An Alternative Solution Method for Continuous-Time Heterogeneous Agent Models with Aggregate Shocks*

Nobuhide Okahata†

July, 2018

*** Preliminary and Incomplete Draft ***

Please Do Not Circulate without Author's Permission

Abstract

The increasing availability of micro data has led researchers to develop increasingly rich heterogeneous agent models. Solving these models involves nontrivial computational costs. The continuous-time solution method proposed by Ahn, Kaplan, Moll, Winberry, and Wolf (NBER Macroeconomics Annual 2017, volume 32) is dramatically fast, making feasible the solution of heterogeneous agent models with aggregate shocks by applying local perturbation and dimension reduction. While this computational innovation contributes enormously to expanding the research frontier, the essential reliance on the local linearization limits a class of problems researchers can investigate to the one where certainty equivalence with respect to aggregate shocks holds. This implies that it may be unsuitable for analyzing models where large aggregate shocks exist or nonlinearity matters. To resolve this issue, I propose an alternative solution method for continuous-time heterogeneous agent models with aggregate shocks by extending the Backward Induction method originally developed for discrete time models by Reiter (2010). The proposed method is nonlinear and global with respect to both idiosyncratic and aggregate shocks. I apply this method to solve a Krusell and Smith (1998) economy and evaluate its performance along two dimensions: accuracy and computation speed. I find that the proposed method is accurate even with large aggregate shocks and high curvature without surrendering computation speed (the baseline economy is solved within a few seconds). This new method is also applied to a model with recursive utility and an Overlapping Generations (OLG) model, and it is able to solve both models quickly and accurately.

Keywords: Heterogeneous agents, Backward induction, Continuous-time
JEL classification: C63, E24, E32

*The author especially appreciates Benjamin Moll, Sehyoun Ahn and Thomas Winberry for their helpful comments and suggestions. The author also appreciates Aubhik Khan for his several long discussions motivating this work. I also thanks valuable comments from Julia K. Thomas, Dean Corbae, Kyle Dempsey, Taisuke Nakata, Shin-ichi Nishiyama, Takayuki Tsuruga, Takeki Sunakawa and participants at workshops at the Ohio State University and Kobe University.

†Department of Economics, The Ohio State University. Email: nobuhide.okahata@gmail.com.

1 Introduction

Over the past decades, researchers have incorporated rich microeconomic heterogeneity empirically documented by micro data into macroeconomic models to derive new economic insights and policy implications. However, this still involves nontrivial computational costs.

The continuous-time solution methods for macroeconomic models based on the mean field game have eased many computational difficulties in solving heterogeneous agent models.[1] In particular, Ahn et al. (2017) first propose a continuous-time solution method for models with aggregate shocks. Their method linearizes the dynamic system around the steady state while simultaneously applying a dimension reduction method, and achieves dramatically fast and accurate computation of heterogeneous-agent business cycle models. While this computational innovation contributes enormously to expanding macroeconomic research frontiers, the essential reliance on linearization limits a class of problems researchers address to the one where certainty equivalence with respect to aggregate shocks holds.[2]

The purpose of this paper is to propose an alternative solution method for continuous-time heterogeneous agent models with aggregate shocks and to examine its accuracy and speed. Specifically, the proposed method applies the Backward Induction method originally developed by Reiter (2010) for discrete-time heterogeneous agent models[3] to a continuous-time setting. Compared with the linear approximation method by Ahn et al. (2017), the proposed method provides a global solution with respect to both idiosyncratic and aggregate shocks.

The proposed method involves three steps: in the first one, an infinite dimensional cross-sectional distribution of agents is replaced with a finite number of moments characterizing the distribution itself, over which an approximate value function and a Hamiltonian-Jacobi-Bellman (HJB) equation are defined. This approximation reduces the computational difficulty inherent in solving heterogeneous agent models with aggregate shocks: a cross-sectional distribution of agents, an infinite dimensional object, is included in the aggregate state variable of the value function and is computationally intractable. As a result, the standard numerical methods such as the finite difference method can be used to solve for the value function. However, before the actual computation, time derivatives of the moments

[1] See Achdou et al. (2017). The paper is one of the pioneering works to establish the continuous-time approach for solving heterogeneous agent macroeconomic models without aggregate shocks.

[2] Krusell (2017) also points out this in his comments on Ahn et al. (2017).

[3] The application of Reiter (2010) includes Gornemann et al. (2016), Khan (2017), and Kim (2017).

in the HJB equation also needs to be pinned down as a term in the HJB equation.

The second step consists of selecting a proxy distribution for the agents to be used for the computation of the time derivatives of moments. These are analogous to the perceived law of motion of the moments in discrete-time models. Krusell and Smith (1998) formulate the corresponding object as a linear function of aggregate state variables and update it through a costly simulation step. Instead, the proposed method explicitly computes the law of motion by integrating individual decision rules over the proxy distribution, thus generating a non-parametric function of the aggregate state variables. Moreover, the law of motion and the value function are updated simultaneously through a value function iteration. By avoiding the costly simulation step, this procedure allows a significant reduction in the computation time. The proxy distribution, defined over each point in the aggregate state, needs to satisfy a moment-consistency condition: the moments of the proxy distribution are equal to the values of the endogenous aggregate state variables in that point.

The third step is to refine and update the proxy distribution to make it reflect the dynamics of the model. For this purpose, the model economy is simulated to generate a sequence of the aggregate shock and an evolution of the distribution. Based on the simulation results, a reference distribution is constructed. As such distribution may not satisfy the moment-consistency condition, the updated proxy distribution is explicitly solved subject to this condition.

The merits of the proposed method compared to the linearized solution method can be characterized as follows: First, the proposed method can study the implications of aggregate shocks and risk aversion on the economic dynamics while the linearized solution method preserves certainty equivalence with respect to aggregate shocks. Second, the representation of the system through approximate aggregate states enables researchers to obtain a global solution easily while the linearized solution method evaluates the value function along the time instead of aggregate state variables, and this makes it challenging to implement a higher-order approximation of the model dynamics. Third, in relation to that the proposed method solves models globally, approximation errors are kept small even with relatively large aggregate shocks. Finally, while the linearized solution method imposes symmetric dynamics for both positive and negative exogenous aggregate shocks, the proposed method can capture these asymmetric responses to exogenous aggregate shocks of different signs.

I use three models — the standard Krusell and Smith (1998) model, a version of the Krusell and Smith model with Epstein-Zin Preferences, and an Overlapping Generations (OLG) model — as laboratories to evaluate the proposed method along two dimensions;

the accuracy of solutions and computational speed.

As a result, I find that the proposed method achieves high accuracy even when I consider the case with large aggregate shocks or the case with more risk averse households. For example, in a benchmark Krusell and Smith (1998) economy, the method achieves about 15 times higher accuracy as measured by den Haan (2010) statistics than the perturbation method proposed by Ahn et al. (2017) when the variance of the aggregate shock is large. Using the version of the Krusell and Smith (1998) model with Epstein-Zin preferences, I find that even when the household's risk aversion is high and the model is thus highly nonlinear, the method is still able to solve the model accurately.

Moreover, the proposed method can accommodate cases with nonlinear individual decision rules or where the wealth distribution in the economy exhibits such rich movement with respect to aggregate shocks that simple approximate aggregation or local projection do not work.[4] In a continuous-time OLG model with aggregate shocks, I find that the proposed method yields an accurate solution.

I also find that by building on the continuous-time formulation, the computation speed of the proposed method is still fast compared to the local linearization method. This speed advantage arises because we can avoid a simulation step to update the aggregate law of motion of the moments while simultaneously enjoying the speed gains brought about by the continuous-time formulation. For example, in the fastest case, the proposed approach can solve the Krusell and Smith (1998) model in about a second in Matlab without any parallelization or code tuning. This speed gain enables researchers to tackle projects with a high computational burden of high-dimensional heterogeneous agent macroeconomic models such as OLG models, multi-country open economy models, or full likelihood estimation of heterogeneous agent business cycle models.

Given these features, while the linearized solution method in continuous time still provides a powerful and useful tool for many problems, the proposed method could be best characterized as a complement to the existing toolbox of continuous-time numerical solution methods.

The rest of the paper is organized as follows. In Section 2, I set up the model economy. In Section 3, I lay out numerical settings and algorithm. Section 4 reports results of the baseline economy and compares results with Ahn et al. (2017). Section 5 reports the setup and results for the extended examples. Section 6 concludes.

[4]For example, Krueger and Kubler (2004) argues that approximate aggregation and linear forecasting rules as in Krusell and Smith (1998) do not yield an accurate solution to OLG models.

Related Literature

In this subsection, I summarize the recent developments in solution techniques to heterogeneous agent models with aggregate shocks and discuss their relationships with the proposed method.

One of the popular approaches to solving heterogeneous agent macroeconomic models is the projection and perturbation method by Reiter (2009).[5] This method solves models globally in individual states, but only provides local approximations in aggregate states.[6] Ahn et al. (2017) apply this approach to continuous-time models. Winberry (2018) instead approximates the distribution with a flexible parametric family and reduces the dimensionality by considering a finite set of parameters. Then, he perturbs the parameterized system for the aggregate dynamics to solve and estimate a lumpy investment model in discrete time. On the other hand, Anderson et al. (2012), Evans (2015), and Bhandari et al. (2017) propose a method of taking Taylor expansions repeatedly around arbitrary points in the state space, not just around the steady state. This approach adopts the perturbation theory developed by Fleming (1971) and Fleming and Soner (1986) using small noise expansions with respect to a one-dimensional parameterization of uncertainty. This method also yields a global solution with respect to the aggregate shocks around an arbitrary point in the state space, but this essentially local expansion method could lead to difficulties in handling binding borrowing constraints and non-convexities in individual decision rules in a discrete time setting. By comparison, as discussed in Achdou et al. (2017), the continuous-time formulation in this paper avoids these problems.

Another stream of research proposes alternative algorithms that do not rely on perturbation, but rather replace the computationally costly stochastic simulation algorithm in Krusell and Smith (1998). For example, den Haan et al. (2010) propose the Explicit Aggregation algorithm, which calculates a law of motion of the aggregate capital by explicitly integrating individual saving decision rules. Young (2010) proposes a Non-stochastic Simulation approach that uses policy functions (together with their interpolation on the individual state space) and reallocates the mass in distribution.[7] The method proposed

[5]For one of the earliest application of the local perturbation method to heterogeneous agent models, see Campbell (1998) for firm heterogeneity. For the recent application of the linearization methods to the model of household heterogeneity, see McKay and Reis (2016). For the application to the model with firm heterogeneity, see Mongay and Williams (2016).

[6]For a related but different approach, Childers (2018) instead develops a perturbation theory over a functional space to analyze the dynamics of infinite dimensional objects.

[7]See Terry (2017) for a comprehensive study on the properties of these different solution methods applied to the lumpy investment model by Khan and Thomas (2008).

in this paper, which is an application of Reiter (2010), can be classified into this stream of research as it also relies on an alternative algorithm to avoid a simulation step by using a proxy distribution to construct a law of motion of aggregate states. Compared to the other methods, the one by Reiter (2010) can accommodate flexible and general representations of distributions without relying on the parameterization of the distribution or specific model assumptions. Extending Reiter (2010) to continuous-time settings provides researchers with a further computational advantage while it can still handle a distribution as flexibly as in discrete time.

Apart from the specific numerical solution methods described above, an appropriate choice between local and linear methods on the one hand and global and nonlinear methods on the other is still a nontrivial issue for researchers. Recent work by Boppart et al. (2017) proposes to utilize the linearity of the impulse response functions with respect to aggregate shocks when researchers are interested in perfect foresight or linearized solution. This linearity property is satisfied when the model to be analyzed is a suitable candidate for the perfect foresight method or linearization solution. Otherwise, a global and nonlinear solution would be required. The simple implementation helps researchers quickly choose the appropriate numerical solution methods between linearization techniques such as Ahn et al. (2017) and global and nonlinear methods such as the proposed method in this paper.

2 Baseline Model and the Backward Induction Method

The first part of this section is devoted to the description of a baseline model. To compare the results with ones by Ahn et al. (2017), I follow their formulation and parameterization.[8] In the second part of the section, the Backward Induction method adapted to the continuous-time environment is illustrated.

2.1 Environment

Assume a continuum of households with an infinite planning horizon in the economy. At every instance at time $t \geq 0$, they consume $c(t) \geq 0$, and either borrow or save $a(t) \subseteq \mathbb{R}$. They gain the interest $r(t)a(t)$ if $a(t) \geq 0$ or pay it for their borrowings if $a(t) < 0$. This paper considers an incomplete market environment, where households can only trade non-contingent bonds, and their borrowings are subject to the borrowing constraint: $a(t) \geq -\underline{a}$,

[8] For the parameterization, Ahn et al. (2017) follows den Haan et al. (2010).

with $\underline{a} \geq 0$. They supply a fixed amount of labor $l(t) = \bar{l}$ and draw an idiosyncratic labor productivity $z(t)$, whose stochastic process[9] follows a two-state Poisson process with arrival rates λ_L and λ_H, and $z(t) \in \{z_L, z_H\} = \mathbf{z} \subseteq \mathbb{R}^+$, where $z_L < z_H$. They inelastically supply an efficiency unit of labor $z(t)\bar{l}$ at each point in time to a firm, and earn the labor income $w(t)z(t)\bar{l}$. The population ratios of people with z_L and z_H are defined as μ_L and μ_H, respectively, implying that the average labor productivity is $\bar{z} = \mu_L z_L + \mu_H z_H$. The household with z_L receives an unemployment insurance $\bar{b}$ financed through the labor income tax τ. Household savings and effective labor are used as inputs in the representative firm's production function. The instantaneous utility function $u(\cdot)$ is a CRRA function, and the discount rate is $\rho \geq 0$.

Households maximize their expected lifetime utilities subject to the budget constraints:

$$\begin{aligned} v = \max_{\{c(t)\}_{t=0}^{\infty}} \mathbb{E}_0 \int_0^{\infty} e^{-\rho t} u(c(t)) dt \qquad (1) \\ \text{s.t.} \quad \dot{a}(t) = \left[r(t)a(t) + (1-\tau)w(t)z(t)\bar{l} + \mathbb{I}_z \bar{b} - c(t) \right], \quad a(t) \geq \underline{a} \\ \text{where} \quad u(c(t)) = \frac{c(t)^{1-\sigma}}{1-\sigma}, \quad \sigma \geq 0 \end{aligned}$$

v is the value for the household, and $\mathbb{E}_0$ is the expectation operator over the realizations of idiosyncratic and aggregate shocks from time $t = 0$ onward, while σ is the utility curvature parameter. $\dot{a}(t) \equiv da(t)/dt$ represents savings, and $\mathbb{I}_z$ is an indicator function taking value 1 if $z(t) = z_L$ and 0 otherwise.

There is a representative firm in the economy that produces an output $Y(t)$ using capital $K(t)$ and labor $L(t)$ at every t according to a Cobb-Douglas production function: $Y(t) = A(t)K(t)^{\alpha}L(t)^{1-\alpha}$, with $0 \leq \alpha \leq 1$. $A(t) \in \mathbb{R}^+$ is the aggregate productivity, and $Z(t) = \log(A(t))$ follows an Ornstein-Uhlenbeck (OU) process:

$$dZ(t) = -\eta Z(t)dt + \sigma dW(t) \qquad (2)$$

where $dW(t)$ is the innovation to a standard Brownian motion, η is the mean reversion parameter, while σ is the volatility parameter. Note that, alternatively, a Poisson process for $A(t)$ could be used. I also consider the case where $A(t)$ follows a two-state Poisson process[10] with arrival rates Ω_L and Ω_H, and $A(t) \in \{A_L, A_H\} = \mathbf{A} \subseteq \mathbb{R}^+$, with $A_L < A_H$.

[9] The idiosyncratic shock process is independent of aggregate shocks, while it is not the case in the original Krusell and Smith (1998) model. For the comparison with Ahn et al. (2017), I use their formulation.

[10] In this paper, I adopt both the OU process used by Ahn et al. (2017) and a two-state Poisson process.

The results for this specification are reported in Section 4 and 5.

Each household is characterized by different levels of net assets $a(t)$ and idiosyncratic productivity $z(t)$. $\mathbf{g}_t$ is defined as a probability distribution function of households at time t over the individual state variables $(a(t), z(t)) \subseteq [\underline{a}, \infty) \times \mathbf{z}$. This $\mathbf{g}_t$ adds up to 1, corresponding to the normalized measure of the overall population in the economy:

$$\begin{aligned} &\mathbf{g}_t(a(t), z(t)) \geq 0 \quad \text{for all } a(t) \in [\underline{a}, \infty), \text{ and } z(t) \in \mathbf{z} \\ \textit{and} \quad &\int \sum_{z(t)\in\mathbf{z}} \mathbf{g}_t(a(t), z(t))da = 1 \end{aligned} \tag{3}$$

Moreover, the evolution of distribution $\dot{\mathbf{g}}_t$ is obtained from a function $\mathbf{G}$ induced by the of households' optimal saving decisions:

$$\dot{\mathbf{g}}_t = \mathbf{G}(\mathbf{g}_t, Z(t)) \tag{4}$$

The interest and wage rates are competitively determined to clear the factor markets:

$$r(t) = \alpha e^{Z(t)} K(t)^{\alpha-1} L(t)^{1-\alpha} - \delta, \quad w(t) = (1-\alpha) e^{Z(t)} K(t)^{\alpha} L(t)^{-\alpha} \tag{5}$$

where $\delta > 0$ is the depreciation rate.
The market clearing conditions for the capital and labor markets are, respectively:

$$K(t) = \int \sum_{z(t)\in\mathbf{z}} a(t)\mathbf{g}_t(a(t), z(t))da, \quad L_t = \int \sum_{z(t)\in\mathbf{z}} z(t)l(t)\mathbf{g}_t(a(t), z(t))da = \bar{z}\bar{l} \tag{6}$$

In addition, the government budget constraint satisfies:

$$w(t)\bar{b} \int \mathbf{g}_t(a(t), z_L)da + G^s(t) = \tau \int \sum_{z(t)\in\mathbf{z}} w(t)z(t)l(t)\mathbf{g}_t(a(t), z(t))da \tag{7}$$

where $G^s(t)$ is the government spending[11]

The equations (1) - (7) fully characterize the dynamics of the system.

In the following two subsections, in order to give motivation toward establishing the proposed numerical solution via the Backward Induction method, I first describe the state

For the latter case, I calibrate the process to the same persistence and variance as the OU process.

[11] In the numerical exercise in Section 4, $G^s(t)$ is set to 0 for all t. In a general case where the idiosyncratic shock depends on the aggregate shock, $\mu_L$varies over time and $G^s(t)$ should be adjusted so that equation (7) holds.

space representation of the dynamic system and illustrate computational challenge in the model with aggregate shocks. Then, I briefly illustrate the linearization approach by Ahn et al. (2017) and clarify the difference between this method with the proposed method.

State Space Representation[12]

I reformulate the dynamic problem above to a state space representation using the value function defined over individual states and aggregate states, the Hamilton-Jacobi-Bellman (HJB) equation, and the Kolmogorov forward equation. In this case, the HJB equation includes an infinite dimensional object $\mathbf{g}$ and its Frechét derivative, which are computationally intractable. Thus, as in discrete-time modes, some approximation is needed.

In a state space representation, the equilibrium conditions of the dynamic system are expressed in terms of individual states (a, z), where $a \in [\underline{a}, \infty)$ and $z \in \mathbf{z}$, and aggregate states $(\mathbf{g}, Z)$, where $\mathbf{g}$ is the distribution of households and Z is the aggregate shock.

Let us define the interest and wage rates as functions of the aggregate state $(\mathbf{g}, Z)$.

$$
\begin{aligned}
&r(\mathbf{g}, Z) = \alpha e^{Z} K(\mathbf{g})^{\alpha-1} L(\mathbf{g})^{1-\alpha} - \delta, \quad w(\mathbf{g}, Z) = (1-\alpha) e^{Z} K(\mathbf{g})^{\alpha} L(\mathbf{g})^{-\alpha} \\
&K(\mathbf{g}) = \int \sum_{z \in \mathbf{z}} a \mathbf{g}(a, z) da \\
&L(\mathbf{g}) = \int \sum_{z \in \mathbf{z}} z \bar{l} \mathbf{g}(a, z) da = \bar{z} \bar{l}
\end{aligned}
$$

Then, let us define the Kolmogorov forward operator K_z which operates on $\mathbf{g}$:

$$
\frac{d\mathbf{g}_t(a, z)}{dt} = (K_z g_t)(a, z) \equiv -\partial_a [s(a, z; \mathbf{g}_t, Z) \mathbf{g}_t(a, z)] - \lambda_z \mathbf{g}_t(a, z) + \lambda_{z'} \mathbf{g}_t(a, z')
$$

where $s(a, z; \mathbf{g}_t, Z)$ is a saving policy function for the household with individual state (a, z) and aggregate state $(\mathbf{g}_t, Z)$, λ_z takes value λ_H if $z = z_H$ and λ_L if $z = z_L$, and $z, z' \in \mathbf{z}$ is such that $z \neq z'$. Here, K_z corresponds to the function $\mathbf{G}$ in the previous subsection. More specifically, given $\mathbf{g}_t$, K_z generates the time derivative of $\mathbf{g}_t$, $d\mathbf{g}_t/dt$, and associates it with the households' optimal saving decisions $s(a, z; \mathbf{g}_t, Z)$ and the stochastic process of z (i.e. Kolmogorov forward equation). This corresponds to a law of motion of a distribution in discrete time, and it expresses the net flow of households over (a, z).

With these definitions, I derive the HJB equation where the aggregate productivity follows an OU process. Let $v(a, z; \mathbf{g}, Z)$ be a value function with individual state (a, z) and

[12] The notation and the explanation in this part follow those in Appendix A of Ahn et al. (2017).

aggregate state $(\mathbf{g}, Z)$. Then, the HJB equation is defined as follows:

$$\begin{aligned}\rho v(a,z;\mathbf{g},Z) = \max_c\ & u(c) + \partial_a v(a,z;\mathbf{g},Z)\ [r(\mathbf{g},Z)a + (1-\tau)w(\mathbf{g},Z)z\bar{l} + \mathbb{I}_z\bar{b} - c] \\ & + \lambda_z[v(a,z';\mathbf{g},Z) - v(a,z;\mathbf{g},Z)] \\ & + \partial_Z v(a,z;\mathbf{g},Z)(-\eta Z) + \frac{1}{2}\partial_{ZZ}v(a,z;\mathbf{g},Z)\sigma^2 \\ & + \int \sum_{z'\in\mathbf{z}} \frac{Dv(a,z;\mathbf{g},Z)}{D\mathbf{g}(a,z)} K_z(\mathbf{g})(a,z')da \end{aligned} \tag{8}$$

for all $z, z' \in \mathbf{z}$ such that $z \neq z'$. Moreover, $Dv(a,z;\mathbf{g},Z)/D\mathbf{g}(a,z)$ is a functional derivative of v with respect to $\mathbf{g}$ at the individual state (a,z).

Time-Dependent Representation

As $\mathbf{g}$ and $\frac{Dv(a,z;\mathbf{g},Z)}{D\mathbf{g}(a,z)}$ are infinite-dimensional objects, the HJB equation (8) is computationally intractable. In order to deal with this problem, Ahn et al. (2017) adopt a time-dependent representation of the dynamic system by considering a time-dependent value function $v(a,z;t)$ instead of the value function $v(a,z;\mathbf{g},Z)$ in the state space representation. Accordingly, the HJB equation becomes:

$$\begin{aligned}\rho v(a,z;t) = \max_{c(t)}\ & u(c(t)) + \partial_a v(a,z;t)[r(t)a + (1-\tau)w(t)z\bar{l} + \mathbb{I}_z\bar{b} - c(t)] \\ & \lambda_z[v(a,z';t) - v(a,z;t)] + \frac{1}{dt}\mathbb{E}_t[dv(a,z;t)]\end{aligned} \tag{9}$$

for all $z, z' \in \mathbf{z}$ such that $z \neq z'$ and $t \geq 0$.
Moreover, the corresponding Kolmogorov forward equation is:

$$\frac{d\mathbf{g}_t(a,z)}{dt} = -\partial_a[s(a,z;t)\mathbf{g}_t(a,z)] - \lambda_z\mathbf{g}_t(a,z) + \lambda_{z'}\mathbf{g}_t(a,z') \tag{10}$$

where $s(a,z;t)$ is a saving policy function for the household with individual state (a,z) at time t obtained by solving (9).

The dynamics are characterized by the HJB equation (9) and the Kolmogorov Forward equation (10), the aggregate shock process (2), and the market clearing conditions (3), (5), (6), and the government budget constraint (7). They study the dynamics by applying first-order approximation of the system around a steady state with respect to the aggregate state variables. Moreover, they further apply a dimension reduction method to describe the movement of the cross-sectional distribution $\mathbf{g}_t$ through the movements of bases of smaller dimension.

Ahn et al. (2017) show how to combine the perturbation method for discrete-time settings by Reiter (2009) with powerful continuous-time solution methods. However, the implementation of higher-order approximations of the dynamics is not straightforward in their approach. Instead, it may be possible to construct the dynamic system by using aggregate shocks and the bases themselves as aggregate state variables and then take higher-order approximation. However, this approach still suffers from substantive computation costs if a dimension of bases is large.[13] Thus, I instead turn to an approximate aggregate state space representation and apply the Backward Induction method by Reiter (2010) to propose an alternative, nonlinear, and global solution method.

2.2 The Backward Induction Method

In this subsection, I explain the proposed method that applies the Backward Induction method originally developed by Reiter (2010) in a discrete-time setting to continuous-time models. More in detail, the proposed method involves three steps: the first step consists of replacing an infinite dimensional cross-sectional distribution of agents with a finite number of moments characterizing the distribution itself and constructing the approximate value function and the HJB equation defined over these moments. The second step is to select a proxy distribution over the aggregate state space and use it to derive the law of motion of the moments. The third step involves refining and updating the proxy distribution so to reflect the model dynamics. Hereafter, I provide more detailed explanations on the sub-components of the proposed method.

(1) Approximate Aggregate State Space Representation

Differently from the linearized solution method by Ahn et al. (2017), the proposed method replaces an infinite dimensional distribution of households in the aggregate state space with a vector of finite numbers of statistics $\mathbf{m}$ (the "moments") characterizing the distribution. These approximate aggregate state variables are used to build the approximate HJB equation, with the value function defined over the approximate aggregate state space.[14] The

[13]Another possible strategy would be to use a parameterized distribution and take higher-order approximation of the system. Indeed, Winberry (2018) reports a second-order perturbation of the heterogeneous firm business cycle model of Khan and Thomas (2008). However, as discussed in Winberry (2018), this approach also suffers from huge computational costs when the a model has high dimensional state space.

[14]The proposed method can be essentially viewed as an application of approximate aggregation to continuous-time heterogeneous agent models in the spirit of Krusell and Smith (1998).

solution to this equation provides the global solution for both approximate aggregate and individual state variables.

Let $\mathbf{m}$ be a vector of finite statistics as a function of the cross-sectional distribution of households $\mathbf{g}(a, z)$. In this paper, the aggregate capital is used for the moments $\mathbf{m}$ ($\mathbf{m} = K$), and the explanation of the method is adjusted accordingly. While the examples in this paper contain just one variable for $\mathbf{m}$, it is generally possible to include more than one such as higher-order moments of the distribution.[15]

The proposed method uses the moments $\mathbf{m}$ to define the approximate aggregate state space $(\mathbf{m}, Z)$ instead of the true aggregate state space $(\mathbf{g}, Z)$ and to construct the approximate dynamic system over $(\mathbf{m}, Z)$. Consequently, the aggregate state variables $(\mathbf{g}, Z)$ in equation (8) are replaced with $(\mathbf{m}, Z)$. Moreover, the last term in (8), which is also a computationally intractable object, becomes the usual derivative of the value function with respect to the moment $\mathbf{m}$ multiplied by the time derivative of the moment $\dot{\mathbf{m}}$, and does not include the infinite-dimensional functional derivative. The approximate HJB equation to be solved numerically is:

$$\begin{aligned}\rho v(a, z; \mathbf{m}, Z) = \max_c \ & u(c) + \partial_a v(a, z; \mathbf{m}, Z) \cdot [r(\mathbf{m}, Z)a + (1-\tau)w(\mathbf{m}, Z)z\bar{l} + \mathbb{I}_z \bar{b} - c] \\ & + \lambda_z [v(a, z'; \mathbf{m}, Z) - v(a, z; \mathbf{m}, Z)] + \partial_m v(a, z; \mathbf{m}, Z) \cdot \frac{\partial \mathbf{m}}{\partial t} \\ & + \partial_Z v(a, z; \mathbf{m}, Z)(-\eta Z) + \frac{1}{2}\partial_{ZZ} v(a, z; \mathbf{m}, Z)\sigma^2 \end{aligned} \tag{11}$$

for all $z, z' \in \mathbf{z}$ such that $z \neq z'$.

The dynamics are now characterized by the HJB equation (11), the Kolmogorov forward equation (10), the aggregate shock process (2), the market clearing conditions (3), (5), and (6), and the government budget constraint (7).

At this point, in order to apply the standard numerical methods such as the finite difference method, the time derivative of the moment $\dot{\mathbf{m}}$ needs to be pinned down.

(2) Computing the Law of Motion of the Moment and Proxy Distribution

Law of Motion of the Moment: the time derivative of the moment $\dot{\mathbf{m}}$ is analogous to the perceived law of motion of the moment in discrete-time heterogeneous agent models.

[15] This point is also mentioned by Reiter (2010) for discrete-time models.

Here, $\dot{\mathbf{m}}$ is assumed to be a function of the approximate aggregate state $(\mathbf{m}, Z)$:

$$\dot{\mathbf{m}} = \hat{\mathbf{G}}(\mathbf{m}, Z) \tag{12}$$

for some function $\hat{\mathbf{G}}$. For example, in discrete time, Krusell and Smith (1998) formulate the corresponding function of $\hat{\mathbf{G}}$ as a linear function of the approximate aggregate state $(\mathbf{m}, Z)$.

The proposed method instead follows the different approach in formulating $\hat{\mathbf{G}}$. As in discrete-time models, the aggregate law of motion needs to be consistent with optimal choices of households obtained by solving (11). As the aggregate capital is now chosen as an approximate state variable, the function $\hat{\mathbf{G}}$ satisfying this consistency is thus given by integrating the individual saving decision rules over some probability distribution corresponding to a cross-sectional distribution of households over (a, z):

$$\frac{\partial \mathbf{m}}{\partial t} = \int \sum_{z \in \mathbf{z}} s^*(a, z; \mathbf{m}, Z)\mathbf{g}^p(a, z; \mathbf{m}, Z)da \tag{13}$$

where $s^*(a, z; \mathbf{m}, Z)$ is a saving policy function obtained by solving (11) and $\mathbf{g}^p(a, z; \mathbf{m}, Z)$ is some probability distribution function of a cross-sectional distribution of households over (a, z) given $\mathbf{m}$ and z such that:

$$\begin{aligned} & \mathbf{g}^p(a, z; \mathbf{m}, Z) \geq 0 \text{ for all } a \in [\underline{a}, \infty), \text{ and } z \in \mathbf{z} \\ \textit{and} \quad & \int \sum_{z \in \mathbf{z}} \mathbf{g}^p(a, z; \mathbf{m}, Z)da = 1 \end{aligned} \tag{14}$$

This $\mathbf{g}^p$ is specifically called proxy distribution. The detailed explanation and the additional condition are discussed in the next subsection.

Additionally, $\hat{\mathbf{G}}$ in the proposed method is a nonparametric function of the aggregate state variables, and it allows a more flexible representation of the law of motion compared to the linear forecasting rule in Krusell and Smith (1998).

Proxy Distribution: the proxy distribution is defined over each point in the aggregate state space. Let us define a function called distribution selection function (DSF). The DSF maps a point in the aggregate space $(\mathbf{m}, Z)$ to a cross-sectional distribution of households over (a, z) which is a proxy distribution $\mathbf{g}^p(a, z; \mathbf{m}, Z)$:

$$DSF : (\mathbf{m}, Z) \rightarrow \mathbf{g}^p(a, z; \mathbf{m}, Z) \tag{15}$$

In other words, the DSF defines a family of cross-sectional distributions parameterized by $(\mathbf{m}, Z)$. As explained in the previous subsection, the proxy distribution is used to compute the law of motion of the moments in equation (13). In addition to condition (14) for $\mathbf{g}^p$ to be a probability distribution function, the proxy distribution needs to satisfy a cross-sectional moment consistency condition, that is, the statistics characterizing $\mathbf{g}^p(a, z; \mathbf{m}, Z)$ have to be consistent with $\mathbf{m}$ for every $(\mathbf{m}, Z)$:

$$\int \sum_{z \in \mathbf{z}} a\mathbf{g}^p(a, z; \mathbf{m}, Z)da = \mathbf{m} \tag{16}$$

Furthermore, the marginal probability of z equals to the population ratios of people with z:

$$\int \mathbf{g}^p(a, z_L; \mathbf{m}, Z)da = \mu_L, \quad \int \mathbf{g}^p(a, z_H; \mathbf{m}, Z)da = \mu_H \tag{17}$$

(3) Main Routine

The main routine of the proposed method simultaneously solves for the value function and the law of motion of the moments. This main routine involves following two computation procedures. First, given the distribution selection function, the proxy distribution and $\dot{\mathbf{m}}$, equation (11) is solved at each point in $(\mathbf{m}, Z)$ to obtain the updated value function and the saving policy function $s^*(a, z; \mathbf{m}, Z)$. Then, $\dot{\mathbf{m}}$ is computed and updated according to equation (13) given the newly obtained s^* and the proxy distribution. These steps are iterated until $v(a, z; \mathbf{m}, Z)$ and $\dot{\mathbf{m}} = \hat{G}(\mathbf{m}, Z)$ converge. Solving this nested fixed point problem provides the value and policy function defined over both the individual and the (approximate) aggregate state spaces, as well as the law of motion of the moment.

One of the advantages of using the Backward Induction method is that this approach can avoid the costly simulation step to solve for the stochastic equilibria. By combining this advantage with the benefits of continuous-time formulation, the proposed method achieves a significant reduction in the computation time, which is comparable to the one of the linearized solution method by Ahn et al. (2017).

(4) Refinement of the Proxy Distribution

The main routine is computed given a specific distribution selection function and a family of proxy distributions. Based on the results of the main routine, the choice of the proxy distribution may be refined. Specifically, using the value and policy functions and the law of motion of the moment obtained, the model economy is simulated to generate time series of aggregate capital, aggregate productivity shock, and the distribution of households. This data is used to construct a new family of distributions of households over the individual state space given $(\mathbf{m}, Z)$. These distributions are called reference distributions. While these reference distributions reflect the model solution, they are not necessarily moment consistent. Thus, an additional step is required to solve for a moment-consistent proxy distribution, which is used to compute step (3) again. Then, steps (3) and (4) are iterated until the predetermined criteria for the accuracy are satisfied.

In the next section, I discuss the numerical implementation of the proposed method.

3 Numerical Implementation

This section discusses the parameterization of the baseline model. The computation settings and the related algorithm are also described.

3.1 Parameterization

Table 1 summarizes the parameter values used in the numerical experiments. In order to make a comparison between the linearized solution method in Ahn et al. (2017) and the proposed method, these are mostly taken from their paper and the published replication codes. Their work also refers to the parameter values in den Haan et al. (2010). The unit of the model time period is a quarter. The rate of time preference parameter ρ is set to 0.01, while the coefficient of relative risk aversion σ is set to 1.0. As for the rate of capital depreciation, I select $\delta = 0.025$, with a capital share $\alpha = 0.36$. As noted in Ahn et al. (2017), I choose the same levels of idiosyncratic labor productivity z_L and z_H as in den Haan et al. (2010). For the aggregate productivity parameters, the same values as in Ahn et al. (2017) are chosen in the OU process case. Given the quarterly persistence of the aggregate shock of $corr(\log A_{t+1}, \log A_t) = 0.75$, which is approximately the same as in den Haan et al. (2010), the persistence parameter η is set to 0.25. The volatility parameter σ is set to 0.007 (0.7%) in the baseline model. In Section 4, the numerical results based on

Table 1: Parameter Values for the Baseline Model

Parameters	**Calibrated Values**
σ : Relative Risk Aversion	1.0
ρ : Rate of time preference	0.01
δ : Rate of capital depreciation	0.025
α : Capital Share	0.36
z: Idiosyncratic Productivity	$z_L = 0.0$ and $z_H = 1.0$
η : Persistence of the OU process	0.25
σ : Volatility of the OU process	0.007
$\bar{b}$: Social Security	0.15
$\underline{a}$: Borrowing constraint	0.0

different values of the volatility parameter are reported. In the Poisson process case, the Tauchen method is used to construct the corresponding two-state Markov process, which is then converted to a Poisson process in continuous time, given the same degree of the quarterly persistence and the volatility of the aggregate shock as in the case with the OU process. Thus, the values for the aggregate productivity A_L and A_H vary depending on the parameters characterizing the stochastic process for the aggregate productivity shock.

As for the parameters of the social security transfer and the borrowing constraint, I choose the same values as in den Haan et al. (2010) and Ahn et al. (2017).

3.2 Numerical Settings and Algorithm

In this subsection, I illustrate the computation settings and the numerical algorithm of the proposed method. For the computation grids, equispaced grids are used for the state variables in the baseline case. As for the individual endogenous state variable, the net asset holdings of the household a, I choose a grid $\{a_i\}_{i \in I}$ with $I = 100$, $a_1 = 0$ and $a_{100} = 100$. For the aggregate endogenous state variable, the aggregate capital in this paper, I choose a grid $\{m_j\}_{j \in J}$ with a total number of points J and the support of the grid reflecting model parameters and dynamics.[16] For the aggregate shock Z that follows the OU process, the equispaced grid $\{Z_n\}_{n \in N}$ is chosen, while the idiosyncratic shock takes two values (high or low) following the Poisson process. For the tolerance parameters for the value function and $\dot{\mathbf{m}}$ updating in the main routine, 10^{-5} and 10^{-6} are selected respectively.

When solving the partial differential equation and implementing value function itera-

[16] Robustness of the results are checked under the different discretization schemes.

tion, I adopt the finite difference method and apply the implicit method.[17] Furthermore, in order to calculate the partial derivative of the value function, I apply the upwind method depending on the signs of drifts for state variables.

I summarize the numerical algorithm for the proposed method below.

Outline of the Algorithm

- **(a) Construct the Proxy Distribution**

- **(b) Solve for the Value Functions and the Law of Motion of the Moments**

  1. Given the proxy distribution and the law of motion of the moments, solve for an updated value function and a policy function using the finite difference method.
  2. Given the proxy distribution and the updated policy functions, update the law of motion of the moments for each grid point of $(\mathbf{m}, Z)$
  3. Iterate steps 1 and 2 until the value function and the law of motion of the moments converge.

- **(c) Simulation and Refinement of the Proxy Distribution**

  1. Simulate the model economy for T periods. Discretize the time dimension by a step length dt. Then, there are T^s time steps for the simulation.[18] Given the initial aggregate capital and the distribution of households, update the distribution using the Kolmogorov forward equation at each time step.[19] Save the time series data for aggregate shock, aggregate capital, and distribution.
  2. Construct the reference distributions as weighted averages of the simulated distributions of households given the realization of the aggregate shocks.
  3. Update the Proxy distribution by solving a constrained optimization problem under cross-sectional moment consistency conditions.

- **(d) Repeat Step (b) and Step (c) until Convergence.**

Next section reports the results of the numerical experiments.

[17] The explicit method is also used. The numerical solution converges if CFL condition is satisfied. and the results are not changed.

[18] The length of a time step is set to $dt = 0.25$ for both the linearized solution case and the proposed method case. Thus, the total simulation time T equals to $T^s \times dt$.

[19] The implicit method is used to update the distribution at each time step numerically.

Figure 1: Simulated Paths of the Aggregate Capital $K(t)$

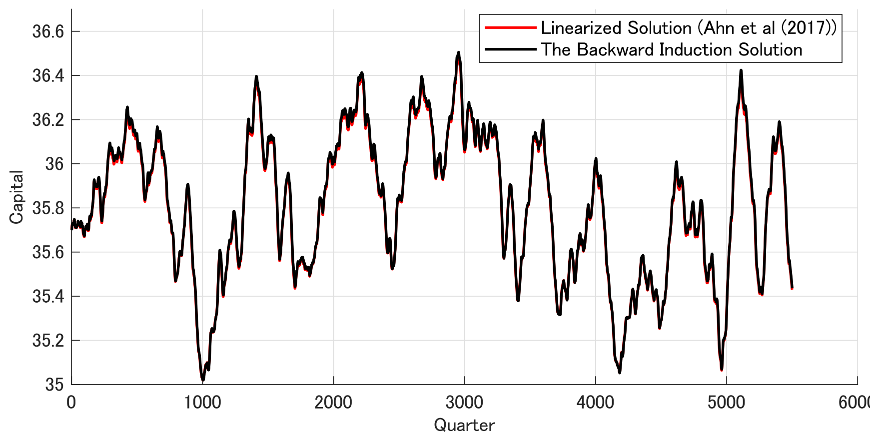

Notes: The simulated paths of the aggregate capital of the Baseline model which is solved by two different methods. The red line corresponds to the simulated path of $K(t)$ obtained by the linearized solution method in Ahn et al. (2017). The black line corresponds to the simulated $K(t)$ obtained by the proposed solution method in this paper. The initial capital, aggregate shock realizations, and parameters are the same in both the linearized solution case and the Backward Induction case. The total simulation steps are 5,500.

4 Results

In this section, I outline the main numerical results from the baseline model, by also comparing them to those obtained through the linearized solution method by Ahn et al. (2017).[20]

Accuracy Measure

To check for the accuracy of the numerical methods, I adopt the maximum approximation error metric proposed by den Haan (2010). The error metric is calculated by generating two time series of aggregate capital stocks from the simulation: the first one is the multi-step forecast capital $\{K_t\}_{t\in[0,T]}$ generated from equation (13) for all t, while the second one corresponds to the simulated $\{K_t^*\}_{t\in[0,T]}$ obtained by solving equations (11) and (13) and updating the distribution according to the individual households' policy functions for each simulation step.

[20]For generating the results from the application of the linearized solution method, I reuse and modify the replication codes for Ahn et al. (2017) published in the GitHub repository

Table 2: Comparison of Maximum den Haan Error (%)

St. Dev Shocks (%)	Ahn et al. (2017)	Proposed Method (%)
0.1	0.001	0.005
0.7	0.049	0.026
1.0	0.118	0.043
5.0	3.282	0.22

Notes: The second column replicates the error statistics in Table 2 in Ahn et al. (2017) which applies the linearized solution method. The third column reports the error statistics of solving the baseline models with the proposed method. Refer to the main text for the definition and calculation of den Haan statistics.

Thus, the approximation error metric ϵ_t is computed as follows:

$$\epsilon_t = 100 \times \max_{t \in [0,T]} |\log K_t - \log K_t^*| \tag{18}$$

for each simulation step t.

As highlighted in Ahn et al. (2017), den Haan (2010) originally used this measure to investigate the accuracy of the forecast rules in Krusell and Smith (1998). Ahn et al. (2017) discuss the dynamics of $\mathbf{g}_t$ obtained through the linearized method are analogous to the ones from the forecast rule. In the proposed method, the nonparametric law of motion of moments $\dot{\mathbf{m}} = \mathbf{G}(\mathbf{m}, A)$ is the comparable object.

Results

The number of total simulation steps is set to 5,500, with the first 500 ones dropped as a burn-in period when computing the accuracy measures. For the linearization method, I follow the calculation reported in Ahn et al. (2017) and replicate their statistics.

Figure 1 shows the simulated paths of the aggregate capital $K(t)$ when the Baseline model is solved using both the linearized solution method and the proposed method in the paper. The same parameter values and the realized values of the aggregate shocks are used for both cases. The figure shows that the proposed method closely tracks the path generated by the linearization method.[21]

Table 2 reports the maximum approximation errors of the two numerical solutions for different standard deviation parameters of the aggregate productivity shock. The linearization method is able to solve the models with relatively small shocks quite accurately. For example, in the models with the standard deviation of the shock as small as 0.1% and 0.7%,

[21]The average absolute percentage difference between two time series is less than 0.02 %.

Figure 2: Internal Consistency Check for the Proposed Method

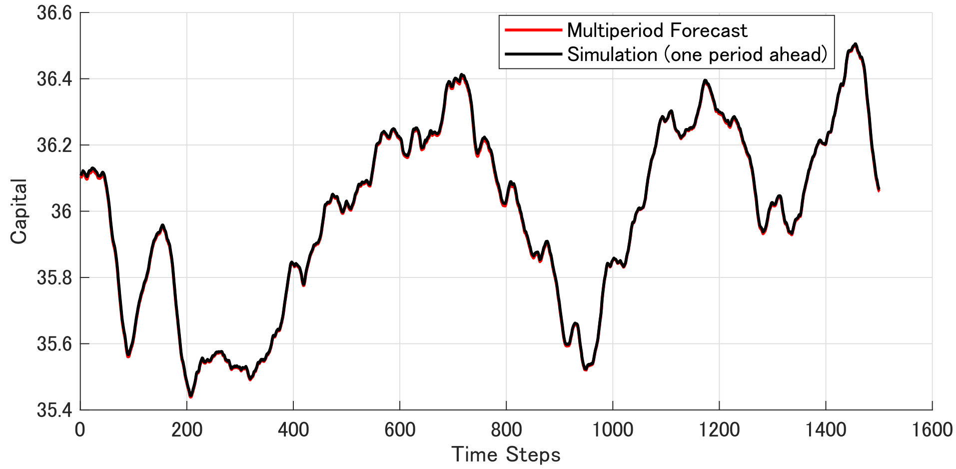

Notes: The red line corresponds to the multi-step forecast of aggregate capital $\{K_t\}_{t\in[0,T]}$ generated by using (13) for all t. The black line corresponds to the series $\{K_t^*\}_{t\in[0,T]}$ generated by solving (11) and (13) and updating the distribution of households.

the maximum approximation error is at most 0.049%. However, the approximation error increases with the variance of the aggregate shock. For example, Ahn et al. (2017) report that the maximum error reaches as high as 3.282%. This reflects the property that their method is a local approximation with respect to the aggregate shock, and the accuracy of the solution decreases when large shocks hit the economy.

On the other hand, the proposed method solves the model accurately even when the standard deviation of the aggregate shock is large. For example, when the standard deviation is 5.0%, the maximum approximation error is 0.22%, which is about 15 times more accurate than the linearized solution method. This demonstrates that this global solution method captures the aggregate risk and the model dynamics accurately even when the economy is far from the steady state. Furthermore, the proposed method solves the model with the similar degree of accuracy with linearization method when the shock size is small or moderate. For example, when the standard deviation is 1.0%, the proposed method achieves 0.04% of maximum approximation error, which is almost the same as the one obtained through the linearized solution method. While the error gets larger compared to the linearization method when the aggregate shock variance is small, this mainly comes from the numerical errors due to the specific choice of discretization or interpolation schemes.[22]

[22]The refinements of the grids generally improve the accuracy measures.

Table 3: Comparison of Error Statistics and Computation Time

	Maximum % Error	**Average % Error**	**Time**
Ahn et al. (2017)	0.049	-	0.25 sec
Proposed Method			
Poisson Process	0.036	0.009	1.2 sec
OU Process	0.026	0.009	9.6 sec

Notes: The first row replicates the result from Ahn et al. (2017) using the linearized solution method. The second row reports results using the proposed method. The computation time is the sum of time spent for computing steady state, derivatives, and linear system in Table 1 in Ahn et al. (2017) for the linearized solution method or the sum of time spent for computing the proxy distribution and stochastic equilibria for the proposed method.

Figure 2 shows part of two time series of aggregate capital obtained from the simulation with the proposed method (simulation time steps between 1,500 and 3,000): the multi-step forecast of aggregate capital $\{K_t\}_{t\in[0,T]}$ generated from equation (13) for all t and the simulated path of the aggregate capital $\{K_t^*\}_{t\in[0,T]}$ generated by solving equations (11) and (13), and updating the distribution at each simulation step. As explained in the previous subsection, these two series are used to compute the approximation error metric proposed by den Haan (2010). Figure 2 shows that two time series move quite closely, thus providing an illustration of the accuracy of the proposed method.

Table 3 summarizes the maximum and the average approximation errors, as well as the computation time for both methods. The linearized solution method provides solutions in a quarter of a second, while the proposed method is slower, although still very quick (within 2 seconds in the fastest case when a two-state Poisson process is used for the stochastic process of the aggregate shock in the baseline model). Even when the OU process is adopted, the stochastic equilibria of the baseline economy are obtained in less than 10 seconds. This speed gain comes from combining the merits of the continuous-time formulation and of the proposed method. As explained in Achdou et al. (2017), the continuous-time formulation simplifies many computational difficulties such as root findings or binding borrowing constraints. In addition, the adaptation of the Backward Induction method further enables researchers to avoid a costly simulation step to find the solutions of heterogeneous-agent business cycle models.

Overall, while the linearized solution method in continuous time still provides a powerful and useful tool to solve many problems, the proposed method can be seen as a complement to the existing toolbox of continuous-time numerical solution methods, as it enables the investigation of nonlinear and global solutions within a reasonable computation time.

Remarks on the Other Characteristics of the Proposed Method

In this subsection, I make several remarks on the characteristics of the proposed method other than accuracy and computation speed.

One of the remarkable contributions of Ahn et al. (2017) is the dramatic reduction of the programming costs by providing the standardized procedures of the linearization and dimension reduction, as well as the accompanying open-source toolbox. Once the steady-state version of the model is solved, it is possible to analyze the model with aggregate uncertainty very easily. In this respect, the proposed method requires more programming time because of two additional routines: one to solve the HJB equation with aggregate state variables, and the other one to solve for the moment-consistent proxy distribution.

The weakness of the proposed method is the necessity to deal with increasing dimensionality of the value function resulting from increasing the number of endogenous and exogenous state variables. While this issue also emerges in discrete-time solution methods, it causes an additional problem in continuous-time methods: the computational advantage of utilizing sparse matrix when updating the value function will be decreased as the cost of solving a system of linear equations becomes large.[23] Thus, models with many aggregate shocks such as the one by Smets and Wouter (2007) would require additional efforts to address high dimensionality.[24][25]

5 Other Examples

As an illustration of the merits of the proposed method, I consider two examples: a baseline economy with Epstein and Zin (1991) preferences and an OLG economy with aggregate shocks. Note that, in this section, I use a two-state Poisson process to model the aggregate productivity shock.

[23]When the system of linear equations is solved by the direct solution method, the sparse structure of the matrix is lost while implementing the LU decomposition. This problem is called "fill-in" problem.

[24]Using non-tensor-product grids such as adaptive sparse grids would be an option to address the dimensionality problem.

[25]For the other promising approach, Duarte (2018) recently developed a method which applied theories and computational libraries of the machine learning to solve a high-dimensional HJB equation. His paper studies the economy populated with multiple agents while this paper studies the economy with a continuum of agents.

5.1 Baseline model with Epstein and Zin (1991) Preferences

In this subsection, I modify the baseline model described in Section 2 by replacing the CRRA utility function with Epstein and Zin (1991) (EZ) preferences. This introduces more non-linearity in the model and serves as an effective example to illustrate the merits of nonlinear solution methods. For the continuous-time formulation of EZ preferences, I follow Duffie and Zin (1992).

Epstein and Zin (1991) Preferences in Continuous Time

To introduce Epstein and Zin (1991) preferences in continuous time, let us assume that v_t is the continuation value for a household at time t:

$$v_t = \mathbb{E}_t\Big[\int_t^\infty f(c_s, v_s)ds\Big]$$

where f is a utility aggregator for consumption c_t, and $\mathbb{E}_t$ is a conditional expectation operator at time t.

For $\sigma \neq 1$, f takes the following functional form:

$$f(c,v) = \frac{\rho}{1-\sigma}(1-\gamma)v\Big[\Big(\frac{c}{((1-\gamma)v)^{\frac{1}{1-\gamma}}}\Big)^{1-\sigma} - 1\Big]$$

Whereas for $\sigma = 1$ we have:

$$f(c,v) = \beta(1-\gamma)v\Big[\Big(\log c - \frac{1}{1-\gamma}\log((1-\gamma)v)\Big)\Big]$$

where ρ is the rate of time preference, σ is the (inverse of) elasticity of intertemporal substitution (EIS), γ is the parameter for relative risk aversion (RRA), c is consumption, and v is the value for the household.

Now, let us consider the model without aggregate productivity shocks. The HJB equation takes the following form:

$$0 = \max_c\ f(c,v) + \partial_a v(a,z)(ra - c + wz\bar{n}) + \lambda_z[v(a,z') - v(a,z)]$$

for all $z, z' \in \mathbf{z}$ such that $z \neq z'$.

The optimal consumption is derived from the first-order condition for the HJB equation

Table 4: Results on Numerical Solution and Errors (Epstein & Zin Preference)

	Maximum % Error	**Average % Error**	**Time**
Proposed Method			
EIS=1.5 and RRA=3	0.077	0.02	16.05 sec
EIS=1.5 and RRA=6	0.096	0.019	17.7 sec

Notes: The computation time for the proposed method is the sum of computation time for computing the proxy distribution and stochastic equilibria.

with respect to consumption given v:

$$\frac{\partial f(c,v)}{\partial c} = \partial_a v$$

Then, let us turn to the model with aggregate productivity shocks. The corresponding approximate recursive formulation takes the following form:

$$\begin{aligned} 0 = \max_c\ & f(c,v) + \partial_a v(a,z;\mathbf{m},A)\cdot[r(\mathbf{m},A)a + (1-\tau)w(\mathbf{m},A)z\bar{l} + \mathbb{I}_z\bar{b} - c] \\ & + \lambda_z[v(a,z';\mathbf{m},A) - v(a,z;\mathbf{m},A)] \\ & + \Omega_A[v(a,z';\mathbf{m},A') - v(a,z;\mathbf{m},A)] + \partial_m v(a,z;\mathbf{m},A)\cdot\frac{\partial\mathbf{m}}{\partial t} \end{aligned}$$

for all $z, z' \in \mathbf{z}$ such that $z \neq z'$ and $A, A' \in \mathbf{A}$ such that $A \neq A'$. Also, Ω_A becomes Ω_H if $A = A_H$ and Ω_L if $A = A_L$.

The next subsection describes numerical results obtained by the proposed method.

Numerical Results for the Model with Epstein and Zin Preferences

Following the same steps as in Section 2 and 4, I apply the Backward Induction method to compute the stochastic equilibria of the model with the Epstein and Zin preferences and simulate the model economy over 5,500 time steps.

Table 4 reports the results of the simulation with different parameter values. While this case features more curvature, the proposed method solves the models quickly and accurately. Specifically, it takes less than 20 seconds for stochastic equilibria. Moreover, there is no significant accuracy loss even though the model features high curvature, and the approximation error is as small as the baseline models with less curvature. This suggests that the solution method accurately captures the nonlinearity of individual decision rules and the associated change in the value after the aggregate shock realization and change in the endogenous aggregate state variables.

5.2 Overlapping Generations Economy with Aggregate Shocks

In this subsection, I consider the example of an OLG economy with aggregate shocks. In this case, as pointed out in Krueger and Kubler (2004), a simple approximation aggregation and Krusell and Smith (1998) algorithm do not work well and accuracy loss becomes larger. Instead, the continuous-time Backward Induction method still produces accurate solutions very fast.

In order to setup the OLG economy in continuous time, I follow the mean field game literature on player entry and exit.[26] Each agent now carries additional individual state variable $J \in [J_{start}, J_{end}] \subseteq R^{+}$ which corresponds to the age of the agent. The agents in the economy distribute over the tuples $(a(t), z(t), J)$. The age variable increases by the same amount in accordance with a unit of the time step; in other words, agents move deterministically along the age variable. Finally, although many parameters are carried over from the baseline model, but in the later, several new parameters are introduced.

The OLG Model in Continuous Time

First, I illustrate a sequential formulation of the agent optimization problem.

There is a continuum of agents with finite horizon. The value for the agent of age J at time t_0 is as follows:

$$
\begin{aligned}
v_{t_0} = & \max_{\{c(t)\}_{t=t_0}^{t=J_{end}}} \mathbb{E}_{t_0} \int_{t_0}^{J_{end}} e^{-\rho(t-t_0)} u(c(t)) dt \\
\text{s.t. } & u(c(t)) = \frac{c(t)^{1-\sigma}}{1-\sigma}, \sigma \geq 0 \\
& \dot{a}(t) = \left[r(t)a(t) + w(t)z(t)\bar{l}f(J) + \tau(J) - c(t) \right], a(t) \geq \underline{a}
\end{aligned}
$$

where $f : \mathbb{R}^{+} \to \mathbb{R}^{+}$ is the age-specific labor productivity, while $\tau : \mathbb{R}^{+} \to \mathbb{R}^{+}$ is the age-specific tax and transfer. The remaining notation is same as the baseline model.

In the OLG model, each household differs in age J, as well as in the levels of the net assets $a(t)$ and idiosyncratic productivity $z(t)$. Accordingly, I redefine $\mathbf{g}_t$ as a probability distribution of households at time t over the individual state variables $(a(t), z(t), J)$ and

[26] For example, see Lachapelle et al. (2011), Djehiche et al. (2017), and Bauso (2016).

the probability distribution adds up to one:

$$\mathbf{g}_t(a(t), z(t), J) \geq 0 \quad for\ all\ a(t) \in [\underline{a}, \infty), z(t) \in \mathbf{z}, and\ J \in [J_{start}, J_{end}],$$

$$and \quad \int\int \sum_{z(t)\in\mathbf{z}} \mathbf{g}_t(a(t), z(t), J)da = 1$$

Moreover, the evolution of distribution $\dot{\mathbf{g}}_t$ is obtained from a function $\mathbf{G}^{OLG}$ induced by the optimal saving decisions of households and the aggregate productivity shock $A(t)$:

$$\dot{\mathbf{g}}_t = \mathbf{G}^{OLG}(\mathbf{g}_t, A(t))$$

The important difference from the baseline model is that the households die at age J_{end} and thus exit the economy. At the same time, newly born households enter the economy. To preserve the total mass of households, at each point in time, the measure of households with age J_{end} should be the same as the measure of households with age J_{start}. Moreover, in this example, bequests from the dead to the new generation are not considered. Thus, this example specifically considers the case where households of age J_{start} starts with $a = 0$, while those of age J_{end} die with $a = 0$. These conditions can be summarized as follows[27]:

$$\sum_{z\in\mathbf{Z}} \mathbf{g}_t(0, z(t), J_{start}) = \int \sum_{z\in\mathbf{Z}} \mathbf{g}_t(a(t), z(t), J_{end})da$$

Furthermore, agents are assumed to be uniformly distributed over age J. Consequently, the mass of agents at each age J is $n_J = 1/J$.

Given the definition of $\mathbf{g}_t$, the labor and rental markets equilibrium conditions become:

$$r(t) = \alpha A(t)K(t)^{\alpha-1}L(t)^{1-\alpha}, \quad w(t) = (1-\alpha)A(t)K(t)^{\alpha}L(t)^{-\alpha}$$

$$K(t) = \int\int \sum_{z\in\mathbf{z}} a\mathbf{g}_t(a, z, J)dadJ$$

$$L(t) = \int\int \sum_{z\in\mathbf{z}} \bar{l}zf(J)\mathbf{g}_t(a, z, J)dadJ$$

The next subsection describes the approximate aggregate state space representation of the model which is used in the computation.

[27]The example here considers the case where all the households with age J_{end} has $a = 0$ when they die as a result of the terminal condition of the value function, while it is possible that they die with $a > 0$ in a different case where bequests exist, for example. The condition after this footnote allows for both cases.

Approximate Aggregate State Space Representation

As in the baseline model, the aggregate capital is used to approximate the aggregate endogenous variable and the moments $\mathbf{m}$.

In this case, the HJB equation takes the following form:

$$
\begin{aligned}
\rho v(a, z, J; \mathbf{m}, A) = \max_{c} \ & u(c) \\
& + \partial_a v(a, z, J; \mathbf{m}, A)(r(\mathbf{m}, A)a - c + w(\mathbf{m}, A)zf(J)\bar{l} + \tau(J)) \\
& + \partial_J v(a, z, J; \mathbf{m}, A)\Delta_{age} + \partial_m v(a, z, J; \mathbf{m}, A) \cdot \frac{\partial \mathbf{m}_t}{\partial t} \\
& + \lambda_z [v(a, z', J; \mathbf{m}, A) - v(a, z, J; \mathbf{m}, A)] \\
& + \Omega_A [v(a, z, J; \mathbf{m}, A') - v(a, z, J; \mathbf{m}, A)]
\end{aligned}
$$

and $\Delta_{age} \equiv \lim_{\Delta \to 0} \frac{\Delta_J}{\Delta} = 1$ for all $a \in \mathbb{R}^+$, $\mathbf{m} \in \mathbb{R}^+$, $J \in [J_{start}, J_{end})$, $z, z' \in \mathbf{z}$ such that $z \neq z'$ and $A, A' \in \mathbf{A}$ such that $A \neq A'$.

Furthermore, at the age of $J = J_{end}$, households die and leave the economy. I assume that there is a utility flow associated with leaving by consuming cash on hand: $c_{end} = \tau(J_{end})$. This translates into the following boundary condition on the value function:

$$
v(a, z, J_{end}; \mathbf{m}, A) = u(c_{end})
$$

The real interest and wage rates are given by:

$$
r(\mathbf{m}, A) = \alpha A \mathbf{m}^{\alpha - 1} (\bar{z}\bar{l})^{1-\alpha} - \delta, \quad w(\mathbf{m}, A) = (1-\alpha) A \mathbf{m}^{\alpha} (\bar{z}\bar{l})^{-\alpha}
$$

Then, let us compute a nonparametric law of motion of the moments $\dot{\mathbf{m}} = \hat{\mathbf{G}}_{OLG}(\mathbf{m}, A)$. Similarly to the baseline model, the law of motion of the moments needs to be consistent with the individual policies:

$$
\frac{\partial \mathbf{m}}{\partial t} = \hat{\mathbf{G}}_{OLG}(\mathbf{m}, A) = \int \int \sum_{z \in \mathbf{z}} s^*(a, z, J; \mathbf{m}, A) \mathbf{g}^p_{OLG}(a, z, J; \mathbf{m}, A) da dJ
$$

where $s^*(a, z, J; \mathbf{m}, A) \equiv r(\mathbf{m}, A)a - c + w(\mathbf{m}, A)zf(J)\bar{l} + \tau(J)$.

$\mathbf{g}^p_{OLG}(a, z, J; \mathbf{m}, A)$ is the proxy distribution of the cross-sectional distribution of households

Figure 3: Cohort Average of Asset, Consumption, & Income in Steady State OLG

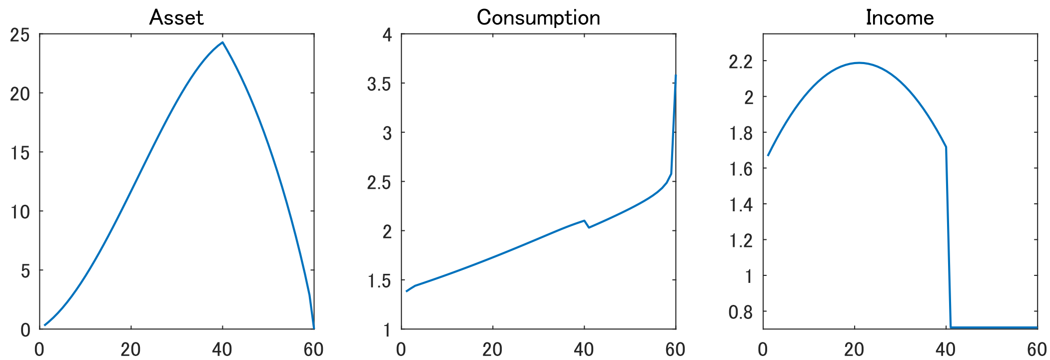

Notes: The horizontal axis represents the age. The vertical axis represents the cohort average value for each variable in the steady state of OLG economy. Income here refers to the sum of the labor income (adjusted by the age-specific labor productivity) and the age-specific tax and the transfer.

over (a, z, J) given $\mathbf{m}$ and A such that:

$$\begin{aligned} &\mathbf{g}^p_{OLG}(a, z, J; \mathbf{m}, A) \geq 0 \ for\ all\ a \in [\underline{\text{a}}, \infty), z \in \mathbf{z}, and\ J \in [J_{start}, J_{end}], \\ &\int\int \sum_{z\in\mathbf{Z}} \mathbf{g}^p_{OLG}(a, z, J; \mathbf{m}, A) da dJ = 1, \\ and \quad &\int \sum_{z\in\mathbf{Z}} \mathbf{g}^p_{OLG}(a, z, J; \mathbf{m}, A) da = n_J \ for\ all\ J \in [J_{start}, J_{end}] \end{aligned}$$

This $\mathbf{g}^p_{OLG}$ also needs to be consistent with the approximate aggregate variables:

$$\int\int a \sum_{z\in\mathbf{z}} \mathbf{g}^p_{OLG}(a, z, J; \mathbf{m}, A) da dJ = \mathbf{m}$$

With these, I use the proposed method to solve for stochastic equilibria.

Numerical Results for the OLG Model with and without Aggregate Shocks

This subsection introduces the numerical results for the OLG economy.

Before showing them, I briefly overview parameters and numerical settings in the economy. Unless otherwise stated, I use the same parameter values as the baseline economy. A household in the economy lives about 60 years, thus $J_{start} = 1$ and $J_{end} = 60$. I choose an equispaced grid for J with 60 grid points. Moreover, I denote $J_{ret} = 40$ as the retirement age. I choose the age-specific labor productivity $f(J)$ and the age-specific tax and a transfer $\tau(J)$ such that the income profile is hump-shaped as in Figure 3. $f(J)$ is a quadratic function when $J < J_{ret}$ and it becomes 0 when $J_{ret} \leq J$. $\tau(J)$ becomes is labor income tax

Table 5: Approximation Errors and Computation Time (OLG economy)

	Maximum % Error	**Average % Error**	**Time**
Continuous Time OLG	0.4	0.09	82.8 sec

Notes: The computation time is the sum of computation time for computing proxy distribution and stochastic equilibria.

when $J < J_{ret}$ and a transfer when $J_{ret} \leq J$. For the net assets grid, I choose $I = 80$ with $a_1 = 0$ and $a_I = 65$. As for the approximate endogenous aggregate variable (the aggregate capital), I select the upper and lower bounds of the given model parameters.[28]

Figure 3 shows the cross-sectional cohort averages of net assets, consumption, and income in the steady state. Here, income is defined as the sum of labor income (adjusted by the age-specific labor productivity), age-specific tax, and transfer.[29] It is hump-shaped and drops after the household retires. Such shape of the income trajectory is arbitrarily chosen, and thus it does not correspond exactly to the results reported by the previous empirical literature on the household income.[30] With this income profile and a finite life span, the net asset profile also exhibits a hump-shaped path across age. As a consequence, the younger households increase their net asset holdings, while the older and retired households start to draw down their assets to compensate the decrease in income. When households die and exit the economy, they spend all their assets, and there is nothing left to the off-springs.

Next, I introduce results for the numerical solution of the OLG model with the aggregate shock, following the same steps outlined in the previous sections. I solve for value function $v(a, z, J; \mathbf{m}, A)$ and the law of motion of the moments $\dot{\mathbf{m}} = \mathbf{G}^{OLG}(\mathbf{m}, A)$ given the proxy distribution. Note that the proxy distribution can be refined after the simulation step.[31]

Table 5 reports the error statistics and the computation time from the results obtained by using the proposed method. As pointed out in Krueger and Kubler (2004), the OLG

[28] For other parameters in the current version, I set $\sigma = 2$, $z_L = 0.9, z_H = 1.1$, and $\lambda_z = 0.5$.

[29] The choice of parameters and model specification results in this specific consumption path. Depending on the parameter values, the consumption profile of each cohort can look quite different.

[30] As for the idiosyncratic shock process, I currently work on the extension to introduce realistic income processes with kurtosis and skewness in the model.

[31] When a simulated path for the aggregate capital and the cross-sectional distribution is computed, the Kolmogorov forward Equation is modified to include birth and death of households. Specifically, let $\tilde{\mu}$ be inflows and outflows of households with $a = 0$ from the economy between time t and $t + dt$. Then, corresponding Kolmogorov forward Equation using the implicit method is following:

$$\mathbf{g}_{t+dt} = (\mathbf{I} - \mathbf{A}^T)^{-1}\mathbf{g}_t dt + \tilde{\mu}$$

economy features a richer movement of the cross-sectional distribution of households with respect to aggregate shocks; thus, a simple approximation algorithm and algorithm as proposed in Krusell and Smith (1998) may not work. However, as the proposed method can accommodate a movement of cross-distribution movement by the flexible representation of the proxy distribution and the law of motion of moments, this method achieves high accuracy. This is confirmed by the fact that the maximum approximation error in den Haan (2010) is still as small as 0.4%. On the other hand, the computation time increases compared to the baseline model in Section 2 because increment in dimensionality makes the computation of systems of linear equation costly. Nonetheless, the appropriate choice of the initial proxy distribution and a proper algorithm to find the most suitable one can significantly speed up computation. Moreover, the parallelization could help hasten the computation when a high dimensional dynamic system is considered.

6 Conclusion

This paper proposes a global and nonlinear solution method for continuous-time heterogeneous agent models developed by extending the Backward Induction method originally developed by Reiter (2010) for discrete time models. The main findings are as follows: Compared to the linearized solution method, the proposed method achieves a higher degree of accuracy even with the large volatility of the aggregate shock, and the computational speed remains fast compared to the linearized solution method developed in Ahn et al. (2017).

For the further study, these merits naturally lead to an application of the proposed method to revisit the implications of aggregate shocks on the economy. Specifically, the computational advantage of the proposed method enables researchers to address fundamental macroeconomic questions such as the asymmetric business cycle dynamics, the cost of the business cycle, and asset pricing implications on real variables even in models with a fairy rich heterogeneity of households or firms. For example, the proposed method can be extended to study the implications of aggregate shock on liquid and illiquid asset models in Kaplan et al. (2018).

While the proposed method could serve as an attractive complement to the existing methods, many practical issues in the numerical method make further development of continuous-time solution method itself looks a very fruitful research path. For example, as a common weak point of continuous-time solution method, increasing dimensions would become a non-negligible factor for losing computation speed, and this may force researchers to give up quantitatively relevant model elements. From this perspective, incorporation of dimension reduction technique such as the proper orthogonal decomposition method or the sparse grid method can further enhance the capability of the continuous-time approach.

References

Achdou, Y., J. Han, J.-M. Lasry, P.-L. Lions, and B. Moll (2017). "Income and Wealth Distribution in Macroeconomics: A Continuous-Time Approach." Working Paper 23732, National Bureau of Economic Research.

Ahn, S., G. Kaplan, B. Moll, T. Winberry, and C. Wolf (2017). *When Inequality Matters for Macro and Macro Matters for Inequality.* University of Chicago Press.

Anderson, E. W., L. P. Hansen, and T. J. Sargent (2012). "Small noise methods for risk-sensitive/robust Economies." *Journal of Economic Dynamics and Control 36*(4), 468 – 500.

Bauso, D. (2016). *Game Theory with Engineering Applications.* Society for Industrial and Applied Mathematics.

Bhandari, A., D. Evans, M. Golosov, and T. J. Sargent (2017). "Inequality, Business Cycles and Fiscal-Monetary Policy." University of Minnesota working paper.

Boppart, T., P. Krusell, and K. Mitman (2017). "Exploiting MIT shocks in Heterogeneous-Agent Economies: The Impulse Response as a Numerical Derivative." Working Paper 24138, National Bureau of Economic Research.

Campbell, J. R. (1998). "Entry, Exit, Embodied Technology, and Business Cycles." *Review of Economic Dynamics 1*(2), 371 – 408.

Childers, D. (2018). "Solution of Rational Expectations Models with Function Valued States." Carnegie Mellon University working paper.

Djehiche, B., A. Tcheukam, and H. Tembine (2017). A mean-field game of evacuation in multilevel building. *IEEE Transactions on Automatic Control 62*(10), 5154–5169.

Duarte, V. (2018). "Sectoral Reallocation and Endogenous Risk-Aversion: Solving Macro-Finance Models with Machine Learning." MIT Sloan School of Management working paper.

Duffie, D. and L. G. Epstein (1992). "Stochastic Differential Utility." *Econometrica 60*(2), 353–394.

Epstein, L. G. and S. E. Zin (1991). "Substitution, Risk aversion, and the Temporal Behavior of Consumption and Asset Returns: An Empirical Analysis." *Journal of Political Economy 99*(2), 263–286.

Evans, D. G. (2015). *Perturbation theory with heterogeneous agents: Theory and applications*. Ph. D. thesis, New York University.

Fleming, W. H. (1971). "Stochastic Control for Small Noise Intensities." *SIAM Journal on Control 9*(3), 473–517.

Fleming, W. H. and P. E. Souganidis (1986). "Asymptotic Series and the Method of Vanishing Viscosity." *Indiana University Mathematics Journal 35*(2), 425–447.

Gornemann, N., K. Kuester, and M. Nakajima (2016). "Doves for the Rich, Hawks for the Poor? Distributional Consequences of Monetary Policy." CEPR Discussion Papers 11233, C.E.P.R. Discussion Papers.

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

den Haan, W. J., K. L. Judd, and M. Juillard (2010). "Computational suite of models with heterogeneous agents: Incomplete markets and aggregate uncertainty." *Journal of Economic Dynamics and Control 34*(1), 1 – 3.

den Haan, W. J. and P. Rendahl (2010). "Solving the incomplete markets model with aggregate uncertainty using explicit aggregation." *Journal of Economic Dynamics and Control 34*(1), 69 – 78.

Kaplan, G., B. Moll, and L. G. Violante (2018). "Monetary Policy According to HANK." *Forthcoming in American Economic Review*.

Khan, A. (2017). "Large Recessions in an Overlapping Generations Model with Unemployment Risk." Ohio State University working paper.

Khan, A. and J. K. Thomas (2008). "Idiosyncratic Shocks and the Role of Nonconvexities in Plant and Aggregate Investment Dynamics." *Econometrica 76*(2), 395–436.

Kim, H. (2017). "Inequality, Portfolio Choice, and the Business Cycle." Concordia University working paper.

Krueger, D. and F. Kubler (2004). "Computing equilibrium in olg models with stochastic production." *Journal of Economic Dynamics and Control 28*(7), 1411 – 1436.

Krusell, P. (2017). "AKMWW: a weapon of Mass Creation?," comment on Ahn, Kaplan, Moll, Winberry and Wolf (2017) "When inequality matters for Macro and Macro matters for Inequality," *NBER Macroeconomics Annual*.

Krusell, P. and J. Anthony A. Smith (1998). "Income and Wealth Heterogeneity in the Macroeconomy." *Journal of Political Economy 106*(5), pp. 867–896.

Lachapelle, A. and M.-T. Wolfram (2011). "On a mean field game approach modeling congestion and aversion in pedestrian crowds." *Transportation research part B: methodological 45*(10), 1572–1589.

McKay, A. and R. Reis (2016). "The Role of Automatic Stabilizers in the U.S. Business Cycle." *Econometrica 84*(1), 141–194.

Mongey, S. and J. Williams (2016). "Firm Dispersion and Business Cycles: Estimating Aggregate Shocks using Panel Data." *Working paper, Department of Economics, NYU*.

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

Reiter, M. (2010). "Solving the incomplete markets model with aggregate uncertainty by backward induction." *Journal of Economic Dynamics and Control 34*(1), 28 – 35.

Smets, F. and R. Wouters (2007). "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *The American Economic Review 97*(3), 586–606.

Terry, S. J. (2017). "Alternative Methods for Solving Heterogeneous Firm Models." *Journal of Money, Credit and Banking 49*(6), 1081–1111.

Winberry, T. (2018). "A Method for Solving and Estimating Heterogeneous Agent Macro Models." *Forthcoming in Quantitative Economics*.

Young, E. R. (2010). "Solving the incomplete markets model with aggregate uncertainty using the krusell–smith algorithm and non-stochastic simulations." *Journal of Economic Dynamics and Control 34*(1), 36 – 41.