients with respect to all network parameters.

Grid-based backward induction requires $20^5 = 3.2$ million evaluations per iteration for a 5-dimensional problem with 20 points per dimension. The neural network evaluates 128 randomly sampled states per iteration instead, achieving tractability in 5 dimensions where grid methods become infeasible, at reduced accuracy (Euler errors of -2 to -3 vs. -4 to -6 for grid methods on lower-dimensional problems). The tradeoff: contraction mappings guarantee geometric convergence, while neural network optimization provides only asymptotic consistency and may converge to local minima. Across all three calibrations and initializations tested, the Rao-Blackwellized AIO estimator did not converge to spurious solutions; we have not characterized all local minima beyond these experiments.

5 Results and Validation

We evaluate both architectures on three calibrations spanning annual, stylized, and quarterly frequencies, each testing a different region of the patience condition parameter space. The stylized calibration ($\mathbf{B} = 0.9$) sits far from the patience boundaries and produces sharp buffer-stock dynamics; the baseline ($\mathbf{BR} \approx 1$) sits near the growth impatience boundary where both methods face challenges; the quarterly calibration tests whether the approach survives re-parameterization. This section documents the accuracy comparison (Table 3), the extrapolation behavior beyond the training domain (Figure 9), and the initialization robustness that matters for iterative equilibrium computation. Table 2 reports all parameter values.

5.1 Parameter Calibration

The baseline calibration uses annual frequency with deterministic discount factor $\mathbf{B} = 0.96$ (4.2% time preference rate), deterministic risk-free rate $\mathbf{R} = 1.04$ (4% annual return), and coefficient of relative risk aversion $\gamma = 2$. The product $\mathbf{BR} = 0.9984 \approx 1$ indicates near golden rule dynamics, and the limiting MPC

under perfect foresight (all shocks at their mean value of 1) is $\underline{\kappa} = 1 - \Phi_R = 1 - \mathbf{R}^{(1-\gamma)/\gamma}\mathbf{B}^{1/\gamma} \approx 0.039$. Income process parameters follow empirical estimates of permanent and transitory earnings variation [Abowd and Card, 1989]: the permanent income standard deviation $\sigma_\psi = 0.06$ and transitory standard deviation $\sigma_\theta = 0.20$ match estimates from DeBacker et al. [2013]. The per-period unemployment probability $\wp = 0.005$ is chosen to ensure weak return impatience holds; this value is small enough that its effect on buffer-stock behavior is primarily through the patience conditions rather than through expected income levels. The deterministic income growth rate $\mathbf{G} = 1.03$ captures long-run real income growth. The preference and discount factor shock parameters ($\sigma_\beta = \sigma_\mu = 0.02$, $\rho_\beta = \rho_\mu = 0.20$) stress-test the method's scalability to high-dimensional state spaces.

A robustness check with these volatilities set to zero (reducing the state space from five to three dimensions) produces consumption functions that differ by less than 1% at mean state values, confirming that the solution method handles the additional state variables without degradation. The moderation bounds extend to all five dimensions with closed-form expressions (Lemma 3.1), so the architectural guarantees hold regardless of dimensionality.

To highlight buffer-stock dynamics more sharply for the main accuracy comparison in Table 3, the stylized calibration raises the time preference rate to 11.1% ($\mathbf{B} = 0.9$) and sets innovation standard deviations to $\sigma \approx 2\%$ for all shocks. At mean state values, the limiting MPC under perfect foresight is $\underline{\kappa} \approx 0.07$.

At quarterly frequency, the discount and return parameters become $\mathbf{B} = 0.99$ and $\mathbf{R} = 1.01$ (annual equivalents approximately 0.96 and 1.04). Shock standard deviations scale by $1/\sqrt{4}$ for quarterly frequency; the resulting values appear in Table 2.

All calibrations satisfy the buffer-stock existence conditions: finite value of autarky Proposition 3.2, weak return impatience Proposition 3.2, return impatience Proposition 3.2, and growth impatience Proposition 3.2. The stylized and quarterly calibrations also satisfy strong growth impatience Proposition 3.2, guaranteeing individual target wealth existence. The baseline calibration, with $\mathbf{BR} \approx 1$, lies near the boundary of strong growth impatience and exhibits a pseudo-target but no individual target for some states in the ergodic distribution. The unemployment probability $\wp = 0.005$ is of the same order of magnitude as the quarterly permanent-income-loss rate in Carroll [1997] (our baseline is annual; the value is chosen primarily to satisfy the patience conditions); it is also sufficient to satisfy weak return impatience Proposition 3.2 across all calibrations, ensuring the moderation bounds are well-defined.

5.2 Converged Solutions

Figure 3 confirms that the unconstrained network recovers the concave buffer-stock shape predicted by Theorem 3.2. Figure 4 shows the training dynamics for both approaches.[18] The Euler residual R_1 and complementarity residual R_2 are combined with weight $v = 1$ in the loss function Paragraph, giving equal importance to both conditions. Residuals R_1 and R_2 contribute approximately equally to the total loss at convergence.

Figure 5 demonstrates that the converged moderation policy respects the analytical bounds from Proposition 3.1 at every wealth level, with the pessimist-optimist gap narrowing as precautionary motives weaken.

Figure 6 shows the moderation ratio ω increasing monotonically with wealth, consistent with the theoretical prediction (Remark 5.1) that precautionary motives weaken as the buffer stock grows.

5.3 Shock Sensitivity Analysis

Beyond evaluating policies at mean state values, we examine how consumption and moderation ratios respond to shock uncertainty by evaluating at mean $\pm$ 1 standard deviation for all shocks simultaneously. The stylized calibration uses innovation standard deviations of $\sigma \approx 2\%$ for most shocks. Combined with persistence parameters ($\rho = 0.2$ for preference and return shocks, $\rho = 0.9$ for persistent income), ergodic standard deviations reach 2–5%. State normalization in the neural network input layer maps $\pm 1\sigma$ shocks to ± 1.0 in the network's input space.

[18] The reported residual norm is $\sqrt{\Xi(\mathbf{W})}$, the square root of the combined loss Paragraph, which represents the RMS deviation from Euler equation and complementarity conditions. We also compute the standard Euler equation error $\log_{10}|1 - c/c^*|$ from Judd [1998], where c^* is the Euler-implied consumption from 5,000 Monte Carlo draws at each wealth point (at mean exogenous states, unconstrained points only). See Table 3 for the comparison.

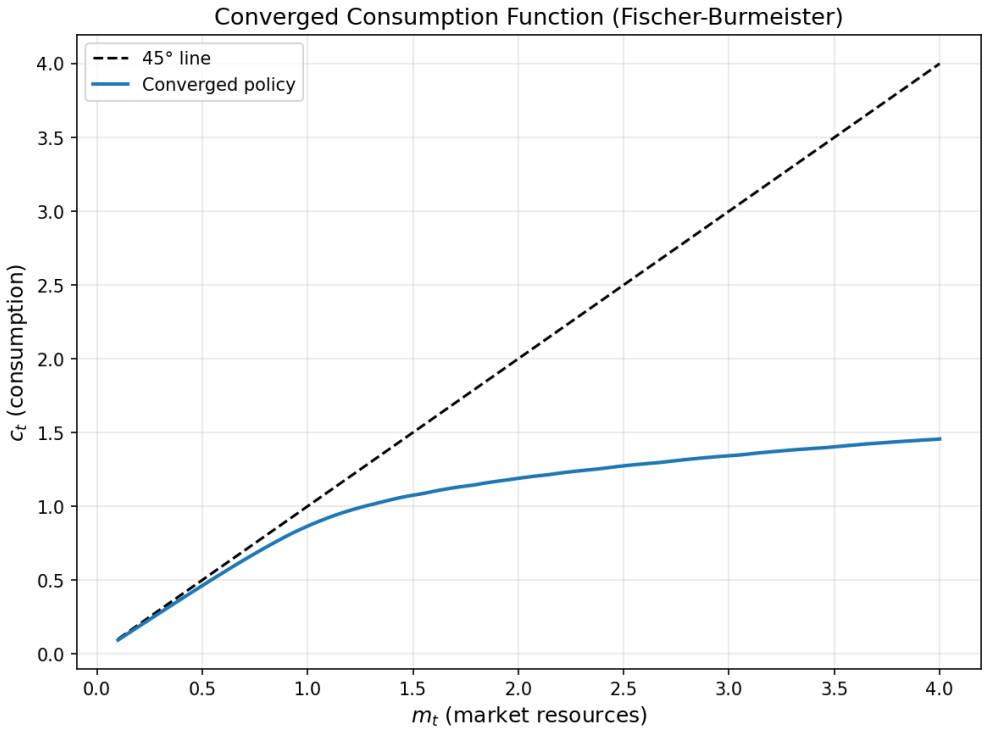


Figure 3: Even without structural constraints, the unconstrained network recovers concavity and monotonicity predicted by Theorem 3.2, confirming that the Euler residual objective identifies the correct basin when initialized at $\varsigma_0 = 0.5$. Consumption is strictly increasing, strictly concave, and lies below the 45-degree line at all wealth levels. Policy evaluated at mean state values after 50,000 iterations.

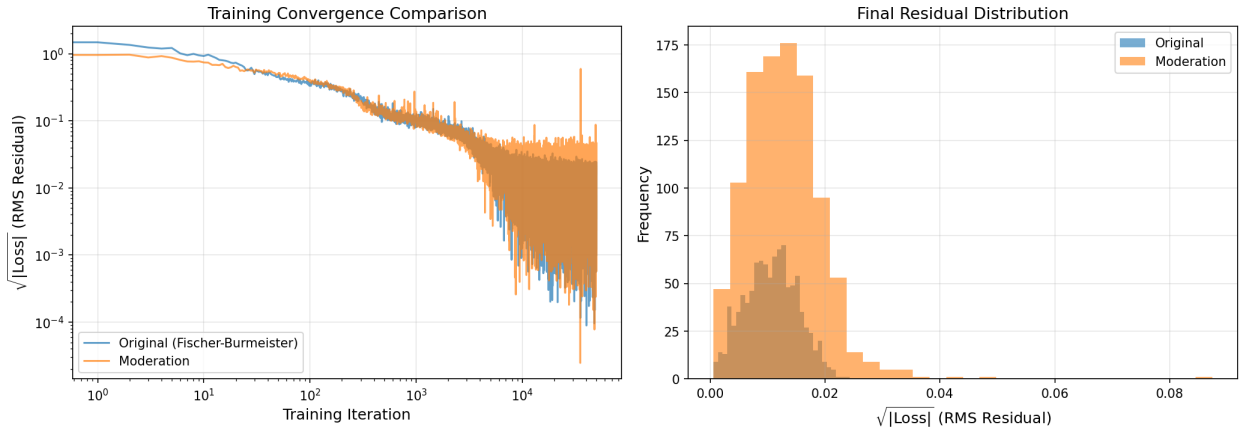


Figure 4: The moderation approach converges to roughly half the final loss of the unconstrained approach (5.2×10^{-5} vs. 1.3×10^{-4}), with smoother convergence because the bounded $\omega \in [0, 1]$ output prevents the large early-training residuals that cause gradient spikes in the unconstrained architecture. Left: RMS residual vs. iteration (log-log). Right: residual distribution over the final 1,000 iterations.

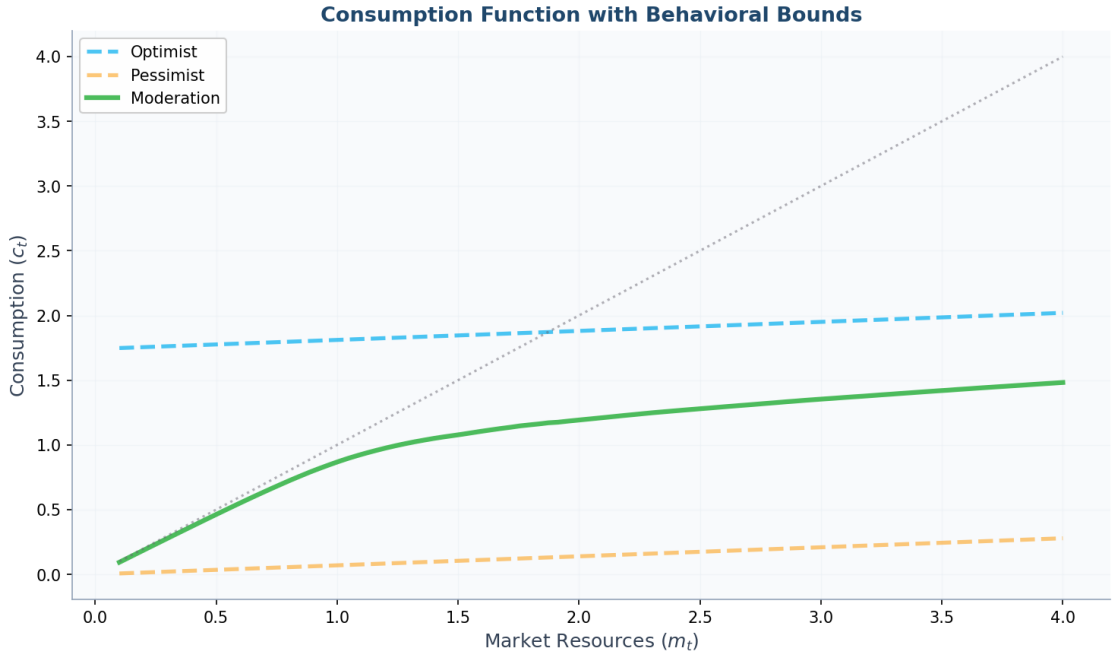


Figure 5: The moderation approach guarantees consumption between analytical bounds from Proposition 3.1 at every wealth level, evaluated at mean state values. The converged policy (green solid) lies strictly between the optimist bound (light blue dashed, mean future shocks) and pessimist bound (yellow dashed, permanent unemployment), with the gap narrowing at higher wealth as precautionary motives weaken.

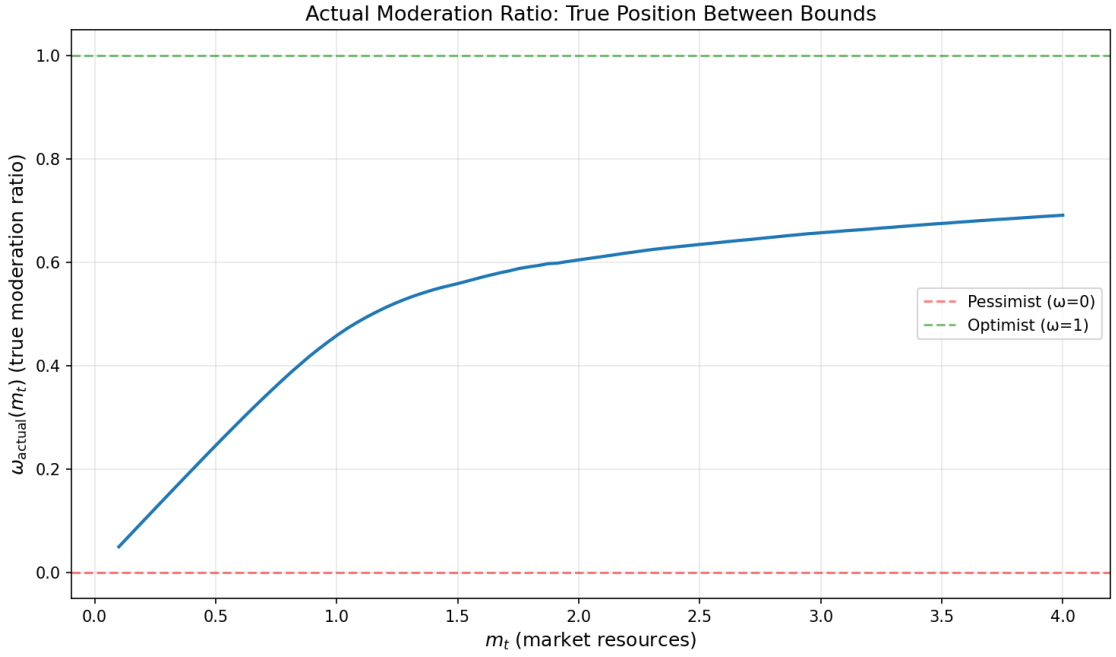


Figure 6: The moderation ratio $\omega(m_t, s_t)$ rises monotonically with wealth (Remark 5.1), confirming that precautionary motives weaken as the buffer stock grows: by $m \approx 4$, the consumer behaves three-quarters of the way toward the optimist rule. At low wealth, strong precautionary motives keep consumption close to the pessimist bound ($\omega = 0$, red dashed); as wealth increases, $\omega \to 1$ (green dashed) asymptotically. Evaluated at mean state values.

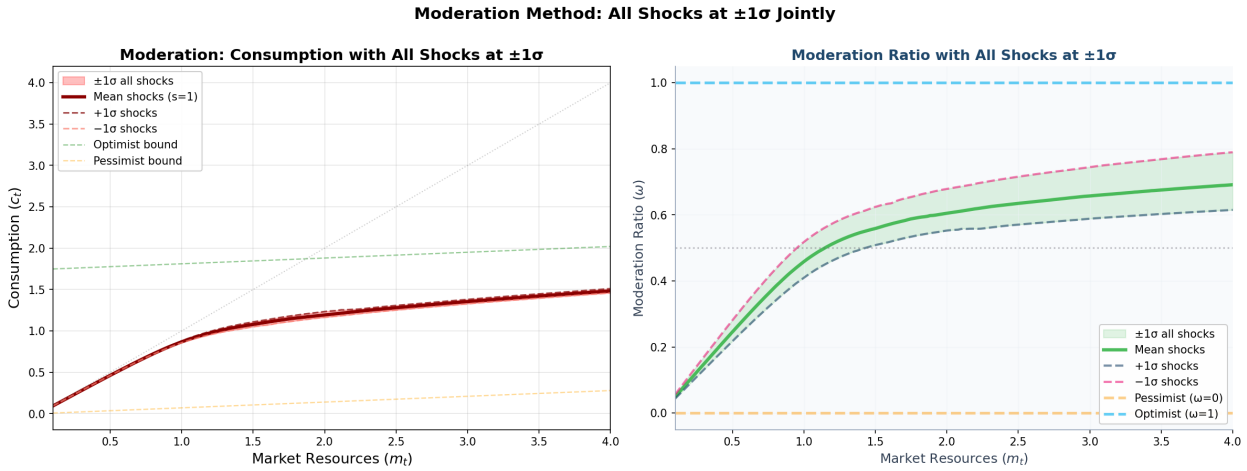


Figure 7: Policy responses to simultaneous $\pm 1\sigma$ shocks in all state variables. Left panel: Consumption with shock confidence intervals (shaded red band). Mean consumption (dark red solid) lies between $+1\sigma$ (maroon dashed) and -1σ (salmon dashed) responses. Analytical bounds (green and orange dashed) provide theoretical reference points. The 45-degree line (gray dotted) marks the borrowing constraint. Right panel: Moderation ratio with shock intervals (shaded green band). Mean ω (green solid) increases with wealth as precautionary motives weaken. Positive shocks ($+1\sigma$, dark blue dashed) push ω toward the optimist bound; negative shocks (-1σ, pink dashed) push toward the pessimist bound. Both panels evaluated across 200 wealth points.

Figure 7 quantifies the sensitivity of consumption and the moderation ratio to simultaneous $\pm 1\sigma$ perturbations across all state variables. The proportional impact of shocks on consumption decreases with wealth. At $m \approx 0.2$, moving all shocks $\pm 1\sigma$ simultaneously induces $\pm 15\%$ consumption variation; by $m \approx 4.0$ the swing halves to $\pm 5\%$. The proportional impact decreases because precautionary motives weaken as the buffer stock grows, so the same shock magnitude produces a smaller consumption response. Positive shocks push ω toward the optimist bound while negative shocks push toward the pessimist, confirming state-contingent adaptation. Throughout this range, consumption stays within the analytical bounds, even under the combined $\pm 1\sigma$ perturbation that shifts all five exogenous states simultaneously.

5.4 Validation of Theoretical Properties

Remark 5.1 (Moderation Approach Guarantees)**.** *The moderation approach inherits three theoretical guarantees from its construction: bound validity, monotonicity, and asymptotic behavior.*

Bound validity. *Since $\omega \in [0,1]$ via sigmoid activation and the min formula (32) maps ω into $[\underline{\kappa} m, \min(m, \underline{\kappa}(m + \bar{h}))]$, consumption satisfies $\underline{c} \leq c \leq \bar{c}$ at every state. The network can only approximate the true solution, not arbitrary functions.*

Monotonicity. *Under return impatience Proposition 3.2 and growth impatience Proposition 3.2, the moderation ratio $\omega(m_t, s_t)$ increases with wealth because $c^{*\prime}(m) > \underline{\kappa}$ at all finite wealth levels, giving $d\omega/dm > 0$ above the cusp.*

Asymptotic behavior. *The log-odds prior (36) drives $\omega \to 1$ at high wealth through the $\gamma \log(m/m^*)$ term, matching the theoretical decay rate $1 - \omega \propto m^{-\gamma}$ from Carroll [2011]. Consumption approaches $\underline{\kappa}(m + \bar{h})$ from below, recovering the optimist rule Theorem 3.6. Near the borrowing constraint, ω stays bounded away from 1, reflecting strong precautionary motives. Extrapolation errors are bounded by the gap $\bar{c} - \underline{c}$, which narrows as wealth grows.*

Concavity. *Above the cusp ($m > m^*$), the min formula inherits concavity from the trained ω function since $c'' = \underline{\kappa}\bar{h} \cdot \omega''$, and $\omega'' \leq 0$ follows from prudence ($u''' > 0$) via Carroll and Kimball [1996]. Below the cusp, concavity is a property of the converged solution guaranteed by the same prudence argument, not of the architecture for arbitrary weights.*

Figure 8 confirms that return impatience holds throughout the ergodic state space, with $\underline{\kappa}$ concentrated near 0.07. We validated theoretical properties on 1,000 randomly sampled states from the ergodic distribution. Return impatience Proposition 3.2 holds uniformly: $\Phi_R(s_t) \in (0.925, 0.935)$ across all samples, a narrow range reflecting the modest persistence ($\rho = 0.2$) of return and preference shocks. Growth impatience Proposition 3.2 holds as well, with $\Phi_G(s_t) \in (0.88, 0.92)$. Consistent with the MPC bounds Lemma 3.3, these ranges translate to a limiting MPC of $\underline{\kappa}(s_t) \in (0.065, 0.075)$ with mean 0.070 and standard deviation 0.0024; at the other extreme, the upper MPC bound $\bar{\kappa}(s_t) \approx 0.12$ governs behavior near the borrowing constraint.

Both orderings $\bar{h} \geq \underline{h}$ and $\bar{c} \geq \underline{c}$ hold universally, as the theory requires. Under the stylized calibration ($\mathbf{B} = 0.9$, $\mathbf{R} = 1.04$), the cusp point $m^* = \underline{\kappa}\bar{h}/(1-\underline{\kappa}) \approx 1.88$ (with $\underline{\kappa} \approx 0.07$ and $\bar{h} \approx 25$) separates low-wealth and high-wealth regimes; roughly 46% of the ergodic wealth distribution falls below this threshold. Strong growth impatience Proposition 3.2 holds, guaranteeing an individual target at $\check{m} \approx 1.8$ where $\mathbb{E}_t[m_{t+1}] = m_t$. As expected from the theoretical ordering $\tilde{m} < \hat{m} < \check{m}$ (Theorem 3.4), the pseudo-target for aggregate balanced growth lies lower, at $\hat{m} \approx 1.5$. That the target $\check{m}$ lies just below the cusp m^* means most agents in the ergodic distribution experience both binding-constraint regimes during their lifetime, exercising both branches of the min formula. Consumption from the trained network satisfies $\underline{c} \leq c \leq \bar{c}$ at every sampled state. The moderation ratio at target wealth, $\omega(\check{m}, s_t) \approx 0.65$, is strictly interior as predicted by Remark 3.2. The trained network places ω at this value because stationarity at $\check{m}$ pins down the consumption level: if ω were near 1, the consumer would spend like an optimist and wealth would fall below target; near 0, pessimist-level consumption would push wealth above target.

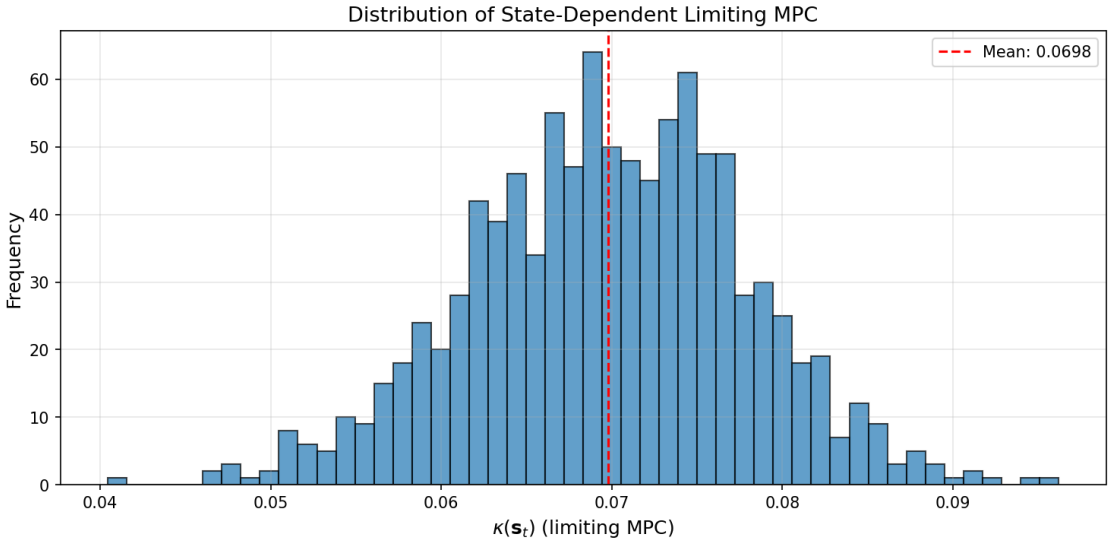


Figure 8: Distribution of state-dependent limiting marginal propensity to consume $\underline{\kappa}(s_t)$ across 1,000 randomly sampled states. The narrow concentration around $\underline{\kappa} \approx 0.07$ (red dashed) confirms that the Return Impatience Condition $\Phi_R^* < 1$ holds throughout the ergodic state space, and that the moderation floor $\underline{\kappa}m$ barely varies with the exogenous state, making the architectural guarantee robust rather than state-contingent.

5.5 Extrapolation Performance

Beyond the training range, the moderation approach keeps consumption within $[\underline{c}, \bar{c}]$:

Remark 5.2 (Extrapolation Properties)**.** *The moderation approach maintains theoretical consistency outside the training domain by construction. For any wealth level $m_t \in \mathbb{R}_+$, consumption satisfies $c(m_t, s_t) \in [\underline{c}(m_t, s_t), \bar{c}(m_t, s_t)]$ regardless of whether m_t was included in the training set. The unconstrained (MMW) approach guarantees $c \in [0, m]$ through the sigmoid architecture, which suffices within the training domain but provides no structural guidance outside it: the consumption share may drift toward 1 (near-zero saving) or toward 0 (near-zero consumption) in untrained regions.*

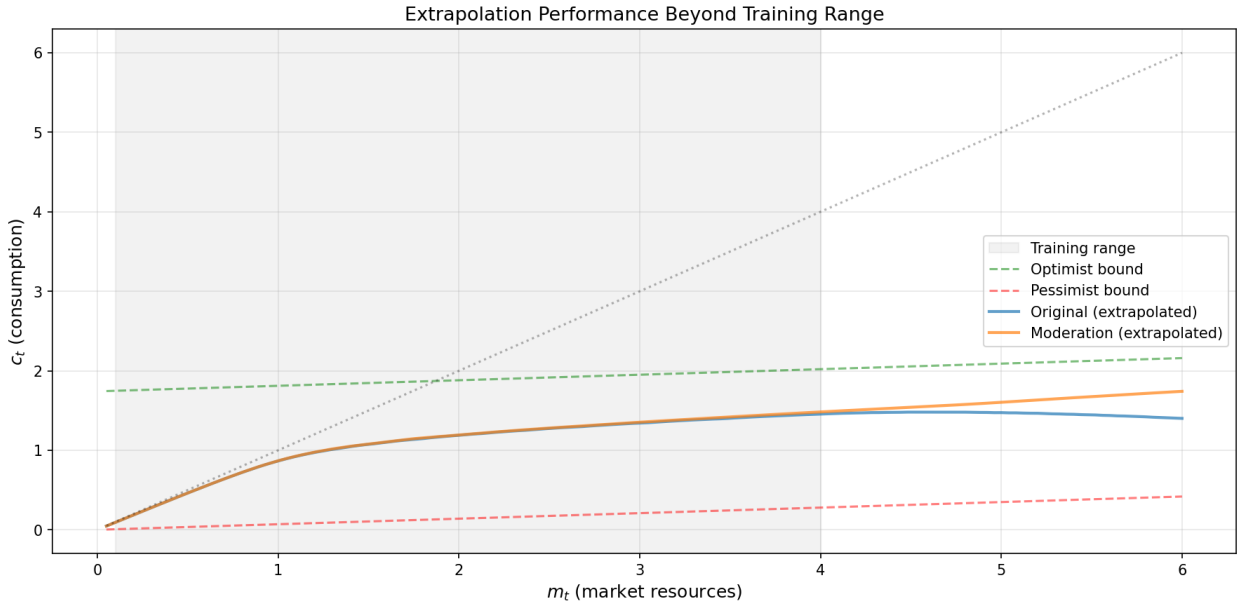


Figure 9: Extrapolation performance beyond the training wealth range (shaded gray region, $m \in [0.1, 4.0]$). The moderation approach (orange solid) respects analytical bounds (green and red dashed) in untrained regions ($m < 0.1$ and $m > 4.0$), maintaining consumption strictly between pessimist and optimist rules. Outside the training domain, the unconstrained (MMW) approach (blue solid) lacks structural guidance: the consumption share ς drifts toward 1 at high wealth (negligible saving), a natural consequence of extrapolating a learned function beyond its training support. At $m = 6.0$, the unconstrained network saves less than 2% of resources while the moderation network saves 28%, consistent with the optimist bound.

Table 3: Euler equation errors: unconstrained vs. moderation (stylized calibration, 50,000 iterations)

Architecture	Mean $\log_{10}$ EE	Median $\log_{10}$ EE	Max $\log_{10}$ EE	Mean AIO loss (iter. 49k-50k)
Unconstrained (MMW)	-2.05	-2.05	-1.54	1.3×10^{-4}
Moderation (MOD)	-2.75	-2.64	-1.92	5.2×10^{-5}

Figure 9 shows this contrast directly: at $m = 10$, well beyond the training range of $[0.1, 4.0]$, the unconstrained network's consumption share drifts toward $\varsigma \approx 1$ (near-zero saving). The moderation network stays within $[\underline{c}, \bar{c}]$ regardless of wealth, because the bounds are architectural rather than learned. For policy counterfactuals requiring evaluation at rarely observed states, this architectural guarantee matters more than marginal improvements in training-domain accuracy.

5.6 Comparison of Architectures

Table 3 compares the unconstrained (MMW) and moderation approaches on the stylized calibration, each trained for 50,000 iterations with learning rate $\alpha = 10^{-3}$, batch size $n = 128$, and the Rao-Blackwellized AIO estimator. Both architectures use identical hidden layers (3 layers, 32 units, ReLU); only the output transformation differs.

The moderation approach achieves lower Euler equation errors than the unconstrained method across all metrics (-2.75 mean vs. -2.05, -2.64 median vs. -2.05, -1.92 max vs. -1.54) and lower final loss (5.2×10^{-5} vs. 1.3×10^{-4}), an improvement of roughly 0.7 orders of magnitude in mean accuracy. The improvement stems from the γ-slope prior and the $\log(m/m^*)$ normalization: because the prior provides the correct asymptotic convergence rate, the network's capacity goes entirely to approximating the interior of the consumption function rather than rediscovering asymptotic behavior that theory already determines. The structural

guarantees ($c \in [\underline{c}, \bar{c}]$ by construction) come with an accuracy gain, not a penalty. At median wealth ($m \approx 1.8$), the moderation error of $10^{-2.64} \approx 0.2\%$ of consumption corresponds to roughly \$100 per year for a household with \$50,000 in permanent income.

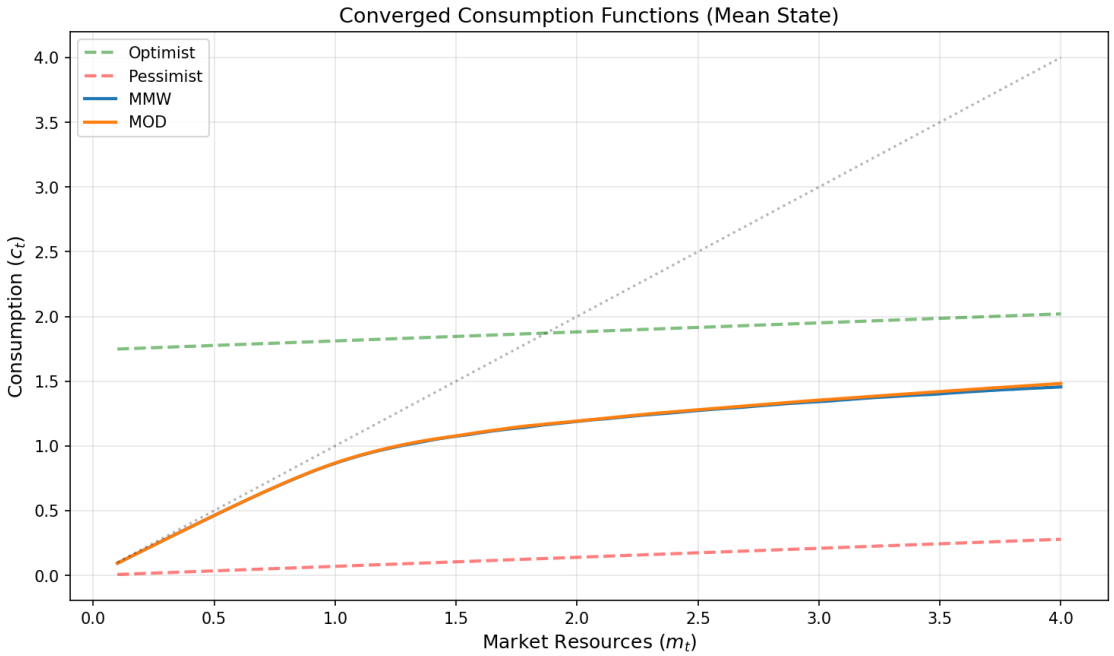


Figure 10: Converged consumption functions for the unconstrained (MMW) and moderation (MOD) architectures (mean state values, 50,000 iterations). The moderation approach finds the correct concave buffer-stock solution between the pessimist and optimist bounds. The unconstrained approach also converges to a reasonable solution when initialized at $\varsigma_0 = 0.5$, but converges to near-zero consumption when initialized at $\varsigma_0 = \underline{\kappa} \approx 0.07$, a degenerate local minimum the moderation architecture cannot reach.

5.6.1 Initialization Sensitivity

When initialized at the pessimist consumption share ($\varsigma_0 = \underline{\kappa} \approx 0.07$), the unconstrained network converges to a degenerate solution with near-zero consumption across all wealth levels. This solution is a local minimum of the loss function: when $c \approx 0$ uniformly, the consumption ratio $c_{t+1}/c_t \approx 1$ and the Euler kernel approximately satisfies $\mathcal{K}_{t+1} \approx 1$, so the Euler residual R_1 is small everywhere. The Fischer-Burmeister residual R_2 is also small because $\varsigma \approx 0$ implies large savings ($1 - \varsigma \gg 0$) and the multiplier λ can vanish. The optimizer is trapped because the gradient of the loss with respect to network parameters nearly vanishes: every policy in the neighborhood of the degenerate solution is also near-zero, creating a wide flat basin rather than a sharp minimum. The loss is low, but the solution is economically wrong: the agent hoards all wealth and never consumes.

Figure 10 shows results with $\varsigma_0 = 0.5$, which places the unconstrained network in the basin of attraction of the correct solution. The moderation approach avoids this issue by construction: $\omega = 0$ maps to pessimist consumption $\underline{c} = \underline{\kappa} m > 0$, not to zero, so the degenerate basin is unreachable.

In heterogeneous-agent computation, this robustness is not a convenience but a requirement. When the consumption policy is re-trained as aggregate conditions change, there is no guarantee that a previously good initialization remains in the correct basin after an equilibrium update. Figure 11 shows that the Rao-Blackwellized AIO estimator stabilizes convergence for both architectures, preventing the variance spikes that would otherwise require clipping or warmup schedules.

5.7 Robustness Considerations

The baseline calibration sits near the boundary of strong growth impatience ($\mathbf{BR} \approx 1$), a region where both traditional and neural network methods face challenges. Raising $\mathbf{B}$ toward the GIC boundary ($\Phi_G^* \to 1$) or

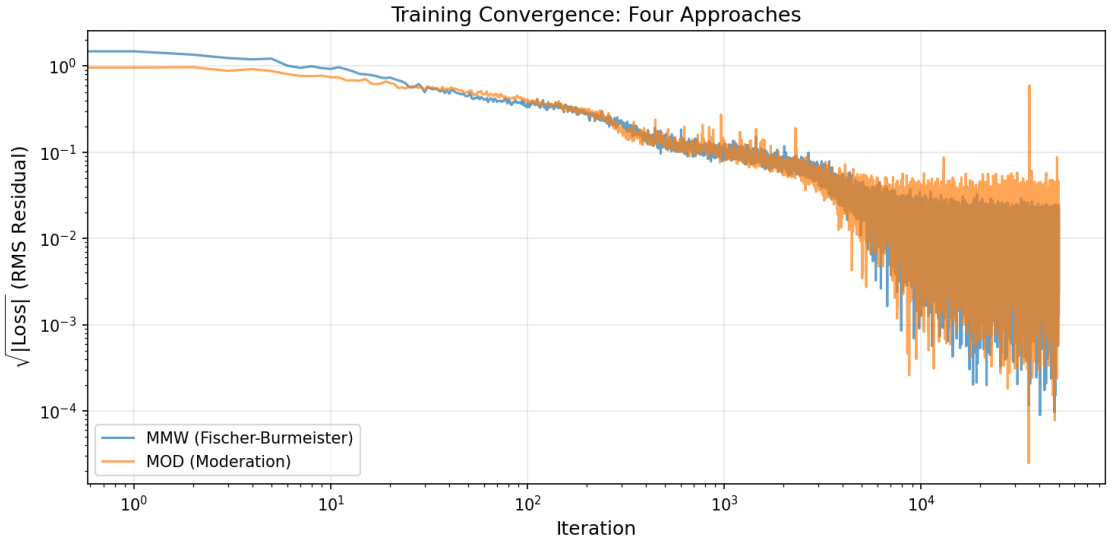


Figure 11: Training convergence for unconstrained (MMW) and moderation (MOD) approaches (RMS residual on log-log scale). Both architectures stabilize by iteration 10,000; the moderation approach reaches roughly half the final loss. The Rao-Blackwellized AIO estimator prevents the variance spikes that would otherwise require residual clipping or warmup schedules.

lowering $\mathbf{R}$ toward the RIC boundary ($\Phi_R^* \to 1$) slows convergence in both the contraction mapping and neural network training: the consumption function flattens, reducing the signal in Euler residuals. Both approaches converge for all three calibrations in Table 2, though we did not measure how convergence speed varies with distance from the patience boundaries. The baseline calibration ($\mathbf{BR} \approx 0.9984$) already sits near the boundary and provides a partial test, but a systematic map from patience factor margins to required training iterations remains open.

The primary modification that ensures training stability is the Rao-Blackwellized AIO estimator ((40)). Integrating out the discrete unemployment component analytically eliminates the variance spikes that occur when rare zero-income events ($\theta = 0$, probability $\wp \approx 0.005$) enter the stochastic product $R_1^{(1)} \cdot R_1^{(2)}$. Without Rao-Blackwellization, the product can exceed its typical magnitude by several orders when both independent draws happen to include unemployment (probability $\wp^2 \approx 2.5 \times 10^{-5}$), producing unstable gradients during early training. A secondary stabilization uses a floor on next-period market resources $m_{t+1} \geq \epsilon$ (we use $\epsilon = 10^{-8}$) to prevent numerical overflow in the CRRA marginal utility $c_{t+1}^{-\gamma}$ when unemployment drives resources near zero. Neither modification introduces bias at convergence.

6 Conclusion

Neural networks trained to solve consumption-saving problems with only a budget constraint ($c \in [0, m]$) lack structural guarantees outside the training domain: the consumption share can drift toward one (near-zero saving) at extreme wealth levels, and from certain initializations the network can converge to a degenerate near-zero consumption solution that satisfies the Euler equation but has no economic content. A simple output-layer reparameterization addresses both. Replacing the budget-constrained sigmoid output with a min formula that maps the network's ratio $\omega \in [0, 1]$ into the interval $[\underline{\kappa} m, \min(m, \underline{\kappa}(m + \bar{h}))]$ guarantees three constraints by construction: a positive MPC floor, the optimist upper bound, and the borrowing constraint. The moderation architecture achieves lower Euler equation errors than the unconstrained benchmark (Table 3) while correctly tracking analytical bounds outside the training range. The Rao-Blackwellized AIO estimator eliminates the variance spikes from rare unemployment events that would otherwise destabilize training. We extend patience conditions and MPC bounds from the canonical two-shock model [Carroll and Shanker, 2024] to state-dependent AR(1) processes with closed-form expressions that make this structural embedding possible.

Euler equation errors fall in the $\log_{10}$ range of -2 to -3, one to two orders of magnitude above the -4 to -6 range typical of endogenous gridpoints solutions [Carroll, 2006] on lower-dimensional problems. The comparison is not directly controlled: EGM solves a one- or two-dimensional problem, while the moderation network solves a five-dimensional one, so the accuracy gap overstates the neural network disadvantage on equal-dimensional problems. The gap also reflects a modest network architecture (3 hidden layers, 32 units); wider networks (128+ units per layer), curriculum training, or hybrid approaches that use the neural network for interpolation and analytical bounds for extrapolation are natural candidates for closing it, though we have not tested these systematically.

The moderation approach requires analytical bounds, restricting it to settings where CRRA homogeneity enables normalization by permanent income and yields linear perfect-foresight consumption rules. Portfolio choice with multiple assets, durable goods, habit formation, Stone-Geary minimum consumption, and Epstein-Zin preferences all break this structure because human wealth is no longer a simple geometric series and the MPC bounds no longer separate additively. For such settings, the unconstrained approach remains preferable.

For heterogeneous-agent macroeconomic models [Aiyagari, 1994, Huggett, 1993, Krusell and Smith, 1998], the moderation approach addresses a specific bottleneck: when simulating large populations, agents inevitably reach extreme wealth levels where training data is sparse. Embedded bounds guarantee sensible consumption at these states without dense grids in the tails. The 60-second training time enables rapid equilibrium iteration, and permanent income normalization connects naturally to aggregation methods [Harmenberg, 2021]. Two extensions would sharpen these results. First, closing the Euler error gap to -4 or better, as discussed above. Second, characterizing how convergence degrades as patience conditions approach their boundaries ($\Phi_R^* \to 1$ or $\Phi_G^* \to 1$): the consumption function flattens near these boundaries, reducing the signal in Euler residuals, and the near-unit-root behavior of the wealth process makes ergodic sampling less informative. A systematic map from patience factor margins to required network capacity and training iterations would guide practitioners in choosing architecture size for specific calibrations. The agent at three times target wealth who saved almost nothing under the unconstrained architecture now saves at approximately the rate the analytical optimist bound prescribes. Two output-layer changes made the difference: the min formula and the log-odds prior.

References

S. Rao Aiyagari. Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, 109(3):659–684, 1994. doi:10.2307/2118417.

Mark Huggett. The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of Economic Dynamics and Control*, 17(5-6):953–969, 1993. doi:10.1016/0165-1889(93)90024-M.

Per Krusell and Anthony A. Smith, Jr. Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy*, 106(5):867–896, 1998. doi:10.1086/250034.

Lilia Maliar, Serguei Maliar, and Pablo Winant. Deep learning for solving dynamic economic models. *Journal of Monetary Economics*, 122:76–101, 2021. doi:10.1016/j.jmoneco.2021.07.004. URL https://www.sciencedirect.com/science/article/abs/pii/S0304393221000799.

Miles S. Kimball. Precautionary saving in the small and in the large. *Econometrica*, 58(1):53–73, 1990. doi:10.2307/2938334.

Christopher D. Carroll and Akshay Shanker. Theoretical Foundations of Buffer Stock Saving. *Quantitative Economics*, 2024. URL https://github.com/econ-ark/BufferStockTheory. Forthcoming.

Marlon Azinovic, Luca Gaegauf, and Simon Scheidegger. Deep Equilibrium Nets. *International Economic Review*, 63(4):1471–1525, 2022. doi:10.1111/iere.12575.

Jiequn Han, Yucheng Yang, and Weinan E. Deepham: A Global Solution Method for Heterogeneous Agent Models with Aggregate Shocks. *arXiv preprint arXiv:2112.14377*, 2021. doi:10.2139/ssrn.3990409. URL https://arxiv.org/abs/2112.14377.

Marina Azinovic, Luca Gaegauf, and Simon Scheidegger. Economics-Inspired Neural Networks with Stabilizing Homotopies. *arXiv preprint arXiv:2303.14802*, 2023. doi:10.48550/arXiv.2303.14802.

Julien Pascal. Artificial neural networks to solve dynamic programming problems: A bias-corrected Monte Carlo operator. *Journal of Economic Dynamics and Control*, 162:104853, 2024. doi:10.1016/j.jedc.2024.104853.

Julian Ashwin, Paul Beaudry, and Martin Ellison. Neural network learning for nonlinear economies. *Journal of Monetary Economics*, 149:103723, 2025. doi:10.1016/j.jmoneco.2024.103723.

Hui Chen, Yuhan Cheng, Yanchu Liu, and Ke Tang. Teaching Economics to the Machines. Working Paper 34713, National Bureau of Economic Research, 2024. URL https://www.nber.org/papers/w34713.

Jesús Fernandez-Villaverde, Galo Nuno, and Jesse Perla. Taming the Curse of Dimensionality: Quantitative Economics with Deep Learning. Working Paper 33117, National Bureau of Economic Research, 2024.

Jack Schechtman. An income fluctuation problem. *Journal of Economic Theory*, 12(2):218–241, 1976. doi:10.1016/0022-0531(76)90075-2.

Christopher D. Carroll. Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis. *The Quarterly Journal of Economics*, 112(1):1–55, 1997. doi:10.1162/003355397555109. URL https://www.econ2.jhu.edu/people/ccarroll/papers/BufferStockTheory/.

Qingyin Ma, John Stachurski, and Alexis Akira Toda. The income fluctuation problem and the evolution of wealth. *Journal of Economic Theory*, 187:105003, 2020. doi:10.1016/j.jet.2020.105003.

Lilia Maliar, Serguei Maliar, and Kenneth L. Judd. Numerically stable and accurate stochastic simulation approaches for solving dynamic economic models. *Quantitative Economics*, 2(2):173–210, 2011. doi:10.3982/QE14.

Kenneth L. Judd. *Numerical Methods in Economics*. MIT Press, Cambridge, MA, 1998. doi:10.7551/mitpress/4904.001.0001.

Christopher D. Carroll. The method of endogenous gridpoints for solving dynamic stochastic optimization problems. *Economics Letters*, 91(3):312–320, 2006. doi:10.1016/j.econlet.2005.09.013. URL https://www.econ2.jhu.edu/people/ccarroll/EndogenousGridpoints.pdf.

Lilia Maliar and Serguei Maliar. Envelope condition method versus endogenous grid method for solving dynamic programming problems. *Economics Letters*, 120(2):262–266, 2013. doi:10.1016/j.econlet.2013.04.031.

Sydney C. Ludvigson and Alexander Michaelides. Does buffer-stock saving explain the smoothness and excess sensitivity of consumption? *American Economic Review*, 91(3):631–647, 2001. doi:10.1257/aer.91.3.631.

Milton Friedman. *A Theory of the Consumption Function*. Princeton University Press, Princeton, NJ, 1957. doi:10.1515/9780691188485.

John F. Muth. Optimal properties of exponentially weighted forecasts. *Journal of the American Statistical Association*, 55(290):299–306, 1960. doi:10.1080/01621459.1960.10482064.

Jack Schechtman and Vera L. S. Escudero. Some results on "an income fluctuation problem". *Journal of Economic Theory*, 16(2):151–166, 1977. doi:10.1016/0022-0531(77)90003-5.

Stephen P. Zeldes. Optimal consumption with stochastic income: Deviations from certainty equivalence. *The Quarterly Journal of Economics*, 104(2):275–298, 1989. doi:10.2307/2937848.

John M. Abowd and David Card. On the covariance structure of earnings and hours changes. *Econometrica*, 57(2):411–445, 1989. doi:10.2307/1912561.

Thomas E. MaCurdy. The use of time series processes to model the error structure of earnings in a longitudinal data analysis. *Journal of Econometrics*, 18(1):83–114, 1982. doi:10.1016/0304-4076(82)90096-3.

Juan Pablo Rincón-Zapatero and Carlos Rodrı
guez-Palmero. Existence and uniqueness of solutions to the Bellman equation in the unbounded case. *Econometrica*, 71(5):1519–1555, 2003. doi:10.1111/1468-0262.00457.

Qingyin Ma, John Stachurski, and Alexis Akira Toda. Unbounded dynamic programming via the Q-transform. *Journal of Mathematical Economics*, 100:102652, 2022. doi:10.1016/j.jmateco.2022.102652.

Fernando Alvarez and Nancy L. Stokey. Dynamic programming with homogeneous functions. *Journal of Economic Theory*, 82(1):167–189, 1998. doi:10.1006/jeth.1998.2431.

V. Filipe Martins-da Rocha and Yiannis Vailakis. Existence and uniqueness of a fixed point for local contractions. *Econometrica*, 78(3):1127–1141, 2010. doi:10.3982/ECTA7920.

Anna Jaśkiewicz and Andrzej S. Nowak. Discounted dynamic programming with unbounded returns: Application to economic models. *Journal of Mathematical Analysis and Applications*, 378(2):450–462, 2011. doi:10.1016/j.jmaa.2010.08.073.

Huiyu Li and John Stachurski. Solving the income fluctuation problem with unbounded rewards. *Journal of Economic Dynamics & Control*, 45:353–365, 2014. doi:10.1016/j.jedc.2014.06.003.

Gary Chamberlain and Charles A. Wilson. Optimal intertemporal consumption under uncertainty. *Review of Economic Dynamics*, 3(3):365–395, 2000. doi:10.1006/redy.2000.0098.

Adam Szeidl. Invariant Distribution in Buffer-Stock Saving and Stochastic Growth Models. 2013.

George 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.

John H. Boyd. Recursive utility and the Ramsey problem. *Journal of Economic Theory*, 50(2):326–345, 1990. doi:10.1016/0022-0531(90)90006-6.

Jorge Durán. Discounting long run average growth in stochastic dynamic programs. *Economic Theory*, 22 (2):395–413, 2003. doi:10.1007/s00199-002-0316-5.

Christopher D. Carroll and Miles S. Kimball. On the concavity of the consumption function. *Econometrica*, 64(4):981–992, 1996. doi:10.2307/2171853.

Karl Harmenberg. Aggregating heterogeneous-agent models with permanent income shocks. *Journal of Economic Dynamics & Control*, 129:104185, 2021. doi:10.1016/j.jedc.2021.104185.

Christopher D. Carroll. Solution Methods for Microeconomic Dynamic Stochastic Optimization Problems. techreport, Johns Hopkins University, 2011. URL `https://www.econ2.jhu.edu/people/ccarroll/SolvingMicroDSOPs`.

Andreas Fischer. A special Newton-type optimization method. *Optimization*, 24(3-4):269–284, 1992. doi:10.1080/02331939208843795.

Diederik P. Kingma and Jimmy Ba. Adam: A Method for Stochastic Optimization. In *Proceedings of the 3rd International Conference on Learning Representations (ICLR)*, 2015. URL `https://arxiv.org/abs/1412.6980`.

Jason DeBacker, Bradley Heim, Vasia Panousi, Shanthi Ramnath, and Ivan Vidangos. Rising inequality: Transitory or persistent? New evidence from a panel of U.S. tax returns. *Brookings Papers on Economic Activity*, 2013(1):67–142, 2013. doi:10.1353/eca.2013.0004.