A Global Solution Method for HACT Models with Aggregate Risk

Preliminary and incomplete draft
All comments are welcome

Niklas Bonnmann*, Elisabeth Pröhl†

March 20, 2025

Abstract

Heterogeneous agent models in continuous time (HACT) have become a workhorse in macroeconomics, but incorporating aggregate risk remains a major computational challenge. Existing methods often rely on local approximations or restrict models to quasi-stationary settings, limiting their ability to capture nonlinear macroeconomic dynamics. This paper introduces a novel approach that globally solves HACT models with aggregate risk by leveraging the master equation. Our approach transforms this non-standard partial differential equation (PDE) into a high-dimensional yet standard PDE using Polynomial Chaos expansions, enabling the application of advanced but off-the-shelf deep learning techniques from the applied mathematics literature. Specifically, we adapt the Deep BSDE method to solve this equation efficiently. To demonstrate its applicability, we solve a two-asset HACT model with aggregate risk featuring adjustment costs and collateral constraints. By overcoming a key computational barrier, our method enables the study of fully nonlinear heterogeneous agent models, opening new avenues for macroeconomic research.

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 and Tinbergen Institute, Roetersstraat 11, 1018 WB Amsterdam, Netherlands, e.proehl@uva.nl.

We thank Simon Scheidegger for helpful discussions.

1 Introduction

Heterogeneous agent models have seen a remarkable revival in macroeconomic research following the influential work of Achdou, Han, Lasry, Lions, and Moll (2022) and Kaplan, Moll, and Violante (2018). While these models are particularly well-suited for analyzing critical questions related to inequality, only the recent methodological advances in solving these complex models propelled them to the forefront in macroeconomics. In particular, Achdou et al. (2022) leverage continuous-time modeling to use a finite-difference partial differential equation (PDE) solver to compute these models efficiently. However, a key methodological roadblock remains: incorporating aggregate risk. The core computational obstacle lies in the fact that the stochastic cross-sectional household distribution becomes an infinite-dimensional state variable, influencing both prices and policy functions. Although there has been 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 inherent, for instance, in macro-finance or environmental economics. In a recent essay, Moll (2024) highlights the difficulty of solving heterogeneous agent models with aggregate risk as one reason to move away from rational expectations. This paper directly addresses these challenges by providing a novel computational method that globally solves a general heterogeneous agent model in continuous time (HACT) with aggregate risk.[2]

Our method is designed to handle general HACT models where the aggregate Brownian motion can enter the budget constraint both indirectly and directly. Moreover, it accommodates non-stationary dynamics, non-trivial borrowing constraints and market clearing conditions, and, many state variables. We rely on the *master equation* which integrates the Kolmogorov Forward equation (KFE) with the Hamilton-Jacobi-Bellman equation to ensure the consistency of the cross-sectional distribution's law of motion. We solve this equation using neural networks. Al-

[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] While we acknowledge the significant value of expectation formation models that are micro-founded, we believe that their investigation should be primarily motivated by empirical evidence rather than the computational challenges associated with a different class of models. Furthermore, it is not unlikely that also the computational complexity of expectation formation models increases in future research. Thus, methods that solve complex models like ours may also be useful for models not following the rational expectations hypothesis.

though we are not the first to do so (see Gu, Laurière, Merkel, and Payne, 2024; Gopalakrishna, Payne, and Gu, 2024; Huang, 2024; Payne, Rebei, and Yang, 2024), the generality of our setting is novel. The existing work focuses on models with deterministic KFE which are are quasi-stationary, i.e., their optimal policies and prices are state-dependent but not time-dependent. In contrast, our method extends to fully stochastic settings and thus, allows for the incorporation of richer aggregate nonlinearities. To illustrate that this extension is important in practice, note that a stochastic KFE would already arise when incorporating idiosyncratic risk into the macro-finance model by Brunnermeier and Sannikov (2014). By drawing on various insights from applied mathematics – many of which have already proven valuable in different contexts of economics – we provide a robust and scalable solution to a long-standing computational challenge.

Our contributions are twofold. First, we heuristically derive the master equation for this general HACT model. The rigorous mathematical theory on *mean-field games* underlying this derivation has been established by Cardaliaguet, Delarue, Lasry, and Lions (2019); Carmona and Delarue (2018a,b). In contrast to these rigorous results, our goal is to provide an accessible derivation for economists who are familiar with the basics of continuous-time stochastic processes. While Bilal (2023) carefully derives this equation for quasi-stationary models, we are the first to do so for cases where the KFE is stochastic, enabling the study of genuinely non-stationary models. The key of deriving the master equation in this case is an extension of Itô's lemma. Originally, Itô's lemma yields the dynamics of functions of stochastic processes, whereas, its extension defines the dynamics of functions of stochastic processes and stochastic distributions.

The master equation is a non-standard partial differential equation (PDE) because it contains distributional derivatives of the value function. These derivatives are functional derivatives that evaluate how the value function would change given a marginal change in the cross-sectional distribution. Our second contribution lies in transforming this non-standard PDE into a high-dimensional yet standard PDE for which we adapt off-the-shelf solution methods from applied mathematics. The key is to leverage the equivalence of three different mathematical definitions of distributional derivatives that vary in how the cross-sectional distribution is represented, by a p.d.f., a c.d.f. or a random variable. Economically, this does not make any difference. However, representing distributional derivatives using random variables allows us to

exploit a general approximation theorem. Any square-integrable random variable can be approximated perfectly by a projection on a stochastic family of orthogonal polynomials when the degree of the polynomials approaches infinity. This projection is called a Polynomial Chaos expansion and was introduced as an approximation of the cross-sectional distribution in a discrete-time model by Proehl (2019). When projecting the cross-sectional distribution, the distributional derivatives in the master equation transform into derivatives w.r.t. the projection coefficients. This representation, hence, results in a high-dimensional yet standard PDE allowing us to adapt solution methods recently developed in the applied mathematics literature. We predominantly rely on the Deep BSDE method by E, Han, and Jentzen (2017) that uses neural networks to solve high-dimensional PDEs after converting them into a system of stochastic differential equations (SDE). Simulating these SDEs creates the training data to optimize an appropriate loss function with stochastic gradient descent. Since the neural networks essentially approximate the agents' optimal controls, this algorithm falls into the category of reinforcement learning. Lastly, we modify the Deep BSDE method according to insights by Nüsken and Richter (2021); Huang, Negyesi, and Oosterlee (2025) to improve the convergence for our strongly coupled system of forward backward stochastic differential equations (FBSDE).

Lastly, we illustrate this method with a standard two-asset HACT model with aggregate risk featuring a one-period zero-coupon bond and illiquid capital. We evaluate how hedging motives and the wealth distribution are affected by different collateral constraints and agent-specific endogenous labor choice.Overall, by overcoming a long-standing computational hurdle, this paper provides a foundation for future research into fully nonlinear heterogeneous agent models with aggregate risk – offering new insights into macroeconomic dynamics that were previously infeasible to explore.

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, follows the Deep BSDE method developed by E et al. (2017). This method reformulates the PDE as a system of stochastic differential equations and solves it using deep 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. 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 (2019) uses it as a discretization scheme for heterogeneous agent models. The key idea is to follow Proehl (2024) in rewriting recursive equilibrium in terms of square-integrable random variables, the space of which is fully spanned by polynomial chaos. Our approach is the first, to our knowledge, to integrate this discretization scheme into a deep learning framework for solving HACT models. Existing methods primarily rely on discretization using finite agents, finite-element methods, or a combination with neural network approximations.

Outline. The paper proceeds as follows. The next section introduces a general heterogeneous agent real business cycle model 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. Lastly, Section 6 applies the proposed algorithm to study the interaction of three sources of inequality – as considered by Kaplan et al. (2018) – in a macro-finance model with aggregate risk: uninsurable income risk, non-trivial borrowing constraints, and risky asset illiquidity. Section 7 concludes.

2 A General HACT Model

We introduce a general heterogeneous agent real business cycle model to facilitate the use of the methodology introduced in this paper for a range of applications.

Environment. Time is continuous $t \in [0, \infty)$. The economy is populated by a continuum of agents $i \in [0, 1]$. Suppose that all stochastic processes are defined on a generic filtered probability space $(\Omega, \Sigma, \mathcal{F}_{t\geq 0}, \mathbb{P})$. The aggregate risk is modeled by an aggregate d^{ag}-dimensional Brownian motion $\{B_t^{ag}\}_{t\geq 0}$. Similarly, agent i's idiosyncratic risk is given by a d^{id}-dimensional Brownian motion $\{B_t^i\}_{t\geq 0}$. This process is i.i.d. across agents $i \in [0, 1]$. Furthermore, we assume that the entries of the vectors B_t^{ag} and B_t^i are pairwise independent for any $i \in [0, 1]$.[3]

Household Problem. All agents face the same optimization problem. At any

[3]Nevertheless, correlation between the aggregate and idiosyncratic risk can be introduced as is illustrated in our example model in Section 6.

instant $t \geq 0$, agent $i \in [0,1]$ controls their N^x-dimensional vector of agent-specific state variables $\tilde{x}_t^i$ with an N^y-dimensional vector of controls y_t^i to maximize their flow utility $u(y_t^i, \tilde{x}_t^i)$ which is discounted at the constant rate ρ. Agents optimize while observing an N^a-dimensional vector of aggregate state variables a_t. This optimization is subject to a N^e-dimensional function of state constraints $e(t, \tilde{x}_t^i, a_t) \geq 0$. 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. For that reason, they play the same role as the Brownian innovations when stating the household problem. Overall, agent $i \in [0,1]$ optimizes

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

where the drift and dispersion functions are given by $\tilde{\mu} : \mathbb{R}^{N^y+N^x+N^a+N^r} \to \mathbb{R}^{N^x}$, $\tilde{\sigma}^o : \mathbb{R}^{N^y+N^x+N^a+N^r} \to \mathbb{R}^{N^x \times d^o}$ with $o \in \{id, ag\}$, $\eta : \mathbb{R}^{N^a+N^r} \to \mathbb{R}^{N^a}$ and $\zeta : \mathbb{R}^{N^a+N^r} \to \mathbb{R}^{N^a \times d^{ag}}$. Note that the state vector $\tilde{x}_t^i$ can contain exogenous states as well. In that case, the corresponding drift and dispersion entries are constant across agent-specific variables $(y_t^i, \tilde{x}_t^i)$. The states' aggregate and idiosyncratic covariance matrices are defined as

$$\begin{aligned}
\tilde{\Sigma}^o\left(y^i, \tilde{x}^i, a, r\right) =& \tilde{\sigma}^o\left(y^i, \tilde{x}^i, a, r\right) \tilde{\sigma}^o\left(y^i, \tilde{x}^i, a, r\right)^T \in \mathbb{R}^{N^x \times N^x}, o \in \{id, ag\} \qquad (1) \\
\tilde{\Sigma}^{xa}\left(y^i, \tilde{x}^i, a, r\right) =& \tilde{\sigma}^{ag}\left(y^i, \tilde{x}^i, a, r\right) \zeta(a, r)^T \in \mathbb{R}^{N^x \times N^a} \\
\Sigma^a =& \zeta(a, r)\zeta(a, r)^T \in \mathbb{R}^{N^a \times N^a}.
\end{aligned}$$

Example 1 (Krusell and Smith (1998))**.** *The continuous-time version of this model fits our framework as follows. The agent-specific state $\tilde{x}_t^i = [l_t^i, k_t^i]$ consists of agent i's exogenous labor productivity and her endogenous capital savings. The agent's controls solely consist of her consumption $y_t^i = c_t^i$. Furthermore, each agent observes the aggregate TFP shock $a_t = z_t$. The budget constraint determining the savings' dynamics also depends on the aggregate marginal poductivities $r_t = [\partial F_t/\partial L, \partial F_t/\partial K]$ that the agents take as given. Lastly, the constraint is defined by $e(t, \tilde{x}_t^i, a_t) = k_t^i - \bar{k}$.*

It is well known (see e.g. Achdou et al., 2022; Bilal, 2023) that the state constraint can be transformed into a state boundary condition. To do so, we first assume sufficient regularity of the state constraint and then rewrite the state dynamics.

Assumption 2 (State Constraint Regularity)**.** *Assume that the state constraint function* $e : [0, \infty) \times \mathbb{R}^{N^x+N^a} \to \mathbb{R}^{N^e}$ *is differentiable once in time and twice in the states. Moreover, suppose that* $N^e \leq N^x$*. W.l.o.g. divide the agent-specific state into its unconstrained and constrained part, i.e.,* $\tilde{x}^i = [\tilde{x}^{i,uc}, \tilde{x}^{i,c}]^T$ *where* T *denotes the matrix transpose. Assume that the state constraint function* e *is invertible for* $\tilde{x}^{i,c}$*, i.e.,* $e^{-1}(t, \tilde{x}_t^{i,uc}, \tilde{e}, a_t) = \{\tilde{x}_t^{i,c} \in \mathbb{R}^{N^e} | e(t, \tilde{x}_t^{i,uc}, \tilde{x}_t^{i,c}, a_t) = \tilde{e}\}$ *is a singleton for any* $\tilde{e} \in \mathbb{R}^{N^e}$*.*

The constraint $\tilde{e}_t^i = e(t, \tilde{x}_t^i, a_t)$ follows the dynamics

$$d\tilde{e}_t^i = \psi\left(t, y_t^i, \tilde{x}_t^i, a_t, r_t\right) dt + \upsilon^{id}\left(y_t^i, \tilde{x}_t^i, a_t, r_t\right) dB_t^i + \upsilon^{ag}\left(y_t^i, \tilde{x}_t^i, a_t, r_t\right) dB_t^{ag},$$

where the j-th rows of the drift and dispersion functions ψ_j, υ_j^o, $o \in \{id, ag\}$, with $j \in \{1, \ldots, N^e\}$ follow trivially from Itô's formula

$$\begin{aligned}
\psi_j =& e_t^j + \sum_{n=1}^{N^a} e_{a_n}^j \eta_n + \frac{1}{2} \sum_{n,n'=1}^{N^a} e_{a_n a_{n'}}^j \Sigma_{nn'}^a + \sum_{n=1}^{N^x} e_{\tilde{x}_n^i}^j \tilde{\mu}_n \\
&+ \frac{1}{2} \sum_{n,n'=1}^{N^x} \sum_{o \in \{id,ag\}} e_{\tilde{x}_n^i \tilde{x}_{n'}^i}^j \tilde{\Sigma}_{nn'}^o + \frac{1}{2} \sum_{n=1}^{N^x} \sum_{n'=1}^{N^a} e_{\tilde{x}_n^i a_{n'}}^j \tilde{\Sigma}_{nn'}^{xa} \\
\upsilon_j^{id} =& \sum_{n=1}^{N^x} e_{\tilde{x}_n^i}^j \tilde{\sigma}_n^{id} \\
\upsilon_j^{ag} =& \sum_{n=1}^{N^a} e_{a_n}^j \zeta_n + \sum_{n=1}^{N^x} e_{\tilde{x}_n^i}^j \tilde{\sigma}_n^{ag},
\end{aligned}$$

with the covariance matrices as in (1) and where the subscripts of e^j denote its derivatives. Function arguments have been dropped for notational convenience. Note that $e^j, e_t^j \in \mathbb{R}$, $e_a^j \in \mathbb{R}^{N^a}$, $e_{\tilde{x}^i}^j \in \mathbb{R}^{N^x}$, $e_{aa}^j \in \mathbb{R}^{N^a \times N^a}$, $e_{\tilde{x}^i a}^j \in \mathbb{R}^{N^x \times N^a}$ and $e_{\tilde{x}^i \tilde{x}^i}^j \in \mathbb{R}^{N^x \times N^x}$. Considering Assumption 2, the individual optimization problem can be

restated as

$$\begin{aligned}
\max_{\{y_t^i\}_{t\geq 0}} \mathbb{E}_0 &\left[\int_0^\infty e^{-\rho t} u\left(y_t^i, \tilde{x}_t^{i,uc}, \tilde{x}_t^{i,c}\right) dt\right] \qquad (2)\\
\text{s.t. } d\tilde{x}_t^{i,uc} &= \tilde{\mu}^{uc}\left(y_t^i, \tilde{x}_t^{i,uc}, \tilde{x}_t^{i,c}, a_t, r_t\right) dt + \tilde{\sigma}^{id,uc}\left(y_t^i, \tilde{x}_t^{i,uc}, \tilde{x}_t^{i,c}, a_t, r_t\right) dB_t^i \\
&\quad + \tilde{\sigma}^{ag,uc}\left(y_t^i, \tilde{x}_t^{i,uc}, \tilde{x}_t^{i,c}, a_t, r_t\right) dB_t^{ag} \\
\tilde{x}_t^{i,c} &= e^{-1}\left(t, \tilde{x}_t^{i,uc}, \tilde{e}_t^i, a_t\right) \\
da_t &= \eta\left(a_t, r_t\right) dt + \zeta\left(a_t, r_t\right) dB_t^{ag} \\
d\tilde{e}_t^i &= \psi\left(t, y_t^i, \tilde{x}_t^{i,uc}, \tilde{x}_t^{i,c}, a_t, r_t\right) dt + \upsilon^{id}\left(y_t^i, \tilde{x}_t^{i,uc}, \tilde{x}_t^{i,c}, a_t, r_t\right) dB_t^i \\
&\quad + \upsilon^{ag}\left(y_t^i, \tilde{x}_t^{i,uc}, \tilde{x}_t^{i,c}, a_t, r_t\right) dB_t^{ag} \\
\tilde{e}_t^i &\geq 0,
\end{aligned}$$

where superscript uc denotes the first $N^x - N^e$ rows of the respective vectors and matrices. Given this restated optimization problem, we can now substitute the agent-specific states by $x_t^i = [\tilde{x}_t^{i,uc}, \tilde{e}_t^i]^T$ with dynamics

$$dx_t^i = \mu\left(t, y_t^i, x_t^i, a_t, r_t\right) dt + \sigma^{id}\left(y_t^i, x_t^i, a_t, r_t\right) dB_t^i + \sigma^{ag}\left(y_t^i, x_t^i, a_t, r_t\right) dB_t^{ag} \qquad (3)$$

with

$$\begin{aligned}
\mu\left(t, y_t^i, x_t^i, a_t, r_t\right) &= \begin{bmatrix} \tilde{\mu}^{uc}\left(y_t^i, \tilde{x}_t^{i,uc}, e^{-1}(t, x_t^i, a_t), a_t, r_t\right) \\ \psi\left(t, y_t^i, \tilde{x}_t^{i,uc}, e^{-1}(t, x_t^i, a_t), a_t, r_t\right) \end{bmatrix} \\
\sigma^o\left(y_t^i, x_t^i, a_t, r_t\right) &= \begin{bmatrix} \tilde{\sigma}^{o,uc}\left(y_t^i, \tilde{x}_t^{i,uc}, e^{-1}(t, x_t^i, a_t), a_t, r_t\right) \\ \upsilon^o\left(y_t^i, \tilde{x}_t^{i,uc}, e^{-1}(t, x_t^i, a_t), a_t, r_t\right) \end{bmatrix}, o \in \{id, ag\}.
\end{aligned}$$

This yields the following idiosyncratic and aggregate covariance matrices

$$\begin{aligned}
\Sigma^o\left(y^i, x^i, y^{i'}, x^{i'}, a, r\right) &= \sigma^o\left(y^i, x^i, a, r\right) \sigma^o\left(y^{i'}, x^{i'}, a, r\right)^T \in \mathbb{R}^{N^x \times N^x}, o \in \{id, ag\} \qquad (4)\\
\Sigma^{xa}\left(y^i, x^i, a, r\right) &= \sigma^{ag}\left(y^i, x^i, a, r\right) \zeta(a, r)^T \in \mathbb{R}^{N^x \times N^a} \\
\Sigma^a(a, r) &= \zeta(a, r)\zeta(a, r)^T \in \mathbb{R}^{N^a \times N^a}.
\end{aligned}$$

Optimal Controls. The household optimization problem is a standard optimal control problem which can be solved by dynamic programming.[4] Define the value

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

function by

$$J(t, x_t^i, a_t) = \max_{\{y_s^i\}_{s\geq t}} \mathbb{E}_t \left[\int_t^\infty e^{-\rho(s-t)} u\left(y_t^i, \tilde{x}_t^{i,uc}, e^{-1}\left(t, x_t^i, a_t\right)\right) ds \right]. \tag{5}$$

The corresponding Hamilton-Jacobi-Bellman (HJB) equation for the value function $J(t, x^i, a)$ and its boundary conditions can be derived using standard dynamic programming results. Define the 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

$$\begin{aligned}\mathcal{H}^{HJB}[J, y^i, r] = &- \rho J + u + J_t + \sum_{n=1}^{N^a} J_{a_n}\eta_n + \sum_{n=1}^{N^x} J_{x_n^i}\mu_n + \frac{1}{2}\left[\sum_{n,n'=1}^{N^a} J_{a_n a_{n'}}\Sigma^a_{nn'}\right. \\ &\left. + \sum_{n,n'=1}^{N^x}\sum_{o\in\{id,ag\}} J_{x_n^i x_{n'}^i}\Sigma^o_{nn'} + \sum_{n=1}^{N^x}\sum_{n'=1}^{N^a} J_{x_n^i a_{n'}}\Sigma^{xa}_{nn'}\right]\end{aligned} \tag{6}$$

where the subscripts of J denote its derivatives. Function arguments have been dropped for notational convenience. The covariance matrices are given in (4).

Define the operator for the terminal boundary condition $\mathcal{H}^{TBC}[J](x^i, a)$ mapping the function J into a function of (x^i, a) by

$$\mathcal{H}^{TBC}[J](x^i, a) = \lim_{T\to\infty} J(T, x^i, a). \tag{7}$$

The state constraint results in an additional state boundary condition (see e.g. Achdou et al., 2022; Bilal, 2023). To guarantee that the constraint holds at all times, the drift of $\tilde{e}_t^i$ in (2) at $\tilde{e}_t^i = 0$ must be nonnegative and the dispersion terms must vanish, i.e., for any $[t, y^i, \tilde{x}^{i,uc}, a, r] \in \mathbb{R}^{1+N^y+N^x-N^e+N^a+N^r}$, it must hold that

$$\begin{aligned}0 \leq &\psi\left(t, y^i, \tilde{x}^{i,uc}, e^{-1}(t, \tilde{x}^{i,uc}, 0, a), a, r\right) \\ 0 = &\upsilon^{id}\left(y^i, \tilde{x}^{i,uc}, e^{-1}(t, \tilde{x}^{i,uc}, 0, a), a, r\right) \\ 0 = &\upsilon^{ag}\left(y^i, \tilde{x}^{i,uc}, e^{-1}(t, \tilde{x}^{i,uc}, 0, a), a, r\right).\end{aligned}$$

This implies that part of the HJB equation vanishes at the state boundary $\tilde{e}^i = 0$. This vanishing part constitutes the additional boundary condition. Therefore, we define the operator for the state boundary condition $\mathcal{H}^{SBC}[J, y^i, r](t, \tilde{x}^{i,uc}, a)$ mapping the

functions $[J, y^i, r]$ into a function of $(t, \tilde{x}^{i,uc}, a)$ by

$$\begin{aligned}
&\mathcal{H}^{SBC}[J, y^i, r] = \\
&\lim_{\tilde{e}^i \to 0} \sum_{n=1}^{N^e} J_{\tilde{e}^i_n} \min(0, \psi_n) + \sum_{n=1}^{N^x - N^e} \sum_{n'=1}^{N^e} \sum_{o \in \{id, ag\}} J_{\tilde{x}^{i,uc}_n \tilde{e}^i_{n'}} \Sigma^o_{n, N^x - N^e + n'} \\
&\quad + \frac{1}{2} \left[\sum_{n,n'=1}^{N^e} \sum_{o \in \{id, ag\}} J_{\tilde{e}^i_n \tilde{e}^i_{n'}} \Sigma^o_{N^x - N^e + n, N^x - N^e + n'} + \sum_{n=1}^{N^e} \sum_{n'=1}^{N^a} J_{\tilde{e}^i_n a_{n'}} \Sigma^{xa}_{N^x - N^e + n, n'} \right]
\end{aligned} \tag{8}$$

The HJB equation and its corresponding boundary conditions are then given by

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

The corresponding optimal controls feature the same function arguments as the operator $\mathcal{H}^{HJB}[J, y^i, r]$ such that $y^{i,*}_t = y^{i,*}(t, x^i_t, a_t, r_t)$. The w-th element of the optimal control vector $y^{i,*} = [y^{i,*}_1, \ldots, y^{i,*}_{N^y}]^T$, $w \in \{1, \ldots, N^y\}$, satisfies the first-order condition

$$\begin{aligned}
0 =& \frac{\partial}{\partial y^i_w} \mathcal{H}^{HJB}[J, y^i, r](t, x^i, a) \\
=& \frac{\partial}{\partial y^i_w} \left\{ u + \sum_{n=1}^{N^x} J_{x^i_n} \mu_n + \frac{1}{2} \left[\sum_{n,n'=1}^{N^x} \sum_{o \in \{id, ag\}} J_{x^i_n x^i_{n'}} \Sigma^o_{nn'} + \sum_{n=1}^{N^x} \sum_{n'=1}^{N^a} J_{x^i_n a_{n'}} \Sigma^{xa}_{nn'} \right] \right\}.
\end{aligned} \tag{9}$$

Market Clearing. The vector of relative returns r_t is determined by market clearing. The optimally controlled states $x^{i,*}_t$ differ across agents due to the idiosyncratic risk, which we summarize by the cross-sectional cumulative distribution function $M_t \in \mathcal{P}(\mathbb{R}^{N^x})$, i.e., $M_t : \mathbb{R}^{N^x} \to [0, 1]$. Naturally, the cross-sectional density m_t is given by $M_t(x) = \int_{z \leq x} m_t(z) dz$. Market clearing aggregates the agents' states and choices according to this distribution. Thus, the relative returns r^* that clear markets are defined by

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

where $\mathcal{H}^{MC}[r, y^i](t, M, a)$ is the market-clearing operator mapping the functions r and

y into a function of (t, M, a). This implies $r_t^* = r^*(t, M_t, a_t)$. We assume that the market clearing conditions have the following structure.

Assumption 3 (Market Clearing Regularity)**.** *Suppose that the market clearing operator $\mathcal{H}^{MC}[r, y^i](t, M, a) \in \mathbb{R}^{N^r}$ for any $(t, M, a) \in [0, \infty) \times \mathcal{P}(\mathbb{R}^{N^x}) \times \mathbb{R}^{N^a}$, i.e., there are as many conditions as market instruments. Additionally, assume that the n-th market clearing condition $\mathcal{H}_n^{MC}[r, y^i](t, M, a) = 0$, $n \in \{1, \dots, N^r\}$, is equivalent to either of the following equations.*

(*i*) *Explicit market clearing: $r_n^* = \mathcal{M}_n^{ex}(t, M, a)$ for some function $\mathcal{M}^{ex} : [0, \infty) \times \mathcal{P}(\mathbb{R}^{N^x}) \times \mathbb{R}^{N^a} \to \mathbb{R}$*

(*ii*) *Implicit market clearing: $0 = \mathcal{M}_n^{im}[r^*, y^{i,*}](t, M, a)$ for some function $\mathcal{M}^{im} : [0, \infty) \times \mathcal{P}(\mathbb{R}^{N^x}) \times \mathbb{R}^{N^a} \to \mathbb{R}$ or operator $\mathcal{M}_n^{im}$ acting on r^* and/or y^i*

Example 4. *In Krusell and Smith (1998), the labor and capital market clearing condition can be posed in the explicit form. In an endowment economy similar to Huggett (1993) with a bond and a stock in zero and unit net supply, respectively, the bond market clearing condition is in implicit form – though without a direct dependence on r^* or $y^{i,*}$ –, whereas, the stock market clearing condition can be converted into explicit form by aggregating over the total financial wealth yielding the stock price.*

Consistent Law of Motion. The cross-sectional distribution M_t used for market clearing needs to be consistent with the agents' optimal choices. Agent i's optimally controlled state $x_t^{i,*} = [\tilde{x}_t^{i,uc,*}, \tilde{e}_t^{i,*}]$ is defined by its dynamics as in (3) such that $y_t^i = y^{i,*}(t, x_t^{i,*}, a_t, r_t)$ solves (9). 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[5]

$$
\begin{aligned}
dm_t(x) = \sum_{n=1}^{N^x} \Bigg\{ & -\frac{\partial}{\partial x_n} \left[\mu_n \left(t, y^{i,*}(t, x, a_t, r_t), x, a_t, r_t \right) m_t(x) \right] \\
& + \frac{1}{2} \sum_{n'=1}^{N^x} \sum_{o \in \{id, ag\}} \frac{\partial^2}{\partial x_n \partial x_{n'}} \left[\Sigma_{nn'}^o \left(y^{i,*}(t, x, a_t, r_t), x, a_t, r_t \right) m_t(x) \right] \Bigg\} dt \\
& - \sum_{n=1}^{N^x} \frac{\partial}{\partial x_n} \left[\sigma_n^{ag} \left(y^{i,*}(t, x, a_t, r_t), x, a_t, r_t \right) m_t(x) \right] dB_t^{ag}
\end{aligned}
\tag{10}
$$

[5]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).

for any $x \in \mathbb{R}^{N^x}$. The dynamics of the distribution sum over the first- and second-order variations stemming from changes in each element of the agent-specific state vector. The Kolmogorov forward equation is stochastic. This follows from the aggregate Brownian shock in (3) which is the same for all agents and, thus, does not wash out when taking the cross-sectional average. We refer to this property ($\sigma^{ag} > 0$) as *direct* aggregate risk. This paper, hence, 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 this case, aggregate Brownian shocks only affect the dynamics of aggregate states – such as TFP in Krusell and Smith (1998) – and the Kolmogorov forward equation has a stochastic drift but no volatility term. Furthermore, the framework in this paper allows for controlled volatility in agent-specific states and, thus, for asset pricing applications.[6]

To derive the master equation in the following section, we also need the dynamics of the marginal distribution $M_t(x_n) = \lim_{x_{-n}\to\infty} M_t(x)$ where $M_t(x) = \int_{z\leq x} m_t(z)dz$. Note that $x_n \in \mathbb{R}$ denotes the n-th element of the state vector $x \in \mathbb{R}^{N^x}$, whereas, $x_{-n} \in \mathbb{R}^{N^x-1}$ excludes the n-th element from the state vector. Given (10), the dynamics of the marginals read as follows.

Proposition 5 (Consistent Law of Motion for the Marginal Distributions)**.** *Given that the process of cross-sectional p.d.f.s $\{m_t(x)\}_{t\geq 0, x\in\mathbb{R}^{N^x}}$ follows (10), the dynamics of the marginal c.d.f.s $\{M_t(x_n)\}_{t\geq 0, x_n\in\mathbb{R}}$ are given by*

$$
\begin{aligned}
dM_t(x_n) = \Bigg\{ &- \int_{\mathbb{R}^{N^x-1}} \left[\mu_n\left(t, y^{i,*}(t,z,a_t,r_t), z, a_t, r_t\right) m_t(z)\right]\Big|_{z_n=x_n} dz_{-n} \\
&+ \frac{1}{2}\sum_{n'=1}^{N^x}\sum_{o\in\{id,ag\}} \int_{\mathbb{R}^{N^x-1}} \left[\frac{\partial}{\partial z_{n'}}\Sigma^o_{nn'}\left(y^{i,*}(t,z,a_t,r_t), z, a_t, r_t\right) m_t(z)\right]\Big|_{z_n=x_n} dz_{-n} \Bigg\} dt \\
&- \left\{\int_{\mathbb{R}^{N^x-1}} \left[\sigma^{ag}_n\left(y^{i,*}(t,z,a_t,r_t), z, a_t, r_t\right) m_t(z)\right]\Big|_{z_n=x_n} dz_{-n}\right\} dB^{ag}_t.
\end{aligned} \tag{11}
$$

Working with the dynamics of the marginal c.d.f.s has computational advantages over the marginal densities. Apart from reducing the need to compute derivatives, it also allows for a more straightforward numerical approximation of the master equation.

[6] In existing work in economics, Huang (2024) is the only one to consider controlled volatility.

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* $\{r_t^*\}_{t\geq 0}$*, households choose their controls* $y_t^{i,*} = y^{i,*}(t, x^i, a)$ *to optimize their value function* $J(t, x^i, a)$ *such that*

$$0 = \frac{\partial}{\partial y_n^i}\mathcal{H}^{HJB}[J, y^{i,*}, r^*](t, x^i, a),\ n \in \{1, \dots, N^y\}, \tag{12}$$

$$0 = \mathcal{H}^{HJB}[J, y^{i,*}, r^*](t, x^i, a) \tag{13}$$

$$0 = \mathcal{H}^{TBC}[J](x^i, a) \tag{14}$$

$$0 = \mathcal{H}^{SBC}[J, y^{i,*}, r^*](t, \tilde{x}^{i,uc}, a), \tag{15}$$

where $\mathcal{H}^{HJB}$*,* $\mathcal{H}^{TBC}$ *and* $\mathcal{H}^{SBC}$ *are defined as in (6), (7) and (8), respectively.*

(*ii*) *The cross-sectional density* $\{m_t\}_{t\geq 0}$ *of the optimally controlled individual states follows the consistent law of motion in (10) such that the cross-sectional marginal distribution functions* $\{M_t\}_{t\geq 0}$ *follow the consistent laws of motion in (11).*

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

$$0 = \mathcal{H}^{MC}[r^*, y^{i,*}](t, M, a). \tag{16}$$

3 Deriving the Master Equation

The master equation essentially combines conditions (*i*) and (*ii*) in Definition 6 so that equilibrium can be described by a set of equations $0 = \mathcal{H}[J, y^{i,*}, r^*]$. This section derives the master equation for the general heterogeneous agent model with aggregate risk in a heuristic manner. The goal is to make the rigorous derivations in Carmona and Delarue (2018a,b) and Cardaliaguet et al. (2019) accessible for a wider quantitative economics audience and offer a recipe to apply the methodology in various contexts.

Let us first shed light on how the master equation combines household optimality and the consistent law of motion. The HJB equation in (13) ensures that the value function in equilibrium is a (local) martingale and agents smooth consumption. When setting up the agents' optimization problem, we postulated that they take the returns

as given. What does this mean in detail? Agents are fully rational. They know how the returns are determined through the market clearing equation in (16). Agents also have full knowledge of the cross-sectional distribution and its consistent law of motion in (10). However, each agent is assumed infinitesimal so that a single agent's deviation from the recursive policy function $y^{i,*}(t, x_t^i, a_t)$ would not change the distribution or returns. Thus, in deriving the first-order condition for the optimal controls $y_t^{i,*}$ in (12), 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 the exogenous variation stemming from the aggregate and idiosyncratic Brownian motion and the variation stemming from changes in the cross-sectional distribution taken as given by the agent. In order to ensure consistency of the distribution's law of motion with agents' choices, this latter variation needs to be considered separately. This is achieved by specifying the equilibrium value function in further detail by setting $J(t, x_t^i, a_t) := \mathcal{J}(t, M_t, x_t^i, a_t)$.

The master equation then results from rewriting the HJB equation in (13) in terms of $\mathcal{J}$. Note that the HJB operator defined in (6) subtracts the time-derivative of the value function in (5) which equals $\rho J(t, x_t^i, a_t) - u(y_t^i, x_t^i)$ from the drift of the Itô process $J(t, x_t^i, a_t)$. Thus, we will subsequently define the master operator equivalently by subtracting the time-derivative of the value function in (5) from the drift of the stochastic process $\mathcal{J}(t, M_t, x_t^i, a_t)$. This, however, is easier said than done because the distribution M_t is not a standard Itô process such that one cannot simply apply Itô's formula to derive the drift of $\mathcal{J}$. The following subsection explains how Itô's calculus can be adjusted to accommodate the distribution process.

3.1 The Cross-Sectional Distribution as a Stochastic Process

It is important to note that the aggregate distribution $M_t(x)$ is a stochastic process that is different from the aggregate state process a_t which is a standard Itô process. For any fixed time point t, a_t is a real vector, whereas, $M_t(x)$ is a real-valued function. Thus, the distribution does not only depend on time, but also on its argument $x \in \mathbb{R}^{N^x}$. Compared to the state a_t which is a stochastic process with the standard time index set $t \in [0, \infty)$, the distribution $M_t(x)$ and, similarly, the density $m_t(x)$ are generalized stochastic processes with the larger index set $(t, x) \in [0, \infty) \times \mathbb{R}^{N^x}$.[7]

[7] Generalized stochastic processes are called random fields which have been introduced in economics by Proehl (2024) in the context of the discrete-time heterogeneous-agent RBC model as in Krusell and Smith (1998).

With this extended index set, the standard Itô formula is no longer sufficient to derive the dynamics of a real-valued function of time and the distribution process $f(t, M_t)$. A suitable expansion requires well-defined distributional derivatives of the function f. First, note that we can define f in terms of a function of the density $\tilde{f}(t, m_t) = \tilde{f}(t, dM_t/dx) =: f(t, M_t)$.

Definition 7 (Distributional Derivatives (see Cardaliaguet et al., 2019, Definitions 2.2.1, 2.2.2 and Section 2.2.2)). *Let $\mathbf{m}_t(x) = [m_t(x_1), \dots, m_t(x_{N^x})]$ denote the vector of marginal densities $m_t(x_n) = \int m_t(x_n, x_{-n}) dx_{-n}$. Then, the derivative of $\tilde{f}$ w.r.t. the density is defined by the Riesz representation given by*

$$\lim_{s \to 0} \frac{\tilde{f}(t, m_t + s(m'_t - m_t)) - \tilde{f}(t, m_t)}{s} = \int_{\mathbb{R}^{N^x}} \tilde{f}_m(t, m_t(x))^T \left(\mathbf{m}'_t(x) - \mathbf{m}_t(x)\right) dx,$$

where $\tilde{f}_m(t, m_t(x)) \in \mathbb{R}^{N^x}$. The derivative w.r.t. the c.d.f. f_M is then defined by

$$f_{M_n}(t, M_t(x)) := \frac{d}{dx_n} \tilde{f}_{m_n}(t, m_t(x)).$$

The second-order derivatives are similarly defined by

$$\lim_{s \to 0} \frac{\tilde{f}_m(t, (m_t + s(m'_t - m_t))(x)) - \tilde{f}_m(t, m_t(x))}{s} = \int_{\mathbb{R}^{N^x}} \tilde{f}_{mm}(t, m_t(x, x'))^T \left(\mathbf{m}'_t(x') - \mathbf{m}_t(x')\right) dx',$$

where $\tilde{f}_{mm}(t, m_t(x, x')) \in \mathbb{R}^{N^x \times N^x}$ such that

$$f_{M_n M_{n'}}(t, M_t(x, x')) := \frac{d^2}{dx_n dx'_{n'}} \tilde{f}_{m_n m_{n'}}(t, m_t(x, x')).$$

The first-order derivative may differ for various arguments x of the distribution, i.e., it evaluates how $\tilde{f}$ changes when the density only changes at the argument x and nowhere else. Similarly, the second-order derivative evaluates how $\tilde{f}$ changes when the density changes at x_n and $x'_{n'}$ but nowhere else.

Given $\tilde{f}_m$ and $\tilde{f}_{mm}$, the definition of the derivatives w.r.t. the c.d.f. are judiciously chosen to ensure that the second-order Taylor expansions of $\tilde{f}(t, m_t) =: f(t, M_t)$ coincide. The second-order Taylor expansion of $\tilde{f}$ follows from the definition of the

derivative w.r.t. the density which implies

$$\tilde{f}(t,m'_t) - \tilde{f}(t,m_t) = \int_0^1 \sum_{n=1}^{N^x} \int_{\mathbb{R}} \tilde{f}_{m_n}(t, m_{t,s}(x))(m'_t(x_n) - m_t(x_n))dx_n ds$$
$$+ \frac{1}{2}\int_0^1 \sum_{n,n'=1}^{N^x} \int_{\mathbb{R}^2} \tilde{f}_{m_n m_{n'}}(t, m_{t,s}(x,x'))(m'_t(x'_{n'}) - m_t(x'_{n'}))$$
$$(m'_t(x_n) - m_t(x_n))dx'_{n'}dx_n ds$$

with $m_{t,s} = m_t + s(m'_t - m_t)$. Integrating by parts once in the first term and twice in the second term yields

$$\tilde{f}(t,m'_t) - \tilde{f}(t,m_t) = \tag{17}$$
$$- \int_0^1 \sum_{n=1}^{N^x} \int_{\mathbb{R}} f_{M_n}(t, M_{t,s}(x))(M'_t(x_n) - M_t(x_n))dx_n ds$$
$$+ \frac{1}{2}\int_0^1 \sum_{n,n'=1}^{N^x} \int_{\mathbb{R}^2} f_{M_n M_{n'}}(t, M_{t,s}(x,x'))(M'_t(x'_{n'}) - M_t(x'_{n'}))$$
$$(M'_t(x_n) - M_t(x_n))dx'_{n'}dx_n ds$$
$$=: f(t,M'_t) - f(t,M_t)$$

with $d/dx M_{t,s}(x) = m_{t,s}(x)$. Heuristically, adding a change in the time argument $\Delta t = t' - t$ to this second-order Taylor expansion and setting $M'_t = M_{t'}$ such that $(M'_t(x_n) - M_t(x_n)) = (M_{t'}(x_n) - M_t(x_n)) = \Delta M_t(x_n)$ rationalizes the Itô expansion of a function $f(t, M_t)$. Note that $\Delta M_t(x_n)$ becomes $dM_t(x_n)$ following (11) when the time change is infinitesimal. If M_t is the cross-sectional c.d.f. consistent with the idiosyncratic state following (3), the appropriate expansion of a function $f(t, M_t)$

reads[8]

$$
\begin{aligned}
&df(t, M_t) = \\
&\Bigg\{ f_t(t, M_t) + \sum_{n=1}^{N^x} \int_{\mathbb{R}^{N^x}} f_{M_n}(t, M_t(x)) \mu_n \left(t, y^{i,*}(t, x, a_t, r_t), x, a_t, r_t\right) m_t(x) dx \\
&+ \frac{1}{2} \sum_{n,n'=1}^{N^x} \sum_{o \in \{id, ag\}} \int_{\mathbb{R}^{N^x}} f_{M_n x_{n'}}(t, M_t(x)) \Sigma^{o,*}_{nn'}(t, x, x) m_t(x) dx \\
&+ \frac{1}{2} \sum_{n,n'=1}^{N^x} \int_{\mathbb{R}^{2N^x}} f_{M_n M_{n'}}(t, M_t(x, x')) \Sigma^{ag,*}_{nn'}(t, x, x') m_t(x) m_t(x') dx dx' \Bigg\} dt \\
&+ \sum_{n=1}^{N^x} \left\{ \int_{\mathbb{R}^{N^x}} f_{M_n}(t, M_t(x)) \sigma^{ag}_n \left(y^{i,*}(t, x, a_t, r_t), x, a_t, r_t\right) m_t(x) dx \right\} dB^{ag}_t
\end{aligned} \tag{18}
$$

with $f_{Mx}(t, M_t(x)) \in \mathbb{R}^{N^x \times N^x}$ and the optimally controlled covariance function

$$
\Sigma^{o,*}(t, x, x') = \sigma^{ag} \left(y^{i,*}(t, x, a_t, r_t), x, a_t, r_t\right) \sigma^{ag} \left(y^{i,*}(t, x', a_t, r_t), x', a_t, r_t\right)^T \tag{19}
$$

with $o \in \{id, ag\}$. Note that the term containing f_{Mx} results from combining the first-order term in the Taylor expansion in (17) with the second term of the marginal c.d.f. dynamics in (11) and applying integration by parts once more. This term measures the cross-variation stemming from changes to the distribution and changes to the realization of the current draw of this distribution.

Lastly, let us give some practical examples of the derivations of this section. The following example illustrates how the distributional derivatives can be computed for a concrete example function.

Example 8 (Distributional Derivatives of a Production Function)**.** *Consider the model in Krusell and Smith (1998) where we have $x^i_t = [l^i_t, k^i_t]$, i.e, $N^x = 2$. Denote the Cobb-Douglas production function with elasticity $\alpha \in (0, 1)$ by $F(L_t, K_t) = L_t^{1-\alpha} K_t^{\alpha}$. Define the aggregates by*

$$
L_t = \int_{\mathbb{R}^2} l m_t(l, k) dl dk = \int_{\mathbb{R}^2} l dM_t(l, k)
$$

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

and

$$K_t = \int_{\mathbb{R}^2} k m_t(l,k) dl dk = \int_{\mathbb{R}^2} k dM_t(l,k).$$

Note that they can be written as functions of either m_t *or* M_t*, i.e.,* $\tilde{f}(t,m_t) = f(t,M_t) = F(L_t,K_t)$*. Then, the derivative of the production function reads*

$$\lim_{s\to 0} \frac{\tilde{f}(t,m_t + s(m_t' - m_t)) - \tilde{f}(t,m_t)}{s} = \int_{\mathbb{R}^2} \underbrace{\begin{bmatrix} F_L(L_t,K_t)l \\ F_K(L_t,K_t)k \end{bmatrix}^T}_{=\tilde{f}_m(t,m_t(l,k))^T} (\mathbf{m}_t'(l,k) - \mathbf{m}_t(l,k)) dl dk$$

such that

$$f_M(t,M_t(l,k)) = \begin{bmatrix} F_L(L_t,K_t) \\ F_K(L_t,K_t) \end{bmatrix}.$$

It follows that $f_{Mx}(t,M_t(k,l)) = 0$ *and*

$$\tilde{f}_{mx}(t,m_t(l,k)) = \begin{bmatrix} F_L(L_t,K_t) & 0 \\ 0 & F_K(L_t,K_t) \end{bmatrix}.$$

The second-order derivatives are given by

$$\tilde{f}_{mm}(t,m_t(l,k,l',k')) = \begin{bmatrix} F_{LL}(L_t,K_t)ll' & F_{LK}(L_t,K_t)lk' \\ F_{KL}(L_t,K_t)kl' & F_{KK}(L_t,K_t)kk' \end{bmatrix}$$

and

$$f_{MM}(t,M_t(l,k,l',k')) = \begin{bmatrix} F_{LL}(L_t,K_t) & F_{LK}(L_t,K_t) \\ F_{KL}(L_t,K_t) & F_{KK}(L_t,K_t) \end{bmatrix}.$$

Existing work on the master equation in economics utilizes $\tilde{f}_m$ (Schaab, 2020; Bilal, 2023; Gu et al., 2024). Contrary to that, we work with f_M and equation (18). Both are equivalent and can be easily translated. We illustrate this translation with the model with indirect aggregate risk featured in Bilal (2023).

Example 9 (Case Without Direct Aggregate Risk in Bilal (2023))**.** *In this model,* $N^x = d^{ag} = 1, \zeta > 0$ *and* $\sigma^{ag} = 0$*. In the case without direct aggregate risk (*$\sigma^{ag} = 0$*),*

the stochastic integral in (18) is obsolete such that

$$
\begin{aligned}
& df(t, M_t) \\
= & \left\{ f_t(t, M_t) + \int_{\mathbb{R}} f_M(t, M_t(x)) \mu\left(t, y^{i,*}(t, x, a_t, r_t), x, a_t, r_t\right) m_t(x) dx \right. \\
& \left. + \frac{1}{2} \int_{\mathbb{R}} f_{Mx}(t, M_t(x)) \Sigma^{id,*}(t, x, x) m_t(x) dx \right\} dt \\
= & \left\{ \tilde{f}_t(t, m_t(x)) + \int_{\mathbb{R}} \frac{d}{dx} \tilde{f}_m(t, m_t(x)) \mu\left(t, y^{i,*}(t, x, a_t, r_t), x, a_t, r_t\right) m_t(x) dx \right. \\
& \left. + \frac{1}{2} \int_{\mathbb{R}} \frac{d^2}{dx^2} \tilde{f}_m(t, m_t(x)) \Sigma^{id,*}(t, x, x) m_t(x) dx \right\} dt \\
= & \left\{ \tilde{f}_t(t, m_t(x)) + \int_{\mathbb{R}} \tilde{f}_m(t, m_t(x)) \underbrace{\frac{d}{dx}\left[-\mu\left(t, y^{i,*}(t, x, a_t, r_t), x, a_t, r_t\right) m_t(x)\right]}_{=:\alpha_t^m(x)} dx \right. \\
& \left. + \frac{1}{2} \int_{\mathbb{R}} \tilde{f}_m(t, m_t(x)) \underbrace{\frac{d^2}{dx^2}\left[\Sigma^{id,*}(t, x, x) m_t(x)\right]}_{=:\varsigma_t^m(x)} dx \right\} dt
\end{aligned}
$$

where the second equality follows from Definition 7 and the third equality results from integration by parts. It follows that

$$
\frac{d}{dt} \tilde{f}_t(t, m_t(x)) = \tilde{f}_t(t, m_t(x)) + \int_{\mathbb{R}} \tilde{f}_m(t, m_t(x)) \left[\alpha_t^m(x) + \frac{1}{2} \varsigma_t^m(x)\right] dx
$$

where the Kolmogorov forward equation reads $\frac{d}{dt} m_t(x) = \alpha_t^m(x) + \frac{1}{2} \varsigma_t^m(x)$. *Bilal (2023) uses this equation and sets* $\tilde{f}$ *equal to the value function to derive the master equation.*

Apart from illustrating their equivalence, this example demonstrates the advantages of working with f_M rather than $\tilde{f}_m$ to derive the master equation. The former is simply multiplied by the drift and covariance terms of the state dynamics, whereas, the latter requires the computation of derivatives which can be numerically cumbersome.

3.2 Derivation of the Master Equation

The master equation results from rewriting the HJB equation in (13) in terms of $\mathcal{J}$. Akin to the HJB operator in (6), we define the master operator to subtract the time-derivative of the value function in (5), which equals $\rho\mathcal{J}(t, M_t, x_t^i, a_t) - u(y_t^i, x_t^i)$, from the drift of the stochastic process $\mathcal{J}(t, M_t, x_t^i, a_t)$. We have seen in the previous subsection that Itô's formula does not apply to functions of the random field $\{M_t(x)\}_{t\geq 0, x\in\mathbb{R}^{N^x}}$ but appropriate extensions have been derived in Carmona and Delarue (2018b) and Cardaliaguet et al. (2019). Since $\mathcal{J}(t, M_t, x_t^i, a_t)$ is a function of time, the distribution with marginals following (11) and the states as in (2), the expansion in equation (18) needs to be extended to[9]

$$d\mathcal{J}\left(t, M_t, x_t^i, a_t\right) = d\mathcal{A}_{tX}\left(t, M_t, x_t^i, a_t\right) + d\mathcal{A}_{M}\left(t, M_t, x_t^i, a_t\right) + d\mathcal{A}_{MX}\left(t, M_t, x_t^i, a_t\right).$$

The first term summarizes the variation stemming from t and the state processes $\{x_t^i, a_t\}_{t\geq 0}$. It follows from the standard Itô formula, i.e.,

$$\begin{aligned}
&d\mathcal{A}_{tX}\left(t, M_t, x_t^i, a_t\right) = \\
&\Bigg\{\mathcal{J}_t\left(t, M_t, x_t^i, a_t\right) + \sum_{n=1}^{N^x}\mathcal{J}_{x_n^i}\left(t, M_t, x_t^i, a_t\right)\mu_n\left(t, y_t^{i,*}, x_t^i, a_t, r_t^*\right) \\
&+\frac{1}{2}\sum_{n,n'=1}^{N^x}\sum_{o\in\{id,ag\}}\mathcal{J}_{x_n^i x_{n'}^i}\left(t, M_t, x_t^i, a_t\right)\Sigma_{nn'}^{o,*}\left(t, x_t^i, x_t^i\right) \\
&+\sum_{n=1}^{N^a}\mathcal{J}_{a_n}\left(t, M_t, x_t^i, a_t\right)\eta_n\left(a_t, r_t^*\right) + \frac{1}{2}\sum_{n,n'=1}^{N^a}\mathcal{J}_{a_n a_{n'}}\left(t, M_t, x_t^i, a_t\right)\Sigma_{nn'}^{a}\left(a_t, r_t^*\right) \\
&+\frac{1}{2}\sum_{n=1}^{N^x}\sum_{n'=1}^{N^a}\mathcal{J}_{x_n^i a_{n'}}\left(t, M_t, x_t^i, a_t\right)\Sigma_{nn'}^{xa}\left(y_t^{i,*}, x_t^i, a_t, r_t^*\right)\Bigg\}dt \\
&+\left\{\sum_{n=1}^{N^x}\mathcal{J}_{x_n^i}\left(t, M_t, x_t^i, a_t\right)\sigma_n^{id}\left(y_t^{i,*}, x_t^i, a_t, r_t^*\right)\right\}dB_t^i \\
&+\left\{\sum_{n=1}^{N^x}\mathcal{J}_{x_n^i}\left(t, M_t, x_t^i, a_t\right)\sigma_n^{ag}\left(y_t^{i,*}, x_t^i, a_t, r_t^*\right) + \sum_{n=1}^{N^a}\mathcal{J}_{a_n}\left(t, M_t, x_t^i, a_t\right)\zeta_n\left(a_t, r_t^*\right)\right\}dB_t^{ag}
\end{aligned}$$

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

with the covariance matrices $\Sigma^{\cdot}$ as in (4) and $\Sigma^{\cdot,*}$ as in (19). The other two terms, on the other hand, account for the dependence on the random field $\{M_t(x)\}_{t\geq 0, x\in\mathbb{R}^{N^x}}$. Thus, $d\mathcal{A}_M\left(t, M_t, x_t^i, a_t\right)$ results from the expansion in (18) such that

$$
\begin{aligned}
&d\mathcal{A}_M\left(t, M_t, x_t^i, a_t\right) = \\
&\Bigg\{\sum_{n=1}^{N^x}\int_{\mathbb{R}^{N^x}} \mathcal{J}_{M_n}\left(t, M_t(x), x_t^i, a_t\right)\mu_n\left(t, y_t^{i,*}, x, a_t, r_t^*\right) m_t(x)dx \\
&+\frac{1}{2}\sum_{n,n'=1}^{N^x}\int_{\mathbb{R}^{2N^x}} \mathcal{J}_{M_nM_{n'}}\left(t, M_t(x,x'), x_t^i, a_t\right)\Sigma_{nn'}^{ag,*}(t,x,x')m_t(x)m_t(x')dxdx' \\
&+\frac{1}{2}\sum_{n,n'=1}^{N^x}\sum_{o\in\{id,ag\}}\int_{\mathbb{R}^{N^x}} \mathcal{J}_{M_nx_{n'}}\left(t, M_t(x), x_t^i, a_t\right)\Sigma_{nn'}^{o,*}(t,x,x)m_t(x)dx\Bigg\}dt \\
&+\sum_{n=1}^{N^x}\left\{\int_{\mathbb{R}^{N^x}} \mathcal{J}_{M_n}\left(t, M_t(x), x_t^i, a_t\right)\sigma_n^{ag}\left(y_t^{i,*}, x, a_t, r_t^*\right) m_t(x)dx\right\}dB_t^{ag},
\end{aligned}
$$

whereas, the last term $d\mathcal{A}_{MX}\left(t, M_t, x_t^i, a_t\right)$ 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, i.e,

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

Subtract the time-derivative of the value function in (5) and collecting the drift terms of this expansion yields the master operator. We define it by $\mathcal{H}^M[\mathcal{J}, y^i, r](t, M, x^i, a)$

mapping the functions $[\mathcal{J}, y^i, r]$ into a function of (t, M, x^i, a) such that

$$
\begin{aligned}
&\mathcal{H}^M[\mathcal{J}, y^i, r](t, M, x^i, a) = \qquad (20)\\
&-\rho\mathcal{J} + u\left(y^i, x^i\right) + \mathcal{J}_t + \sum_{n=1}^{N^x} \int_{\mathbb{R}^{N^x}} \mathcal{J}_{M_n}\left(t, M(x), x^i, a\right) \mu_n\left(t, y^i, x, a, r\right) m(x)dx\\
&+ \frac{1}{2}\sum_{n,n'=1}^{N^x} \int_{\mathbb{R}^{2N^x}} \mathcal{J}_{M_n M_{n'}}\left(t, M(x, x'), x^i, a\right) \Sigma^{ag,*}_{nn'}(t, x, x')m(x)m(x')dxdx'\\
&+ \frac{1}{2}\sum_{n,n'=1}^{N^x} \sum_{o\in\{id,ag\}} \int_{\mathbb{R}^{N^x}} \mathcal{J}_{M_n x_{n'}}\left(t, M(x), x^i, a\right) \Sigma^{o,*}_{nn'}(t, x, x)m(x)dx\\
&+ \frac{1}{2}\sum_{n=1}^{N^x}\sum_{n'=1}^{N^a} \int_{\mathbb{R}^{N^x}} \mathcal{J}_{M_n a_{n'}}\left(t, M_t(x), x^i, a\right) \Sigma^{xa}_{nn'}\left(y^i, x, a, r\right) m(x)dx\\
&+ \frac{1}{2}\sum_{n,n'=1}^{N^x} \int_{\mathbb{R}^{N^x}} \mathcal{J}_{x^i_n M_{n'}}\left(t, M(x), x^i, a\right) \Sigma^{ag,*}_{nn'}\left(t, x^i, x\right) m(x)dx\\
&+ \sum_{n=1}^{N^a} \mathcal{J}_{a_n}\eta_n\left(a, r\right) + \frac{1}{2}\sum_{n,n'=1}^{N^a} \mathcal{J}_{a_n a_{n'}}\Sigma^{a}_{nn'}\left(a, r\right) + \sum_{n=1}^{N^x} \mathcal{J}_{x^i_n}\mu_n\left(t, y^i, x^i, a, r\right)\\
&+ \frac{1}{2}\sum_{n,n'=1}^{N^x} \sum_{o\in\{id,ag\}} \mathcal{J}_{x^i_n x^i_{n'}}\Sigma^{o,*}_{nn'}\left(t, x^i, x^i\right) + \frac{1}{2}\sum_{n=1}^{N^x}\sum_{n'=1}^{N^a} \mathcal{J}_{x^i_n a_{n'}}\Sigma^{xa}_{nn'}\left(y^i, x^i, a, r\right)
\end{aligned}
$$

with the covariance matrices $\Sigma^{\cdot}$ as in (4) and $\Sigma^{\cdot,*}$ as in (19). When inserting the optimal controls $y^{i,*}$ that solve the first-order conditions in (12), the master equation sets the master operator to zero. Note that $\mathcal{J}_o = J_o$, $o \in \{x^i, a, x^i x^i, aa, x^i a\}$, such that the first two terms and the last two lines of the master operator equal the terms in the HJB operator in (6). The additional terms are due to introducing the consistent law of motion of the distribution into the equation, i.e, splitting the time derivative J_t into the variation stemming from the exogenous Brownian innovations given by $\mathcal{J}_t$ and the variation stemming from changes in the distribution summarized by the distributional derivatives. Note that in models solely featuring *indirect* aggregate risk ($\zeta > 0, \sigma^{ag} = 0$) as in, e.g., Bilal (2023); Gu et al. (2024), it holds that $\mathcal{J}_t = 0$ because exogenous aggregate Brownian innovations only affect the agents-specific state indirectly through the aggregate state. In our more general setting that includes *direct* aggregate risk ($\sigma^{ag} > 0$), however, the time derivative $\mathcal{J}_t$ is non-negligible as the agent-specific state variable is directly affected by exogenous Brownian innovations.

Overall, the master equation allows us to characterize the equilibrium as follows.

Proposition 10 (Characterization of the Recursive Equilibrium)**.** *For any trajectory of consistent cross-sectional distributions* $\{M_t\}_{t\geq 0}$*, the recursive equilibrium of Definition 6 consists of a set of controls* $\{y_t^{i,*}\}_{t\geq 0, i\in[0,1]}$ *and returns* $\{r_t^*\}_{t\geq 0}$ *such that the following conditions hold.*

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

$$0 = \frac{\partial}{\partial y_n^i}\mathcal{H}^{HJB}[\mathcal{J}, y^{i,*}, r^*](t, M, x^i, a),\ n \in \{1, \dots, N^y\}, \tag{21}$$

$$0 = \mathcal{H}^{M}[\mathcal{J}, y^{i,*}, r^*](t, M, x^i, a) \tag{22}$$

$$0 = \mathcal{H}^{TBC}[\mathcal{J}](M, x^i, a) \tag{23}$$

$$0 = \mathcal{H}^{SBC}[\mathcal{J}, y^{i,*}, r^*](t, M, \tilde{x}^{i,uc}, a), \tag{24}$$

with $\mathcal{H}^{HJB}$ *as in (6),* $\mathcal{H}^{M}$ *as in (20),* $\mathcal{H}^{TBC}$ *as in (7) and* $\mathcal{H}^{SBC}$ *as in (8).*

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

$$0 = \mathcal{H}^{MC}[r^*, y^{i,*}](t, M, a). \tag{25}$$

4 Approximating the Equilibrium Conditions

The master equation in (22) is a non-standard PDE due to the distributional derivatives. Thus, to solve the equilibrium conditions (21)-(25), we first propose an approximation of the distributional derivatives. Importantly, this approximation allows us to approximate the master equation with a standard PDE.

4.1 Approximating the Distributional Derivatives

It is useful to note that the distribution function $M_t(.) = \mathbb{P}(x_t^i \leq .) = p[x_t^i](.)$ is a function of the random variable x_t^i. In this context, however, x_t^i does not denote a specific outcome for one particular agent i but rather all possible outcomes that agent i could have at time t. To emphasize this difference between realized and hypothetical state, we keep x_t^i to denote the realization of agent i's state variable at time t and

denote the collection of possible outcomes at time t by χ_t^i. Thus, we substitute $M_t(.) = p[\chi_t^i](.)$. Therefore, the value function $\mathcal{J}(t, M_t, x_t^i, a_t) = \mathcal{J}(t, p[\chi_t^i], x_t^i, a_t) = \tilde{\mathcal{J}}(t, \chi_t^i, x_t^i, a_t)$ can be rewritten in terms of the state random variable χ_t^i.[10] As an Itô process, the agent-specific state in (3) is a square-integrable stochastic process implying that χ_t^i for a fixed time point t is a square-integrable random variable, i.e., $\chi_t^i \in L^2(\Omega, \mathcal{F}_t, \mathbb{P}; \mathbb{R}^{N^x})$. Since the state random variable lives on a Hilbert space, the derivative of a function of that random variable $\tilde{f}(t, \chi_t^i)$ is given by the standard Fréchet derivative $\tilde{f}_{\chi^i}$.

Definition 11 (Distributional Derivatives II (see Carmona and Delarue, 2018a, Definition 5.22))**.** *The derivative w.r.t. a square-integrable random variable is defined by the Fréchet derivative $\tilde{f}_{\chi^i}$ given by*[11]

$$\lim_{s \to 0} \frac{\tilde{f}(t, \chi_t^i + s(\chi_t^{i'} - \chi_t^i)) - \tilde{f}(t, \chi_t^i)}{s} = \int_{\mathbb{R}^{N^x}} \tilde{f}_{\chi^i}(t, \chi_t^i, x)^T (x' - x) d\mathbb{P}\left(\{\chi_t^i \leq x\} \cap \{\chi_t^{i'} \leq x'\}\right)$$

for any $\chi_t^i, \chi_t^{i'} \in L^2(\Omega, \mathcal{F}_t, \mathbb{P}; \mathbb{R}^{N^x})$. The second-order derivative is similarly defined by

$$\lim_{s \to 0} \frac{\tilde{f}_{\chi^i}(t, \chi_t^i + s(\chi_t^{i'} - \chi_t^i), x) - \tilde{f}_{\chi^i}(t, \chi_t^i, x)}{s} = \int_{\mathbb{R}^{N^x}} \tilde{f}_{\chi^i \chi^i}(t, \chi_t^i, x, x')(x'' - x') d\mathbb{P}\left(\{\chi_t^i \leq x'\} \cap \{\chi_t^{i'} \leq x''\}\right).$$

As for the other distributional derivatives, we have that $\tilde{f}_{\chi^i}(t, \chi_t^i, x) \in \mathbb{R}^{N^x}$ and $\tilde{f}_{\chi^i \chi^i}(t, \chi_t^i, x, x') \in \mathbb{R}^{N^x \times N^x}$. Moreover, there is a direct link between derivatives w.r.t. the c.d.f. and derivatives w.r.t. the corresponding random variable.

Proposition 12 (Cardaliaguet et al. (2019, Proposition A.1.2))**.** *If a function of the distribution $f(t, M_t) = \tilde{f}(t, \chi_t^i)$ is differentiable w.r.t. the c.d.f., then it is differentiable w.r.t. the corresponding random variable and $f_M(t, M_t(x)) = \tilde{f}_{\chi^i}(t, \chi_t^i, x)$. Similarly, the second-order derivatives agree, i.e., $f_{MM}(t, M_t(x, x')) = \tilde{f}_{\chi^i \chi^i}(t, \chi_t^i, x, x')$.*

[10] The fact that functions of distributions can be rewritten as functions of random variables has also been exploited in Proehl (2024) in a discrete-time context. As that paper provides a detailed account of this correspondence, we only give a high-level summary here.

[11] The equation to define $\tilde{f}_{\chi^i}$ is equivalent to Cardaliaguet et al. (2019, Equation (A.5)).

The advantage of working with $\tilde{f}(t, \chi_t^i)$ and its derivatives lies in the fact that one can approximate the random variable χ_t^i with polynomial chaos, a stochastic variant of orthogonal polynomial projection. Importantly, this approximation converges when the order of the polynomial approximation increases. This discretization technique has first been proposed for heterogeneous agent models by Proehl (2019) in a discrete-time framework. At a fixed time point t, any element of the multivariate state random variable $(\chi_t^i)_n \in L^2(\Omega, \Sigma, \mathbb{P}; \mathbb{R})$ can be written as

$$(\chi_t^i)_n = \sum_{v=0}^{\infty} \varphi_{t,n}^v \phi^v(\xi) \tag{26}$$

where $\xi \in L^\infty(\Omega, \Sigma, \mathbb{P}; \mathbb{R}^{N^\xi})$ is an N^ξ-dimensional vector of independent random variables for which all moments are finite. The corresponding family of orthogonal polynomials $\{\phi^v\}_{v=0}^\infty$ satisfies

$$\mathbb{1}_{\{v=u\}} b^v = \int_{\mathbb{R}^{N^\xi}} \phi^v(\xi)\phi^u(\xi) m^\xi(\xi) d\xi \tag{27}$$

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

$$\varphi_{t,n}^v = \frac{1}{b^v} \int_{\mathbb{R}} \int_{\mathbb{R}^{N^\xi}} x_n \phi^v(e) d\mathbb{P}\left(\{(\chi_t^i)_n \leq x_n\} \cap \{\xi \leq e\}\right). \tag{28}$$

The seed random variables ξ in the Polynomial Chaos representation (26) are fixed exogenously. There is a trade-off to be considered for that choice. On the one hand, one can choose ξ to equal a standard random variable, e.g., normal or uniform. In that case, the truncation error in (26) will become small only for higher-order polynomials. However, the family of polynomials is then standard and does not need to be constructed. On the other hand, one can tailor ξ to a distribution close to the one expected in the model. In that case, truncating (26) at lower-order will already achieve a good approximation but the family of polynomials needs to be constructed separately. The latter is what is done in Proehl (2019) as this paper relies on a tensor-grid discretization of the state space. In contrast, we choose the former because our numerical method relies on neural nets and therefore, suffers less from the curse of dimensionality. In (26), we fix $\xi \sim \mathcal{N}(0, 1)$. Furthermore, we assume this standard univariate distribution to be i.i.d. for all states $n \in \{1, \ldots, N^x\}$. The corresponding family of orthogonal polynomials are the Hermite polynomials (see e.g.

Ernst et al., 2012; Xiu and Karniadakis, 2002). This implies that $b^v = \sqrt{2\pi}\, v!$ in (27) and $\phi^0(\xi) = 1$ such that $\varphi^0_{t,n}$ in (28) yields the mean of the state random variable $(\chi^i_t)_n$, i.e., $\varphi^0_{t,n} = \mathbb{E}[(\chi^i_t)_n]$.

Using this polynomial projection technique, the multivariate random variable χ^i_t can be represented by a set of projection coefficient series $(\{\varphi^v_{1,t}\}^\infty_{v=0}, \ldots, \{\varphi^v_{N^x,t}\}^\infty_{v=0})$. Therefore, we can rewrite any function $\tilde{f}(t, \chi^i_t) = \hat{f}(t, \{\varphi^v_{1,t}\}^\infty_{v=0}, \ldots, \{\varphi^v_{N^x,t}\}^\infty_{v=0})$. Given (28), each projection coefficient in $\{\varphi^v_{n,t}\}^\infty_{v=0}$ is a function of the corresponding element $(\chi^i_t)_n$ of the multivariate state random variable. Hence, the distributional derivative can be computed using the chain rule.

Proposition 13. *When the multivariate state random variable is represented by Polynomial Chaos as in (26), the distributional derivative of a function defined as* $\tilde{f}(t, \chi^i_t) = \hat{f}(t, \{\varphi^v_{1,t}\}^\infty_{v=0}, \ldots, \{\varphi^v_{N^x,t}\}^\infty_{v=0})$ *is given by*

$$\tilde{f}_{\chi^i_n}(t, \chi^i_t, x(e)) = \sum_{v=0}^{\infty} \frac{\partial}{\partial \varphi^v_n} \hat{f}(t, \{\varphi^v_{1,t}\}^\infty_{v=0}, \ldots, \{\varphi^v_{N^x,t}\}^\infty_{v=0}) \frac{\phi^v(e)}{b^v},$$

where b^v *as in (27) and* $x_n(e) = \sum^\infty_{v=0} \varphi^v_{t,n} \phi^v(e)$. *The second-order derivatives are given by*

$$\tilde{f}_{\chi^i_n \chi^i_{n'}}(t, \chi^i_t, x(e), x(e')) = \sum_{v,v'=0}^{\infty} \frac{\partial^2}{\partial \varphi^v_n \partial \varphi^{v'}_{n'}} \hat{f}(t, \{\varphi^v_{1,t}\}^\infty_{v=0}, \ldots, \{\varphi^v_{N^x,t}\}^\infty_{v=0}) \frac{\phi^v(e)}{b^v} \frac{\phi^{v'}(e')}{b^{v'}}$$

$$\tilde{f}_{\chi^i_n x_{n'}}(t, \chi^i_t, x(e)) = \sum_{v=0}^{\infty} \frac{\frac{\partial}{\partial \varphi^v_n} \hat{f}(t, \{\varphi^v_{1,t}\}^\infty_{v=0}, \ldots, \{\varphi^v_{N^x,t}\}^\infty_{v=0}) \frac{\partial}{\partial e} \phi^v(e)}{b^v \sum^\infty_{v'=0} \varphi^{v'}_{n',t} \frac{\partial}{\partial e} \phi^{v'}(e)}.$$

Remark. *The distributional derivative* $\tilde{f}_{\chi^i_n}$ *is a random variable itself with projection coefficients* $\{\hat{f}_{\varphi^v_n}/b^v\}^\infty_{v=0}$.

Numerically, the series in (26) will be truncated at V_n meaning that we set $\{\varphi^v_{t,n}\}_{v>V_n} = 0$. Proehl (2019) 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. This convergence result likely carries over to the framework at hand which we will check numerically for our example model.

4.2 Approximated Equilibrium Conditions

Overall, we can rewrite the value function by

$$\begin{aligned}\mathcal{J}\left(t, M_t, x_t^i, a_t\right) =& \tilde{\mathcal{J}}\left(t, \chi_t^i, x_t^i, a_t\right) = \hat{\mathcal{J}}\left(t, \{\varphi_{1,t}^v\}_{v=0}^{\infty}, \ldots, \{\varphi_{N^x,t}^v\}_{v=0}^{\infty}, x_t^i, a_t\right) \\ \approx& \hat{\mathcal{J}}\left(t, \{\varphi_{1,t}^v\}_{v=0}^{V_1}, \ldots, \{\varphi_{N^x,t}^v\}_{v=0}^{V_{N^x}}, x_t^i, a_t\right).\end{aligned}$$

Note that the derivatives of $\mathcal{J}$ and $\hat{\mathcal{J}}$ w.r.t. $\{t, x^i, a, x^i x^i, aa, x^i a\}$ coincide. Combining Propositions 12 and 13 allows us to approximate the master operator in (20) with a standard PDE for the function $\hat{\mathcal{J}}(t, \{\varphi_1^v\}_{v=0}^{V_1}, \ldots, \{\varphi_{N^x}^v\}_{v=0}^{V_{N^x}}, x^i, a)$. Define the approximated master operator $\mathcal{H}^{aM}[\hat{\mathcal{J}}, y^i, r](t, \{\varphi_1^v\}_{v=0}^{V_1}, \ldots, \{\varphi_{N^x}^v\}_{v=0}^{V_{N^x}}, x^i, a)$ mapping functions $[\hat{\mathcal{J}}, y^i, r]$ into a function of $(t, \{\varphi_1^v\}_{v=0}^{V_1}, \ldots, \{\varphi_{N^x}^v\}_{v=0}^{V_{N^x}}, x^i, a)$ by

$$\begin{aligned}&\mathcal{H}^{aM}[\hat{\mathcal{J}}, y^i, r](t, \{\varphi_1^v\}_{v=0}^{V_1}, \ldots, \{\varphi_{N^x}^v\}_{v=0}^{V_{N^x}}, x^i, a) = \qquad (29)\\
&-\rho\hat{\mathcal{J}} + u\left(y^i, x^i\right) + \hat{\mathcal{J}}_t + \sum_{n=1}^{N^x}\sum_{v=0}^{V_n}\frac{\hat{\mathcal{J}}_{\varphi_n^v}}{b^v}\int_{\mathbb{R}^{N^\xi}}\mu_n\left(t, y^i, x(e), a, r\right)\phi^v(e)m^\xi(e)de\\
&+\frac{1}{2}\sum_{n,n'=1}^{N^x}\sum_{v=0}^{V_n}\sum_{v'=0}^{V_{n'}}\frac{\hat{\mathcal{J}}_{\varphi_n^v\varphi_{n'}^{v'}}}{b^v b^{v'}}\int_{\mathbb{R}^{2N^\xi}}\Sigma_{nn'}^{ag,*}(t, x(e), x(e'))\phi^v(e)m^\xi(e)de\phi^{v'}(e')m^\xi(e')de'\\
&+\frac{1}{2}\sum_{n,n'=1}^{N^x}\sum_{o\in\{id,ag\}}\sum_{v=0}^{V_n}\frac{\hat{\mathcal{J}}_{\varphi_n^v}}{b^v}\int_{\mathbb{R}^{N^\xi}}\Sigma_{nn'}^{o,*}(t, x(e), x(e))\frac{\frac{\partial}{\partial e}\phi^v(e)}{\sum_{v'=0}^{V_{n'}}\varphi_{n'}^{v'}\frac{\partial}{\partial e}\phi^{v'}(e)}m^\xi(e)de\\
&+\frac{1}{2}\sum_{n=1}^{N^x}\sum_{n'=1}^{N^a}\sum_{v=0}^{V_n}\frac{\hat{\mathcal{J}}_{\varphi_n^v a_{n'}}}{b^v}\int_{\mathbb{R}^{N^\xi}}\Sigma_{nn'}^{xa}\left(y^i, x(e), a, r\right)\phi^v(e)m^\xi(e)de\\
&+\frac{1}{2}\sum_{n,n'=1}^{N^x}\sum_{v=0}^{V_n}\frac{\hat{\mathcal{J}}_{\varphi_n^v x_{n'}^i}}{b^v}\int_{\mathbb{R}^{N^\xi}}\Sigma_{nn'}^{ag,*}\left(t, x^i, x(e)\right)\phi^v(e)m^\xi(e)de\\
&+\sum_{n=1}^{N^a}\hat{\mathcal{J}}_{a_n}\eta_n(a, r) + \frac{1}{2}\sum_{n,n'=1}^{N^a}\hat{\mathcal{J}}_{a_n a_{n'}}\Sigma_{nn'}^{a}(a, r) + \sum_{n=1}^{N^x}\hat{\mathcal{J}}_{x_n^i}\mu_n\left(t, y^i, x^i, a, r\right)\\
&+\frac{1}{2}\sum_{n,n'=1}^{N^x}\sum_{o\in\{id,ag\}}\hat{\mathcal{J}}_{x_n^i x_{n'}^i}\Sigma_{nn'}^{o,*}\left(t, x^i, x^i\right) + \frac{1}{2}\sum_{n=1}^{N^x}\sum_{n'=1}^{N^a}\hat{\mathcal{J}}_{x_n^i a_{n'}}\Sigma_{nn'}^{xa}\left(y^i, x^i, a, r\right)\end{aligned}$$

where b^v as in (27), $x_n(e) = \sum_{v=0}^{V_n}\varphi_n^v\phi^v(e)$, the covariance matrices $\Sigma^{\cdot}$ as in (4) and $\Sigma^{\cdot,*}$ as in (19). We can now define the approximated equilibrium.

Definition 14 (Approximated Recursive Equilibrium)**.** *Suppose that the agent-specific*

state random variable is approximated by polynomial chaos as in (26). Then, for any trajectory of consistent cross-sectional distributions $\{M_t\}_{t\geq 0}$*, the approximated recursive equilibrium consists of a set of controls* $\{y_t^{i,*}\}_{t\geq 0, i\in[0,1]}$ *and returns* $\{r_t^*\}_{t\geq 0}$ *such that the following conditions hold.*

(*i*) *Given the returns* $\{r_t^*\}_{t\geq 0}$*, households optimize their approximated value function* $\hat{\mathcal{J}}(t, \{\varphi_1^v\}_{v=0}^{V_1}, \ldots, \{\varphi_{N^x}^v\}_{v=0}^{V_{N^x}}, x^i, a)$ *by choosing their optimal control functions* $y_t^{i,*} = y^{i,*}(t, \{\varphi_1^v\}_{v=0}^{V_1}, \ldots, \{\varphi_{N^x}^v\}_{v=0}^{V_{N^x}}, x^i, a)$ *such that for any* $n \in \{1, \ldots, N^y\}$

$$0 = \frac{\partial}{\partial y_n^i}\mathcal{H}^{HJB}[\hat{\mathcal{J}}, y^{i,*}, r^*]\left(t, \{\varphi_1^v\}_{v=0}^{V_1}, \ldots, \{\varphi_{N^x}^v\}_{v=0}^{V_{N^x}}, x^i, a\right) \tag{30}$$

$$0 = \mathcal{H}^{aM}[\hat{\mathcal{J}}, y^{i,*}, r^*]\left(t, \{\varphi_1^v\}_{v=0}^{V_1}, \ldots, \{\varphi_{N^x}^v\}_{v=0}^{V_{N^x}}, x^i, a\right) \tag{31}$$

$$0 = \mathcal{H}^{TBC}[\hat{\mathcal{J}}]\left(\{\varphi_1^v\}_{v=0}^{V_1}, \ldots, \{\varphi_{N^x}^v\}_{v=0}^{V_{N^x}}, x^i, a\right) \tag{32}$$

$$0 = \mathcal{H}^{SBC}[\hat{\mathcal{J}}, y^{i,*}, r^*]\left(t, \{\varphi_1^v\}_{v=0}^{V_1}, \ldots, \{\varphi_{N^x}^v\}_{v=0}^{V_{N^x}}, \tilde{x}^{i,uc}, a\right), \tag{33}$$

where $\mathcal{H}^{HJB}$ *as in (6),* $\mathcal{H}^{aM}$ *as in (29),* $\mathcal{H}^{TBC}$ *as in (7) and* $\mathcal{H}^{SBC}$ *as in (8).*

(*ii*) *Given the controls* $\{y_t^{i,*}\}_{t\geq 0, i\in[0,1]}$*, the returns that clear markets are given by functions* $r_t^* = r^*(t, \{\varphi_1^v\}_{v=0}^{V_1}, \ldots, \{\varphi_{N^x}^v\}_{v=0}^{V_{N^x}}, a)$ *such that*

$$0 = \mathcal{H}^{MC}[r^*, y^{i,*}]\left(t, \{\varphi_1^v\}_{v=0}^{V_1}, \ldots, \{\varphi_{N^x}^v\}_{v=0}^{V_{N^x}}, a\right). \tag{34}$$

Note that this set of equations simplifies for some example models. For instance, whenever (30) and (34) yield closed-form solutions for y^* and r^*, respectively, we can insert this directly into the approximated master equation and the boundary equations.

Example 15. *The returns in Krusell and Smith (1998) are explicitly given by*

$$r_t = \begin{bmatrix} \frac{\partial}{\partial L}F\left(z_t, L_t, K_t\right) \\ \frac{\partial}{\partial K}F\left(z_t, L_t, K_t\right) \end{bmatrix} = \begin{bmatrix} \frac{\partial}{\partial L}F\left(z_t, \mathbb{E}\left[l_t^i\right], \mathbb{E}\left[k_t^i\right]\right) \\ \frac{\partial}{\partial K}F\left(z_t, \mathbb{E}\left[l_t^i\right], \mathbb{E}\left[k_t^i\right]\right) \end{bmatrix} = \begin{bmatrix} \frac{\partial}{\partial L}F\left(z_t, \varphi_{t,l}^0, \varphi_{t,k}^0\right) \\ \frac{\partial}{\partial K}F\left(z_t, \varphi_{t,l}^0, \varphi_{t,k}^0\right) \end{bmatrix}$$

such that the market clearing condition (34) can be inserted directly into equations (30)-(33). In contrast, the bond return in Huggett (1993) is implicitly given by the market clearing condition

$$0 = \mathbb{E}\left[b_t^i\right] = \mathbb{E}\left[\varphi_{t,b}^0\right]$$

and thus, cannot be inserted.

5 Solving for the Approximated Equilibrium

Next, we outline a numerical method to compute global, grid-free approximations of the equilibrium functions $(D_{x^i}\hat{\mathcal{J}}, D^2\hat{\mathcal{J}}, r^*)$ in Definition 14, where D_{x^i} denotes the *spatial* gradient operator with regards to idiosyncratic states x^i and D^2 denotes the full spatial Hessian operator. We recover the equilibrium policy function $y^{i,*}$ through the optimality condition (9) from $(D_{x^i}\hat{\mathcal{J}}, D^2\hat{\mathcal{J}})$ by assuming sufficient regularity of the state dynamics in (3) such that a unique, closed-form control map $\mathcal{C}$ exists.

Assumption 16 (Hamiltonian Regularity)**.**
The drift and dispersion functions μ, σ^o with $o \in \{id, ag\}$, and ζ are such that

(i) the HJB-operator $\mathcal{H}^{HJB}$ in (6) is convex in y^i;

(ii) (30) has a unique, closed-form solution in terms of the partial gradient $D_{x^i}\hat{\mathcal{J}} : [0,\infty) \times \mathbb{R}^N \to \mathbb{R}^{N^x}$ and the full Hessian $D^2\hat{\mathcal{J}} : [0,\infty) \times \mathbb{R}^N \to \mathbb{R}^{N\times N}$, where $N := N^x + N^a + \sum_{n=1}^{N^x} V_n$ is the number of spatial dimensions, i.e., there exists a unique closed-form operator $\mathcal{C}[D_{x^i}\hat{\mathcal{J}}, D^2\hat{\mathcal{J}}](t, x^i, a, \varphi, r) \in \mathbb{R}^{N^y}$ such that $\mathcal{C}[D_{x^i}\hat{\mathcal{J}}, D^2\hat{\mathcal{J}}] = y^{i,*}$, where $\varphi := \left(\{\varphi_1^v\}_{v=0}^{V_1}, \ldots, \{\varphi_{N^x}^v\}_{v=0}^{V_{N^x}}\right)$.[12]

Remark. *The operator $\mathcal{C}$ defines a function $\tilde{y}^{i,*} : \left([0,\infty) \times \mathbb{R}^N \times \mathbb{R}^{N^r}\right) \times \mathbb{R}^{N^x} \times \mathbb{R}^{N\times N} \to \mathbb{R}^{N^y}$ which maps values in the image $im(D_{x^i}\hat{\mathcal{J}}, D^2\hat{\mathcal{J}}) \in \mathbb{R}^{N^x} \times \mathbb{R}^{N\times N}$ into the optimal control $y^{i,*} \in \mathbb{R}^{N^y}$, for every $(t, x^i, a, \varphi, r) \in [0,\infty) \times \mathbb{R}^N \times \mathbb{R}^{N^r}$.*

We outline our solution method in two steps. First, we derive a probabilistic representation of recursive equilibrium. Second, we use this probablistic representation to deploy methods from the applied mathematics literature to jointly approximate $(D_{x^i}\hat{\mathcal{J}}, D^2\hat{\mathcal{J}}, r^*)$ such that $\hat{\mathcal{J}}$ satisfies (30)-(33) and r^* satisfies market clearing (34).

[12] Note that $y^{i,*}$ only depends on the parts of the Hessian w.r.t. x^i or a, since the cross-sectional dynamics do not matter for optimality. We omit denoting this explicitly since it does not matter for the design of our algorithm.

5.1 Probablistic Representation of the Approximated Recursive Equilibrium

Following the Deep BSDE method (E et al., 2017; Han et al., 2018), we reformulate the *analytical* task of finding (the gradient and Hessian of) $\hat{\mathcal{J}}$ as a *probabilistic* task of solving backward stochastic differential equations (BDSEs). Since we allow for controlled direct volatility, the PDE solved by $\hat{\mathcal{J}}$ is fully nonlinear and the associated BSDEs are of *second*-order. Such a nonlinear setting is considered in the Deep *2*BSDE method further developed by Beck et al. (2019), which forms the building block of the algorithm developed herein.

To outline the reformulation of the solution $\hat{\mathcal{J}}$ to (30)-(33), assume that some price function $r^*(t, \varphi, a)$ exists such that (34) is satisfied $\forall t \in [0, \infty)$. Then the functions $\mathcal{C}[D_{x^i}\hat{\mathcal{J}}, D^2\hat{\mathcal{J}}]$ and r^* can be substituted into the approximated master equation (31) such that it constitutes a fully-nonlinear, yet, standard PDE of dimension $1 + N$.[13] The Deep BSDE method formulates such PDEs equivalently as BSDEs through a generalization of the famous Feynman-Kac formula (see, e.g., Pardoux and Peng, 1992). Loosely speaking, this corresponds to applying Pontryagin's maximum principle yielding a *probabilistic* rather than *analytical* representation of optimal control problems (Yong and Zhou, 1999).[14] However, since the optimality condition (30) does not take into account the dynamics of the cross-sectional distribution, we cannot view the *master* equation (31) as the *HJB* equation of a standard stochastic control problem. Consequently, we cannot directly employ variations of the Deep BSDE method which are based on the stochastic maximum principle (e.g., Ji et al., 2022; Huang et al., 2025). Rather, we treat (31) as a generic fully-nonlinear PDE. Through the nonlinear Feynman-Kac formula in Cheridito et al. (2007, Theorem 4.10), it is related to the following system of strongly coupled forward and second-order backward

[13] The PDE is fully-nonlinear whenever $y^{i,*}$ depends on $D^2\hat{\mathcal{J}}$, that is, whenever agent-specific volatility is controlled, i.e., $\exists m \in \{id, ag, xa\}$ such that $\partial\Sigma^m/\partial y^i \neq 0$.

[14] We refer to Huang (2024) for an illustration in the context of economic models. Furthermore, note that under Assumption 16.i the second-order adjoint equation in the more general case treated by Yong and Zhou (1999) becomes irrelevant (see p. 141 and Proof of Ch.3 Theorem 5.4).

stochastic differential equations (FBDSE),[15]

$$
\begin{aligned}
x_{n,t}^{i,*} =& x_{n,0}^{i} + \int_0^t \mu_n(y_s^{i,*}, x_s^{i,*}, a_s, r_s^*)ds + \sum_{j=1}^{d^{id}} \int_0^t \sigma_{nj}^{id}(y_s^{i,*}, x_s^{i,*}, a_s, r_s^*)dB_{j,s}^{i} \qquad (35)\\
&+ \sum_{j=1}^{d^{ag}} \int_0^s \sigma_{nj}^{ag}(y_s^{i,*}, x_s^{i,*}, a_s, r_s^*)dB_{j,s}^{ag},\\
a_{n,t} =& a_{n,0} + \int_0^t \eta_n\left(a_s, r_s^*\right) ds + \sum_{j=1}^{d^{ag}} \int_0^t \zeta_{nj}\left(a_s, r_s^*\right) dB_{j,s}^{ag}, \qquad (36)\\
\varphi_{n,t}^{v} =& \varphi_{n,0}^{v} + \int_0^t \Bigg\{ \int_{\mathbb{R}^{N^\xi}} \frac{\phi^v(e)}{b^v} \mu_n\left(y^{i,*}(s, \varphi_s, x(e), a_s), x(e), a_s, r_s^*\right) m^\xi(e)de \qquad (37)\\
&+ \frac{1}{2} \sum_{o\in\{id,ag\}} \sum_{n'=1}^{N^x} \int_{\mathbb{R}^{N^\xi}} \frac{\frac{\partial}{\partial e}\phi^v(e)}{b^v \sum_{v'=1}^{V_{n'}} \varphi_{s,n'}^{v'} \frac{\partial}{\partial e}\phi^{v'}(e)} \Sigma_{nn'}^{o,*}(s, x(e), x(e)) m^\xi(e)de \Bigg\} ds\\
&+ \sum_{j=1}^{d^{ag}} \int_0^t \left\{ \int_{\mathbb{R}^{N^\xi}} \frac{\phi^v(e)}{b^v} \sigma_{nj}^{ag,*}\left(y^{i,*}(s, \varphi_s, x(e), a_s), x(e), a_s, r_s^*\right) m^\xi(e)de \right\} dB_{j,s}^{ag},\\
Y_t =& -\int_t^\infty \left\{\rho Y_s - u(y_s^{i,*}, x_s^{i,*})\right\} ds - \sum_{j=1}^{d^{id}} \sum_{n'=1}^{N^x} \int_t^\infty Z_{n',s}\sigma_{n'j}^{id}(y_s^{i,*}, x_s^{i,*}, a_s, r_s^*)dB_{j,s}^{i} \qquad (38)\\
&- \sum_{j=1}^{d^{ag}} \int_t^\infty \Bigg\{ \sum_{n'=1}^{N^x} Z_{n',s}\sigma_{n'j}^{ag}(y_s^{i,*}, x_s^{i,*}, a_s, r_s^*) + \sum_{n'=1}^{N^a} Z_{N^x+n',s}\zeta_{n'j}\left(a_s, r_s^*\right)\\
&+ \sum_{n'=1}^{N^x} \sum_{v'=0}^{V_{n'}} Z_{(n',v'),s} \int_{\mathbb{R}^{N^\xi}} \frac{\phi^{v'}(e)}{b^{v'}} \sigma_{n'j}^{ag,*}\left(y^{i,*}(s, \varphi_s, x(e), a_s), x(e), a_s, r_s^*\right) m^\xi(e)de \Bigg\} dB_{j,s}^{ag},\\
Z_{m,t} =& -\int_t^\infty A_{m,s}ds - \sum_{n'=1}^{N^x} \int_t^\infty \Gamma_{mn',s}dx_{n',s}^{i,*} - \sum_{n'=1}^{N^a} \int_t^\infty \Gamma_{m(N^a+n'),s}da_{n',s} \qquad (39)\\
&- \sum_{n'=1}^{N^x} \sum_{v'=0}^{V_{n'}} \int_t^\infty \Gamma_{m(n',v'),s}d\varphi_{n',s}^{v'},
\end{aligned}
$$

where $n \in \{1, \dots N^x\}$, $v \in \{0, \dots, V_n\}$, $m \in \{1, \dots, N\}$, $r_t^* = r^*(t, \varphi_t, a_t)$, and $y_t^{i,*} = \tilde{y}^{i,*}(t, x_t^{i,*}, a_t, \varphi_t, r_t^*, Z_{1:N^x,t}, \Gamma_t)$ is the optimal control of a specific household as defined in the Remark to Assumption 16. The dependence on Z_t and Γ_t foreshadows the equivalence result below and implies the strong coupling of the FBSDEs. Specifically, the forward SDEs (36)-(37) encode the aggregate dynamics of a specific economy with initial state (a_0, φ_0) and shock history $\{B_t^{ag}\}_{t\geq 0}$. The forward SDE (35) in turn

[15] Here, we index over the $\sum_{n=1}^{N^x} V_n$ projection coefficients by (n, v) denoting φ_n^v's position.

encodes the idiosyncratic state dynamics of a specific household with initial state x_0^i and shock history $\{B_t^i\}_{t\geq 0}$, which populates this economy. These forward SDEs depend on the backward SDEs (38)-(39) encoding the dynamics of the current value Y_t and the vector of spatial marginal values Z_t of this specific household. That is, $Y : [0,\infty)\times\Omega \to \mathbb{R}$ as well as $Z, A : [0,\infty)\times\Omega \to \mathbb{R}^N$ and $\Gamma : [0,\infty]\times\Omega \to \mathbb{R}^{N\times N}$ are household-specific $\mathcal{F}$-adapted stochastic processes. Defining the approximated value function $\hat{\mathcal{J}}$ as in (5), one can verify, using Itô's lemma twice, that indeed[16]

$$Y_t = \hat{\mathcal{J}}(t,\varphi_t,x_t^{i,*},a_t), \qquad Z_t = D\hat{\mathcal{J}}(t,\varphi_t,x_t^{i,*},a_t), \tag{40}$$

$$A_t = \mathcal{B}[D\hat{\mathcal{J}}](t,\varphi_t,x_t^{i,*},a_t), \qquad \Gamma_t = D^2\hat{\mathcal{J}}(t,\varphi_t,x_t^{i,*},a_t), \tag{41}$$

where $D\hat{\mathcal{J}} : [0,\infty)\times\mathbb{R}^N \to \mathbb{R}^N$ denotes the *full* spatial gradient of $\hat{\mathcal{J}}$ and the functional operator $\mathcal{B}$ is defined in Appendix B.1. In what follows, it is important to differentiate between the functions defined on the entire state space $[0,\infty)\times\mathbb{R}^N$ on the right-hand side of equations (40)-(41), and their values at $(t,\varphi_t,x_t^{i,*},a_t)$, which are random variables.

It remains to impose additional conditions on the processes $\{Y_t,Z_t,A_t,\Gamma_t\}_{t\geq 0}$ in (35)-(39) such that they also satisfy the transversality and state boundary conditions. In light of (40), the transversality condition (32) is equivalent to the terminal condition $\lim_{T\to\infty}[Y_T,Z_T]^T = 0$. The state boundary condition (33), which ensures by smooth pasting that the unconstrained HJB equation (13) converges to the constrained HJB equation at the boundary, is imposed by reflecting the dynamics $d\tilde{e}_t^i$ at $\tilde{e}^i = 0$. Specifically, the last N^e elements of dx_t^i are rewritten such that, at the boundary, the drift is nonnegative and the dispersion terms vanish, i.e.,

$$\begin{aligned} d\tilde{e}_t^i =& \mathbb{1}_{\tilde{e}_t^i>0}\psi(t,y_t^i,\tilde{x}_t^i,a_t,r_t)dt + \mathbb{1}_{\tilde{e}_t^i=0}\max\left(\psi(t,y_t^i,\tilde{x}_t^i,a_t,r_t),0\right)dt \\ &+ \mathbb{1}_{\tilde{e}_t^i>0}\left\{\upsilon^{id}\left(y_t^i,\tilde{x}_t^i,a_t,r_t\right)dB_t^i + \upsilon^{ag}\left(y_t^i,\tilde{x}_t^i,a_t,r_t\right)dB_t^{ag}\right\}. \end{aligned} \tag{42}$$

Given a function r^* satisfying (34) $\forall t\in[0,\infty)$, the processes $\{Y_t,Z_t,A_t,\Gamma_t\}_{t\geq 0}$ which solve the FBSDEs with the terminal condition $\lim_{T\to\infty}[Y_T,Z_T]^T = 0$, therefore, equivalently represent recursive equilibrium. E et al. (2017) were first to view the task of computing approximations of the functions in (40)-(41) as a deep reinforcement learning (RL) problem.

[16] For the necessary regularity assumptions underlying this result we refer to Cheridito et al. (2007).

5.2 Baseline Algorithm

A host of methods employs this insight to solve high-dimensional PDEs with varying properties (see E et al., 2021, for a review). Broadly speaking, these methods vary along 4 important dimensions: *(i)* how the RL problem is defined; *(ii)* whether the FBSDEs are temporally discretized in time as a backward or forward system; *(iii)* which subset of the functions in (40)-(41) is approximated with neural nets; *(iv)* how the RL problem is numerically implemented. Ideally, these algorithmic choices must not only be consistent with each other, but also be suitable to properties of the FBSDEs at hand, as well as retain convergence results established in the literature.

Here, our numerical algorithm must accommodate five critical properties. First, the FBSDEs are of second order since (31) is fully nonlinear. Second, the FBSDEs are strongly coupled because the drift and volatility term of $dx_t^{i,*}$ depend on Y_t and Z_t through the *individual* choices $y_t^{i,*}$. Both properties are due to controlled volatility as a natural feature of macro-finance models. To accomodate them we modify the Deep 2BSDE method by Beck et al. (2019) according to insights by Nüsken and Richter (2021); Huang et al. (2025). Third, the initial state (x_0^i, a_0, φ_0) is a $\mathbb{R}^N$-valued random variable distributed according to a calibrated empirical distribution $\tilde{M}_0(x, a, \varphi)$.[17] We adopt an appropriate modification of the Deep BSDE method, that is already proposed in the original paper by Han et al. (2018). Fourth, the cross-sectional dynamics (37) and market-clearing condition (34) depend on the policy function $\mathcal{C}[D_{x^i}\hat{\mathcal{J}}, D^2\hat{\mathcal{J}}]$ with subdomain $\mathbb{R}^{N^x}$ – not on a single realized policy $y_t^{i,*}$. Fifth, the FBSDEs are only well-defined given a price function r^*, which is, however, not known in closed form under Assumption 3.ii. The last two properties manifest the culprit of solving HACT models with aggregate risk and implicit market clearing conditions. To accommodate both features we augment the 2BSDE method by Beck et al. (2019) as detailed below.

***(i)* Reinforcement Learning Problem.** Intuitively, one can think of the RL problem introduced by Beck et al. (2019) as finding suitably defined functions for the initial (marginal) values $\vartheta^{Y_0}(x_0^i, a_0, \varphi_0) = Y_0$ and $\vartheta^{Z_0}(x_0^i, a_0, \varphi_0) = Z_0$, as well as for their drift and volatility functions $\vartheta^{A}(t, x_t^{i,*}, a_t, \varphi_t) = A_t$ and $\vartheta^{\Gamma}(t, x_t^{i,*}, a_t, \varphi_t) = \Gamma_t$, $\forall t \in [0, \infty)$. These functions must imply paths of $\{Y_t, Z_t\}_{t\geq 0}$ that minimize some

[17] This c.d.f. follows from a calibrated distribution $\tilde{M}_0(x, a) \sim (x', a')$ implying a jointly random vector of projection coefficients φ discretizing $\tilde{M}_0(x, a') \sim x'$ for every realization of a'.

"loss" functional relative to an exogenous target quantity. Although this functional needs to be chosen such that the optimum coincides with the solution of the FBSDEs, its specification leaves many degrees of freedom. For *decoupled* FBSDEs, Beck et al. (2019) employ the mean squared error of the terminal condition on Y_T as the exogenous target quantity. For *coupled* FBSDEs, recent findings (Andersson et al., 2023; Huang et al., 2025) show that solely relying on the terminal condition on Y_T leads to convergence issues. Therefore, we follow Huang et al. (2025), who develop a robust adaption for strongly coupled second-order FBSDEs, by augmenting the objective functional in (43) with the mean squared error of the terminal condition $\lim_{T\to\infty} Z_T = 0$. Furthermore, we follow Nüsken and Richter (2021), who show that using the error variance rather than the mean squared error of terminal conditions leads to more robust finite-sample approximations.

Because we need to track distributional dynamics and accommodate implicit market clearing, we augment the RL problem of Beck et al. (2019) along two dimensions. First, distributional dynamics depend on *all* household choices $y^{i,*} = \mathcal{C}[D_{x^i}\hat{\mathcal{J}}, D^2\hat{\mathcal{J}}]$, i.e., on the gradient and Hessian *functions* with subdomain $\mathbb{R}^{N^x}$. By the equivalence relation in (41), $D^2\hat{\mathcal{J}}(t, x_t^{i,*}, a_t, \varphi_t) = \vartheta^\Gamma(t, x_t^{i,*}, a_t, \varphi_t)$ is part of the RL problem's controlled functions. The gradient $D\hat{\mathcal{J}}(t, x_t^{i,*}, a_t, \varphi_t)$, however, is not a control in the RL problem of Beck et al. (2019); Huang et al. (2025), as its realized evolution $\{Z_t\}_{t\geq 0}$ serves as the predicted output. To evaluate the distributional dynamics (37), we therefore add the function $\vartheta^{Z_x}(t, x_t^{i,*}, a_t, \varphi_t) = D_{x^i}\hat{\mathcal{J}}(t, x_t^{i,*}, a_t, \varphi_t)$ as an additional function control and use its evaluated mean squared deviation from the realized processes $\{Z_{1:N^x,t}\}_{t\geq 0}$ over the time interval $[0, \infty)$ as the corresponding objective.

Second, implicit market clearing (Assumption 3.ii) introduces the unknown return function $\vartheta^r(t, a_t, \varphi_t) = r_t^*$ as an additional control of the RL problem (43). Specifically, household choices $y^{i,*} = \mathcal{C}[D_{x^i}\hat{\mathcal{J}}, D^2\hat{\mathcal{J}}](r_t^*)$, which take returns as given, must clear markets at all times, i.e., $\mathcal{M}_n^{im}[r^*, y^{i,*}](t, a_t, \varphi_t) = 0, \ \forall t \in [0, \infty), \forall n \in \{1, \dots, N^r\}$.[18] We employ the mean squared error for this market-clearing condition over the time interval $[0, \infty)$ as the corresponding objective. That is, optimal household behavior and the return function, which depend on the controlled functions $\vartheta := (\vartheta^{Y_0}, \vartheta^{Z_0}, \vartheta^{Z_x}, \vartheta^A, \vartheta^\Gamma, \vartheta^r)$, must jointly determine a path $\{a_t, \varphi_t\}_{t=0}^\infty$ through (37), which is consistent with market clearing. Meanwhile, explicit mar-

[18] In the Huggett (1993) economy (Example 4 and 15) the bond market clearing condition is simply $\mathcal{M}_b^{im}[y^{i,*}](t, a_t, \varphi_t) = \varphi_{t,b}^0 = 0$ as it does not depend on household choices directly.

ket clearing (Assumption 3.i) significantly simplifies the RL problem because then $r^* = \vartheta^r(t,a,\varphi) = \mathcal{M}^{ex}(t,a,\varphi)$ is known a priori. Therefore, we only outline our algorithm for the more general implicit case.

$$
\begin{cases}
\min_\vartheta \lim_{T\to\infty} \left\{ \mathbb{V}\mathrm{ar}\left[Y_T(\vartheta)\right] + \sum_{n=1}^N \mathbb{V}\mathrm{ar}\left[Z_{n,T}(\vartheta)\right] \right\} \\
\quad +\mathbb{E}\left[\int_0^\infty \left\| Z_{1:N^x,t}(\vartheta) - \vartheta^{Z_x}(t, x_t^{i,*}, a_t, \varphi_t) \right\|^2 dt \right] \\
\quad +\mathbb{E}\left[\int_0^\infty \left\| \mathcal{M}^{im}\left[r_t^*(\vartheta), \mathcal{C}[D_{x^i}\hat{\mathcal{J}}, D^2\hat{\mathcal{J}}](r_t^*(\vartheta)) \right](t, a_t(\vartheta), \varphi_t(\vartheta)) \right\|^2 dt \right], \\
\text{subject to (35)-(39),(42) with} \\
Y_0 = \vartheta^{Y_0}(x_0^i, a_0, \varphi_0), Z_0 = \vartheta^{Z_0}(x_0^i, a_0, \varphi_0), \text{ for a given } (x_0^i, a_0, \varphi_0) \sim \tilde{M}_0 \\
A_t = \vartheta^A(t, x_t^{i,*}, a_t, \varphi_t), \Gamma_t = \vartheta^\Gamma(t, x_t^{i,*}, a_t, \varphi_t), r_t^* = \vartheta^r(t, a_t, \varphi_t), \ \forall t \in [0,\infty) \\
D_{x^i}\hat{\mathcal{J}}(t,\cdot,a_t,\varphi_t) = \vartheta^{Z_x}(t,\cdot,a_t,\varphi_t), D^2\hat{\mathcal{J}}(t,\cdot,a_t,\varphi_t) = \vartheta^\Gamma(t,\cdot,a_t,\varphi_t)
\end{cases}
\tag{43}
$$

The terms in the objective of the RL problem (43) operationalize the three recursive equilibrium conditions of Definition 6 within its probabilistic representation derived above. The first two terms ensure *(i)* individual household optimality, the third term *(ii)* its consistency with the law of motion of the cross-sectional distribution, and the last term *(iii)* market clearing. Furthermore, it is important to note that (43) is a *global* problem in the sense that $\lim_{T\to\infty}\{Y_T(\vartheta), Z_T(\vartheta)\}$ and $\{Z_t(\vartheta), a_t(\vartheta), \varphi_t(\vartheta)\}_{t\geq 0}$ depend on the system's entire dynamics as controlled by ϑ. Hence, although one may think about this method as a sort of shooting algorithm, all loss terms can be reformulated into divergences in (controlled) probability measures over the entire time path, rather than just at terminal time (Nüsken and Richter, 2021). In light of this, the RL problem may be viewed as an importance sampling task. Its aim is to minimize variation in the terminal value of sampled households, deviations between their optimal behavior and that consistent with the distributional dynamics, as well as from market clearing – i.e., to minimize violations of equilibrium.

***(ii)* Temporal Discretization.** Since the FBSDEs are coupled, we discretize the FBSDEs into a forward Euler-Murayama scheme (Andersson et al., 2023; Huang et al., 2025). This also allows us to initialize simulated paths under a calibrated empirical distribution $\tilde{M}_0(x,a,\varphi)$. In order to handle the time domain $[0,\infty)$, we follow Friedl et al. (2023) and adopt the transformation of Traeger (2014) mapping unbounded physical time $t \in [0,\infty)$ into transformed time $\tau(t) = 1 - \exp(-\nu t) \in [0,1), \nu > 0$,

such that the inverse map $t(\tau_k) := \tau^{-1}(\tau_k)$ exists. The unit interval is then partitioned into K intervals $\cup_{k=0}^{K}[\tau_k, \tau_{k+1}) = [0,1)$ of equal length $1/K$, where $\tau_0 := 0$ and $\tau_K := 1$. Time is, hence, discretized into K intervals of increasing *physical* length $\Delta t_k := t(\tau_k) - t(\tau_{k-1}) = \frac{1}{\nu}\ln\left(1 + \frac{1/K}{1-\tau_k}\right)$. Under this discretization the FBSDEs turn into a forward Euler-Murayama scheme in transformed time, i.e., for $X_0 \sim \bar{M}_0$

$$\begin{bmatrix} X_{\tau_{k'}} \\ Y_{\tau_{k'}} \\ Z_{\tau_{k'}} \end{bmatrix} = \begin{bmatrix} X_0 \\ Y_0 \\ Z_0 \end{bmatrix} + \sum_{k=0}^{k'-1} F^{EM}\left(\tau_k, \begin{bmatrix} X_{\tau_k} \\ Y_{\tau_k} \\ Z_{\tau_k} \end{bmatrix}, \begin{bmatrix} A_{\tau_k} \\ \Gamma_{\tau_k} \\ r^*_{\tau_k} \end{bmatrix}; \mathcal{C}\begin{bmatrix} D_{x^i}\hat{\mathcal{J}}(t(\tau_k), \dot{X}_{\tau_k}), \\ D^2\hat{\mathcal{J}}(t(\tau_k), \dot{X}_{\tau_k}) \end{bmatrix}(r^*_{\tau_k}), \begin{bmatrix} \Delta t_k \\ \Delta B_{\tau_k} \end{bmatrix}\right), \tag{44}$$

where $X_{\tau_k} := [x^{i,*}_{\tau_k}, a_{\tau_k}, \varphi_{\tau_k}]^T$, $\dot{X}_{\tau_k} := [\cdot, a_{\tau_k}, \varphi_{\tau_k}]^T$, $\Delta B_{\tau_k} := \{\{\Delta B^i_{j,\tau_k}\}_{j=1}^{d^{id}}, \{\Delta B^{ag}_{j,\tau_k}\}_{j=1}^{d^{ag}}\}$ with $\Delta B^{o\in\{i,ag\}}_{j,\tau_k} \overset{i.i.d.}{\sim} \mathcal{N}(0, \Delta t_k)$, and the function F^{EM} is defined in Appendix B.2.

***(iii)* Neural Network Approximation.** Numerical implementation of the RL problem relies on deep learning, i.e., on parametrizing ϑ by neural nets.[19] Specifically, we closely follow Beck et al. (2019) and approximate the functions ϑ^A and ϑ^Γ at different time increments $\tau_k \in [0,1)$ with separate multilayer (deep) neural networks $\mathrm{N}^{\theta^A}_{\tau_k} : \mathbb{R}^N \to \mathbb{R}^N$ and $\mathrm{N}^{\theta^\Gamma}_{\tau_k} : \mathbb{R}^N \to \mathbb{R}^{N\times N}$. Analogously, the gradient ϑ^{Z_x} and return functions ϑ^r are parametrized by separate nets $\mathrm{N}^{\theta^{Z_x}}_{\tau_k} : \mathbb{R}^N \to \mathbb{R}^{N^x}$ and $\mathrm{N}^{\theta^r}_{\tau_k} : \mathbb{R}^{N-N^x} \to \mathbb{R}^{N^r}$. Lastly, the neural nets $\mathrm{N}^{\theta^{Y_0}}_0 : \mathbb{R}^N \to \mathbb{R}$ and $\mathrm{N}^{\theta^{Z_0}}_0 : \mathbb{R}^N \to \mathbb{R}^N$ parametrize the initial (marginal) values ϑ^{Y_0} and ϑ^{Z_0}. Given a network hyper parametrization $\Theta \ni \theta := \left[\theta^{Y_0}, \theta^{Z_0}, \theta^{Z_x}, \theta^A, \theta^\Gamma, \theta^r\right]^T$, the discretized RL problem is

$$\begin{cases} \min_{\theta\in\Theta} \mathbb{V}\mathrm{ar}\left[Y_1(\theta)\right] + \sum_{n=1}^{N} \mathbb{V}\mathrm{ar}\left[Z_{n,1}(\theta)\right] \\ \qquad +\mathbb{E}\left[\sum_{k=1}^{K}\left\|Z_{1:N^x,\tau_k}(\theta) - \mathrm{N}^{\theta^Z}_{\tau_k}(X_{\tau_k})\right\|^2\right] \\ \qquad +\mathbb{E}\left[\sum_{k=1}^{K}\left\|\mathcal{M}^{im}\left[r^*_{\tau_k}(\theta), \mathcal{C}[D_{x^i}\hat{\mathcal{J}}, D^2\hat{\mathcal{J}}](r^*_{\tau_k}(\theta))\right](t(\tau_k), \dot{X}_{\tau_k}(\theta))\right\|^2\right], \\ \text{s.t. (44) with } Y_0 = \mathrm{N}^{\theta^{Y_0}}(X_0), Z_0 = \mathrm{N}^{\theta^{Z_0}}(X_0), \text{ for a given } X_0 \sim \tilde{M}_0 \\ A_{\tau_k} = \mathrm{N}^{\theta^A}_{\tau_k}(X_{\tau_k}), \Gamma_{\tau_k} = \mathrm{N}^{\theta^\Gamma}_{\tau_k}(X_{\tau_k}), r^*_{\tau_k} = \mathrm{N}^{\theta^r}_{\tau_k}(X_{\tau_k}), \ \forall k \in \{0,\dots,K-1\} \\ D_{x^i}\hat{\mathcal{J}}(t(\tau_k), \dot{X}_{\tau_k}) = \mathrm{N}^{\theta^Z}_{\tau_k}(\dot{X}_t), D^2\hat{\mathcal{J}}(t(\tau_k), \dot{X}_{\tau_k}) = \mathrm{N}^{\theta^\Gamma}_{\tau_k}(\dot{X}_t) \end{cases}. \tag{45}$$

***(iv)* Iterative Scheme.** The Deep (2)BSDE method implements the RL prob-

[19] For introductions to neural networks we refer to, e.g., Maliar et al. (2021); Azinovic et al. (2022).

lem in (45) through mini-batch stochastic gradient descent. In the present context, Nüsken and Richter (2021) call such schemes – perhaps more instructively – "iterative diffusion optimization". It involves iteratively *1.)* estimating *all* expectations and variances given a parametrization $\theta_{e,i}$ through batch Monte Carlo simulation of the diffusion process (44); and *2.)* updating $\theta_{e,i}$ by the "Adam"-variant of gradient descent optimization (Kinga et al., 2015; E et al., 2017, Appendix A). To speed up training, Beck et al. (2019) also apply batch normalization (Ioffe and Szegedy, 2015) and a decaying learning rate schedule $\{\lambda_i\}_{i=1}^{K_{\text{batch}}}$ (e.g., Glorot and Bengio, 2010). Our adaption of their method is summarized by the pseudo-code in Algorithm 1. Appendix C provides further details on the network architecture and the Adam optimizer.

Algorithm 1 Approximation of $(D_{x^i}\hat{\mathcal{J}}, D^2\hat{\mathcal{J}}, r^*)$

Input: Hyper parametrization $M_{\text{batch}}, M_{\text{train}}, M_{\text{epochs}}, K, \nu, \{\lambda_i\}_{i=1}^{K_{\text{batch}}}, \Theta$;
Calibration $\tilde{M}_0 \in \mathcal{P}(\mathbb{R}^N)$

Output: Neural network approximations

$$\{\mathrm{N}_{\tau_k}^{\theta_*^{Z_x}}, \mathrm{N}_{\tau_k}^{\theta_*^{\Gamma}}, \mathrm{N}_{\tau_k}^{\theta_*^{r}}\}_{k=0}^{K-1} \approx \{D_{x^i}\hat{\mathcal{J}}(t(\tau_k), X), D^2\hat{\mathcal{J}}(t(\tau_k), X), r^*(t(\tau_k), X)\}_{k=0}^{K-1}$$

Initialize a network parametrization: $\theta_{0,K_{\text{batch}}+1} \in \Theta$

Draw M_{batch} states $X_0 \sim \tilde{M}_0$ and $M_{\text{train}} \times K$ Brownian increments ΔB_{τ_k}

for $e = 1, \dots, M_{\text{epochs}}$ (sequentially) **do**

  Initialize network parameters from previous epoch: $\theta_{e,1} = \theta_{e-1,M_{\text{batch}}+1}$

  for $i = 1, \dots, K_{\text{batch}} := M_{\text{train}}/M_{\text{batch}} \in \mathbb{N}$ (sequentially) **do**

    for $1 \le m \le M_{\text{batch}}$ (in parallel) **do**

      Initialize X_0^m from mth of the M_{batch} draws of X_0 and set

$$Y_0^m = \mathrm{N}_0^{\theta_{e,i}^{Y_0}}(X_0^m), \quad Z_0 = \mathrm{N}_0^{\theta_{e,i}^{Z_0}}(X_0^m)$$

      for $k = 0, \dots, K-1$ (sequentially) **do**

        Set network approximation:

$$\begin{bmatrix} A_{\tau_k}^m \\ \Gamma_{\tau_k}^m \\ r_{\tau_k}^* \end{bmatrix} = \begin{bmatrix} \mathrm{N}_{\tau_k}^{\theta_{e,i}^{A}}(X_{\tau_k}^m) \\ \mathrm{N}_{\tau_k}^{\theta_{e,i}^{\Gamma}}(X_{\tau_k}^m) \\ \mathrm{N}_{\tau_k}^{\theta_{e,i}^{r}}(X_{\tau_k}^m) \end{bmatrix}, \quad \begin{bmatrix} D_{x^i}\hat{\mathcal{J}}(t(\tau_k), \dot{X}_{\tau_k}^m) \\ D^2\hat{\mathcal{J}}(t(\tau_k), \dot{X}_{\tau_k}^m) \end{bmatrix} = \begin{bmatrix} \mathrm{N}_{\tau_k}^{\theta_{e,i}^{Z_x}}(\dot{X}_{\tau_k}^m) \\ \mathrm{N}_{\tau_k}^{\theta_{e,i}^{\Gamma}}(\dot{X}_{\tau_k}^m) \end{bmatrix}$$

        Compute next values according to Euler-Murayama scheme (44):[20]

$$\begin{bmatrix} X_{\tau_{k+1}}^m \\ Y_{\tau_{k+1}}^m \\ Z_{\tau_{k+1}}^m \end{bmatrix} = \begin{bmatrix} X_{\tau_k}^m \\ Y_{\tau_k}^m \\ Z_{\tau_k}^m \end{bmatrix} + F^{EM}\left(\tau_k, \begin{bmatrix} X_{\tau_k}^m \\ Y_{\tau_k}^m \\ Z_{\tau_k}^m \end{bmatrix}, \begin{bmatrix} A_{\tau_k}^m \\ \Gamma_{\tau_k}^m \\ r_{\tau_k}^* \end{bmatrix}; \mathcal{C}\begin{bmatrix} D_{x^i}\hat{\mathcal{J}}(t(\tau_k), \dot{X}_{\tau_k}^m), \\ D^2\hat{\mathcal{J}}(t(\tau_k), \dot{X}_{\tau_k}^m) \end{bmatrix}(r_{\tau_k}^*), \begin{bmatrix} \Delta t_k \\ \Delta B_{\tau_k} \end{bmatrix}\right)$$

[20] In the application to the example model, we furthermore "enforce" market clearing by setting $\varphi_{b,\tau_k}^{0,m} = 0 \ \forall k, m$ (see Example 15) in order to avoid convergence issues associated with training samples that violate market clearing (Azinovic et al., 2022). The deviation from this condition enters the last term of the loss function.

13: **end for**
14: **end for**
15: Compute an estimation of the objective (loss) function in (45):

$$\begin{aligned}\Phi(\theta_{e,i}) :=& \frac{\sum_{m=1}^{M_{\text{batch}}}}{M_{\text{batch}}-1}\left(Y_1^m - \sum_{m'=1}^{M_{\text{batch}}} Y_1^{m'}\right)^2 + \frac{\sum_{n=1}^{N}\sum_{m=1}^{M_{\text{batch}}}}{M_{\text{batch}}-1}\left(Z_{n,1}^m - \sum_{m'=1}^{M_{\text{batch}}} Z_{n,1}^{m'}\right)^2 \\ &+ \frac{1}{M_{\text{batch}}}\sum_{m=1}^{M_{\text{batch}}}\sum_{k=1}^{K}\sum_{n=1}^{N^x}\left(Z_{n,\tau_k}^m - \mathrm{N}_{n,\tau_k}^{\theta_{e,i}^{Z_x}}(X_{\tau_k}^m)\right)^2 \\ &+ \frac{1}{M_{\text{batch}}}\sum_{m=1}^{M_{\text{batch}}}\sum_{k=1}^{K}\sum_{n=1}^{N^r}\left(\mathcal{M}_n^{im}\left[\mathrm{N}_{\tau_k}^{\theta_{e,i}^r}, \mathcal{C}\left[\mathrm{N}_{\tau_k}^{\theta_{e,i}^{Z_x}}, \mathrm{N}_{\tau_k}^{\theta_{e,i}^{\Gamma}}\right](\mathrm{N}_{\tau_k}^{\theta_{e,i}^r})\right](t(\tau_k), \dot{X}_{\tau_k}^m)\right)^2\end{aligned} \tag{46}$$

16: Update $\theta_{e,i} = [\theta_{e,i}^{Y_0}, \theta_{e,i}^{Z_0}, \theta_{e,i}^{Z_x}, \theta_{e,i}^{A}, \theta_{e,i}^{\Gamma}, \theta_{e,i}^{r}]^T$ by "Adam" with learning rate $\lambda_i > 0$

$$\theta_{e,i+1} = \theta_{e,i} - \lambda_i \text{Adam}\left(D_{\theta_{e,i}}\Phi(\theta_{e,i})\right) \tag{47}$$

17: **end for**
18: **end for**
19: Return $\theta_* = \theta_{M_{\text{epochs}},M_{\text{batch}}+1}$

Remark. *The K_{batch} mini-batches sample over the initial distribution $\tilde{M}_0 \in \mathcal{P}(\mathbb{R}^N)$ and trajectories of $\{X, Y, Z\}$ to estimate the objective (loss) function $\Phi(\theta_{e,i})$ in step 15.*

The algorithm has some desirable properties for solving HACT models with aggregate risk. First, the simulation step ensures a global, grid-free approximation since each optimization step is based on sampled data rather than a pre-determined grid. As pointed out by Azinovic et al. (2022), this approximates the ergodic set of the recursive equilibrium under $\theta_{e,i}$. This avoids computation in regions of the state space, which are not visited in equilibrium. The method, therefore, caters to equilibrium state spaces with irregularly shaped (non-hypercubic) geometries (see e.g., Maliar et al., 2011). Second, the algorithm can be augmented with a parallelization scheme to speed up its numerical execution (see e.g., Azinovic et al., 2022, Appendix A.9).

Third, the algorithm benefits from the well-known computational efficiency of neural nets. Since they represent a sequence of matrix operations, the computational complexity in their evaluation and automatic differentiation is of a similar magnitude (Azinovic et al., 2022). The Deep BSDE method exploits this regularity in sampling many batches and in evaluating the gradient $D_{\theta_{e,i}}$ in (47) to handle high-dimensional PDEs. Here, the high dimensionality is due to the $\sum_{n=1}^{N^x} V_n$ projection coefficients discretizing the household distribution. These projection coefficients track the cross-sectional evolution of the economy – the methodological roadblock in HACT models

with aggregate risk. Since Polynomial Chaos expansions can capture the (tail of) distributions particularly well if $\sum_{n=1}^{N^x} V_n$ is large, we view its combination with deep PDE solvers as a natural strategy to solve such models.

6 A Macro-Finance Example

Using this solution method, we study the interaction between uninsurable income risk, borrowing constraints, and asset illiquidity in a macro-finance model with aggregate risk. Although this interaction has gained much attention since the influential work by Kaplan et al. (2018), it has not yet been extensively investigated under aggregate risk. Therefore, the effect of variations in households' self-insurance and hedging motives against risk are not fully understood. We first inspect a canonical model with inelastic labor supply in which self-insurance and hedging can only be achieved through saving-investment decisions. In an extension, we study the effects of endogenous household-specific labor choice and of a collateral rather than absolute borrowing constraint.

6.1 Environment

We consider the two-asset economy in Krusell and Smith (1997) – an extension of Examples 1 and 8 – with illiqudity in productive capital as in Kaplan et al. (2018). Specifically, individual worker-households hold illiquid capital and rent it to a perfectly competitive representative firm producing $F(L_t, K_t) - \delta K_t$ units of a single good, $\delta > 0$.[21] Furthermore, we consider *direct* aggregate risk shocking household's individual capital rather than an aggregate variable, e.g., TFP. We have that $d^{id} = d^{ag} = 1$, $N^x = 3$, and $N^a = 0$. Therefore, we drop any dependence on a_t.

Household Problem. Each household is characterized by its labor productivity $e^{z_t^i}$ earning the wage rate w_t, liquid bond holdings b_t^i paying risk-free rate r_t^b, and illiquid capital holdings k_t^i paying the risky return r_t^k, i.e., $x_t^i = [z_t^i, b_t^i, k_t^i]$, $r_t = [w_t, r_t^b, r_t^k]$. Log-labor productivity z_t^i follows an idiosyncratic Ornstein-Uhlenbeck process which is correlated with aggregate shocks by $\omega \in [-1, 1]$. Illiquid capital is subject to k_t^i-proportional aggregate shocks and convex adjustment costs $\Psi(\iota_t^i) = \psi_0|\iota_t^i| + \psi_1|\iota_t^i|^{\psi_2}$ per ι_t^i invested (divested) units ($\psi_0, \psi_1 > 0$; $\psi_2 > 1$). Households have CRRA-preferences over consumption, $u(c_t^i) = (c_t^i)^{1-\gamma}/(1-\gamma)$, inelastically supply labor, and

[21] Therefore, there are no firm profits to be distributed.

choose controls $y_t^i = [c_t^i, \iota_t^i]$ subject to a borrowing and short-selling constraint, i.e., $e(t, x_t^i) = [b_t^i - \underline{b}, k_t^i]$. Hence, the household problem is

$$\begin{aligned} \max_{\{y_t^i\}_{t\geq 0}} \mathbb{E}_0 \left[\int_0^\infty e^{-\rho t} (c_t^i)^{1-\gamma}/(1-\gamma) dt \right] \\ \text{s.t. } dz_t^i &= \kappa_z(\bar{z} - z_t^i)dt + \sqrt{1-\omega^2}\sigma_z dB_t^i + \omega\sigma_z dB_t^{ag}, \\ db_t^i &= \left\{ w_t e^{z_t^i} + b_t^i r_t^b + k_t^i(r_t^k - \iota_t^i - \Psi(\iota_t^i)) - c_t^i \right\} dt, \\ dk_t^i &= \iota_t^i k_t^i dt + k_t^i \sigma_K dB_t^{ag}, \\ [b_t^i - \underline{b}, k_t^i] &\geq 0. \end{aligned} \tag{48}$$

Optimal Controls. The optimal controls $y_t^{i,*}$ satisfy the first-order condition

$$[c_t^{i,*},\, 0] = [(\mathcal{J}_b)^{-\frac{1}{\gamma}},\, -k_t^i(1 + \Psi_\iota(\iota_t^{i,*}))\mathcal{J}_b + k_t^i \mathcal{J}_k], \tag{49}$$

such that households' optimal investment policies feature an inaction region, i.e., $0 \in \text{im}(\iota^{i,*})$, due to the proportional component of adjustment costs ($\psi_0 > 0$): 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, which here is normalized to one, is higher than the minimum marginal cost ψ_0 of doing so.[22]

Proposition 17 (Optimal Investment Policy)**.** *For $k_t^i > 0$ a.s., if (49) holds, then*

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

Furthermore, it holds then that the set of optimal wealth allocations without adjustment costs is in the inaction region, i.e., assuming $k_t^i + b_t^i \neq 0$ a.s.,

$$\frac{k_t^i}{k_t^i + b_t^i} = \underbrace{-\frac{\mathcal{V}_k}{\mathcal{V}_{kk}(k_t^i + b_t^i)} \frac{r_t^k - r_t^b}{\sigma_K^2}}_{\text{(i) Myopic diversification demand}} \underbrace{-\frac{\mathcal{V}_{zk}}{\mathcal{V}_{kk}(k_t^i + b_t^i)} \frac{\omega\sigma_z}{\sigma_K}}_{\text{(ii) Hedging demand}} \Rightarrow \iota_t^{i,*} = 0, \tag{50}$$

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

where $\mathcal{V}$ denotes households' value function in an equivalent economy with $\Psi(\iota_t^i) = 0$.

Therefore, the inaction region increases in the minimum marginal cost ψ_0. Furthermore, Davis and Norman (1990) prove that, without convex costs and labor income (risk; $\bar{z}, \psi_1, \sigma_z = 0$), the optimally controlled illiquid wealth share $\{k_t^i/(k_t^i + b_t^i)\}_{t\geq 0}$ follows a reflected diffusion inside a ψ_0-dependent wedge around the unconstrained optimum in (50). We conjecture and will numerically verify that this property extends to the present model:[23] Households' illiquid capital demand comprises of *(i)* their myopic diversification demand, *(ii)* their hedging demand due to labor income risk, and *(iii)* a tolerated deviation from their sum due to adjustment costs $\Psi(\cdot)$.

Market Clearing. To ease notation, the market-clearing operator $\mathcal{H}^{MC}$ encodes optimal factor inputs as well as clearing of the labor, bond, and capital market,

$$\mathcal{H}^{MC}[y_t^{i,*}, r_t^*](t, M_t) = \begin{bmatrix} F_L\left(\int_{\mathbb{R}^3} e^z dM_t(x), \int_{\mathbb{R}^3} k dM_t(x)\right) - w_t^* \\ \int_{\mathbb{R}^3} b dM_t(x) \\ F_K\left(\int_{\mathbb{R}^3} e^z dM_t(x), \int_{\mathbb{R}^3} k dM_t(x)\right) - \delta - r_t^{k,*} \end{bmatrix}, \tag{51}$$

such that the condition (25) yields explicit expressions of w_t^* and $r_t^{k,*}$, which can be substituted into the master operator $\mathcal{H}^M$ in (20), and defines the risk-free rate $r_t^{b,*}$ as a fixed point. The sequence of fixed points $\{r_t^{b,*}\}_{t\geq 0}$ and household controls $\{y_t^{i,*}\}_{t\geq 0, i\in[0,1]}$ jointly constitute recursive equilibrium as specified in Proposition 10.

Extension. In an extension, we consider a collateral rather than absolute borrowing constraint $[b_t^i + \pi k_t^i, k_t^i] \geq 0$, where $\pi \in (0, 1]$ is the pledgeability parameter. Furthermore, we assume that households no longer supply labor inelastically. They choose individually optimal working hours $l^{i,*}(t, x_t^i, M_t) \in y_t^{i,*}$ maximizing expected discounted flow utility $u(c_t^i, l_t^i) = u(c_t^i) - \varepsilon_1 (l_t^i)^{1-\varepsilon_2}/(1-\varepsilon_2)$, where $\varepsilon_1 > 0$ is a preference shifter and $1/\varepsilon_2 \geq 0$ is the Frisch elasticity.

6.2 Calibration and Results

Work in progress.

[23] Isaenko (2005) considers the same convex adjustment costs, as we do here, for a finite-horizon, complete-market case without intermediate consumption concluding that the shape of the inaction region is independent of the degree of convexity $\psi_2 - 1 > 0$.

7 Conclusions

This paper addresses a longstanding computational challenge in solving heterogeneous agent models with aggregate risk by introducing a novel global solution method. While existing approaches either rely on local approximations or focus on quasi-stationary settings with deterministic KFEs, our method extends to fully stochastic environments, accommodating richer aggregate nonlinearities. This setting is, for instance, also applicable to fully non-stationary models crucial in environmental economics, e.g. Friedl et al. (2023). By leveraging insights from applied mathematics – particularly the master equation and polynomial chaos expansion – we transform an inherently infinite-dimensional problem into a high-dimensional yet tractable PDE.

Our approach not only expands the computational feasibility of rational expectations models but also provides a scalable and flexible framework for studying macroeconomic dynamics under aggregate uncertainty. By adapting the Deep BSDE method, we demonstrate how neural networks can effectively approximate agents' optimal controls, offering a reinforcement learning-based solution to forward-backward stochastic differential equations.

To illustrate the power of our method, we applied it to a two-asset HACT model with aggregate risk, showing how collateral constraints and endogenous labor choices impact the wealth distribution and hedging motives. This framework lays the foundation for future research into nonlinear heterogeneous agent models, enabling new insights in macroeconomics, macro-finance, and environmental economics. By overcoming key computational hurdles, our work contributes to a more precise and comprehensive understanding of macroeconomic risk and policy design.

Appendix A Proofs

A.1 Proof of Proposition 5

Proof. The dynamics of the c.d.f. $\{M_t(x)\}_{t\geq 0, x\in\mathbb{R}^{N^x}}$ with $M_t(x) = \int_{z\leq x} m_t(z)dz$ follow from integrating (10) over $z \leq x \in \mathbb{R}^{N^x}$,

$$
\begin{aligned}
dM_t(x) =& \int_{z\leq x} dm_t(z)dz \\
=& \sum_{n=1}^{N^x} \Bigg\{ - \int_{z\leq x} \frac{\partial}{\partial x_n} \left[\mu_n\left(y^{i,*}(t,x,a_t,r_t),x,a_t,r_t\right) m_t(x)\right]\big|_{z_n=x_n} dz \\
&+ \frac{1}{2} \sum_{o\in\{id,ag\}} \int_{z\leq x} \frac{\partial^2}{\partial x_n^2} \left[\Sigma^o_{nn'}\left(y^{i,*}(t,x,a_t,r_t),x,a_t,r_t\right) m_t(x)\right]\big|_{z_n=x_n} dz \\
&+ \sum_{n'=1}^{n-1} \sum_{o\in\{id,ag\}} \int_{z\leq x} \frac{\partial^2}{\partial x_n \partial x_{n'}} \left[\Sigma^o_{nn'}\left(y^{i,*}(t,x,a_t,r_t),x,a_t,r_t\right) m_t(x)\right]\big|_{z_{n,n'}=x_{n,n'}} dz \Bigg\} dt \\
&- \sum_{n=1}^{N^x} \left\{ \int_{z\leq x} \frac{\partial}{\partial x_n} \left[\sigma^{ag}_n\left(y^{i,*}(t,x,a_t,r_t),x,a_t,r_t\right) m_t(x)\right]\big|_{z_n=x_n} dz \right\} dB^{ag}_t \\
=& \sum_{n=1}^{N^x} \Bigg\{ - \int_{z_{-n}\leq x_{-n}} \left[\mu_n\left(y^{i,*}(t,z,a_t,r_t),z,a_t,r_t\right) m_t(z)\right]\big|_{z_n=x_n} dz_{-n} \\
&+ \frac{1}{2} \sum_{o\in\{id,ag\}} \int_{z_{-n}\leq x_{-n}} \frac{\partial}{\partial x_n} \left[\Sigma^o_{nn}\left(y^{i,*}(t,z,a_t,r_t),z,a_t,r_t\right) m_t(z)\right]\big|_{z_n=x_n} dz_{-n} \\
&+ \sum_{n'=1}^{n-1} \sum_{o\in\{id,ag\}} \int_{z_{-n,n'}\leq x_{-n,n'}} \left[\Sigma^o_{nn'}\left(y^{i,*}(t,z,a_t,r_t),z,a_t,r_t\right) m_t(z)\right]\big|_{z_{n,n'}=x_{n,n'}} dz_{-n,n'} \Bigg\} dt \\
&- \sum_{n=1}^{N^x} \left\{ \int_{z_{-n}\leq x_{-n}} \left[\sigma^{ag}_n\left(y^{i,*}(t,z,a_t,r_t),z,a_t,r_t\right) m_t(z)\right]\big|_{z_n=x_n} dz_{-n} \right\} dB^{ag}_t.
\end{aligned}
$$

where the second equality exploits symmetry of the Hessian of $\Sigma^o_{nn'}m_t(x)$, $o \in \{id, ag\}$ and the third equality follow from partial integration in dimensions n and n'. The dynamics of the marginal c.d.f.s $\{M_t(x_n)\}_{t\geq 0, x_n\in\mathbb{R}}$ in (11) follow from $M_t(x_n) = \lim_{x_{-n}\to\infty} M_t(x)$, i.e., from marginalizing out x_{-n} under the differential operator. □

A.2 Proof of Proposition 13

Proof. The distributional derivatives of $\tilde{f}(t,\chi^i_t) = \hat{f}(t,\{\varphi^v_{1,t}\}_{v=0}^{\infty},\ldots,\{\varphi^v_{N^x,t}\}_{v=0}^{\infty})$ follow from the chain rule and applying Definition 7 to each projection coefficient in

$\{\varphi^v_{n,t}\}_{v=0}^{\infty}$, $1 \leq n \leq N^x$, as a function of the corresponding element $(\chi^i_t)_n$ of the multivariate state random variable in the RHS of (28),

$$\begin{aligned}
\tilde{f}_{\chi^i_n}(t, \chi^i_t, x(e)) &= \sum_{v=0}^{\infty} \frac{\partial}{\partial \varphi^v_n} \hat{f}(t, \{\varphi^v_{1,t}\}_{v=0}^{\infty}, \ldots, \{\varphi^v_{N^x,t}\}_{v=0}^{\infty}) \frac{\partial \varphi^v_{n,t}}{\partial \chi^i_n} \\
&= \sum_{v=0}^{\infty} \frac{\partial}{\partial \varphi^v_n} \hat{f}(t, \{\varphi^v_{1,t}\}_{v=0}^{\infty}, \ldots, \{\varphi^v_{N^x,t}\}_{v=0}^{\infty}) \frac{\phi^v(e)}{b^v}, \\
\tilde{f}_{\chi^i_n \chi^i_{n'}}(t, \chi^i_t, x(e), x(e')) &= \sum_{v,v'=0}^{\infty} \frac{\partial}{\partial \varphi^{v'}_{n'}} \left\{ \frac{\partial}{\partial \varphi^v_n} \hat{f}(t, \{\varphi^v_{1,t}\}_{v=0}^{\infty}, \ldots, \{\varphi^v_{N^x,t}\}_{v=0}^{\infty}) \frac{\phi^v(e)}{b^v} \right\} \frac{\partial \varphi^v_{n',t}}{\partial \chi^{i\prime}_n} \\
&= \sum_{v,v'=0}^{\infty} \frac{\partial^2}{\partial \varphi^v_n \partial \varphi^{v'}_{n'}} \hat{f}(t, \{\varphi^v_{1,t}\}_{v=0}^{\infty}, \ldots, \{\varphi^v_{N^x,t}\}_{v=0}^{\infty}) \frac{\phi^v(e)}{b^v} \frac{\phi^{v'}(e')}{b^{v'}}, \\
\tilde{f}_{\chi^i_n x_{n'}}(t, \chi^i_t, x(e)) &= \sum_{v=0}^{\infty} \frac{d}{dx_{n'}(e)} \left\{ \frac{\partial}{\partial \varphi^v_n} \hat{f}(t, \{\varphi^v_{1,t}\}_{v=0}^{\infty}, \ldots, \{\varphi^v_{N^x,t}\}_{v=0}^{\infty}) \frac{\phi^v(e)}{b^v} \right\} \\
&= \sum_{v=0}^{\infty} \frac{\frac{\partial}{\partial \varphi^v_n} \hat{f}(t, \{\varphi^v_{1,t}\}_{v=0}^{\infty}, \ldots, \{\varphi^v_{N^x,t}\}_{v=0}^{\infty})}{b^v} \frac{d\phi^v(e)}{dx_{n'}(e)} \\
&= \sum_{v=0}^{\infty} \frac{\frac{\partial}{\partial \varphi^v_n} \hat{f}(t, \{\varphi^v_{1,t}\}_{v=0}^{\infty}, \ldots, \{\varphi^v_{N^x,t}\}_{v=0}^{\infty})}{b^v \frac{dx_{n'}(e)}{de}} \frac{\partial}{\partial e} \phi^v(e) \\
&= \sum_{v=0}^{\infty} \frac{\frac{\partial}{\partial \varphi^v_n} \hat{f}(t, \{\varphi^v_{1,t}\}_{v=0}^{\infty}, \ldots, \{\varphi^v_{N^x,t}\}_{v=0}^{\infty}) \frac{\partial}{\partial e} \phi^v(e)}{b^v \sum_{v'=0}^{\infty} \varphi^{v'}_{n',t} \frac{\partial}{\partial e} \phi^{v'}(e)}.
\end{aligned}$$

□

A.3 Proof of Proposition 17

Proof. The proof follows the same steps as Isaenko (2005). If $k^i_t > 0$,

$$\begin{aligned}
(49) \;\Rightarrow\; & -\mathcal{J}_b \left(1 + \psi_0 \operatorname{sgn}(\iota^{i,*}_t) + \psi_1 \psi_2 |\iota^{i,*}_t|^{\psi_2 - 1} \operatorname{sgn}(\iota^{i,*}_t)\right) + \mathcal{J}_k = 0, \\
\overset{\iota^{i,*}_t \neq 0}{\Rightarrow} \;& \psi_0 + \psi_1 \psi_2 |\iota^{i,*}_t|^{\psi_2 - 1} = |\mathcal{J}_k / \mathcal{J}_b - 1|,
\end{aligned}$$

which has the solution

$$|\iota^{i,*}_t| = \left(\frac{|\mathcal{J}_k / \mathcal{J}_b - 1| - \psi_0}{\psi_1 \psi_2} \right)^{\frac{1}{\psi_2 - 1}},$$

as long as $\psi_0 < |\mathcal{J}_k/\mathcal{J}_b - 1|$. Otherwise, the first-order condition is satisfied by $\iota_t^{i,*} = 0$. The first expression in Proposition 17 follows from collecting cases.
The implication in equation (50) follows from $\mathcal{V} \geq \mathcal{J}$ due to (the absence of) adjustment costs. That is, if the wealth allocation under adjustment costs equals the optimal wealth allocation in the frictionless case, no additional utility can be gained by deviating from it. The optimal wealth allocation in the frictionless case follows from the first-order condition in the associated HJB equation with control variable $\theta_t^i \equiv \frac{k_t^i}{k_t^i + b_t^i}$, which can be freely chosen if $\Psi(\iota_t^i) = 0$ such that $\mathcal{V}_k = \mathcal{V}_b$, i.e., from

$$\rho\mathcal{V} = \max_{c_t^i, \theta_t^i} \Big\{ \mathcal{V}_t + u(c_t^i) + [w_t e^{l_t^i} + (k_t^i + b_t^i) r_t^b + \theta_t^i (k_t^i + b_t^i)(r_t^k - r^b) - c_t^i] \mathcal{V}_b \\ + \frac{\sigma_K^2}{2} \left(\theta_t^i\right)^2 (k_t^i + b_t^i)^2 \mathcal{V}_{kk} + \pi_l (\bar{l} - l_t^i) \mathcal{V}_l + \frac{\sigma_l^2}{2} \mathcal{V}_{ll} + \theta_t^i (k_t^i + b_t^i) \omega \sigma_l \sigma_K \mathcal{V}_{lk} \Big\}.$$

□

A.4 Alternative Derivation of $\mathcal{H}^{aM}$ in (29)

Let each set of projection coefficients $\{\varphi_{t,n}^v\}_{v=0}^{V_n}$ be defined by the set of time-invariant functions $\{\varphi_n^v(\chi_t^i)\}_{v=0}^{V_n}$ in the RHS of (28). Equivalently, they are defined by functions $\{\tilde{\varphi}_n^v(M_t)\}_{v=0}^{V_n}$ of the corresponding c.d.f. $M_t \sim \chi_t^i$. By Proposition 12 and the Proof of Proposition 13 each coefficient's distributional derivatives are, therefore, defined as

$$\frac{\partial}{\partial M_m} \tilde{\varphi}_n^v(M_t(x)) = \frac{\partial}{\partial \chi_m^i} \varphi_n^v(\chi_t^i, x(e)) = \begin{cases} \frac{\phi^v(e)}{b^v}, & \text{if } m = n \\ 0, & \text{if } m \neq n \end{cases}$$

$$\frac{d}{dx_{n'}} \frac{\partial}{\partial M_m} \tilde{\varphi}_n^v(M_t(x)) = \frac{d}{dx_{n'}} \frac{\partial}{\partial \chi_m^i} \varphi_n^v(\chi_t^i, x(e)) = \begin{cases} \frac{\frac{\partial}{\partial e}\phi^v(e)}{b^v \sum_{v'=1}^{V_{n'}} \varphi_{t,n'}^{v'} \frac{\partial}{\partial e}\phi^{v'}(e)}, & \text{if } m = n \\ 0, & \text{if } m \neq n \end{cases}$$

$$\frac{\partial^2}{\partial M_m \partial M_{m'}} \tilde{\varphi}_n^v(M_t(x, x')) = \frac{\partial^2}{\partial \chi_m^i \partial \chi_{m'}^i} \varphi_n^v(\chi_t^i, x(e), x(e')) = 0.$$

Applying the generalized Itô's formula in (18), each project coefficient follows

$$
\begin{aligned}
d\tilde{\varphi}_n^v(M_t) =& \Bigg\{ \int_{\mathbb{R}^{N^x}} \frac{\partial}{\partial M_n} \tilde{\varphi}_n^v(M_t(x)) \mu_n \left(y^{i,*}(t,x,a_t,r_t), x, a_t, r_t \right) m_t(x) dx \\
&+ \frac{1}{2} \sum_{o \in \{id, ag\}} \sum_{n'=1}^{N^x} \int_{\mathbb{R}^{N^x}} \frac{d}{dx_{n'}} \frac{\partial}{\partial M_n} \tilde{\varphi}_n^v(M_t(x)) \Sigma_{nn'}^{o,*}(t,x,x) m_t(x) dx \Bigg\} dt \\
&+ \left\{ \int_{\mathbb{R}^{N^x}} \frac{\partial}{\partial M_n} \tilde{\varphi}_n^v(M_t(x)) \sigma_n^{ag} \left(y^{i,*}(t,x,a_t,r_t), x, a_t, r_t \right) m_t(x) dx \right\} dB_t^{ag} \\
=& \Bigg\{ \int_{\mathbb{R}^{N^\xi}} \frac{\phi^v(e)}{b^v} \mu_n \left(y^{i,*}(t,x(e),a_t,r_t), x(e), a_t, r_t \right) m^\xi(e) de \\
&+ \frac{1}{2} \sum_{o \in \{id, ag\}} \sum_{n'=1}^{N^x} \int_{\mathbb{R}^{N^\xi}} \frac{\frac{\partial}{\partial e} \phi^v(e)}{b^v \sum_{v'=1}^{V_{n'}} \varphi_{t,n'}^{v'} \frac{\partial}{\partial e} \phi^{v'}(e)} \Sigma_{nn'}^{o,*}(t,x(e),x(e)) m^\xi(e) de \Bigg\} dt \\
&+ \left\{ \int_{\mathbb{R}^{N^\xi}} \frac{\phi^v(e)}{b^v} \sigma_n^{ag,*} \left(y^{i,*}(t,x(e),a_t,r_t), x(e), a_t, r_t \right) m^\xi(e) de \right\} dB_t^{ag}.
\end{aligned}
$$

The approximated master operator $\mathcal{H}^{aM}$ in (29) subtracts $\rho \hat{\mathcal{J}} - u(y_t^{i,*}, x_t^i)$ from $\frac{\mathbb{E}[d\hat{\mathcal{J}}]}{dt}$, which follows from applying Itô's lemma to $\hat{\mathcal{J}} \left(t, \{\varphi_{1,t}^v\}_{v=0}^{V_1}, \ldots, \{\varphi_{N^x,t}^v\}_{v=0}^{V_{N^x}}, x_t^i, a_t \right)$ for $d\varphi_{n,t}^v = d\tilde{\varphi}_n^v(M_t)$. To arrive at the third line of (29), note that if $\sigma_n^{ag,*}(t,x)\phi^v(e) \in L^1(\Omega, \mathcal{F}_t, \mathbb{P}, \mathbb{R}^{N^\xi})$ $\forall v \leq V_n, \forall n \leq N^x$, where $\sigma_n^{ag,*}(t,x) \equiv \sigma_n^{ag} \left(y^{i,*}(t,x,a_t,r_t), x, a_t, r_t \right)$, it holds that

$$
\begin{aligned}
&\int_{\mathbb{R}^{N^\xi}} \sigma_n^{ag,*}(t,x(e)) \phi^v(e) m^\xi(e) de < \infty, \ \forall v \leq V_n, \forall n \leq N^x \\
&\Rightarrow \left(\int_{\mathbb{R}^{N^\xi}} \sigma_n^{ag,*}(t,x(e)) \phi^v(e) m^\xi(e) de \right) \left(\int_{\mathbb{R}^{N^\xi}} \sigma_{n'}^{ag,*}(t,x(e')) \phi^{v'}(e') m^\xi(e') de' \right) \\
&= \int_{\mathbb{R}^{N^\xi}} \int_{\mathbb{R}^{N^\xi}} \sigma_n^{ag,*}(t,x(e)) \sigma_{n'}^{ag,*}(t,x(e')) \phi^v(e) m^\xi(e) de \phi^{v'}(e') m^\xi(e') de' \\
&= \int_{\mathbb{R}^{2N^\xi}} \Sigma_{nn'}^{ag,*}(t,x(e),x(e')) \phi^v(e) m^\xi(e) de \phi^{v'}(e') m^\xi(e') de'.
\end{aligned}
$$

Appendix B Expressions

B.1 Expressions in Section 5.1

$$\mathcal{B}[f](t,\varphi_t,x_t^i,a_t) = \frac{\partial f}{\partial t}(t,\varphi_t,x_t^i,a_t) + \frac{1}{2}\sum_{n,n'=1}^{N^a} f_{a_n a_{n'}} \Sigma^a_{nn'}(a,r) \tag{52}$$

$$+ \frac{1}{2}\sum_{n,n'=1}^{N^x}\sum_{v=0}^{V_n}\sum_{v'=0}^{V_{n'}} \frac{f_{\varphi_n^v \varphi_{n'}^{v'}}}{b^v b^{v'}} \int_{\mathbb{R}^{2N^\xi}} \Sigma^{ag,*}_{nn'}(t,x(e),x(e'))\phi^v(e)m^\xi(e)de\phi^{v'}(e')m^\xi(e')de'$$

$$+ \frac{1}{2}\sum_{n=1}^{N^x}\sum_{n'=1}^{N^a}\sum_{v=0}^{V_n} \frac{f_{\varphi_n^v a_{n'}}}{b^v} \int_{\mathbb{R}^{N^\xi}} \Sigma^{xa}_{nn'}\left(y^i,x(e),a,r\right)\phi^v(e)m^\xi(e)de$$

$$+ \frac{1}{2}\sum_{n,n'=1}^{N^x}\sum_{v=0}^{V_n} \frac{f_{\varphi_n^v x_{n'}^i}}{b^v} \int_{\mathbb{R}^{N^\xi}} \Sigma^{ag,*}_{nn'}\left(t,x^i,x(e)\right)\phi^v(e)m^\xi(e)de$$

$$+ \frac{1}{2}\sum_{n,n'=1}^{N^x}\sum_{o\in\{id,ag\}} f_{x_n^i x_{n'}^i}\Sigma^{o,*}_{nn'}\left(t,x^i,x^i\right) + \frac{1}{2}\sum_{n=1}^{N^x}\sum_{n'=1}^{N^a} f_{x_n^i a_{n'}}\Sigma^{xa}_{nn'}\left(y^i,x^i,a,r\right).$$

B.2 Expressions in Section 5.2

The function F^{EM} in equation (44) is defined such that

$$
\begin{aligned}
x^{i,*}_{n,\tau_{k'}} =& x^i_{n,0} + \sum_{k=0}^{k'-1} F^{x,uc}_n \left(\tau_k, x^{i,*}_{\tau_k}, a_{\tau_k}, \varphi_{\tau_k}, Z_{\tau_k}, \Gamma_{\tau_k}; \Delta t_k, \Delta B_{\tau_k}\right), \\
x^{i,*}_{n',\tau_{k'}} =& x^i_{n',0} + \sum_{k=0}^{k'-1} F^{x,c}_{n''} \left(\tau_k, x^{i,*}_{\tau_k}, a_{\tau_k}, \varphi_{\tau_k}, Z_{\tau_k}, \Gamma_{\tau_k}; \Delta t'_k, \Delta B_{\tau_k}\right), \\
& \qquad \text{where } n \in \{1, \dots, N^x - N^e\},\ n' = N^x - N^e + n'',\ n'' \in \{1, \dots, N^e\}, \\
a_{n,\tau_{k'}} =& a_{n,0} + \sum_{k=0}^{k'-1} F^a_n \left(\tau_k, a_{\tau_k}, \varphi_{\tau_k}; \Delta t_k, \{\Delta B^{ag}_{j,\tau_k}\}_{j=1}^{d^{ag}}\right) \\
\varphi^v_{n,\tau_{k'}} =& \varphi^v_{n,0} + \sum_{k=0}^{k'-1} F^{\varphi,v}_n \left(\tau_k, a_{\tau_k}, \varphi_{\tau_k}, D_{x^i}\hat{\mathcal{J}}, D^2\hat{\mathcal{J}}; \Delta t_k, \{\Delta B^{ag}_{j,\tau_k}\}_{j=1}^{d^{ag}}\right) \\
Y_{\tau_{k'}} =& Y_0 + \sum_{k=0}^{k'-1} F^Y \left(\tau_k, x^{i,*}_{\tau_k}, a_{\tau_k}, \varphi_{\tau_k}, Z_{\tau_k}, \Gamma_{\tau_k}, D_{x^i}\hat{\mathcal{J}}, D^2\hat{\mathcal{J}}; \Delta t_k, \Delta B_{\tau_k}\right) \\
Z_{m,\tau_{k'}} =& Z_{m,0} + \sum_{k=0}^{k'-1} \Bigg\{ A_{m,\tau_k} \Delta t_k + \sum_{n'=1}^{N^x} \Gamma_{mn',t} \Delta x^{i,*}_{n',\tau_k} + \sum_{n'=1}^{N^a} \Gamma_{m(N^a+n'),\tau_k} \Delta a_{n',\tau_k} \\
& + \sum_{n'=1}^{N^x} \sum_{v'=0}^{V_{n'}} \Gamma_{m(n'v'),\tau_k} \Delta \varphi^{v'}_{n',\tau_k} \Bigg\},
\end{aligned}
$$

where

$$
\begin{aligned}
F_n^{x,uc}(\cdot) =& G_n^{x,uc}\left(\tau_k, x_{\tau_k}^{i,*}, a_{\tau_k}, \varphi_{\tau_k}, Z_{\tau_k}, \Gamma_{\tau_k}\right)\Delta t_{k+1} + \sum_{j=1}^{d^{id}} H_{nj}^{x,id,uc}\left(\tau_k, x_{\tau_k}^{i,*}, a_{\tau_k}, \varphi_{\tau_k}, Z_{\tau_k}, \Gamma_{\tau_k}\right)\Delta B_{j,\tau_{k+1}}^{i},\\
&+ \sum_{j=1}^{d^{ag}} H_{nj}^{x,ag,uc}\left(\tau_k, x_{\tau_k}^{i,*}, a_{\tau_k}, \varphi_{\tau_k}, Z_{\tau_k}, \Gamma_{\tau_k}\right)\Delta B_{j,\tau_{k+1}}^{ag},\ n \in \{1,\dots,N^x - N^e\}\\
F_{n'}^{x,c}(\cdot) =& G_{n''}^{x,c}\left(\tau_k, x_{\tau_k}^{i,*}, a_{\tau_k}, \varphi_{\tau_k}, Z_{\tau_k}, \Gamma_{\tau_k}, \Delta t_{k+1}, \{\Delta B_{j,\tau_{k+1}}^{i}\}_{j=1}^{d^{id}}, \{\Delta B_{j,\tau_{k+1}}^{ag}\}_{j=1}^{d^{ag}}\right)\Delta t_{k+1}\\
&+ \sum_{j=1}^{d^{id}} H_{n''j}^{x,id,c}\left(\tau_k, x_{\tau_k}^{i,*}, a_{\tau_k}, \varphi_{\tau_k}, Z_{\tau_k}, \Gamma_{\tau_k}, \Delta t_{k+1}, \{\Delta B_{j,\tau_{k+1}}^{i}\}_{j=1}^{d^{id}}, \{\Delta B_{j,\tau_{k+1}}^{ag}\}_{j=1}^{d^{ag}}\right)\Delta B_{j,\tau_{k+1}}^{i}\\
&+ \sum_{j=1}^{d^{ag}} H_{n''j}^{x,ag,c}\left(\tau_k, x_{\tau_k}^{i,*}, a_{\tau_k}, \varphi_{\tau_k}, Z_{\tau_k}, \Gamma_{\tau_k}, \Delta t_{k+1}, \{\Delta B_{j,\tau_{k+1}}^{i}\}_{j=1}^{d^{id}}, \{\Delta B_{j,\tau_{k+1}}^{ag}\}_{j=1}^{d^{ag}}\right)\Delta B_{j,\tau_{k+1}}^{ag},\\
&\qquad\qquad n' = N^x - N^e + n'',\ n'' \in \{1,\dots,N^e\}\\
F_n^{a}(\cdot) =& G_n^{a}\left(\tau_k, a_{\tau_k}, \varphi_{\tau_k}\right)\Delta t_{k+1} + \sum_{j=1}^{d^{ag}} H_{nj}^{a}\left(\tau_k, a_{\tau_k}, \varphi_{\tau_k}\right)\Delta B_{j,\tau_{k+1}}^{ag}\\
F_n^{\varphi}(\cdot) =& G_n^{\varphi,v}\left(\tau_k, a_{\tau_k}, \varphi_{\tau_k}, D_{x^i}\hat{\mathcal{J}}, D^2\hat{\mathcal{J}}\right)\Delta t_{k+1} + \sum_{j=1}^{d^{ag}} H_{nj}^{\varphi,v}\left(\tau_k, a_{\tau_k}, \varphi_{\tau_k}, D_{x^i}\hat{\mathcal{J}}, D^2\hat{\mathcal{J}}\right)\Delta B_{j,\tau_{k+1}}^{ag}\\
F_n^{Y}(\cdot) =& G^{Y}\left(\tau_k, x_{\tau_k}^{i,*}, Y_{\tau_k}, Z_{\tau_k}, \Gamma_{\tau_k}\right)\Delta t_{k+1} + \sum_{j=1}^{d^{id}} H_{j}^{Y,id}\left(\tau_k, x_{\tau_k}^{i,*}, a_{\tau_k}, \varphi_{\tau_k}, Z_{\tau_k}, \Gamma_{\tau_k}\right)\Delta B_{j,\tau_{k+1}}^{i}\\
&+ \sum_{j=1}^{d^{ag}} H_{j}^{Y,ag}\left(\tau_k, x_{\tau_k}^{i,*}, a_{\tau_k}, \varphi_{\tau_k}, Z_{\tau_k}, \Gamma_{\tau_k}, D_{x^i}\hat{\mathcal{J}}, D^2\hat{\mathcal{J}}\right)\Delta B_{j,\tau_{k+1}}^{ag}
\end{aligned}
$$

and in turn the drift and diffusion functions G and H are defined with $t(\tau_k) := \tau^{-1}(\tau_k)$ as

$$
\begin{aligned}
G_n^{x,uc}(\cdot) &= \tilde{\mu}_n^{uc}\left(y^{i,*}(t(\tau_k), \varphi_{\tau_k}, x_{\tau_k}^{i,*}, a_{\tau_k}; Z_{\tau_k}, \Gamma_{\tau_k}, r^*(t(\tau_k), \varphi_{\tau_k}, a_{\tau_k})), x_{\tau_k}^{i,*}, a_{\tau_k}, r^*(t(\tau_k), \varphi_{\tau_k}, a_{\tau_k})\right),\\
H_{nj}^{x,o,uc}(\cdot) &= \tilde{\sigma}_{nj}^{o,uc}\left(y^{i,*}(t(\tau_k), \varphi_{\tau_k}, x_{\tau_k}^{i,*}, a_{\tau_k}; Z_{\tau_k}, \Gamma_{\tau_k}, r^*(t(\tau_k), \varphi_{\tau_k}, a_{\tau_k})), x_{\tau_k}^{i,*}, a_{\tau_k}, r^*(t(\tau_k), \varphi_{\tau_k}, a_{\tau_k})\right),\\
&\qquad\qquad n \in \{1,\dots,N^x - N^e\}
\end{aligned}
$$

$$
\begin{aligned}
G^{x,c}_{n'}(\cdot) =& \mathbb{1}_{x^{i,*}_{n'}=0} \max\Big\{ \psi_{n'}\left(y^{i,*}(t(\tau_k),\varphi_{\tau_k},x^{i,*}_{\tau_k},a_{\tau_k};Z_{\tau_k},\Gamma_{\tau_k},r^*(t(\tau_k),\varphi_{\tau_k},a_{\tau_k})),\right.\\
&\left. x^{i,*}_{\tau_k},a_{\tau_k},r^*(t(\tau_k),\varphi_{\tau_k},a_{\tau_k}),0\right)\Big\}\\
&+\mathbb{1}_{x^{i,*}_{n}>0}\Big\{\mathbb{1}_{E_{n'}\geq 0}\psi_{n'}\big(y^{i,*}(t(\tau_k),\varphi_{\tau_k},x^{i,*}_{\tau_k},a_{\tau_k};Z_{\tau_k},\Gamma_{\tau_k},r^*(t(\tau_k),\varphi_{\tau_k},a_{\tau_k})),\\
&x^{i,*}_{\tau_k},a_{\tau_k},r^*(t(\tau_k),\varphi_{\tau_k},a_{\tau_k})\big)-\mathbb{1}_{E_{n'}<0}\frac{x^{i,*}_{n,\tau_k}}{\Delta t_{k+1}}\Big\}\\
H^{x,o,c}_{n'}(\cdot) =& \mathbb{1}_{x^{i,*}_{n''}>0}\mathbb{1}_{E_{n'}\geq 0}\upsilon^o_{n'}\left(y^{i,*}(t(\tau_k),\varphi_{\tau_k},x^{i,*}_{\tau_k},a_{\tau_k};Z_{\tau_k},\Gamma_{\tau_k},r^*(t(\tau_k),\varphi_{\tau_k},a_{\tau_k})),x^{i,*}_{\tau_k},a_{\tau_k},r^*(t(\tau_k),\varphi_{\tau_k},a_{\tau_k})\right)\\
&\text{where } E_{n'}=E_{n'}\left(t(\tau_k),x^{i,*}_{\tau_k},a_{\tau_k},\varphi_{\tau_k},Z_{\tau_k},\Gamma_{\tau_k},\Delta t_{k+1},\{\Delta B^i_{j,\tau_{k+1}}\}_{j=1}^{d^{id}},\{\Delta B^{ag}_{j,\tau_{k+1}}\}_{j=1}^{d^{ag}}\right) \text{ with}\\
E_{n'}(\cdot) =& x^{i,*}_{n'',\tau_k}+\psi_{n'}\left(y^{i,*}(t(\tau_k),\varphi_{\tau_k},x^{i,*}_{\tau_k},a_{\tau_k};Z_{\tau_k},\Gamma_{\tau_k},r^*(t(\tau_k),\varphi_{\tau_k},a_{\tau_k})),x^{i,*}_{\tau_k},a_{\tau_k},r^*(t(\tau_k),\varphi_{\tau_k},a_{\tau_k})\right)\Delta t_{k+1}\\
&+\sum_{j=1}^{d^{id}}\upsilon^{id}_{n'j}\left(y^{i,*}(t(\tau_k),\varphi_{\tau_k},x^{i,*}_{\tau_k},a_{\tau_k};Z_{\tau_k},\Gamma_{\tau_k},r^*(t(\tau_k),\varphi_{\tau_k},a_{\tau_k})),x^{i,*}_{\tau_k},a_{\tau_k},r^*(t(\tau_k),\varphi_{\tau_k},a_{\tau_k})\right)\Delta B^i_{j,\tau_{k+1}}\\
&+\sum_{j=1}^{d^{ag}}\upsilon^{ag}_{n'j}\left(y^{i,*}(t(\tau_k),\varphi_{\tau_k},x^{i,*}_{\tau_k},a_{\tau_k};Z_{\tau_k},\Gamma_{\tau_k},r^*(t(\tau_k),\varphi_{\tau_k},a_{\tau_k})),x^{i,*}_{\tau_k},a_{\tau_k},r^*(t(\tau_k),\varphi_{\tau_k},a_{\tau_k})\right)\Delta B^{ag}_{j,\tau_{k+1}}\\
&\qquad\qquad n''=N^x-N^e+n',\ n'\in\{1,\dots,N^e\}\\
G^a_n(\cdot) =& \eta_n\left(a_{\tau_k},r^*(t(\tau_k),\varphi_{\tau_k},a_{\tau_k})\right),\qquad H^a_{nj}(\cdot)=\zeta_{nj}\left(a_{\tau_k},r^*(t(\tau_k),\varphi_{\tau_k},a_{\tau_k})\right)\\
G^{\varphi,v}_n(\cdot) =& \Big\{\int_{\mathbb{R}^{N^\xi}}\frac{\phi^v(e)}{b^v}G^{x,uc}_n\left(\tau_k,x(e),a_{\tau_k},\varphi_{\tau_k},D_{x^i}\hat{\mathcal{J}}(t(\tau_k),x(e),\cdot),D^2\hat{\mathcal{J}}(t(\tau_k),x(e),\cdot)\right)m^\xi(e)de\\
&+\frac{1}{2}\sum_{o\in\{id,ag\}}\sum_{n'=1}^{N^x}\int_{\mathbb{R}^{N^\xi}}\frac{\frac{\partial}{\partial e}\phi^v(e)}{b^v\sum_{v'=1}^{V_{n'}}\varphi^{v'}_{t,n'}\frac{\partial}{\partial e}\phi^{v'}(e)}\Sigma^{o,*}_{nn'}(t(\tau_k),x(e),x(e);D_{x^i}\hat{\mathcal{J}},D^2\hat{\mathcal{J}})m^\xi(e)de\Big\},\\
&\qquad\qquad n\in\{1,\dots,N^x-N^e\}\\
G^{\varphi,v}_{n'}(\cdot) =& \Big\{\int_{\mathbb{R}^{N^\xi}}\frac{\phi^v(e)}{b^v}G^{x,c}_{n''}\left(\tau_k,x(e),a_{\tau_k},\varphi_{\tau_k},D_{x^i}\hat{\mathcal{J}}(t(\tau_k),x(e),\cdot),D^2\hat{\mathcal{J}}(t(\tau_k),x(e),\cdot)\right)m^\xi(e)de\\
&+\frac{1}{2}\sum_{o\in\{id,ag\}}\sum_{n'''=1}^{N^x}\int_{\mathbb{R}^{N^\xi}}\frac{\frac{\partial}{\partial e}\phi^v(e)}{b^v\sum_{v'''=1}^{V_{n'''}}\varphi^{v'''}_{t,n'''}\frac{\partial}{\partial e}\phi^{v'''}(e)}\Sigma^{o,*}_{n'n'''}(t(\tau_k),x(e),x(e);D_{x^i}\hat{\mathcal{J}},D^2\hat{\mathcal{J}})m^\xi(e)de\Big\},\\
&\qquad\qquad n'=N^x-N^e+n'',\ n''\in\{1,\dots,N^e\}\\
H^{\varphi,v}_{nj}(\cdot) =& \int_{\mathbb{R}^{N^\xi}}\frac{\phi^v(e)}{b^v}H^{x,ag,uc}_{nj}\left(\tau_k,x(e),a_{\tau_k},\varphi_{\tau_k},D_{x^i}\hat{\mathcal{J}}(t(\tau_k),x(e),\cdot),D^2\hat{\mathcal{J}}(t(\tau_k),x(e),\cdot)\right)m^\xi(e)de,\\
&\qquad\qquad n\in\{1,\dots,N^x-N^e\}\\
H^{\varphi,v}_{n'j}(\cdot) =& \int_{\mathbb{R}^{N^\xi}}\frac{\phi^v(e)}{b^v}H^{x,ag,c}_{n''j}\left(\tau_k,x(e),a_{\tau_k},\varphi_{\tau_k},D_{x^i}\hat{\mathcal{J}}(t(\tau_k),x(e),\cdot),D^2\hat{\mathcal{J}}(t(\tau_k),x(e),\cdot)\right)m^\xi(e)de,\\
&\qquad\qquad n'=N^x-N^e+n'',\ n''\in\{1,\dots,N^e\}
\end{aligned}
$$

$$
\begin{aligned}
G^{Y}(\cdot) =& \rho Y_{\tau_k} - u\left(y^{i,*}(t(\tau_k), \varphi_{\tau_k}, x^{i,*}_{\tau_k}, a_{\tau_k}; Z_{\tau_k}, \Gamma_{\tau_k}, r^*(t(\tau_k), \varphi_{\tau_k}, a_{\tau_k})), x^{i,*}_{\tau_k}\right), \\
H^{Y,id}_j(\cdot) =& \sum_{n=1}^{N^x - N^e} Z_{n,\tau_k} H^{x,id,uc}_{nj}\left(\tau_k, x^{i,*}_{\tau_k}, a_{\tau_k}, \varphi_{\tau_k}, Z_{\tau_k}, \Gamma_{\tau_k}\right) \\
&+ \sum_{n=1}^{N^e} Z_{n,\tau_k} H^{x,id,c}_{nj}\left(\tau_k, x^{i,*}_{\tau_k}, a_{\tau_k}, \varphi_{\tau_k}, Z_{\tau_k}, \Gamma_{\tau_k}\right), \\
H^{Y,ag}_j(\cdot) =& \sum_{n'=1}^{N^x} \sum_{v'=0}^{V_n} Z_{N^x + N^a + (n',v'),\tau_j} H^{\varphi,v'}_{n'j}\left(\tau_k, a_{\tau_k}, \varphi_{\tau_k}, D_{x^i}\hat{\mathcal{J}}, D^2\hat{\mathcal{J}}\right) \\
&+ \sum_{n'=1}^{N^a} Z_{N^x + n',\tau_k} H^{a}_{n'j}\left(\tau_k, a_{\tau_k}, \varphi_{\tau_k}\right) + \sum_{n'=1}^{N^x - N^e} Z_{n',\tau_k} H^{x,ag,uc}_{n'j}\left(\tau_k, x^{i,*}_{\tau_k}, a_{\tau_k}, \varphi_{\tau_k}, Z_{\tau_k}, \Gamma_{\tau_k}\right) \\
&+ \sum_{n'=1}^{N^e} Z_{n',\tau_k} H^{x,ag,c}_{n'j}\left(\tau_k, x^{i,*}_{\tau_k}, a_{\tau_k}, \varphi_{\tau_k}, Z_{\tau_k}, \Gamma_{\tau_k}\right)
\end{aligned}
$$

Appendix C Details on Algorithm 1

The goal of this section is to describe details of the baseline algorithm developed in section 5, which are not necessary to follow its main idea, but important for its practical application – specifically, the network architecture with batch normalization and the Adam-optimizer with decaying learning rates. We try to offer a concise yet self-contained description without assuming more than a basic knowledge of deep learning. Although this comes at the expense of the generality and mathematical rigor of the formulation in Beck et al. (2019, Framework 3.2), we adopt some of their notation and, thus, deviate somewhat from the simplified notation in the main text.

C.1 Network Architecture with Batch Normalization

The network architecture follows Beck et al. (2019, Fig. 5). Specially, all intermediate-time neural networks $\mathrm{N}^{\theta^j}_{\tau_k}$, $j \in \{Z_x, A, \Gamma, r\}$ and intial-time neural networks $\mathrm{N}^{\theta^j}_{0}$, $j \in \{Y_0, Z_0\}$ are fully connected with two hidden layers with respectively $\mathfrak{N}^j_1 \in \mathbb{N}$ and $\mathfrak{N}^j_2 \in \mathbb{N}$ neurons. Input and output layer dimensions are determined by the respective domain and range. Moreover, we also adopt rectifier activation functions at the input and hidden layers $\mathcal{R}_{o\in\mathbb{N}}(x \in \mathbb{R}^o) := [\max\{x_1, 0\}, \ldots, \max\{x_o, 0\}]^T$ as a common choice for deep neural networks (see e.g., Glorot et al., 2011). Furthermore, these neural networks are trained using batch normalization (Ioffe and Szegedy, 2015) *before* activation. The basic idea is to normalize inputs (here, the idiosyncratic or aggregate

state variables) $X^m_{\tau_k}$ and neurons $L^{l,j,m}_{\tau_k}$, $l \in \{1,2\}$, using data from observation m's mini-batch $i \in \{1,\ldots,K_{\text{batch}}\}$ at time increment τ_k, to regularize optimization; i.e.,

$$\text{BN}^{\theta^j,i}(X^m_{\tau_k}) = \boldsymbol{\gamma}^{j,i}_{\tau_k} \odot \frac{X^m_{\tau_k} - \hat{\boldsymbol{\mu}}^i_{\tau_k}}{\hat{\boldsymbol{\sigma}}^i_{\tau_k}} + \boldsymbol{\beta}^{j,i}_{\tau_k}, \quad \text{BN}^{\theta^j,l,i}(L^{l,m}_{\tau_k}) = \boldsymbol{\gamma}^{j,l,i}_{\tau_k} \odot \frac{X^m_{\tau_k} - \hat{\boldsymbol{\mu}}^{j,l,i}_{\tau_k}}{\hat{\boldsymbol{\sigma}}^{j,l,i}_{\tau_k}} + \boldsymbol{\beta}^{j,l,i}_{\tau_k},$$

where operations are element-wise, $\hat{\boldsymbol{\mu}}^i_{\tau_k}, \hat{\boldsymbol{\sigma}}^i_{\tau_k} \in \mathbb{R}^N$ ($\mathbb{R}^{N-N^x}$ for $j = r$) are the mean and standard deviation vectors of the inputs, and $\hat{\boldsymbol{\mu}}^{l,j,i}_{\tau_k}, \hat{\boldsymbol{\sigma}}^{l,j,i}_{\tau_k} \in \mathbb{R}^{\mathfrak{N}^j_l}$, $l \in \{1,2\}$ those of the neurons of net j across the mini-batch at iteration i and time increment τ_k. The vectors $\boldsymbol{\gamma}^{j,i}_{\tau_k}, \boldsymbol{\beta}^{j,i}_{\tau_k} \in \mathbb{R}^N$ ($\mathbb{R}^{N-N^x}$ for $j = r$) and $\boldsymbol{\gamma}^{j,l,i}_{\tau_k}, \boldsymbol{\beta}^{j,l,i}_{\tau_k} \in \mathbb{R}^{\mathfrak{N}^j_l}$, $l \in \{1,2\}$ are trainable parameters at iteration i. *In training,* each deep neural net, therefore, depends on an iteration's parameter vector $\theta^j_i \in \mathbb{R}^{\nu_j}$, and on its mini-batch data – which we donote by the additional superscript i. Defining the affine linear transformation $M^{\theta^j,v}_{p,q} : \mathbb{R}^q \to \mathbb{R}^p$, $v, p, q \in \mathbb{N}$, which applies network weights and biases, as

$$M^{\theta^j,v}_{p,q}(x \in \mathbb{R}^q) = \left[\left(\sum_{i=1}^{q} \theta^j_{v+i} x_i \right) + \theta^j_{v+pq+1}, \ldots, \left(\sum_{i=1}^{q} \theta^j_{v+(p-1)q+i} x_i \right) + \theta^j_{v+pq+p} \right]^T$$

the neural nets can, hence, for $k \in \{0, \dots, K-1\}$ with $\tau_0 = 0$, be defined as

$$
\begin{aligned}
\mathrm{N}_0^{\theta^{Y_0},i} =& M_{1,\mathfrak{N}_2^{Y_0}}^{\theta^{Y_0},2N+\mathfrak{N}_1^{Y_0}(N+3)+\mathfrak{N}_2^{Y_0}(\mathfrak{N}_1^{Y_0}+3)} \circ \mathcal{R}_{\mathfrak{N}_2^{Y_0}} \circ \mathrm{BN}^{\theta^{Y_0},2,i} \circ M_{\mathfrak{N}_2^{Y_0},\mathfrak{N}_1^{Y_0}}^{\theta^{Y_0},2N+\mathfrak{N}_1^{Y_0}(N+3)} \\
& \circ \mathcal{R}_{\mathfrak{N}_1^{Y_0}} \circ \mathrm{BN}^{\theta^{Y_0},1,i} \circ M_{\mathfrak{N}_1^{Y_0},N}^{\theta^{Y_0},2N} \circ \mathrm{BN}^{\theta^{Y_0},i}(X_{\tau_k}^m), \\
& \nu_{Y_0} = 2N + \mathfrak{N}_1^{Y_0}(N+3) + \mathfrak{N}_2^{Y_0}(\mathfrak{N}_1^{Y_0}+4) + 1; \\
\mathrm{N}_0^{\theta^{Z_0},i} =& M_{N,\mathfrak{N}_2^{Z_0}}^{\theta^{Z_0},2N+\mathfrak{N}_1^{Z_0}(N+3)+\mathfrak{N}_2^{Z_0}(\mathfrak{N}_1^{Z_0}+3)} \circ \mathcal{R}_{\mathfrak{N}_2^{Z_0}} \circ \mathrm{BN}^{\theta^{Z_0},2,i} \circ M_{\mathfrak{N}_2^{Z_0},\mathfrak{N}_1^{Z_0}}^{\theta^{Z_0},2N+\mathfrak{N}_1^{Z_0}(N+3)} \\
& \circ \mathcal{R}_{\mathfrak{N}_1^{Z_0}} \circ \mathrm{BN}^{\theta^{Z_0},1,i} \circ M_{\mathfrak{N}_1^{Z_0},N}^{\theta^{Z_0},2N} \circ \mathrm{BN}^{\theta^{Z_0},i}(X_{\tau_k}^m), \\
& \nu_{Z_0} = N(\mathfrak{N}_2^{Y_0}+3) + \mathfrak{N}_1^{Z_0}(N+3) + \mathfrak{N}_2^{Z_0}(\mathfrak{N}_1^{Z_0}+3); \\
\mathrm{N}_{\tau_k}^{\theta^{Z_x},i} =& M_{N^x,\mathfrak{N}_2^{Z_x}}^{\theta^{Z_x},(k+1)\left(2N+\mathfrak{N}_1^{Z_x}(N+3)+\mathfrak{N}_2^{Z_x}(\mathfrak{N}_1^{Z_x}+3)\right)} \circ \mathcal{R}_{\mathfrak{N}_2^{Z_x}} \circ \mathrm{BN}^{\theta^{Z_x},2,i} \circ M_{\mathfrak{N}_2^{Z_x},\mathfrak{N}_1^{Z_x}}^{\theta^{Z_x},k\mathfrak{N}_2^{Z_x}(\mathfrak{N}_1^{Z_x}+3)+(k+1)\left(2N+\mathfrak{N}_1^{Z_x}(N+3)\right)} \\
& \circ \mathcal{R}_{\mathfrak{N}_1^{Z_x}} \circ \mathrm{BN}^{\theta^{Z_x},1,i} \circ M_{\mathfrak{N}_1^{Z_x},N}^{\theta^{Z_x},k\left(\mathfrak{N}_1^{Z_x}(N+3)+\mathfrak{N}_2^{Z_x}(\mathfrak{N}_1^{Z_x}+3)\right)+(k+1)2N} \circ \mathrm{BN}^{\theta^{Z_x},i}(X_{\tau_k}^m), \\
& \nu_{Z_x} = K\left(2N + \mathfrak{N}_1^{Z_x}(N+3) + \mathfrak{N}_2^{Z_x}(\mathfrak{N}_1^{Z_x}+3) + N^x(\mathfrak{N}_2^{Z_x}+1)\right); \\
\mathrm{N}_{\tau_k}^{\theta^{A},i} =& M_{N,\mathfrak{N}_2^{A}}^{\theta^{A},(k+1)\left(2N+\mathfrak{N}_1^{A}(N+3)+\mathfrak{N}_2^{A}(\mathfrak{N}_1^{A}+3)\right)} \circ \mathcal{R}_{\mathfrak{N}_2^{A}} \circ \mathrm{BN}^{\theta^{A},2,i} \circ M_{\mathfrak{N}_2^{A},\mathfrak{N}_1^{A}}^{\theta^{A},k\mathfrak{N}_2^{A}(\mathfrak{N}_1^{A}+3)+(k+1)\left(2N+\mathfrak{N}_1^{A}(N+3)\right)} \\
& \circ \mathcal{R}_{\mathfrak{N}_1^{A}} \circ \mathrm{BN}^{\theta^{A},1,i} \circ M_{\mathfrak{N}_1^{A},N}^{\theta^{A},k\left(\mathfrak{N}_1^{A}(N+3)+\mathfrak{N}_2^{A}(\mathfrak{N}_1^{A}+3)\right)+(k+1)2N} \circ \mathrm{BN}^{\theta^{A},i}(X_{\tau_k}^m), \\
& \nu_{A} = K\left(N(\mathfrak{N}_2^{A}+3) + \mathfrak{N}_1^{A}(N+3) + \mathfrak{N}_2^{A}(\mathfrak{N}_1^{A}+3)\right); \\
\mathrm{N}_{\tau_k}^{\theta^{\Gamma},i} =& M_{N^2,\mathfrak{N}_2^{\Gamma}}^{\theta^{\Gamma},(k+1)\left(2N+\mathfrak{N}_1^{\Gamma}(N+3)+\mathfrak{N}_2^{\Gamma}(\mathfrak{N}_1^{\Gamma}+3)\right)} \circ \mathcal{R}_{\mathfrak{N}_2^{\Gamma}} \circ \mathrm{BN}^{\theta^{\Gamma},2,i} \circ M_{\mathfrak{N}_2^{\Gamma},\mathfrak{N}_1^{\Gamma}}^{\theta^{\Gamma},k\mathfrak{N}_2^{\Gamma}(\mathfrak{N}_1^{\Gamma}+3)+(k+1)\left(2N+\mathfrak{N}_1^{\Gamma}(N+3)\right)} \\
& \circ \mathcal{R}_{\mathfrak{N}_1^{\Gamma}} \circ \mathrm{BN}^{\theta^{\Gamma},1,i} \circ M_{\mathfrak{N}_1^{\Gamma},N}^{\theta^{\Gamma},k\left(\mathfrak{N}_1^{\Gamma}(N+3)+\mathfrak{N}_2^{\Gamma}(\mathfrak{N}_1^{\Gamma}+3)\right)+(k+1)2N} \circ \mathrm{BN}^{\theta^{\Gamma},i}(X_{\tau_k}^m), \\
& \nu_{\Gamma} = K\left(N(2+N(\mathfrak{N}_2^{\Gamma}+1)) + \mathfrak{N}_1^{\Gamma}(N+3) + \mathfrak{N}_2^{\Gamma}(\mathfrak{N}_1^{\Gamma}+3)\right); \\
\mathrm{N}_{\tau_k}^{\theta^{r},i} =& M_{N^r,\mathfrak{N}_2^{r}}^{\theta^{r},(k+1)\left(2(N-N^x)+\mathfrak{N}_1^{r}((N-N^x)+3)+\mathfrak{N}_2^{r}(\mathfrak{N}_1^{r}+3)\right)} \circ \mathcal{R}_{\mathfrak{N}_2^{r}} \circ \mathrm{BN}^{\theta^{\Gamma},2,i} \circ M_{\mathfrak{N}_2^{r},\mathfrak{N}_1^{r}}^{\theta^{r},k\mathfrak{N}_2^{r}(\mathfrak{N}_1^{r}+3)+(k+1)\left(2(N-N^x)+\mathfrak{N}_1^{r}((N-N^x)\right.} \\
& \circ \mathcal{R}_{\mathfrak{N}_1^{r}} \circ \mathrm{BN}^{\theta^{\Gamma},1,i} \circ M_{\mathfrak{N}_1^{r},N-N^x}^{\theta^{r},k\left(\mathfrak{N}_1^{r}((N-N^x)+3)+\mathfrak{N}_2^{r}(\mathfrak{N}_1^{r}+3)\right)+(k+1)2(N-N^x)} \circ \mathrm{BN}^{\theta^{\Gamma},i}(X_{\tau_k}^m), \\
& \nu_{r} = K\left(2(N-N^x) + \mathfrak{N}_1^{r}((N-N^x)+3) + \mathfrak{N}_2^{r}(\mathfrak{N}_1^{r}+3) + N^r(\mathfrak{N}_2^{r}+1)\right).
\end{aligned}
$$

In evaluation, when no mini-batch moments $\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\sigma}}$ are available, these are replaced with their exponentially weighted averages across iterations $i \in \{1, \dots, K_{\text{batch}}\}$ defined as

$$
\hat{\hat{\boldsymbol{\mu}}}_{\tau_k}^{(j,l)} = \sum_{i=1}^{K_{\text{batch}}} \hat{\boldsymbol{\mu}}_{\tau_k}^{(j,l,)i} \left(\tilde{\beta}_1\right)^{K_{\text{batch}}+1-i}, \quad \hat{\hat{\boldsymbol{\sigma}}}_{\tau_k}^{(j,l)} = \sum_{i=1}^{K_{\text{batch}}} \hat{\boldsymbol{\sigma}}_{\tau_k}^{(j,l,)i} \left(\tilde{\beta}_2\right)^{K_{\text{batch}}+1-i},
$$

where the decay rates $\tilde{\beta}_1, \tilde{\beta}_2 \in (0,1)$ are non-trainable hyper-parameters. The trainable parameter vector is $\theta := [\theta^{Y_0}, \theta^{Z_0}, \theta^{Z_x}, \theta^A, \theta^\Gamma, \theta^r]^T \in \Theta = \mathbb{R}^\nu$, $\nu := \sum_j \nu_j$ for $j \in \{Y_0, Z_0, Z_x, A, \Gamma, r\}$ as specified above.

C.2 Adam optimizer with Learning rate schedule

The Adam (adaptive moment estimation) algorithm (Kinga et al., 2015) has become popular to make optimization faster, more stable, and better suited for deep learning. It is an adaption of gradient descent and combines *momentum*, i.e., descending by an exponentially weighted moving (leaky) average of the gradient $\boldsymbol{g}_i := D_{\theta_i}\Phi(\theta_i) \in \mathbb{R}^\nu$, and *adaptive* learning, i.e., scaling individual parameters' learning rate by a leaky average of its variance $\boldsymbol{g}_i^2 \in \mathbb{R}^\nu$ as a proxy for the scale of the corresponding element of the Hessian (Duchi et al., 2011). Specifically, given a hyperparametric initialization $\boldsymbol{v}_0, \boldsymbol{s}_0$, the leaky averages of the gradient and its variance are recursively defined for $i \in \{1, \ldots, K_{\text{batch}}\}$ as

$$\boldsymbol{v}_i = \tilde{\tilde{\beta}}_1 \boldsymbol{v}_{i-1} + (1 - \tilde{\tilde{\beta}}_1)\boldsymbol{g}_i; \qquad \boldsymbol{s}_i = \tilde{\tilde{\beta}}_2 \boldsymbol{s}_{i-1} + (1 - \tilde{\tilde{\beta}}_2)\boldsymbol{g}_i^2$$

such that the (rescaled) gradient update in equation (47) reads as

$$\theta_{i+1} = \theta_i - \lambda_i \left[\sqrt{\frac{\boldsymbol{s}_i}{1 - \left(\tilde{\tilde{\beta}}_2\right)^i}} + \varepsilon \right]^{-1} \frac{\boldsymbol{v}_i}{1 - \left(\tilde{\tilde{\beta}}_1\right)^i},$$

where the decay rates $\tilde{\tilde{\beta}}_1, \tilde{\tilde{\beta}}_2 \in (0,1)$ are again non-trainable hyper-parameters and the small number $\varepsilon = 10^{-6}$ (by custom) avoids division by zero. In addition to the adaptivity of the Adam optimizer, Beck et al. (2019) employ an exponentially discretely decaying learning rate (staircase) schedule such that

$$\lambda_i = \lambda_0 \left[\tilde{\lambda}\right]^{\left\lfloor \frac{i}{K_{\text{step}}} \right\rfloor},$$

where the initial learning and decay rates $\lambda_0, \tilde{\lambda} \in (0,1)$ as well as the decay step size $K_{\text{step}} \leq K_{\text{batch}}$ are non-trainable hyper-parameters.

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.

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.

John Duchi, Elad Hazan, and Yoram Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of machine learning research*, 12(7), 2011.

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.

Weinan E, Jiequn Han, and Arnulf Jentzen. Algorithms for solving high dimensional pdes: from nonlinear monte carlo to machine learning. *Nonlinearity*, 35(1):278, 2021.

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.

Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256. JMLR Workshop and Conference Proceedings, 2010.

Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep sparse rectifier neural networks. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 315–323. JMLR Workshop and Conference Proceedings, 2011.

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.

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

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

Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *Proceedings of The 32nd International Conference on Machine Learning (ICML)*, 2015.

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

Shaolin Ji, Shige Peng, Ying Peng, and Xichuan Zhang. Solving stochastic optimal control problem via stochastic maximum principle with deep learning method. *Journal of Scientific Computing*, 93(1):30, 2022.

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

D Kinga, Jimmy Ba Adam, et al. A method for stochastic optimization. In *International conference on learning representations (ICLR)*, volume 5, page 6. San Diego, California;, 2015.

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.

Serguei Maliar, Lilia Maliar, and Kenneth Judd. Solving the multi-country real business cycle model using ergodic set methods. *Journal of Economic Dynamics and Control*, 35(2):207–228, 2011.

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.

Etienne Pardoux and Shige Peng. Backward stochastic differential equations and quasilinear parabolic partial differential equations. In *Stochastic Partial Differential Equations and Their Applications*, volume 176, pages 200–217. Springer, 1992.

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*, 2019.

Elisabeth Proehl. Existence and uniqueness of recursive equilibria with aggregate and idiosyncratic risk. *SSRN Working Paper 3250651*, 2024.

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.

Jiongmin Yong and Xun Yu Zhou. *Stochastic controls: Hamiltonian systems and HJB equations*, volume 43. Springer Science & Business Media, 1999.