Aggregate Uncertainty, Repeated Transition Method, and the Aggregate Cash Cycle*

Hanbaek Lee†

August 13, 2022

(click here for the latest version)

Abstract

This paper develops and tests a novel algorithm that solves heterogeneous agent models with aggregate uncertainty. The algorithm is based on the ergodic theorem: if a simulated path of the aggregate shock is long enough, all the possible aggregate allocations are realized, which allows to fully recover rationally expected future outcomes at each point on the path. This method solves the nonlinear dynamic stochastic general equilibrium globally with a high degree of accuracy. Furthermore, the market-clearing prices and the expected aggregate states are directly computed at each point on the path without relying on a parametric law of motion. Using the algorithm, I analyze a heterogeneous-firm business cycle model where firms are subject to an external financing cost and hoard cash as a buffer stock. In the model, due to the missing general equilibrium effect on cash, the aggregate fluctuations in cash and consumption feature significant nonlinearity and state dependence. Based on the model, I discuss the business cycle implications of the corporate cash holdings.

*This paper is based on my Ph.D. dissertation at the University of Pennsylvania. I am extremely grateful to Jesús Fernández-Villaverde, Dirk Krueger, Andrew Abel, and Frank Schorfheide for their invaluable guidance and support. I also thank Florin Bilbiie, Christopher Carroll, Eunseong Ma, Alessandro Peri, José-Víctor Ríos-Rull, Ludwig Straub, and Guanyi Yang for insightful comments and discussions. All errors are my own.

†University of Cambridge. Email: hanbaeklee1@gmail.com

Contents

1 Introduction

In this paper, I introduce an algorithm that solves a heterogeneous agent model with aggregate uncertainty that is free from the law of motion specification. I name the algorithm as repeated transition method.

Under the rational expectation, heterogeneous agents are aware of the true law of motion in the aggregate states and make a correct prediction on the future aggregate state. In contrast, there is no specific form of the law of motion known to a researcher. And it is computationally costly to track the evolution of a distribution that is an infinite-dimensional object. To overcome this problem, Krusell and Smith (1998) suggested a log-linear prediction rule of the finite number of moments of the individual state distribution as an approximation to the true law of motion. Afterward, numerous research papers in the literature have found this prediction rule gives a surprisingly accurate approximation to the true law of motion in the broad class of heterogeneous agent models with aggregate uncertainty.

Still, there are macroeconomic environments where the log-linear rule does not apply. A dynamics of aggregate allocations subject to occasionally binding constraints are an example of such cases (Fernandez-Villaverde et al., 2020). Also, history dependence in the investment dynamics, as in Lee (2022), makes it difficult to approximate the true law of motion using the log-linear specification. According to Krusell and Smith (1997) and Krusell and Smith (1998), these problems can be handled by tracking more moments of the state distribution. However, the functional form of the prediction rule and selection of the moments still remain as an open-ended problem.

The repeated transition method overcomes these problems by recursive approximation to the evolution of true state distributions on a single simulated path of aggregate shocks (in-sample simulation). The method does not depend on the parametric form of the prediction rule because the market-clearing prices and the expected allocations are directly computed at each point on the path. Once the approximation is completed, I estimate the best-fitting non-parametric/parametric law of motion from the in-sample simulation. Using this law of motion, I extrapolate the stochastic dynamics of allocations over the out-of-sample simulated paths of the aggregate shocks. Lastly, I check the validity of the law of motion by comparing the model's solution over the out-of-sample simulated paths of the aggregate shocks based on the estimated law of motion and the extrapolated aggregate allocations.

The key step in the repeated transition method is to build a correct time-specific expected future value function in each period by combining value functions in the simulation history that share the same aggregate states. For example, if an economy is located at time t, for each possible aggregate shock realization $s \in S$ in $t+1$, I find a period τ_s in the simulation history where the aggregate states are the closest to the aggregate state of period $t+1$, including the aggregate shock s. Then, I combine the value functions from these periods $\{\tau_s\}_{s\in S}$ to construct the expected value function at $t+1$. Theoretically, if the simulation path is infinitely long, there exists the period τ_s where the aggregate allocations are perfectly identical to period $t+1$ with an aggregate shock realization s with probability one. Therefore, the true expected value function can be constructed from this approach. However, in practice, due to the finite length of the simulation path, often there is no exact period τ_s in the simulation history that shares the same aggregate allocations including a shock realization s as in period $t+1$. Therefore, I approximate the expected value function by interpolating value functions from periods that closely mimics period $t+1$ for each aggregate shock realization.

The repeated transition method builds upon the method utilizing perfect-foresight impulse response suggested by Boppart et al. (2018). In the paper, aggregate allocations' impulse responses are obtained from the transition dynamics induced from MIT shocks to the steady-state distribution. Then, the law of motion of aggregate allocations is locally approximated around the steady-state. Therefore, the method assumes certainty equivalence between the expected deterministic path and the expected path when the uncertainty is present. In contrast, the repeated transition method does not assume certainty equivalence and globally solves the model. And it directly computes aggregate allocations and market-clearing prices in each period on the simulation path without specifying the law of motion.

Therefore, the repeated transition algorithm is distinguished from the solution methods based on perturbation and linearization (Reiter, 2009; Boppart et al., 2018; Winberry, 2018; Childers, 2018; Auclert et al., 2019). As this method utilizes a single path of simulated aggregate shock that is long enough to fully represents the stochastic process, its approach is closely related to Kahou et al. (2021). Kahou et al. (2021) utilizes the fact that a whole economy's dynamics can be characterized by solving a finite number of agents' problesm on a single Monte Carlo draw of individual shocks under the permutation-invariance condition. And the law of motion is nonlinearly computed using the deep-learning algorithm. Instead of the law of motion being characterized as an equilibrium object, the repeated transition algorithm computes the path of equilibrium allocations at each point on the simulated path. Then the law of motion can be backed out from the time series of the realized allocations. This method relies only on relatively simple computational techniques but computes highly accurate solutions. Also, the algorithm is widely applicable as the algorithm does not rely on the particular characteristics of the problem presented in this paper.

The repeated transition algorithm outperforms the algorithm of Krusell and Smith (1997) in models with non-trivial market-clearing conditions and nonlinear aggregate dynamics in terms of accuracy and computation time. However, for the models with log-linear aggregate dynamics without a non-trivial market-clearing condition, such as the model of Krusell and Smith (1998), the repeated transition method does not work as fast as Krusell and Smith (1998) algorithm.

Using the repeated transition method, I study a business cycle implication of corporate cash holdings in a heterogeneous-firm business cycle model. In the model, firms face a convex external financing cost, so they have a precautionary motivation to hoard cash. Cash is assumed to be an internal asset of a firm. Thus, it is not traded across firms and discounted at a different rate than the interest rate in the equity market. The rate is exogenously given as a parameter in the model. Due to these features of cash, the dynamics of aggregate cash holdings in the model become highly nonlinear; there is no general equilibrium force to flatten the dynamics of aggregate cash holdings. On top of this nonlinearity, the market-clearing condition in the model is not trivial, as in Khan and Thomas (2008). Despite these difficulties in computation, the repeated transition method solves the model efficiently and accurately.

In the model, lagged aggregate cash holding significantly lowers the consumption volatility. This model prediction is empirically supported by consumption heteroskedasticity on lagged cash holding, and this empirical pattern is observed only after the early 1980s.[1] The fact that the corporate cash holding has dramatically increased after the early 1980s partly explains why conditional heteroskedasticity is observed only after the early 1980s. Then I show that the smoothing effect of cash holding on consumption occurs only when the negative aggregate productivity shock hits the economy. This validates the model's main mechanism where cash holding gives insurance to households' dividend income against the negative productivity shock.

Roadmap Section 2 explains the repeated transition method based on the model in Krusell and Smith (1998). Section 3 compares the computation results of the repeated transition method with other methods existing in the literature. Section 4 introduces a heterogeneous-firm business cycle model where firms save cash. Section 5 discusses the business cycle implication of corporate cash holdings predicted by the model compared to the observations from the data. Section 6 concludes. Other detailed figures and tables are included in appendices.

2 Repeated transition method

2.1 A model for algorithm introduction: Krusell and Smith (1998)

I explain the repeated transition method based on the heterogeneous agent model with aggregate uncertainty in Krusell and Smith (1998). In this section, I briefly introduce the basic environment of the model.

[1] The result is robust over other choices of the cutoff year around 1980.

A measure one of ex-ante homogenous households consumes and saves. At the beginning of a period, a household is given wealth a_t and an idiosyncratic labor supply shock z_t. Households are aware of the distribution of households Φ_t, the aggregate productivity shock A_t, and how the aggregate states evolve in the future $G(A_t, \Phi_t)$. The idiosyncratic shock and the aggregate shock follow the stochastic Markov processes elaborated in Krusell and Smith (1998). Households are subject to a borrowing constraint $a_{t+1} \geq 0$, as in the standard incomplete market model. I close the model by introducing a representative firm producing output from a constant returns-to-scale production function. The recursive formulation of the model is as follows:

$$
\begin{aligned}
\text{(Household)} \quad & v(a, s; S, \Phi) = \max_{c, a'} \quad log(c) + \beta \mathbb{E}(v(a', s'; S', \Phi')) \\
& \text{s.t.} \quad c + a' = w(S, \Phi) z(s) + (1 + r(S, \Phi)) a \\
& \quad a' \geq 0, \quad \Phi' = G(\Phi, S) \\
\text{(Production sector)} \quad & \max_{K, L} \quad A(S) K^{\alpha} L^{1-\alpha} - w(S, \Phi) L - (r(S, \Phi) + \delta) K \\
\text{(Market clearing)} \quad & \hat{K}(S, \Phi) = \int a d\Phi(a, z; S, \Phi) \\
& \hat{L}(S, \Phi) = \int z d\Phi(a, z; S, \Phi) \\
\text{(Shock processes)} \quad & \mathbb{P}(s', S' | s, S) = \pi_{sS, s'S'}, \quad s, s' \in \{u, e\}, \quad S, S' \in \{B, G\}
\end{aligned}
$$

$$
\pi := \begin{bmatrix}
\pi_{uB,uB} & \pi_{uB,eB} & \pi_{uB,uG} & \pi_{uB,eG} \\
\pi_{eB,uB} & \pi_{eB,eB} & \pi_{eB,uG} & \pi_{eB,eG} \\
\pi_{uG,uB} & \pi_{uG,eB} & \pi_{uG,uG} & \pi_{uG,eG} \\
\pi_{eG,uB} & \pi_{eG,eB} & \pi_{eG,uG} & \pi_{eG,eG}
\end{bmatrix} = \begin{bmatrix}
0.525 & 0.350 & 0.03125 & 0.09375 \\
0.035 & 0.84 & 0.0025 & 0.1225 \\
0.09375 & 0.03125 & 0.292 & 0.583 \\
0.0099 & 0.1151 & 0.0245 & 0.8505
\end{bmatrix}
$$

where $s = u$ means unemployed idiosyncratic state, $z = 0.25$. $s = e$ means employed idiosyncratic state, $z = 1$. $S = B$ indicates a bad aggregate state, $A = 0.99$. $S = G$ indicates a good aggregate state, $A = 1.01$.

2.2 Algorithm

I simulate a single path of aggregate TFP shocks, $\mathbb{A} = \{A_t\}_{t=0}^{T}$, from the aggregate transition matrix π^A. The aggregate transition matrix is as follows:

$$\pi^A = \begin{bmatrix} \pi_{B,B} & \pi_{B,B} \\ \pi_{G,B} & \pi_{G,B} \end{bmatrix} = \begin{bmatrix} 0.875 & 0.125 \\ 0.125 & 0.875 \end{bmatrix}$$

For the brevity of notation, I define a price vector $p_t := (w_t, r_t)$. The repeated transition method is based on the following statements:

1. If the true time-specific prices, $\{p_t\}_{t=0}^{T}$ are known, the true dynamic path of value function $\{v_t\}_{t=0}^{T-BurnIn}$ can be approximated by solving the problem from backward starting from $t = T$ with an initial guess v_T^0.[2]

2. If the true time-specific value functions, $\{v_t\}_{t=0}^{T}$ are known, optimal inter-temporal policy functions, $\{a_{t+1}\}_{t=0}^{T-1}$ can be obtained. Then, the true dynamic path of distribution $\{\Phi_t\}_{t=BurnIn}^{T}$ can be obtained by evolving the initial guess Φ_0^0 forward using the optimal inter-temporal policy functions, $\{a_{t+1}\}_{t=0}^{T-1}$.

3. If true dynamic paths of the value functions and the distributions, $\{v_t, \Phi_t\}_{t=0}^{T}$ are known, the time-specific prices, $\{p_t\}_{t=0}^{T}$ can be obtained from the market-clearing conditions.

By these three statements, I can approximate the true allocations, $(p_t, v_t, \Phi_t)_{t=BurnIn}^{T-BurnIn}$ from the following simulation chain:

1. Given nth guess on the price vector, $\{p_t^{(n)}\}_{t=0}^{T}$, compute $\{\widehat{v}_t^*\}_{t=0}^{T}$.

2. Given the value functions $\{\widehat{v}_t^*\}_{t=0}^{T}$, obtain the inter-temporal policy functions, and compute $\{\widehat{\Phi}_t^*\}_{t=0}^{T}$.

3. Given $\{\widehat{v}_t^*, \widehat{\Phi}_t^*\}_{t=0}^{T}$, compute $\{\widehat{p}_t^*\}_{t=0}^{T}$ from the market-clearing conditions.

[2] To make an agent correctly expect a one-period-ahead value function for each future shock realization, I use an interpolation method which will be explained later.

4. Evaluate the following criterion.

$$\sup_{BurnIn \le t \le T - BurnIn} ||p_t^{(n)} - \widehat{p}_t^*||_\infty < tol$$

If the criterion is not satisfied, update the guess $\{p_t^{(n+1)}\}_{t=0}^T$.

The detailed algorithm is explained below with the pseudo code. The convergence of the algorithm hinges on the stability of the equilibrium transition path. If the stability is not guaranteed, then the repeated transition method does not work. One important thing to note is the algorithm uses only a single simulated path of the aggregate shocks as if they are given parameters. I call this path to be fitted as in-sample path.

Once the algorithm converges, the approximated law of motion is parametrically or non-parametrically estimated from the simulated path (in-sample). Then, using the estimated law of motion, the stochastic equilibrium allocations for the out-of-sample paths are obtained.

The pseudo code for the repeated transition method is as follows:

1. Discretize the aggregate TFP shock process by S states.[3]

2. Simulate a path of aggregate TFP shocks, $\mathbb{A} = \{A_t\}_{t=0}^T$. This path is going to be repeatedly used in the following steps without any change. This is the in-sample path.

3. Guess on the paths of the prices, the value functions, and the state distributions, $\{\widehat{p}_t^{(n)}, \widehat{v}_t^{(n)}, \widehat{\Phi}_t^{(n)}\}_{t=0}^T$.

4. Solve the model backward in the following sub-steps:

   (a) Make a partition $\tilde{T}(s)$ of simulation paths grouped by the realized aggregate TFP level: $\tilde{T}(s) = \{\tau | A_\tau = A_s\} \subseteq \{0, 1, 2, ..., T\}$ for $s \in \{1, 2, 3, ...S\}$.

   (b) For each $s \in \{1, 2, 3, ...S\}$ find $\{\omega_{j,s} | \sum \omega_{j,s} = 1, j \in \tilde{T}(s)\}$ such that $\widehat{\Phi}_{t+1}^{(n)} = \sum_{j \in \tilde{T}(s)} \omega_{j,s} \widehat{\Phi}_j^{(n)}$. That is, find a set of weights that interpolates $\Phi_{t+1}^{(n)}$ using $\{\widehat{\Phi}_j^{(n)} | j \in$

[3]This discretization step is unnecessary. However, for practical illustration, I describe the pseudo code based on the discretization.

$\tilde{T}(s)\}$. The uniqueness of the weight set $\{\omega_{j,s}\}$ is not guaranteed. However, for the given model, I track only the first moment $\widehat{K}_{t+1}$ in the spirit of Krusell and Smith (1998). Therefore, I can find the unique $\{\omega_{j1,s}, \omega_{j2,s}\}$ such that $\widehat{K}_{t+1} = \widehat{K}_{j1}\omega_{j1,s} + \widehat{K}_{j2}\omega_{j2,s}$, and $\omega_{j1,s} + \omega_{j2,s} = 1$ with $\{\widehat{K}_{j1}, \widehat{K}_{j2}\}$ being the nearest combination to $\widehat{K}_{t+1}$. I set $\omega_{m,s} = 0$ for $m \notin \{j1, j2\}$.

(c) Approximate the true future value function $\mathbb{E}v_{t+1}(\cdot,\cdot) := \mathbb{E}v(\cdot,\cdot; A_{t+1}, \Phi_{t+1})$ as follows:

$$\mathbb{E}v_{t+1}(\cdot,\cdot) \cong \mathbb{E}\left[\sum_{j\in\tilde{T}(s)} \widehat{v}_j^{(n)}(\cdot,\cdot)\omega_{j,s}\right] = \mathbb{E}\left[\widehat{v}_{j1}^{(n)}(\cdot,\cdot)\omega_{j1,s} + \widehat{v}_{j2}^{(n)}(\cdot,\cdot)\omega_{j2,s}\right]$$

In this step, the value function is linearly interpolated. As the value function is the most smooth object in the equilibrium allocations, this step incurs only a small approximation error if the elements of $\{K_j | j \in \tilde{T}(s)\}$ are closely located to each other.

(d) Solve the problem for given t. Then I obtain the solution $\{\widehat{v}_t^*, \widehat{a}_{t+1}^*\}$

(e) If $t > 0$, update $t' = t - 1$ and go back to step 4a. If the algorithm arrives at $t = 0$, the time specific inter-temporal policy functions $\{\widehat{a}_{t+1}^*\}_{t=0}^T$ are obtained. Also, the sequence of implied value functions $\{\widehat{v}_t^*\}_{t=0}^T$ are obtained.

5. Using $\{\widehat{a}_{t+1}^*\}_{t=0}^T$, simulate the distribution for $t = 1, 2, 3, ...T$ starting from $\widehat{\Phi}_0^{(n)}$. From the simulation, I can get an implied sequence of state distributions $\{\widehat{\Phi}_t^*\}_{t=0}^T$, where the initial distribution satisfies $\widehat{\Phi}_0^* = \widehat{\Phi}_0^{(n)}$.[4]

6. Using $\{\widehat{\Phi}_t^*\}_{t=0}^T$, all the aggregate allocations over the whole path such as $\{\widehat{K}_t^*\}_{t=0}^T$ can be obtained. Using the market-clearing condition, compute $\{\widehat{p}_t^*\}_{t=0}^T$.[5]

[4] In this step, I use the non-stochastic simulation method (Young, 2010).

[5] This step directly computes market-clearing prices even for a model with non-trivial market-clearing conditions. In Section 3, I use this algorithm to solve Khan and Thomas (2008) where the marginal value of consumption needs to be computed in the external loop of the model due to the non-trivial market-clearing condition. I found this technique significantly saves computation time. Further discussion on the computational gain is in Section 3.

7. From the obtained prices, check the distance between the implied sequence of $\{\widehat{p}_t^*\}_{t=0}^T$ and the guess on the prices over the whole path $\{\widehat{p}_t^{(n)}\}_{t=0}^T$ using the following metric:

$$\sup_{BurnIn \leq t \leq T-BurnIn} ||\widehat{p}_t^{(n)} - \widehat{p}_t^*||_\infty < tol$$

   Note that the distance is measured from simulations except for the burn-in periods. If the distance is smaller than the tolerance level, the algorithm is converged. Otherwise, make the following updates on the guess:[6]

$$\widehat{p}_t^{(n+1)} = \widehat{p}_t^{(n)}\psi_1 + \widehat{p}_t^*(1-\psi_1)$$

$$\widehat{v}_t^{(n+1)} = \widehat{v}_t^{(n)}\psi_2 + \widehat{v}_t^*(1-\psi_2)$$

$$\widehat{\Phi}_t^{(n+1)} = \widehat{\Phi}_t^{(n)}\psi_3 + \widehat{\Phi}_t^*(1-\psi_3)$$

   for $\forall t \in \{0, 1, 2, 3, ..., T\}$.

   With the updated guess $\{\widehat{p}_t^{(n+1)}, \widehat{v}_t^{(n+1)}, \widehat{\Phi}_t^{(n+1)}\}_{t=0}^T$, go to step 4.

(ψ_1, ψ_2, ψ_3) are the control parameters of convergence speed in the algorithm. If ψ_i is high, then the algorithm conservatively updates the guess, leading to slow convergence speed. If the equilibrium dynamics are almost linear due to strong general equilibrium effect as in Krusell and Smith (1998), I found setting ψ_i around 0.8 guarantees convergence without much sacrifice in the convergence speed. However, if a model is highly nonlinear, as in the baseline model in Section 4, the convergence speed needs to be controlled to be much slower than the one in the linear models. This is because the nonlinearity can lead to a sudden jump in the realized allocations during the iteration if a new guess is too dramatically changed from

[6] In highly nonlinear aggregate dynamics, I have found that the log-convex combination updating rule marginally dominates the standard convex combination updating rule in terms of convergence speed. The log-convex combination rule is as follows:

$$log(\widehat{p}_t^{(n+1)}) = log(\widehat{p}_t^{(n)})\psi_1 + log(\widehat{p}_t^*)(1-\psi_1)$$

the last guess. Potentially, heterogeneous updating rule $\psi_i \neq \psi_j$ $(i \neq j)$ might be helpful. However, without a particular reason to do so, I assume homogenous weights throughout whole computations in this paper.

Note that in step 4c, I assume the value function's local linearity in aggregate states is given. If the local linearity of value function is significantly violated at an aggregate state, the approximation breaks down, so another approximation is needed.[7] However, this is not a common concern in the broad class of problems because the value function is generally smooth and locally linear along the aggregate states, while the policy functions might not be the case.

After the equilibrium allocations are computed over the in-sample path $\mathbb{A}$, I estimate the implied law of motion from the in-sample allocations. The law of motion can potentially take any nonlinear form. Then, using the fitted law of motion, equilibrium allocations are computed over out-of-sample paths of simulated aggregate shocks.

3 Computation accuracy

This section compares the equilibrium allocations obtained from the repeated transition method and the method in Krusell and Smith (1998). In the computation, parameters are set as in the benchmark model in Krusell and Smith (1998) without idiosyncratic shocks in the patience parameter β. For both of the algorithms, I stopped when the largest absolute difference between the simulated average capital stock and the expected average capital stock is less than 10^{-6}.

In the converged solution, the mean squared difference in the solutions between the repeated transition method and Krusell and Smith (1998) algorithm is around $2 * 10^{-4}$. It takes around 30 minutes for the repeated transition method to converge under the convergence speed parameter $\psi_1 = \psi_2 = \psi_3 = 0.8$, while it takes around 20 mins for Krusell and

[7] If there is a kink point in the value function along the individual states, the algorithm can be modified to include backward steps in the endogenous grid method.

Smith (1998) algorithm.[8] The convergence speed might change depending on the updating weight.

Figure 1 plots the expected path (Predicted) and the simulated path (Realized) of aggregate capital K_t obtained from the repeated transition method and the simulated path from Krusell and Smith (1998). [9] As can be seen from all three lines hardly distinguished from each other, the repeated transition method computes almost identical equilibrium allocations as Krusell and Smith (1998) algorithm at a slower speed. This is because the model in Krusell and Smith (1998) features linear dynamics of aggregate capital. Thus, their algorithm with the log-linear law of motion can compute the solution both fast and accurately.

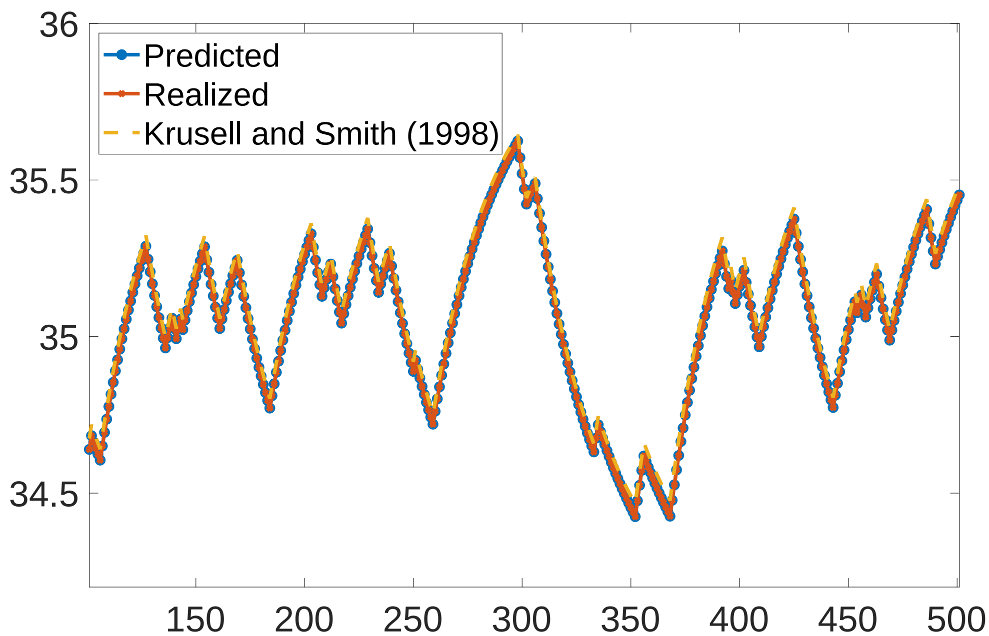

Figure 1: Computed dynamics in aggregate wealth (Krusell and Smith, 1998)

However, the repeated transition method outperforms Krusell and Smith (1997) algorithm when the market-clearing condition is non-trivial, as in the model of Khan and Thomas (2008).[10] This is because the non-trivial market-clearing condition requires an extra loop to

[8]This computation is done in 2015 MacBook Pro laptop with a 2.2 GHz quad-core processor

[9]This figure is motivated from the fundamental accuracy plot suggested in Den Haan (2010).

[10]Krusell and Smith (1997) algorithm is a variant of the algorithm in Krusell and Smith (1998), which is applicable to models with non-trivial market-clearing conditions. Khan and Thomas (2008) uses this algorithm.

find an exact market-clearing condition in each iteration.

I solve Khan and Thomas (2008) using both the repeated transition method and the Krusell and Smith (1998) algorithm with an external loop for non-trivial market-clearing condition. I stopped the iteration when the following criterion is satisfied:

$$\max\{\sup_t\{||p_t^{(n)} - p_t^{(n+1)}||\}, \sup_t\{||K_t^{(n)} - K_t^{(n+1)}||\}\} < 10^{-6}$$

Figure 2 plots the dynamics of price p_t and aggregate capital stock K_t computed from the repeated transition method and Krusell and Smith (1998) algorithm. For the allocations computed from the repeated transition method, both the predicted value and the realized values are reported. As shown from the figure, all three lines display almost identical dynamics of the price and the aggregate allocations. The mean squared difference in the solutions between the repeated transition method and Khan and Thomas (2008) is less than 10^{-5}.

In the computation of repeated transition method, I use $\psi_1 = \psi_2 = \psi_3 = 0.9$ for speed of convergence. The reason for using this conservative updating rule is because the model in Khan and Thomas (2008) features a strong general equilibrium effect; dramatic updates in the price might lead to divergence. The repeated transition method took around 20 minutes to converge on average, while Krusell and Smith (1998) algorithm converged in around 30 minutes on average. The convergence speed might change depending on the updating weight.

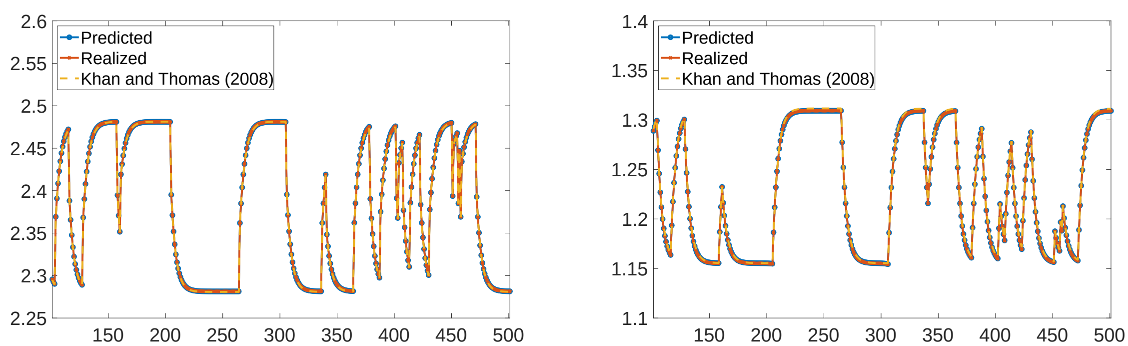

(a) Price p_t $(= 1/C_t)$

(b) Aggregate capital stocks K_t

Figure 2: Computed dynamics in aggregate capital stocks (Khan and Thomas, 2008)

In the next section, I will compare the algorithm performance between the recursive transition method and Krusell and Smith (1998) algorithm using a model with nonlinear dynamics. The previous comparisons were made for linear aggregate dynamic models, where Krusell and Smith (1998) algorithm can make a successful approximation to true aggregate dynamics. However, in nonlinear models, the accurate approximation might be hard to achieve for Krusell and Smith (1998) algorithm, while the repeated transition method successfully makes a convergence between predicted allocations and realized allocations.

4 Baseline model

In this section, I introduce a heterogeneous-firm business cycle model to study the role of corporate cash holdings on aggregate consumption.

There is a continuum of measure one of ex-ante homogenous firms that hoard cash and produces business outputs. For simplicity of the model, I assume all the firms are homogenous in terms of their capital stocks normalized at one. And the depreciation rate is assumed to be zero. At the beginning of each period, a firm i is given with a cash holding $ca_{i,t}$ and an idiosyncratic productivity level $z_{i,t}$. All firms rationally expect the future and are aware of the full distribution of each firm-level allocation and the law of motion of the aggregate states.

The business output is produced by the following Cobb-Douglas production function:

$$f(n_{i,t}, z_{i,t}; A_t) = z_{i,t} n_{i,t}^{\gamma} A_t$$

where $n_{i,t}$ is a labor demand; $\gamma < 1$ is a span of control parameter; a capital stock $k_{i,t}$ is normalized as just unity; $z_{i,t}$ and A_t are idiosyncratic and aggregate productivities, respectively. Each firm needs to pay a fixed operation cost $\xi > 0$ in each period.

The idiosyncratic and aggregate productivity shock processes, $\{z_{i,t}\}, \{A_t\}$ are specified as follows:

$$z_{i,t+1} = \rho_z z_{i,t} + \epsilon_{i,t+1}, \quad \epsilon_{i,t+1} \sim_{i.i.d} N(0, \sigma_z)$$
$$A_{t+1} = \rho_A A_t + \tilde{\epsilon}_{t+1}, \quad \tilde{\epsilon}_{t+1} \sim_{i.i.d} N(0, \sigma_A)$$

For computation, both of the shock processes are discretized by the Tauchen method.[11]

A firm earns operating profit and decides how much to distribute as a dividend $d_{i,t}$ to equity holders (a representative household). The remaining part in the operating profit after dividend payout is used to adjust cash holding, $ca_{t+1}/(1+r^{ca}) - ca_t$. The future cash holding is discounted at an internal discount rate $r^{ca} > 0$ as cash is not traded in the market across the firms. r^{ca} is an exogenous parameter and assumed to be lower than market interest rate r_t. Cash holding level is assumed to be non-negative $ca_t \geq 0$. Thus, the model imposes a standard incomplete market assumption as in Aiyagari (1994).

If a dividend is determined to be negative, then a firm is issuing equity, which incurs extra pecuniary cost $C(d_{i,t})$ (Jermann and Quadrini, 2012; Riddick and Whited, 2009). This equity issuance cost is specified as follows:

$$C(d_{i,t}) := \frac{\mu}{2}\mathbb{I}\{d_{i,t} < 0\}d_{i,t}^2$$

Thus, the net dividend is $d_{i,t} - \frac{\mu}{2}\mathbb{I}\{d_{i,t} < 0\}d_{i,t}^2$. It is worth noting that this net dividend function belongs to $\mathbb{C}^1$ class as it smoothly changes the slope at $d_{i,t} = 0$ without a kink.

If there is no equity financing cost, holding cash is not the desired option for an equity holders because it is more expensive than receiving the dividend $\left(\frac{1}{1+r^{ca}} > \frac{1}{1+r_t}\right)$. However, due to the presence of equity financing cost, a firm has a precautionary motivation to hoard cash. They save cash for the case when their business is in trouble (low z_t or low A_t), and by doing so, they reduce equity financing costs in their difficult times. In the corporate finance

[11] I use three grid points where each neighboring points are apart by one standard deviation around the mean for both processes.

literature, there has been a rich set of empirical evidence for corporates' dividend smoothing behavior (Leary and Michaely, 2011; Bliss et al., 2015). Especially, Leary and Michaely (2011) empirically showed that cash-rich firms smoothen their dividend significantly more than the others.

The recursive formulation of a firm's problem is as follows:

$$
\begin{aligned}
\text{(Firm)} \quad J(ca, z; A, \Phi) &= \max_{ca', d} \quad d - C(d) + \frac{1}{1 + r(A, \Phi)} \mathbb{E}(J(ca', z'; A', \Phi')) \\
\text{s.t.} \quad & d + \frac{ca'}{1 + r^{ca}} = \pi(z; A, \Phi) + ca \\
& ca' \geq 0, \quad \Phi' = G(\Phi, A) \\
\text{(Operating profit)} \quad \pi(z; A, \Phi) &:= \max_{n} zAn^{\gamma} - w(A, \Phi)n - \xi \\
\text{(Equity issuance cost)} \quad C(d) &:= \frac{\mu}{2} \mathbb{I}(d < 0) d^2
\end{aligned}
$$

where J is the value function of a firm; ca and z are cash holding and idiosyncratic productivity as an individual state variable; A is the aggregate productivity; Φ is the distribution of the individual state variables; w and r are wage and interest rate which are functions of aggregate state variables (A, Φ).

I close the model by introducing a stand-in household that holds equity as wealth and saves on equity. The household consumes and supplies labor and rationally expects the future aggregate states. The income sources of the household are labor income and dividend from equity holding.

The recursive formulation of the representative household's problem is as follows:

$$
\begin{aligned}
V(a; \Phi, A) &= \max_{c, a', l_H} \quad log(c) - \eta l_H + \beta \mathbb{E}^{A'} V(a'; \Phi', A') \\
\text{s.t.} \ & c + \frac{a'}{1 + r(\Phi, A)} = w(\Phi, A) l_H + a \\
& G(a, \Phi) = \Phi', \quad G_A(A) = A'
\end{aligned}
$$

where V is the value function of the household; a is wealth; c is consumption; a' is a future saving level; l_H is labor supply; w is wage, and r is the real interest rate. The household is holding the equity of firms as their wealth.

The recursive competitive equilibrium is defined based on the following market-clearing conditions:

$$\text{(Labor market)} \quad l_H(A, \Phi) = \int n(ca, z; A, \Phi) d\Phi$$
$$\text{(Equity market)} \quad a(A, \Phi) = \int J(ca, z; A, \Phi) d\Phi$$

The model does not assume a centralized market for cash holding. Therefore, r^{ca} is not endogenously determined at the market. This is a realistic assumption as a firm's cash holding is not tradable across firms. I interpret this setup as the cash holding return is determined by each firm's idiosyncratic financing status independently from the centralized capital market condition. r^{ca} is the average level of the idiosyncratic financing cost.[12]

5 Quantitative analysis

In this section, I quantitatively analyze the recursive competitive equilibrium allocations computed from the repeated transition method. For easier computation, I first normalize the firm's value function by contemporaneous consumption c_t following Khan and Thomas (2008). I define the consumption good price $p_t := 1/c_t$, so the normalized value function is $\tilde{J}_t = p_t J_t$. From the intra-temporal and inter-temporal optimality conditions of households, I have $w_t = \eta / p_t$ and $r_t = p_{t+1}/p_t$. Thus, p_t is the only price to characterize the equilibrium. The following analysis will focus on the dynamics of p_t and the aggregate cash holdings (the first moment of the distribution of cash holding).

[12] Fot simplicity, the model is abstract from the heterogeneity in the financing cost.

5.1 Calibration

The model's key parameters are the external financing cost parameter μ and the operating cost parameter ξ. The external financing cost is identified from the aggregate-level corporate cash holding-to-consumption ratio. In the moment calculation, the aggregate cash holding is obtained from the Flow of Funds.[13] Consumption is from the National Income and Product Accounts (NIPA).[14] In the model, as μ increases, the corporate cash holding-to-consumption ratio increases due to increasing precautionary motivation. The key identifying moment of the operating cost parameter is the dispersion of the cash holdings among corporates. For this, I use the time-series average of the cross-sectional standard deviation of cash holding normalized by the cross-sectional average of the cash holding.[15] As operating cost increases, the dispersion of cash increases in the model. Additionally, labor disutility cost η is calibrated to have a representative household spend a third of its hours on the labor supply. The calibrated results are summarized in Table 1. The other fixed parameters are summarized in Appendix A.1.

Parameters	Target Moments	Data	Model	Level
μ	Corporate cash holding/Consumption	17.7	17.9	0.33
ξ	Avg. of $sd(Cash_{i,t})/mean(Cash_{i,t})$	1.5	1.5	0.42
η	Labor supply hours	0.33	0.37	12.3

Table 1: Calibration target and parameters

5.2 Nonlinear business cycle

Using the repeated transition method, I compute the recursive competitive equilibrium allocations over the simulated path of aggregate shocks. In the algorithm, the interpolation of the value function (step 4b) is based on the first moment of the cash distributions (the aggregate cash holding level) following Krusell and Smith (1998) (hereafter, KS algorithm).

[13]The detailed definition of aggregate cash holding is available in Appendix A.2.
[14]In this ratio, the consumption includes both durable and non-durable consumptions.
[15]To rule out extreme outliers, I winsorize the cash holdings distribution at the top 90th percentile.

The aggregate cash holding level follows highly nonlinear dynamics in the computed outcome because the general equilibrium effect does not strongly affect each firm's cash holding demand. The price of cash holding is r^{ca} which is exogenously determined in the model because the cash holding is not allowed to be traded across the firms. In the setup where the cash is traded across the firms, the opportunity cost of cash holding $(r_t - r_t^{ca})$ shrinks close to zero. So, the aggregate cash holding is predicted to be higher on average in the alternative setup. I check this point using the computed result from the prototype KS algorithm instead of the repeated transition method.[16]

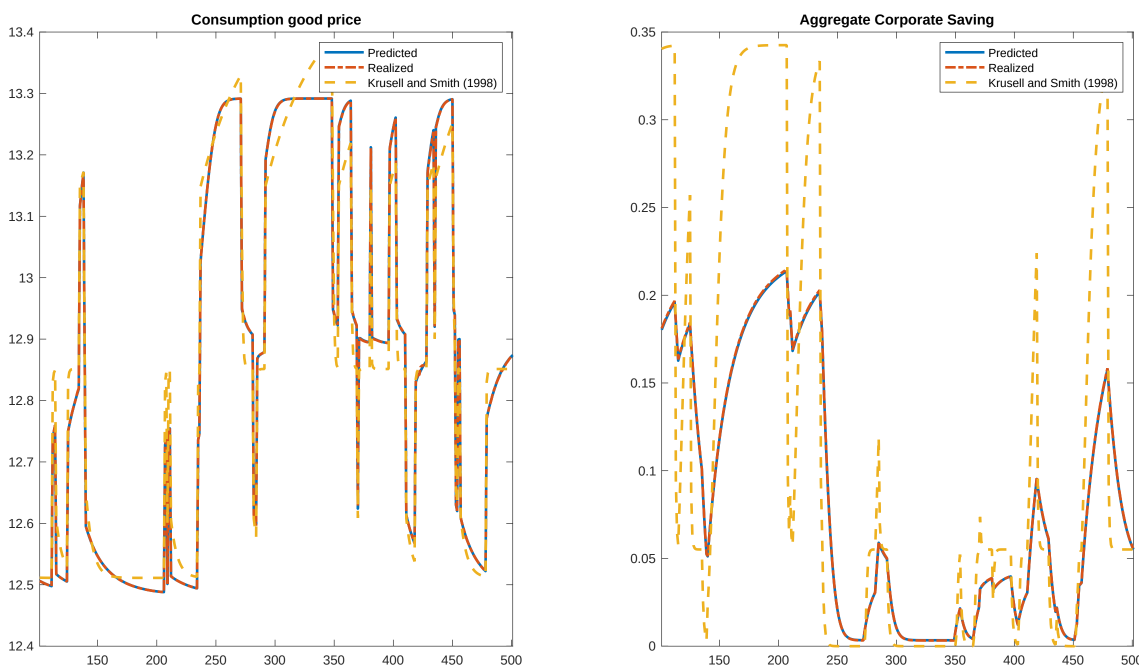

Figure 3: Aggregate fluctuations in the economy

Figure 3 plots a part of the simulated path of consumption good price and aggregate corporate saving obtained from both the repeated transition method and the prototype. The solid line plots expected allocations in the repeated transition method, and dash-dotted line plots simulated allocations in the repeated transition method. The dashed line represents the

[16]The prototype refers to the method of tracking the first moment of the state distribution, and the predicting prices based on the first moment as in Khan and Thomas (2008).

dynamics of the allocations in the prototype KS algorithm. As can be seen from the aggregate corporate saving in the right-hand side figure, the average corporate saving is higher in the KS algorithm than the repeated transition method. This is because the prototype KS algorithm assumes log-linearity in the law of motion of aggregate corporate saving and assumes that internal cash holding is linearly affected by the real interest rate.

To determine which prediction is the correct approximation to the true dynamics, I first evaluate the goodness of fitness R^2 and mean-squared error between expected dynamics and simulated dynamics on the newly simulated shock path (out-of-sample path). KS algorithm immediately gives the parametric form of the law of motion after the algorithm converges. In contrast, the repeated transition method gives the sequence of allocations which requires an extra step to fit the sequences into a parametric/non-parametric law of motion.

The repeated transition method gives R^2 of 0.9999 and mean squared error of 10^{-6} for both consumption good price and aggregate cash holding dynamics. On the other hand, the KS algorithm gives the following law of motion and goodness of fitness:[17]

$$log(CA_{t+1}) = -0.8238 + 0.9755 * log(CA_t), \quad \text{if } A_t = A_1, \text{ and } R^2 = 0.9788, \ MSE = 1.0464$$

$$log(CA_{t+1}) = -2.0397 + 0.2963 * log(CA_t), \quad \text{if } A_t = A_2, \text{ and } R^2 = 0.5532, \ MSE = 0.6598$$

$$log(CA_{t+1}) = -0.1787 + 0.8332 * log(CA_t), \quad \text{if } A_t = A_3, \text{ and } R^2 = 0.9854, \ MSE = 0.0098$$

$$log(p_t) = 2.5741 - 0.0008 * log(CA_t), \quad \text{if } A_t = A_1, \text{ and } R^2 = 0.5470, \ MSE = 0.0000$$

$$log(p_t) = 2.5508 - 0.0009 * log(CA_t), \quad \text{if } A_t = A_2, \text{ and } R^2 = 0.3410, \ MSE = 0.0000$$

$$log(p_t) = 2.5221 - 0.0042 * log(CA_t), \quad \text{if } A_t = A_3, \text{ and } R^2 = 0.8974, \ MSE = 0.0000$$

The log-linear rule of the prototype KS algorithm relies on the prices' smoothing effect on the dynamics of aggregate allocations. For example, when there is a surge of cash holding demand, the price of cash holding goes up to mitigate the surge, and vice versa for the case of decreasing cash holding demand. In numerous applications in the literature, this flattening

[17] The aggregate productivity shock is discretized by three grid points.

force from the general equilibrium has been proved to be powerful enough to guarantee the log-linear specification as the true law of motion of aggregate variable. One example is Khan and Thomas (2008) where the micro-level lumpiness is smoothed out by real interest rate dynamics. However, in the baseline model of this paper, the general equilibrium effect is missing for the cash holding demand. Thus, the log-linear prediction rule fails to capture the true law of motion in the recursive competitive equilibrium.

On top of the nonlinearity, there is another feature in the model that makes the prototype KS algorithm cannot simply address: there is a non-trivial market-clearing condition with respect to consumption good price p_t. Krusell and Smith (1997) suggested an algorithm to solve this problem by considering an external loop in the algorithm that solves market-clearing price p_t in each iteration. This algorithm is known to successfully solve the log-linear models with non-trivial market-clearing conditions such as Khan and Thomas (2008). However, due to the extra loop in each iteration, the algorithm entails high computation cost. In the repeated transition method, the price and allocations are explicitly computed at each point on the simulated path in every iteration. Therefore, the method does not require an extra loop for computing market-clearing price, so it saves great amount computation time. In the baseline model, computation time is reduced by factor of 2.[18]

5.3 Discussion: Model prediction and empirical evidence

In this section, I analyze the role of corporate cash holdings on the aggregate fluctuations using the baseline model and support the model prediction from the empirical evidence.

To investigate the role of the corporate cash holding on consumption dynamics, I analyze how the consumption volatility changes over the average lagged cash holding level. First, I residualize the aggregate consumption time-series by the recent four lagged consumptions

[18]The KS algorithm takes around one hour to compute a converged solution when the simulation length is T = 500 and the cross-sectional grid of cash holding is 50 points. However, in the repeated transition method, it takes only around 30 minutes to make a convergence. For the fair comparison, the initial guess of the KS algorithm is from the log-linear relationship implied in the initial guess of the repeated transition method.

after taking a log.

$$log(C_t) = \rho_1 log(C_{t-1}) + \rho_2 log(C_{t-2}) + \rho_3 log(C_{t-3}) + \rho_4 log(C_{t-4}) + \epsilon_t, \quad \epsilon_t \sim N(0, \sigma)$$

Then, I run the regression of the logged absolute-valued residuals on the average lagged cash holdings for the periods with $\Delta log(C_t) > 0$ (positive consumption growth) and $\Delta log(C_t) < 0$ separately (negative consumption growth).[19]

$$log(\hat{\sigma}_t) = \rho log(\overline{Cash}_{t-1}) + \eta_t, \quad \eta_t \sim N(0, \sigma_\eta)$$
$$\text{s.t. } \overline{Cash}_{t-1} = \frac{1}{4}\sum_{i=1}^{4} Cash_{t-i}$$

Table 2 reports the regression results. The residual standard deviation is negatively correlated with the average lagged cash holding in the periods with the negative consumption growth. Conversely, the residual standard deviation is positively correlated with the average lagged cash holding in the periods with the positive consumption growth. The volatility of consumption decreases by 1.1% when the lagged aggregate cash holding increases by 1% for the periods with the negative consumption growth. This relationship is visualized by a scatter plot in Figure 4.

	Dependent variable:	
	$log(\hat{\sigma}_t)$ (%)	
	Neg.	Pos.
	(1)	(2)
$log(\overline{Cash}_{t-1})$(%)	-1.075^{***}	1.694^{***}
	(0.286)	(0.337)
Constant	Yes	
Observations	197	204
R^2	0.068	0.111
Note:	*p<0.1; **p<0.05; ***p<0.01	

Table 2: Heteroskedasticity of consumption conditional on average lagged cash holding in the model

[19]The residualized consumptions are normalized by the unconditional standard deviation of the residuals.

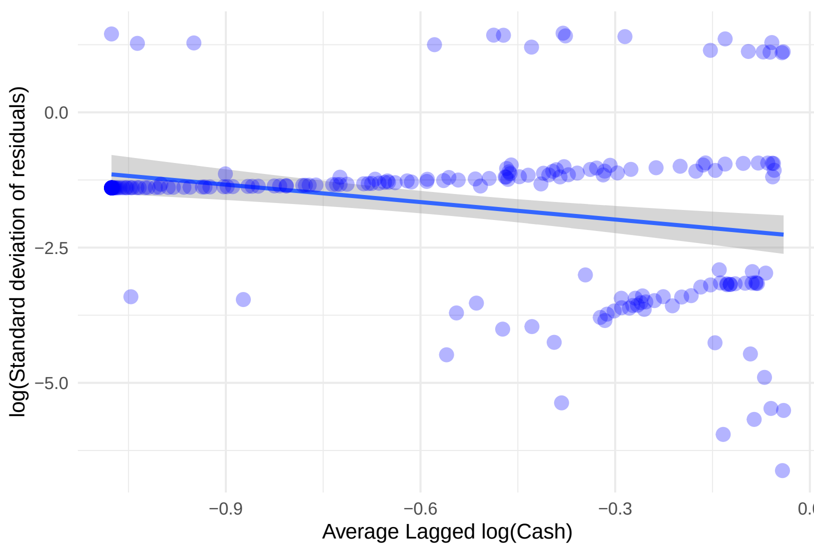

Figure 4: Scatter plot of logged residual standard deviation and average lagged cash holding conditional on $\Delta log(C_t) < 0$

Therefore, the aggregate cash holding gives a consumption buffer against a negative aggregate shock by smoothing the dividend stream in the simulated data. I support this model prediction from the macro-level data. The data is the quarterly frequency and covers from 1951 to 2018. Consumption and the total dividend of the corporate sector are from BEA National Income and Product Accounts (NIPA); the aggregate cash holding and the total asset holding are obtained from the Flow of Funds. I normalize the aggregate cash holding and dividend by the total asset holding. The aggregate consumption is detrended by HP-filter with a smoothing parameter at 1600.

Table 3 reports the regression results of conditional heteroskedasticity, using the empirical counterparts of the model variables. First, the consumption is residualized using the autoregressive process up to the fourth order.[20] The residualized consumption is regressed on average lagged normalized cash and dividend separately for pre-1980 periods and post-1980 periods. The reason for separating the two periods is because the corporate cash holding has increased dramatically after 1980, which made pre-1980 and post-1980 periods starkly

[20]As in the model counterpart, the residualized consumptions are normalized by unconditional standard deviation of the residuals.

different in terms of the size of corporate cash holdings.[21]

	Dependent variable:			
	$\hat{\sigma}_t$ (%)			
	Pre 1980	Pre 1980	Post 1980	Post 1980
	(1)	(2)	(3)	(4)
Cash_{t-1} (%)	0.558		-0.947^{**}	
	(0.448)		(0.412)	
Dividends_{t-1} (%)		0.828		-0.967^{***}
		(0.607)		(0.351)
Constant	Yes	Yes	Yes	yes
Observations	107	107	156	156
R^2	0.015	0.017	0.033	0.047
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 3: Sensitivity of consumption to aggregate TFP shock contingent on corporate cash holdings

As can be seen from Table 3, the residualized consumption display heteroskedasticity conditional on aggregate cash holding during the post-1980 periods. The greater the lagged aggregate cash holding is, the weaker responsiveness consumption displays to an exogenous aggregate shock. The same interpretation can be made to the aggregate dividend as well. These empirical results are consistent with the model prediction.

However, the model diverges from the data when it comes to the pre-1980 periods. The possible explanation for this result is that before 1980, corporate cash holding was not large enough to play an important role in dividend smoothing. Therefore, an increase in cash holding did not help consumption smoothing in the pre-1980 periods.

Figure 5 plots the scatter plot of the residualized consumption's standard deviation as a function of lagged aggregate cash holding (panel (a) and (b)), and as a function of lagged aggregate dividend (panel (c) and (d)) separately for pre-1980 and post-1980 periods. A significant negative relationship is observed from the post-1980 periods.

I further investigate whether it is a negative aggregate shock or a positive aggregate shock that drives the conditional heteroskedasticity of consumption. Here I use variation in

[21]The result is robust over other choices of the cutoff year around 1980.

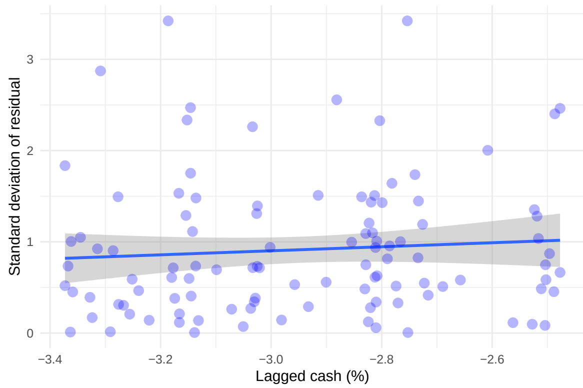

(a) Cash ratio (%) before 1980

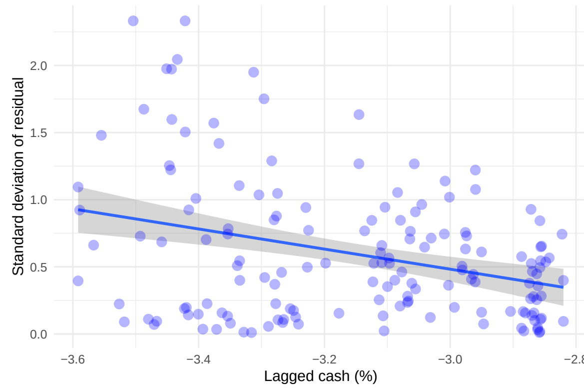

(b) Cash ratio (%) after 1980

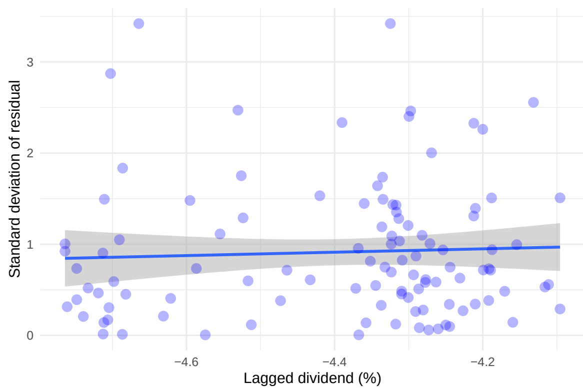

(c) Dividend ratio (%) before 1980

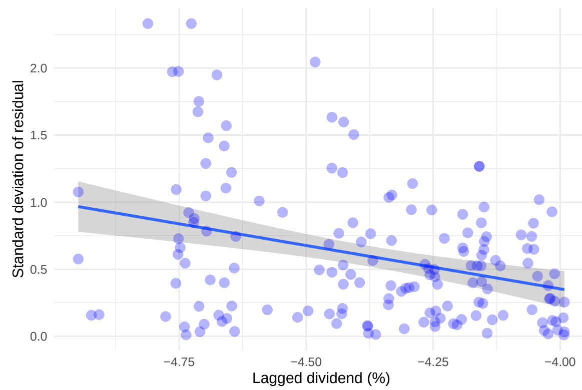

(d) Dividend ratio (%) after 1980

Figure 5: Conditional heteroskedasticity of consumption growth rate (%) before and after 1980

the Solow residual (TFP) as an aggregate shock. The TFP time-series is fitted into $AR(1)$ process to obtain the innovation in TFP, and I group observations into the positive innovation period and the negative innovation period based on the sign of TFP innovation in each period. Then, I run the following regression:

$$\frac{\Delta C_t}{C_t} = \beta_0 + \beta_1 \text{TFP Innovation}_t + \beta_2 \text{TFP Innovation}_t \times Cash_{t-1} + X_t + \epsilon_t$$

where X_t is a vector of control variables including $Cash_t$ and $Dividend_t$; TFP innovation$_t$ is normalized by its standard deviation. The coefficient of interest is β_2. If cash holding buffers consumption response, the sign of β_2 would be negative.

Table 4 reports the regression coefficients, β_1 and β_2, with standard errors in the bracket.

	Dependent variable:			
	$\Delta C_t/C_t$ (%) before 1980		$\Delta C_t/C_t$ (%) after 1980	
	Neg.	Pos.	Neg.	Pos.
	(1)	(2)	(3)	(4)
TFP Innovation$_t$ ($s.d.\%$)	-0.001	0.010*	0.013***	0.005
	(0.005)	(0.005)	(0.003)	(0.004)
TFP Innovation$_t$ $\times$ Cash$_{t-1}$ (%)	0.115	-0.090	-0.186^{**}	-0.086
	(0.088)	(0.102)	(0.079)	(0.090)
Control	Yes	Yes	Yes	yes
Constant	Yes	Yes	Yes	yes
Observations	59	53	79	77
R^2	0.264	0.409	0.409	0.186
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 4: Sensitivity of consumption: cash

As can be seen from the third column of the table, the significant consumption smoothing effect is observed only for negative TFP innovation during post-1980 periods. A similar result is obtained when the TFP innovation term interacts with the lagged dividend, as reported in Table 5. Therefore, I conclude that the model prediction of the consumption smoothing effect of corporate cash holding towards the negative aggregate shock is empirically supported from the data.

	Dependent variable:			
	$\Delta C_t/C_t$ (%) before 1980		$\Delta C_t/C_t$ (%) after 1980	
	Neg.	Pos.	Neg.	Pos.
	(1)	(2)	(3)	(4)
TFP Innovation$_t$ ($s.d.\%$)	0.006	0.009	0.015***	-0.001
	(0.008)	(0.006)	(0.003)	(0.004)
TFP Innovation$_t$ $\times$ Dividends$_t$ (%)	-0.098	-0.268	-0.696^{***}	0.101
	(0.636)	(0.463)	(0.241)	(0.282)
Control	Yes	Yes	Yes	Yes
Constant	Yes	Yes	Yes	Yes
Observations	59	53	79	77
R^2	0.241	0.404	0.429	0.177
Note:				*p<0.1; **p<0.05; ***p<0.01

Table 5: Sensitivity of consumption: dividend

6 Conclusion

This paper develops and introduces a novel algorithm to solve heterogeneous-agent models with aggregate uncertainty, which I name as repeated transition method. This method iteratively updates agents' expectations on the future path of aggregate states from the transition dynamics on a single path of simulated shocks. The algorithm runs until the expected path converges to the simulated path. In each iteration, market-clearing prices and aggregate allocations are explicitly computed at each period on the simulation path. Therefore, the method does not rely on a parametric form of the law of motion or an external loop for non-trivial market-clearing conditions.

Then, I introduce a heterogeneous-firm business cycle model where firms face a convex external financing cost and hoard cash out of precautionary motivation. Using the model, I study the business cycle implication of corporate cash holding. Cash is assumed to be an internal asset of a firm; thus, not traded across firms; and discounted at a different rate than the real interest rate in the equity market. The model features highly nonlinear dynamics of aggregate cash holdings due to the absence of general equilibrium force on the aggregate cash holding. I found the repeated transition method solves the problem more efficiently and more accurately than the existing global methods. The model predicts that the more outstanding corporate cash holding lowers the consumption volatility. This model prediction is supported by macro-level evidence of consumption heteroskedasticity conditional on the lagged aggregate cash holding.

References

Ahn, S., G. Kaplan, B. Moll, T. Winberry, and C. Wolf (2018): "When Inequality Matters for Macro and Macro Matters for Inequality," *NBER Macroeconomics Annual*, 32, 1–75, _eprint: https://doi.org/10.1086/696046.

Aiyagari, S. R. (1994): "Uninsured Idiosyncratic Risk and Aggregate Saving," *The Quarterly Journal of Economics*, 109, 659–684, publisher: Oxford University Press.

Auclert, A., B. Bardóczy, M. Rognlie, and L. Straub (2019): "Using the Sequence-Space Jacobian to Solve and Estimate Heterogeneous-Agent Models," Tech. Rep. w26123, National Bureau of Economic Research, Cambridge, MA.

Bliss, B. A., Y. Cheng, and D. J. Denis (2015): "Corporate payout, cash retention, and the supply of credit: Evidence from the 2008–2009 credit crisis," *Journal of Financial Economics*, 115, 521–540.

Boppart, T., P. Krusell, and K. Mitman (2018): "Exploiting MIT shocks in heterogeneous-agent economies: the impulse response as a numerical derivative," *Journal of Economic Dynamics and Control*, 89, 68–92.

Chien, Y., H. Cole, and H. Lustig (2012): "Is the Volatility of the Market Price of Risk Due to Intermittent Portfolio Rebalancing?" *American Economic Review*, 102, 2859–2896.

——— (2016): "Implications of heterogeneity in preferences, beliefs and asset trading technologies in an endowment economy," *Review of Economic Dynamics*, 20, 215–239.

Childers, D. (2018): "Solution of Rational Expectations Models with Function Valued States," *Working Paper*, 126.

DEN HAAN, W. J. (1997): "SOLVING DYNAMIC MODELS WITH AGGREGATE SHOCKS AND HETEROGENEOUS AGENTS," *Macroeconomic Dynamics*, 1, 355–386, publisher: Cambridge University Press.

Den Haan, W. J. (2010): "Assessing the accuracy of the aggregate law of motion in models with heterogeneous agents," *Journal of Economic Dynamics and Control*, 34, 79–99.

Duffie, D., J. Geanakoplos, A. Mas-Colell, and A. McLennan (1994): "Stationary Markov Equilibria," *Econometrica*, 62, 745.

Fernandez-Villaverde, J., S. Hurtado, and G. Nuno (2020): "Financial Frictions and the Wealth Distribution," *Working Paper*, 69.

Jermann, U. and V. Quadrini (2012): "Macroeconomic Effects of Financial Shocks," *American Economic Review*, 102, 238–271.

Kahou, M. E., J. Fernandez-Villaverde, J. Perla, and A. Sood (2021): "Exploiting Symmetry in High-Dimensional Dynamic Programming," *Working Paper*, 51.

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

Krueger, D. and H. Lustig (2010): "When is market incompleteness irrelevant for the price of aggregate risk (and when is it not)?" *Journal of Economic Theory*, 145, 1–41.

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

Krusell, P. and A. A. Smith (1997): "INCOME AND WEALTH HETEROGENEITY, PORTFOLIO CHOICE, AND EQUILIBRIUM ASSET RETURNS," *Macroeconomic Dynamics*, 1.

Leary, M. T. and R. Michaely (2011): "Determinants of Dividend Smoothing: Empirical Evidence," *Review of Financial Studies*, 24, 3197–3249.

Lee, H. (2022): "Striking While the Iron Is Cold: Fragility after a Surge of Lumpy Investments," *Working Paper*, 68.

Maliar, L., S. Maliar, and F. Valli (2010): "Solving the incomplete markets model with aggregate uncertainty using the Krusell–Smith algorithm," *Journal of Economic Dynamics and Control*, 34, 42–49.

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

——— (2010): "Approximate and Almost- Exact Aggregation in Dynamic Stochastic Heterogeneous-Agent Models," *IHS Working Paper*, 258.

Riddick, L. A. and T. M. Whited (2009): "The Corporate Propensity to Save," *The Journal of Finance*, 64, 1729–1766.

Winberry, T. (2018): "A method for solving and estimating heterogeneous agent macro models," *Quantitative Economics*, 9, 1123–1151.

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, 36–41.

A Appendix

A.1 Fixed Parameters

The fixed parameters are set at the following levels:

$$
\begin{aligned}
&\text{(Span of control)} && \gamma = 0.7; \\
&\text{(Corporate saving technology)} && r^{ca} = 0.038; \\
&\text{(Idiosyncratic shock persistence)} && \rho_z = 0.90; \\
&\text{(Idiosyncratic shock volatility)} && \sigma_z = 0.053; \\
&\text{(Aggregate shock persistence)} && \rho_A = 0.95; \\
&\text{(Aggregate shock volatility)} && \sigma_A = 0.007; \\
&\text{(Household's discount factor)} && \beta = 0.985.
\end{aligned}
$$

These fixed parameters are chosen at a reasonable level based on the literature.

A.2 Definition: Aggregate cash holding from the Flow of Funds

The aggregate cash holding is defined as sum of following items in the Flow of Funds:

- (FL103091003) Foreign deposits
- (FL103020000) Checkable deposits and currency
- (FL103030003) Time and savings deposits
- (FL103034000) Money market fund shares
- (LM103064203) Mutual fund shares
- (FL102051003) Security repurchase agreements
- (FL103069100) Commercial paper
- (LM103061103) Treasury securities

A.3 Cash holding and dividend

	Dependent variable:	
	Dividends$_t$ (%)	
	Neg.	Pos.
	(1)	(2)
Cash$_{t-1}$ (%)	0.095***	0.210***
	(0.011)	(0.028)
TFP Control	Yes	Yes
Constant	Yes	Yes
Observations	112	156
R^2	0.395	0.278
Note:	*p<0.1; **p<0.05; ***p<0.01	

Table A.1: Correlation between dividend and cash