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

Hanbaek Lee†

November 1, 2022

(click here for the latest version)

Abstract

This paper develops and tests a novel algorithm that globally solves nonlinear dynamic stochastic general equilibrium models with a high degree of accuracy. Models with heterogeneous agents can also be accurately solved using this method. 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. 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.

Keywords: Nonlinear business cycle, heterogeneous agents, stochastic dynamic programming, monotone function, state dependence, firm dynamics.

JEL codes: E32, C63, D25

*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, Xiang Fang, Zhen Huo, Miles Kimball, Eunseong Ma, Alessandro Peri, José-Víctor Ríos-Rull, Ludwig Straub, Guanyi Yang, the session participants at the Korean Economic Association International Conference and the Computational Economics and Finance Conference at Dallas, and the seminar participants at Keio University and the University of Colorado Boulder for insightful comments and discussions. All errors are my own.

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

Contents

1 Introduction

In this paper, I introduce an algorithm that solves a heterogeneous agent model with aggregate uncertainty without specifying the law of motion. I name the algorithm repeated transition method. This method provides a breakthrough to solve a broad class of heterogeneous-agent models that feature rich nonlinear aggregate dynamics accurately and globally. The methodology is also useful for representative-agent models with highly nonlinear aggregate dynamics, as highly accurate global solution can be obtained.

Under the rational expectation, heterogeneous agents are aware of the true law of motion in the aggregate states and correctly predict the future aggregate state. In contrast, there is no specific form of the law of motion known to a researcher. Moreover, 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 that 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. The dynamics of aggregate allocations subject to occasionally binding constraints are an example of such cases (Fernandez-Villaverde et al., 2020). Also, the history dependence of 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 an open-ended problem.

The repeated transition method overcomes these problems by relying on the theoretical fact of the ergodic theorem: if a simulated path of stationary shock process is long enough, all the possible allocations should be realized on the simulation path. This fact implies that state-contingent future allocations are obtainable somewhere on the simulation path as a realized outcome. Then, by properly identifying which period has the corresponding outcome to each of the expected future states, an agent's rational expectation at any time on the simulated path can be fully recovered. In the identifying step for the corresponding periods for expected future outcomes, the law of motion does not need to be specified: the only information needed for this step is a measure of similarity among the aggregate states across the periods.

For example, consider an agent is at time t, and a macroeconomist needs to come up with the rationally expected value function of period $t+1$. For each possible exogenous aggregate shock realization $s \in S$ in $t+1$, I find a period $\widetilde{t}_s + 1$ in the simulation history where the endogenous aggregate states are the closest to the ones in period $t+1$, and the aggregate shock realization is s. Then, I combine the value functions from these periods $\{\widetilde{t}_s + 1\}_{s \in S}$ to construct the expected future value function. Due to the ergodic theorem, if the simulation path is long enough, there almost surely exists a period $\widetilde{t}_s + 1$ where the endogenous aggregate allocations such as the distribution of individual states are perfectly identical to the ones in period $t+1$ among the periods where the aggregate shock realization is s. Therefore, the expected value function can be correctly constructed by combining these state-contingent value functions on the path.

In the implementation of the algorithm, due to the finite length of the simulation path, it is often difficult to find period $\widetilde{t}_s + 1$ that shares the exactly identical endogenous aggregate allocations to the ones in period $t+1$. I overcome this hurdle by approximating the expected future value function through interpolating the value functions of the periods which closely mimic period $t+1$ in terms of the endogenous aggregate states for each exogenous aggregate shock realization.

The repeated transition method is a sequence-space-based methodology, where the market-clearing prices and the expected allocations are directly computed at each point on the path (in-sample simulation). 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 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 by 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 aggregate 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' problems 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 method also provides an accurate global solution for the representative-agent models with aggregate uncertainty, when the model features highly nonlinear dynamics in the allocations. The application to the representative-agent model smoothly follows once the endogenous aggregate state variable is set to be the aggregate allocations.[1]

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 works only at a similar speed 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).

[1] In the application, the repeated transition method utilizes sufficient statistics of the distribution of the individual states, which helps the seamless application of the algorithm to the representative agent model.

Despite these difficulties in computation, the repeated transition method solves the model efficiently and accurately.

In the model, lagged aggregate cash holding significantly lowers consumption response to the negative productivity shock and intensifies consumption response to the positive productivity shock. This is because the corporate cash holding behavior helps consumption smoothing through their dividend smoothing behavior. Especially, the corporate cash stock gives an asymmetrically stronger insurance effect toward the negative TFP shock than the consumption boosting effect when the positive TFP shock hits. This model prediction of state-dependence is empirically supported by the data couterpart, and this empirical pattern is observed only after the early 1980s.[2] The fact that the corporate cash holding has dramatically increased after the early 1980s partly explains why such significant nonlinear cash-holding effect is observed only after the early 1980s.

Roadmap Section 2 explains the repeated transition method based on the model in Krusell and Smith (1998). Section 3 validates the accuracy of the repeated transition method by comparing the computed outcome with the existing well-known results 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.

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 $G(\Phi_t, A_t, A_{t+1})$. The idiosyncratic shock and the aggregate shock follow the stochastic Markov processes elaborated in

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

Krusell and Smith (1998). The Markov process is specified by the transition matrix π

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

In each element of the matrix, the first index indicate the current individual state $s \in \{u, e\}$, where u indicates an unemployed status and e indicates an employed status; the second index indicate the current aggregate state $S \in \{B, G\}$ where B indicates a bad aggregate productivity state and G indicates a good aggregate productivity state. The third and fourth indices are the future individual and aggregate states, respectively. For example, $\pi_{uB,uB}$ implies a transition probability that an unemployed worker stays unemployed in the next period when the economy is bad and stays bad in the future period.

The income sources of a household are labor income and capital income. The budget constraint of the household is as follows:

$$c_t + a_{t+1} = w_t z_t + (1 + r_t)a_t$$

The wage w_t and capital rent r_t is determined at the competitive input factor market. 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:

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

All the variables with an apostrophe indicate variables in the future period. Following the original model assumption, $z = 0.25$ when $s = u$ and $z = 1$ when $s = e$. If $S = B$, I assume

$A = 0.99$, and when $S = G$, $A = 1.01$.[3]

2.2 Intuition behind the methodology

In this section, I explain the basic intuition behind the methodology. For this, I first briefly describe the methodology. Suppose we simulate T periods of aggregate shocks $\{A_t\}_{t=0}^T$, and hypothetically the simulated path is long enough to make almost all the possible equilibrium allocations happen on the simulated path.[4] Suppose I start from the initial guess of three time series: 1) value functions, $\{V_t^{(0)}\}_{t=0}^T$, 2) distributions of individual states $\{\Phi_t^{(0)}\}_{t=0}^T$, and 3) prices $\{p_t^{(0)}\}_{t=0}^T$. Using these guesses, I solve the allocations backward from the terminal period T, and simulate the economy forward using the solution. The forward simulation generates the time series of the distribution of individual states and prices. Using these, I update the guess to move on to the next iteration, $\{V_t^{(1)}, \Phi_t^{(1)}, p_t^{(1)}\}_{t=0}^T$.

Now, suppose I've run the $(n-1)$th iteration and that I am now at the nth iteration at period t after solving the problem backward from the terminal period T until $(t+1)$ period. On the simulated aggregate state path, suppose period $t+1$ features that $S_{t+1} = G$ ($A_{t+1} = 1.01$). For a problem of an agent at t, I need to construct a rationally expected future value function $\mathbb{E}_t \widetilde{V}_{t+1}$. This step is problematic for the economist because only $V_{t+1}(\cdot, S = G)$ is available from the backward solution, while $V_{t+1}(\cdot, S = B)$ is not. This is because only one shock is realized in a period. I define this unobserved value $V_{t+1}(\cdot, S = B)$ as a counterfactual conditional value function.

In the standard state-space-based approach, this problem is handled by replacing the time index with the distribution or sufficient statistics and specifying a law of motion in these aggregate states. The counterfactual conditional value function is obtained by interpolating unconditional value functions at the predicted future state. Therefore, depending on the accuracy of this predicted future state, the accuracy of the solution is determined. However, before obtaining the solution and simulating the economy based on the solution, it is hardly understandable whether the law of motion is correctly specified or not. Importantly, if the law of motion is turned out to be incorrect, a researcher needs to come up with a new guess about the law of motion, which is subject to almost an infinite degree of freedom. To summarize the difficulties in this step in two aspects, one is about which statistics to include in the law of motion, and the other is about what parametric

[3] For brevity, I omit the explanation of the other parameter levels.

[4] In theory, an infinitely-long simulation might need to be considered, but for the illustrative purpose, I consider T-period long simulation. Later in the application, a long-enough finite simulation is used as an approximation for the infinitely-long ergodic simulation.

forms to choose for the law of motion. Unless the aggregate dynamics is well-known to be log-linear, as in Krusell and Smith (1998), this problem cannot be easily resolved.

Then, I consider a new approach where the counterfactual conditional value function is obtained from the value function of another period $\widetilde{t}+1$ in which the endogenous aggregate state is exactly the same as the period $t+1$ but the counterfactual shock is realized:

$$\Phi^{(n)}_{\widetilde{t}+1} = \Phi^{(n)}_{t+1}$$
$$S_{\widetilde{t}+1} = B \neq G = S_{t+1}$$

Then, all the aggregate states of the realized state of period $\widetilde{t}+1$ are identical to the ones in the counterfactual state of period $t+1$. Thus, we are given that

$$V^{(n)}_{\widetilde{t}+1}(\cdot, S = B) = V^{(n)}_{t+1}(\cdot, S = B).$$

Importantly, $V^{(n)}_{\widetilde{t}+1}(\cdot, S = B)$ is observed factual conditional value function available in the nth iteration. As we have both $V^{(n)}_{t+1}(\cdot, S = G)$ and $V^{(n)}_{t+1}(\cdot, S = B)(= V^{(n)}_{\widetilde{t}+1}(\cdot, S = B))$, the rationally expected future value function $\mathbb{E}_t \widetilde{V}_{t+1}$ can be correctly constructed. Even if an aggregate shock process is discretized finer than two grid points, we can construct the rationally expected future value function using the same process. Due to the ergodic theorem, if a simulated path is long enough, the existence of such period $\widetilde{t}+1$ is almost surely guaranteed.

In this new approach, a law of motion is not needed to construct the rational expected future value function. As long as the period $\widetilde{t}+1$ that contains the information about the counterfactual realization of $t+1$ is identified, the problem can be solved. For this step, tracking $\{\Phi^{(n)}_t\}_{t=0}^T$ is important, as it allows us to identify the period $\widetilde{t}+1$. In the following section, I elaborate on the detailed steps to implement the repeated transition method.

2.3 Algorithm

I simulate a single path of exogenous aggregate TFP shocks for a long-enough period T, $\mathbb{A} = \{A_t\}_{t=0}^T$, using the aggregate transition matrix π^A. So, we also have the time series of the corresponding aggregate states, $\mathbb{S} = \{S_t\}_{t=0}^T$, where $S_t \in \{B, G\}$. The aggregate transition matrix is as follows:[5]

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

For the brevity of notation, I define a price vector $p_t := (w_t, r_t)$. I define a time partition $\mathcal{T}(S)$ that groups periods with the same shock as follows.

$$\mathcal{T}_S := \{\tau | S_\tau = S\} \subseteq \{0, 1, 2, ..., T\} \text{ for } S \in \{B, G\}.$$

The pseudo algorithm of the repeated transition method is as follows:

Step 1. Guess on the paths of the prices, the value functions, and the state distributions, $\{V_t^{(n)}, \Phi_t^{(n)}, p_t^{(n)}\}_{t=0}^T$.[6]

Step 2. Solve the model backward from the terminal period T in the following sub-steps. The explanation is based on an arbitrary period t. Without a loss of generality, I assume $S_t = G$ and $S_{t+1} = G$:

2-a. Find $\widetilde{t}+1$ where the endogenous aggregate allocation in period is identical to the one in period $t+1$, but the shock realization is different from period $t+1$:

$$\widetilde{t}+1 = \arg\inf_{\tau \in \mathcal{T}_B} ||\Phi_\tau^{(n)} - \Phi_{t+1}^{(n)}||_\infty,$$

2-b. Compute the expected future value function as follows:

$$\mathbb{E}_t \widetilde{V}_{t+1} = \pi_{G,G} V_{t+1} + \pi_{G,B} V_{\widetilde{t}+1}$$

2-c. Using $\mathbb{E}_t \widetilde{V}_{t+1}$ and $p_t^{(n)}$, solve the individual agent's problem at the period t. Then, I obtain the solution $\{V_t^*, a_{t+1}^*\}$
After the taking these sub-steps for $\forall t$, $\{V_t^*, a_{t+1}^*\}_{t=0}^T$ are available.

Step 3. Using $\{a_{t+1}^*\}_{t=0}^T$, simulate forward the time series of the distribution of the individual states $\{\Phi_t^*\}_{t=0}^T$ starting from $\Phi_0^* = \Phi_0^{(n)}$.[7]

Step 4. Using $\{\Phi_t^*\}_{t=0}^T$, all the aggregate allocations over the whole path such as $\{K_t^*\}_{t=0}^T$ can be obtained. Using the market-clearing condition, compute the time series of the

[5]The transition matrix is from Krusell and Smith (1998).
[6]In practice, I use the stationary equilibrium allocations for all periods as the initial guess.
[7]In this step, I use the non-stochastic simulation method (Young, 2010).

implied prices $\{p_t^*\}_{t=0}^T$.[8]

Step 5. Check the distance between the implied prices and the guessed prices.

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

Note that the distance is measured after excluding the burn-in periods at the beginning and the end of the simulation path. This is an adjustment to handle a potential bias from the imperfect guesses on the terminal period's value function $V_T^{(n)}$ and the initial period's distribution $\Phi_0^{(n)}$.

If the distance is smaller than the tolerance level, the algorithm is converged. Otherwise, I make the following updates on the guess:[9]

$$p_t^{(n+1)} = p_t^{(n)}\psi_1 + p_t^*(1-\psi_1)$$

$$V_t^{(n+1)} = V_t^{(n)}\psi_2 + V_t^*(1-\psi_2)$$

$$\Phi_t^{(n+1)} = \Phi_t^{(n)}\psi_3 + \Phi_t^*(1-\psi_3)$$

for $\forall t \in \{0,1,2,3,...,T\}$. With the updated guess $\{V_t^{(n+1)}, \Phi_t^{(n+1)}, p_t^{(n+1)}\}_{t=0}^T$, I go back to Step 1.

(ψ_1, ψ_2, ψ_3) are the parameters of convergence speed in the algorithm. If ψ_i is high, then the algorithm conservatively updates the guess, leaving the algorithm to converge slowly. If the equilibrium dynamics are almost linear, as in Krusell and Smith (1998), I found uniformly setting ψ_i around 0.8 guarantees convergence at a fairly high 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

[8] It is worth noting that the prices here are not the market-clearing prices that are determined from the interactions between demand and supply. Rather, the prices in this algorithm are the prices implied by the market-clearing condition. 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.

[9] 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:

$$log(p_t^{(n+1)}) = log(p_t^{(n)})\psi_1 + log(p_t^*)(1-\psi_1)$$

allocations during the iteration if a new guess is too dramatically changed from the last guess. A heterogeneous updating rule $\psi_i \neq \psi_j$ $(i \neq j)$ is also helpful in cases where the dynamics of certain allocations are particularly more nonlinear than the others.

As can be seen from the convergence criterion in Step 5, the algorithm stops only when the expected allocation paths are close enough to the simulated allocation paths. Therefore, once the convergence is achieved, the accuracy of the solution is guaranteed. If the accuracy is measured in R^2 or in the mean-squared errors, as in Krusell and Smith (1998), the repeated transition method features R^2 of unity, and its mean-squared error becomes negligibly different than zero.

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.

2.4 A sufficient statistics approach

In the algorithm explained in the previous section, Step 2-a is the most demanding step as it needs to find a period $\widetilde{t}+1$ that is identical to period $t+1$ in terms of the distribution. Therefore, the similarity of the distributions across the periods needs to be measured, which is a computationally costly process.

However, if there are sufficient statistics that can perfectly represent a period's endogenous aggregate state, such as aggregate capital, the computational efficiency can be substantially improved.[10] This is because we can find period $\widetilde{t}+1$ by only comparing the distance between these sufficient statistics instead of the distributions. For example, in Krusell and Smith (1998), the aggregate capital is the sufficient statistics, which makes Step 2-a easier:

$$\widetilde{t}+1 = \arg\inf_{\tau \in \mathcal{T}_B} ||K_{\tau}^{(n)} - K_{t+1}^{(n)}||_{\infty},$$

As the algorithm relies on the ergodic theorem, a sufficiently long period of simulation is needed for accurate convergence. However, in practice, the simulation still ends in finite periods. Therefore, the period $\widetilde{t}+1$ that shares exactly identical sufficient statistics as period t+1 might not exist. For this hurdle, the following adjusted versions of Step 2-a and Step 2-b help improve the accuracy of the solution:

[10] Under which condition the sufficient statistics approach can be used is discussed in Section 2.5

2-a′. Find $\widetilde{t}^{up}+1$ where the endogenous aggregate allocation is closest to the one in period $t+1$ from above, but the shock realization is different from period $t+1$:

$$\widetilde{t}^{up}+1 = \arg \inf_{\tau\in\mathcal{T}_B,\, K_\tau^{(n)}\geq K_{t+1}^{(n)}} ||K_\tau^{(n)} - K_{t+1}^{(n)}||_\infty,$$

Similarly, find $\widetilde{t}^{dn}+1$ where the endogenous aggregate allocation is closest to the one in period $t+1$ from below, but the shock realization is different from period $t+1$:

$$\widetilde{t}^{dn}+1 = \arg \inf_{\tau\in\mathcal{T}_B,\, K_\tau^{(n)}< K_{t+1}^{(n)}} ||K_\tau^{(n)} - K_{t+1}^{(n)}||_\infty,$$

Then, we have $K^{(n)}_{\widetilde{t}^{up}+1}$ and $K^{(n)}_{\widetilde{t}^{dn}+1}$ that are closest to $K^{(n)}_{t+1}$ from above and below, respectively. Using these two, we can compute the weight ω to be used in the convex combination of value functions in the next step:

$$\omega = \frac{K^{(n)}_{t+1} - K^{(n)}_{\widetilde{t}^{dn}+1}}{K^{(n)}_{\widetilde{t}^{up}+1} - K^{(n)}_{\widetilde{t}^{dn}+1}}$$

2-b′. Compute the expected future value function as follows:

$$\mathbb{E}_t\widetilde{V}_{t+1} = \pi_{G,G}V_{t+1} + \pi_{G,B}\left(\omega V_{\widetilde{t}^{up}+1} + (1-\omega)V_{\widetilde{t}^{dn}+1}\right)$$

Step 2-a′ and Step 2-b′ construct a synthetic counterfactual conditional value function by the convex combination of the two value functions that are for the most similar periods to period $t+1$. These adjusted steps help accurately solve the problem in relatively short periods of simulation. For example, the model in Krusell and Smith (1998) can be accurately solved using only $T = 500$ periods of simulation (except for 100 burn-in periods in the beginning and the end of the simulation path).

2.5 A sufficient condition for the sufficient statistics

In this section, analyze under which condition the sufficient statistic can replace the entire distribution in the repeated transition method to allow the sufficient statistics approach (Section 2.4). In Krusell and Smith (1998), the law of motion in the entire distribution is sharply approximated by the law of motion in the aggregate capital stock. This is one example where a sufficient statistic can completely represent the infinite-dimensional object.

Likewise, various research in the literature has considered sufficient statistics to overcome the curse of dimensionality, but there has been limited theoretical understanding about when we can use such approximation. Proposition 1 provides a sufficient condition for those approaches in the application of the repeated transition method.

Proposition 1 (A sufficient condition for the sufficient statistics)**.**
For a sufficiently large T, if there exists a time series of an aggregate allocation $\{x_t\}_{t=0}^{T}$ such that for each time partition $\mathcal{T}_S = \{t|S_t = S\}, \forall S \in \{B, G\}$,

$$V_t(a,z) \text{ is strictly monotone in } x_t \text{ for } \forall(a,z), \forall t \in \mathcal{T}_S,$$

then x_t is the sufficient statistics of the endogenous aggregate state Φ_t for $\forall t$.

Proof.
See Appendix. ■

Proposition 1 states that if a time series $\{x_t\}_{t=0}^{T}$ monotonically ranks the level of value function for each individual state, x_t is the sufficient statistic of time period t. Intuition behind the proposition is as follows. Suppose a situation where a researcher is searching for a value function to build a rationally expected future value function. If a time index of the correct counterfactual period to use τ is given to a researcher, then the researcher can easily identify which value function to use, as all value functions are indexed by time. So, in this case, V_τ is trivially the one to use. Suppose $\{x_t\}_{t=0}^{T}$ satisfies the sufficient condition in Proposition 1.

Now instead of τ, suppose the level of x_τ is known to the researcher. Then, similar to the prior situation where τ is known, the researcher can identify which value function to use because the ranking information uniquely pins down the corresponding value function due to the strict monotonicity. For example, if two periods τ_0 and τ_1 share the same level of x_t, thus $x_{\tau_0} = x_{\tau_1}$, then the strict monotonicity says $V_{\tau_0} = V_{\tau_1}$. If this is not the case ($V_{\tau_0} \neq V_{\tau_1}$), then either the ergodicity or strict monotonicity assumption is violated, and this is the key idea of the proof.

To summarize the theoretical results in this section, once the ranking information is known, we can exactly pin down the value function to use. In the application steps using the interpolation (2-a′ and 2-b′ in Section 2.4), the strict monotonicity of value functions in the sufficient statistic makes it feasible to smoothly interpolate the value functions along the sufficient statistic.

The sufficient condition provides a theoretical ground to understand how a sufficient statistics approach work in the repeated transition method. In the quantitative analysis of the baseline model in Section **??**, the monotonicity is quantitatively validated for the converged solution. However, the sufficient condition is not constructive for the algorithm as it cannot be checked prior to the implementation: the condition can be verified only after the solution converges. Also, the sufficient statistics in the sequence-space-based approach do not imply that these statistics are only allocations to be considered in the law of motion in the state-space-based approach. This is because the former may not include sufficient information about the inter-temporal dynamics in the endogenous aggregate state variables.[11]

3 Accuracy of the repeated transition method

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 20 minutes for the repeated transition method to converge under the convergence speed parameter $\psi_1 = \psi_2 = \psi_3 = 0.8$; it takes around 20 mins for Krusell and Smith (1998) algorithm.[12] 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).[13] The expected path refers to $\{V_t^{(n)}, \Phi_t^{(n)}, p_t^{(n)}\}_{t=0}^{T}$ in Section 2.3, and the simulated path indicates $\{V_t^*, \Phi_t^*, p_t^*\}_{t=0}^{T}$. 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 similar speed. This is because the model in Krusell and Smith (1998) features log-linear dynamics of aggregate capital. Thus, their algorithm with the log-linear law of motion can compute

[11] When I fit the nonlinear aggregate dynamics of sufficient statistics obtained from the repeated transition method to the parametric/non-parametric law of motions in Section **??**, the fittest specification includes multiple lagged terms of the sufficient statistics. However, the sufficient statistics for each time period in the repeated transition method is just a single-dimensional aggregate allocation.

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

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

the solution both fast and accurately.

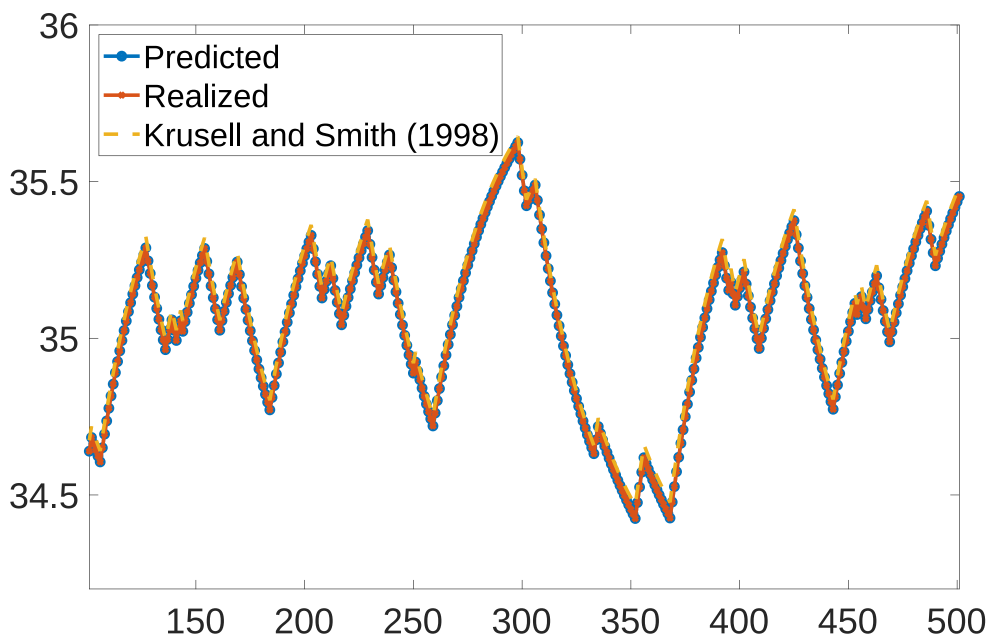

Figure 1: Computed dynamics in aggregate wealth (Krusell and Smith, 1998)

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).[14] This is because the non-trivial market-clearing condition requires an extra loop to 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:[15]

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

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

[14] 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.

[15] The terminal condition is slightly different from the one in Step 5 of Section 2.3. Likewise, the terminal condition can be flexibly adjusted based on different combinations of $V_t^{(n)}$, $\Phi_t^{(n)}$, and $p_t^{(n)}$.

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 3 to 4 hours on average. The convergence speed might change depending on the updating weight.

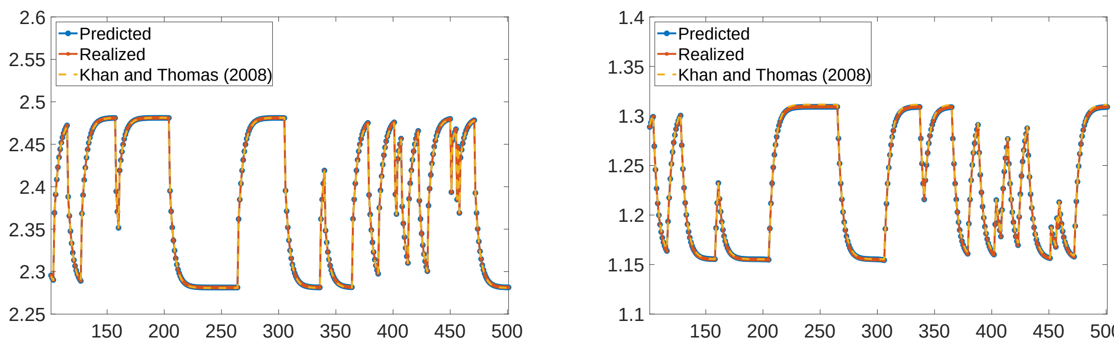

(a) Price p_t $(= 1/C_t)$

(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 analyze the business cycle implication of the rising corporate cash holding, using the heterogeneous-firm real business cycle model. In the model, the dynamics of firm-level and aggregate-level cash holding becomes highly nonlinear due to lack of general equilibrium effect on the corporate cash holding, as cash is assumed to be not priced in the market. The model is solved using the repeated transition method.

4.1 Technology

There is a continuum of measure one of ex-ante homogenous firms that hoard cash and produces business outputs. For simplicity, I assume a firm operates using only labor input. This can be understood as equivalent to a setup where a firm uses both capital and labor inputs but the optimal capital demand decision is already internalized in the labor demand decision. Consistent with this explanation, I set the labor share (equivalent to the span of control parameter) at $\gamma = 0.85$ in the quantitative analysis, which is greater than the standard labor share and captures internalized capital demand decision.

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

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

where n_{it} is a labor demand; $\gamma < 1$ is the span of control parameter; z_{it} and A_t are idiosyncratic and aggregate productivities, respectively. Each firm needs to pay a fixed operation cost $\xi > 0$ in each period.

The logged idiosyncratic productivity shock process $\{z_{it}\}$ follows an AR(1) process:

$$log(z_{it+1}) = \rho_z log(z_{it}) + \epsilon_{it+1}, \quad \epsilon_{it+1} \sim_{i.i.d} N(0, \sigma_z)$$

For computation, the idiosyncratic productivity process is discretized by the Tauchen method.[16] The stochastic aggregate productivity process is from (Krusell and Smith, 1998):[17]

$$\Gamma_A = \begin{bmatrix} 0.8750 & 0.1250 \\ 0.1250 & 0.8750 \end{bmatrix}$$
$$A_t \in \{0.99, 1.01\}.$$

[16] I discretize it using equally-spaced nine grid points within the two-standard deviation range around the mean.

[17] The repeated transition method works for a finer discretization than two grid points. However, to preserve the symmetry between the corporate cash-holding model and the household saving model (Krusell and Smith, 1998), I assume the same aggregate productivity process.

4.2 External financing cost

A firm earns operating profit and decides how much to distribute as a dividend d_{it} 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 features a standard incomplete market assumption with the borrowing limit as in Aiyagari (1994).

The corporate finance literature has theoretically and empirically investigated the reason why firms hoard cash. Among the well-known reasons, the precautionary motivation about binding financial constraint in the future has pointed out as a driver of corporate cash

If a dividend is determined to be negative, then a firm is financing through an external sources, which incurs extra pecuniary cost $C(d_{it})$ (Jermann and Quadrini, 2012; Riddick and Whited, 2009). This external financing cost is specified as follows:

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

Thus, the net dividend is $d_{it} - \frac{\mu}{2}\mathbb{I}\{d_{it} < 0\}d_{it}^2$. It is worth noting that this net dividend function belongs to $\mathbb{C}^2$ class as it smoothly changes the slope at $d_{it} = 0$ without a kink. Therefore, the standard analysis used in the models with concave differentiable utility of households smoothly applies to the model.

If there is no external financing cost, hoarding cash is not the desired option for a firm 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 the external financing cost, a firm has a precautionary motivation to hoard cash, saving for rainy days (low z_t or low A_t). In the corporate finance 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 equilibrium patterns in this model can match these empirical patterns.

4.3 Resursive formulation

At the beginning of each period, a firm i is given with a cash holding ca_{it} and an idiosyncratic productivity level z_{it}. Thus, the individual state variable s_{it} is as follows:

$$s_{it} = \{ca_{it}, z_{it}\}$$

All firms rationally expect the future and are aware of the full distribution of the firm-level state variables. The aggregate state variable S_t is as follows:

$$S_t = \{A_t, \Phi_t\}$$

where A_t is the aggregate productivity and Φ_t is the distribution of the individual state variable s_{it}.

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

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

where J is the value function of a firm; ca and z are cash holding and idiosyncratic productivity, respectively; 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 $S = \{A, \Phi\}$.

4.4 Equilibrium

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:

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

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 q is the state-contigent equity price. The household is holding the equity of firms as their wealth.

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

$$\begin{aligned} \text{(Labor market)} \quad l_H(S) &= \int n(ca,z;S) d\Phi \\ \text{(Equity market)} \quad a(A,\Phi) &= \int (J(ca,z;S) + C(d(ca,z;S))) d\Phi \end{aligned}$$

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.[18]

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).

[18] 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.[19] Consumption is from the National Income and Product Accounts (NIPA).[20] 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.[21] 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/Output (%)	10.00	9.28	0.40
ξ	Consumption/Output (%)	66.00	64.02	0.15
η	Labor supply hours	0.33	0.34	3.90

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 2.3) is based on the first moment of the cash distributions (the aggregate cash holding level) following Krusell and Smith (1998) (hereafter, KS algorithm). 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

[19] The detailed definition of aggregate cash holding is available in Appendix A.2.

[20] In this ratio, the consumption includes both durable and non-durable consumptions.

[21] To rule out extreme outliers, I winsorize the cash holdings distribution at the top 90th percentile.

$(r_t - r_t^{ca})$ 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.[22]

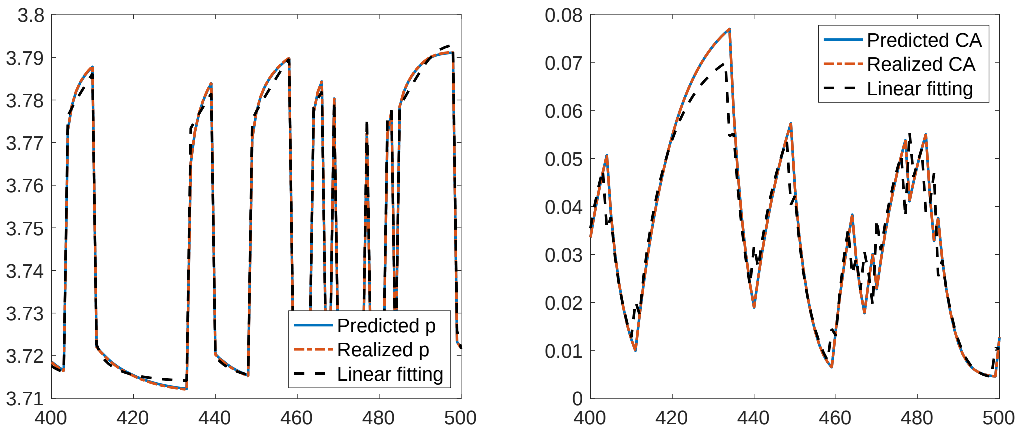

Figure 3: Aggregate fluctuations in the economy

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 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,

[22]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).

the KS algorithm gives the following law of motion and goodness of fitness:[23]

$$log(CA_{t+1}) = -0.8238 + 0.9755 * log(CA_t), \quad \text{if } A_t = A_1, \text{ and } R^2 = 0.9788, \ MSE = 1.0464$$
$$log(CA_{t+1}) = -2.0397 + 0.2963 * log(CA_t), \quad \text{if } A_t = A_2, \text{ and } R^2 = 0.5532, \ MSE = 0.6598$$
$$log(CA_{t+1}) = -0.1787 + 0.8332 * log(CA_t), \quad \text{if } A_t = A_3, \text{ and } R^2 = 0.9854, \ MSE = 0.0098$$
$$log(p_t) = 2.5741 - 0.0008 * log(CA_t), \quad \text{if } A_t = A_1, \text{ and } R^2 = 0.5470, \ MSE = 0.0000$$
$$log(p_t) = 2.5508 - 0.0009 * log(CA_t), \quad \text{if } A_t = A_2, \text{ and } R^2 = 0.3410, \ MSE = 0.0000$$
$$log(p_t) = 2.5221 - 0.0042 * log(CA_t), \quad \text{if } A_t = A_3, \text{ and } R^2 = 0.8974, \ MSE = 0.0000$$

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 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.[24]

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

[24]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.

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 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).[25]

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

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.

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

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

	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

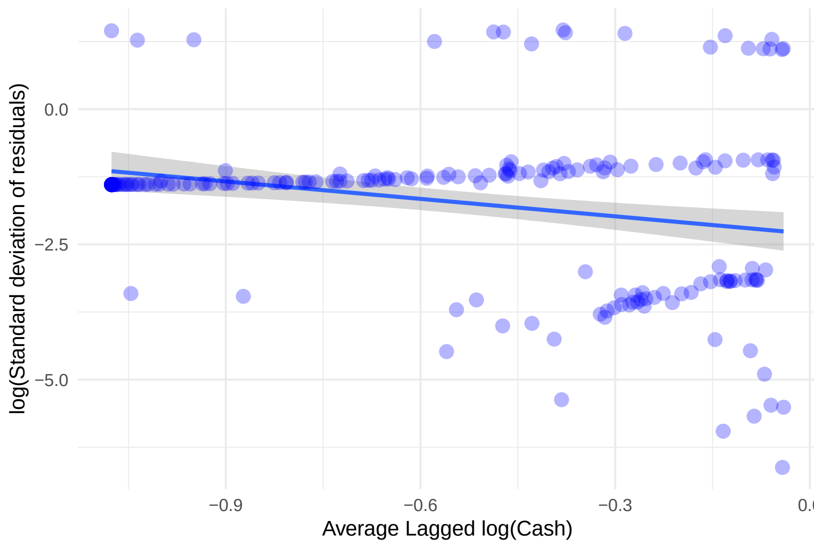

Figure 4: Scatter plot of logged residual standard deviation and average lagged cash holding conditional on $\Delta log(C_t) < 0$

the autoregressive process up to the fourth order.[26] 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 different in terms of the size of corporate cash holdings.[27]

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

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

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

	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

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 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.

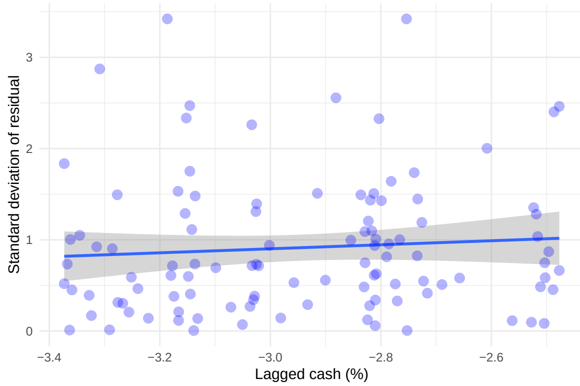

(a) Cash ratio (%) before 1980

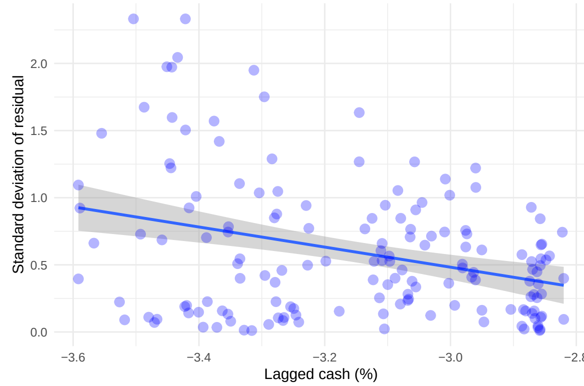

(b) Cash ratio (%) after 1980

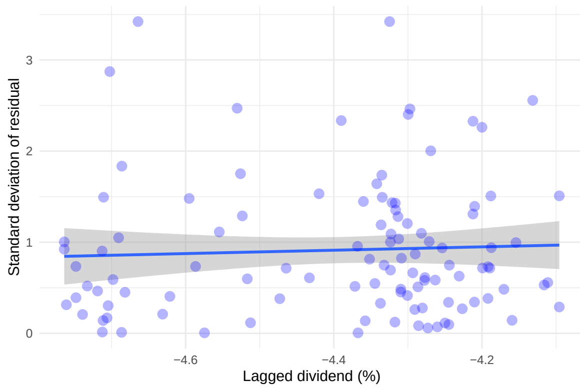

(c) Dividend ratio (%) before 1980

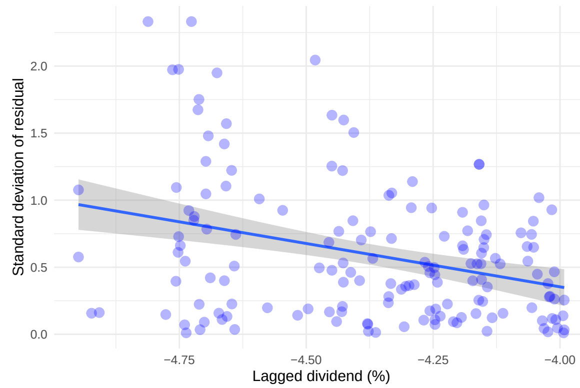

(d) Dividend ratio (%) after 1980

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

	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

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

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

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