A Global Solution Method for HACT Models with Aggregate Risk

Preliminary and incomplete draft
All comments are welcome

Niklas Bonnmann*, Elisabeth Pröhl†

March 25, 2026

Abstract

We develop a method to compute global solutions to continuous-time heterogeneous agent (HACT) models with aggregate risk, non-stationarity, and financial constraints. When the cross-sectional distribution evolves stochastically, the associated master equation describing equilibrium becomes second order. First, we provide a heuristic derivation of this second-order master equation and show how its distributional derivatives can be approximated using a polynomial chaos expansion, thereby transforming it into a standard, albeit high-dimensional, partial differential equation (PDE). Second, building on the Deep BSDE approach for high-dimensional PDEs, we develop a general-equilibrium extension – GE-Deep BSDE – that jointly solves for value functions, equilibrium prices, and distributional consistency using multi-objective deep learning. Applying the method to a two-asset model, we show that while total factor productivity shocks fail to generate sufficient risk premia, stochastic capital depreciation with collateral constraints produces endogenous risk and helps reconcile microeconomic wealth dispersion with macro-finance evidence.

Keywords: Dynamic stochastic general equilibrium, Incomplete markets, Heterogeneous agents, Aggregate uncertainty, Numerical solutions, Reinforcement learning, Deep neural networks, Polynomial chaos

JEL Classification: C62, C63, D31, D52, E21

*Vrije Universiteit Amsterdam and Tinbergen Institute, De Boelelaan 1105, 1081 HV Amsterdam, Netherlands, n.m.bonnmann@vu.nl.

†University of Amsterdam, Roetersstraat 11, 1018 WB Amsterdam, Netherlands, e.proehl@uva.nl.

We thank Daniel Albuquerque, Léonard Bocquet, Ji Huang, Kenneth Judd, Greg Kaplan, Walter Pohl, Simon Scheidegger and Yucheng Yang for valuable discussions. We also thank seminar participants at the University of Amsterdam, Center for Advanced Study Norway, Tinbergen Institute and VU Amsterdam for their helpful comments, as well as conference participants, in 2025, at the Conference on Deep Learning for Dynamic Stochastic Models, ES World Congress, SITE Workshop on Asset Pricing, ECONDAT Spring Meeting and T2M.

1 Introduction

Heterogeneous-agent models have become central to modern macroeconomics. They are particularly well suited for studying distributional dynamics, but their broader applicability has long been constrained by computational difficulties. Continuous-time formulations, as in Achdou, Han, Lasry, Lions, and Moll (2022), have substantially reduced these barriers by leveraging efficient finite-difference methods. Incorporating aggregate risk, however, remains a considerable obstacle. In the presence of aggregate shocks, the cross-sectional distribution of households evolves stochastically and becomes an infinite-dimensional state variable that influences equilibrium prices and individual policy functions, giving rise to a severe curse of dimensionality.

Although there has been significant progress in computing models with aggregate risk using local approximations around the steady state,[1] this approach is less suitable for models with aggregate nonlinearities. More recently, a growing literature has tackled the curse of dimensionality of global solutions by applying machine learning techniques.[2] These methods, however, abstract from non-stationary settings and financial constraints. As a result, analyzing heterogeneous agent models that combine aggregate nonlinearities, non-stationarity, and financial constraints with business-cycle risk remains a major challenge.[3] Non-stationarity, for example, is intrinsic to climate-economic models where temperature is an upward-trending state variable, while aggregate nonlinearities and financial frictions play an important role in macro-finance models that generate endogenously arising crises dynamics.

This paper addresses these challenges by introducing a novel computational method that delivers global solutions to a broad class of continuous-time heterogeneous agent models with aggregate risk. Our approach accommodates non-stationary dynamics, handles nontrivial borrowing constraints and market clearing conditions, and high-dimensional state spaces. It is based on the master equation, which combines the Hamilton–Jacobi–Bellman (HJB) and Kolmogorov forward equations (KFE) to ensure consistency between individual decision-making and the evolution of the cross-sectional distribution. We solve this equation using machine-learning techniques. We

[1] A non-exhaustive list includes Ahn et al. (2018); Boppart et al. (2018); Bayer and Luetticke (2020); Auclert et al. (2021); Bilal (2023).

[2] See Gu et al. (2024); Gopalakrishna et al. (2024); Huang (2024); Payne et al. (2024); Yang et al. (2025).

[3] Moll (2024) even names the difficulty of solving heterogeneous agent models with aggregate risk as one reason to move away from rational expectations.

illustrate this method with a two-asset HACT model with aggregate risk, featuring a one-period zero-coupon bond and illiquid capital. Kaplan, Moll, and Violante (2018) show that this environment generates an empirically realistic wealth distribution. However, despite its success in matching microeconomic moments, the model fails to reproduce key macro-finance facts: when aggregate risk takes the form of total factor productivity shocks, the equity premium is too low. We show that this disconnect is alleviated when aggregate risk instead enters through stochastic capital depreciation combined with a collateral constraint. In this setting, capital is not only illiquid but also risky, and the collateral constraint generates endogenous risk in the spirit of Brunnermeier and Sannikov (2014).

Introducing aggregate risk through stochastic capital depreciation substantially increases the theoretical complexity of the model in two aspects: The master equation becomes second order because the law of motion of the wealth distribution is itself stochastic, and, the equation features distributional derivatives that are difficult to approximate numerically. Thus, our first contribution is of a theoretical nature and tackles these two challenges. First, we provide a heuristic derivation of the second-order master equation, guided by the rigorous theory of mean-field games developed by Cardaliaguet, Delarue, Lasry, and Lions (2019); Carmona and Delarue (2018a,b). Our objective is to make this derivation accessible to economists familiar with the basics of continuous-time stochastic processes. While Bilal (2023) carefully derives the master equation for models with deterministic KFEs, we are the first to do so when the KFE is stochastic. The key step is an extension of Itô's lemma: whereas the classical result characterizes the dynamics of functions of stochastic processes, its extension characterizes the dynamics of functions defined over both stochastic processes and stochastic distributions. Second, we show how the distributional derivatives appearing in the master equation can be approximated using standard derivatives on the real line. To this end, we represent the cross-sectional distribution via a polynomial chaos expansion, an approach first introduced in a discrete-time heterogeneous-agent setting by Proehl (2025a). This projection transforms the distributional derivatives into derivatives with respect to a finite-dimensional vector of projection coefficients, thereby reducing the master equation to a standard, albeit high-dimensional, partial differential equation (PDE).

Our second contribution is computational. We develop a numerical method to solve the master equation jointly with the market-clearing conditions. The approach

builds on the Deep BSDE method of E, Han, and Jentzen (2017); Han, Jentzen, and E (2018), which is designed to solve high-dimensional partial differential equations. The method reformulates the PDE as a system of laws of motion for the state variables, the value function, and its first and second derivatives, which are simulated forward from unknown initial conditions. These initial objects, i.e., the value function and its slope and curvature, are recovered by minimizing violations of the terminal boundary condition implied by the simulated terminal values. In this sense, Deep BSDE can be viewed as a shooting algorithm that approximates unknown functions with neural networks and updates them via stochastic gradient descent. In our setting, the value function's slope and curvature also determine agents' optimal policies, placing the algorithm in the class of reinforcement-learning methods. In general equilibrium (GE), prices such as bond returns must be determined jointly with agents' policies, and equilibrium requires consistency between individual decision rules and the evolution of the cross-sectional distribution. We incorporate these features by extending the Deep BSDE framework in two ways. First, we expand the set of unknown functions to include equilibrium prices. Second, we discipline these additional neural networks by penalizing market-clearing violations and distributional inconsistency in the loss function. We employ multi-objective deep learning techniques to ensure stable convergence. We refer to the resulting algorithm as the GE–Deep BSDE method.

Related Literature. This paper contributes to multiple strands of literature, foremost among them the development of numerical algorithms for solving heterogeneous agent models with aggregate risk. Traditional approaches address the challenge of an infinite-dimensional state space by either relying on a bounded-rationality approach to reduce dimensionality (see Den Haan, 2010, for an overview), or on local perturbations around a stationary equilibrium as in, e.g., Reiter (2009); Ahn, Kaplan, Moll, Winberry, and Wolf (2018); Boppart, Krusell, and Mitman (2018); Bayer and Luetticke (2020); Auclert, Bardóczy, Rognlie, and Straub (2021). Grid-based methods solving the rational expectations equilibrium globally suffer from the course of dimensionality. To overcome this limitation, a growing body of research – including this paper – leverages grid-free, simulation-based deep learning as a powerful tool for tackling (very) high-dimensional optimization problems, see e.g., Scheidegger and Bilionis (2019); Maliar, Maliar, and Winant (2021); Azinovic, Gaegauf, and Scheidegger (2022). Another class of methods employs neural networks under bounded-rationality to approximate the perceived law of motion of aggregate variables more flexibly as in

Fernández-Villaverde, Hurtado, and Nuno (2023), or to parametrize a mapping from a finite-agent discretization to a lower-dimensional finite-moment representation of the cross section as in Han, Yang, and E (2021).

More specifically, our method belongs to the class of algorithms for models posed in continuous time. Popularized in macroeconomics by Achdou et al. (2022), their analytical tractability has not only been exploited in a local perturbation method as in Bilal (2023), but it also forms the basis of several global solution methods – either using value function iteration with adaptive-sparse grids as in Schaab (2020), or using deep learning. More concretely, the recursive equilibrium in these models satisfies a fundamental necessary condition on the value function: an infinite-dimensional, non-standard PDE known as the master equation. Within the deep learning literature on solving this equation, two main approaches have emerged. The first employs techniques from Sirignano and Spiliopoulos (2018) and Raissi, Perdikaris, and Karniadakis (2019), where a neural network approximates the value function by minimizing the mean-squared error of the PDE and its boundary conditions, see e.g., Gopalakrishna (2021); Duarte, Duarte, and Silva (2024); Fan, Qiao, Jiao, Gu, Li, and Lu (2024); Gu, Laurière, Merkel, and Payne (2024); Gopalakrishna, Payne, and Gu (2024); Payne, Rebei, and Yang (2024). The second approach, which we adopt in this paper, employs reinforcement learning. To the best of our knowledge, this approach has so far only been adopted in the macroeconomic literature by Han, Yang, and E (2021) in discrete time, and Huang (2024) in continuous time for models with full information rational expectations. Yang, Wang, Schaab, and Moll (2025) use reinforcement learning to solve models with bounded rationality. Our method goes beyond these methods as we accommodate non-stationarity and financial frictions. More broadly, Pham and Warin (2022, 2023) use an adaption of the Deep BSDE method proposed in Huré, Pham, and Warin (2020) to solve mean-field games, though without applications to macroeconomics.

Finally, this paper contributes to the macroeconomic literature that employs polynomial chaos expansions as a numerical tool. While its application for parameter sensitivity analysis is closest to its origin in engineering and physics as in Harenberg, Marelli, Sudret, and Winschel (2019); Fehrle, Heiberger, and Huber (2025), Proehl (2025a) uses it as a discretization scheme for heterogeneous agent models. Our approach is the first, to our knowledge, to integrate this discretization scheme into a deep learning framework for solving HACT models.

Outline. The paper proceeds as follows. The next section introduces our example model with stochastic capital depreciation and a general model setting nesting a range of applications. For this general model, Section 3 derives the master equation, which we then approximate using polynomial chaos in Section 4. In Section 5, we describe how a global solution of the approximated master equation can be computed using an adaption of the Deep BSDE method. The last section explores the economic insights from our example model.

2 A General HACT Model Setup

Before stating our fully general model framework, we introduce a well-known benchmark model: the two-asset economy in Krusell and Smith (1997) with illiquidity in productive capital as in Kaplan et al. (2018). We adjust this benchmark model to better illustrate the advantages of our solution method. Instead of introducing the aggregate risk through a total factor productivity shock as in Bilal (2023); Ahn et al. (2018), we model it as an aggregate stochastic depreciation shock that affects the households' capital stock directly.

2.1 An Example Model

Firm. Individual households hold capital and rent it to a perfectly competitive representative firm aggregating to capital supply K_t. Households also provide one unit of labor each with an exogenously given individual labor productivity aggregating to labor supply L_t. The firm produces $F(L_t, K_t) - \delta K_t$ units of a single good after accounting for deterministic firm-level capital depreciation at a rate $\delta > 0$. The firm optimally pays an aggregate rental rate r_t and wage rate w_t to households given by the marginal productivities net of depreciation.

Household Problem. There is a continuum of households with measure one. Each household i has idiosyncratic labor productivity $e^{z_t^i}$ earning the wage rate w_t. This labor income is consumed and saved in a liquid and an illiquid asset. The liquid bond pays the risk-free rate r_t^b on household i's bond holdings b_t^i. Household i's capital holdings k_t^i are illiquid as they are subject to convex adjustment costs $\Psi(\iota_t^i) = \psi_0|\iota_t^i| + \psi_1|\iota_t^i|^{\psi_2}$ per unit of in- or divestment ι_t^i with ψ_0, $\psi_1 > 0$ and $\psi_2 > 1$. Furthermore, households incur aggregate depreciation shocks B_t^{ag} proportional to

their capital holdings. Log-labor productivity z_t^i follows an idiosyncratic Ornstein-Uhlenbeck process which is correlated with the aggregate shocks B_t^{ag} by $\omega \in [-1, 1]$. Households have CRRA-preferences. Furthermore, they face a borrowing constraint $k_t^i \geq 0$ and a collateral constraint $b_t^i \geq \bar{b} - \phi_t k_t^i$ where $\bar{b} \leq 0$ is constant and $\phi_t \in [0, 1]$ is a deterministic process. Overall, the household problem reads

$$\begin{aligned} \max_{\{c_t^i, \iota_t^i\}_{t\geq 0}} \; & \mathbb{E}_0 \left[\int_0^\infty e^{-\rho t} \frac{(c_t^i)^{1-\gamma}}{(1-\gamma)} dt \right] \\ \text{s.t. } & dz_t^i = \kappa_z(\bar{z} - z_t^i)dt + \sqrt{1-\omega^2}\sigma_z dB_t^i + \omega\sigma_z dB_t^{ag}, \\ & dk_t^i = (r_t^k + \iota_t^i)k_t^i dt + k_t^i \sigma_K dB_t^{ag}, \\ & db_t^i = \left\{ w_t e^{z_t^i} + b_t^i r_t^b - k_t^i(\iota_t^i + \Psi(\iota_t^i)) - c_t^i \right\} dt, \\ & b_t^i \geq \bar{b} - \phi_t k_t^i, \end{aligned} \tag{1}$$

where B_t^{id} denotes the idiosyncratic Brownian innovations that are i.i.d. across households. Note that the law of motion of capital automatically ensures $k_t^i \geq 0$ due to its geometric Brownian motion structure.

Brunnermeier and Sannikov (2014) argue within a representative agent framework that, at the aggregate level, whether aggregate shocks affect the capital stock through stochastic depreciation or total factor productivity (TFP) is equivalent. We show in this paper that this equivalence breaks down in our heterogeneous agent economy. Although both approaches result in risky returns on capital, stochastic depreciation adds an additional layer of risk as the value of capital stock becomes risky. One can interpret this as households basing their decisions on their capital's market value rather than its book value. However, this modeling choice implies a stochastic law of motion of the cross-sectional distribution,[4] a case that has not yet been considered in the existing HACT macroeconomics literature (e.g. Kaplan et al., 2018; Ahn et al., 2018; Bilal, 2023; Gu et al., 2024). This choice, therefore, illustrates the generality of our modeling framework. In addition, we model the collateral constraint parameter ϕ_t as a deterministic process. This serves as an illustration of a policy experiment. To the best of our knowledge, we are the first to compute a transition along a deterministic time trend in a heterogeneous-agent economy with aggregate risk.

Market Clearing. In equilibrium, the labor, capital and bond market must

[4]The Kolmogorov forward equation in our framework is a stochastic differential equation rather than an ordinary differential equation.

clear, i.e.,

$$\begin{aligned} w_t^* =& F_L\left(\int e^z dM_t(z,k,b), \int k dM_t(z,k,b)\right) \\ r_t^{k,*} =& F_K\left(\int e^z dM_t(z,k,b), \int k dM_t(z,k,b)\right) - \delta \\ 0 =& \int b dM_t(z,k,b), \end{aligned} \tag{2}$$

where $\{M_t\}_{t\geq 0}$ denotes the cross-sectional distribution of household-specific variables. In equilibrium, the distribution's law of motion must be consistent with the households' optimal choices.

Treatment of the Collateral Constraint. To enable the computation of the constrained problem, we substitute the state variable of bond holdings b_t^i with the slackness of the constraint $e_t = b_t^i + \phi_t k_t^i - \bar{b}$. Thus, the collateral constraint is equivalent to $e_t \geq 0$. Furthermore, we substitute the dynamics db_t^i with the dynamics of the substitute state[5]

$$de_t^i = \left\{ w_t e^{z_t^i} + b_t^i r_t^b + k_t^i \left[\frac{d}{dt}\phi_t + \phi_t r_t^k - (1-\phi_t)\iota_t^i - \Psi(\iota_t^i)\right] - c_t^i \right\} dt + \phi_t k_t^i \sigma_K dB_t^{ag}$$

which follows from Itô's lemma. The household optimization problem in (1) is equivalent to

$$\begin{aligned} \max_{\{c_t^i,\iota_t^i\}_{t\geq 0}} \ & \mathbb{E}_0\left[\int_0^\infty e^{-\rho t}\frac{(c_t^i)^{1-\gamma}}{(1-\gamma)}dt\right] \\ \text{s.t. } & dz_t^i = \kappa_z(\bar{z}-z_t^i)dt + \sqrt{1-\omega^2}\sigma_z dB_t^i + \omega\sigma_z dB_t^{ag}, \\ & dk_t^i = (r_t^k + \iota_t^i)k_t^i dt + k_t^i \sigma_K dB_t^{ag}, \\ & de_t^i = \left\{ w_t e^{z_t^i} + (e_t^i + \phi_t k_t^i - \bar{b}) r_t^b + k_t^i \left[\frac{d}{dt}\phi_t + \phi_t r_t^k - (1-\phi_t)\iota_t^i - \Psi(\iota_t^i)\right] - c_t^i \right\} dt \\ & \qquad + \phi_t k_t^i \sigma_K dB_t^{ag} \\ & e_t^i \geq 0. \end{aligned} \tag{3}$$

Lastly, the bond market clearing condition adjusts to $0 = \int (e + \phi_t k - \bar{b}) dM_t(z,k,e)$. It is well known (see Achdou et al., 2022) that the state constraint of this rewritten problem can be transformed into state boundary conditions. These conditions restrict

[5] If the deterministic process ϕ_t has jumps, then $\int_{t^*-\Delta t}^{t^*+\Delta t} k_t^i \frac{d}{dt}\phi_t dt$ should be interpreted as $k_{t^*}^i \Delta\phi_{t^*}$ when $\Delta t \to 0$ with the process changing by $\Delta\phi_{t^*} = \phi_{t^*} - \phi_{t^*-}$ at the jump time point t^*.

the optimal controls in the constrained region to ensure that the volatility of de_t^i is zero and its drift is nonnegative whenever $e_t^i = 0$, i.e.,

$$\begin{aligned} 0 &= \lim_{e_t^i \to 0} \phi_t k_t^i \sigma_K \\ 0 &= \lim_{e_t^i \to 0} \min \left\{ 0, w_t e^{z_t^i} + b_t^i r_t^b + k_t^i \left[\frac{d}{dt}\phi_t + \phi_t r_t^k - (1-\phi_t)\iota_t^{i,*} - \Psi\left(\iota_t^{i,*}\right) \right] - c_t^{i,*} \right\} \end{aligned} \tag{4}$$

where the asterisk denotes the optimal controls.

Optimal Controls. Denote the value function by $V(t, z_t^i, k_t^i, e_t^i) = e^{-\rho t} J(t, z_t^i, k_t^i, e_t^i)$. We drop function arguments for notational convenience and denote function derivatives by a subscript, e.g. $J_{k^i} = \frac{\partial}{\partial k^i} J$. In the unconstrained region, standard dynamic programming arguments yield the first-order conditions for the optimal controls

$$\begin{aligned} 0 =& c_t^{i,*} - \left(J_e\right)^{-\frac{1}{\gamma}} \\ 0 =& J_k - \left(1 - \phi_t + \Psi_\iota\left(\iota_t^{i,*}\right)\right) J_e. \end{aligned} \tag{5}$$

Note that households' optimal investment policies feature an inaction region which is due to the cost function's linear term. Households only adjust their holdings of illiquid capital if the absolute difference between their marginal rate of substitution between both assets and the cost of capital $1 - \phi_t$ is higher than the minimum marginal cost ψ_0 of doing so,[6] i.e.,

$$\iota_t^{i,*} = \begin{cases} \left(\frac{J_k - (1-\phi_t+\psi_0)J_e}{\psi_1 \psi_2 J_e}\right)^{\frac{1}{\psi_2 - 1}} & \text{if} \quad \frac{J_k}{J_e} - (1-\phi_t) > \psi_0 \\ 0 & \text{if} \quad \left|\frac{J_k}{J_e} - (1-\phi_t)\right| \leq \psi_0 \\ -\left(\frac{(1-\phi_t-\psi_0)J_e - J_k}{\psi_1 \psi_2 J_e}\right)^{\frac{1}{\psi_2 - 1}} & \text{if} \quad (1-\phi_t) - \frac{J_k}{J_e} > \psi_0. \end{cases}$$

The inaction region widens with the minimum marginal cost ψ_0. The collateral constraint directly reduces the cost of capital by the fraction ϕ_t that may serve as collateral. Thus, the collateral constraint cushions the negative effect of adjustment costs when investing and amplifies it when divesting. Overall, portfolio choice theory tells us that households' illiquid capital demand consists of *(i)* their myopic diversification demand stemming from aggregate risk, *(ii)* their hedging demand due to idiosyncratic labor income risk, and *(iii)* a tolerated deviation from the sum of the former two due

[6] This is a well-established result, see e.g., Constantinides (1986) or Davis and Norman (1990).

to adjustment costs.[7]

State Boundary Conditions. Given the optimal controls, we can derive the state boundary conditions from (4). If $\sigma_K > 0$ and $\phi_t > 0$, we need that $\lim_{e_t^i \to 0} k_t^i = 0$. Given the dynamics of k_t^i in (3), this implies a restriction on $\iota_t^{i,*}$ which translates into a boundary condition for J_{k^i} due to (5), i.e.[8]

$$\lim_{e_t^i \to 0} \iota_t^{i,*} = -\infty \Leftrightarrow \lim_{e_t^i \to 0} J_{k^i} = -\infty. \tag{6}$$

This boundary condition illustrates that the collateral constraint imposes an additional risk that becomes larger the closer an agent is to the boundary. The agent faces rollover risk for their liquid bond debt due to the riskiness of the collateral. Agents hedge this risk with the illiquid asset. This additional (negative) hedging demand starts to dominate any other motive for holding the illiquid asset the closer the agent is to the boundary. Furthermore, the state boundary condition implies at $e_t^i = 0$ that

$$c_t^{i,*} \leq w_t e^{z_t^i} + \bar{b} r_t^b \Leftrightarrow J_{e^i} \geq \left(w_t e^{z_t^i} + \bar{b} r_t^b \right)^{-\gamma}.$$

When negative shocks move an agent closer to the constraint, the agent not only reduces consumption and borrowing, but also starts selling capital even though it is illiquid. Thus, a collateral constraint amplifies the price impact of a large or persistent negative aggregate shock due to the fire-sale boundary condition in (6). This amplification effect occurs only because the aggregate shock enters the capital dynamics directly. The effect would not exist if $\sigma_K = 0$ and the aggregate risk entered through an aggregate TFP process.

[7] In the special case without convex costs and labor income risk ($\bar{z}, \psi_1, \sigma_z = 0$), Davis and Norman (1990) prove that the optimal risky wealth share $\{k_t^i/(k_t^i + b_t^i)\}_{t \geq 0}$ when risky wealth is illiquid follows a reflected diffusion inside a ψ_0-dependent wedge around the optimal risky wealth share when it is liquid, i.e., without adjustment costs. Furthermore, Isaenko (2005) considers convex adjustment costs in a finite-horizon, complete-market model without intermediate consumption concluding that the shape of the inaction region is independent of the degree of convexity $\psi_2 - 1 > 0$.

[8] The capital dynamics result in $\log(k_t^i) = \log(k_0^i) + \int_0^t (r_s^k + \iota_s^{i,*} - \sigma_K^2/2) ds + \sigma_K B_t^{ag}$ which gives rise to the state boundary condition. Note that total investment of the constrained agents vanishes because $\lim_{e_t^i \to 0} \iota_t^{i,*} k_t^i \approx \lim_{k_t^i \to 0} k_t^i \log(k_t^i) = 0$.

2.2 The General Framework

Time is continuous $t \in [0, \infty)$. The economy is populated by a continuum of agents $i \in [0, 1]$. All stochastic processes are defined on a generic filtered probability space $(\Omega, \Sigma, \mathcal{F}_{t\geq 0}, \mathbb{P})$. Aggregate risk is modeled by an aggregate d^{ag}-dimensional Brownian motion $\{B_t^{ag}\}_{t\geq 0}$. Agent i's idiosyncratic risk is given by a d^{id}-dimensional Brownian motion $\{B_t^i\}_{t\geq 0}$ that is i.i.d. across agents $i \in [0, 1]$. We assume that the entries of the vectors B_t^{ag} and B_t^i are pairwise independent for any $i \in [0, 1]$.[9]

2.2.1 The Agent's Optimization Problem

At any instant $t \geq 0$, agent $i \in [0, 1]$ controls their N^x-dimensional vector of agent-specific state variables x_t^i with an N^y-dimensional vector of controls y_t^i to maximize their flow utility $u(y_t^i, x_t^i)$ which is discounted at the constant rate ρ. The state vector $x_t^i = [x_t^{i,uc}, e_t^i]$ contains $N^e \leq N^x$ constrained states with $e_t^i \geq 0$ for any $t \geq 0$. Agents optimize while observing an N^a-dimensional vector of aggregate state variables a_t. Furthermore, all agents have access to N^r market instruments whose prices have potentially risky relative returns r_t. When optimizing, agents take these returns as given. They are treated as known equilibrium functions $r_t = r(t, a_t)$. Overall, agent $i \in [0, 1]$ optimizes

$$\begin{aligned} \max_{\{y_t^i\}_{t\geq 0}} \mathbb{E}_0 &\left[\int_0^\infty e^{-\rho t} u\left(y_t^i, x_t^i\right) dt\right] \qquad (7) \\ \text{s.t. } dx_t^i &= \mu\left(t, y_t^i, x_t^i, a_t, r_t\right) dt + \sigma^{id}\left(t, y_t^i, x_t^i, a_t, r_t\right) dB_t^i + \sigma^{ag}\left(t, y_t^i, x_t^i, a_t, r_t\right) dB_t^{ag} \\ da_t &= \eta\left(t, a_t, r_t\right) dt + \zeta\left(t, a_t, r_t\right) dB_t^{ag} \\ e_t^i &\geq 0 \, \forall t \geq 0 \end{aligned}$$

where $\mu(t, y_t^i, x_t^i, a_t, r_t) \in \mathbb{R}^{N^x}$, $\sigma^o(t, y_t^i, x_t^i, a_t, r_t) \in \mathbb{R}^{N^x \times d^o}$ for $o \in \{id, ag\}$, $\eta(t, a_t, r_t) \in \mathbb{R}^{N^a}$ and $\zeta(t, a_t, r_t) \in \mathbb{R}^{N^a \times d^{ag}}$. The states' aggregate and idiosyncratic covariance

[9] Nevertheless, correlation between the aggregate and idiosyncratic risk can be introduced through a covariance matrix as is illustrated in our example model in Section 2.1.

matrices are defined as

$$\begin{aligned}\Sigma^o\left(t,y^i,x^i,y^{i'},x^{i'},a,r\right) =&\sigma^o\left(t,y^i,x^i,a,r\right)\sigma^o\left(t,y^{i'},x^{i'},a,r\right)^T \in \mathbb{R}^{N^x\times N^x} \\ \Sigma^{xa}\left(t,y^i,x^i,a,r\right) =&\sigma^{ag}\left(t,y^i,x^i,a,r\right)\zeta(t,a,r)^T \in \mathbb{R}^{N^x\times N^a} \\ \Sigma^a(t,a,r) =&\zeta(t,a,r)\zeta(t,a,r)^T \in \mathbb{R}^{N^a\times N^a}.\end{aligned} \tag{8}$$

Example. *In our example model, the aggregate and idiosyncratic Brownian motions are both one-dimensional. Households choose consumption and investment, that is,* $y_t^i = [c_t^i, \iota_t^i]^T$. *The agent-specific state variables are* $x_t^i = [z_t^i, k_t^i, e_t^i]^T$ *and there are no aggregate variables* a_t, *i.e.,* $0 = a_0 = \eta(t,a_t,r_t) = \zeta(t,a_t,r_t)$. *The return vector features both assets* $r_t = [r_t^k, r_t^b]^T$. *This implies the drift vector and dispersion matrices*

$$\mu\left(t,y_t^i,x_t^i,a_t,r_t\right) = \begin{bmatrix} \kappa_z(\bar{z}-z_t^i) \\ (r_t^k+\iota_t^i)k_t^i \\ w_t e^{z_t^i} + \left(e_t^i - \phi_t k_t^i + \bar{b}\right) r_t^b + k_t^i\left[\frac{d}{dt}\phi_t + \phi_t r_t^k - (1-\phi_t)\iota_t^i - \Psi(\iota_t^i)\right] - c_t^i \end{bmatrix}$$

$$\sigma^{id}\left(t,y_t^i,x_t^i,a_t,r_t\right) = \begin{bmatrix}\sqrt{1-\omega^2}\sigma_z \\ 0 \\ 0\end{bmatrix}, \quad \sigma^{ag}\left(t,y_t^i,x_t^i,a_t,r_t\right) = \begin{bmatrix}\omega\sigma_z \\ k_t^i\sigma_K \\ \phi_t k_t^i \sigma_K\end{bmatrix}.$$

Remark. *Similarly to the collateral constraint in our example model, models with a non-trivial state constraint can typically be transformed to fit our general setting by substituting the distance to the binding constraint as a state variable.*

2.2.2 The Agent's Optimality Conditions

Agent i's optimization problem can be solved by dynamic programming.[10] Define the value function by

$$V(t,x_t^i,a_t) = \max_{\{y_s^i\}_{s\geq t}} \mathbb{E}_t\left[\int_t^\infty e^{-\rho s} u\left(y_s^i, x_s^i\right) ds\right]. \tag{9}$$

The Hamilton-Jacobi-Bellman (HJB) equation is stated in terms of the indirect utility function $J(t,x^i,a) = e^{\rho t}V(t,x^i,a)$ such that

$$0 = \max_{y^i}\left\{u\left(y^i,x^i\right) - \rho J(t,x^i,a) + \frac{\mathbb{E}_t\left[dJ(t,x^i,a)\right]}{dt}\right\}, \tag{10}$$

[10] See for instance Björk (2009, Chapter 19).

where the last term represents the drift term of Itô's formula applied to the function J. To fully state the HJB equation and the related boundary conditions, we denote the respective partial differential equations using operators. We first define the HJB operator $\mathcal{H}^{HJB}[J, y^i, r](t, x^i, a)$ mapping the functions $[J, y^i, r]$ into a function of (t, x^i, a) by

$$\begin{aligned}\mathcal{H}^{HJB}[J, y^i, r] =& u - \rho J + J_t + \sum_{n=1}^{N^a} J_{a_n} \eta_n + \sum_{n=1}^{N^x} J_{x_n^i} \mu_n + \frac{1}{2} \left[\sum_{n,n'=1}^{N^a} J_{a_n a_{n'}} \Sigma^a_{nn'} \right. \\ &\left. + \sum_{n,n'=1}^{N^x} \sum_{o \in \{id, ag\}} J_{x_n^i x_{n'}^i} \Sigma^o_{nn'} + \sum_{n=1}^{N^x} \sum_{n'=1}^{N^a} J_{x_n^i a_{n'}} \Sigma^{xa}_{nn'} \right]\end{aligned} \tag{11}$$

where the subscripts of J denote its derivatives. Function arguments have been dropped for notational convenience. The covariance matrices are given in (8).

We define the state boundary condition to guarantee that the state constraint $e_t^i \geq 0$ holds at all times. This means that the drift of e_t^i in (7) at the boundary $e_t^i = 0$ must be nonnegative and the dispersion terms must vanish. Let

$$de_t^i = \psi\left(t, y_t^i, x_t^i, a_t, r_t\right) dt + \upsilon^{id}\left(t, y_t^i, x_t^i, a_t, r_t\right) dB_t^i + \upsilon^{ag}\left(t, y_t^i, x_t^i, a_t, r_t\right) dB_t^{ag}.$$

At the boundary $e_t^i = 0$, it must hold that

$$\begin{aligned}0 \leq& \psi\left(t, y_t^i, [x_t^{i,uc}, 0], a_t, r_t\right) \\ 0 =& \upsilon^{id}\left(t, y_t^i, [x_t^{i,uc}, 0], a_t, r_t\right) = \upsilon^{ag}\left(t, y_t^i, [x_t^{i,uc}, 0], a_t, r_t\right).\end{aligned}$$

As the drift and volatility depend either explicitly or implicitly on the optimal controls, these conditions ultimately translate into constraints on the optimal controls. We define the operator for the state boundary condition $\mathcal{H}^{SBC}[y^i, r](t, x^{i,uc}, a)$ mapping the functions $[y^i, r]$ into a function of $(t, x^{i,uc}, a)$ by

$$\mathcal{H}^{SBC}[y^i, r] = \lim_{e^i \to 0} \sum_{n=1}^{N^e} |\min(0, \psi_n)| + \sum_{n=1}^{N^e} \sum_{o \in \{id, ag\}} \Sigma^o_{N^x - N^e + n, N^x - N^e + n} \tag{12}$$

with Σ^o as in (8). Setting this state boundary condition to zero ensures that each of its terms vanishes.

Since we consider an infinite horizon problem, the HJB equation does not feature

a terminal boundary condition like the bequest function would constitute in a finite horizon setting. Nevertheless, our numerical algorithm relies on a restriction at the end of the time horizon. Thus, we consider the transversality condition.[11] Define the operator for the terminal boundary condition by $\mathcal{H}^{TBC}[J](x^i, a)$ mapping the function J into a function of (x^i, a) such that

$$\mathcal{H}^{TBC}[J, y^i](x^i, a) = \lim_{T\to\infty} \mathbb{E}\left[e^{-\rho T} J(T, x_T^i, a_T)\middle| x_0^i = x^i, a_0 = a\right] \tag{13}$$

where x_T^i is the state occurring at time T when starting at a specific initial value and following the control y^i. Overall, the HJB equation and its boundary conditions are given by

$$\begin{aligned} 0 =& \max_{y^i} \left\{\mathcal{H}^{HJB}[J, y^i, r](t, x^i, a)\right\} \\ =& \mathcal{H}^{TBC}[J, y^i](x^i, a) = \mathcal{H}^{SBC}[J, y^i, r](t, x^{i,uc}, a). \end{aligned} \tag{14}$$

The optimal controls $y^{i,*}$ feature the same function arguments as the operator $\mathcal{H}^{HJB}$ such that $y_t^{i,*} = y^{i,*}[J, r](t, x_t^i, a_t)$. This notation signifies that, given $[J, r]$, the agents' optimal choices are functions of time and the state variables $y^{i,*}(t, x_t^i, a_t)$. The w-th element of the optimal control vector $y^{i,*} = [y_1^{i,*}, \ldots, y_{N^y}^{i,*}]^T$, $w \in \{1, \ldots, N^y\}$, satisfies the first-order condition

$$\begin{aligned} 0 =& \frac{\partial}{\partial y_w^i} \mathcal{H}^{HJB}[J, y^i, r](t, x^i, a) \\ =& \frac{\partial}{\partial y_w^i} \left\{ u + \sum_{n=1}^{N^x} J_{x_n^i} \mu_n + \frac{1}{2} \left[\sum_{n,n'=1}^{N^x} \sum_{o\in\{id,ag\}} J_{x_n^i x_{n'}^i} \Sigma_{nn'}^o + \sum_{n=1}^{N^x} \sum_{n'=1}^{N^a} J_{x_n^i a_{n'}} \Sigma_{nn'}^{xa} \right] \right\}. \end{aligned} \tag{15}$$

Therefore, taking the returns r as given, the indirect utility function $J[r](t, x_t^i, a_t)$ and the optimal control function $y^{i,*}[r](t, x_t^i, a_t)$ are jointly determined by the HJB equation and boundary conditions in (14) and the first-order condition in (15). To simplify the our numerical algorithm, we make the following assumption.

Assumption 1 (Optimal controls in closed form)**.** *Denote the spatial gradient of the indirect utility function by* $DJ = [J_{x_1^i}, \ldots, J_{x_{N^x}^i}, J_{a_1}, \ldots, J_{a_{N^a}}]^T$. *Let* D^2J *denote the corresponding Hessian matrix collecting all second-order spatial derivatives. Let*

[11] See e.g. Duffie (2001, p.213) for an explanation.

$\mathcal{M}^y[DJ, D^2J, r](t, x^i, a)$ *denote an operator mapping the returns, the gradient and the Hessian into a real-valued* N^y*-dimensional function of* (t, x^i, a)*. Assume that the first-order conditions in (15) are given in closed form such that*

$$y^{i,*}[r](t, x_t^i, a_t) = \mathcal{M}^y[DJ, D^2J, r](t, x_t^i, a_t). \tag{16}$$

Example. *Households in our example model take wages* w_t*, returns* $[r_t^b, r_t^k]$ *and the collateral policy* ϕ_t *as given. The HJB equation and boundary conditions are given by*

$$\begin{aligned}
0 =& -\rho J + \frac{(c^{i,*})^{1-\gamma}}{1-\gamma} + J_t + \kappa_z(\bar{z} - z^i)J_{z^i} + \left(r_t^k + \iota^{i,*}\right)k^i J_{k^i} + \Big\{w_t e^{z^i} + \left(e^i - \phi_t k^i + \bar{b}\right) r_t^b \\
&+ k^i\left[\frac{d}{dt}\phi_t + \phi_t r_t^k - (1-\phi_t)\iota^{i,*} - \Psi(\iota^{i,*})\right] - c^{i,*}\Big\} J_{e^i} + \frac{1}{2}\left[\sigma_z^2 J_{z^i z^i} + (k^i)^2\sigma_K^2 J_{k^i k^i}\right] \\
&+ \frac{1}{2}\phi_t^2(k^i)^2\sigma_K^2 J_{e^i e^i} + \omega\sigma_z k^i \sigma_K J_{z^i k^i} + \omega\sigma_z\phi_t k^i\sigma_K J_{z^i e^i} + \phi_t(k^i)^2\sigma_K^2 J_{k^i e^i} \\
0 =& \lim_{T\to\infty} J(T, z^i, k^i, e^i) \\
0 =& \lim_{e^i\to 0}\Big\{\phi_t^2(k^i)^2\sigma_K^2 + \Big|\min\Big(0, \Big[w_t e^{z^i} + \left(e^i - \phi_t k^i + \bar{b}\right) r_t^b \\
&\qquad + k^i\left[\frac{d}{dt}\phi_t + \phi_t r_t^k - (1-\phi_t)\iota^{i,*} - \Psi(\iota^{i,*})\right] - c^{i,*}\Big]\Big)\Big|\Big\}.
\end{aligned}$$

The first-order conditions correspond to (5) such that the optimal controls satisfy Assumption 1 with

$$\begin{aligned}
c^{i,*} =& (J_{e^i})^{-\frac{1}{\gamma}} \qquad (17)\\
\iota^{i,*} =& \left(\frac{\max\{0, J_{k^i} - (1-\phi_t+\psi_0)J_{e^i}\}}{\psi_1\psi_2 J_{e^i}}\right)^{\frac{1}{\psi_2 - 1}} - \left(\frac{\max\{0, (1-\phi_t-\psi_0)J_{e^i} - J_{k^i}\}}{\psi_1\psi_2 J_{e^i}}\right)^{\frac{1}{\psi_2-1}}.
\end{aligned}$$

2.2.3 Aggregate Quantities

The optimally controlled states $x_t^{i,*}$ differ across agents due to the idiosyncratic risk, which we summarize by the cross-sectional cumulative distribution function $M_t : \mathbb{R}^{N^x} \to [0,1]$.[12] Naturally, the cross-sectional density m_t is given by $M_t(x) = \int_{x^i \le x} m_t(x^i)dx^i$. Market clearing aggregates the agents' states and choices according to this distribution. The relative returns r^* clear markets such that

$$0 = \mathcal{H}^{MC}[r^*, y^{i,*}](t, M, a), \tag{18}$$

[12]Note that we use the notation M_t interchangeably to denote the c.d.f. and a Borel-measure in $\mathcal{P}(\mathbb{R}^{N^x})$ as these are equivalent representations of the cross-sectional distribution in our framework.

where $\mathcal{H}^{MC}[r, y^i](t, M, a)$ is the market-clearing operator mapping the functions $[r, y^i]$ into a function of (t, M, a). We assume that the market clearing conditions have the following structure.

Assumption 2 (Market Clearing Regularity)**.** *Suppose that there are as many market-clearing conditions as market instruments. Let $\mathcal{M}_n^r[y^i](t, M^1, \ldots, M^{N^x}, a)$ denote an operator mapping the controls y^i into a real-valued function of $(t, M^1, \ldots, M^{N^x}, a)$ where M^n denotes the n-th marginal distribution $M^n(x_n) = \int_0^{x_n} \int m(z_n, z_{-n}) dz_{-n} dz_n$. Assume that the n-th market clearing condition $0 = \mathcal{H}_n^{MC}[r, y^i](t, M, a)$ adheres to one of the following two equations.*

(*i*) *Explicit market clearing: $r_n^* = \mathcal{M}_n^r(t, M^1, \ldots, M^{N^x}, a)$, i.e., returns only depend on the current states, not the optimal controls*

(*ii*) *Implicit market clearing: $0 = \mathcal{M}_n^r[y^i](t, M^1, \ldots, M^{N^x}, a)$*

Remark. *With this assumption, we limit market clearing conditions to depend on the marginal distributions of the state variables. Thus, agents solely have to forecast the evolutions of these marginal distributions to determine prices such that the dependency structure of the joint distribution M is superfluous for price determination.*

The assumption implies that, given a law of motion of the cross-sectional distribution consistent with the agents' optimal control function $y^{i,*}(t, x_t^i, a_t)$, the equilibrium return is a function $r_t^* = r^*(t, M_t^1, \ldots, M_t^{N^x}, a_t)$.

Example. *The first two market clearing conditions of our example model in (2) explicitly define the wage and rental rate $[w_t, r_t^k]$ to clear the labor and capital market, whereas, the last condition implicitly defines the bond return r_t^b to clear the bond market. All three conditions solely depend on the marginal distributions since we can rewrite them by*

$$\mathcal{H}_n^{MC}[r, y^i](t, M) = \begin{bmatrix} F_L\left(\int e^z dM_t^z(z), \int k dM_t^k(k)\right) - w_t \\ F_K\left(\int e^z dM_t^z(z), \int k dM_t^k(k)\right) - \delta - r_t^k \\ \int e dM_t^e(e) - \phi_t \int k dM_t^k(k) + \overline{b} \end{bmatrix}.$$

Note that agent i's bond holdings are given by $b_t^i = e_t^i - \phi_t k_t^i + \overline{b}$.

The cross-sectional distribution M_t must be consistent with the agents' optimal choices. Agent i's optimally controlled state $x_t^{i,*}$ is defined by its dynamics in (7)

where, taking r_t as given, $y_t^i = y^{i,*}[r_t](t, x_t^{i,*}, a_t)$ solves (15). The cross-sectional c.d.f. consistent with this optimally controlled state is given by $M_t(x) = \mathbb{P}(x_t^{i,*} \leq x)$. The consistent law of motion of the corresponding density equals the Kolmogorov forward equation[13]

$$
\begin{aligned}
dm_t(x) = \sum_{n=1}^{N^x} \Bigg\{ & -\frac{\partial}{\partial x_n} \left[\mu_n \left(t, y^{i,*}[r_t](t,x,a_t), x, a_t, r_t \right) m_t(x) \right] \\
& + \frac{1}{2} \sum_{n'=1}^{N^x} \sum_{o \in \{id, ag\}} \frac{\partial^2}{\partial x_n \partial x_{n'}} \left[\Sigma^o_{nn'} \left(t, y^{i,*}[r_t](t,x,a_t), x, a_t, r_t \right) m_t(x) \right] \Bigg\} dt \\
& - \sum_{n=1}^{N^x} \frac{\partial}{\partial x_n} \left[\sigma_n^{ag} \left(t, y^{i,*}[r_t](t,x,a_t), x, a_t, r_t \right) m_t(x) \right] dB_t^{ag}
\end{aligned} \tag{19}
$$

for any $x \in \mathbb{R}^{N^x}$. The Kolmogorov forward equation is stochastic, i.e., it features a stochastic integral. This follows from the aggregate Brownian shock that directly enters the dynamics of x^i. Because this shock is the same for all agents, it does not wash out when taking the cross-sectional average. We refer to this property as *direct* aggregate risk ($\sigma^{ag} > 0$). This paper therefore generalizes the work of Bilal (2023) and Gu et al. (2024) which focuses on models with *indirect* aggregate risk ($\zeta > 0, \sigma^{ag} = 0$). In that case, aggregate Brownian shocks only affect the dynamics of aggregate states such that the Kolmogorov forward equation has a stochastic drift but no stochastic integral. In addition, the framework in this paper allows for controlled volatility in agent-specific states which is important for asset pricing applications.[14]

3 Deriving the Master Equation

A recursive equilibrium in this model combines (i) agent optimality, (ii) a consistent law of motion of the cross-sectional distribution and (iii) market clearing. The master equation combines (i) and (ii) so that equilibrium can be fully described by a set of equations $0 = \mathcal{H}[J, y^{i,*}, r^*]$. This section derives the master equation in a heuristic manner. The goal is to make the rigorous derivations in Cardaliaguet et al. (2019)

[13] This corresponds to Equation (2.4) in Carmona and Delarue (2018b) who derive it in Subsection 2.1.2 on p. 110-112. In Cardaliaguet et al. (2019), this equation corresponds to Equation (4.2).

[14] To the best of our knowledge, Huang (2024) is the only other work to consider controlled volatility in a HA setting.

and Carmona and Delarue (2018a,b) accessible for a wider quantitative economics audience and offer a recipe to apply the methodology in various contexts.

Each agent is assumed infinitesimal, so that the deviation of a single agent from the optimal policy function $y^{i,*}(t, x_t^i, a_t)$ does not change the cross-sectional distribution or the returns. Thus, in deriving the first-order condition for the optimal controls $y_t^{i,*}$ in (15), it suffices for agents to assume that the value function J solely depends on (t, x_t^i, a_t). The dependency on t summarizes any deterministic time variation and the time variation induced by the evolution of the cross-sectional distribution in (19) and changes in the returns determined through market clearing in (18). Incorporating the consistent law of motion into the HJB equation is, therefore, achieved by substituting t with $[t, M_t]$ in the value function $J(t, M_t, x_t^i, a_t)$.

The master equation then results from adjusting the HJB equation in (10) accordingly. The substitution of t with $[t, M_t]$ affects the term $\mathbb{E}_t[dJ]/dt$. Normally, Itô's formula delivers the dynamics of $dJ(t, x_t^i, a_t)$ such that[15]

$$\begin{aligned} dJ =& J_t dt + (J_{x^i})^T dx_t^i + (J_a)^T da_t \\ &+ \frac{1}{2}\left[(dx_t^i)^T J_{x^i x^i} dx_t^i + (da_t)^T J_{aa} da_t\right] + (dx_t^i)^T J_{x^i a} da_t. \end{aligned}$$

Substituting t with $[t, M_t]$ requires the extension of this formula with distributional derivatives. However, the cross-sectional c.d.f. M_t is not a standard Itô process such that one cannot simply apply Itô's formula to derive the dynamics of $dJ(t, M_t, x_t^i, a_t)$.[16] To derive the term $\mathbb{E}_t[dJ(t, M_t(x), x_t^i, a_t)]/dt$ in the HJB equation (10), we first need to define the derivatives of the function J w.r.t. the distribution M_t.

[15] We use matrix-vector notation here to highlight that Itô's formula can be interpreted as a second-order Taylor approximation, i.e., $DJ = [J_{x^i}, J_a]^T$ denotes the spatial gradient and the corresponding Hessian matrix is $D^2 J = \begin{bmatrix} J_{x^i x^i} & J_{x^i a} \\ J_{x^i a} & J_{aa} \end{bmatrix}$.

[16] For instance, while the standard Itô process of the aggregate state a_t is a real vector at any fixed time point t, the distribution $M_t(x)$ is a real-valued function instead. The distribution $M_t(x)$ and, similarly, the density $m_t(x)$ are generalized stochastic processes with a larger index set $(t, x) \in [0, \infty) \times \mathbb{R}^{N^x}$. Generalized stochastic processes are called random fields which have been introduced in economics by Proehl (2025b) in the context of the discrete-time heterogeneous-agent RBC model as in Krusell and Smith (1998).

3.1 Distributional Derivatives

Definition 3 (Distributional Derivatives (see Cardaliaguet et al., 2019, Definitions 2.2.1, 2.2.2 and Section 2.2.2)). *Consider a function $f : \mathcal{P}(\mathbb{R}^{N^x}) \to \mathbb{R}$. The distributional derivative of f w.r.t. M is defined by the Riesz representation given by*

$$\lim_{s \to 0} \frac{f(M + s(M' - M)) - f(M)}{s} = \int_{\mathbb{R}^{N^x}} f_M[M](x) \left(m'(x) - m(x)\right) dx,$$

where $m = \frac{d}{dx} M$ and $m' = \frac{d}{dx} M'$ denote the respective densities. The intrinsic distributional derivative is defined by

$$f_{x_n M}[M](x) := \frac{\partial}{\partial x_n} f_M[M](x), \, \forall n \in \{1, \ldots, N^x\}.$$

The second-order derivatives are similarly defined by

$$\lim_{s \to 0} \frac{f_M[M + s(M' - M)](x) - f_M[M](x)}{s} = \int_{\mathbb{R}^{N^x}} f_{MM}[M](x, x') \left(m'(x') - m(x')\right) dx'$$

and

$$f_{x_n x'_{n'} MM}[M](x, x') := \frac{\partial^2}{\partial x_n \partial x'_{n'}} f_{MM}[M](x, x'), \, \forall n, n' \in \{1, \ldots, N^x\}.$$

The first-order derivative $f_M[M](x)$ is an operator mapping the c.d.f. M into a real-valued function of x. This should be interpreted as the change of f if the density $m(.) = \frac{d}{dx} M(.)$ only changes at the outcome x and nowhere else. Thus, $f_M[M](x)$ can be interpreted as the change in the function f when the number of agents that currently observe the outcome x changes. Similarly, the second-order derivative evaluates how f changes when the density $m(.)$ changes at x and x' but nowhere else. Hence, $f_{MM}[M](x, x')$ can be interpreted as the change in the function f when the two groups of agents who observe the outcome x and x', respectively, both change in size.

Cardaliaguet et al. (2019) not only define the derivative of a function f w.r.t. M but also the intrinsic derivative which offers yet another interpretation of distributional derivatives. The c.d.f. $M_t(x) = \mathbb{P}(x_t^i \leq x)$ essentially counts how many agents $i \in [0, 1]$ observe a state x_t^i lower than or equal to x at time t, i.e., x_t^i is the random variable of possible outcomes that any agent could have at time

t.[17] More generally, any distribution function is a function of its corresponding random variable $M = \mathbb{P}(x^i \leq .) =: \mathcal{L}(x^i)$ such that we can rewrite a function $f(M) = f(\mathcal{L}(x^i)) =: \tilde{f}(x^i)$.[18] The Fréchet derivative of $\tilde{f}$ w.r.t. the random variable x^i delivers an alternative definition of distributional derivatives, see Carmona and Delarue (2018a, Definition 5.22) or Cardaliaguet et al. (2019, Equation (A.5)). Furthermore, Cardaliaguet et al. (2019, Proposition A.1.2) show that this alternative definition coincides with the intrinsic distributional derivative given in Definition 3. Therefore, $f_{x_n M}[M](x)$ can be interpreted as the change in the function f when all agents currently observing the state $x = [x_n, x_{-n}]$ switch to a different outcome $\tilde{x} = [\tilde{x}_n, x_{-n}]$ in the n-th element of the state vector. Similarly, $f_{x_n x'_{n'} MM}[M](x, x')$ can be interpreted as the change in the function f when all agents that currently observe x switch to some $\tilde{x} = [\tilde{x}_n, x_{-n}]$ and all agents that currently observe x' switch to some $\tilde{x}' = [\tilde{x}'_{n'}, x'_{-n'}]$.

3.2 Distributional Dynamics

Given the distributional derivatives, we can now derive the evolution of a function $f(M_t)$ over time. Fix two distributions M and M' and define the function $g : \mathbb{R} \to \mathbb{R}$ by $g(v) = f(M + v(M' - M))$. Due to Definition 3, we obtain

$$\frac{d}{dv} g(v) = \int_{\mathbb{R}^{N^x}} f_M[M + v(M' - M)](x) \left(m'(x) - m(x)\right) dx$$
$$\frac{d^2}{dv^2} g(v) = \int_{\mathbb{R}^{2N^x}} f_{MM}[M + v(M' - M)](x, x') \left(m'(x') - m(x')\right) \left(m'(x) - m(x)\right) dx' dx.$$

The second-order Taylor approximation $g(1) = g(0) + \frac{d}{dv} g(0) + \frac{1}{2} \frac{d^2}{dv^2} g(0)$ results in

$$f(M') - f(M) = \int_{\mathbb{R}^{N^x}} f_M[M](x) \left(m'(x) - m(x)\right) dx + \frac{1}{2} \int_{\mathbb{R}^{2N^x}} f_{MM}[M](x, x') \left(m'(x') - m(x')\right) \left(m'(x) - m(x)\right) dx' dx.$$

[17] As an Itô process, the agent-specific state in (7) is a square-integrable stochastic process implying that x^i_t for a fixed time point t is a square-integrable random variable, i.e., $x^i_t \in L^2(\Omega, \mathcal{F}_t, \mathbb{P}; \mathbb{R}^{N^x})$.

[18] The fact that functions of distributions can be rewritten as functions of random variables has also been exploited in Proehl (2025b) in a discrete-time context. We refer to this paper for a detailed account of this correspondence.

Evaluating this equation at the specific distributions $M = M_t$ and $M' = M_{t+dt}$ and taking the limit of $dt \to 0$ yields

$$\begin{aligned} df(M_t) = &\int_{\mathbb{R}^{N^x}} f_M[M_t](x) dm_t(x) dx \\ &+ \frac{1}{2} \int_{\mathbb{R}^{2N^x}} f_{MM}[M_t](x, x') dm_t(x') dm_t(x) dx' dx. \end{aligned} \tag{20}$$

Alternatively, we can express this expansion in terms of intrinsic distributional derivatives.

Proposition 4 (Itô's Formula for Distributions)**.** *Equation (20) is equivalent to*[19]

$$\begin{aligned} df(M_t) = &\left\{ \sum_{n=1}^{N^x} \int_{\mathbb{R}^{N^x}} f_{x_n M}[M_t](x) \mu_n \left(t, y^{i,*}[r_t](t, x, a_t), x, a_t, r_t\right) m_t(x) dx \right. \\ &+ \frac{1}{2} \sum_{n,n'=1}^{N^x} \sum_{o \in \{id, ag\}} \int_{\mathbb{R}^{N^x}} f_{x_{n'} x_n M}[M_t](x) \Sigma^{o,*}_{nn'}(t, x, x) m_t(x) dx \\ &\left. + \frac{1}{2} \sum_{n,n'=1}^{N^x} \int_{\mathbb{R}^{2N^x}} f_{x_n x'_{n'} MM}[M_t](x, x') \Sigma^{ag,*}_{nn'}(t, x, x') m_t(x) m_t(x') dx dx' \right\} dt \\ &+ \sum_{n=1}^{N^x} \left\{ \int_{\mathbb{R}^{N^x}} f_{x_n M}[M_t](x) \sigma^{ag}_n \left(t, y^{i,*}[r_t](t, x, a_t), x, a_t, r_t\right) m_t(x) dx \right\} dB^{ag}_t \end{aligned} \tag{21}$$

with the optimally controlled covariance function for $o \in \{id, ag\}$

$$\Sigma^{o,*}(t, x, x') = \sigma^o \left(t, y^{i,*}[r_t](t, x, a_t), x, a_t, r_t\right) \sigma^o \left(t, y^{i,*}[r_t](t, x', a_t), x', a_t, r_t\right)^T. \tag{22}$$

The mechanics behind the proof of this equivalence is illustrated most clearly in the model in which aggregate risk enters through TFP as in Bilal (2023).

Example 5 (Case Without Direct Aggregate Risk as in Bilal (2023); Krusell and Smith (1998))**.** *This model translates into our general set up with* $N^x = 2$, $d^{id} =$

[19] This follows from Theorem 4.14 in Carmona and Delarue (2018b) and Theorem A.1 in Cardaliaguet et al. (2019).

$d^{ag} = 1$, $\zeta > 0$ *and* $\sigma^{ag} = 0$. *The household problem reads*

$$\begin{aligned}
&\max_{\{c_t^i\}_{t\geq 0}} \mathbb{E}_0 \left[\int_0^\infty e^{-\rho t} \frac{(c_t^i)^{1-\gamma}}{(1-\gamma)} dt \right] \\
&\quad s.t.\ dz_t^i = \kappa_z(\bar{z} - z_t^i)dt + \sigma_z dB_t^i, \\
&\qquad dk_t^i = \left\{ w_t e^{z_t^i} + r_t^k k_t^i - c_t^i \right\} dt, \\
&\qquad da_t = \kappa_a(\bar{a} - a_t)dt + \sigma_a dB_t^{ag}, \\
&\qquad k_t^i \geq 0,
\end{aligned}$$

where $\{e^{a_t}\}_{t\geq 0}$ *is the total factor productivity entering the production function. The aggregate wage and return are given by*

$$\begin{aligned}
r_t^k &= \alpha e^{a_t} \left(\frac{\int e^z m_t(z,k) dk dz}{\int k m_t(z,k) dk dz} \right)^{1-\alpha} - \delta \\
w_t &= (1-\alpha) e^{a_t} \left(\frac{\int k m_t(z,k) dk dz}{\int e^z m_t(z,k) dk dz} \right)^{\alpha}.
\end{aligned}$$

Denote the idiosyncratic variables by $x = [z, k]$ *and returns by* $r = [r^k, w]$ *in line with the notation of our general setup. Then, the Kolmogorov forward equation reads*

$$dm_t(x) = \left\{ \sum_{n=1}^{2} \frac{d}{dx_n} \left[-\mu_n \left(c^{i,*}[r_t](x, a_t), x, r_t \right) m_t(x) \right] + \frac{1}{2} \frac{d^2}{dz^2} \left[\sigma_z^2 m_t(x) \right] \right\} dt.$$

As this law of motion does not contain a stochastic integral, the second-order term in

(20) is obsolete such that

$$
\begin{aligned}
df(M_t) &= \int_{\mathbb{R}^{N^x}} f_M[M_t](x) dm_t(x) dx \\
&= \left\{ \int_{\mathbb{R}^2} f_M[M_t](x) \sum_{n=1}^{2} \frac{d}{dx_n} \left[-\mu_n \left(c^{i,*}[r_t](x, a_t), x, r_t \right) m_t(x) \right] dx \right. \\
&\quad \left. + \frac{1}{2} \int_{\mathbb{R}^2} f_M[M_t](x) \frac{d^2}{dz^2} \left[\sigma_z^2 m_t(x) \right] dx \right\} dt \\
&= \left\{ \int_{\mathbb{R}} \sum_{n=1}^{2} \frac{d}{dx_n} \left[f_M[M_t](x) \right] \mu_n \left(c^{i,*}[r_t](x, a_t), x, r_t \right) m_t(x) dx \right. \\
&\quad \left. + \frac{1}{2} \int_{\mathbb{R}} \frac{d^2}{dz^2} \left[f_M[M_t](x) \right] \sigma_z^2 m_t(x) dx \right\} dt \\
&= \left\{ \int_{\mathbb{R}} \sum_{n=1}^{2} f_{x_n M}[M_t](x) \mu_n \left(c^{i,*}[r_t](x, a_t), x, r_t \right) m_t(x) dx \right. \\
&\quad \left. + \frac{1}{2} \int_{\mathbb{R}} f_{zzM}[M_t](x) \sigma_z^2 m_t(x) dx \right\} dt
\end{aligned}
$$

where the third equality results from integration by parts and the last equality follows from Definition 3.

The typical approach to the master equation in economics is to use the representation in terms of distributional derivatives (see Schaab, 2020; Bilal, 2023; Gu et al., 2024). The example illustrates that this requires the computation of derivatives of the drift and dispersion terms which can be numerically cumbersome. Instead, we use the intrinsic distributional derivatives which are simply multiplied with the state's drift and covariance terms.

3.3 The Master Equation

The master equation results from adjusting the last term of the HJB equation in (10), i.e., $\mathbb{E}_t[dJ]/dt$. To derive $dJ(t, M_t, x_t^i, a_t)$, the standard Itô formula is combined with Itô's formula for distributions given in (21) resulting in[20]

$$
\begin{aligned}
dJ\left(t, M_t, x_t^i, a_t\right) =& d\mathcal{A}_{t,x^i,a}\left(t, M_t, x_t^i, a_t\right) + d\mathcal{A}_M\left(t, M_t, x_t^i, a_t\right) \\
&+ d\mathcal{A}_{M,x^i,a}\left(t, M_t, x_t^i, a_t\right).
\end{aligned} \tag{23}
$$

[20] This follows from Theorem 4.17 in Carmona and Delarue (2018b).

The first term summarizes the variation stemming from t and the state processes $\{x_t^i, a_t\}_{t\geq 0}$. The other two terms account for the dependence on the distribution $\{M_t(x)\}_{t\geq 0, x\in\mathbb{R}^{N^x}}$. The term $d\mathcal{A}_M$ results from the expansion in (21), whereas, the term $d\mathcal{A}_{M,x^i,a}$ takes the interaction of the distribution $\{M_t(x)\}_{t\geq 0, x\in\mathbb{R}^{N^x}}$ with the state processes $\{x_t^i, a_t\}_{t\geq 0}$ into account yielding

$$
\begin{aligned}
&d\mathcal{A}_{M,x^i,a}\left(t, M_t, x_t^i, a_t\right) = \\
&\frac{1}{2}\Bigg\{\sum_{n,n'=1}^{N^x}\int_{\mathbb{R}^{N^x}} J_{x_n^i x_{n'} M}[M_t]\left(x; t, x_t^i, a_t\right)\Sigma_{nn'}^{ag,*}\left(t, x_t^i, x\right) m_t(x)dx \\
&+\sum_{n=1}^{N^x}\sum_{n'=1}^{N^a}\int_{\mathbb{R}^{N^x}} J_{a_{n'} x_n M}[M_t]\left(x; t, x_t^i, a_t\right)\Sigma_{nn'}^{xa}\left(t, y_t^{i,*}, x, a_t, r_t^*\right) m_t(x)dx\Bigg\}dt
\end{aligned}
$$

with $\Sigma^{ag,*}$ as in (22) and Σ^{xa} as in (8). Collecting the drift terms $\mathbb{E}_t[dJ]/dt$ yields the master operator. We define it by $\mathcal{H}^M[J, y^i, r](t, M, x^i, a)$ mapping the functions $[J, y^i, r]$ into a function of (t, M, x^i, a) such that

$$
\begin{aligned}
&\mathcal{H}^M[J, y^i, r](t, M, x^i, a) = u - \rho J + J_t \qquad (24)\\
&+\sum_{n=1}^{N^x}\int_{\mathbb{R}^{N^x}} J_{x_n M}[M]\left(x; t, x^i, a\right)\mu_n\left(t, y^i, x, a, r\right) m(x)dx \\
&+\frac{1}{2}\sum_{n,n'=1}^{N^x}\sum_{o\in\{id,ag\}}\int_{\mathbb{R}^{N^x}} J_{x_{n'} x_n M}[M]\left(x; t, x^i, a\right)\Sigma_{nn'}^{o,*}(t, x, x)m(x)dx \\
&+\frac{1}{2}\sum_{n,n'=1}^{N^x}\int_{\mathbb{R}^{2N^x}} J_{x_n x'_{n'} MM}[M]\left(x, x'; t, x^i, a\right)\Sigma_{nn'}^{ag,*}(t, x, x')m(x)m(x')dxdx' \\
&+\frac{1}{2}\sum_{n=1}^{N^x}\sum_{n'=1}^{N^a}\int_{\mathbb{R}^{N^x}} J_{a_{n'} x_n M}[M]\left(x; t, x^i, a\right)\Sigma_{nn'}^{xa}\left(t, y^i, x, a, r\right) m(x)dx \\
&+\frac{1}{2}\sum_{n,n'=1}^{N^x}\int_{\mathbb{R}^{N^x}} J_{x_n^i x_{n'} M}[M]\left(x; t, x^i, a\right)\Sigma_{nn'}^{ag,*}\left(t, x^i, x\right) m(x)dx \\
&+\sum_{n=1}^{N^a} J_{a_n}\eta_n\left(t, a, r\right) + \sum_{n=1}^{N^x} J_{x_n^i}\mu_n\left(t, y^i, x^i, a, r\right) + \frac{1}{2}\sum_{n,n'=1}^{N^a} J_{a_n a_{n'}}\Sigma_{nn'}^{a}\left(t, a, r\right) \\
&+\frac{1}{2}\sum_{n,n'=1}^{N^x}\sum_{o\in\{id,ag\}} J_{x_n^i x_{n'}^i}\Sigma_{nn'}^{o,*}\left(t, x^i, x^i\right) + \frac{1}{2}\sum_{n=1}^{N^x}\sum_{n'=1}^{N^a} J_{x_n^i a_{n'}}\Sigma_{nn'}^{xa}\left(t, y^i, x^i, a, r\right)
\end{aligned}
$$

with the covariance matrices $\Sigma^{\cdot}$ as in (8) and $\Sigma^{\cdot,*}$ as in (22). When inserting the optimal controls $y^{i,*}$ that solve the first-order conditions in (15), the master equation sets the master operator to zero. Overall, we can now define the equilibrium as follows.

Definition 6 (Recursive Equilibrium)**.** *The recursive equilibrium of this heterogeneous agent economy is a set of controls* $\{y_t^{i,*}\}_{t\geq 0, i\in[0,1]}$*, returns* $\{r_t^*\}_{t\geq 0}$ *and cross-sectional distributions* $\{M_t\}_{t\geq 0}$ *such that the following conditions hold.*

(*i*) *Given the returns* $\{r_t^*\}_{t\geq 0}$*, agents choose their controls* $y_t^{i,*} = y^{i,*}(t, M, x^i, a)$ *to optimize their indirect utility function* $J(t, M, x^i, a)$ *such that*

$$0 = \frac{\partial}{\partial y_n^i}\mathcal{H}^{HJB}[J, y^{i,*}, r^*](t, M, x^i, a),\ n \in \{1, \dots, N^y\}, \tag{25}$$

$$0 = \mathcal{H}^{M}[J, y^{i,*}, r^*](t, M, x^i, a) \tag{26}$$

$$0 = \mathcal{H}^{TBC}[J, y^{i,*}](M, x^i, a) \tag{27}$$

$$0 = \mathcal{H}^{SBC}[J, y^{i,*}, r^*](t, M, x^{i,uc}, a), \tag{28}$$

with $\mathcal{H}^{HJB}$ *as in (11),* $\mathcal{H}^{M}$ *as in (24),* $\mathcal{H}^{TBC}$ *as in (13) and* $\mathcal{H}^{SBC}$ *as in (12).*

(*ii*) *Given the controls* $\{y_t^{i,*}\}_{t\geq 0, i\in[0,1]}$*, the returns* $r_t^* = r^*(t, M^1, \dots, M^{N^x}, a)$ *clear markets such that*

$$0 = \mathcal{H}^{MC}[r^*, y^{i,*}](t, M, a). \tag{29}$$

4 Approximating the Equilibrium Conditions

Given Assumption 2, the equilibrium returns are determined by the marginal distributions only, that is, $r^* = r^*(t, M^1, \dots, M^{N^x}, a)$. Therefore, we conjecture and numerically verify that the optimal controls of agents $y_t^{i,*} = y^{i,*}(t, M^1, \dots, M^{N^x}, x^i, a)$ and their value function $J(t, M^1, \dots, M^{N^x}, x^i, a)$ solely depend on the marginal distributions as well. To approximate the equilibrium conditions (25)-(29), we first simplify the master equation in (26) in accordance with our conjecture. The simplified master equation remains a non-standard PDE due to the marginal distributional derivatives. Thus, in a second step, we propose an approximation of the marginal distributional derivatives. Importantly, this allows us to approximate the simplified master equation with a standard PDE.

4.1 Simplifying the Master Equation

Sklar's theorem[21] tells us that any joint distribution function can be rewritten as a joint distribution function of uniform variables, i.e., $M_t(x) = C_t(u_x)$ where the uniform variable $u_x = [M_t^1(x_1), \ldots, M_t^{N^x}(x_{N^x})]$ consists of the marginal distributions $M_t^n(x_n) = \int m_t(x) dx_{-n}$ with $x = [x_n, x_{-n}]$. The distribution C_t is called a copula. By differentiation, the joint density can be rewritten using the copula density and the marginal densities such that

$$m_t(x) = c_t(u) \cdot m_t^1(x_1) \cdot \ldots \cdot m_t^{N^x}(x_{N^x}). \tag{30}$$

Therefore, we can separate the joint density into the marginal densities and the dependency structure. The joint density changes due to changes in the copula and/or changes in each marginal density. The distributional derivative f_M can, hence, be further refined with the corresponding partial distributional derivatives.

Definition 7 (Partial Distributional Derivatives)**.** *Consider a function $f : \mathcal{P}(\mathbb{R}^{N^x}) \mapsto \mathbb{R}$ with the total distributional derivative $f_M[M](x)$ as defined in Definition 3. Set $u_z = [M^1(z_1), \ldots, M^{N^x}(z_{N^x})]$ and let*

$$M(x) = \int_{z \leq x} m(z) dz = \int_{z \leq x} c(u_z) \cdot m^1(z_1) \cdot \ldots \cdot m^{N^x}(z_{N^x}) dz,$$

where c denotes the copula density and $M^n(x_n) = \int_{z_n \leq x_n} m^n(z_n) dz_n$ denotes the n-th marginal distribution of $M \in \mathcal{P}(\mathbb{R}^{N^x})$. Then, the partial distributional derivatives w.r.t. the copula and the marginal distributions are defined by

$$\begin{aligned}
&\int_{\mathbb{R}^{N^x}} f_M[M](x)(m'(x) - m(x)) dx \\
&= \int_{\mathbb{R}^{N^x}} \left\{ f_C[M](x) \left(c'\left(u_x\right) - c\left(u_x\right) \right) \prod_{n=1}^{N^x} m^n(x_n) \right. \\
&\quad \left. + \sum_{n=1}^{N^x} f_{M^n}[M](x_n) \left(m^{n'}(x_n) - m^n(x_n) \right) c\left(u_x\right) \prod_{j \neq n, j=1}^{N^x} m^j(x_j) \right\} dx.
\end{aligned}$$

[21] See for instance Jaworski et al. (2010, Theorem 1.4.1).

The second-order derivatives are defined by

$$
\begin{aligned}
&\int_{\mathbb{R}^{N^x}} f_{MM}[M](x,x')(m'(x') - m(x'))dx' \\
&= \int_{\mathbb{R}^{N^x}} \Bigg\{ f_{CC}[M](x,x') \left(c'\left(u_{x'}\right) - c\left(u_{x'}\right)\right) \prod_{n=1}^{N^x} m^n(x'_n) \\
&\quad + 2\sum_{j=1}^{N^x} f_{M^jC}[M](x,x'_j) \left(m^{j'}(x'_j) - m^j(x'_j)\right) c\left(u_{x'}\right) \prod_{k\neq j,k=1}^{N^x} m^k(x'_k) \\
&\quad + \sum_{n,j=1}^{N^x} f_{M^jM^n}[M](x_n,x'_j) \left(m^{j'}(x'_j) - m^j(x'_j)\right) c\left(u_{x'}\right) \prod_{k\neq j,k=1}^{N^x} m^k(x'_k) \Bigg\} dx'.
\end{aligned}
$$

The intrinsic derivatives are given as before, i.e., $f_{x_nC}[M](x) = \frac{\partial}{\partial x_n} f_C[M](x)$ *and* $f_{x_nM^j}[M](x_j) = \frac{\partial}{\partial x_n} f_{M^j}[M](x_j)$ *etc. for any* $n, j \in \{1,\dots,N^x\}$.

Remark. *Note that spliting the joint distribution into marginals and dependency structure results in zero intrinsic derivatives for the marginals, i.e.,* $f_{x_nM^j}[M](x_j) = 0$ *whenever* $n \neq j$ *and* $f_{x_nx'_{n'}M^jC}[M](x,x'_j) = 0$ *whenever* $n' \neq j$. *Furthermore,* $f_{x_nx'_{n'}M^jM^k}[M](x_k,x'_j) = 0$ *whenever* $n \neq k$ *and* $n' \neq j$.

Similarly to Section 3, the master equation can be derived in terms of the partial distributional derivatives. This extended master equation can then be simplified with the following conjecture.

Assumption 8 (Dependency Conjecture)**.** *We conjecture that the agents' equilibrium value function solely depends on the marginal distributions and not on the dependency structure given by the copula, i.e., we assume that all partial distributional derivatives of the agents' value function w.r.t. the copula are zero.*

We argue that this assumption is reasonable because returns only depend on the marginal distributions due to Assumption 2. Ultimately, agents forecast the evolution of returns to optimize their controls for which the dependency structure is irrelevant.

Theorem 9 (Simplified Master Equation)**.** *Set* $x = [x_n, x_{-n}] \in \mathbb{R}^{N^x}$ *for some* $n \in \{1,\dots,N^x\}$. *Denote the conditional cross-sectional expectation* $\mathbf{E}[f(x)|x_n] = \int f(x)\frac{m(x)}{m^n(x_n)}dx_{-n}$ *by* $\overline{f}(x_n)$. *Given Assumption 8, the master equation operator in (24) solely relies on the partial distributional derivatives w.r.t. the marginal distribu-*

tions. It reads

$$\mathcal{H}^M[J, y^i, r](t, M, x^i, a) = u - \rho J + J_t \tag{31}$$
$$+ \sum_{n=1}^{N^x} \int_{\mathbb{R}} J_{x_n M^n}[M]\left(x_n; t, x^i, a\right) \overline{\mu_n}\left(t, y^i, x_n, a, r\right) m^n(x_n) dx_n$$
$$+ \frac{1}{2} \sum_{n=1}^{N^x} \sum_{o \in \{id, ag\}} \int_{\mathbb{R}} J_{x_n x_n M^n}[M]\left(x_n; t, x^i, a\right) \overline{\Sigma_{nn}^{o,*}}\left(t, x_n, x_n\right) m^n(x_n) dx_n$$
$$+ \frac{1}{2} \sum_{n,n'=1}^{N^x} \int_{\mathbb{R}^2} J_{x_n x'_{n'} M^{n'} M^n}[M]\left(x_n, x'_{n'}; t, x^i, a\right) \Sigma_{nn'}^{M,*}(t, x_n, x'_{n'}) m^{n'}(x'_{n'}) m^n(x_n) dx'_{n'} dx_n$$
$$+ \frac{1}{2} \sum_{n=1}^{N^a} \sum_{n'=1}^{N^x} \int_{\mathbb{R}} J_{a_n x_{n'} M^{n'}}[M]\left(x_{n'}; t, x^i, a\right) \Sigma_{nn'}^{aM,*}\left(t, x_{n'}\right) m^{n'}(x_{n'}) dx_{n'}$$
$$+ \frac{1}{2} \sum_{n,n'=1}^{N^x} \int_{\mathbb{R}} J_{x_n^i x_{n'} M^{n'}}[M]\left(x_{n'}; t, x^i, a\right) \Sigma_{nn'}^{x^i M,*}(t, x^i, x_{n'}) m^{n'}(x_{n'}) dx_{n'}$$
$$+ \sum_{n=1}^{N^a} J_{a_n} \eta_n\left(t, a, r\right) + \sum_{n=1}^{N^x} J_{x_n^i} \mu_n\left(t, y^i, x^i, a, r\right) + \frac{1}{2} \sum_{n,n'=1}^{N^a} J_{a_n a_{n'}} \Sigma_{nn'}^{a}\left(t, a, r\right)$$
$$+ \frac{1}{2} \sum_{n,n'=1}^{N^x} \sum_{o \in \{id, ag\}} J_{x_n^i x_{n'}^i} \Sigma_{nn'}^{o,*}\left(t, x^i, x^i\right) + \frac{1}{2} \sum_{n=1}^{N^x} \sum_{n'=1}^{N^a} J_{x_n^i a_{n'}} \Sigma_{nn'}^{xa}\left(t, y^i, x^i, a, r\right)$$

with the covariance matrices $\Sigma^{\cdot}$ as in (8) and $\Sigma^{\cdot,*}$ as in (22), (43) and (44).

Remark. *Although the simplified master equation operator solely depends on the partial distributional derivatives w.r.t. the marginal distributions, the equation overall does rely on the dependency structure of the joint distribution because the copula density features in the conditional expectations*

$$\overline{f}(x_n) = \mathbf{E}[f(x)|x_n] = \int_{\mathbb{R}^{N^x - 1}} f(x) c(u_x) \prod_{j \neq n, j=1}^{N^x} m^j(x_j) dx_{-n}.$$

Thus, approximating the copula density is still a significant step in the numerical method. The details of its practical implementation can be found in Section 1.2 of the online appendix.

4.2 Approximating the Marginal Distributional Derivatives

The distributional terms in the simplified master equation in (31) feature the intrinsic partial distributional derivatives w.r.t. the marginal distributions. Each marginal distribution $M_t^n = \mathbb{P}(x_{n,t}^i \leq .) =: \mathcal{L}(x_{n,t}^i)$ can be written as a function of the marginal random variable.[22] Furthermore, any square-integrable random variable can be represented by projecting it onto a family of stochastic orthogonal polynomials, a method called polynomial chaos. This discretization technique has first been used in the context of heterogeneous agent models by Proehl (2025a). At a fixed time point t, the state $x_{n,t}^i$ can be represented as

$$x_{n,t}^i = \sum_{v=0}^{\infty} \varphi_{n,t}^v \phi^v(\xi_n) \tag{32}$$

where $\xi_n \in L^\infty(\Omega, \Sigma, \mathbb{P}; \mathbb{R})$ is a continuous random variable for which all moments are finite and grow slowly enough.[23] The corresponding family of orthogonal polynomials $\{\phi^v\}_{v=0}^\infty$ satisfies

$$\mathbb{1}_{\{v=u\}} b^v = \int_{\mathbb{R}} \phi^v(\xi_n)\phi^u(\xi_n) m^{\xi_n}(\xi_n) d\xi_n \tag{33}$$

such that $b^v > 0$ for any $v \in \{0, 1, \ldots\}$. The projection coefficients $\{\varphi_{n,t}^v\}_{v=0}^\infty$ are defined by

$$\varphi_{t,n}^v = \frac{1}{b^v} \int_{\mathbb{R}} x_n \phi^v(\xi(x_n)) m_t^n(x_n) dx_n = \frac{1}{b^v} \int_{\mathbb{R}} x_n \phi^v(\xi(x_n)) d\mathbb{P}(x_t^i \leq x_n) \tag{34}$$

with $\xi(x_n) = (M^{\xi_n})^{-1}(M^n(x_n))$. The second equality illustrates that the projection coefficients map the random variable $x_{n,t}^i$ into $\mathbb{R}$, i.e., we can write $\varphi_{t,n}^v = \varphi^v[x_{n,t}^i]$.

Overall, we can represent the marginal distributions

$$M_t^n = \mathcal{L}\left(x_{n,t}^i\right) = \mathcal{L}\left(\sum_{v=0}^{\infty} \varphi_{n,t}^v \phi^v(\xi)\right) = g\left(\varphi^0[x_{n,t}^i], \ldots, \varphi^\infty[x_{n,t}^i]\right)$$

by a function g of projection coefficients defined as in (34). Hence, any function of a marginal distribution is a function of the corresponding projection coefficients, i.e.,

[22] As an Itô process, the agent-specific state in (7) is a square-integrable stochastic process implying that $x_{n,t}^i$ for a fixed time point t is a square-integrable random variable, i.e., $x_{n,t}^i \in L^2(\Omega, \mathcal{F}_t, \mathbb{P}; \mathbb{R})$.

[23] See Ernst et al. (2012, Theorem 3.4) for the exact requirements on ξ to ensure convergence of the sequence in (32).

$f(M^n) = \tilde{f}(\varphi^0[x_n^i], \ldots, \varphi^\infty[x_n^i])$. It follows that the intrinsic marginal distributional derivatives can be stated in terms of the derivatives w.r.t. the projection coefficients.

Lemma 10. *When the state random variables x_n^i, $n \in \{1, \ldots, N^x\}$, are represented by polynomial chaos as in (32), the intrinsic marginal distributional derivatives of a function $f(M^1, \ldots, M^{N^x}) = \tilde{f}\left(\{\varphi_1^v\}_{v=0}^\infty, \ldots, \{\varphi_{N^x}^v\}_{v=0}^\infty\right)$ are given by*

$$f_{x_n M^n}[M](x_n) = \sum_{v=0}^{\infty} \tilde{f}_{\varphi_n^v} \frac{\phi^v\left(\xi(x_n)\right)}{b^v},$$

with $\xi(x_n) = (M^{\xi_n})^{-1}(M^n(x_n))$ and b^v as in (33). The higher-order derivatives are given by

$$\begin{aligned} f_{x_n x_n M^n}[M](x_n) &= \sum_{v=0}^{\infty} \tilde{f}_{\varphi_n^v} \frac{\phi_\xi^v\left(\xi(x_n)\right)}{b^v} \frac{m^n(x_n)}{m^{\xi_n}(\xi(x_n))} \\ f_{x_n x'_{n'} M^{n'} M^n}[M](x_n, x'_{n'}) &= \sum_{v,v'=0}^{\infty} \tilde{f}_{\varphi_n^v \varphi_{n'}^{v'}} \frac{\phi^v\left(\xi(x_n)\right)}{b^v} \frac{\phi^{v'}\left(\xi(x'_{n'})\right)}{b^{v'}}. \end{aligned}$$

Numerically, the polynomial chaos expansion in (32) will be truncated at V_n meaning that we set $\{\varphi_{t,n}^v\}_{v>V_n} = 0$.[24] Overall, using Proposition 10, we can approximate

[24] Proehl (2025a) showed in a discrete-time framework that, when increasing the truncation order V_n of the polynomial projection, the policy function of the truncated representation converges to the true policy function exponentially.

the simplified master equation operator in (31) by[25]

$$
\begin{aligned}
&\mathcal{H}^M[J, y^i, r](t, M, x^i, a) \approx u - \rho J + J_t \qquad (35)\\
&+ \sum_{n=1}^{N^x} \sum_{v=0}^{V_n} \frac{J_{\varphi_n^v}}{b^v} \int_{\mathbb{R}} \phi^v\left(\xi(x_n)\right) \overline{\mu_n}\left(t, y^i, x_n, a, r\right) m^n(x_n) dx_n \\
&+ \frac{1}{2} \sum_{n=1}^{N^x} \sum_{o \in \{id, ag\}} \sum_{v=0}^{V_n} \frac{J_{\varphi_n^v}}{b^v} \int_{\mathbb{R}} \phi_\xi^v\left(\xi(x_n)\right) \frac{m^n(x_n)}{m^{\xi_n}(\xi(x_n))} \overline{\Sigma_{nn}^{o,*}}\left(t, x_n, x_n\right) m^n(x_n) dx_n \\
&+ \frac{1}{2} \sum_{n,n'=1}^{N^x} \sum_{v=0}^{V_n} \sum_{v'=0}^{V_{n'}} \frac{J_{\varphi_n^v \varphi_{n'}^{v'}}}{b^v b^{v'}} \int_{\mathbb{R}^2} \phi^v\left(\xi(x_n)\right) \phi^{v'}\left(\xi(x'_{n'})\right) \Sigma_{nn'}^{\mathrm{M},*}(t, x_n, x'_{n'}) m^{n'}(x'_{n'}) m^n(x_n) dx'_{n'} dx_n \\
&+ \frac{1}{2} \sum_{n=1}^{N^a} \sum_{n'=1}^{N^x} \sum_{v=0}^{V_{n'}} \frac{J_{a_n \varphi_{n'}^v}}{b^v} \int_{\mathbb{R}} \phi^v\left(\xi(x_{n'})\right) \Sigma_{nn'}^{aM,*}\left(t, x_{n'}\right) m^{n'}(x_{n'}) dx_{n'} \\
&+ \frac{1}{2} \sum_{n,n'=1}^{N^x} \sum_{v=0}^{V_{n'}} \frac{J_{x_n^i \varphi_{n'}^v}}{b^v} \int_{\mathbb{R}} \phi^v\left(\xi(x_{n'})\right) \Sigma_{nn'}^{x^i \mathrm{M},*}(t, x^i, x_{n'}) m^{n'}(x_{n'}) dx_{n'} \\
&+ \sum_{n=1}^{N^a} J_{a_n} \eta_n\left(t, a, r\right) + \sum_{n=1}^{N^x} J_{x_n^i} \mu_n\left(t, y^i, x^i, a, r\right) + \frac{1}{2} \sum_{n,n'=1}^{N^a} J_{a_n a_{n'}} \Sigma_{nn'}^{a}\left(t, a, r\right) \\
&+ \frac{1}{2} \sum_{n,n'=1}^{N^x} \sum_{o \in \{id, ag\}} J_{x_n^i x_{n'}^i} \Sigma_{nn'}^{o,*}\left(t, x^i, x^i\right) + \frac{1}{2} \sum_{n=1}^{N^x} \sum_{n'=1}^{N^a} J_{x_n^i a_{n'}} \Sigma_{nn'}^{xa}\left(t, y^i, x^i, a, r\right)
\end{aligned}
$$

with $\xi(x_n) = (M^{\xi_n})^{-1}(M^n(x_n))$ and the covariance matrices $\Sigma^{\cdot}$ as in (8) and $\Sigma^{\cdot,*}$ as in (22), (43) and (44).

Remark. *Equation (35) illustrates the advantages of approximating the marginal distributions with polynomial chaos. This approximation substitutes the cross-sectional distribution M in the state space with the $(\sum_{n=1}^{N^x} V_n)$-dimensional real vector of projection coefficients $[\varphi_n^{v_n}]_{n \in \{1,\dots,N^x\}, v_n \in \{0,\dots,V_n\}}$. Therefore, the distributional derivatives can be approximated with standard derivatives of a real-valued function. Hence, we approximate the master equation with a standard, yet high-dimensional, partial differential equation.*

The seed random variables ξ_n, $n \in \{1, \dots, N^x\}$, in the Polynomial Chaos representation (32) are set exogenously. This choice presents a trade-off between the complexity of the corresponding polynomials and the speed of convergence of the se-

[25] To simplify notation we drop the distinction between a function of projection coefficients $\tilde{f}$ and a function of the distribution f.

ries in (32). For instance, one can choose ξ_n to equal a standard random variable, e.g., a standard normal.[26] Although the truncation error in (32) typically only becomes small at higher-order polynomials in that case, the family of polynomials is standard and does not need to be constructed separately. Alternatively, one can tailor ξ_n to a distribution close to the one expected in the model. In that case, truncating (32) at lower order already achieves a good approximation but the family of polynomials needs to be constructed separately. The latter is the method chosen in Proehl (2025a) as this paper relies on a tensor-grid discretization of the state space. In contrast, we choose the former approach because our numerical method relies on neural nets and thus, suffers less from the curse of dimensionality. We outline the details of the practical implementation of the polynomial chaos approximation in Section 1.1 of the online appendix.

5 Solving for the Approximated Equilibrium

The recursive equilibrium is computed by solving the first-order conditions for the optimal controls, the approximated master equation, its corresponding terminal and state boundary conditions, and market clearing. Assumptions 1 and 2 deliver the optimal controls $y^{i,*}$ and the returns associated with the explicit market clearing conditions $r^{*,ex}$ in closed form. Thus, we can reduce the number of equations by inserting these closed-form solutions directly into the approximated master equation and its boundary conditions. Overall, the main equation to be solved is the approximated master equation in (35). Previous work in computational economics using the master equation like Gu et al. (2024); Han et al. (2021); Huang (2024) directly approximates the indirect utility function and often relies on automated differentiation to obtain its derivatives. Since our master equation is of second order and high dimension, this approach is infeasible. Instead, we base our solution method on the Deep BSDE method. It was first developed by E et al. (2017); Han et al. (2018) to solve high-dimensional first-order PDEs and has been extended by Beck et al. (2019) to accommodate second-order PDEs. This method can be interpreted as a shooting algorithm that is best illustrated in the model without direct aggregate risk as in

[26] Xiu and Karniadakis (2002) present the possible standard distributions.

Example 5.[27]

Example (Case Without Direct Aggregate Risk as in Bilal (2023); Krusell and Smith (1998))**.** *The household problem provides us with the dynamics for the agent-specific and aggregate variables* $[dx_t^i, da_t]$*. Given these, the cross-sectional distribution can be approximated with the projection coefficients* $[\varphi_{n,t}^{v_n}]_{n\in\{1,2\},v_n\in\{0,\dots,V_n\}}$*. The returns* $[r_t^k, w_t]$ *are explicitly defined given the distribution's approximation. All of these variables can be simulated forward given their initial values. Additionally, one can derive the dynamics for agent i's indirect utility* $Y_t^i = J(t, [\varphi_{n,t}^{v_n}]_{n\in\{1,2\},v_n\in\{0,\dots,V_n\}}, x_t^i, a_t)$ *and marginal utility* $Z_t^i = DJ(t, [\varphi_{n,t}^{v_n}]_{n\in\{1,2\},v_n\in\{0,\dots,V_n\}}, x_t^i, a_t)$*. The terminal boundary condition pins down the terminal values of these but the initial values* $[Y_0^i, Z_0^i]$ *are unknown. The essential idea behind the Deep BSDE method is to approximate the initial utility and its marginal, then simulate forward to evaluate the terminal boundary condition and repeat this step until the approximation of* $[Y_0^i, Z_0^i]$ *minimizes the error in the terminal condition.*

This illustration distills the main idea of the method but abstracts important details which will be presented in the following. Let us first summarize the standard Deep BSDE method before we explain how to adapt it to our equilibrium setting.

5.1 The Deep BSDE Method

The Deep BSDE method exploits that PDEs are closely related to backward stochastic differential equations (BSDE). Similar to the state dynamics in (7), the indirect utility function J as well as its gradient DJ can be represented by stochastic differential equations. The difference lies in the direction. The state dynamics move forward given the initial state whereas the dynamics of J and DJ move backwards given the terminal boundary condition in (27). Given this system of forward backward stochastic differential equations (FBSDE), the Deep BSDE method approximates the unknown functions with neural nets that are solved for by a reinforcement learning algorithm. Let me summarize both steps in more detail.

1. *FBSDE system:* Define the state space of distributional projection coefficients, agent-specific and aggregate states by $X_t^i = ([\varphi_{n,t}^{v_n}]_{n\in\{1,\dots,N^x\},v_n\in\{0,\dots,V_n\}}, x_t^i, a_t)$.

[27] Although the method minimizes the error in the terminal boundary condition, Nüsken and Richter (2021) show that this error measure is equivalent to a measure of the cumulative value function error along the transition path. Thus, the Deep BSDE method goes beyond a pure shooting method which also explains its superior convergence properties.

Define the stochastic processes of the indirect utility $Y_t^i = J(t, X_t^i)$, its gradient $Z_t^i = DJ(t, X_t^i)$ and Hessian $\Gamma_t^i = D^2J(t, X_t^i)$.[28] Lastly, define the process $A_t^i = \mathcal{A}[DJ](t, X_t^i)$ by the infinitesimal generator mapping the gradient into a function of (t, X_t^i).[29] Cheridito et al. (2007) show that a PDE like the approximated master equation in (35) is equivalent to a system of FBSDEs containing the dynamics $[dX_t^i, dY_t^i, dZ_t^i]$ where the last two components constitute the backward equations which depend on the unknowns $[Y_0^i, Z_0^i, \{\Gamma_t^i\}_{t\geq 0}, \{A_t^i\}_{t\geq 0}]$. If the unknowns were given, the FBSDE system could be discretized with a forward scheme and easily simulated.

2. *Reinforcement learning algorithm:* The unknowns of the FBSDE system are given by functions

$$Y_0^i = \vartheta^{Y_0}(X_0^i), \quad Z_0^i = \vartheta^{Z_0}(X_0^i),$$
$$\Gamma_t^i = \vartheta^{\Gamma}(t, X_t^i), \quad A_t^i = \vartheta^{A}(t, X_t^i)$$

that are approximated by neural nets. Given a preliminary approximation, the FBSDE system is simulated forward to evaluate the loss function. The loss is defined by the mean-square error of the PDE's terminal condition in the simulated data. The neural net approximation of the unknowns is then updated using stochastic gradient descent. This step is repeated until the loss function is minimized.

5.2 General Equilibrium Extension: The GE-Deep BSDE Method

The standard Deep BSDE method has been designed to solve high-dimensional PDEs. However, our equilibrium problem presents additional complications. Thus, we extend the standard method to accommodate the cross-sectional interaction of our heterogeneous agent model, the state boundary condition that takes care of the state constraint, and the implicit market clearing equation defining the remaining returns. We introduce these adjustments following the two steps of the standard Deep BSDE method.

[28] Here, $DJ = [J_{X_1^i}, \ldots, J_{X_{N^X}^i}]^T$ with $N^X = \sum_{j=1}^{N^x} V_j + N^x + N^a$. The Hessian contains all second-order derivatives w.r.t. X^i.

[29] This process is defined precisely in Proposition 11.

5.2.1 The System of FBSDEs

Let us first derive the system of forward backward stochastic differential equations in our setting.[30]

Proposition 11 (System of FBSDEs)**.** *Define the aggregate state process by a vector* $X_t^a = ([\varphi_{n,t}^{v_n}]_{n\in\{1,\dots,N^x\},v_n\in\{0,\dots,V_n\}}, a_t)$ *and agent* i*'s state process by the* N^X*-dimensional vector* $X_t^i = ([\varphi_{n,t}^{v_n}]_{n\in\{1,\dots,N^x\},v_n\in\{0,\dots,V_n\}}, x_t^i, a_t)$*. Suppose that Assumptions 1 and 2 hold. Denote the equilibrium returns by* $r_t^* = [r_t^{*,ex}, r_t^{*,im}]$ *where* $r_t^{*,ex}$ *is given by Assumption 2 (i). Suppose that the implicit returns are given by a function* $r_t^{*,im} = r^{*,im}(t, X_t^a)$*. Then, a solution to the approximated master equation with terminal boundary condition*

$$\begin{aligned}0 =&\mathcal{H}^M[J, y^{*,i}, r^*](t, X^i) = \mathcal{H}^M[J, \mathcal{M}^y[DJ, D^2J, r^*], r^*](t, X^i)\\ 0 =&\mathcal{H}^{TBC}[J, y^{*,i}](X^i) = \mathcal{H}^{TBC}[J, \mathcal{M}^y[DJ, D^2J, r^*]](X^i)\end{aligned}$$

with $\mathcal{H}^M$ *as in (35) and* $\mathcal{H}^{TBC}$ *as in (13) is equivalent to a solution of the system of FBSDEs for the processes* $\{X_t^i, Y_t^i, Z_t^i\}_{t\geq 0}$ *defined as follows. The forward equation of the optimally controlled state dynamics is given by the initial data* X_0^i *and dynamics* $dX_t^i = ([d\varphi_{n,t}^{v_n}]_{n\in\{1,\dots,N^x\},v_n\in\{0,\dots,V_n\}}, dx_t^i, da_t)$ *with* (dx_t^i, da_t) *as in (7) and*

$$\begin{aligned}d\varphi_{n,t}^{v_n} =&\frac{1}{b^v}\left\{\int_{\mathbb{R}}\phi^v\left(\xi(x_n)\right)\overline{\mu_n}\left(t, y_t^{i,*}, x_n, a_t, r_t^*\right)m^n(x_n)dx_n\right.\\ &\left.+\frac{1}{2}\sum_{o\in\{id,ag\}}\int_{\mathbb{R}}\phi_\xi^v\left(\xi(x_n)\right)\frac{m^n(x_n)}{m^{\xi_n}(\xi(x_n))}\overline{\Sigma_{nn}^{o,*}}\left(t, x_n, x_n\right)m^n(x_n)dx_n\right\}dt\\ &+\frac{1}{b^v}\left\{\int_{\mathbb{R}}\phi^v\left(\xi(x_n)\right)\overline{\sigma_n^{ag}}\left(t, y_t^{i,*}, x_n, a_t, r_t^*\right)m^n(x_n)dx_n\right\}dB_t^{ag}\end{aligned}\tag{36}$$

with $y_t^{i,*} = \mathcal{M}^y[Y_t, \Gamma_t, r_t^*]$ *and* $\xi(x_n) = (M^{\xi_n})^{-1}(M^n(x_n))$ *and the covariance matrix* $\Sigma^{o,*}$ *as in (22). The backward equations are given by the terminal condition* $\lim_{T\to\infty}\mathbb{E}_0[e^{-\rho T}Y_T^i] = 0$ *and the dynamics*

$$\begin{aligned}dY_t^i =&-f(t, X_t^i, Y_t^i, Z_t^i, \Gamma_t^i)dt + (Z_t^i)^T dX_t^i\\ dZ_t^i =&A_t^i dt + \Gamma_t^i dX_t^i\end{aligned}\tag{37}$$

[30] Proposition 11 follows directly from Cheridito et al. (2007, Theorem 4.10) so that we omit the proof.

with

$$f(t, X_t^i, Y_t^i, Z_t^i, \Gamma_t^i) = u_t^i - \rho Y_t^i + \sum_{n=1}^{N^X} Z_{n,t}^i \frac{\mathbb{E}\left[dX_{n,t}^i\right]}{dt}$$

where $\mathbb{E}[.]/dt$ *delivers the drift of the process and* u_t^i *denotes the utility given the agent-specific state* X_t^i*. The equivalence between PDE and FBSDE yields*

$$\begin{aligned}
Y_t^i =& J(t, X_t^i) \\
Z_t^i =& DJ(t, X_t^i) \\
\Gamma_t^i =& D^2 J(t, X_t^i) \\
A_t^i =& \frac{\partial}{\partial t} DJ(t, X_t^i) + \frac{1}{2} \sum_{n,n'=1}^{N^X} \frac{\partial^2}{\partial X_{n'}^i \partial X_n^i} DJ(t, X_t^i) \frac{\mathbb{E}\left[dX_{n,t}^i dX_{n',t}^i\right]}{dt},
\end{aligned}$$

where $\mathbb{E}[.]/dt$ *collects the cross-variation of the two processes.*

Remark. *This result effectively translates the consistent law of motion of the cross-sectional density in (19) into the consistent law of motion of projection coefficients that discretize the cross-sectional distribution in (36).*

To simulate the system of FBSDEs $[dX_t^i, dY_t^i, dZ_t^i]$, we descretize it with a forward Euler-Murayama scheme following Andersson et al. (2023); Huang et al. (2025). However, we have to adjust this scheme slightly because our problem features an infinite instead of a finite horizon. We follow Friedl et al. (2023) and adopt the transformation of Traeger (2014) who maps the infinite time horizon $t \in [0, \infty)$ into a finite time interval through $\tau(t) = 1 - \exp(-\nu t) \in [0, 1)$ with $\nu > 0$. The Euler-Murayame scheme is then applied w.r.t. the finite time interval using the inverse transformation $t(\tau) = -\log(1-\tau)/\nu$. An equispaced grid $\tau_d = \frac{d}{D}$, $d \in \{0, 1, \ldots, D-1\}$ hence translates into a temporal discretization $t_d = t(\tau_d)$. In addition, we approximate the terminal time point with $t_D = \tau(1-\varepsilon)$ with $\varepsilon > 0$. Note that, given some initial data $\{X_0^i, Y_0^i, Z_0^i\}$, a set of simulated data $\{X_{t_d}^i, Y_{t_d}^i, Z_{t_d}^i\}_{d \in \{0,1,\ldots,D\}}$ represents the time series of states, utility and marginal utility for one agent. The details of the discretization scheme can be found in Section 2 of the online appendix.

5.2.2 The Reinforcement Learning Algorithm

The reinforcement learning algorithm approximates a set of unknown functions with neural nets. This approximation is found by minimizing a loss function using stochas-

tic gradient descent on simulated data from the FBSDE system. We adapt the standard procedure of the Deep BSDE method along multiple dimensions.

Unknown Functions. We extend the set of unknown functions to be approximated. In addition to the unknowns $[Y_0^i, Z_0^i, \{\Gamma_t^i\}_{t\geq 0}, \{A_t^i\}_{t\geq 0}]$, we also approximate the implicit equilibrium returns $\{r_t^{*,im}\}_{t\geq 0}$. Furthermore, note that the simulated data of $\{Z_t^i, r_t^{*,ex}\}_{t\geq 0}$ from Proposition 11 and the approximated unknowns $\{\Gamma_t, r_t^{*,im}\}_{t\geq 0}$ directly deliver $y^{*,i}(t, X_t^i) = \mathcal{M}^y[Z_t^i, \Gamma_t^i, r_t^*](t, X_t^i)$ due to Assumption 1, i.e., agent i's simulated trajectory of optimal controls. However, to simulate the projection coefficients of the consistent cross-sectional distribution in (36), we need to compute the expectation across all agents' optimal controls. For that sole purpose, we also approximate $\{Z_t^i\}_{t\geq 0}$ as an unknown function.[31] Overall, the functions that are approximated with neural nets consist of

$$\begin{aligned} Y_0^i &= \vartheta^{Y_0}(X_0^i), \quad Z_0^i = \vartheta^{Z_0}(X_0^i), \quad \Gamma_t^i = \vartheta^{\Gamma}(t, X_t^i), \\ A_t^i &= \vartheta^{A}(t, X_t^i), \quad Z_t^i = \vartheta^{Z}(t, X_t^i), \quad r_t^{*,im} = \vartheta^{r}(t, X_t^a). \end{aligned}$$

Simulation. Given an empirically calibrated initial state X_0^i and the approximated initial values $[\vartheta^{Y_0}(X_0^i), \vartheta^{Z_0}(X_0^i)]$, the discretized version of the FBSDE system $[dX_t^i, dY_t^i, dZ_t^i]$ from Proposition 11 is simulated forward. The approximated functions ϑ^{Γ} and ϑ^{A} define the dynamics dZ_t^i in (37) whereas the functions ϑ^{Γ} and ϑ^{r} feature in the law of motions of $[dX_t^i, dY_t^i]$ due to the optimal controls $y_t^{*,i} = \mathcal{M}^y[Z_t^i, \vartheta^{\Gamma}(t, X_t^i), r_t^{*,ex}, \vartheta^{r}(t, X_t^a)]$ and the direct dependence of dX_t^i on $r_t^* = [r_t^{*,ex}, \vartheta^{r}(t, X_t^a)]$. The approximated function ϑ^{Z} features in the dynamics of the projection coefficients in (36) as we integrate over the optimal control as a function of x, i.e., $y_t^{*,i} = \mathcal{M}^y[\vartheta^{Z}(t, x), \vartheta^{\Gamma}(t, x), r_t^{*,ex}, \vartheta^{r}(t, X_t^a)]$.

We use the simulation of the FBSDE system to enforce the state boundary condition. The simulated data serves as the training data for the approximated solution and hence, all equilibrium conditions should hold at all times. If the simulated data of the constrained state $e_t^i \geq 0$ violate the constraint at some time point $t_{d'}$, we modify the simulated data around that time point to enforce the constraint and then continue the simulation with the modified data. This includes not only setting $e_{t_{d'}}^i \geq 0$ but also modifying the optimal controls $[y_{t_{d'-1}}^{*,i}, y_{t_{d'}}^{*,i}]$ and the gradient $[Z_{t_{d'-1}}^i, Z_{t_{d'}}^i]$ such

[31] In fact, we do not need to approximate the full gradient $DJ = [[J_{\varphi_n^{v_n}}]_{n\in\{1,\dots,N^x\}, v_n\in\{0,\dots,V_n\}}, J_{x^i}, J_a]^T$ with an additional unknown function but only J_{x^i} because the optimal controls solely depend on the derivatives w.r.t. the agent-specific states.

that the drift and volatility of $de^i_{t_{d'}}$ are nonnegative and vanish, respectively.

Example. *In our example model, the state boundary is enforced as follows. Choose a small* $\epsilon > 0$. *Set* $k^i_{t_{d'}} = \epsilon$ *to let the volatility of the constrained state vanish.*[32] *Set* $\iota^{i,*}_{t_{d'-1}}$ *to exactly the value that results in* $k^i_{t_{d'}} = \epsilon$ *in the discretized FBSDE. Set the simulated marginal utility* J_k *to a large value, i.e.,* $Z^i_{n_k,t_{d'-1}} = 1/\epsilon$. *Recompute* $[X^i_{t_{d'}}, Y^i_{t_{d'}}, Z^i_{t_{d'}}]$ *with these new values at* $t_{d'-1}$. *If the constrained state is still in violation* $e^i_{t_{d'}} < 0$, *set* $e^i_{t_{d'}} = \epsilon$. *Furthermore, set* $c^{i,*}_{t_{d'}} = \min\{c^{i,*}_{t_{d'}}, w_{t_{d'}} e^{z^i_{t_{d'}}} + \bar{b} r^b_{t_{d'}}\}$ *and* $Z^i_{n_e,t_{d'}} = \max\{Z^i_{n_e,t_{d'}}, (w_{t_{d'}} e^{z^i_{t_{d'}}} + \bar{b} r^b_{t_{d'}})^{-\gamma}\}$ *which corresponds to* J_{e^i}.

Similarly to enforcing the state boundary in the simulated data, we also enforce the market clearing condition. The zeroth-order projection coefficient φ^0_n is the cross-sectional mean of the state variable x^i_n. When market clearing is implicit, the simulated zeroth-order coefficients are corrected to satisfy the market clearing equation.

Example. *In our example, we approximate the states* $[\exp(z^i), k^i, e^i]$ *with polynomial chaos. Thus, the corresponding aggregates are given by the zeroth-order projection coefficients* $[\varphi^0_{e^z}, \varphi^0_k, \varphi^0_e]$. *These are simulated within the FBSDE system. The simulated return on capital and the wage are then calculated by inserting these projection coefficients into the corresponding marginal productivity. We enforce the implicit market clearing for the bond by correcting the simulated value of* φ^0_e *accordingly, i.e.,* $\varphi^0_e = -\bar{b} + \phi_t \varphi^0_k$.

Lastly, we modify the sampling in order to enable the approximation of the terminal boundary condition. Denote the simulated terminal value by $Y^i_\infty = \mathbb{E}_0[e^{-\rho t_D} Y^i_{t_D}]$. We sample the various agent i's initial states X^i_0 from the initial cross-sectional distribution using quasi-Monte Carlo sampling. To compute the expectation in Y^i_∞, we duplicate all agent i's initial states X^i_0 multiple times and compute the corresponding $Y^i_{t_D}$ for different trajectories of the exogenous shocks. The simulated terminal value Y^i_∞ is then calculated as the corresponding Monte Carlo average. We simulate $Z^i_\infty = \mathbb{E}_0[e^{-\rho t_D} Z^i_{t_D}]$ equivalently.

Loss Function and Gradient Descent. Originally, the Deep BSDE method defines the loss by the mean-square error of the PDE's terminal condition in the

[32] We do not set $k^i_{t_{d'}} = 0$ because agents avoid the fire-sale boundary at all cost in equilibrium. In our model, zero capital is an absorbing state, i.e., capital would stay at zero forever. As capital is log-normal, this is a zero probability event.

simulated data which reads $\frac{1}{I}\sum_{i=1}^{I}(Y_\infty^i - 0)^2$. It has been shown that this loss function may lead to convergence issues by Nüsken and Richter (2021) when the FBSDE system is high-dimensional and by Huang et al. (2025) when the FBSDE system is strongly coupled. Both properties apply to our setting. Thus, we follow Nüsken and Richter (2021); Huang et al. (2025) and use the variance of the terminal condition for both Y^i and its gradient Z^i to improve convergence. This yields a loss function with multiple components: the empirical variances $\frac{1}{I-1}\sum_{i=1}^{I}(Y_\infty^i - \overline{Y_\infty})^2$ and $\frac{1}{I-1}\sum_{i=1}^{I}(Z_\infty^i - \overline{Z_\infty})^2$ where $\overline{Y_\infty}$ and $\overline{Z_\infty}$ denote the empirical averages across i. The loss resulting from the terminal condition disciplines the approximation of the traditional unknown functions $[\vartheta^{Y_0}, \vartheta^{Z_0}, \vartheta^{\Gamma}, \vartheta^{A}]$. As we added two more unknowns to be approximated by neural nets, we need to extend the loss function to discipline the two additional functions $[\vartheta^{Z}, \vartheta^{r}]$. We add the mean-square error of ϑ^Z compared to the simulated Z_t^i as a third component and the mean-square error of the implicit market clearing equation $0 = \mathcal{M}^r[y_t^{i,*}]$ as a fourth component. The total loss function consists of

$$\mathcal{L}(\vartheta) = \begin{bmatrix} \frac{1}{I-1}\sum_{i=1}^{I}\left(Y_\infty^i - \overline{Y_\infty}\right)^2 \\ \frac{1}{I-1}\sum_{i=1}^{I}\left(Z_\infty^i - \overline{Z_\infty}\right)^2 \\ \frac{1}{I}\sum_{i=1}^{I}\sum_{d=1}^{D}\left(\vartheta^Z(t_d, X_{t_d}^i) - Z_{t_d}^i\right)^2 \\ \frac{1}{I}\sum_{i=1}^{I}\sum_{d=0}^{D}\left(\mathcal{M}^r\left[y_{t_d}^{i,*}\right]\right)^2 \end{bmatrix}.$$

To summarize, the first two components evaluate individual optimality, the third component evaluates the cross-sectional distribution's consistency with individual optimality, and, the last component evaluates market clearing. Minimizing the loss function, therefore, balances multiple objectives. Applying stochastic gradient descent to a simple weighted average of these multiple objectives generally does not converge, because the gradients of the different objectives may point in conflicting descent directions. To avoid these issues, we apply the correlation-reduced stochastic multi-objective gradient manipulation (CR-MOGM) by Zhou et al. (2022) in addition to projecting conflicting gradients (PCGrad) according to Yu et al. (2020). This approach adjusts the weights of the different loss components when computing the stochastic gradient descent step to find a Pareto optimum that balances all objectives. Chen et al. (2025) provide a comprehensive survey on multi-objective deep learning.

Apart from the modifications described, we closely follow Beck et al. (2019) to implement the neural net approximation and the mini-batch stochastic gradient descent to minimize the loss function. The details of this implementation can be found

in Section 3 of the online appendix.

6 Economic Insights

This is work in progress. We are currently finalizing the implementation and expect results in late Spring 2026.

Appendix A Proofs

A.1 Proof of Proposition 4

Proof. Note that the Kolmogorov forward equation in (19) states the distributional dynamics in the form of $dm_t(x) = \sum_{n=1}^{N^x} \frac{\partial}{\partial x_n} dD_t^n(x)$ with

$$\begin{aligned} dD_t^n(x) = \Big\{ & -\mu_n\left(t, y^{i,*}[r_t](t,x,a_t), x, a_t, r_t\right) m_t(x) \\ & + \frac{1}{2} \sum_{n'=1}^{N^x} \sum_{o \in \{id, ag\}} \frac{\partial}{\partial x_{n'}} \left[\Sigma^o_{nn'}\left(t, y^{i,*}[r_t](t,x,a_t), x, a_t, r_t\right) m_t(x) \right] \Big\} dt \\ & - \sigma_n^{ag}\left(t, y^{i,*}[r_t](t,x,a_t), x, a_t, r_t\right) m_t(x) dB_t^{ag}. \end{aligned} \tag{38}$$

Therefore, we can apply integration by parts in (20) once in the first term and twice in the second term so that the dynamics of $f(M_t)$ can be stated in terms of the intrinsic distributional derivatives

$$\begin{aligned} df(M_t) = & -\int_{\mathbb{R}^{N^x}} \sum_{n=1}^{N^x} f_{x_n M}[M_t](x) dD_t^n(x) dx \\ & + \frac{1}{2} \int_{\mathbb{R}^{2N^x}} \sum_{n,n'=1}^{N^x} f_{x_n x'_{n'} MM}[M_t](x,x') dD_t^n(x) dD_t^{n'}(x') dx' dx. \end{aligned} \tag{39}$$

Inserting the definition of $dD_t^n(x)$ yields the appropriate expansion in (21). Note that the term containing $f_{x_{n'} x_n M}[M_t](x) = \frac{\partial}{\partial x_{n'}} f_{x_n M}[M_t](x)$ results from applying integration by parts once more to the second drift term of $dD_t^n(x)$ in (38). □

A.2 Proof of Theorem 9

Proof. Recall from (20), that the distributional terms in the master equation result from the expansion

$$\begin{aligned} df(M_t) = &\int_{\mathbb{R}^{N^x}} f_M[M_t](x) dm_t(x) dx \\ &+ \frac{1}{2} \int_{\mathbb{R}^{2N^x}} f_{MM}[M_t](x, x') dm_t(x') dm_t(x) dx' dx. \end{aligned}$$

with $dm_t(x)$ as in the Kolmogorov forward equation in (19). Given (30), the joint distribution's dynamics are given by

$$dm_t(x) = dc_t(u) \cdot \prod_{j=1}^{N^x} m_t^j(x_j) + \sum_{n=1}^{N^x} dm_t^n(x_n) \cdot c_t(u) \cdot \prod_{j \neq n, j=1}^{N^x} m_t^j(x_j). \tag{40}$$

The last N^x summands can be determined by computing the Kolmogorov forward equation of the marginal distributions $dm_t^n(x_n) = \int dm_t(x) dx_{-n}$ resulting in

$$\begin{aligned} & dm_t^n(x_n) \cdot c_t(u) \cdot \prod_{j \neq n, j=1}^{N^x} m_t^j(x_j) \qquad (41) \\ = & \Bigg\{ -\frac{\partial}{\partial x_n} \left(\mathbf{E}_t \left[\mu_n \left(t, y^{i,*}[r_t](t, x, a_t), x, a_t, r_t \right) \middle| x_n \right] m_t(x) \right) \\ & + \frac{1}{2} \sum_{o \in \{id, ag\}} \frac{\partial^2}{\partial x_n^2} \left(\mathbf{E}_t \left[\Sigma_{nn}^{o,*} \left(t, x, x \right) \middle| x_n \right] m_t(x) \right) \Bigg\} dt \\ & - \frac{\partial}{\partial x_n} \left(\mathbf{E}_t \left[\sigma_n^{ag} \left(t, y^{i,*}[r_t](t, x, a_t), x, a_t, r_t \right) \middle| x_n \right] m_t(x) \right) dB_t^{ag} \\ = & \frac{\partial}{\partial x_n} \Bigg[\Bigg\{ -\overline{\mu_n} \left(t, y_t^{i,*}, x_n, a_t, r_t \right) m_t(x) + \frac{1}{2} \sum_{o \in \{id, ag\}} \frac{\partial}{\partial x_n} \left[\overline{\Sigma_{nn}^{o,*}} \left(t, x_n, x_n \right) m_t(x) \right] \Bigg\} dt \\ & - \overline{\sigma_n^{ag}} \left(t, y_t^{i,*}, x_n, a_t, r_t \right) m_t(x) dB_t^{ag} \Bigg] \\ =: & \frac{\partial}{\partial x_n} dD_t^{\text{marg},n}(x_n) \end{aligned}$$

with $\Sigma^{o,*}$ as in (22). In light of (40) and (19), we obtain the copula dynamics

$$
\begin{aligned}
& dc_t(u) \cdot \prod_{j=1}^{N^x} m_t^j(x_j) \qquad (42)\\
=& \sum_{n=1}^{N^x} \frac{\partial}{\partial x_n} \Big[\Big\{ - \left(\mu_n \left(t, y^{i,*}[r_t](t,x,a_t), x, a_t, r_t \right) - \overline{\mu_n} \left(t, y_t^{i,*}, x_n, a_t, r_t \right) \right) m_t(x) \\
& + \frac{1}{2} \sum_{o \in \{id, ag\}} \frac{\partial}{\partial x_n} \left[\left(\Sigma_{nn}^{o,*}(t,x,x) - \overline{\Sigma_{nn}^{o,*}}(t,x_n,x_n) \right) m_t(x) \right] \\
& + \sum_{n'=1}^{n-1} \sum_{o \in \{id, ag\}} \frac{\partial}{\partial x_{n'}} \left[\Sigma_{nn'}^{o,*}(t,x,x)\, m_t(x) \right] \Big\} dt \\
& - \left(\sigma_n^{ag} \left(t, y^{i,*}[r_t](t,x,a_t), x, a_t, r_t \right) - \overline{\sigma_n^{ag}} \left(t, y_t^{i,*}, x_n, a_t, r_t \right) \right) m_t(x) dB_t^{ag} \Big] \\
=:& \sum_{n=1}^{N^x} \frac{\partial}{\partial x_n} dD_t^{\mathrm{cop},n}(x).
\end{aligned}
$$

Overall, we can rewrite the Kolmogorov forward equation using (40)-(42) such that

$$
dm_t(x) = \sum_{n=1}^{N^x} \frac{\partial}{\partial x_n} \left[dD_t^{\mathrm{cop},n}(x) + dD_t^{\mathrm{marg},n}(x_n) \right].
$$

To derive the master equation in terms of partial distributional derivatives, insert Definition 7 into (20) to obtain

$$
\begin{aligned}
df(M_t) = \int_{\mathbb{R}^{N^x}} \Bigg\{ & f_C[M](x) \cdot dc_t(u) \cdot \prod_{j=1}^{N^x} m_t^j(x_j) \\
& + \sum_{n=1}^{N^x} f_{M^n}[M](x_n) \cdot dm_t^n(x_n) \cdot c_t(u) \cdot \prod_{j\neq n, j=1}^{N^x} m_t^j(x_j) \Bigg\} dx \\
& + \frac{1}{2} \int_{\mathbb{R}^{2N^x}} \Bigg\{ f_{CC}[M](x,x') \cdot dc_t(u_{x'}) \cdot \prod_{j=1}^{N^x} m_t^j(x'_j) \cdot dc_t(u_x) \cdot \prod_{k=1}^{N^x} m_t^k(x_k) \\
& + 2 \sum_{n'=1}^{N^x} f_{M^{n'}C}[M](x, x'_{n'}) \cdot dm_t^{n'}(x'_{n'}) \cdot c_t(u_{x'}) \cdot \prod_{j\neq n', j=1}^{N^x} m_t^j(x'_j) \cdot dc_t(u_x) \cdot \prod_{k=1}^{N^x} m_t^k(x_k) \\
& + \sum_{n,n'=1}^{N^x} f_{M^{n'}M^n}[M](x_n, x'_{n'}) \cdot dm_t^{n'}(x'_{n'}) \cdot c_t(u_{x'}) \cdot \prod_{j\neq n', j=1}^{N^x} m_t^j(x'_j) \\
& \cdot dm_t^n(x_n) \cdot c_t(u_x) \cdot \prod_{k\neq n, k=1}^{N^x} m_t^k(x_k) \Bigg\} dx' dx. \\
= \int_{\mathbb{R}^{N^x}} \Bigg\{ & f_C[M](x) \cdot \sum_{n=1}^{N^x} \frac{\partial}{\partial x_n} dD_t^{\mathrm{cop},n}(x) + \sum_{n=1}^{N^x} f_{M^n}[M](x_n) \cdot \frac{\partial}{\partial x_n} dD_t^{\mathrm{marg},n}(x_n) \Bigg\} dx \\
& + \frac{1}{2} \int_{\mathbb{R}^{2N^x}} \Bigg\{ f_{CC}[M](x,x') \cdot \sum_{n'=1}^{N^x} \frac{\partial}{\partial x'_{n'}} dD_t^{\mathrm{cop},n'}(x') \cdot \sum_{n=1}^{N^x} \frac{\partial}{\partial x_n} dD_t^{\mathrm{cop},n}(x) \\
& + 2 \sum_{n'=1}^{N^x} f_{M^{n'}C}[M](x, x'_{n'}) \cdot \frac{\partial}{\partial x'_{n'}} dD_t^{\mathrm{marg},n'}(x'_{n'}) \cdot \sum_{n=1}^{N^x} \frac{\partial}{\partial x_n} dD_t^{\mathrm{cop},n}(x) \\
& + \sum_{n,n'=1}^{N^x} f_{M^{n'}M^n}[M](x_n, x'_{n'}) \cdot \frac{\partial}{\partial x'_{n'}} dD_t^{\mathrm{marg},n'}(x'_{n'}) \cdot \frac{\partial}{\partial x_n} dD_t^{\mathrm{marg},n}(x_n) \Bigg\} dx' dx,
\end{aligned}
$$

where the last equality follows from inserting (41)-(42). Integration by parts delivers the following expansion

$$
\begin{aligned}
df(M_t) = &- \sum_{n=1}^{N^x} \int_{\mathbb{R}^{N^x}} \left\{ f_{x_n C}[M](x) dD_t^{\text{cop},n}(x) + f_{x_n M^n}[M](x_n) dD_t^{\text{marg},n}(x_n) \right\} dx \\
&+ \sum_{n,n'=1}^{N^x} \int_{\mathbb{R}^{2N^x}} \left\{ \frac{1}{2} f_{x_n x'_{n'} CC}[M](x,x') dD_t^{\text{cop},n'}(x') dD_t^{\text{cop},n}(x) \right. \\
&+ f_{x_n x'_{n'} M^{n'} C}[M](x, x'_{n'}) dD_t^{\text{marg},n'}(x'_{n'}) dD_t^{\text{cop},n}(x) \\
&\left. + \frac{1}{2} f_{x_n x'_{n'} M^{n'} M^n}[M](x_n, x'_{n'}) dD_t^{\text{marg},n'}(x'_{n'}) dD_t^{\text{marg},n}(x_n) \right\} dx' dx.
\end{aligned}
$$

Note that the cross-variations define the following covariance matrices

$$
\begin{aligned}
\Sigma^{\text{C},*}_{nn'}(t,x,x') :=& \frac{dD_t^{\text{cop},n'}(x') dD_t^{\text{cop},n}(x)}{dt} \frac{1}{m_t(x') m_t(x)} \\
=& \left(\sigma^{ag}_{n'} \left(t, y^{i,*}[r_t](t,x',a_t), x', a_t, r_t \right) - \overline{\sigma^{ag}_{n'}} \left(t, y_t^{i,*}, x'_{n'}, a_t, r_t \right) \right) \\
& \cdot \left(\sigma^{ag}_{n} \left(t, y^{i,*}[r_t](t,x,a_t), x, a_t, r_t \right) - \overline{\sigma^{ag}_{n}} \left(t, y_t^{i,*}, x_n, a_t, r_t \right) \right) \\
\Sigma^{\text{MC},*}_{nn'}(t,x,x'_{n'}) :=& \frac{dD_t^{\text{marg},n'}(x'_{n'}) dD_t^{\text{cop},n}(x)}{dt} \frac{1}{m_t(x') m_t(x)} \\
=& \overline{\sigma^{ag}_{n'}} \left(t, x'_{n'}, a_t, r_t \right) \left(\sigma^{ag}_{n} \left(t, y^{i,*}[r_t](t,x,a_t), x, a_t, r_t \right) - \overline{\sigma^{ag}_{n}} \left(t, y_t^{i,*}, x_n, a_t, r_t \right) \right) \\
\Sigma^{\text{M},*}_{nn'}(t,x_n,x'_{n'}) :=& \frac{dD_t^{\text{marg},n'}(x'_{n'}) dD_t^{\text{marg},n}(x_n)}{dt} \frac{1}{m_t(x') m_t(x)} \\
=& \overline{\sigma^{ag}_{n'}} \left(t, y_t^{i,*}, x'_{n'}, a_t, r_t \right) \overline{\sigma^{ag}_{n}} \left(t, y_t^{i,*}, x_n, a_t, r_t \right).
\end{aligned}
\tag{43}
$$

Inserting the definitions of $D_t^{\text{cop},n}$ and $D_t^{\text{marg},n}$ and using integration by parts once more, this allows us to rewrite the expansion as

$$
\begin{aligned}
&df(M_t) = \\
&\sum_{n=1}^{N^x} \int_{\mathbb{R}^{N^x}} \Big\{ -f_{x_n C}[M](x) \left(\mu_n \left(t, y^{i,*}[r_t](t,x,a_t), x, a_t, r_t \right) - \overline{\mu_n} \left(t, y_t^{i,*}, x_n, a_t, r_t \right) \right) \\
&+ \frac{1}{2} \sum_{o \in \{id, ag\}} f_{x_n x_n C}[M](x) \left(\Sigma_{nn}^{o,*}(t,x,x) - \overline{\Sigma_{nn}^{o,*}}(t,x_n,x_n) \right) \\
&+ \sum_{n'=1}^{n-1} \sum_{o \in \{id, ag\}} f_{x_{n'} x_n C}[M](x) \Sigma_{nn'}^{o,*}(t,x,x) + f_{x_n M^n}[M](x_n) \overline{\mu_n} \left(t, y_t^{i,*}, x_n, a_t, r_t \right) \\
&+ \frac{1}{2} \sum_{o \in \{id, ag\}} f_{x_n x_n M^n}[M](x_n) \overline{\Sigma_{nn}^{o,*}}(t,x_n,x_n) \Big\} m_t(x) dx\, dt \\
&+ \sum_{n,n'=1}^{N^x} \int_{\mathbb{R}^{2N^x}} \Big\{ \frac{1}{2} f_{x_n x'_{n'} CC}[M](x,x') \Sigma_{nn'}^{\text{C},*}(t,x,x') \\
&+ f_{x_n x'_{n'} M^{n'} C}[M](x, x'_{n'}) \Sigma_{nn'}^{\text{MC},*}(t,x,x'_{n'}) \\
&+ \frac{1}{2} f_{x_n x'_{n'} M^{n'} M^n}[M](x_n, x'_{n'}) \Sigma_{nn'}^{\text{M},*}(t,x_n,x'_{n'}) \Big\} m_t(x') m_t(x) dx' dx\, dt \\
&+ \sum_{n=1}^{N^x} \Big\{ \int_{\mathbb{R}^{N^x}} f_{x_n C}[M](x) \left(\sigma_n^{ag} \left(t, y^{i,*}[r_t](t,x,a_t), x, a_t, r_t \right) - \overline{\sigma_n^{ag}} \left(t, y_t^{i,*}, x_n, a_t, r_t \right) \right) m_t(x) dx \\
&+ \int_{\mathbb{R}^{N^x}} f_{x_n M^n}[M](x_n) \overline{\sigma_n^{ag}} \left(t, y_t^{i,*}, x_n, a_t, r_t \right) m_t(x) dx \Big\} dB_t^{ag}
\end{aligned}
$$

As in (23), the master equation combines the standard Itô terms and the distributional terms of the expansion in the previous equation as well as the cross-terms of the former two. Define the additional covariance matrices

$$
\begin{aligned}
\Sigma_{nn'}^{x^i \text{C},*}(t,x^i,x) =& \sigma_n^{ag} \left(t, y^{i,*}[r_t](t,x^i,a_t), x^i, a_t, r_t \right) \\
& \cdot \left(\sigma_{n'}^{ag} \left(t, y^{i,*}[r_t](t,x,a_t), x, a_t, r_t \right) - \overline{\sigma_{n'}^{ag}} \left(t, y_t^{i,*}, x_{n'}, a_t, r_t \right) \right) \\
\Sigma_{nn'}^{x^i \text{M},*}(t,x^i,x_{n'}) =& \sigma_n^{ag} \left(t, y^{i,*}[r_t](t,x^i,a_t), x^i, a_t, r_t \right) \overline{\sigma_{n'}^{ag}} \left(t, y_t^{i,*}, x_{n'}, a_t, r_t \right) \\
\Sigma_{nn'}^{aC,*}(t,x) =& \zeta(t,a_t,r_t) \left(\sigma_{n'}^{ag} \left(t, y^{i,*}[r_t](t,x,a_t), x, a_t, r_t \right) - \overline{\sigma_{n'}^{ag}} \left(t, y_t^{i,*}, x_{n'}, a_t, r_t \right) \right) \\
\Sigma_{nn'}^{aM,*}(t,x_{n'}) =& \zeta(t,a_t,r_t) \overline{\sigma_{n'}^{ag}} \left(t, y_t^{i,*}, x_{n'}, a_t, r_t \right).
\end{aligned} \tag{44}
$$

Then, the cross term in (23) is given by

$$
\begin{aligned}
&d\mathcal{A}_{M,x^i,a}\left(t, M_t, x_t^i, a_t\right) = \\
&\frac{1}{2}\Bigg\{ \sum_{n,n'=1}^{N^x} \int_{\mathbb{R}^{N^x}} J_{x_n^i x_{n'} C}[M_t]\left(x; t, x_t^i, a_t\right) \Sigma_{nn'}^{x^i \mathrm{C},*}(t, x^i, x) m_t(x) dx \\
&+ \sum_{n,n'=1}^{N^x} \int_{\mathbb{R}^{N^x}} J_{x_n^i x_{n'} M^{n'}}[M_t]\left(x_{n'}; t, x_t^i, a_t\right) \Sigma_{nn'}^{x^i \mathrm{M},*}(t, x^i, x_{n'}) m_t(x) dx \\
&+ \sum_{n=1}^{N^a} \sum_{n'=1}^{N^x} \int_{\mathbb{R}^{N^x}} J_{a_n x_{n'} C}[M_t]\left(x; t, x_t^i, a_t\right) \Sigma_{nn'}^{aC,*}(t, x)\, m_t(x) dx \\
&+ \sum_{n=1}^{N^a} \sum_{n'=1}^{N^x} \int_{\mathbb{R}^{N^x}} J_{a_n x_{n'} M^{n'}}[M_t]\left(x_{n'}; t, x_t^i, a_t\right) \Sigma_{nn'}^{aM,*}(t, x_{n'})\, m_t(x) dx \Bigg\} dt.
\end{aligned}
$$

Combining all the terms of dJ according to (23) and collecting its drift delivers the master equation in terms of partial distributional derivatives. Setting all partial derivatives w.r.t. the copula $J_{\ldots C}[M_t] = 0$ inline with Assumption 8 delivers the simplifies master equation and concludes the proof. □

A.3 Proof of Lemma 10

Before we proceed to the proof of this lemma, let us introduce two more concepts from the mean field game literature that are relevant here. As an Itô process, the agent-specific state in (7) is a square-integrable stochastic process implying that, for a fixed time point t, x_t^i is a square-integrable random variable, i.e., $x_t^i \in L^2(\Omega, \mathcal{F}_t, \mathbb{P}; \mathbb{R}^{N^x})$. The space of square-integrable random variables is a Hilbert space such that the derivative of a function of the random variable $f(x^i)$ is given by the standard Fréchet derivative.

Definition 12 (Distributional Derivatives II (see Carmona and Delarue, 2018a, Definition 5.22)). *The Fréchet derivative of a function $f : L^2(\Omega, \mathcal{F}_t, \mathbb{P}; \mathbb{R}^{N^x}) \to \mathbb{R}$ is given by the the total derivative $f_{x^i}[x^i](x) = \sum_{n=1}^{N^x} f_{x_n^i}[x^i](x_n)$ with the partial derivatives*

given by[33]

$$\lim_{s\to 0}\frac{f(x^i+s(x^{i'}-x^i))-f(x^i)}{s}=\mathbb{E}^{x,x'}\left[\sum_{n=1}^{N^x}f_{x_n^i}[x^i](x_n)(x'_n-x_n)\right],$$

where the expectation is taken over $x,x'\in L^2(\Omega,\mathcal{F}_t,\mathbb{P};\mathbb{R}^{N^x})$ *where* x *has the same distribution as* x^i *and* x' *has the same distribution as* $x^{i'}$*. The second-order derivative is similarly defined by*

$$\begin{aligned}\lim_{s\to 0}\frac{f_{x^i}[x^i+s(x^{i'}-x^i)](x)-f_{x^i}[x^i](x)}{s}=&\mathbb{E}^{\tilde{x},\tilde{x}'}\left[\sum_{n=1}^{N^x}f_{x_n^ix^i}[x^i](x,\tilde{x}_n)(\tilde{x}'_n-\tilde{x}_n)\right]\\=&\mathbb{E}^{\tilde{x},\tilde{x}'}\left[\sum_{n,n'=1}^{N^x}f_{x_n^ix_{n'}^i}[x^i](x_{n'},\tilde{x}_n)(\tilde{x}'_n-\tilde{x}_n)\right],\end{aligned}$$

where the expectation is taken over $\tilde{x},\tilde{x}'\in L^2(\Omega,\mathcal{F}_t,\mathbb{P};\mathbb{R}^{N^x})$ *where* $\tilde{x}$ *has the same distribution as* x^i *and* $\tilde{x}'$ *has the same distribution as* $x^{i'}$*.*

There is a direct link between the intrinsic partial distributional derivatives w.r.t. the distribution M and derivatives w.r.t. the corresponding random variable x^i.

Proposition 13 (Cardaliaguet et al. (2019, Proposition A.1.2))**.** *With the distribution* $M(x)=\mathbb{P}(x^i\leq x)$*, we can rewrite a function of the distribution in terms of the random variable, i.e.,* $f(M)=\tilde{f}(x^i)$*. If* f *is differentiable w.r.t. the distribution, then it is differentiable w.r.t. the corresponding random variable and the intrinsic distributional derivatives equal* $f_{x_nM}[M](x)=\tilde{f}_{x_n^i}[x^i](x)$*. Similarly, the second-order derivatives agree* $f_{x_nx'_{n'}MM}[M](x,x')=\tilde{f}_{x_n^ix_{n'}^i}[x^i](x,x')$*.*

Corollary 14. *Given Definition 7 of the partial distributional derivatives, the previous result can be restated by*

$$\begin{aligned}\tilde{f}_{x_n^i}[x^i](x)=&f_{x_nM}[M](x)\\=&f_{x_nC}[M](x)+\sum_{n'=1}^{N^x}f_{x_nM^{n'}}[M](x_{n'})\\=&f_{x_nC}[M](x)+f_{x_nM^n}[M](x_n)\end{aligned}$$

[33]The equation to define f_{x^i} is equivalent to Cardaliaguet et al. (2019, Equation (A.5)).

because $f_{x_n M^{n'}}[M](x_{n'}) = 0$ *when* $n' \neq n$. *For our indirect utility function, the equality further simplifies to* $\tilde{J}_{x_n^i}[x^i](x) = J_{x_n M^n}[M](x_n)$ *because Assumption 8 holds. Similarly, the second-order derivatives read*

$$\begin{aligned}\tilde{f}_{x_n^i x_{n'}^i}[x^i](x,x') =& f_{x_n x'_{n'} MM}[M](x,x') \\ =& f_{x_n x'_{n'} CC}[M](x,x') + f_{x_n x'_{n'} M^n C}[M](x_n,x') \\ &+ f_{x_n x'_{n'} C M^{n'}}[M](x,x'_{n'}) + f_{x_n x'_{n'} M^n M^{n'}}[M](x_n,x'_{n'}).\end{aligned}$$

Given Assumption 8, this implies that $\tilde{J}_{x_n^i x_{n'}^i}[x^i](x,x') = J_{x_n x'_{n'} M^{n'} M^n}[M](x_n,x'_{n'})$.

Proof of Lemma 10. Consider the polynomial chaos representation (32) of the n-th element of the random variable $x^i \in L^2(\Omega,\mathcal{F}_t,\mathbb{P};\mathbb{R}^{N^x})$. The corresponding projection coefficient given in (34) defines a function $\varphi^v : L^2(\Omega,\mathcal{F}_t,\mathbb{P};\mathbb{R}) \to \mathbb{R}$ with

$$\varphi^v(x_n^i) = \frac{1}{b^v}\int_{\mathbb{R}} x_n \phi^v(\xi(x_n)) d\mathbb{P}(x_t^i \leq x_n) = \frac{1}{b^v}\mathbb{E}\left[x_n^i \phi^v(\xi_n)\right]$$

where $\xi(x_n) = (M^{\xi_n})^{-1}(M^n(x_n))$ and b^v as in (33). Its Fréchet derivative according to Definition 12 reads $\varphi^v_{x_n^i}[x^i](x_n) = \phi^v(\xi(x_n))/b^v$. Its second-order derivative is $\varphi^v_{x_n^i x_n^i}[x^i](x_n,x'_n) = 0$. Let us now consider a function of projection coefficients $\tilde{f}(\{\varphi_1^v\}_{v=0}^\infty,\ldots,\{\varphi_{N^x}^v\}_{v=0}^\infty)$. Any $\varphi_n^v = \varphi^v(x_n^i)$, $v \in \{0,1,\ldots\}$, such that the distributional derivative in Definition 12 of the function $\tilde{f}$ is

$$\tilde{f}_{x_n^i}[x^i](x_n) = \sum_{v=0}^\infty \tilde{f}_{\varphi_n^v} \cdot \varphi^v_{x_n^i}[x^i](x_n) = \sum_{v=0}^\infty \frac{\tilde{f}_{\varphi_n^v}}{b^v}\phi^v(\xi(x_n)).$$

The second-order derivative follows accordingly

$$\begin{aligned}\tilde{f}_{x_n^i x_{n'}^i}[x^i](x_n,x_{n'}) &= \sum_{v,v'=0}^\infty \tilde{f}_{\varphi_n^v \varphi_{n'}^{v'}} \cdot \varphi^v_{x_n^i}[x^i](x_n) \cdot \varphi^{v'}_{x_{n'}^i}[x^i](x_{n'}) \\ &= \sum_{v,v'=0}^\infty \frac{\tilde{f}_{\varphi_n^v \varphi_{n'}^{v'}}}{b^v b^{v'}}\phi^v(\xi(x_n))\phi^{v'}(\xi(x_{n'})).\end{aligned}$$

The equivalence of Proposition 13 and Corollary 14 delivers the first- and second-order

derivatives of the lemma. Lastly,

$$\begin{aligned} f_{x_n x_n M^n}[M](x_n) =& \frac{\partial}{\partial x_n} f_{x_n M^n}[M](x_n) = \frac{\partial}{\partial x_n} \tilde{f}_{x_n^i}[x^i](x_n) \\ =& \sum_{v=0}^{\infty} \frac{\tilde{f}_{\varphi_n^v}}{b^v} \phi_\xi^{\prime v}(\xi(x_n)) \frac{\partial}{\partial x_n} \xi(x_n) \end{aligned}$$

where $\partial/\partial x_n \xi(x_n)$ is the derivative of the quantile function which, due to the inverse function theorem, equals one over the density evaluated at the quantile for continuous distributions, i.e.,

$$\frac{\partial}{\partial x_n}(M^{\xi_n})^{-1}(M^n(x_n)) = \frac{1}{m^{\xi_n}(\xi(x_n))} \frac{\partial}{\partial x_n} M^n(x_n) = \frac{m^n(x_n)}{m^{\xi_n}(\xi(x_n))}.$$

This concludes the proof. □

References

Yves Achdou, Jiequn Han, Jean-Michel Lasry, Pierre-Louis Lions, and Benjamin Moll. Income and wealth distribution in macroeconomics: A continuous-time approach. *The Review of Economic Studies*, 89(1):45–86, 04 2022.

SeHyoun Ahn, Greg Kaplan, Benjamin Moll, Thomas Winberry, and Christian Wolf. When inequality matters for macro and macro matters for inequality. *NBER Macroeconomics Annual*, 32:1–75, 2018.

Kristoffer Andersson, Adam Andersson, and Cornelis W Oosterlee. Convergence of a robust deep fbsde method for stochastic control. *SIAM Journal on Scientific Computing*, 45(1):A226–A255, 2023.

Adrien Auclert, Bence Bardóczy, Matthew Rognlie, and Ludwig Straub. Using the sequence-space jacobian to solve and estimate heterogeneous-agent models. *Econometrica*, 89(5):2375–2408, 2021.

Marlon Azinovic, Luca Gaegauf, and Simon Scheidegger. Deep equilibrium nets. *International Economic Review*, 63(4):1471–1525, 2022.

Christian Bayer and Ralph Luetticke. Solving discrete time heterogeneous agent models with aggregate risk and many idiosyncratic states by perturbation. *Quantitative Economics*, 11(4), 2020.

Christian Beck, Weinan E, and Arnulf Jentzen. Machine learning approximation algorithms for high-dimensional fully nonlinear partial differential equations and second-order backward stochastic differential equations. *Journal of Nonlinear Science*, 29:1563–1619, 2019.

Adrien Bilal. Solving heterogeneous agent models with the master equation. *NBER Working Paper 31103*, 2023.

Tomas Björk. *Arbitrage Theory in Continuous Time*. Oxford University Press, Oxford, 3 edition, 2009.

Timo Boppart, Per Krusell, and Kurt Mitman. Exploiting MIT shocks in heterogeneous-agent economies: the impulse response as a numerical derivative. *Journal of Economic Dynamics and Control*, 89:68–92, 2018.

Markus K. Brunnermeier and Yuliy Sannikov. A macroeconomic model with a financial sector. *American Economic Review*, 104(2):379–421, February 2014.

Pierre Cardaliaguet, François Delarue, Jean-Michel Lasry, and Pierre-Louis Lions. *The Master Equation and the Convergence Problem in Mean Field Games*, volume 201 of *Annals of Mathematics Studies*. Princeton University Press, 2019.

René Carmona and François Delarue. *Probabilistic Theory of Mean Field Games with Applications I: Mean Field FBSDEs, Control, and Games*, volume 83 of *Probability Theory and Stochastic Modelling*. Springer, 2018a.

René Carmona and François Delarue. *Probabilistic Theory of Mean Field Games with Applications II: Mean Field Games with Common Noise and Master Equations*, volume 84 of *Probability Theory and Stochastic Modelling*. Springer, 2018b.

Weiyu Chen, Baijiong Lin, Xiaoyuan Zhang, Xi Lin, Han Zhao, Qingfu Zhang, and James T. Kwok. Gradient-based multi-objective deep learning: Algorithms, theories, applications, and beyond. *arXiv 2501.10945 Working Paper*, 2025.

Patrick Cheridito, H Mete Soner, Nizar Touzi, and Nicolas Victoir. Second-order backward stochastic differential equations and fully nonlinear parabolic pdes. *Communications on Pure and Applied Mathematics*, 60(7):1081–1110, 2007.

George M Constantinides. Capital market equilibrium with transaction costs. *Journal of Political Economy*, 94(4):842–862, 1986.

Mark HA Davis and Andrew R Norman. Portfolio selection with transaction costs. *Mathematics of Operations Research*, 15(4):676–713, 1990.

Wouter J Den Haan. Comparison of solutions to the incomplete markets model with aggregate uncertainty. *Journal of Economic Dynamics and Control*, 34(1):4–27, 2010.

Victor Duarte, Diogo Duarte, and Dejanir H Silva. Machine learning for continuous-time finance. *The Review of Financial Studies*, 37(11):3217–3271, 2024.

Darrell Duffie. *Dynamic Asset Pricing Theory*. Princeton University Press, 3 edition, 2001.

Weinan E, Jiequn Han, and Arnulf Jentzen. Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. *Communications in mathematics and statistics*, 5(4): 349–380, 2017.

Oliver G. Ernst, Antje Mugler, Hans-Jörg Starkloff, and Elisabeth Ullmann. On the convergence of generalized polynomial chaos expansions. *ESAIM: Mathematical Modelling and Numerical Analysis*, 46(02):317–339, 2012.

Benjamin Fan, Edward Qiao, Anran Jiao, Zhouzhou Gu, Wenhao Li, and Lu Lu. Deep learning for solving and estimating dynamic macro-finance models. *Computational Economics*, pages 1–37, 2024.

Daniel Fehrle, Christopher Heiberger, and Johannes Huber. Polynomial chaos expansion: Efficient evaluation and estimation of computational models. *Computational Economics*, pages 1–64, 2025.

Jesús Fernández-Villaverde, Samuel Hurtado, and Galo Nuno. Financial frictions and the wealth distribution. *Econometrica*, 91(3):869–901, 2023.

Aleksandra Friedl, Felix Kübler, Simon Scheidegger, and Takafumi Usui. Deep uncertainty quantification: with an application to integrated assessment models. *Working Paper*, 2023.

Goutham Gopalakrishna. ALIENs and continuous time economies. *Swiss Finance Institute Research Paper No. 21-34*, 2021.

Goutham Gopalakrishna, Jonathan Payne, and Zhouzhou Gu. Asset pricing, participation constraints, and inequality. *Working Paper*, November 2024.

Zhouzhou Gu, Mathieu Laurière, Sebastian Merkel, and Jonathan Payne. Global solutions to master equations for continuous time heterogeneous agent macroeconomic models. *arXiv preprint arXiv:2406.13726*, 2024.

Jiequn Han, Arnulf Jentzen, and Weinan E. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.

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.

Daniel Harenberg, Stefano Marelli, Bruno Sudret, and Viktor Winschel. Uncertainty quantification and global sensitivity analysis for economic models. *Quantitative Economics*, 10(1):1–41, 2019.

Ji Huang. Breaking the curse of dimensionality in heterogeneous-agent models: A deep learning-based probabilistic approach. *SSRN Working Paper 4649043*, 2024.

Zhipeng Huang, Balint Negyesi, and Cornelis W Oosterlee. Convergence of the deep bsde method for stochastic control problems formulated through the stochastic maximum principle. *Mathematics and Computers in Simulation*, 227:553–568, 2025.

Côme Huré, Huyên Pham, and Xavier Warin. Deep backward schemes for high-dimensional nonlinear pdes. *Mathematics of Computation*, 89(324):1547–1579, 2020.

Sergei Isaenko. Portfolio choice under convex transaction costs. *Working Paper*, 2005.

Piotr Jaworski, Fabrizio Durante, Wolfgang Karl Härdle, and Tomasz Rychlik. *Copula Theory and Its Applications: Proceedings of the Workshop Held in Warsaw, 25-26 September 2009*. Lecture Notes in Statistics. Springer, 2010.

Greg Kaplan, Benjamin Moll, and Giovanni L Violante. Monetary policy according to HANK. *American Economic Review*, 108(3):697–743, 2018.

Per Krusell and Anthony A. Smith, Jr. Income and wealth heterogeneity, portfolio choice, and equilibrium asset returns. *Macroeconomic Dynamics*, 1(2):387–422, 1997.

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

Lilia Maliar, Serguei Maliar, and Pablo Winant. Deep learning for solving dynamic economic models. *Journal of Monetary Economics*, 122:76–101, 2021.

Benjamin Moll. The trouble with rational expectations in heterogeneous agent models: A challenge for macroeconomics. *Working Paper*, 2024.

Nikolas Nüsken and Lorenz Richter. Solving high-dimensional hamilton–jacobi–bellman pdes using neural networks: perspectives from the theory of controlled diffusions and measures on path space. *Partial differential equations and applications*, 2(4):48, 2021.

Jonathan Payne, Adam Rebei, and Yucheng Yang. Deep learning for search and matching models. *Available at SSRN*, 2024.

Huyên Pham and Xavier Warin. Mean-field neural networks-based algorithms for mckean-vlasov control problems. *arXiv preprint arXiv:2212.11518*, 2022.

Huyên Pham and Xavier Warin. Mean-field neural networks: learning mappings on wasserstein space. *Neural Networks*, 168:380–393, 2023.

Elisabeth Proehl. Approximating equilibria with ex-post heterogeneity and aggregate risk. *SSRN Working Paper 2620937*, 2025a.

Elisabeth Proehl. Existence and uniqueness of recursive equilibria with aggregate and idiosyncratic risk. *SSRN Working Paper 3250651*, 2025b.

Maziar Raissi, Paris Perdikaris, and George E Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational physics*, 378:686–707, 2019.

Michael Reiter. Solving heterogeneous-agent models by projection and perturbation. *Journal of Economic Dynamics and Control*, 33(3):649–665, 2009.

Andreas Schaab. Micro and macro uncertainty. *SSRN Working Paper 4099000*, 2020.

Simon Scheidegger and Ilias Bilionis. Machine learning for high-dimensional dynamic stochastic economies. *Journal of Computational Science*, 33:68–82, 2019.

Justin Sirignano and Konstantinos Spiliopoulos. Dgm: A deep learning algorithm for solving partial differential equations. *Journal of computational physics*, 375: 1339–1364, 2018.

Christian P Traeger. A 4-stated dice: Quantitatively addressing uncertainty effects in climate change. *Environmental and Resource Economics*, 59:1–37, 2014.

Dongbin Xiu and George E. Karniadakis. The Wiener-Askey polynomial chaos for stochastic differential equations. *SIAM Journal on Scientific Computing*, 24(2): 619–644, 2002.

Yucheng Yang, Chiyuan Wang, Andreas Schaab, and Benjamin Moll. Structural reinforcement learning for heterogeneous agent macroeconomics. *Working Paper*, 2025.

Tianhe Yu, Saurabh Kumar, Abhishek Gupta, Sergey Levine, Karol Hausman, and Chelsea Finn. Gradient surgery for multi-task learning. *Advances in Neural Information Processing Systems*, 33:5824–5836, 2020.

Shiji Zhou, Wenpeng Zhang, Jiyan Jiang, Wenliang Zhong, Jinjie GU, and Wenwu Zhu. On the convergence of stochastic multi-objective gradient manipulation and beyond. *Advances in Neural Information Processing Systems*, 35:38103–38115, 2022.