Why Don't the Poor Own Stocks? A Theory of Endogenous Hedging and Limited Participation *

Mingzhe Li †‡

Wednesday 25th June, 2025

Abstract

This paper develops a continuous-time heterogeneous agent framework that reveals a novel mechanism—Endogenous Aggregate Risk Aversion (EARA)—to explain portfolio choice and wealth inequality. This framework shows that when households face common shocks, they optimally adjust their portfolios to hedge the endogenous risk to their own stochastic continuation value. This EARA strategic adjustment generates powerful, wealth-dependent heterogeneity in risk-taking. For households near borrowing constraints, the adjustment is so strong that it compels them to hold zero risky assets, providing an endogenous rationalization for limited participation without ad-hoc assumptions. Conversely, the adjustment's influence diminishes with wealth, causing the portfolios of affluent households to converge toward the classical baseline ratio. By endogenously sorting agents into different risk-return profiles, this model generates non-degenerate wealth distributions and demonstrates a new channel through which aggregate risk shapes inequality.

*The author declares no specific research support or funding for this project. The research was conducted as part of the author's academic duties at the College of Economics and Management, Shihezi University. The author gratefully acknowledges intellectual contributions from Benjamin Moll, Jesús Fernández-Villaverde, and Samuel Hurtado. The author has no financial relationships or other potential conflicts of interest to disclose.

†Email: mingzheli@stu.shzu.edu.cn, College of Economics and Management, Shihezi University

‡Replication code for this paper is available on GitHub: https://github.com/MZLi-econ/EARA.git.

1 Introduction

This paper introduces the concept of Endogenous Aggregate Risk Aversion (EARA)—a strategic adjustment in portfolio allocation that endogenously arises when households internalize their exposure to common economic shocks. We demonstrate that this previously unrecognized single mechanism provides a unified and endogenous explanation for the significant gap between theory and empirical evidence in macroeconomics and finance: the limited participation of low-wealth households in risky asset markets—a stark contrast to standard portfolio theory, which predicts a constant, positive allocation to risky assets regardless of wealth. The EARA framework resolves the core contradiction in standard continuous-time heterogeneous agent (CTHA) models, which have struggled to simultaneously generate these key patterns.

The EARA mechanism stems from a fundamental departure from standard models. In the standard incomplete market model of Aiyagari (1994), Bewley (1986), and Huggett (1993) in continuous time frameworks, an agent's optimal problem is a deterministic object. However, when the economy is subject to aggregate uncertainty and each household faces common shocks, a household's value function evolves into a random field, governed by a backward stochastic partial differential equation rather than a classical HJB equation. Households understand that their lifetime utility is itself a source of risk driven by common shocks. The EARA represents the wealth-dependent consumption-portfolio optimization required to hedge this endogenous risk to their own continuation value. It is not a feature of preferences, but a consequence of optimal behavior in a world with aggregate risk.

The theoretical power of the EARA is channeled through a wealth-dependent strategic adjustment, which alters the classical portfolio allocation. We show this adjustment, Ω^*, can be expressed as a function of what we term Endogenous Relative Risk Aversion (ERRA), denoted by η. Specifically, $\Omega^* = \frac{1}{a\sigma_s}\frac{\eta-1}{\eta}$, where $\eta = -a\frac{v_{aa}}{v_a}$. Through the envelope condition, we link this effective risk aversion directly to observable economic behavior: $\eta = \epsilon_{c,a} \cdot \gamma$, where $\epsilon_{c,a}$ is the elasticity of consumption with respect to wealth and γ is the household's CRRA preference parameter.

For households near the borrowing constraint, we prove that ERRA $\eta(a)$ governing portfolio choice approaches zero. This singular behavior causes the EARA strategic adjustment, Ω^*, to diverge to negative infinity. This creates a powerful endogenous hedging demand that overwhelms the positive myopic demand for the risky asset, driving the optimal share, ω^*, to its lower bound of zero. This theoretical result provides a compelling rationalization for limited participation without relying on ad-hoc transaction costs, preference heterogeneity, or extreme income-risk correlations.

In contrast, for the wealthiest households, we prove that in the temporary equilibrium, the household consumption-to-wealth ratio still converges to a ratio in a stochastic steady state without income risk, and this ratio is equal to the Merton baseline ratio (Merton, 1969). Theoretical results on the asymptotic convergence of the optimal consumption-to-wealth ratio imply that the consumption-wealth elasticity $\epsilon_{c,a}$ converges to unity, which makes the ERRA parameter η converge to the preference-based CRRA parameter. At the same time, as the wealth factor gradually increases, the EARA adjustment Ω^* approaches zero from above, ensuring its influence becomes negligible. Their portfolio allocation thus endogenously converges to the classical Merton baseline.

This wealth-dependent heterogeneity in portfolio choice becomes a powerful engine for distributional dynamics. By endogenously sorting agents into different risk-return profiles, the EARA mechanism generates non-degenerate wealth distributions even when households are ex-ante identical and labor income risk is purely idiosyncratic. Methodologically, our paper demonstrates how the mathematical theory of "Mean Field Games with Common Noise"(Cardaliaguet et al., 2019), particularly the BSPDE formulation, can be harnessed to unlock new economic insights that are inaccessible within the standard HJB framework. We bridge the gap between abstract mathematical finance theory and core macroeconomic questions of inequality and portfolio choice.

The remainder of the paper formalizes these claims. We develop the full theoretical model, derive the analytical properties of the EARA mechanism at both tails of the wealth distribution, and provide a quantitative analysis to illustrate the mechanism's power. Our findings suggest that the feedback loop between macro aggregate risk and micro individual optimal choice is a first-order determinant of household financial decisions and macroscopic inequality.

1.1 Literature Reviews

Our work is situated at the confluence of three dynamic research areas: CTHA macroeconomics, the mathematical theory of MFG with common noise, and the development of advanced computational methods to solve high-dimensional economic models.

This paper first contributes to the CTHA literature that builds on the seminal Aiyagari-Bewley-Huggett framework. The continuous-time approach, pioneered in this context by Achdou et al. (2022), offers significant advantages over discrete-time formulations by avoiding integer problems and generating sparse equilibrium conditions. Subsequent research, such as Ahn et al. (2018) and Fernández-Villaverde et al. (2023), has extended these models to incorporate multiple assets and financial frictions. However, a fundamental challenge persists: these frameworks often abstract from a meaningful

asset distribution or, when they include it, struggle to endogenously generate empirically documented patterns like the limited participation of low-wealth households in risky-asset markets.

Methodologically, our framework departs from standard CTHA models by formally adopting the structure of a mean-field game with common noise. The mathematical foundations of MFG theory were established by Lasry and Lions (2007), with a substantial literature focusing on systems without aggregate uncertainty. The introduction of common noise, which is essential for studying aggregate risk, represents a major theoretical leap. Carmona, Delarue, et al. (2018) and, most comprehensively, Cardaliaguet et al. (2019) characterized the resulting system of coupled backward-forward stochastic partial differential equations. While this literature provides the rigorous mathematical apparatus, its application to address core macroeconomic questions of portfolio allocation and wealth inequality remains a nascent field. Our paper helps bridge this gap by using the BSPDE formulation to derive new economic insights.

Finally, our research is linked to the rapidly advancing frontier of computational macroeconomics. For models without common noise, a range of powerful techniques now exists, including the finite difference methods used by Achdou et al. (2022), the state-space iteration of Ahn et al. (2018), and the machine learning approaches of Fernández-Villaverde et al. (2023). More recent innovations include the perturbation techniques of Bilal (2023) and the sequence-space Jacobian methods of Bilal and Goyal (2025). However, the presence of common noise in our model renders these methods insufficient. The resulting BSPDEs require novel computational strategies. We follow the pathbreaking work of Han et al. (2018), who pioneered the use of deep learning frameworks to solve high-dimensional parabolic PDEs. Adapting and extending these techniques to the specific structure of heterogeneous agent models under aggregate uncertainty is an active area of research to which our future numerical implementation will contribute.

2 Economic Model

The economy comprises a continuum of infinitely-lived households indexed by $i \in [0, 1]$ with unit mass. Households exhibit constant relative risk aversion (CRRA) preferences over consumption, characterized by the instantaneous utility function:

$$U(c_t) := \frac{c_t^{1-\gamma} - 1}{1 - \gamma},$$

where $\gamma > 0$ is the coefficient of relative risk aversion.

Each household i chooses optimal consumption c_t and portfolio allocation ω_t to maximize expected lifetime utility:

$$\max_{c_t,\omega_t} \mathbb{E}_0 \left[\int_0^\infty e^{-\rho t} U(c_t)\, dt \right],$$

where $\rho > 0$ is the discount rate. The household faces a continuous-time intertemporal budget constraint:

$$da_t = \{[\omega_t(\mu_s - r_f) + r_f]a_t + y_t - c_t\}\, dt + \omega_t a_t \sigma_s dW_s, \tag{1}$$

subject to a borrowing constraint $a_t \geq 0$ and initial wealth $a_0 > 0$. At the borrowing constraint, consumption cannot exceed current labor income: $0 < c_t \leq y_t$.

The household's investment opportunity set consists of two assets: a risky asset with instantaneous return dS_t/S_t and a risk-free asset with return dB_t/B_t. Labor income $y_t = w_t n_t$ constitutes a non-tradable asset that influences the dynamics of the precautionary saving motive.

Financial assets and labor income follow diffusion processes:

$$\begin{aligned} \frac{dS_t}{S_t} =& \mu_s dt + \sigma_s dW_s \\ \frac{dB_t}{B_t} =& r_f dt \\ dy_t =& \mu_y dt + \sigma_y dW_y \end{aligned}$$

with correlation structure $dW_s dW_y = \rho_{sy} dt$. The Wiener process dW represents the common noise affecting all households.

We impose the economically motivated assumption that market risk in risky assets is more "fundamental" and partially drives labor income fluctuations. This implies:

$$\begin{aligned} dW_s =& dW_1 \\ dW_y =& \rho_{sy} dW_1 + \sqrt{1 - \rho_{sy}^2}\, dW_2, \end{aligned} \tag{2}$$

where dW_2 is a standard Wiener process independent of dW_1. Consequently:

$$dW_2 = \frac{1}{\sqrt{1 - \rho_{sy}^2}} (dW_y - \rho_{sy} dW_s). \tag{3}$$

The two-dimensional Wiener process vector is expressed as:

$$d\mathbf{W}_t = (dW_1, dW_2)^\top. \tag{4}$$

The household's dynamic optimization problem yields a backward stochastic partial differential equation (BSPDE) governing the evolution of the value function. Following

the theoretical framework developed by Cardaliaguet et al. (2019), Carmona, Delarue, et al. (2018), and Peng (1992), the value function $v_t(a_t, y_t, g_t, \mathbf{W}_t)$ satisfies:

$$dv = \left\{ u(c_t) + \mathcal{L}[v_t] + \mathbb{E}_t \left[\frac{\partial v_t}{\partial t} \right] - \text{div}(\mathcal{V}_t) \right\} dt + \mathcal{V}_t d\mathbf{W}_t, \tag{5}$$

where $\mathcal{V}_t$ represents the volatility process of the value function, and the infinitesimal generator $\mathcal{L}$ of the underlying state process is given by:

$$\begin{aligned} \mathcal{L}[v_t] \equiv & \{[\omega_t(\mu_s - r_f) + r_f]a_t + y_t - c_t\} \frac{\partial v}{\partial a} + \mu_y \frac{\partial v}{\partial y} \\ & + \frac{1}{2}(\omega_t a_t \sigma_s)^2 \frac{\partial^2 v}{\partial a^2} + \frac{1}{2}\sigma_y^2 \frac{\partial^2 v}{\partial y^2} + \omega a_t \rho_{sy} \sigma_s \sigma_y \frac{\partial^2 v}{\partial a \partial y}. \end{aligned} \tag{6}$$

The term $\mathbb{E}_t[\frac{\partial v_t}{\partial t}]$ captures how distributional changes affect individual value functions through general equilibrium effects:

$$\mathbb{E}_t \left[\frac{\partial v_t}{\partial t} \right] \equiv \iint \frac{\delta v_t}{\delta g_t(a, y')} \frac{\partial g_t(a, y')}{\partial t} da \, dy, \tag{7}$$

where $\frac{\delta v_t}{\delta g_t(a,y')}$ denotes the Fréchet derivative, quantifying the marginal impact of distributional changes on the value function.

The volatility process of the value function, $\mathcal{V}_t$, has components:

$$\begin{aligned} \mathcal{V}_{1,t} =& \omega_t a_t \sigma_s \frac{\partial v_t}{\partial a} + \sigma_y \rho_{sy} \frac{\partial v_t}{\partial y} \\ \mathcal{V}_{2,t} =& \sigma_y \sqrt{1 - \rho_{sy}^2} \frac{\partial v_t}{\partial y}. \end{aligned} \tag{8}$$

Economically, the stochastic term $\mathcal{V}_t d\mathbf{W}_t$ represents the direct channel through which aggregate uncertainty, common to all agents, propagates to each household's lifetime utility. Its presence transforms the value function from a simple deterministic object into a random field, whose evolution is intrinsically stochastic.

We define $-\text{div}(\mathcal{V}_t)$ as the *endogenous aggregate risk aversion* (EARA):

$$\begin{aligned} -\text{div}(\mathcal{V}_t) =& -\left\{ \frac{\partial \mathcal{V}_{1,t}}{\partial a} + \frac{\partial \mathcal{V}_{2,t}}{\partial y} \right\} \\ =& -\left\{ \omega_t \sigma_s \left(a_t \frac{\partial^2 v_t}{\partial a^2} + \frac{\partial v_t}{\partial a} \right) + \sigma_y \rho_{sy} \frac{\partial^2 v_t}{\partial y \partial a} + \sigma_y \sqrt{1 - \rho_{sy}^2} \frac{\partial^2 v_t}{\partial y^2} \right\}, \end{aligned} \tag{9}$$

which appears in the drift component of the BSPDE and constitutes the central mechanism of this paper. This term is a fundamental consequence of optimizing in a mean-field environment with common shocks; it arises because the volatility process $\mathcal{V}_t$ is itself adapted to the common noise filtration. It quantifies the endogenous risk premium required by the household to compensate for the fact that its value function is itself a source of risk. As we will demonstrate in Section 4, a crucial part of this term is

directly affected by the household's control variable ω_t. This premium is derived from the Itô-Kunita-Wentzell formula for random fields (Kunita, 1990).

The presence of common noise introduces a fundamental challenge in analyzing mean field games, as emphasized by Cardaliaguet et al. (2019): the coupled system comprising the stochastic Kolmogorov Forward (SKF) equation and Hamilton-Jacobi-Bellman (HJB) equations becomes itself stochastic. Both the forward equation governing distribution dynamics and the backward equation determining value functions become stochastic, as both the equilibrium distribution g_t and the value function v_t depend on realizations of common shocks $\mathbf{W}_t$.

The SKF equation governing the evolution of the wealth-income distribution g_t takes the form:

$$dg_t = (\mathcal{L}^* g_t)dt - \text{div}(\mathcal{V}_t g_t)d\mathbf{W}_t \tag{10}$$

where $\mathcal{L}^*$ is the adjoint of the operator $\mathcal{L}$:

$$\begin{aligned}\mathcal{L}^*[g_t] \equiv &- \left\{[\omega_t(\mu_s - r_f) + r_f]a_t + y_t - c_t\right\}\frac{\partial g}{\partial a} - \mu_y \frac{\partial g}{\partial y} \\ &+ \frac{1}{2}(\omega_t a_t \sigma_s)^2 \frac{\partial^2 g}{\partial a^2} + \frac{1}{2}\sigma_y^2 \frac{\partial^2 g}{\partial y^2} + \omega_t a_t \rho_{sy}\sigma_s\sigma_y \frac{\partial^2 g}{\partial a \partial y}.\end{aligned} \tag{11}$$

The divergence of $\mathcal{V}_t g_t$ has components:

$$\begin{aligned}\text{div}(\mathcal{V}_{1,t} g_t) =& \frac{\partial}{\partial a}\left(\omega_t a_t \sigma_s g_t\right) + \frac{\partial}{\partial y}\left(\sigma_y \rho_{sy} g_t\right) \\ \text{div}(\mathcal{V}_{2,t} g_t) =& \frac{\partial}{\partial y}\left(\sigma_y \sqrt{1-\rho_{sy}^2}\, g_t\right).\end{aligned} \tag{12}$$

The density satisfies the normalization $\iint g_t(a,y)\,da\,dy = 1$.

2.1 Market Clearing

The model imposes three market clearing conditions that ensure consistency between individual optimization and aggregate resource constraints.

Aggregate capital must equal total household financial assets allocated to risky assets:

$$K_t = \iint \omega_i a_{i,t} g_t(a,y)\,da\,dy. \tag{13}$$

Aggregate labor supply equals labor demand:

$$L_t = \iint n_{i,t} g_t(a,y)\,da\,dy. \tag{14}$$

The risk-free bond market clears with zero net supply:

$$B_t = \iint (1-\omega_{i,t}) a_{i,t} g_t(a,y)\,da\,dy. \tag{15}$$

These conditions ensure that all financial assets are held by households and that aggregate demand for each asset equals its supply.

3 Equilibrium

Under rational expectations, households must forecast the entire cross-sectional distribution g_t to predict equilibrium prices—a computationally intractable problem that exemplifies an extreme form of the curse of dimensionality (Dend Haan, 1996; Fernández-Villaverde et al., 2024; Krusell & Smith, 1998). This modeling approach, while theoretically rigorous, contradicts substantial empirical evidence suggesting that real-world agents employ considerably simpler forecasting heuristics.

Given realistic constraints on information processing capabilities and the inherent complexity of "the big world" (Javed & Sutton, 2024), we adopt a bounded rationality framework wherein households exhibit *sparse attention*. Specifically, we assume households focus exclusively on the first moment, which is the aggregate state variable K_t defined in equation (13), consistent with the sparsity-based approach of Gabaix (2014).

Rather than maintaining perfect foresight regarding K_t's evolution, each household i forms conditional expectations $\tilde{k}_{i,t} = \mathbb{E}_t[k_t|\mathcal{F}_t]$ based on available information $\mathcal{F}_t$. We postulate that households construct subjective mappings from economic shocks to the dynamics of distributional moments.

This leads to our central behavioral assumption: households operate under a *perceived law of motion* (PLM) that takes the form of a McKean-Vlasov equation:

$$d\tilde{k}_{i,t} = \tilde{\mu}_k(X_{i,t}, g_t, \mathbf{W}_t; \boldsymbol{\beta}_{i,t})dt + \tilde{\sigma}_{i,t}\mathbf{B}_{i,t} \tag{16}$$

where $\mathbf{B}_{i,t}$ represents idiosyncratic belief shocks that generate exogenous demand fluctuations. The terminology follows Evans and Honkapohja (2012), and our approach shares conceptual similarities with the approximate aggregation methodology of Krusell and Smith (1998), where macroeconomic dynamics are captured through capital evolution. This framework also connects to recent work by Moll (2024), who demonstrates how households can use price systems to predict distributional moments.

The well-posedness of our McKean-Vlasov formulation follows from the Glivenko-Cantelli theorem (Rachev & Rüschendorf, 2006), and standard results in mean field theory (Cardaliaguet et al., 2019). For large populations, the idiosyncratic shocks $\mathbf{B}_{i,t}$ generate only deterministic drift terms in the aggregate equation (16), eliminating stochastic components at the population level (Achdou et al., 2021).

To maintain analytical tractability while preserving the model's essential economic content, we impose the normalization $\tilde{\sigma}_{k,i} = 0$, effectively removing the stochastic component from the PLM. This simplification allows us to focus on the systematic biases in belief formation while avoiding the technical complications of idiosyncratic beliefs (Anderson et al., 2005; Bhamra & Uppal, 2014; De Long et al., 1990; Jouini & Napp, 2007; Shiller et al., 1984).

Our bounded rationality assumption naturally leads to a temporary equilibrium concept, bridging the macroeconomic literature on limited information processing (Maćkowiak & Wiederholt, 2009, 2015) with microeconomic evidence on wealth-dependent expectation formation (Malmendier & Nagel, 2016).

Definition 1. At time t, a *temporary equilibrium* consists of consumption-portfolio allocations $\{c_{i,t}, \omega_{i,t}\}$, prices $\{r_{f,t}, \mu_{s,t}, \mu_{y,t}\}$, and the distribution $g_t(a, y)$ such that:

1. Each household (i) chooses consumption-portfolio allocation $c_{i,t}, \omega_{i,t}$ to maximize its value function $v_t(a, y, g_t, \mathbf{W}_t)$ given its beliefs and information set $\mathcal{F}_t$;
2. The distribution $g_t(a, y)$ evolves according to the stochastic Kolmogorov forward equation (18);
3. All markets clear at time t.

It is worth noting that temporary equilibrium permits systematic deviations between subjective belief expectations and realized equilibrium prices, relaxing the fixed-point requirement of rational expectations equilibrium. This concept, originally developed by Hicks (1946) and Lindahl (1939) and formalized by Grandmont (1977, 1989), has found renewed application in modern macroeconomic models Farhi and Werning (2019), García-Schmidt and Woodford (2019), Piazzesi and Schneider (2016), and Woodford (2013).

The temporary equilibrium system is characterized by the coupled system of household optimization problems, distributional evolution, and market clearing constraints:

$$\begin{aligned} dv_t =& \max_{c,\omega} \left\{ u(c_t) + \{[\omega_t(\mu_s - r_f) + r_f]a_t + y_t - c_t\} \frac{\partial v}{\partial a} + \frac{1}{2}(\omega_t a_t \sigma_s)^2 \frac{\partial^2 v}{\partial a^2} \right\} dt \\ &+ \left\{ \mu_y \frac{\partial v}{\partial y} + \frac{1}{2}\sigma_y^2 \frac{\partial^2 v}{\partial y^2} + \omega a_t \rho_{sy} \sigma_s \sigma_y \frac{\partial^2 v}{\partial a \partial y} + \tilde{\mu}_k \frac{\partial v}{\partial \tilde{k}} \right\} dt \\ &- \left\{ \omega_t \sigma_s \left(a_t \frac{\partial^2 v_t}{\partial a^2} + \frac{\partial v_t}{\partial a} \right) + \sigma_y \rho_{sy} \frac{\partial^2 v_t}{\partial y \partial a} + \sigma_y \sqrt{1 - \rho_{sy}^2} \frac{\partial^2 v_t}{\partial y^2} \right\} dt \\ &+ \begin{bmatrix} \omega_t a_t \sigma_s \frac{\partial v_t}{\partial a} & \sigma_y \rho_{sy} \frac{\partial v_t}{\partial y} \\ 0 & \sigma_y \sqrt{1 - \rho_{sy}^2} \frac{\partial v_t}{\partial y} \end{bmatrix} d\mathbf{W}_t \end{aligned} \tag{17}$$

$$\begin{aligned} dg_t =& -\left\{ \frac{\partial}{\partial a}\left[\left((\omega_t(\mu_s - r_f) + r_f)\, a_t + y_t - c_t \right) g_t \right] + \frac{\partial}{\partial y}\left(\mu_y g_t \right) + \frac{\partial}{\partial \tilde{k}}\left(\tilde{\mu}_k g_t \right) \right\} dt \\ &+ \left\{ \frac{\partial^2}{\partial a^2}\left(\frac{1}{2}(\omega_t a_t \sigma_s)^2 g_t \right) + \frac{\partial^2}{\partial y^2}\left(\frac{1}{2}\sigma_y^2 g_t \right) + \frac{\partial^2}{\partial a \partial y}\left(\omega a_t \rho_{sy} \sigma_s \sigma_y g_t \right) \right\} dt \\ &- \begin{bmatrix} \frac{\partial}{\partial a}\left(\omega_t a_t \sigma_s g_t \right) & \frac{\partial}{\partial y}\left(\sigma_y \rho_{sy} g_t \right) \\ 0 & \frac{\partial}{\partial y}\left(\sigma_y \sqrt{1 - \rho_{sy}^2} g_t \right) \end{bmatrix} d\mathbf{W}_t. \end{aligned} \tag{18}$$

We establish our analysis around two complementary steady-state concepts that provide analytical benchmarks for the dynamic system.

Definition 2. A *determined steady state* (DSS) is a fixed point of the equilibrium conditions when aggregate volatility vanishes ($\sigma = 0$). The DSS is characterized by the system:

$$\rho v_t = \max_{c,\omega} u(c_t) + \{[\omega_t(\mu_s - r_f) + r_f]a_t + y_t - c_t\} \frac{\partial v}{\partial a} + \mu_y \frac{\partial v}{\partial y} \tag{19}$$

$$0 = -\frac{\partial}{\partial a}\Big[\{[\omega_t(\mu_s - r_f) + r_f]a_t + y_t - c_t\} g_t(a,y)\Big] - \frac{\partial}{\partial y}\Big[\mu_y g_t(a,y)\Big], \tag{20}$$

The DSS serves as the deterministic backbone of our model, providing the baseline equilibrium around which stochastic fluctuations occur.

To implement our BSPDE framework for the household's infinite-horizon problem, we truncate the time horizon at finite T and specify appropriate terminal conditions $v(X_T, g_T, T)$ that capture continuation values beyond T. This specification draws on the concept of stochastic steady states, which represent the long-run distributional properties of the economy under persistent uncertainty.

Following (Schenk-Hoppé & Schmalfuss, 2001), we define points around which the ergodic distribution concentrates as natural candidates for steady-state analysis. Building on (Coeurdacier et al., 2011), we formalize this concept:

Definition 3. A *stochastic steady state* (SSS) is a temporary equilibrium characterized by a time-invariant value function $v^{\mathrm{S}}(X_T, g_T)$ and an invariant distribution $g^{\mathrm{S}}(X_T)$ that solve the equilibrium conditions when the realization of aggregate shocks is zero ($d\mathbf{W}_t = 0$). The SSS represents a fixed point of equilibrium conditions under zero shocks realizations while acknowledging the economy's exposure to future uncertainty.

Assumption 1. Our solution methodology rests on three assumptions regarding terminal value:

1. As $T \to \infty$, the economic system is ergodic and its distribution converges to the stochastic steady state distribution $g^{\mathrm{S}}(X_T)$.

2. The economy faces persistent aggregate shocks $\mathbf{W}_t$, preventing convergence to the DSS. The SSS value function $v^{\mathrm{S}}(X_T, g_T)$ incorporates the risk premia associated with this ongoing uncertainty.

3. The value function satisfies the standard transversality condition (Merton, 1969):

$$\lim_{T\to\infty} \mathbb{E}_t\left[e^{-\rho(T-t)} v(X_T, T)\right] = 0.$$

These assumptions ensure that our finite-horizon approximation captures the essential long-run properties of the infinite-horizon problem. For any state (X, g), the terminal condition for our BSPDE is specified as the SSS value function

$$v(X_T, g_T, \mathbf{W}_T) = v^{\mathrm{S}}(X_T, g_T) \quad \text{as } T \to \infty,$$

where $v^{\mathrm{S}}(X_T, g_T)$ solves the stationary Hamilton-Jacobi-Bellman equation incorporating the endogenous aggregate risk aversion term:

$$\rho v^{\mathrm{S}}(X_T, g_T) = \max_{c,\omega} \left\{ u(c) + \mathcal{L}[v^{\mathrm{S}}(X_T, g_T)] - \mathrm{div}(\mathcal{V}_t) \right\}.$$

Crucially, this specification sets the terminal value to the long-run equilibrium value function that has fully internalized the cost of aggregate risk through the EARA term $-\mathrm{div}(\mathcal{V}_t)$, thereby providing a consistent anchor for the household's optimization problem.

This terminal condition specification embodies several key economic principles. First, it aligns with the fundamental macroeconomic premise that economies exhibit mean-reverting behavior over sufficiently long horizons. While temporary deviations from steady state significantly impact short-to-medium term dynamics, their influence on decision-making diminishes as the planning horizon extends. Consequently, the selection of SSS as the anchor is deemed suitable, as it represents a state of equilibrium that has accurately reflected the long-term, persistent overall risk.

Second, our approach captures the bounded rationality assumption that households form long-term expectations based on the economy's convergence to its stochastic steady state. The random fluctuation term $\mathcal{V}_t d\mathbf{W}_t$ was eliminated because, when defining the SSS, the moment T when the realized shocks were exactly zero was examined. Nevertheless, this does not imply a disregard of the potential risks.

Third, the specification maintains internal consistency with the optimization problem's structure. By anchoring terminal behavior to the SSS value function, $v^{\mathrm{S}}(X_T, g_T)$, we ensure that the solution to the finite-horizon problem correctly inherits the continuation value that fully prices the economy's persistent exposure to aggregate risk. While the stochastic term $\mathcal{V}_t d\mathbf{W}_t$ is absent from the stationary SSS equation (as it describes a fixed point conditional on zero shocks realizations), the EARA in the value function, $-\mathrm{div}(\mathcal{V}_t)$, remains a crucial component. This term permanently adjusts the value function to reflect the ongoing uncertainty, and its inclusion in the terminal condition is essential for the theoretical coherence of our solution.

This methodology provides a computationally tractable approach to solving the household problem while maintaining the essential economic content of the infinite-horizon framework. The resulting solution captures both the short-run dynamics driven by aggregate uncertainty and the long-run convergence properties that discipline agent expectations.

4 Theoretical Results for the Model Under Common Shocks: Portfolio and Inequality

The household's optimization, characterized by the BSPDE (5), yields first-order conditions for consumption and portfolio allocation. The optimal consumption rule remains standard:

$$c^* = (u')^{-1}\left(\frac{\partial v}{\partial a}\right) \tag{21}$$

The optimal portfolio choice, however, is fundamentally altered by the presence of common shocks. Taking the first-order condition with respect to the portfolio share ω_t in the BSPDE (5) yields:

$$\omega_t^* = -\frac{1}{a_t}\left(\frac{\partial^2 v}{\partial a^2}\right)^{-1}\left[\left(\frac{\mu_s - r_f}{\sigma_s^2}\right)\frac{\partial v}{\partial a} + \left(\frac{\rho_{sy}\sigma_y}{\sigma_s}\right)\frac{\partial^2 v}{\partial a \partial y} - \frac{1}{\sigma_s}\left(\frac{\partial^2 v_t}{\partial a^2} + \frac{1}{a_t}\frac{\partial v_t}{\partial a}\right)\right], \tag{22}$$

The final term in the square brackets arises directly from the household's optimization over the endogenous aggregate risk aversion term, $\text{div}(\mathcal{V}_t)$, in their value function. It represents the household's strategic response to the endogenous risk generated by the mean-field interaction.

To isolate this new economic force, we compare equation (22) with the portfolio choice from a classical HJB framework that omits the mean-field terms:

$$\omega_t^C = -\frac{1}{a_t}\left(\frac{\partial^2 v}{\partial a^2}\right)^{-1}\left[\left(\frac{\mu_s - r_f}{\sigma_s^2}\right)\frac{\partial v}{\partial a} + \left(\frac{\rho_{sy}\sigma_y}{\sigma_s}\right)\frac{\partial^2 v}{\partial a \partial y}\right]. \tag{23}$$

The difference between ω_t^* and ω_t^C is driven entirely by the household's response to feedback from future common shocks, which arise from interactions between optimization and distributional evolution. We formalize this novel mechanism as follows:

Definition 4. The *endogenous aggregate risk aversion strategic adjustment* is the component of the optimal portfolio rule that arises from the controllable part of the EARA (9):

$$\Omega_t^*(a_t) \equiv \frac{1}{a_t \sigma_s}\left(\frac{\partial^2 v}{\partial a^2}\right)^{-1}\left(\frac{\partial^2 v_t}{\partial a_t^2} + \frac{1}{a_t}\frac{\partial v_t}{\partial a}\right). \tag{24}$$

The optimal portfolio can thus be expressed as the sum of classical demand and this new strategic adjustment: $\omega_t^* = \omega_t^C + \Omega_t^*$. The EARA strategic adjustment captures how agents actively manage their portfolio to hedge against future common shocks in their own value function. To characterize this mechanism more precisely, we introduce the concept of endogenous relative risk aversion:

Definition 5. The *endogenous relative risk aversion* (ERRA) coefficient is:

$$\eta_t(a) \equiv -a\frac{v_{aa}}{v_a}, \tag{25}$$

where $v_a \equiv \partial v/\partial a$ and $v_{aa} \equiv \partial^2 v/\partial a^2$.

The ERRA $\eta(a)$ measures the household's endogenous aversion to wealth fluctuations in state (a, y), incorporating anticipated responses to future common shocks. Under standard regularity conditions—positive wealth ($a > 0$), positive marginal utility ($v_a > 0$), and concave value function ($v_{aa} < 0$)—we have $\eta_t > 0$.

Substituting the ERRA $\eta(a)$ definition into (24) yields:

$$\Omega_t^* = \frac{1}{a_t \sigma_s} \frac{\eta_t - 1}{\eta_t}. \tag{26}$$

This representation illuminates three crucial properties of the EARA strategic adjustment. First, its magnitude scales inversely with wealth ($1/a$), generating stronger effects for lower-wealth households. Second, higher asset volatility (σ_s) dampens the strategic adjustment's impact. Third, and most importantly, the sign and magnitude of Ω_t^* depend critically on whether ERRA $\eta(a)$ exceeds, equals, or falls below unity.

When $\eta_t = 1$, the EARA strategic adjustment vanishes, and portfolio decisions coincide with classical results (23). This endogenous threshold is crucial to understanding household participation in financial markets.

When $\eta_t > 1$, the strategic adjustment becomes positive, inducing households to increase risky asset holdings—reflecting prudent, forward-looking behavior under common shocks. However, as wealth a increases, two effects dominate. First, the $1/a$ scaling factor ensures the quantitative impact of the EARA adjustment Ω_t^* diminishes rapidly. Second, from an economic standpoint, as financial wealth grows, non-tradable assets like labor income become negligible in relative terms. The household's problem converges to a pure asset allocation problem, where complex hedging demands fade. Consequently, the portfolio allocation converges to the classical Merton baseline, as demonstrated in Corollary 2.

Conversely, when $\eta_t < 1$, the strategic adjustment turns negative, reducing risky asset exposure through a form of endogenous risk aversion distinct from preference-based risk aversion. In this situation, the EARA adjustment Ω_t^* is amplified by both the wealth factor $1/a$ and the ERRA factor $1/\eta_t$, creating a powerful negative force. The economic intuition for this strong effect is a "marginal wealth preservation demand": households with minimal financial wealth prioritize the preservation of their limited buffer stock, making them extremely cautious about exposing it to aggregate risk.

To relate our theoretical results to empirically observable variables, we exploit the envelope condition (21) and derive:

$$v_{aa} = u''(c_t^*) \frac{\partial c_t^*}{\partial a_t} \tag{27}$$

and utilize the standard CRRA specification:

$$\gamma = -c\frac{u''(c_t^*)}{u'(c_t^*)}. \tag{28}$$

Substituting equations (27) and (28) into (25) yields:

$$\begin{aligned}\eta_t(a) =& - a \cdot \frac{v_{aa}}{u'(c_t^*)} \\ =& - a \cdot \frac{\partial c_t^*}{\partial a_t} \cdot \frac{u''(c_t^*)}{u'(c_t^*)} \\ =& \frac{a_t}{c_t^*} \cdot \mathrm{MPC}_a \cdot \gamma,\end{aligned} \tag{29}$$

where $\mathrm{MPC}_a \equiv \partial c_t^*/\partial a_t$ denotes the marginal propensity to consume out of wealth.

Equation (29) reveals that the consumption elasticity with respect to wealth, $\epsilon_{c,a} = (a_t/c_t^*) \cdot \mathrm{MPC}_a$, is not a free parameter but an endogenous outcome determined by the household's optimization under constraints, income dynamics, and risk exposure. Thus, deviations of ERRA η_t from CRRA γ arise entirely from the state-dependent elasticity $\epsilon_{c,a}$, which mediates how microeconomic consumption preferences translate into macroscopic portfolio responses under common shocks.

This framework simultaneously resolves fundamental limitations inherent in continuous-time models, specifically the Aiyagari-Bewley-Huggett (ABH) framework, and within portfolio theory. A significant challenge in ABH economiesis that optimal portfolio allocation in CTHA models without common shocks typically yields $\omega^* = 0$. This leads to the undesirable outcome of wealth distributions degenerating to mass points at borrowing constraints, as documented by (Achdou et al., 2022; Ahn et al., 2018; Fernández-Villaverde et al., 2023).

Our theory addresses this degenerate equilibrium by introducing a risk-bearing asset. Furthermore, even when labor income processes are purely idiosyncratic ($\rho_{sy} = 0$) or entirely free of risk ($\sigma_y = 0$), the strategic adjustment component of the EARA facilitates empirically consistent limited participation at the asset lower bound. This innovation enables the generation of realistic wealth distributions without the need for arbitrary preference heterogeneity, a common reliance in existing models.

4.1 ERRA and Portfolio Choice Near the Borrowing Constraint

We now examine portfolio behavior as households approach their borrowing constraint $\underline{a}$. For agents with wealth a near this lower bound, consumption c^* is determined primarily by labor income y rather than financial wealth, implying $a/c^* \to 0$ as $a \to \underline{a}$.

Proposition 1 (**ERRA at the borrowing constraint**)**.** For any household with finite constant relative risk aversion $\gamma < \infty$ and bounded marginal propensity to consume

$\text{MPC}_a < \infty$, the endogenous relative risk aversion vanishes at the borrowing constraint:

$$\lim_{a \to 0^+} \eta_t(a) = 0. \tag{30}$$

The economic intuition is that as financial wealth becomes negligible relative to human capital, consumption elasticity $\epsilon_{c,a}$ must approach zero, meaning that marginal changes in wealth have a negligible relative impact on total consumption.

Corollary 1 (**Portfolio choice at the borrowing constraint**)**.** The EARA strategic adjustment exhibits singular behavior at the borrowing constraint:

$$\lim_{a \to 0^+} \Omega_t^* = -\infty, \tag{31}$$

leading to optimal portfolio choice:

$$\lim_{a \to 0^+} \omega_t^* = \begin{cases} 0 & \text{under short-selling restrictions} \\ -\infty & \text{without short-selling restrictions.} \end{cases} \tag{32}$$

This result reveals a previously unrecognized mechanism driving portfolio decisions of wealth-constrained households. As wealth approaches the constraint, an endogenous force, unrelated to risk preferences or return expectations, drives these households to minimize risky asset exposure, potentially requiring substantial short positions when institutionally feasible.

This finding contrasts sharply with predictions from conventional Merton models, where households facing positive risk premia should maintain significant risky asset exposure near borrowing constraints unless counteracted by strongly positive correlation between labor income and asset returns. However, empirical evidence documents widespread "limited participation"—households with minimal financial wealth rarely participate in equity markets (Campbell, 2006; Vissing-Jørgensen, 2002). Our model generates this pattern endogenously without requiring controversial parametrizations of the correlation ρ_{sy}.

The underlying economic mechanism operates through "marginal wealth preservation demand." When financial wealth is minimal, households focus intensely on preserving their limited buffer against fluctuations. Under common shocks, this preservation motive manifests as extreme caution toward risky assets, with the mathematical divergence in Ω_t overwhelming conventional portfolio demands as $a \to \underline{a}$.

4.2 ERRA and Portfolio Allocation at High Wealth Levels

Having characterized behavior near borrowing constraints, we now analyze the limiting case as wealth grows arbitrarily large. Consistent with prior literature, which suggests

consumption becomes asymptotically proportional to wealth under CRRA preferences with idiosyncratic shocks (Achdou et al., 2022; Benhabib et al., 2015), we formalize this convergence.

Theorem 1 (Asymptotic Consumption-Wealth Ratio)**.** Let $c^*(a, y, \tilde{k}, t)$ be the optimal consumption policy derived from the solution to the BSPDE (17). Under Assumption 2, the consumption-to-wealth ratio converges in the mean-square sense to a deterministic constant $\bar{k}_{\mathrm{S}}$ as wealth a goes to infinity:

$$\lim_{a\to\infty} \mathbb{E}\left[\left|\frac{c^*(a, y, \tilde{k}, t)}{a} - \bar{k}_{\mathrm{S}}\right|^2\right] = 0,$$

where $\bar{k}_{\mathrm{S}}$ is the optimal consumption-wealth ratio in the SSS optimal problem without labor income or aggregate state dependence, given by:

$$\bar{k}_{\mathrm{S}} = \frac{1}{\gamma}\left(\rho - (1-\gamma)\left(r_f + \frac{(\mu_s - r_f)^2}{2\gamma\sigma_s^2}\right)\right). \tag{33}$$

Proof. See Appendix for the proof regarding the asymptotic behavior of the consumption-wealth ratio. □

Proposition 2 (**ERRA convergence for wealthy households**)**.** For households with positive CRRA coefficient $\gamma > 0$ and satisfing Theorem 1 assumption, the endogenous relative risk aversion converges to the CRRA parameter:

$$\lim_{a\to\infty} \eta_t(a) = \gamma. \tag{34}$$

Proof. Theorem 1 implies that as $a \to \infty$, the consumption function becomes linear in wealth, $c^*(a) \approx \bar{k}_{\mathrm{S}} a$. Thus, the marginal propensity to consume converges to the average propensity to consume:

$$\lim_{a\to\infty} \frac{\partial c^*(a)}{\partial a} = \lim_{a\to\infty} \frac{c^*(a)}{a} = \bar{k}_{\mathrm{S}}.$$

Therefore, the elasticity of consumption with respect to wealth converges to unity:

$$\lim_{a\to\infty} \frac{a}{c^*(a)}\frac{\partial c^*(a)}{\partial a} = \lim_{a\to\infty} \left(\frac{c^*(a)}{a}\right)^{-1} \frac{\partial c^*(a)}{\partial a} = (\bar{k}_{\mathrm{S}})^{-1}\bar{k}_{\mathrm{S}} = 1.$$

Substituting this limit back into the expression for ERRA $\eta(a)$, we obtain the desired result:

$$\lim_{a\to\infty} \eta_t(a) = \gamma \cdot 1 = \gamma.$$

□

Corollary 2 (**Portfolio choice convergence at high wealth levels**). As wealth increases without bound, the EARA strategic adjustment vanishes and portfolio allocation converges to the classical baseline ratio:

$$\lim_{a\to\infty} \Omega_t(a) = 0, \qquad \lim_{a\to\infty} \omega_t^* = \frac{\mu_s - r_f}{\sigma_s^2} \cdot \frac{1}{\gamma}. \tag{35}$$

Proof. The optimal portfolio allocation, ω_t^*, consists of a myopic demand component and a strategic hedging component, $\Omega_t(a)$:

$$\omega_t^*(a) = \frac{\mu_s - r_f}{\sigma_s^2 \eta_t(a)} + \Omega_t(a).$$

The strategic component, $\Omega_t(a)$, hedges against fluctuations in the investment opportunity set and future labor income. As wealth a becomes large, the present value of future labor income becomes negligible relative to financial wealth. Consequently, the hedging motive diminishes, and $\Omega_t(a)$ converges to zero. This is reflected by the wealth-scaling factor in its analytical expression (see Equation (26)), which ensures $\lim_{a\to\infty} \Omega_t(a) = 0$.

For the total allocation, we take the limit as $a \to \infty$:

$$\lim_{a\to\infty} \omega_t^*(a) = \lim_{a\to\infty} \left(\frac{\mu_s - r_f}{\sigma_s^2 \eta_t(a)} \right) + \lim_{a\to\infty} \Omega_t(a).$$

Using the result from the Proposition 2 that $\lim_{a\to\infty} \eta_t(a) = \gamma$ and that the hedging demand vanishes, we find that the portfolio converges to the classic Merton (1969) share:

$$\lim_{a\to\infty} \omega_t^* = \frac{\mu_s - r_f}{\sigma_s^2 \gamma}.$$

□

This convergence reflects the gradual dominance of financial asset returns as the relative importance of human capital income declines. Households need to maintain a certain amount of risk assets to meet their total income needs, even when considering future shocks. The wealth-scaling factor $(1/a)$ in equation (26) ensures quantitative negligibility of the EARA strategic adjustment regardless of finite deviations of ERRA $\eta(a)$ from unity.

The economic implication is profound: portfolio heterogeneity should be concentrated among low-to-moderate wealth households, where EARA strategic adjustment significantly modify allocations relative to standard theory. For wealthy households, conventional portfolio rules provide reasonable approximations, as the wealth-scaling effect renders the EARA strategic adjustment negligible.

These results suggest that wealth distribution dynamics will be driven primarily by standard portfolio mechanisms at the upper tail but by novel EARA strategic adjustment mechanisms at the lower tail, creating endogenous heterogeneity in responses to aggregate shocks across the wealth distribution and potentially explaining differential sensitivity to macroeconomic conditions between wealthy and non-wealthy households.

5 Quantitative Analysis of the EARA Mechanism

To translate our theoretical framework into quantitative predictions, we parameterize the model and visualize the core economic mechanisms. Our parameterization is disciplined by the theoretical properties derived in Section 4. Specifically, we adopt a parsimonious functional form for the household consumption rule, $c(a)$, that is consistent with the theoretical properties derived in the preceding sections. We adopt the form $c(a) = (c_0^p + (k \cdot a)^p)^{1/p}$, which ensures that consumption is sustained by a base level c_0 at low wealth and becomes asymptotically proportional to wealth ($c(a) \approx k \cdot a$) for affluent households, consistent with Theorem 1. This formulation generates a smooth, monotonic consumption elasticity $\epsilon_{c,a}(a)$, which in turn ensures that the ERRA, $\eta(a) = \gamma \cdot \epsilon_{c,a}(a)$, adheres to the boundary conditions established in Propositions 1 and 2. The subsequent analysis illustrates how this endogenous risk aversion shapes portfolio decisions across the wealth distribution.

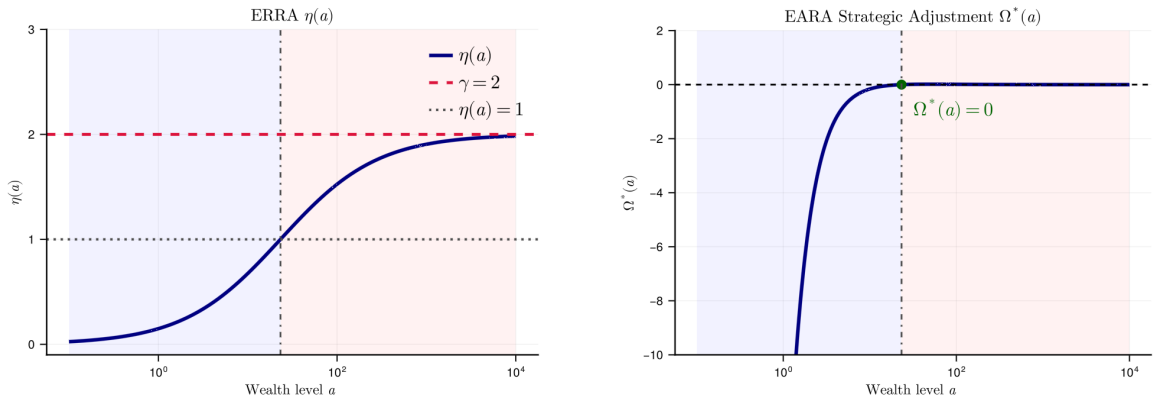
For our baseline simulations, we set the model parameters as follows: the maximum marginal propensity to consume for borrowing-constrained households, $\text{MPC}_{\max}$, is 0.72; the decay rate p is 0.8; the coefficient of CRRA, γ, is 2; the time discount rate, ρ, is 0.05. We assume an annualized expected return for risky assets, μ_s, of 8%, and a risk-free rate, r_f, of 3%. The mean income, y, is normalized to 1.0. Asset volatilities are set at $\sigma_s = 0.30$ for risky assets and $\sigma_y = 0$ for income. For hand-to-mouth households at the lower end of the asset spectrum, consumption c_0 equals their income.

5.1 EARA and its Implications

The cornerstone of our mechanism is the state-dependent nature of households' EARA. Figure 1a plots the ERRA, $\eta(a)$, as a function of household wealth, a. The figure starkly illustrates the theoretical predictions of our model.

First, consistent with Proposition 1, as wealth approaches the borrowing constraint ($a \to 0$), ERRA $\eta(a)$ converges to zero. This occurs because consumption is almost entirely determined by non-financial income as same as wealth is small, making the marginal utility of financial wealth nearly constant with respect to small changes in a. Second, in line with Proposition 2, as wealth becomes large, $\eta(a)$ asymptotically converges to the underlying preference parameter, γ, which is set to 2 in this calibration.

It is worth noting that the transition of $\eta(a)$ across the critical value of unity dictates the sign and impact of the EARA strategic adjustment, $\Omega^*(a) \propto (1 - 1/\eta(a))$. As shown in Figure 1a, for households with wealth below a certain threshold, $\eta(a) < 1$, inducing a negative adjustment that powerfully suppresses demand for risky assets. Conversely, for wealthier households with $\eta(a) > 1$, the adjustment becomes positive, reflecting a



(a) (b)

Figure 1

endogenously hedging behavior against aggregate shocks.

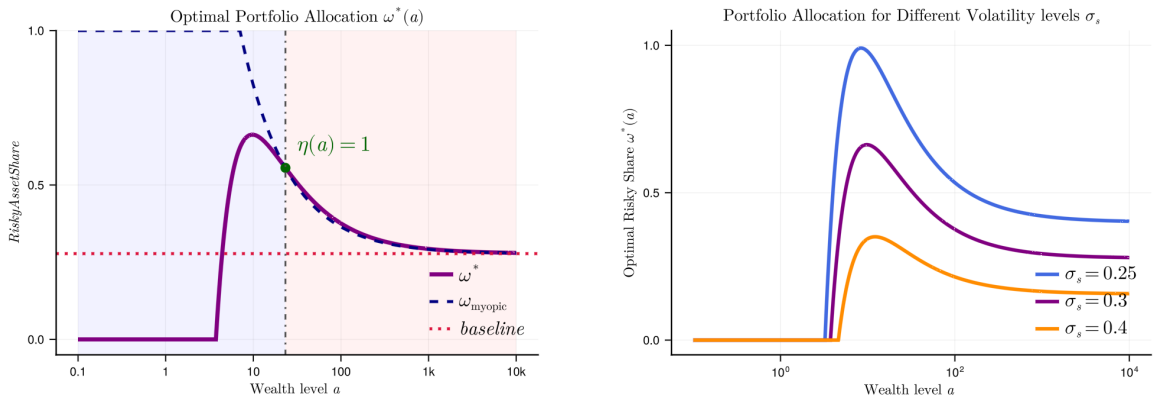
Figure 1b isolates this EARA strategic adjustment term, $\Omega^*(a)$. The singular nature of this term at low wealth levels is apparent. As $a \to 0$, $\eta(a) \to 0$, causing $\Omega^*(a)$ to diverge to $-\infty$. This provides the powerful endogenous force for non-participation in risky assets market. For wealthier households, as $\eta(a)$ approaches γ, the term $(1-1/\eta(a))$ stabilizes. However, the presence of the $1/a$ scaling factor in the definition of $\Omega^*(a)$ ensures that the adjustment term itself vanishes, approaching zero from above.

5.2 Optimal Portfolio Allocation and Endogenous Limited Participation

The state-dependent ERRA $\eta(a)$, channeled through the EARA mechanism, produces a rich pattern of portfolio holdings across the wealth distribution. Figure 2a depicts the optimal risky asset share, $\omega^*(a)$, and contrasts it with the benchmark: the classical Merton (1969, 1973) portfolio choice problem based on no risk in labor income.

The results provide a clear visualization of our central theoretical claims. At low levels of wealth, the EARA strategic adjustment $\Omega^*(a)$ is strongly negative, as shown in Figure 2a. This force overwhelms the myopic demand, driving the optimal allocation $\omega^*(a)$ to the zero lower bound. This provides a novel, endogenous rationale for the empirically documented phenomenon of limited stock market participation among low-wealth households, a direct consequence of Corollary 1. Our model generates this result without imposing exogenous constraints or assuming extreme correlations between labor income and asset returns.

As wealth increases, $\eta(a)$ surpasses one, and the EARA adjustment term becomes positive, leading to a distinctive "hump" shape in the optimal portfolio. For the



(a) (b)

Figure 2

wealthiest households, two effects ensure convergence to the baseline ratio in Corollary 2. First, the EARA strategic adjustment term $\Omega^*(a)$ vanishes due to the $1/a$ scaling factor. Second, $\eta(a)$ converges to γ. Consequently, as predicted by Corollary 2, the optimal portfolio share $\omega^*(a)$ robustly converges to the baseline ratio, demonstrating that standard portfolio theory serves as a good approximation for the affluent but fundamentally mischaracterizes the behavior of the less wealthy facing common shock in risky assets.

Figure 2b further explores the properties of the optimal portfolio allocation by presenting a comparative statics exercise with respect to the volatility of risky assets, σ_s. An increase in the volatility unambiguously lowers the demand for risky assets across all wealth levels, a standard result. Note that the figure illustrates a more subtle point: the wealth threshold at which households begin to participate in risky assets market increases with volatility. This is because higher volatility amplifies the cautionary effect of the EARA mechanism at low wealth levels, making households even more reluctant to hold risky assets.

5.3 Implications for Wealth Dynamics

Finally, we examine how these wealth-dependent portfolio choices translate into wealth accumulation dynamics. Figure 3 plots the expected instantaneous change in wealth, or the wealth drift function, $\mu_a(a) = \mathbb{E}[da]/dt$. The drift function is negative for a significant range of low to moderate wealth levels, indicating an endogenous force pushing these households towards the borrowing constraint. This "poverty trap" dynamic is a direct result of the limited participation induced by the EARA mechanism; by holding only the risk-free asset, these households have lower expected returns.

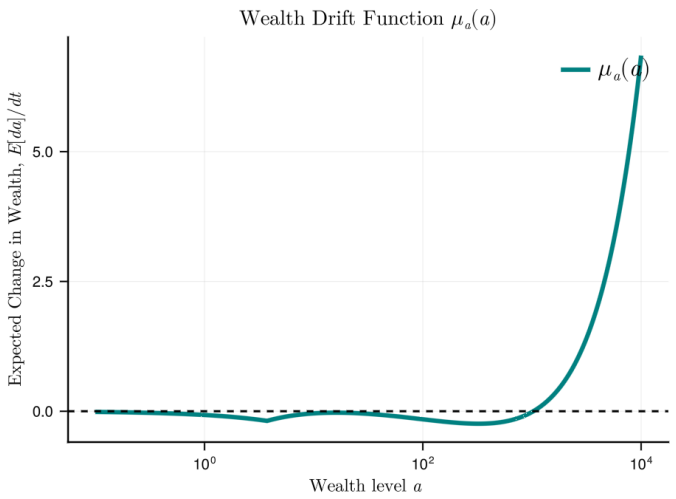


Figure 3

However, at very high wealth levels, the portfolio share converges to the baseline ratio, and the wealth drift becomes strongly positive and convex. This convexity implies that the wealthiest households accumulate wealth at an accelerating rate, creating a powerful mechanism for the emergence of a fat-tailed wealth distribution. The interplay between the poverty trap at the lower assets and the explosive accumulation at the upper assets, both driven by the endogenous portfolio responses to common shocks, provides a novel channel for generating and sustaining wealth inequality.

6 Conclusion

This paper confronts a foundational challenge in modern macroeconomics: the difficulty of constructing models that simultaneously generate realistic wealth inequality and endogenous portfolio heterogeneity in the presence of aggregate risk. We demonstrate that moving from the standard HJB framework to a mean field game with common shocks provides a resolution. This approach, characterized by a coupled system of a BSPDE and an SKF equation, reveals a previously unrecognized mechanism in portfolio choice: the EARA Strategic Adjustment.

Our central theoretical result is that this EARA Strategic Adjustment, which arises purely from the interaction between individual optimization and aggregate dynamics, fundamentally reshapes household risk-taking in a wealth-dependent manner. For households near borrowing constraints, the mechanism generates extreme endogenous relative risk aversion, providing a novel, endogenous explanation for the widespread empirical phenomenon of limited stock market participation. For the wealthiest households, this effect vanishes, and their behavior converges to the predictions of classical portfolio

theory. This differential impact offers a new channel through which aggregate shocks drive wealth distribution dynamics, allowing the model to sustain significant inequality even under standard preference assumptions.

The implications of this framework are broad. For economic theory, it bridges the gap between heterogeneous agent macroeconomics and household finance, establishing the feedback from aggregate risk as a first-order determinant of cross-sectional portfolio choice. For policy analysis, the EARA channel suggests that policies affecting the wealth distribution—such as wealth taxes or social safety nets—will have crucial general equilibrium effects on aggregate risk-taking and financial stability, a margin absent from conventional models.

The primary next step for this research agenda is a full quantitative evaluation. Embedding the EARA mechanism within a real business cycle model and solving the system with advanced numerical methods will allow for a disciplined comparison with U.S. data. A forthcoming version of this paper will quantify the importance of this new channel in matching wealth and portfolio distributions and will explore counterfactual policy experiments.

References

Achdou, Y., Cardaliaguet, P., Delarue, F., Porretta, A., & Santambrogio, F. (2021). *Mean field games: Cetraro, italy 2019* (Vol. 2281). Springer Nature.

Achdou, Y., Han, J., Lasry, J.-M., Lions, P.-L., & Moll, B. (2022). Income and wealth distribution in macroeconomics: A continuous-time approach. *The review of economic studies*, *89*(1), 45–86.

Ahn, S., Kaplan, G., Moll, B., Winberry, T., & Wolf, C. (2018). When inequality matters for macro and macro matters for inequality. *NBER macroeconomics annual*, *32*(1), 1–75.

Aiyagari, S. R. (1994). Uninsured idiosyncratic risk and aggregate saving. *The Quarterly Journal of Economics*, *109*(3), 659–684.

Anderson, E. W., Ghysels, E., & Juergens, J. L. (2005). Do heterogeneous beliefs matter for asset pricing? *The Review of Financial Studies*, *18*(3), 875–924.

Benhabib, J., Bisin, A., & Zhu, S. (2015). The wealth distribution in bewley economies with capital income risk. *Journal of Economic Theory*, *159*, 489–515.

Bewley, T. (1986). Stationary monetary equilibrium with a continuum of independently fluctuating consumers. In W. Hildenbrand & A. Mas-Colell (Eds.), *Contributions to mathematical economics in honor of gerard debreu*. North-Holland.

Bhamra, H. S., & Uppal, R. (2014). Asset prices with heterogeneity in preferences and beliefs. *The Review of Financial Studies*, *27*(2), 519–580.

Bilal, A. (2023). *Solving heterogeneous agent models with the master equation* (tech. rep.). National Bureau of Economic Research.

Bilal, A., & Goyal, S. (2025). *Some pleasant sequence-space arithmetic in continuous time* (tech. rep.). National Bureau of Economic Research.

Campbell, J. Y. (2006). Household finance. *The journal of finance*, *61*(4), 1553–1604.

Cardaliaguet, P., Delarue, F., Lasry, J.-M., & Lions, P.-L. (2019). *The master equation and the convergence problem in mean field games:(ams-201)*. Princeton University Press.

Carmona, R., Delarue, F., et al. (2018). *Probabilistic theory of mean field games with applications i-ii*. Springer.

Coeurdacier, N., Rey, H., & Winant, P. (2011). The risky steady state. *American Economic Review*, *101*(3), 398–401.

De Long, J. B., Shleifer, A., Summers, L. H., & Waldmann, R. J. (1990). Noise trader risk in financial markets. *Journal of political Economy*, *98*(4), 703–738.

Dend Haan, W. J. (1996). Heterogeneity, aggregate uncertainty, and the short-term interest rate. *Journal of Business & Economic Statistics*, *14*(4), 399–411.

Evans, G. W., & Honkapohja, S. (2012). Learning and expectations in macroeconomics. In *Learning and expectations in macroeconomics*. Princeton University Press.

Farhi, E., & Werning, I. (2019). Monetary policy, bounded rationality, and incomplete markets. *American Economic Review*, *109*(11), 3887–3928.

Fernández-Villaverde, J., Hurtado, S., & Nuno, G. (2023). Financial frictions and the wealth distribution. *Econometrica*, *91*(3), 869–901.

Fernández-Villaverde, J., Nuño, G., & Perla, J. (2024). *Taming the curse of dimensionality: Quantitative economics with deep learning* (tech. rep.). National Bureau of Economic Research.

Gabaix, X. (2014). A sparsity-based model of bounded rationality. *The Quarterly Journal of Economics*, *129*(4), 1661–1710.

García-Schmidt, M., & Woodford, M. (2019). Are low interest rates deflationary? a paradox of perfect-foresight analysis. *American Economic Review*, *109*(1), 86–120.

Grandmont, J. M. (1977). Temporary general equilibrium theory. *Econometrica: Journal of the Econometric Society*, 535–572.

Grandmont, J. M. (1989). Temporary equilibrium. In *General equilibrium* (pp. 297–304). Springer.

Han, J., Jentzen, A., & E, W. (2018). Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, *115*(34), 8505–8510.

Hicks, J. R. (1946). Value and capital: An inquiry into some fundamental principles of economic theory. *OUP Catalogue*.

Huggett, M. (1993). The risk-free rate in heterogeneous-agent incomplete-insurance economies. *Journal of economic Dynamics and Control*, *17*(5-6), 953–969.

Javed, K., & Sutton, R. S. (2024). The big world hypothesis and its ramifications for artificial intelligence. *Finding the Frame: An RLC Workshop for Examining Conceptual Frameworks*.

Jouini, E., & Napp, C. (2007). Consensus consumer and intertemporal asset pricing with heterogeneous beliefs. *The Review of Economic Studies*, *74*(4), 1149–1174.

Krusell, P., & Smith, A. A., Jr. (1998). Income and wealth heterogeneity in the macroeconomy. *Journal of political Economy*, *106*(5), 867–896.

Kunita, H. (1990). *Stochastic flows and stochastic differential equations* (Vol. 24). Cambridge university press.

Lasry, J.-M., & Lions, P.-L. (2007). Mean field games. *Japanese journal of mathematics*, *2*(1), 229–260.

Lindahl, E. (1939). *Theory of money and capital*. Allen; Unwin; Holt, Rinehart; Winston.

Maćkowiak, B., & Wiederholt, M. (2009). Optimal sticky prices under rational inattention. *American Economic Review*, *99*(3), 769–803.

Maćkowiak, B., & Wiederholt, M. (2015). Business cycle dynamics under rational inattention. *The Review of Economic Studies*, *82*(4), 1502–1532.

Malmendier, U., & Nagel, S. (2016). Learning from inflation experiences. *The Quarterly Journal of Economics*, *131*(1), 53–87.

Merton, R. C. (1969). Lifetime portfolio selection under uncertainty: The continuous-time case. *The review of Economics and Statistics*, 247–257.

Merton, R. C. (1973). An intertemporal capital asset pricing model. *Econometrica: Journal of the Econometric Society*, 867–887.

Moll, B. (2024). The trouble with rational expectations in heterogeneous agent models: A challenge for macroeconomics. *London School of Economics, mimeo.*

Peng, S. (1992). Stochastic hamilton–jacobi–bellman equations. *SIAM Journal on Control and Optimization*, *30*(2), 284–304.

Piazzesi, M., & Schneider, M. (2016). Housing and macroeconomics. *Handbook of macroeconomics*, *2*, 1547–1640.

Rachev, S. T., & Rüschendorf, L. (2006). *Mass transportation problems: Volume 1: Theory*. Springer Science & Business Media.

Schenk-Hoppé, K. R., & Schmalfuss, B. (2001). Random fixed points in a stochastic solow growth model. *Journal of Mathematical Economics*, *36*(1), 19–30.

Shiller, R. J., Fischer, S., & Friedman, B. M. (1984). Stock prices and social dynamics. *Brookings papers on economic activity*, *1984*(2), 457–510.

Vissing-Jørgensen, A. (2002). Limited asset market participation and the elasticity of intertemporal substitution. *Journal of political Economy*, *110*(4), 825–853.

Woodford, M. (2013). Macroeconomic analysis without the rational expectations hypothesis. *Annu. Rev. Econ.*, *5*(1), 303–346.

A. Proofs of Main Results at the Borrowing Constraint

Proof of Proposition 1

Proof. The endogenous relative risk aversion (ERRA), $\eta_t(a)$, is defined as:

$$\eta_t(a) = \frac{a_t}{c_t^*}\frac{\partial c}{\partial a}\gamma. \tag{36}$$

To prove the proposition, we analyze the limit of each component of the product in (36) as wealth a_t approaches the borrowing constraint from above, which we take to be $a_t \to 0^+$.

The assumptions stated in the proposition are:

1. The coefficient of constant relative risk aversion, γ, is a finite positive constant.
2. The marginal propensity to consume out of wealth, $\frac{\partial c}{\partial a} \equiv \frac{\partial c_t^*}{\partial a_t}$, is bounded, i.e., there exists some $M < \infty$ such that $\frac{\partial c}{\partial a} \leq M$ as $a_t \to 0^+$.

Furthermore, we consider a household with a positive stream of non-financial income, $y_t > 0$. As the household's financial wealth a_t approaches zero, consumption does not fall to zero. Instead, it is sustained by labor income. Thus, optimal consumption c_t^* converges to a strictly positive value, $c^*(0, y_t) > 0$. Consequently, the limit of the wealth-to-consumption ratio is:

$$\lim_{a_t \to 0^+} \frac{a_t}{c_t^*} = \frac{\lim_{a_t \to 0^+} a_t}{\lim_{a_t \to 0^+} c_t^*} = \frac{0}{c^*(0, y_t)} = 0.$$

We can now evaluate the limit of $\eta_t(a)$:

$$\begin{aligned}
\lim_{a_t \to 0^+} \eta_t(a) &= \lim_{a_t \to 0^+} \left(\frac{a_t}{c_t^*}\frac{\partial c}{\partial a}\gamma \right) \\
&= \left(\lim_{a_t \to 0^+} \frac{a_t}{c_t^*} \right) \cdot \left(\lim_{a_t \to 0^+} \frac{\partial c}{\partial a} \right) \cdot \gamma \\
&= (0) \cdot (\text{bounded value}) \cdot (\text{finite constant}) \\
&= 0.
\end{aligned}$$

This establishes that the endogenous relative risk aversion vanishes as the household's wealth approaches the borrowing constraint. □

Proof of Corollary 1

Proof. The proof is presented in two parts, corresponding to the two limits stated in the corollary.

The EARA term is given by equation (26):

$$\Omega_t^*(a_t) = \frac{1}{a_t \sigma_s} \frac{\eta_t(a_t) - 1}{\eta_t(a_t)} = \frac{1}{a_t \sigma_s} \left(1 - \frac{1}{\eta_t(a_t)}\right). \tag{37}$$

We analyze its limit as $a_t \to 0^+$. From Proposition 1, we have $\lim_{a_t \to 0^+} \eta_t(a_t) = 0$. Since for any $a_t > 0$, wealth is positive, consumption is positive, $\frac{\partial c}{\partial a} > 0$ (for a non-Giffen good), and $\gamma > 0$, it follows that $\eta_t(a_t) > 0$. Therefore, the limit is approached from the right, i.e., $\eta_t(a_t) \to 0^+$.

Let's examine the terms in (37):

1. $\lim_{a_t \to 0^+} \frac{1}{a_t \sigma_s} = +\infty$, assuming asset volatility σ_s is a positive constant.
2. $\lim_{a_t \to 0^+} \left(1 - \frac{1}{\eta_t(a_t)}\right)$. Since $\eta_t(a_t) \to 0^+$, the term $1/\eta_t(a_t) \to +\infty$. Thus, the expression tends to $-\infty$.

The limit of the product is therefore:

$$\lim_{a_t \to 0^+} \Omega_t^* = \left(\lim_{a_t \to 0^+} \frac{1}{a_t \sigma_s}\right) \cdot \left(\lim_{a_t \to 0^+} \left(1 - \frac{1}{\eta_t(a_t)}\right)\right) = (+\infty) \cdot (-\infty) = -\infty.$$

This confirms that the EARA term diverges to negative infinity at the borrowing constraint.

The optimal portfolio allocation ω_t^* from equation (22) can be expressed as the sum of the classical portfolio choice ω_t^C (23) and the EARA term Ω_t^* (24).

$$\omega_t^* = \omega_t^C + \Omega_t^*. \tag{38}$$

The classical component can be expressed in terms of ERRA as:

$$\omega_t^C = -\frac{1}{a_t} \left(-\frac{a_t}{\eta_t} \frac{1}{v_a}\right) \left[\left(\frac{\mu_s - r_f}{\sigma_s^2}\right) v_a + \left(\frac{\rho_{sy}\sigma_y}{\sigma_s}\right) v_a y\right] = \frac{1}{\eta_t} \left[\frac{\mu_s - r_f}{\sigma_s^2} + \frac{\rho_{sy}\sigma_y}{\sigma_s} \frac{v_a y}{v_a}\right].$$

The full portfolio is then:

$$\omega_t^* = \frac{1}{\eta_t} \left[\frac{\mu_s - r_f}{\sigma_s^2} + \frac{\rho_{sy}\sigma_y}{\sigma_s} \frac{v_a y}{v_a}\right] + \frac{1}{a_t \sigma_s} \frac{\eta_t - 1}{\eta_t}.$$

To evaluate the limit, we combine terms over a common denominator, $a_t \sigma_s \eta_t$:

$$\omega_t^* = \frac{a_t \sigma_s \left[\frac{\mu_s - r_f}{\sigma_s^2} + \frac{\rho_{sy}\sigma_y}{\sigma_s} \frac{v_a y}{v_a}\right] + \eta_t - 1}{a_t \sigma_s \eta_t}.$$

Let us analyze the limit of the numerator and the denominator as $a_t \to 0^+$.

Numerator: As $a_t \to 0^+$, the first term $a_t\sigma_s[\dots] \to 0$ because the bracketed expression is assumed to be bounded. From Proposition 1, $\eta_t \to 0$. Thus, the numerator converges to:

$$\lim_{a_t \to 0^+} (a_t\sigma_s[\dots] + \eta_t - 1) = 0 + 0 - 1 = -1.$$

Denominator: As $a_t \to 0^+$, both $a_t \to 0$ and $\eta_t \to 0$. The denominator $a_t\sigma_s\eta_t \to 0$. Since $a_t > 0$, $\sigma_s > 0$, and $\eta_t > 0$, the denominator approaches 0 from the positive side (0^+).

Combining these results, the limit of the *unconstrained* optimal portfolio is:

$$\lim_{a_t \to 0^+} \omega^*_{t,\text{unc}} = \frac{-1}{0^+} = -\infty.$$

This shows that the underlying optimal behavior is to take an unboundedly large short position in risky assets.

However, the corollary states the portfolio choice *vanishes*. This result holds under the empirically relevant institutional constraint that households cannot take infinite short positions, often modeled simply as a no-short-sale constraint, $\omega^*_t \geq 0$.

Let $\omega^*_{t,\text{unc}}$ be the unconstrained solution derived above. The constrained optimal portfolio choice is given by:

$$\omega^*_t = \max(0, \omega^*_{t,\text{unc}}).$$

Since we have shown that $\lim_{a_t \to 0^+} \omega^*_{t,\text{unc}} = -\infty$, there exists a $\delta > 0$ such that for all $a_t \in (0, \delta)$, $\omega^*_{t,\text{unc}}(a_t) < 0$. In this region, the optimal constrained choice is the corner solution:

$$\omega^*_t(a_t) = 0 \quad \text{for } a_t \in (0, \delta).$$

Therefore, the limit of the constrained portfolio choice is:

$$\lim_{a_t \to 0^+} \omega^*_t = 0.$$

□

B. Asymptotic Behavior of the Consumption-Wealth Ratio

This appendix provides a rigorous proof for the main result in the text concerning the asymptotic behavior of the household's optimal consumption-wealth ratio. We prove that, as a household's wealth approaches infinity, its consumption-wealth ratio converges in the L^2 sense to a deterministic constant determined by the economy's fundamental

parameters. The proof relies on a scale transformation of the governing Backward Stochastic Partial Differential Equation (BSPDE), modern SPDE regularity theory, and martingale convergence theorems.

B.1 Preliminaries and Statement of the Main Theorem

The agent's value function $v(a, y, \tilde{k}, t)$ is the unique solution to the following BSPDE, given in the main text as equation (17):

$$\begin{aligned} -dv_t = \sup_{c_t, \omega_t} \Big\{ & u(c_t) + \left([\omega_t(\mu_s - r_f) + r_f]a_t + y_t - c_t\right) v_a + \frac{1}{2}(\omega_t a_t \sigma_s)^2 v_{aa} \\ & + \mu_y v_y + \frac{1}{2}\sigma_y^2 v_{yy} + \omega a_t \rho_{sy}\sigma_s\sigma_y v_{ay} + \tilde{\mu}_k v_{\tilde{k}} - \omega_t \sigma_s \left(a_t v_{aa} + v_a\right) \\ & - \sigma_y \rho_{sy} v_{ay} - \sigma_y \sqrt{1 - \rho_{sy}^2} v_{yy} \Big\} \, dt - \mathbf{Z}_t \, d\mathbf{W}_t, \end{aligned}$$

where for brevity we denote partial derivatives by subscripts (e.g., $v_a \equiv \frac{\partial v}{\partial a}$) and the martingale component by $\mathbf{Z}_t$. The optimal policy $c^*(a, y, \tilde{k}, t)$ is derived from the first-order condition $u'(c^*) = v_a$.

Our proof requires standard assumptions on the boundedness and smoothness of the model's coefficients.

Assumption 2 (Regularity of Coefficients)**.** The utility function is of the CRRA form, $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$ with $\gamma > 1$. The market and income process coefficients $(\mu_s, r_f, \sigma_s, \mu_y, \sigma_y, \rho_{sy}, \tilde{\mu}_k)$ are assumed to be uniformly bounded and sufficiently smooth (e.g., globally Lipschitz continuous) in the state variables.

Under this assumption, standard results from the theory of semi-linear SPDEs (see, e.g., Krylov, 1999) guarantee the existence and uniqueness of a solution $v \in C^{2,2,2,1}$ in the state space $(a, y, \tilde{k}, t)$, which possesses sufficient regularity for the following analysis.

It is worth noting that if the state variables were to include distribution measures, this framework would extend to a McKean-Vlasov equation, demanding more stringent regularity assumptions on the coefficients to guarantee the existence and uniqueness of solutions, as discussed in works like Cardaliaguet et al. (2019).

We can now state the main theorem 1 of this appendix.

A.2 Proof of Theorem 1

The proof proceeds in four steps. First, we define a scale-invariant value function and derive the BSPDE it satisfies. Second, we perform an asymptotic analysis on this new BSPDE, showing that its stochastic, income-related, and macro-state-dependent terms

vanish in the limit. Third, we characterize the resulting limiting deterministic HJB equation. Finally, we use the convergence of the value functions to establish the L^2 convergence of the policy functions.

Proof. **Step 1: The Scale Transformation** We introduce a scaling parameter $\xi > 0$ and define a new value function v_ξ which normalizes the original value function v by the agent's scale. The transformation applies only to the individual-specific state variables (a, y), not the aggregate state $\tilde{k}$.

Definition 6. For any $\xi > 0$, the scaled value function $v_\xi : (0, \infty) \times \mathbb{R} \times \mathbb{R} \times [0, T] \to \mathbb{R}$ is defined as:

$$v_\xi(a, y, \tilde{k}, t) := \frac{1}{\xi^{1-\gamma}} v(\xi a, \xi y, \tilde{k}, t).$$

The key to the proof is to understand the dynamics of v_ξ.

Lemma 1. The scaled value function $v_\xi(a, y, \tilde{k}, t)$ satisfies a BSPDE where the coefficients of the labor income process and the magnitude of the martingale term are scaled down by powers of ξ.

Proof of Lemma 1. We derive the dynamics of v_ξ by applying the chain rule to its definition. The partial derivatives of v_ξ relate to those of v as follows:

$$\begin{aligned}
\frac{\partial v_\xi}{\partial a} &= \xi^\gamma v_a(\xi a, \dots), & \frac{\partial^2 v_\xi}{\partial a^2} &= \xi^{\gamma+1} v_{aa}(\xi a, \dots) \\
\frac{\partial v_\xi}{\partial y} &= \xi^\gamma v_y(\xi a, \dots), & \frac{\partial^2 v_\xi}{\partial y^2} &= \xi^{\gamma+1} v_{yy}(\xi a, \dots) \\
\frac{\partial^2 v_\xi}{\partial a \partial y} &= \xi^{\gamma+1} v_{ay}(\xi a, \dots), & \frac{\partial v_\xi}{\partial \tilde{k}} &= \xi^{\gamma-1} v_{\tilde{k}}(\xi a, \dots).
\end{aligned}$$

Let $c_\xi = c/\xi$ and $\omega_\xi = \omega$ be the scaled controls. Substituting these into the original BSPDE 39, evaluating at $(\xi a, \xi y, \tilde{k}, t)$, and multiplying the entire equation by $\xi^{\gamma-1}$ yields the BSPDE for $v_\xi(a, y, \tilde{k}, t)$. The drift term, for instance, transforms as:

$$\xi^{\gamma-1} u(c) = u(c/\xi) = u(c_\xi).$$

And the wealth dynamics term becomes:

$$\begin{aligned}
&\xi^{\gamma-1} \left([\omega(\mu_s - r_f) + r_f](\xi a) + (\xi y) - (\xi c_\xi) \right) v_a(\xi a, \dots) \\
&\qquad = \left([\omega(\mu_s - r_f) + r_f] a + y - c_\xi \right) v_{\xi,a}(a, \dots).
\end{aligned}$$

After transforming all terms, the labor income drift and volatility terms are scaled: $\mu_y \mapsto \mu_y/\xi$ and $\sigma_y \mapsto \sigma_y/\xi$. The martingale term $dM_{\xi,t}$ for v_ξ has coefficients that are also scaled down, which is the subject of the next step. □

Step 2: Asymptotic Analysis of the Scaled BSPDE

We now analyze the BSPDE for v_ξ in the limit as $\xi \to \infty$. We show that all terms that differentiate the agent's problem from the classic Merton problem vanish.

Lemma 2 (Vanishing Martingale Component)**.** Let $M_{\xi,t}$ be the martingale component of the solution to the scaled BSPDE. Then,

$$\lim_{\xi\to\infty} \mathbb{E}\left[\sup_{s\in[0,t]} |M_{\xi,s}|^2\right] = 0.$$

Proof of Lemma 2. The martingale term $dM_{\xi,t}$ is a linear combination of the original stochastic drivers, with coefficients depending on the scaled derivatives. A typical coefficient is the one on $dW_{2,t}$, which includes the term $\frac{\sigma_y}{\xi}\rho_{sy}v_{\xi,y}$. The homogeneity of the value function implies that for some constant K, $|v_y(a,y,\tilde{k},t)| \le Ka^{-\gamma}$. Thus,

$$\left|\frac{\sigma_y}{\xi}\rho_{sy}v_{\xi,y}(a,\dots)\right| = \left|\frac{\sigma_y}{\xi}\rho_{sy}\xi^\gamma v_y(\xi a,\dots)\right| \le \frac{|\sigma_y\rho_{sy}|}{\xi}\xi^\gamma K(\xi a)^{-\gamma} = \frac{K|\sigma_y\rho_{sy}|}{a^\gamma}\frac{1}{\xi}.$$

The squared coefficient decays at a rate of ξ^{-2}. Applying this to all components of the martingale term, the quadratic variation $\langle M_\xi\rangle_t$ satisfies:

$$\mathbb{E}[\langle M_\xi\rangle_t] = \mathbb{E}\left[\int_0^t \|\mathbf{Z}_{\xi,s}\|_F^2 ds\right] \le \frac{Ct}{\xi^2}$$

for some constant C independent of ξ, where $\|\cdot\|_F$ is the Frobenius norm. By the Burkholder-Davis-Gundy (BDG) inequality, there exists a universal constant C_{BDG} such that:

$$\mathbb{E}\left[\sup_{s\in[0,t]} |M_{\xi,s}|^2\right] \le C_{\text{BDG}}\mathbb{E}[\langle M_\xi\rangle_t] \le \frac{C_{\text{BDG}}Ct}{\xi^2}.$$

Taking the limit as $\xi \to \infty$, the right-hand side converges to zero, which proves the lemma. □

The deterministic terms related to labor income ($\mu_y/\xi, \sigma_y/\xi$) and the macro-state ($\tilde{\mu}_k v_{\tilde{k}}/\xi^{1-\gamma}$) also vanish. The former is immediate from the scaling. For the macro-state term, we posit that in the limit of infinite wealth, the agent's value is insensitive to marginal changes in the aggregate state, implying $\lim_{\xi\to\infty} v_{\xi,\tilde{k}} = 0$.

Step 3: Convergence to the Deterministic Benchmark Problem

As $\xi \to \infty$, the BSPDE for v_ξ converges in distribution to a deterministic partial differential equation. This limiting equation is the familiar Hamilton-Jacobi-Bellman (HJB) equation for an infinitely-lived agent with no labor income, corresponding to the SSS optimal problem:

$$\begin{aligned}&\rho v_S(a)\\ =&\sup_{c,\omega}\left\{u(c) + [(\omega(\mu_s - r_f) + r_f)a - c - \omega\sigma_s]v_S'(a) + [\frac{1}{2}(\omega a\sigma_s)^2 - \omega a\sigma_s]v_S''(a)\right\}.\end{aligned} \tag{39}$$

The solution to this HJB equation is $v_S(a) = B\frac{a^{1-\gamma}}{1-\gamma}$, where B is a constant. The associated optimal consumption-wealth ratio is the constant $\bar{k}_{\mathrm{S}}$ given in the statement of Theorem 1.

To derive $\bar{k}_{\mathrm{S}}$, we substitute the proposed solution $v_S(a) = B\frac{a^{1-\gamma}}{1-\gamma}$ and the CRRA utility function $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$ into the HJB equation (39). The first and second derivatives of $v_S(a)$ are:

$$\begin{aligned} v_S'(a) &= Ba^{-\gamma} \\ v_S''(a) &= -\gamma Ba^{-\gamma-1}. \end{aligned}$$

First, we find the optimal consumption c^*. The first-order condition with respect to c is $u'(c) - v_S'(a) = 0$, which gives:

$$c^{-\gamma} = Ba^{-\gamma} \implies c^* = B^{-1/\gamma}a.$$

Let $\bar{k}_{\mathrm{S}} = B^{-1/\gamma}$, so $c^* = \bar{k}_{\mathrm{S}}a$.

Next, we find the optimal portfolio share ω^*. The first-order condition with respect to ω is:

$$[(\mu_s - r_f)a - \sigma_s]v_S'(a) + (\omega\sigma_s^2a^2 - \sigma_s a)v_S''(a) = 0.$$

Substitute the derivatives $v_S'(a)$ and $v_S''(a)$:

$$[(\mu_s - r_f)a - \sigma_s](Ba^{-\gamma}) + (\omega\sigma_s^2a^2 - \sigma_s a)(-\gamma Ba^{-\gamma-1}) = 0.$$

Dividing by $Ba^{1-\gamma}$ (since $B \neq 0$ and $a \neq 0$):

$$\omega^*(a) = \frac{\mu_s - r_f}{\gamma\sigma_s^2} + \frac{\gamma - 1}{\gamma a\sigma_s}.$$

Now, substitute $c^* = \bar{k}_{\mathrm{S}}a$, $\omega^* = \frac{\mu_s - r_f}{\gamma\sigma_s^2} + \frac{\gamma-1}{\gamma a\sigma_s}$, $v_S(a)$, $v_S'(a)$, and $v_S''(a)$ back into the HJB equation (39):

To complete the proof, we substitute the expressions for c^*, ω^*, $v_S(a)$, $v_S'(a)$, and $v_S''(a)$ into the HJB equation (39).

$$\begin{aligned} \rho B\frac{a^{1-\gamma}}{1-\gamma} =& \frac{(\bar{k}_{\mathrm{S}}a)^{1-\gamma}}{1-\gamma} + \left((r_f + \omega^*(\mu_s - r_f))a - \bar{k}_{\mathrm{S}}a - \omega^*\sigma_s\right)Ba^{-\gamma} \\ &+ \left(\frac{1}{2}(\omega^*\sigma_s a)^2 - \omega^* a\sigma_s\right)(-\gamma Ba^{-\gamma-1}). \end{aligned}$$

Divide the entire equation by $Ba^{1-\gamma}$:

$$\begin{aligned} \frac{\rho}{1-\gamma} =& \frac{\bar{k}_{\mathrm{S}}^{1-\gamma}}{B(1-\gamma)} + (r_f + \omega^*(\mu_s - r_f) - \bar{k}_{\mathrm{S}}) - \frac{1}{2}(\omega^*\sigma_s)^2\gamma + \frac{\omega^*\sigma_s}{a}(\gamma - 1) \\ =& \frac{\bar{k}_{\mathrm{S}}^{1-\gamma}}{B(1-\gamma)} + (r_f + \omega^*(\mu_s - r_f) - \bar{k}_{\mathrm{S}}) - \frac{1}{2}(\omega^*\sigma_s)^2\gamma, \qquad \text{as } a \to \infty. \end{aligned}$$

Since $\bar{k}_{\mathrm{S}} = B^{-1/\gamma}$, we have $B = \bar{k}_{\mathrm{S}}^{-\gamma}$. Substitute this into the first term:

$$\begin{aligned}\frac{\rho}{1-\gamma} =& \frac{\bar{k}_{\mathrm{S}}^{1-\gamma}}{\bar{k}_{\mathrm{S}}^{-\gamma}(1-\gamma)} + (r_f + \omega^*(\mu_s - r_f) - \bar{k}_{\mathrm{S}}) - \frac{1}{2}(\omega^*\sigma_s)^2\gamma \\ =& \frac{\bar{k}_{\mathrm{S}}}{1-\gamma} + r_f + \omega^*(\mu_s - r_f) - \bar{k}_{\mathrm{S}} - \frac{1}{2}(\omega^*\sigma_s)^2\gamma.\end{aligned}$$

Combine terms involving $\bar{k}_{\mathrm{S}}$:

$$\begin{aligned}\frac{\rho}{1-\gamma} =& \frac{\bar{k}_{\mathrm{S}} - \bar{k}_{\mathrm{S}}(1-\gamma)}{1-\gamma} + r_f + \omega^*(\mu_s - r_f) - \frac{1}{2}(\omega^*\sigma_s)^2\gamma \\ =& \frac{\gamma\bar{k}_{\mathrm{S}}}{1-\gamma} + r_f + \omega^*(\mu_s - r_f) - \frac{1}{2}(\omega^*\sigma_s)^2\gamma.\end{aligned}$$

Multiply by $(1-\gamma)$:

$$\rho = \gamma\bar{k}_{\mathrm{S}} + (1-\gamma)\left[r_f + \omega^*(\mu_s - r_f) - \frac{1}{2}(\omega^*\sigma_s)^2\gamma\right].$$

Now, substitute $\omega^* = \frac{\mu_s - r_f}{\gamma\sigma_s^2}$ (as $a \to \infty$) into the equation:

$$\begin{aligned}\rho =& \gamma\bar{k}_{\mathrm{S}} + (1-\gamma)\left[r_f + \frac{(\mu_s - r_f)^2}{\gamma\sigma_s^2} - \frac{1}{2}\left(\frac{\mu_s - r_f}{\gamma\sigma_s^2}\right)^2\sigma_s^2\gamma\right] \\ =& \gamma\bar{k}_{\mathrm{S}} + (1-\gamma)\left[r_f + \frac{(\mu_s - r_f)^2}{\gamma\sigma_s^2} - \frac{1}{2}\frac{(\mu_s - r_f)^2}{\gamma^2\sigma_s^4}\sigma_s^2\gamma\right] \\ =& \gamma\bar{k}_{\mathrm{S}} + (1-\gamma)\left[r_f + \frac{(\mu_s - r_f)^2}{\gamma\sigma_s^2} - \frac{1}{2}\frac{(\mu_s - r_f)^2}{\gamma\sigma_s^2}\right] \\ =& \gamma\bar{k}_{\mathrm{S}} + (1-\gamma)\left[r_f + \frac{(\mu_s - r_f)^2}{2\gamma\sigma_s^2}\right].\end{aligned}$$

Finally, solve for $\bar{k}_{\mathrm{S}}$:

$$\gamma\bar{k}_{\mathrm{S}} = \rho - (1-\gamma)\left(r_f + \frac{(\mu_s - r_f)^2}{2\gamma\sigma_s^2}\right)$$

$$\bar{k}_{\mathrm{S}} = \frac{1}{\gamma}\left(\rho - (1-\gamma)\left(r_f + \frac{(\mu_s - r_f)^2}{2\gamma\sigma_s^2}\right)\right).$$

This SSS optimal consumption-wealth ratio is the same as the baseline ratio.

Step 4: Convergence of the Policy Function

The final step is to translate the convergence of the value function $v_\xi \to v_M$ into the convergence of the policy function c^*/a. The optimal consumption policy is given by the first-order condition $u'(c^*) = v_a$, which for CRRA utility implies $c^* = (v_a)^{-1/\gamma}$. For the scaled problem, the policy is $c_\xi^* = (v_{\xi,a})^{-1/\gamma}$.

From the analysis above, we have the convergence $v_\xi \to v_M$ in a sufficiently strong sense (implied by the SPDE regularity theory) that ensures the convergence of the

derivatives, $v_{\xi,a} \to v_{M,a}$. Since $x \mapsto x^{-1/\gamma}$ is a continuous function, the Continuous Mapping Theorem implies:

$$\lim_{\xi\to\infty} c^*_\xi(1, y/\xi, \tilde{k}, t) = (v_{M,a}(1))^{-1/\gamma} = c_M(1) = \bar{k}_{\mathrm{S}} \quad \text{a.s.}$$

To establish the desired convergence, we set the scaling factor ξ equal to the wealth level a:

$$\frac{c^*(a, y, \tilde{k}, t)}{a} = \frac{c^*(\xi \cdot 1, \xi \cdot (y/a), \tilde{k}, t)}{\xi} = c^*_\xi(1, y/a, \tilde{k}, t).$$

Therefore, we have established almost sure convergence:

$$\lim_{a\to\infty} \frac{c^*(a, y, \tilde{k}, t)}{a} = \bar{k}_{\mathrm{S}} \quad \text{a.s.}$$

To strengthen this to L^2 convergence, we must show that the family of random variables $\{(c^*/a)^2\}_{a>0}$ is uniformly integrable. The regularity bounds on the value function v and its derivative v_a (which grows as $a^{-\gamma}$) ensure that c^*/a is uniformly bounded in L^2. By the Dominated Convergence Theorem, almost sure convergence implies L^2 convergence. □