The Inefficient Corporate Cash Buffer and the Nonlinear Business Cycle*

Miguel H. Ferreira † Timo Haber ‡ Hanbaek Lee §

June 2025

Abstract

This paper studies the macroeconomic implications of corporate cash holdings over the business cycle. We develop a heterogeneous firm business cycle model in which firms accumulate cash for precautionary reasons. Firm-level cash accumulation is highly nonlinear due to a satiation point. The nonlinearity survives aggregation, leading to a state-dependent response of aggregate allocations to TFP shocks. In periods of high cash holdings, firms can smooth dividend payouts, thereby dampening the response of aggregate consumption to adverse TFP shocks. Despite this stabilizing role, we show that the decentralized equilibrium level of cash holdings is inefficiently high due to a negative cash externality. We conclude by providing micro-level empirical evidence in support of the model's key mechanisms.

*All errors are our own.

†Queen Mary University of London. migueltchferreira@gmail.com

‡De Nederlandsche Bank. habertimo@gmail.com. Views expressed are those of the authors and do not necessarily reflect official positions of De Nederlandsche Bank or the Eurosystem.

§University of Cambridge. hanbaeklee1@gmail.com

1 Introduction

Corporate cash holdings have trended upwards across several developed economies (Chen et al., 2017). Figure 1 documents this pattern for the United States, showing the ratio of aggregate corporate cash (Flow of Funds) to nominal GDP (NIPA). It is evident that liquid savings by non-financial corporate businesses have grown significantly faster than output over the past 30 years. Although the causes of this rise are increasingly documented, its macroeconomic consequences, especially for business-cycle dynamics, remain more elusive. This paper takes a step toward filling that gap. We develop a business-cycle model with heterogeneous firms that accumulate cash for precautionary reasons and use it to quantify how corporate cash balances shape macroeconomic dynamics.

Through a calibrated version of the model, we show that corporate cash holdings significantly amplify the nonlinearity of aggregate fluctuations, with especially pronounced effects on consumption dynamics. In response to negative productivity shocks, firms with precautionary cash buffers are able to smooth dividend payouts, which in turn dampens the transmission of shocks to household consumption. This dividend-smoothing channel creates asymmetric business cycle responses: recessions are shallower and recoveries more muted when aggregate cash holdings are high. Quantitatively, we find that a cash buffer one standard deviation above the mean reduces the consumption drop in response to a negative TFP shock by 50%.

To clarify the mechanism, we analytically solve a simplified two-period version of the model. We show that firms optimally hold a strictly positive liquidity buffer: reallocating resources into cash in period 1 trades a modest opportunity cost for the option to avoid convex equity-issuance costs in period 2 (as in Hennessy and Whited (2007) and Alfaro et al. (2024)). This precautionary motive operates only up to a natural satiation point, beyond which additional saving provides no further benefit. The resulting cash buffers smooth dividend payouts and, in turn, household income. Consequently, a negative TFP shock leads to a mild consumption decline when firms enter with ample liquidity but a

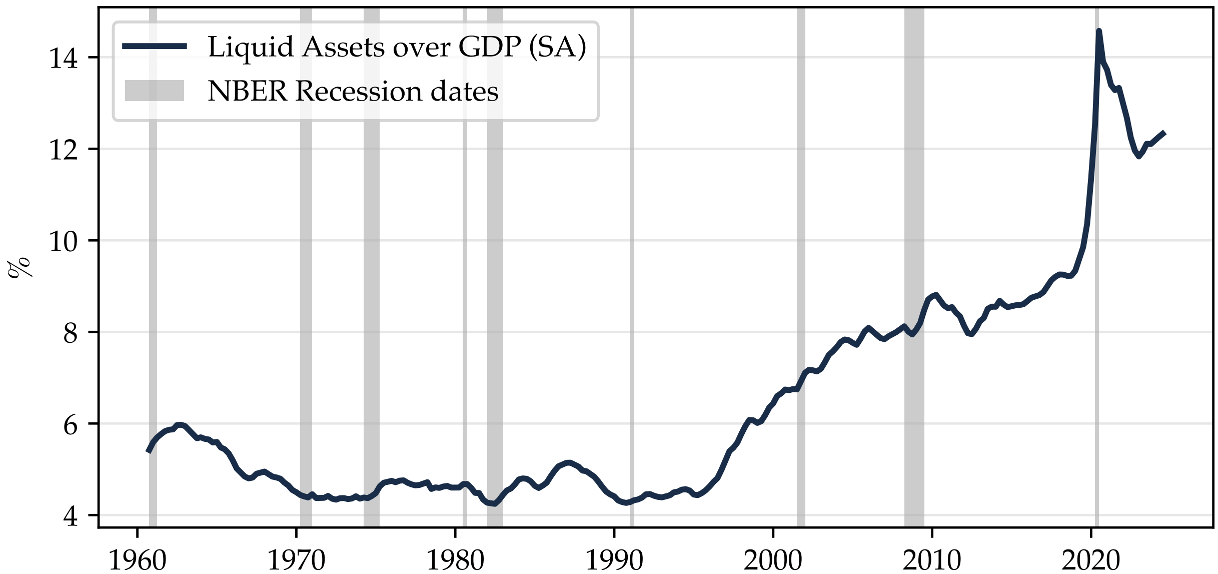

Figure 1: US nonfinancial corporate business liquid asset holdings to GDP

Note. The figure plots the time series of the liquid assets to GDP ratio between 1960Q1 and 2024Q2. The liquid asset stock data is sourced from the Flow of Funds of the Federal Reserve Board, and the nominal GDP is from National Income and Product Accounts (NIPA) from the Bureau of Economic Analysis (BEA). The definition of the measure can be found in Appendix A.1. Seasonal adjustment for liquid assets is done by extracting the trend from a multiplicative seasonal decomposition.

sharper contraction when liquid savings are low. Aggregate consumption is thus highly nonlinear and explicitly state-dependent, varying with the pre-shock distribution of corporate cash. Finally, despite the stabilizing role of cash holdings, we show that because firms do not internalize the effect of their liquidity choices on aggregate consumption, equilibrium cash holdings exceed the efficient benchmark due to the negative cash externality—it is costly to hold cash—making precautionary cash buffers inefficient.

Finally, we provide empirical support for the key mechanisms in our model. Using firm-level data from Compustat, we show that firms with higher cash-to-asset ratios reduce dividend payouts significantly less during recessions, consistent with our model's dividend smoothing mechanism. At the household level, PSID data reveals that increases in dividend income are associated with substantial increases in consumption, even after controlling for labor income, wealth, and fixed effects — confirming the relevance of dividends for household expenditure decisions.

Theoretical model. We begin with a two-period model in which firms produce, pay wages and dividends, and hold cash that earns a return below the risk-free rate. Productivity follows a stochastic process. Firms can also issue equity, subject to a quadratic issuance cost. Households consume and provide labor inelastically. In this simplified setting, we illustrate that firms hold cash out of precautionary motives, to guarantee that they avoid the equity issuance cost even if they suffer a negative productivity shock. In this scenario, household consumption responds less to a first-period TFP shock when firms hold more cash. Higher cash translates into a smaller drop in consumption when a negative TFP shock hits the economy, as dividends become smoother.

Quantitative model. We next extend the two-period model into a dynamic, general equilibrium business cycle model, where firms make decisions each period regarding labor hiring, dividend payouts, and cash accumulation. We follow Alfaro et al. (2024) and assume the return on cash is strictly less than the return on the risk-free asset. In line with the two period model, the key mechanism driving cash holdings is the presence of a convex cost of external financing (*negative dividends*), which creates a precautionary motive: firms hold cash to smooth dividend payments and avoid costly equity issuance. Additionally, we assume that cash is not tradable between firms, further reinforcing the role of internal liquidity buffers.

Although the liquidity preference of firms partially mirrors that of households, there is a fundamental difference due to the structure of their objective function. Unlike the household optimization problem, where marginal utility is strictly concave, the firm's contemporaneous objective is strictly concave only over a limited domain and remains weakly concave globally. This feature arises because internal financing (*positive dividends*) does not trigger adjustment costs, leading to a satiation point in the firm's cash hoarding decision. As a result, the intertemporal cash holdings policy exhibits a flat region, making firm-level cash accumulation strongly nonlinear. This nonlinearity is a defining characteristic that distinguishes our model from standard heterogeneous household frameworks.

Based on this insight, we examine whether micro-level nonlinearity survives aggregation and affects business cycle outcomes. The nature of the market for cash, with cash not being tradable across firms, prevents the fluctuations in the price of cash from being strong enough to wash out all the firm-level nonlinearity, leading to a macro-level nonlinearity in the TFP-driven aggregate fluctuations.

Following a negative TFP shock, the responsiveness of firm-level dividends depends on the level of cash holdings. For example, when a firm is short of cash, a negative TFP shock causes a firm to reduce dividends further than it would do when it has abundant cash stocks. This is due to a firm with little cash needing not only to pay out dividends but save cash out of concern for the future. Due to this precautionary channel, household consumption responds more to negative TFP shocks when the levels of cash holdings are low.

Despite the cash holdings beneficial effect of reducing consumption volatility, both the quantitative and theoretical model show zero socially optimal cash holdings. The cash externality, of reducing household's wealth, dominates the overall decrease in consumption volatility. This result highlights that the true cost of frictions over the business cycle are not on their effect on volatility, but rather on their impact on aggregate ergodic averages.

Empirical validation. Lastly, we empirically validate the two key model micro mechanisms. First, using Compustat data, we illustrate that firms with higher cash holdings present smoother dividends, a result qualitatively and quantitatively similar to when using model generated data. Second, using PSID data, we find households' consumption to be positively correlated with dividends. A $1 increase in dividends is associated with a 7.85 cents increase in consumption. These two empirical tests directly support the model's key mechanisms that generate the nonlinear business cycle.

Related Literature. Our paper relates to three different branches of the literature. First, to the literature studying how financial frictions amplify aggregate shocks. Bernanke et al. (1999) introduced the financial accelerator mechanism, where credit market frictions amplify economic shocks due to entrepreneurs bearing all aggregate risk. Several papers such as Gertler and Kiyotaki (2010), Ottonello and Winberry (2020), Begenau and Salomao (2018), Cloyne et al. (2023), and Ferreira et al. (2023) explored how financial frictions amplify firm-level risk, and this increase in risk can spill over to aggregate consumption in response to negative shocks. These insights underscore that the concentration of risk in leverage firms/entrepreneurs is essential to the propagation of shocks and the response of aggregate variables.

We contribute to this literature by showing that, even absent changes in leverage or borrowing costs, firms' precautionary cash hoarding behavior can generate state-dependent consumption responses to TFP shocks through a dividend smoothing channel.

Second, we contribute to the extensive literature that studies corporate cash hoarding as a function of firm-specific risk, financial constraints, and macroeconomic uncertainty. Opler et al. (1999) established that firms accumulate cash as a hedge against adverse financial conditions, according to the results found by Acharya et al. (2007), who emphasize the role of cash in reducing firms' reliance on costly external funding. Bates et al. (2009) highlighted the growing trend of cash holdings among US firms due to increased volatility in cash flows. Riddick and Whited (2009) argued that firms may accumulate cash even in the absence of financing frictions, as cash holdings optimize investment decisions over time. Chen et al. (2017) further demonstrated that the increase in profits together with dividend stickiness explains the increase in corporate savings.

Our model builds on these insights by linking precautionary cash holdings to macroeconomic amplification, demonstrating how financing frictions can endogenously generate state-dependent consumption dynamics through dividend smoothing.

Third, we relate to the extensive literature on business cycle heterogeneous agent models. While heterogeneous agent models have been widely used to study macroeconomic fluctuations, early work by Aiyagari (1994) and Krusell and Smith (1998) showed that these models exhibit behavior that is largely linear at the aggregate level. Specifically, despite the presence of idiosyncratic shocks, the household savings policy function remains smooth and linear, leading to approximately linear aggregate dynamics. On the firm side, Khan and Thomas (2008) demonstrated that investment policy functions are highly nonlinear due to fixed adjustment costs and capital irreversibility. However, their findings suggest that in a general equilibrium setting, the elasticity of the interest rate offsets firm-level nonlinearities, causing any potential nonlinearities at the micro-level to wash out in the aggregate.

We contribute to this literature by showing that, contrary to prior findings, nonlinear firm behavior—specifically, cash hoarding driven by financing frictions—can survive aggregation and generate significant nonlinearities at the macro level.

Structure The rest of the paper is organized as follows: Section 2 describes and discusses the model mechanisms in a two period setting, Section 3 introduces the business cycle model, in Section 4 we compare the nonlinearity of the cash holdings policy function to a standard Aiyagari (1994) model, Section 5 illustrates the aggregate implications of this micro level nonlinearity, Section 6 presents empirical validation of the micro level mechanisms, and Section 7 concludes.

2 A simple theory on cash holding

This section introduces a simple two-period model that analytically illustrates the main mechanism of the paper and the source of inefficiency. Households, who own firms, receive labor income and dividends in both periods and consume. Firms produce in both periods. In period one, they choose dividend payouts and cash holdings for the next period. In period two, firms distribute all remaining resources as dividends. This model illustrates why firms hold cash and how cash holdings have direct consequences for consumption of households.

Firm The firm enters period 1 with initial cash holdings of n_0. The firm produces $\widetilde{A}_1$ units of output, pays wages w_1, distributes dividends d_1 and decides how much cash n_1 to carry into next period. In period two, the firm produces $\widetilde{A}_2$ units of output, pays wages w_2 and distributes is net worth as dividends d_2 before exiting the market. Productivity in both periods follows the stochastic process:

$$\widetilde{A}_t = \begin{cases} \bar{A} + \Delta & \text{with probability } 1/2, \\ \bar{A} - \Delta & \text{with probability } 1/2. \end{cases} \tag{1}$$

In either period, if the firm issues equity (i.e. $d_t < 0, \forall t = 1,2$), the firm incurs a quadratic equity issuance cost $\mathcal{C}(d)$, which is paid by the shareholder.[1] The problem of the firm is:

$$\max_{d_1,d_2,n_1} \quad d_1 - \mathcal{C}(d_1) + \mathbb{E}\left[m(\widetilde{A}_2)(\widetilde{d}_2 - \mathcal{C}(\widetilde{d}_2))\right], \tag{2}$$

$$\text{s.t.} \quad d_1 = A_1 - w_1 - n_1\beta^n + n_0, \tag{3}$$

$$\widetilde{d}_2 = \widetilde{A}_2 - w_2 + n_1, \tag{4}$$

$$\mathcal{C}(d) = \frac{\mu}{2} d^2 \mathbb{I}(d < 0), \tag{5}$$

where all the variables with a tilde represent random variables and $m(\widetilde{A}_2) = \beta c_1 / \widetilde{c}_2(\widetilde{A}_2)$ is the stochastic discount factor. β^n is the price of cash, assumed to be greater than or equal to the SDF. A wedge between the return on cash and the risk-free rate is introduced by assuming $\beta^n > \mathbb{E}m(\widetilde{A}_2)$ (Cooley and Quadrini, 2001; Alfaro et al., 2024).

The equity-issuance cost $\mathcal{C}(d)$ is the mechanism that gives firms a precautionary motive to stockpile liquidity: by holding cash today they can avoid paying the quadratic cost

[1] We here consider only a quadratic cost function, similar to the quantitative model in Section 3, to guarantee the function is continuously differentiable across the whole domain. Appendix B.1 extends the equity-issuance cost to include a linear component, and Proposition 1 remains valid under this specification.

tomorrow should revenues fall short. The next result formalizes this link between the size of the issuance friction and the cash buffer the firm chooses to carry:

Proposition 1 (Optimal cash holding)**.** *The firm's optimal cash holding n_1^* is (weakly) increasing in the equity-issuance-cost parameter μ and is bounded above by $w_2 + \Delta - A_1$.*

Proof. See Appendix G. ■

Intuitively, a higher μ makes external equity more expensive, so the firm raises its precautionary cash balance; the upper bound reflects the point at which an extra dollar of cash no longer relaxes the period-2 financing constraint.

Household Having established how the equity-issuance cost shapes the firm's cash-holding decision, we now turn to the household in order to trace out the implications for consumption. The household lives for two periods, supplies labor inelastically and receives both labor income and firm dividends in each period. The household's problem is:

$$\max_{c_1, c_2} \quad \log(c_1) + \beta \mathbb{E} \log(\widetilde{c}_2), \tag{6}$$

$$\text{subject to:} \quad c_1 = d_1 + w_1, \tag{7}$$

$$\widetilde{c}_2 = \widetilde{d}_2 + w_2. \tag{8}$$

Although the household ultimately bears the equity-issuance cost $\mathcal{C}(d)$, as the firm deducts it before paying dividends, the cost is immediately rebated to the household and thus does not appear in the budget constraints above. Consequently, $\mathcal{C}(d)$ reallocates resources across time but does not generate a deadweight loss. Its macroeconomic impact arises instead through the firm's precautionary cash behaviour and the resulting path of dividends that enter the household's consumption. The firm's cash buffer transmits directly to households through the dividend stream. The next result shows how a larger initial cash position moderates the consumption impact of a negative productivity shock.

Corollary 1 (The state-dependent consumption response)**.**

The consumption impact of a negative TFP shock weakly decreases in n_0.

$$\frac{\partial(c_1(n_0^H, A_1 - \Delta) - c_1(n_0^L, A_1 - \Delta))}{\partial n_0^H} > 0$$

Proof.

$$\begin{aligned}\frac{\partial\left(C_1(n_0^H, A_1 - \Delta) - C_1(n_0^L, A_1 - \Delta)\right)}{\partial n_0^H} =& \frac{\partial\left(\beta_2^n\left(n_1(n_0^L, A_1 - \Delta) - n_1(n_0^H, A_1 - \Delta)\right) + n_0^H - n_0^L\right)}{\partial n_0^H} \\ =& 1 - \frac{\omega}{\omega^2 E(m(\widetilde{A}_2)) + 1} > 0\end{aligned}$$

As the term $\frac{\omega}{\omega^2 E(m(\widetilde{A}_2))+1}$ is smaller than 1 but larger than 0, the above is positive and smaller than 1. ■

This result shows that higher firm cash holdings smooth consumption by mitigating the effects of negative TFP shocks. Consequently, the consumption response becomes state-dependent, varying with the level of cash holdings.

3 Baseline model

We now proceed to include the mechanism of the preceding section into a quantitative general equilibrium model to study its aggregate consequences. There are three sectors in our economy: a production sector, a household sector and a government. The production sector is populated by a continuum of measure one of ex-ante homogeneous firms that face idiosyncratic productivity shocks. These firms can self-insure against future negative shocks that may require costly external finance using cash, in the spirit of Froot et al. (1993). The household sector features a representative household that consumes, prices the risk-free bond, owns the firms and supplies labor. The government collects taxes from firms and distributes them in a lump-sum fashion to the households. Idiosyncratic risk arises from firm-specific productivity shocks in the production sector, while aggregate risk stems from time-varying fluctuations in total factor productivity (TFP).

3.1 Technology

Firms use labor to produce the final good.[2] The production function has decreasing returns to scale and the following functional form:

$$y_{it} = A_t z_{it} l_{it}^{\gamma}, \tag{9}$$

where y_{it} denotes final good output, l_{it} labor demand, $\gamma < 1$ the span of control parameter, and z_{it} and A_t idiosyncratic and aggregate productivities, respectively. Regarding productivities, we assume that the log of idiosyncratic productivity follows an AR(1) process:

$$\log(z_{it}) = \rho_z \log(z_{it-1}) + \epsilon_{it}, \quad \epsilon_{it} \sim_{i.i.d} N(0, \sigma_z^2), \tag{10}$$

in which σ_z^2 and ρ_z denote the variance of the shock and the persistence of the productivity process, respectively. The stochastic aggregate productivity process is assumed to follow a two-state Markov process.[3] Hence, the transition matrix and set of possible aggregate productivity states are:

$$\Gamma_A = \begin{bmatrix} p(A_{t+1} = A_B | A_t = A_B) & (1 - p(A_{t+1} = A_B | A_t = A_B)) \\ (1 - p(A_{t+1} = A_G | A_t = A_G)) & p(A_{t+1} = A_G | A_t = A_G) \end{bmatrix} \tag{11}$$

$$\text{and} \quad A_t \in \{A_B, A_G\}, \tag{12}$$

[2]The absence of capital in the model is equivalent to a setup where firms use both capital and labor, but optimal capital demand is frictionlessly embedded within the labor demand decision.

[3]This simplification does not affect our results significantly. In particular, the repeated transition method, which is used to solve the model, is able to handle finer discretization than two grid points. For example, Lee (2022) uses a finer discretization (five grid points). However, to simplify the interpretation of recessions and booms, we resort to two grid points only.

where $A_B := 1 - \Delta_A$ and $A_G := 1 + \Delta_A$, and we calibrate Δ_A to match the aggregate output volatility. Due to the presence of aggregate risk and incomplete markets, there will be a time-varying distribution of firms over their idiosyncratic cash holdings n_{it} and productivity status z_{it}. Denote this distribution as Φ and its evolution over time as:

$$\Phi_{t+1}(z_{it+1}, n_{it+1}) = \mathbf{G}(\Phi_t(z_{it}, n_{it}), A_t). \tag{13}$$

where $\mathbf{G}$ denotes the transition operator that maps today's distribution and aggregate productivity into tomorrow's distribution. Finally, for ease of notation, denote the collection of aggregate state variables X_t as:

$$X_t \equiv \{\Phi_t(z_{it}, n_{it}), A_t\}. \tag{14}$$

3.2 Cash holdings and financial frictions

Cash Holdings Firms generate profits that equal revenues net of the wage bill and a fixed operation cost $\xi > 0$ in each period. They then decide how much to distribute as dividends (d_{it}) to their ultimate owners, the households. The portion of earnings retained after dividends (and any financing costs) is allocated to adjusting the firm's cash balance.[4] We assume that the evolution of cash holdings follows:

$$n_{it+1} = n_{it} + q_t^n h_{it}, \tag{15}$$

where n_{it} represents the firm's liquid asset holdings, q_t^n is the price of one unit of liquid asssets, and h_{it} denotes the net investment in liquid assets. Importantly, firms are subject to a no net borrowing constraint on cash, ensuring that cash balances remain non-negative, $n_{it+1} \geq 0$. This constraint mirrors the standard incomplete markets assumption

[4] In what follows we will, with a slight abuse of terminology, refer to liquid assets and cash holdings interchangeably.

with borrowing limits, as in Aiyagari (1994) or Huggett (1993).

Critically, we assume that the return on cash is lower than the risk-free rate, closely following Alfaro et al. (2024) and Cooley and Quadrini (2001). Specifically, we set the price as a multiple κ of the risk-free bond price, such that $q_t^n = \kappa q_t$, where $\kappa > 1$ and q_t is the risk-free bond price, ensuring that cash is not an excessively attractive store of value.

On the aggregate level, the net supply of liquid assets is assumed to follow an exogenous process linked to the price of liquid assets q_t^n. In particular, we again follow Alfaro et al. (2024) and impose a constant elasticity of supply specification:

$$N_{t+1}^S = \mathcal{H}(q_t^n)^{\frac{1}{\zeta}}, \tag{16}$$

where ζ governs the elasticity of the liquid asset supply, and $\mathcal{H} > 0$ is a scaling constant. This assumption captures the responsiveness of liquid asset supply to changes in asset prices.

External financing cost Beyond the non-negativity constraint on cash holdings, we assume that firms face external financing costs when dividends are negative. In such cases, firms incur an additional cost $C(d_{it})$ in the spirit of Jermann and Quadrini (2012) and Riddick and Whited (2009). We assume that this costs takes the following functional form:

$$C(d_{it}) = \frac{\mu}{2}\mathbb{I}\{d_{it} < 0\}d_{it}^2. \tag{17}$$

where μ governs the magnitude of the financing cost, determining how costly it is for firms to raise external funds when dividends are negative. A higher μ implies that even small equity issuances result in significant costs, discouraging firms from relying on external financing and strengthening their precautionary savings motive. The net dividend is given by $d_{it} - \frac{\mu}{2}\mathbb{I}\{d_{it} < 0\}d_{it}^2$. This function is continuously differentiable ($\mathbb{C}^1$) and concave, ensuring smooth adjustment at $d_{it} = 0$ without kinks. Consequently, the standard theory of concave household utility applies seamlessly to the model.[5]

The external financing cost captures firms' limited ability to adjust funding sources in response to financial conditions, as in Jermann and Quadrini (2012). This also aligns with the empirical literature on dividend smoothing where firms avoid drastic payout fluctuations due to managerial incentives and agency considerations (Leary and Michaely, 2011; Bliss et al., 2015). In particular, Leary and Michaely (2011) show that cash-rich firms smooth dividends significantly more than others, a pattern the model replicates.

Absent external financing costs, holding cash would not be optimal since cash earns a lower return than dividends. However, due to these costs, firms hoard cash as a precautionary measure, ensuring liquidity for adverse conditions (e.g., low z_t or low A_t). This precautionary motive leads firms to smooth dividend payouts in equilibrium, consistent with observed corporate behavior.

3.3 Recursive firm problem

We are now able to formulate the firm's problem in recursive form. At the beginning of each period, a firm i is identified by its idiosyncratic states n and z. Furthermore, firms have rational expectations and are aware of the full distribution of the firm-level state variables and its evolution. Using these pieces of the setup, the recursive formulation of

[5] Jermann and Quadrini (2012) highlight that the convexity assumption aligns with empirical findings by Altinkilic and Hansen (2000) and Hansen and Torregrosa (1992), who document that underwriting fees increase at a rising marginal cost as the offering size grows.

the firm's problem can be written as follows:

$$J(n,z;X) = \max_{n',d} \quad d - C(d) + \mathbb{E}\left[m(X;X')J(n',z';X')\right] \tag{18}$$

subject to:

$$d = \pi(z;X)(1-\tau^c) + n - q^n n',$$

$$n' \geq 0,$$

$$\pi(z;X) \equiv \max_{l} zAl^{\gamma} - w(X)l - \xi,$$

$$C(d) = \frac{\mu}{2}\mathbb{I}(d<0)d^2,$$

$$X \equiv \{\Phi, A\},$$

$$\text{and} \quad \Phi' = G(\Phi, A)$$

where J denotes the value function of a firm, d the dividend, m the stochastic discount factor used to price future payoffs, π the operational profits, τ^c the profit tax imposed by the government, w the wage and ξ the fixed production cost.

3.4 Households

The household sectors is populated by a representative household that chooses labor supply l^H, consumption c, bond holdings b and shares that households invest in firms, denoted by s. Thus, the household problem is given by:

$$V^H(b,s;X) = \max_{c,s',b',l^H} \left[\log(c) - \frac{\eta}{1+\frac{1}{\chi}}(l^H)^{1+\frac{1}{\chi}} + \beta\mathbb{E}\left[V^H(b',s';X')\right]\right] \tag{19}$$

subject to:

$$c + q(X)b' + \sum_{A'} \Gamma_{A,A'} \int m(X;X')s'd\Phi' = w(X)l^H + s + b + T(X)$$

$$X \equiv \{\Phi, A\}$$

where η is the labor disutility parameter, and χ is the Frisch elasticity parameter.

3.5 Government

The government sets a tax on profits τ^c and distribute its proceeds in a lump-sum fashion to the households, ensuring a balanced budget every period. Consequently, the government budget constraint reads as follows:

$$\int \tau^c \pi(z; X) d\Phi = T(X) \tag{20}$$

3.6 Competitive equilibrium

Definition 1. *A recursive competitive equilibrium is a set of functions*

$$\left(q, q^n, m, w, T, J, N, D, L, V^H, C, L^H, S, B, N^S, \Phi\right)$$

that solve the firm problem, household problem, government budget constraint, and clear the markets for liquid assets, labor, output and household bond holdings, as described by the following conditions:

1. *Taking q^n, m and w as given, J solves the firm problem described in* (18)*, and (N, D, L) are the associated policy functions for firms.*

2. *V^H solves* (19) *and (C, B, L^H, S) are the associated policy functions for households.*

3. *T solves the government budget constraint described in* (20).

4. *The goods market clears.*

5. *The market for shares clears $s(X) = \int [J(n, z; X) + C(D(n, z; X))] d\Phi$ as the external financing costs and aggregate firm values jointly determine the supply of equity.*

6. *The labor market clears $L^H(X) = \int L(n, z; X) d\Phi$*

7. *The liquid asset market clears* $(N^S(X))' = \int N(n, z; X)d\Phi$

8. *The evolution of the distribution is consistent with policy functions.*

9. *The bond market-clearing condition,* $B(X) = 0$ *is satisfied by Walras's law.*

4 The role of market incompleteness and the financial frictions

In this section, we analyze the individual firm's cash hoarding patterns in the stationary equilibrium. This analysis is essential to understand why the model can feature highly nonlinear dynamics under aggregate uncertainty. Due to the external financing cost, a firm uses cash as a precautionary savings instrument. However, there exists a target cash level beyond which firms do not accumulate further cash. Holding cash beyond this point becomes increasingly costly: once a firm accumulates enough cash to nearly eliminate the risk of future external financing, additional cash holdings yield returns lower than the household discount rate. This logic is formalized in Proposition 2, which establishes the existence of a target cash holding level.[6]

Proposition 2 (The existence of the target cash holding level)**.**
Suppose policy functions are non-trivial: $n'(n, z) > 0$ *and* $d(n, z) > 0$ *for some* $n > 0$*, given* z*. Then, there exists* $\overline{n}(z) > 0$ *such that* $n'(n, z) \leq \overline{n}(z)$ *for* $\forall n \geq 0$*.*

Proof. See Online Appendix. ■

Therefore, the optimal cash holding policy becomes flat once a firm's current cash stock reaches the target level.[7] Figure 2 plots the future cash holding policy function

[6] It is worth noting that the implication of Proposition 2 differs from Proposition 4 in Aiyagari (1993), which suggests that households with excessively high wealth gradually decumulate. In contrast, the target cash result here implies that a firm with an excessively large cash stock immediately reduces it to the target level.

[7] This mirrors the behavior in the consumption buffer stock model (Carroll, 1997). Despite the flat region, the policy function is smooth throughout (of class $\mathbb{C}^1$), with no kinks.

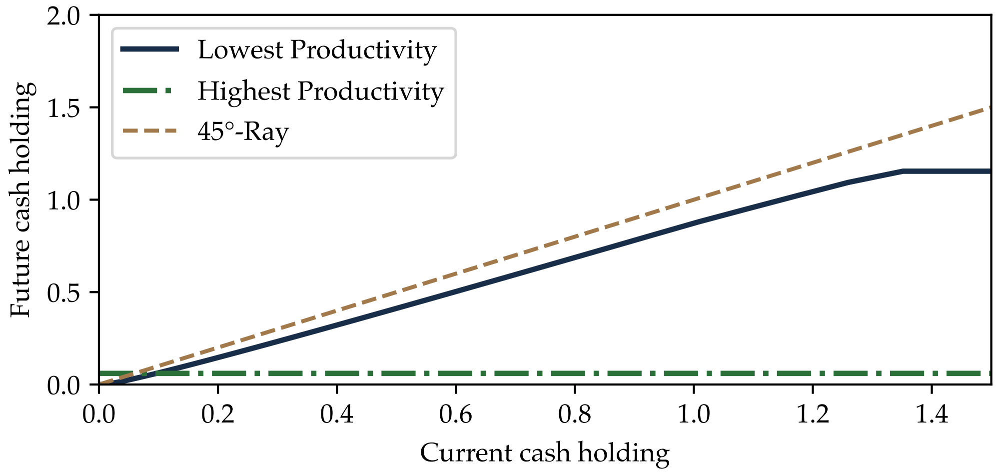

Figure 2: Cash holding policy function of individual firms

Note. The solid line plots the future cash holding policy function of the lowest-productivity firms. The dash-dotted line plots the future cash holding policy function of the highest-productivity firms.

for the lowest and highest productivity firms. For the most productive firms, the target cash level is lowest due to the persistence of high productivity—they are least concerned with future financing constraints. These firms' policy functions cross the 45-degree line, meaning they increase cash holdings until the target is reached. In contrast, the least productive firms face higher target cash levels but generate insufficient profits to save, leading them to reduce future cash holdings; their policy functions lie entirely below the 45-degree line.

The flat region in the cash holding policy function is contrasts sharply with the wealth accumulation pattern of households in Aiyagari (1994). To facilitate comparison, we define liquidity on hand, a firm-side analogue to total resources in the Aiyagari framework:

$$\text{Liquidity on hand}_t := \underbrace{\pi_t}_{\text{Liquidity from operating profit}} + \underbrace{n_t}_{\text{Cash}}.$$

Figure 3 plots the cash holdings and dividend policies in panel (a) and the future liquidity on hand in panel (b) as functions of today's liquidity on hand. This figure is the firm-side

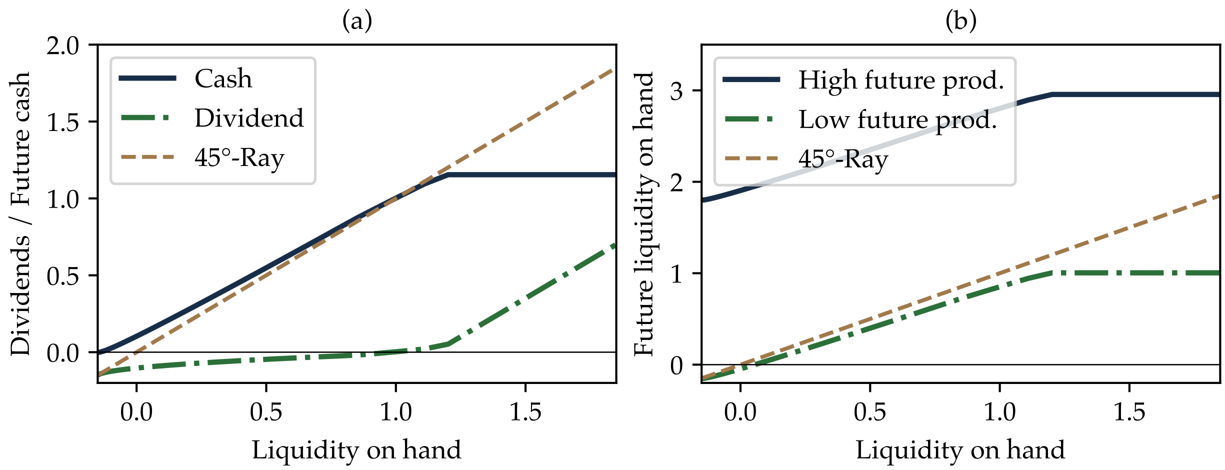

Figure 3: Dividends and cash holdings and liquidity on hand policy functions (when $z = min\mathbb{Z}$)

Notes: This figure is the firm-side counterpart of Figure I in Aiyagari (1994).

counterpart of Figure I in Aiyagari (1994).[8] For a sharp illustration, we only plot the policy functions for a firm at the lowest productivity ($z = min\mathbb{Z}$). Just as in the household case, borrowing constraints bind in certain regions, resulting in flat saving policies near the constraint. However, this region is narrow for firms and barely visible in panel (a) of Figure 3.[9]

If a firm holds little liquidity, additional liquidity is split between dividends and cash holdings, with a bias toward cash. In contrast, once a firm holds ample liquidity, any further increase is paid entirely as dividends — reflected in the dividend function becoming parallel to the 45-degree line in the high-liquidity region. In panel (b), future liquidity on hand also becomes flat at high current liquidity levels, echoing the flat region in the cash policy function from Figure 2.

Next, we show the nonlinear firm-level policy functions induce aggregate nonlinearities in the business cycle. Specifically, the degree of nonlinearity in aggregate cash dy-

[8]Dividend d is the counterpart of consumption c, and future cash holding n' is the counterpart of future wealth a_{t+1}.

[9]A similar kinked saving pattern appears in Krusell and Smith (1998). However, in both Aiyagari (1994) and Krusell and Smith (1998), the fraction of these constrained households is negligibly small. Especially, this is one of the major reasons why the aggregate dynamics in Krusell and Smith (1998) do not feature a nonlinearity.

namics depends on the fraction of firms operating in the flat region of the policy function. Thus, aggregate fluctuations in this model are tightly linked to the endogenous distribution of firm states, Φ.

5 Nonlinear business cycle and endogenous consumption Risk

In this section, we quantitatively analyze the recursive competitive equilibrium allocations computed from the global nonlinear solution method in the sequence space. For tractability, we normalize the firm's value function by contemporaneous consumption c_t following Khan and Thomas (2008). We define the marginal utility price of consumption as $p_t := 1/c_t$ and the normalized value function $\tilde{J}_t := p_t J_t$. From the household's intratemporal and intertemporal optimality conditions, we have $w_t = \eta/p_t$ and $q_t = \beta p_{t+1}/p_t$. Thus, p_t is the only price to characterize the equilibrium.

5.1 Calibration and Solution

We adopt a standard calibration strategy: some parameters are fixed based on the literature, while others are estimated to match empirical moments. The model period is a quarter, so we set the household discount factor β to 0.995, consistent with Chen et al. (2017), implying an annualized real rate of approximately 2% at the non-stochastic steady state. The Frisch labor supply elasticity is set equal to 1, consistent with previous literature, in line with Kaplan et al. (2018).[10] On the firm side, we set the span of control parameter equal to 0.930,

The two key parameters to be calibrated are the external financing cost μ and the operating cost ξ. The external financing cost is identified using the corporate cash-to-output

[10] Chetty et al. (2011) show that micro estimates imply a Frisch elasticity of 0.82. Our value is slightly above this.

Table 1: Fixed Parameters

Parameter	Description	Value	Source
Households			
β	Discount factor	0.995	Chen (2017)
χ	Frisch labor supply elasticity	1.000	Kaplan et al. (2018)
Production			
γ	Span of control	0.930	Standard calibration
κ	Cash return wedge	0.870	Cooley and Quadrini (2001)
ρ_z	Idiosyncratic shock persistence	0.861	Bachmann et al. (2013)
σ_z	Idiosyncratic shock volatility	0.075	Bachmann et al. (2013)
Aggregate			
$p(A_B\|A_B)$	Persistence of low aggregate TFP	0.875	Krusell and Smith (1998)
$p(A_G\|A_G)$	Persistence of high aggregate TFP	0.875	Krusell and Smith (1998)

Note. This table presents the fixed parameters used in the model, along with their sources. For cash return wedge, we back out the ratio wedge κ implied by the calibration used in Cooley and Quadrini (2001): $r/(1/\beta - 1)$. For the idiosyncratic shock volatility σ_z we sum both the idiosyncratic and sector specific shock volatility from Bachmann et al. (2013).

ratio. The cash quarterly data comes from the Flow of Funds (Federal Reserve), and output is measured using GDP from the National Income and Product Accounts (NIPA, BEA).[11] As μ increases, firms accumulate more cash due to heightened precautionary motives, raising the cash-to-output ratio.

The identifying moment of the operating cost parameter ξ is the consumption-to-output ratio. The consumption quarterly data is from NIPA.[12] As operating costs increase, dividend payouts fall, reducing aggregate consumption. The calibrated parameters and the corresponding moments are summarized in Table 2. Data moments are calculates as quarterly averages between 1970 and 2019.

The model on top of matching the targeted moments, equally matches well some untargeted moments. First, it generates a dividend to output ratio of 1.98%, close to the data counterpart of 1.91%. Second, at the micro level, it generates a distribution of cash

[11]See Appendix A.1 for the detailed definition of aggregate cash holdings.

[12]Consumption includes both durable and non-durable consumptions.

Table 2: Calibrated Parameters and Target Moments

Parameters	Description	Data	Model	Calibration
Targeted moments				
μ	Corporate cash holdings relative to output (%)	10.00	9.95	0.08
ξ	Consumption relative to output (%)	66.00	63.63	0.15
η	Hours worked relative to time available (%)	33.00	32.90	13.50
Δ_A	Output volatility (% *p.q.*)	1.45	1.72	0.02
Untargeted moments				
	Dividend to output (%)	1.91	1.98	

Note. This table presents the calibrated parameters along with the corresponding target moments, observed data values, model-implied values, and the level of precision achieved. The observed data values are averages over quarterly data spanning from 1970 to 2019.

holdings to output in line with the micro distribution of cash holdings to sale in the Compustat data.[13] Figure F.11 in Appendix F shows that both in the model and in the data, the correlation between cash-to-output and output is negative.

5.2 Algorithm - Repeated transition method

Despite its relatively parsimonious formulation, our model's computation of recursive competitive equilibrium presents several significant methodological challenges:

1. highly nonlinear aggregate fluctuations,
2. non-trivial market clearing conditions for both labor and consumption goods, and
3. occasionally binding constraints.

We employ the global nonlinear solution method in sequence space developed by Lee (2025), which can efficiently solve this problem. The method solves the problem backward over a long sequence of simulated exogenous aggregate states. During this backward iteration process, conditional expectations are calculated by integrating the realized

[13] For details on the data cleaning procedure, please see Appendix A.2.

value functions from previous *iterations*, thereby eliminating the need to explicitly specify aggregate laws of motion.

Furthermore, the algorithm traces multiple price vectors throughout the iteration process, updating them according to implied price levels derived from market clearing conditions. It is important to distinguish between implied price levels and market clearing prices—the former assumes either supply or demand is exogenously determined by current-iteration guesses. As iterations progress, true market clearing prices emerge asymptotically as both supply and demand converge to equilibrium values. This approach achieves market clearing only at the limit, substantially reducing computational burden.

Our model features non-trivial market clearing for both labor and consumption goods. According to Lee (2025), incorporating even a single non-trivial market clearing condition yields approximately tenfold computational efficiency gains compared to state-space methods such as the Krusell and Smith (1997) algorithm. Given our model's dual market clearing requirements, theoretical computational gains could approach two orders of magnitude.

5.3 Nonlinear business cycle

Using the repeated transition method, we compute the recursive competitive equilibrium allocations over the simulated path of aggregate shocks. The sufficient statistic used is the aggregate cash stock. The dynamics of the aggregate cash stocks are highly nonlinear for two reasons: (1) the individual firm's cash holding policy function becomes flat for high levels of individual cash stocks, as described in Section 4; (2) the general equilibrium effect does not strongly affect each firm's cash holding demand. It is because the wedge between price of cash holding and the risk-free bond is exogenously fixed at κ, as the cash is not allowed to be traded across the firms.

Figure 4 plots a sample of the simulated path of the price p_t (panel (a)) and aggregate

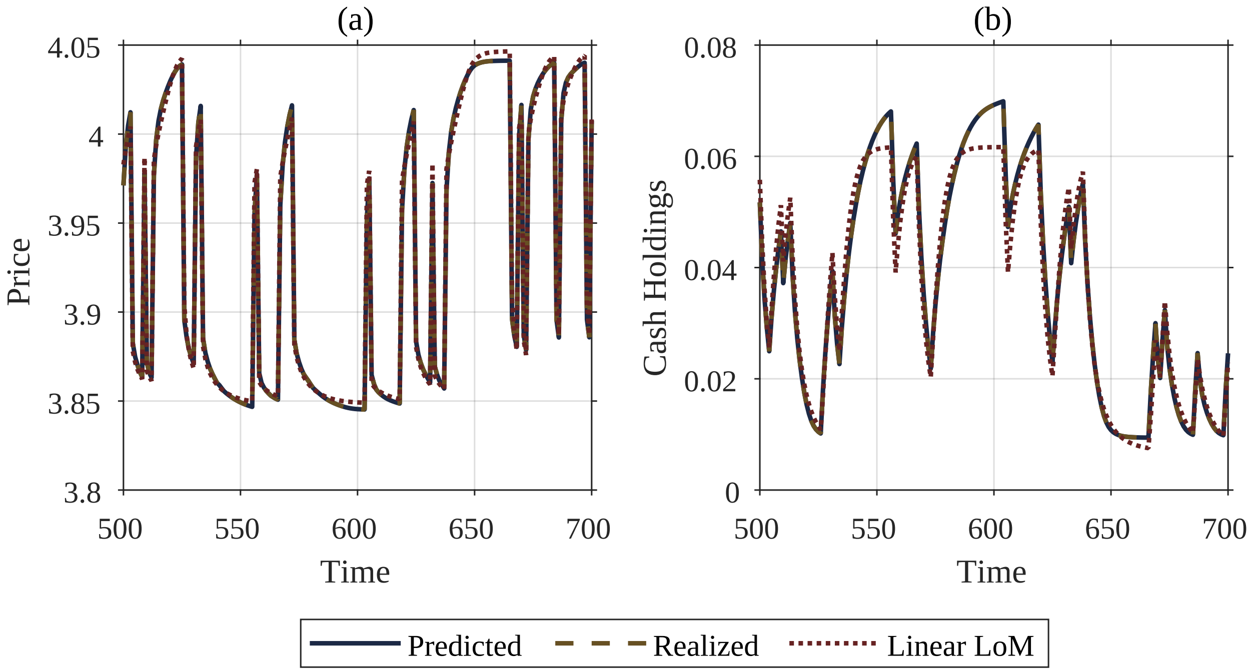

Figure 4: Aggregate fluctuations in the baseline model

Notes: The figure plots the time series of the price p_t the aggregate cash stock N_t in the baseline model. In both panels, the solid line is the predicted time series (n^{th} guess) $\{p_t^{(n)}, N_t^{(n)}\}_{t=400}^{500}$; the dash-dotted line is the realized time series $\{p_t^*, N_t^*\}_{t=400}^{500}$; the dashed line is the predicted time series implied by the linear law of motion.

corporate cash holding N_t (panel (b)) obtained from both the repeated transition method and the log-linear specification of the law of motion. The solid line plots the expected allocations (guess from the n^{th} iteration), and the dash-dotted line plots the realized allocations (simulation based on the policy in $(n+1)^{th}$ iteration) in the repeated transition method. The dashed line represents the dynamics of the allocations in the log-linear specification of the law of motion. To obtain the parameters in the log-linear specification, we fit the equilibrium allocations from the repeated transition method into the log-linear specification, and the result is as follows:

$$log(N_{t+1}) = -0.5473 + 0.8899 * log(N_t), \quad \text{if } A_t = A_B, \text{ and } R^2 = 0.9977, \ MSE = 0.0006$$

$$log(N_{t+1}) = -0.7437 + 0.7331 * log(N_t), \quad \text{if } A_t = A_G, \text{ and } R^2 = 0.9914, \ MSE = 0.0011$$

$$log(p_t) = 1.3545 - 0.0093 * log(N_t), \quad \text{if } A_t = A_B, \text{ and } R^2 = 0.9332, \ MSE = 0.0000$$

$$log(p_t) = 1.3284 - 0.0073 * log(N_t), \quad \text{if } A_t = A_G, \text{ and } R^2 = 0.9573, \ MSE = 0.0000$$

The repeated transition method generates near-perfect alignment between expected and realized paths ($R^2 \approx 1$, MSE $\approx 10^{-6}$). In contrast, the log-linear fits have notably lower R^2 and larger MSEs, indicating substantial model misspecification if linearity is imposed.

One important reason for the nonlinearity is the nature of the market for cash. As cash is not tradable across firms, the dynamics of the price of cash and aggregate cash stocks are smoothed. For example, when there is a surge of cash holding demand, the price of cash does not go up enough to mitigate the surge and vice versa for the case of decreasing cash holding demand. In many of the models in the literature, the flattening force from the general equilibrium has been proven to be powerful enough to guarantee the log-linear specification as the true law of motion. One example is Khan and Thomas (2008), where the micro-level lumpiness is smoothed out by real interest rate dynamics. However, due to the friction in the market for cash, the log-linear prediction rule fails to capture the true law of motion in this paper.

On top of the nonlinearity, there is another complication in the model that the prototype method of Krusell and Smith (1998) cannot simply address: there is a non-trivial market-clearing condition with respect to price p_t. Krusell and Smith (1997) suggests an algorithm to solve this problem by considering an external loop in the algorithm that solves the 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 costs. In contrast, the repeated transition method tracks the implied price instead of the market clearing price on the simulated path. Therefore, the method does not require an extra loop for computing the market-clearing price, so it saves a great amount of computation time.

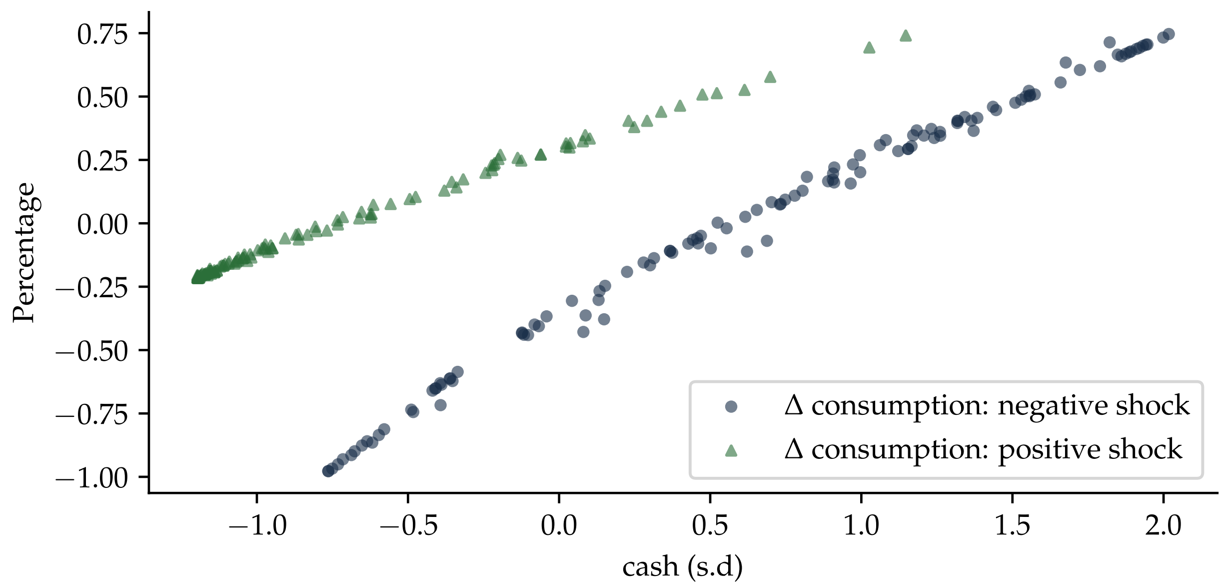

Figure 5: Endogenous state-dependence in the shock responses of consumption

Notes: The figure plots consumption responses (relative to average consumption, in percentages) to negative (blue dots) and positive (green dots) aggregate TFP shocks as a function of lagged aggregate cash stocks on the horizontal axis.

5.4 The endogenous consumption risk and state-dependent consumption responsiveness

This section explores the role of corporate cash in shaping consumption dynamics. In the model, aggregate productivity A_t switches between two states, A_G and A_B, following a persistent Markov process. We define a negative shock as a switch from A_G to A_B and a positive shock as the reverse.

A key mechanism in the model is that firm-level dividend sensitivity to TFP shocks depends on cash holdings. Firms with low cash respond more aggressively to negative shocks by cutting dividends to rebuild buffers, while cash-rich firms do so less. This translates into state-dependent household consumption response as the responsiveness of household consumption becomes dependent on dividends, and consequently on the aggregate cash stock.

Figure 5 plots the relationship between the consumption responsiveness over the aggregate cash stocks separately for negative aggregate shock (panel (a)) and positive aggregate shock (panel (b)). In this model, the magnitude of aggregate shock is uniform at $|A_G - A_B|(= 4\%$ TFP shock).[14] Therefore, if the consumption shock responses are different across the periods, it is due to the endogenous state dependence of the responsiveness rather than the shock magnitude variation. Then, we separately collect the periods of negative aggregate shock and positive aggregate shock and compute the consumption changes ΔC_t as a percentage of the stochastic steady state consumption $\mathbb{E}(C_t)$ and one-period-ahead aggregate cash stock N_{t-1} for each period. As can be seen from Figure 5, the consumption responsiveness decreases in the aggregate cash stock for the negative aggregate shock. From a similar intuition with the opposite direction, the consumption responsiveness increases in the aggregate cash stock for the positive aggregate shock. This asymmetric pattern reflects the role of precautionary motives in firm behavior and the pass-through to household income.

Table 3: Endogenous state-dependence in consumption responses to negative and positive shocks

	Dep. Var.: $log(c_t)$ $(p.p.)$	
	Neg. (1)	Pos. (2)
$Cash_t(s.d.)$	0.605	0.419
	(0.006)	(0.002)
Constant	Yes	Yes
Observations	119	119
R^2	0.988	0.996

Notes: The table reports the results of the regression of consumption responses to a negative and positive aggregate TFP shock on the lagged aggregate cash stocks using the simulated data. The first column is for the negative aggregate TFP shock, and the second is for the positive aggregate TFP shock. The numbers in the brackets are standard errors.

Table 3 reports the regression coefficient when the observations in Figure 5 are fitted into the linear regression. The numbers in the brackets are the standard errors. When the lagged aggregate cash stock increases by one standard deviation, the consumption re-

[14]The aggregate shock is defined as a shift from one productivity to the other.

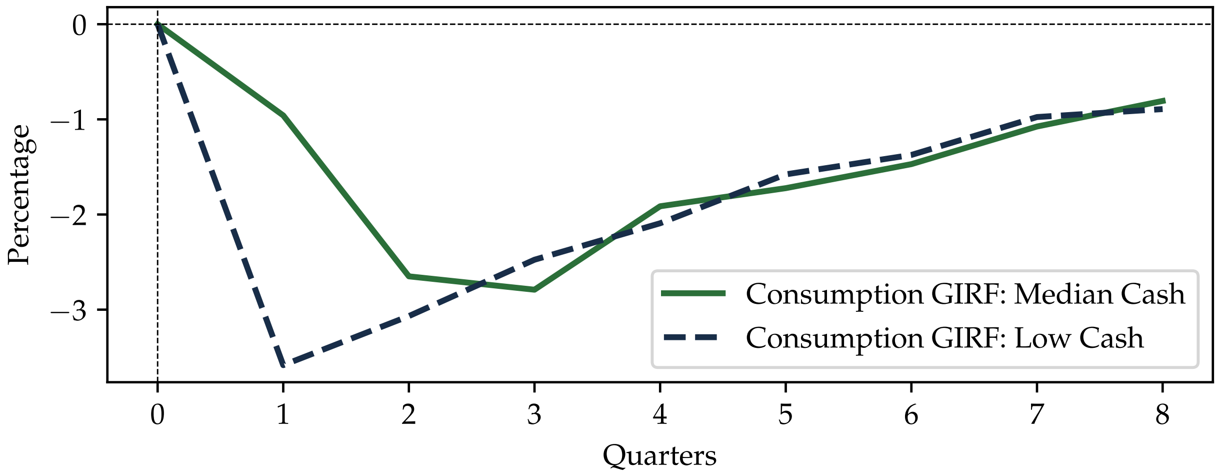

Figure 6: GIRF: Consumption responses to negative shocks for median and low cash

Notes: The figure plots the consumption % change to a negative aggregate TFP shocks when cash is at median level (green solid line) vs lowest level (blue dashed line) over the entire simulation.

sponsiveness to the negative aggregate TFP shock (-4% TFP shock) significantly decreases by 0.17 percentage points. For the positive aggregate TFP shock (+4% TFP shock), the consumption responsiveness decreases by 0.07 percentage points when the lagged aggregate cash stock increases by one standard deviation.

Therefore, the aggregate cash holding gives a consumption buffer against a negative aggregate shock by smoothing the dividend stream in the simulated data. Also, the aggregate cash holding helps a positive productivity shock solely pass down to the consumption. And the consumption buffer effect against the negative TFP shock is significantly stronger than the consumption boosting effect for the positive TFP shock in the model.

The state dependence is equally illustrated in Figure 6, which presents the generalized impulse response function (GIRF) of consumption to a negative TFP shock for different levels of cash. The solid green line presents the response when cash is equal to the median cash level over the entire simulation. The blue dashed line presents the consumption response when cash is at the lowest level over the entire simulation. Consumption drops almost three times more when cash is low vs when cash is at the median value, illustrating not only the state dependence, but equally the quantitatively meaningful effect.

Role of heterogeneity To illustrate the importance of firm heterogeneity in driving the consumption state dependence we compare the baseline model to a representative firm model. Essentially, we shut down the idiosyncratic productivity shock, keeping the market incompleteness and the aggregate productivity shocks. We keep all the parameters fixed at the baseline model value.

Figure 7 compares the results for the baseline, heterogeneous firms model, and the representative firm model. As is illustrated by the figure, the state dependence is considerably stronger in the baseline model. Two reasons explain the difference: 1) in the heterogeneous firms model, what matters, besides the average cash holdings, is the mass of firms that are on either of the flat parts of the cash holdings policy function – the constraint on the bottom part, or the satiation point on the upper part – while in the representative firm model, either of the constraints rarely binds; 2) the cash fluctuations over the cycle are much more muted in the representative firm model, as the firm can more easily insure against aggregate fluctuations. The idiosyncratic productivity component creates an additional layer against which the firms try to insure, and generates consistently higher dispersion and fluctuation of cash holdings over the cycle.

In Appendix F Figure F.12 plots the GIRF to a negative TFP shock in both the baseline and representative firm models, when cash is at the median and at the lowest point over the two simulations. Results again indicate that the state dependence is three times stronger in the heterogeneous agents model.

Efficient benchmark While firms' cash holdings diminish the overall consumption volatility over the business cycle, one can ask if firms are holding the socially optimal amount of cash. We now compare the baseline heterogeneous firms model with the social planner's, which coincides with the standard RBC model, as well as with the representative firm model.[15]

[15] For a detail description of the social planner's problem, please see Appendix E.

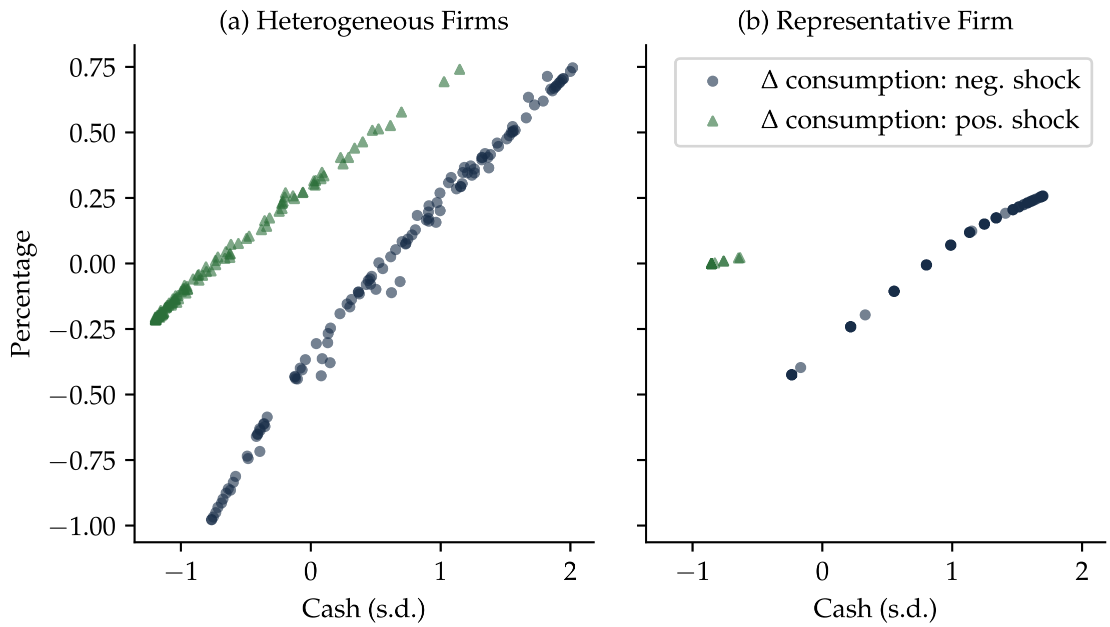

Figure 7: Endogenous state-dependence in consumption responses to negative and positive shocks in the baseline (a) and representative firm (b) models

Notes: The figure plots consumption responses (relative to average consumption, in percentages) to negative (blue dots) and positive (green dots) aggregate TFP shocks as a function of lagged aggregate cash stocks on the horizontal axis. Panel (a) plots the results for the baseline, heterogeneous firms model, while panel (b) plots the results for the representative firm model.

Table 4 shows that the social optimal cash holdings are zero. Despite the increase in consumption volatility with zero cash holdings, there is a negative cash externality, which leads the planner to choose no cash. The negative cash externality takes the form of a negative wealth effect due to the reallocation of resources towards a lower return asset, and dominates the decrease in consumption volatility. This result indicates that the true cost of frictions over of the business cycle is not on volatility is amplified, but rather on the effects on ergodic averages over the cycle. In this case, the representative firm model, which has lower cash holdings, much more closely replicates the social planner solution.

An extension of the theory in Section 2 pins down the exact mechanisms at play. While in one hand higher cash holdings diminish the consumption volatility, there is a cash externality, which reduces the overall household's wealth. As it is costly to hold cash, the social planner strictly prefers zero cash holdings to eliminate the cash externality and

Table 4: Business cycle comparison across models

	Baseline	Social planner	Rep. firm
$\text{std}(\log(C))$	0.0133	0.0188	0.0184
$\text{std}(\log(Y))$	0.0172	0.0119	0.0124
$\text{corr}(\log(C), \log(Y))$	0.9335	1.0000	0.9803
$\mathbb{E}(N)$	0.0333	0.0000	0.0018
$\text{std}(N)$	0.0198	0.0000	0.0021

Notes: The table reports business cycle statistics across the baseline heterogeneous firms model, the social planner, and the representative firm model. The table reports the standard deviation of log consumption, log output and cash holdings, as well as the average cash holdings and the correlation between loc consumption and log output.

increase overall consumption, in line with the results in Table 4. For more details, please see Appendix B.2.

6 Micro evidence of the mechanism

Our proposed mechanism relies on two key factors. First, firms using cash holdings to smooth dividends. Papers such as Opler et al. (1999) and Gao et al. (2013) show that large, publicly listed firms tend to hold more cash, and that these firms use this cash to smooth their dividend over time. Using annual Compustat data, we provide evidence in support of the model proposed mechanism. In particular, we show that firms with larger cash holdings reduce dividend payments less during recessions.

Second, dividends being an important part of households' consumption. Papers such as Baker et al. (2006), Di Maggio et al. (2020) and Bräuer et al. (2022) have documented that consumption responds positively to dividends, more than to other capital gains.[16] We here proceed to test for the correlation between consumption and dividends and how consumption reacts to changes in dividends. To do so, we use micro data from PSID

[16] Notice these papers show this result in a cross-section of households. We show this result holds within household over time.

(2005-2021), which contains bi-annual data on total household expenditure together with dividend and labor income.

6.1 Cash holdings and dividend smoothing

In this section we use annual Compustat data for the period 1980-2019 to show that firms with more liquid assets tend to smooth their dividends.[17] To test for how the firm's dividend issuance during crisis periods depends on cash holdings, we run the following local projection

$$d_{it+h} = \beta \text{rec}_t + \delta n_{it-1} + \gamma \text{rec}_t \times n_{it-1} + \Gamma_h \mathbf{X}_{it-1} + \alpha_i + \epsilon_{it} \tag{21}$$

where d_{it+h} is the log of the real dividends by firm i in year $t+h$, rec_t is an indicator variable that takes the value of 1 whenever real GDP growth is negative, n_{it-1} represents the ratio of cash and short-term investments to total assets, and X_{it-1} is a vector of controls that includes the lag of log real total assets, leverage, and real sales growth.[18] α_i represents the firm fixed effects and standard errors are clustered at the firm level.

Panel (a) on Figure 8 plots the base effect (β) of recessions on dividends on the left panel, and on the right panel the effect of cash holdings on dividends during recessions (γ). For firms with no cash holdings, dividends drop by approximately 2.5 percentage points during a recession, but recover after two years. Importantly, panel (b) illustrates that firms which have a higher cash-to-asset ratio do not reduce their dividend issuance as much after a year of negative GDP growth, as the coefficient is positive and statistically significant in both Year 0 and Year 1. One year after the recession, a 1 percentage point increase in the cash-to-assets ratio mitigates the dividend drop by 0.1 percentage points. Using the distribution of the cash-to-asset ratio in the sample, this implies that a one standard deviation increase in cash holdings reduces the dividend drop by approximately

[17] Appendix A.2 provides the full data cleaning procedure.
[18] All variables except cash are deflated using the CPI.

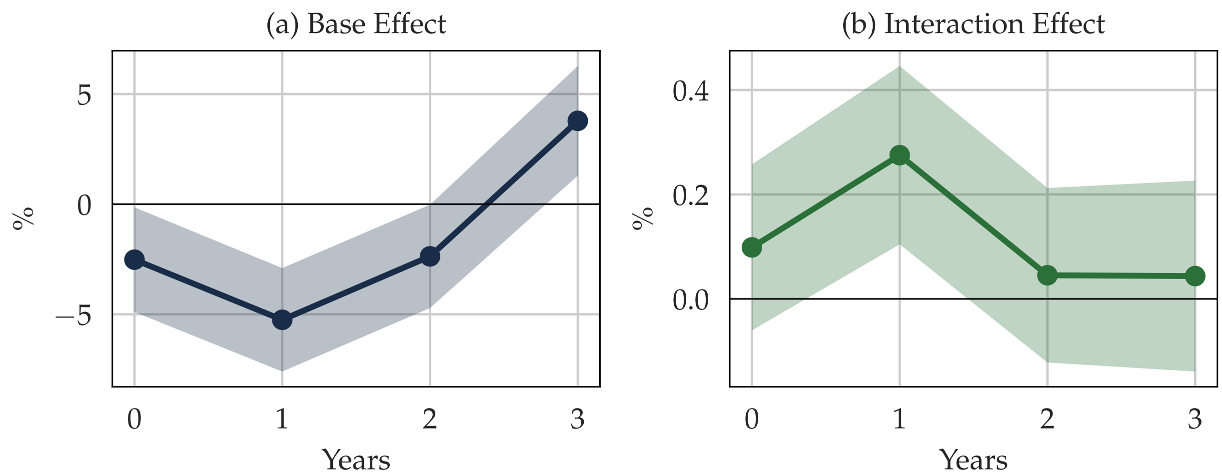

Figure 8: Dynamic dividend reaction to negative GDP and interaction with cash to assets

Note. The figure plots the impact of recessions on panel (a) up to three years after the recession. Panel (b) plots the coefficient associated with the interaction term between cash to total assets ratio and the recession indicator. Shaded bands represent 90% confidence intervals

50%.[19]

This result is supportive of the precautionary mechanism proposed in our model. Firms, to avoid costly equity issuance, hoard cash to be able to smooth their dividend over time and are able to do so during recessions. This is also in line with results presented by Opler et al. (1999) and Gao et al. (2013).

Replication with model simulated data We replicate the empirical exercise using model generated data.[20] Figure 9 plots the overall effect of the impact of a recession on dividends and the coefficient on the interaction with cash. Overall, dividends drop by approximately 3% in the first period of a recession, qualitatively and quantitatively similar to the 2.5% found using Compustat data. Additionally, an increase of cash to output of 1% reduces the dividend decrease by 0.016 percentage points. This implies that a one standard deviation increase in cash holdings reduces the dividend fall by approximately 50%, very

[19]Figure F.14 in Appendix F shows the results when considering the observations when dividends are zero, by using the log of one plus the real dividends. The results are qualitatively and quantitatively similar.

[20]Notice that in the model the firm has no assets besides cash. So we use cash as a % of output. The control variables are just output. Equally, to stay as close to the data as possible, we consider a recession whenever aggregate TFP is in its lower state.

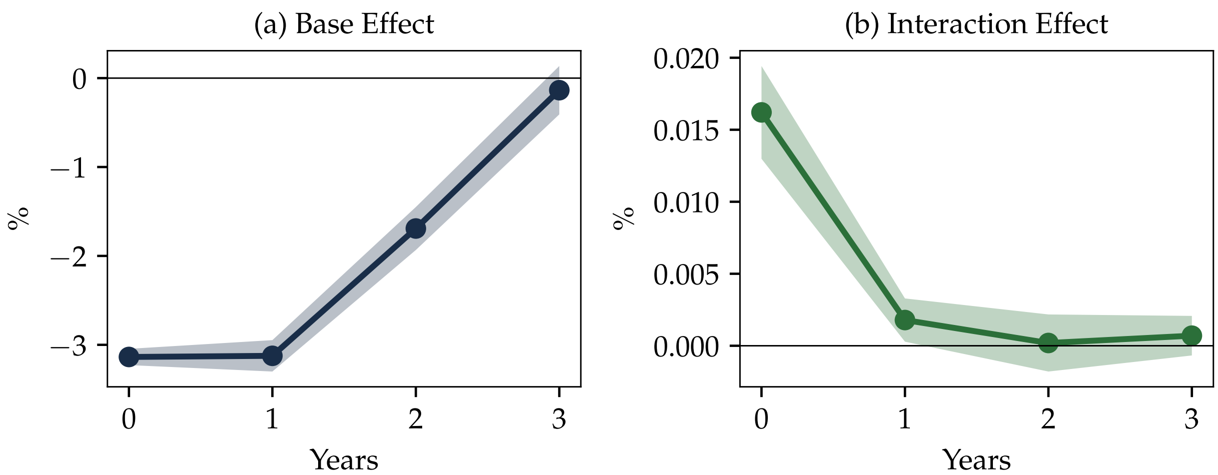

Figure 9: Dynamic dividend reaction to negative GDP and interaction with cash to assets

Note. The figure plots the impact of recessions on panel (a) up to three years after the recession. Panel (b) plots the coefficient associated with the interaction term between cash to total assets ratio and the recession indicator, using model simulated data. Shaded bands represent 90% confidence intervals

similar to the data counterpart.

6.2 Household consumption and dividend income

Our result of corporate dividend state-dependency carrying over into consumption relies on dividends being a key determinant of consumption at the household level. In order to show this connection empirically, we utilize household-level bi-annual data from the PSID between 2005 and 2021.[21] The PSID records dividends received at the household level from 2005 onwards along with total expenditure.[22] Using the PSID, we then run the following regression:

$$c_{it} = \beta d_{it-1} + \Gamma X_{it-1} + \mu_i + \mu_t + \epsilon_{it}, \tag{22}$$

[21] Appendix A.3 describes the cleaning steps involved in the preparation of the PSID panel.

[22] Total expenditures include food (at home, delivered, or eaten out), housing (mortgage or rent, utilities, internet & phone, insurance, property taxes, repairs, furnishings), health (hospital, doctor bills, prescriptions, health insurance), transportation (vehicle loans, leases, down payments, insurance, repairs, gasoline, parking, public transport, taxis), children & education (childcare, education), and personal & recreation (clothing, trips, other recreation expenses).

Table 5: Effects of real dividend income on real expenditure

	Real Expenditure	Real Expenditure	Real Expenditure	Real Expenditure
Real dividend income	0.157***	0.0886**	0.0814*	0.0785*
	(0.0338)	(0.0416)	(0.0428)	(0.0412)
HH covariates		✓	✓	✓
HH FE			✓	✓
Year FE				✓
Observations	46 419	23 138	21 223	21 223

Note. This table shows the effects of a one dollar increase of dividends on real expenditure. Household covariates include financial and business wealth, housing wealth, financial asset income, labor income, age and education. Standard errors clustered at the household level.
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Equation (22) estimates the relationship between household consumption c_{it} and dividend income d_{it-1}, where i indexes households and t indexes time. The coefficient β captures the marginal effect of dividends on consumption. The vector X_{it-1} includes a set of lagged household-level control variables such as labor income, financial income, non-housing wealth, housing wealth, age, and education in order to account for other determinants of consumption. We also include household fixed effects μ_i to control for time-invariant household characteristics, and time fixed effects μ_t to absorb aggregate shocks common to all households. This specification tests whether within-household changes in dividend income are systematically associated with changes in consumption, after accounting for a rich set of controls and fixed effects.

Results are presented in Table 5 and show that an increase in dividend income in year $t - 1$ is positively correlated with expenditure in year t across a number of different specifications. In our strictest specification, with household covariates, fixed effects and time fixed effects, a 1 dollar increase in dividend income is associated with a 7.85 cent increase in real expenditures. This result is in line with papers such as Baker et al. (2006), Di Maggio et al. (2020) and Bräuer et al. (2022), who show for example that consumption responds positively to dividends in the CEX or transaction level data from a German bank.

7 Concluding remarks

Corporate cash holdings have been increasing globally over the past few decades. In this paper, we analyze the effects of this increase on business cycle dynamics. Making use of a heterogeneous-firm business cycle model with financial frictions, in which firms hoard cash out of a precautionary motivation.

We show analytically that the firm's precautionary cash holdings generate inefficiently high levels of cash in the economy. In a stylized two-period model, we prove that the privately optimal cash policy diverges from the socially optimal one: firms hoard excessive cash due to external financing costs while the social planner would prefer zero cash holdings to avoid the pecuniary externality of holding cash.

Our quantitative model shows that precautionary corporate cash holdings significantly amplify and shape the nonlinearity of aggregate fluctuations, especially in consumption. In particular, the dividend smoothing motive induces a state-dependent response: following negative TFP shocks, if cash levels are low, firms cut dividends sharply to preserve internal funds, leading to a stronger drop in aggregate consumption. We find that a one standard deviation increase in cash holdings reduces consumption decrease in response to a negative TFP shock by 5%.

Using micro data for the firms and households, we find empirical evidence in support of model mechanism. First, we find firms with higher cash holdings to have smoother dividends. Second, we find that an increase in dividend income is positively associated with increases in consumption.

References

ACHARYA, V. V., H. ALMEIDA, AND M. CAMPELLO (2007): "Is cash negative debt? A hedging perspective on corporate financial policies," *Journal of financial intermediation*, 16, 515–554.

AIYAGARI, S. R. (1993): "Uninsured idiosyncratic risk and aggregate saving," Working Papers 502, Federal Reserve Bank of Minneapolis.

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

ALFARO, I., N. BLOOM, AND X. LIN (2024): "The Finance Uncertainty Multiplier," *Journal of Political Economy*, 132, 577–615.

ALTINKILIC, O. AND R. S. HANSEN (2000): "Are There Economies of Scale in Underwriting Fees? Evidence of Rising External Financing Costs," *The Review of Financial Studies*, 13, 191–218.

BACHMANN, R., R. J. CABALLERO, AND E. M. R. A. ENGEL (2013): "Aggregate Implications of Lumpy Investment: New Evidence and a DSGE Model," *American Economic Journal: Macroeconomics*, 5, 29–67.

BAKER, M., S. NAGEL, AND J. WURGLER (2006): "The effect of dividends on consumption," .

BATES, T. W., K. M. KAHLE, AND R. M. STULZ (2009): "Why do US firms hold so much more cash than they used to?" *The journal of finance*, 64, 1985–2021.

BEGENAU, J. AND J. SALOMAO (2018): "Firm Financing over the Business Cycle," *The Review of Financial Studies*, 32, 1235–1274.

BERNANKE, B. S., M. GERTLER, AND S. GILCHRIST (1999): "The financial accelerator in a quantitative business cycle framework," *Handbook of macroeconomics*, 1, 1341–1393.

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.

BRÄUER, K., A. HACKETHAL, AND T. HANSPAL (2022): "Consuming dividends," *The Review of Financial Studies*, 35, 4802–4857.

CARROLL, C. D. (1997): "Buffer-Stock Saving and the Life Cycle/Permanent Income Hypothesis," *The Quarterly Journal of Economics*, 112, 56.

CHEN, A. Y. (2017): "External Habit in a Production Economy: A Model of Asset Prices and Consumption Volatility Risk," *The Review of Financial Studies*, 30, 2890–2932.

CHEN, P., L. KARABARBOUNIS, AND B. NEIMAN (2017): "The global rise of corporate saving," *Journal of Monetary Economics*, 89, 1–19.

CHETTY, R., A. GUREN, D. MANOLI, AND A. WEBER (2011): "Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins," *American Economic Review*, 101, 471–75.

CLOYNE, J., C. FERREIRA, M. FROEMEL, AND P. SURICO (2023): "Monetary Policy, Corporate Finance, and Investment," *Journal of the European Economic Association*, jvad009.

COOLEY, T. F. AND V. QUADRINI (2001): "Financial Markets and Firm Dynamics," *American Economic Review*, 91, 1286–1310.

DI MAGGIO, M., A. KERMANI, AND K. MAJLESI (2020): "Stock market returns and consumption," *The Journal of Finance*, 75, 3175–3219.

FERREIRA, M. H., T. HABER, AND C. RÖRIG (2023): "Financial constraints and firm size: Micro-evidence and aggregate implications," Tech. rep., Working Paper.

FROOT, K. A., D. S. SCHARFSTEIN, AND J. C. STEIN (1993): "Risk Management: Coordinating Corporate Investment and Financing Policies," *Journal of Finance*, 48, 1629–1658.

GAO, H., J. HARFORD, AND K. LI (2013): "Determinants of corporate cash policy: Insights from private firms," *Journal of Financial Economics*, 109, 623–639.

GERTLER, M. AND N. KIYOTAKI (2010): "Financial intermediation and credit policy in business cycle analysis," in *Handbook of monetary economics*, Elsevier, vol. 3, 547–599.

HANSEN, R. S. AND P. TORREGROSA (1992): "Underwriter Compensation and Corporate Monitoring," *Journal of Finance*, 47, 1537–1555.

HENNESSY, C. A. AND T. M. WHITED (2007): "How costly is external financing? Evidence from a structural estimation," *The Journal of Finance*, 62, 1705–1745.

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

JERMANN, U. AND V. QUADRINI (2012): "Macroeconomic Effects of Financial Shocks," *American Economic Review*, 102, 238–271.

KAPLAN, G., B. MOLL, AND G. L. VIOLANTE (2018): "Monetary Policy According to HANK," *American Economic Review*, 108, 697–743.

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

KRUSELL, P. AND A. A. SMITH, JR. (1997): "Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns," *Macroeconomic Dynamics*, 1.

——— (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 106, 867–896.

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.

——— (2025): "Global Nonlinear Solutions in Sequence Space and the Generalized Transition Function," *Working paper*.

OPLER, T., L. PINKOWITZ, R. STULZ, AND R. WILLIAMSON (1999): "The determinants and implications of corporate cash holdings," *Journal of financial economics*, 52, 3–46.

OTTONELLO, P. AND T. WINBERRY (2020): "Financial Heterogeneity and the Investment Channel of Monetary Policy," *Econometrica*, 88, 2473–2502.

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

Appendix

A Data Construction and Cleaning

A.1 Definition of liquid assets from the Flow of Funds

The nonfinancial corporate business holdings of liquid assets are defined both as a narrow measure and a broad measure. Liquid assets is defined as the sum of, private foreign deposits (`FL103091003.Q`), checkable deposits and currency (`FL103020005.Q`), total time and savings deposits (`FL103030003.Q`) and money market fund shares (`FL103034000.Q`).

A.2 Compustat Data Construction and Cleaning

We construct a panel of U.S. non-financial, non-utility firms using the Compustat North America annual fundamentals file. The cleaning procedure follows standard practices. The main steps are as follows:

1. **Sample restrictions**: We retain only firm-year observations that satisfy the following:
   - Report in U.S. dollars (`curcd` = USD), use standard data format (`datafmt` = STD), industrial format (`indfmt` = INDL), and consolidated accounting (`consol` = C).
   - Incorporated in the United States (`fic` = USA).
   - Have non-missing accounting dates (`datadate`).
2. **Firm-year identification**: We ensure each firm-year is uniquely identified. For firms with multiple observations in a fiscal year, we keep the latest observation based on reporting month.

3. **Industry exclusions**: We exclude firms in the financial and utility sectors, defined as those with NAICS codes beginning in 22 (utilities), 52 (finance), or 53 (real estate). Multinational firms (NAICS = 999) are also excluded. Observations with missing or overly coarse (e.g., 2-digit) NAICS codes are dropped.

4. **Data validity and trimming**: We apply the following filters to address data quality and outliers:

   - Observations must have strictly positive values for sales and total assets.
   - The cash-to-assets ratio must lie within $[0, 1]$.
   - Observations with negative debt (short or long term) are excluded.
   - Leverage (defined as total debt over assets) is trimmed at the 99th percentile.
   - Dividends-to-assets ratios exceeding 10 are excluded.
   - Sales values are winsorized by dropping the top and bottom 1% of the distribution.

A.3 PSID Data Construction and Cleaning

Our empirical analysis draws on the Panel Study of Income Dynamics (PSID) Family and Individual Files from 2005 to 2021. We construct a panel dataset of U.S. households with a rich set of variables on income, wealth, consumption, and demographics. This appendix summarizes the key steps involved in preparing the dataset used for the regression analysis in Section 6.

Data Sources and Structure

We use the PSID Family Files to construct household-level variables on consumption, wealth, and various forms of income. These files are extracted for each survey year, standardized, and appended. We merge the Family Files with the PSID Individual File to recover detailed demographic characteristics and household composition, particularly the number of dependents and the identity of household heads and spouses.

Cleaning and Harmonization Procedures

- **Expenditure Variables:** We construct total household consumption by aggregating over categories such as food at home and away, rent, utilities, health care, transportation, education, and durable goods. Where expenditures are reported on different time scales (weekly, monthly, annually), we convert all values to annual figures.

- **Income and Dividend Variables:** We compute household-level dividend income as the sum of head and spouse self-reported income from dividends. Careful attention is given to avoid double-counting in cases where joint asset ownership is reported. We similarly construct interest, rent, and trust income, along with earned labor income and transfer payments.

- **Wealth Variables:** Wealth is calculated as the sum of financial (cash, stocks, bonds), real (real estate, vehicles), and business assets, net of outstanding debts such as mortgages and other liabilities.

- **Deflation:** All nominal variables are deflated using the December values of the Consumer Price Index (CPI-U) to reflect constant 2005 dollars. Forward-looking deflation is applied to match the retrospective nature of many income variables in the PSID.

- **Sample Restrictions:** We restrict the sample to households with a valid head, nonmissing expenditure and dividend data, and at least one year of observation between 2005 and 2021. Implausible or top-coded values (e.g., dividends ¿ $10 million) are set to missing, and observations with zero or negative consumption or assets are excluded.

Key Variables for Analysis

The core regression examines the relationship between *real dividend income* and *real household expenditure*. The key variables used include:

- `divtotal_real`: Sum of head and spouse dividend income, deflated.
- `totexpenditure_real`: Total household expenditure, deflated using future CPI.
- `wealth_real`, `hlabinc_real`, `wlabinc_real`: Real wealth and labor income of household members.
- `id`, `year`: Household and year identifiers used for fixed effects.

B Two period model extensions

B.1 Linear equity issuance cost

The firm's problem is to maximize its value, considering current and future dividends net of equity issuance costs. The cost of equity issuance $C(d)$ is given by:

$$C(d) = \begin{cases} \mu_1 d + \frac{1}{2}\mu_2 d^2 & \text{if } d < 0 \\ 0 & \text{if } d \geq 0 \end{cases}$$

where d represents dividends. This cost is incurred only when the firm issues equity (i.e., when dividends are negative).

The firm chooses its cash holding n_1 in period 1. The dividends in period 1 (d_1) and in period 2 ($\tilde{d}_2$) are functions of n_1 and productivity shocks.

The firm's objective is to maximize its value, which depends on its current cash holdings n_1. The firm's value function $V(n_1)$ is given by:

$$V(n_1) = (A_1 - w_1 - n_1\beta^n + n_0) - C(d_1) + \mathbb{E}[m(\tilde{A}_2)(\tilde{d}_2 - C(\tilde{d}_2))]$$

Case 1: n_1^* is such that $d_2^L \geq 0$ (i.e., $n_1^* \geq n_1^{max}$)

Let $n_1^{max} = w_2 - A_1 + \Delta$. In this scenario, $C(d_2^L) = 0$ (since $d_2^L \geq 0$), and thus $C'(d_2^L) = 0$. The first-order condition becomes:

$$-\beta^n + \frac{1}{2}m(A_1 + \Delta) + \frac{1}{2}m(A_1 - \Delta) = -\beta^n + \mathbb{E}[m(\tilde{A}_2)] = 0$$

However, we assume $\beta^n > \mathbb{E}[m(\tilde{A}_2)]$, meaning $-\beta^n + \mathbb{E}[m(\tilde{A}_2)] < 0$. This implies that the firm has a satiation point at $n_1 = n_1^{max}$.

Case 2: $d_i < 0$

The first-order condition is:

$$-\beta^n + \mu_1\beta^n + \mu_2\beta^n d_1 + \mathbb{E}[m(\tilde{A}_2)](1 - \mu_1 - \mu_2 d_2) = 0$$

Solving for n_1^*:

$$n_1 = \frac{\mu_2\left[\beta^n(A_1 - w_1 + n_0) + \mathbb{E}\left(m(\widetilde{A}_2)(w_2 - \widetilde{A}_2)\right)\right] + \left(\beta^n - \mathbb{E}(m(\widetilde{A}_2))\right)(\mu_1 - 1)}{\left(\mathbb{E}(m(\widetilde{A}_2)) + (\beta^n)^2\right)\mu_2}$$

Proposition 3 (Precautionary savings and linear equity issuance cost.)**.**
As $\beta^n > \mathbb{E}(m(\widetilde{A}_2))$ an increase in μ_1 leads to an increase in cash holdings.

B.2 Optimal savings

Social planner The social planner maximizes the present discounted value of the household's utility:

$$\max_{c_1,c_2} \quad \log(c_1) + \beta\mathbb{E}\log(\widetilde{c}_2), \tag{23}$$

$$\text{s.t.} \quad c_1 = A_1 - n_1\beta^n + n_0, \tag{24}$$

$$\widetilde{c}_2 = \widetilde{A}_2 + n_1, \tag{25}$$

$$n_1 \geq 0 \tag{26}$$

Consumption in period 1 equals production net of changes in cash holdings, and consumption in period 2 includes stochastic output and cash. Since the equity issuance cost is rebated, it does not appear in the planner's constraints.

Proposition 4 (Inefficiency of the decentralized cash holdings)**.**

Socially optimal cash holdings are equal to zero.

Proof. The social planner fist order condition implies

$$\beta^n = \beta \mathbb{E} \frac{A_1 - \beta^n n_1 + n_0}{\widetilde{A}_2 + n_1} + \lambda (A_1 - \beta^n n_1 + n_0). \tag{27}$$

Notice the term $\beta \mathbb{E} \frac{A_1 - \beta^n n_1 + n_0}{\widetilde{A}_2 + n_1}$ is equal to the stochastic discount factor and as β^n is strictly lower than the SDF, this implies $\lambda > 0$ and $n_1 = 0$ for any level of n_0. ■

The social planner wants zero cash holdings due to the externality of cash holdings. When firms increase their cash buffers, they depress household wealth by reallocating resources toward lower-return assets, thereby leading to a decline in aggregate consumption. This externality leads to the social planner strictly preferring zero cash to avoid the externality. Despite the cost of holding cash, the firms still prefer strictly positive amounts of cash for precautionary motives as they do not internalize that the equity issuance cost is not a social cost. If firms fully internalized the externality, the decentralized and planner solutions would coincide.

This proposition equally holds for the quantitative model, as the planner's problem is similar due to the presence of the wedge between cash return and the stochastic discount factor and a representative household.

Pecuniary externality The above identified cash externality could also be seen as a pecuniary externality of holding cash, if equity prices were considered. As dividends become on average lower with higher cash holdings, this depresses the equity prices, which leads to a lower level of average consumption. The social planner maximizes the presented discounted value of the household's utility:

$$\max_{c_1,c_2} \quad \log(c_1) + \beta\mathbb{E}\log(\widetilde{c}_2), \tag{28}$$
$$\text{s.t.} \quad c_1 + q_2 = q_1 + w_1,$$
$$\widetilde{c}_2 = q_2 + w_2,$$
$$q_1 = d_1 + \mathbb{E}(m(\widetilde{A_2})d_2),$$
$$q_2 = \mathbb{E}(m(\widetilde{A_2})d_2).$$

Consumption in period 1 equals labor income w_1 net of changes in equity holdings $q_1 - q_2$, and consumption in period 2 includes labor income w_2 and the equity prices q_2 in period 2. Since the equity issuance cost is rebated, it does not appear in the planner's constraints and does not affect equity prices.

Proposition 5 (Inefficiency of the decentralized cash holdings)**.**
Socially optimal cash holdings are equal to zero.

Proof. The social planner fist order condition implies

$$\beta^n = \beta\mathbb{E}\frac{A_1 - \beta^n n_1 + n_0}{\widetilde{A}_2 + n_1} + \lambda(A_1 - \beta^n n_1 + n_0). \tag{29}$$

Notice the term $\beta\mathbb{E}\frac{A_1 - \beta^n n_1 + n_0}{\widetilde{A}_2 + n_1}$ is equal to the stochastic discount factor and as β^n is strictly lower than the SDF, this implies $\lambda > 0$ and $n_1 = 0$ for any level of n_0. ■

The result is still the same as before, but the externality now is a pecuniary externality as it operates via equity prices.

C Note on computation

Then, the equilibrium price p_t is determined from the following variant of the non-trivial market clearing condition:[23]

$$p = \arg_{\widetilde{p}} \left\{ c(\widetilde{p}) - \int \left[d(x; X, \widetilde{p}) + w(X, \widetilde{p}) l(x; X, \widetilde{p}) \right] d\Phi = 0 \right\}.$$

This is a fixed-point problem and computationally costly to solve. Instead, the repeated transition method uses the implied price p^*, which is obtained as follows:

$$\begin{aligned} p^* &= \arg_{\widetilde{p}} \left\{ c(\widetilde{p}) - \int \left[d(x; X, p^{(n)}) + w(X, p^{(n)}) l(x; X, p^{(n)}) \right] d\Phi = 0 \right\} \\ &= 1 / \int \left[d(x; X, p^{(n)}) + w(X, p^{(n)}) l(x; X, p^{(n)}) \right] d\Phi, \end{aligned}$$

where $p^{(n)}$ is the guessed price in the n^{th} iteration. We take the sufficient statistic approach, and the aggregate cash holdings N_t (the first moment of the distribution of cash holding) is the sufficient statistic.

[23]This condition is derived from combining the household's budget constraint and the equity market clearing condition. For a sharp illustration, the time subscript is lifted.

D Recovering the true nonlinear law of motion

In this section, we recover the true law of motion from the converged equilibrium outcomes over the simulated path. Then, we test the validity of the true law of motion by fitting the law of motion into the out-of-sample simulated path.

Specifically, the following laws of motion are studied:

$$N_{t+1} = G_N(N_t, N_{t-1}, N_{t-2}, \ldots, N_{t-n}; A_t)$$
$$p_t = G_p(N_t, N_{t-1}, N_{t-2}, \ldots, N_{t-n}; A_t).$$

Table D.6 reports the goodness of fitness (R^2) of the different specifications. The first five rows report the fitness when only the contemporaneous aggregate cash stock N_t is considered up to different polynomial orders. When a single argument is considered without the higher-order polynomials, R^2 gets as low as 0.8956 for G_N if the contemporaneous productivity state is G.[24] As the more higher-order polynomials are included, the better the fitness becomes. However, the fitness of the specification of G_p stops improving after a certain threshold. This shows that the true law of motion can be recovered only by including further historical allocations.

The bottom seven rows of Table D.6 report the fitness of the law of motion when the additional lagged terms of the aggregate cash stock are considered. Up to the third order polynomials are included for each of lagged terms on top of the polynomial terms of the contemporaneous cash stocks up to the fifth order.[25] As more lagged terms are considered in the law of motion, the fitness improves, especially in G_p. However, only after the polynomials of the seven-period lagged aggregate cash stock are included in the law of motion, the accurate law of motion is recovered.

This exercise shows the substantial nonlinearity of the aggregate fluctuations in this

[24]This pure linear specification is different from the log-linear specification in Section 5.

[25]The results only negligibly change over different order choices.

Table D.6: The fitness of the law of motion across different specifications

	# of lagged	order	Goodness of fitness: R^2			
			CA_{t+1} : *Good*	CA_{t+1} : *Bad*	p_t : *Good*	p_t : *Bad*
Contemp.	0	1	0.8956	0.9452	0.9922	0.9966
	0	2	0.9839	0.9952	0.9927	0.9976
	0	3	0.9973	0.9995	0.9930	0.9976
	0	4	0.9993	0.9999	0.9932	0.9976
	0	5	0.9996	1.0000	0.9933	0.9976
Add. history	1	3	0.9999	1.0000	0.9987	0.9979
	2	3	0.9999	1.0000	0.9997	0.9984
	3	3	0.9999	1.0000	0.9998	0.9987
	4	3	0.9999	1.0000	0.9998	0.9991
	5	3	0.9999	1.0000	0.9998	0.9994
	6	3	0.9999	1.0000	0.9998	0.9996
	7	3	0.9999	1.0000	0.9998	0.9997

Notes: The first column (# of lagged) reports the number of lagged terms included in the specification. The second column (order) reports the highest order of polynomials considered. When lagged terms are included, contemporaneous terms are considered up to the fifth polynomial for all specifications. The third and fourth column reports the R^2 of the law of motion of N_{t+1} for Good and Bad shock realizations. The fifth and sixth column reports the R^2 of the law of motion of p_t for Good and Bad shock realizations.

model. In the repeated transition method, the contemporaneous aggregate cash stock is used as a sufficient statistic of each period's cross-section. However, this does *not* imply that the true aggregate law of motion is a function of only the contemporaneous cash stock. The contemporaneous cash stock is rather a labeling of each period that correctly sorts the rankings of the value functions across the periods in the repeated transition method.

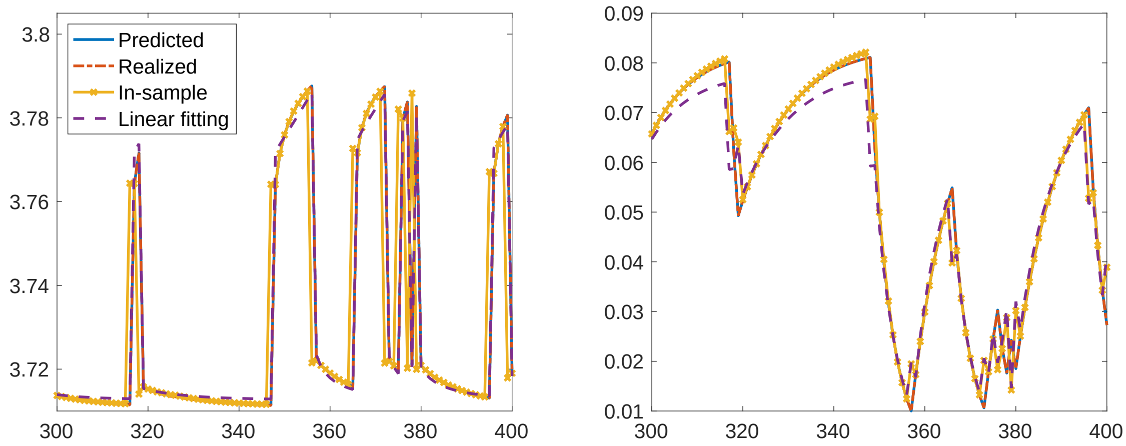

(a) Price p_t $(= 1/C_t)$

(b) Aggregate cash stocks N_t

Figure D.10: Fitting into the out-of-sample path

Notes: The figure plots the time series of the price p_t the aggregate cash stock N_t in the baseline model for an out-of-sample simulated path of the aggregate shocks. In both panels, the solid line is the predicted time series (n^{th} guess) $\{p_t^{(n)}, N_t^{(n)}\}_{t=300}^{400}$; the dash-dotted line is the realized time series $\{p_t^*, N_t^*\}_{t=300}^{400}$; the line with the star tick mark is the predicted time series implied by the exact law of motion recovered from the in-sample path; the dashed line is the predicted time series implied by the linear law of motion.

We validate the recovered law of motion by fitting it on the out-of-sample simulated path.[26] Specifically, we solve the model on another simulated path to obtain the converged equilibrium dynamics using the repeated transition method and compare the dynamics with the implied dynamics in the recovered true law of motion on the in-sample path. Figure D.10 plots p_t (panel (a)) and N_t (panel (b)) for 1) predicted time series (solid line), 2) realized time series (dot-dashed line), 3) time series implied by the recovered in-sample law of motion (solid line with ticks), and 4) time series implied by the linear law of motion (dashed line). The predicted time series and the realized time series are indistinguishably close to each other due to the convergence requirement of the repeated transition method. The time series implied by the recovered law of motion also closely tracks the converged equilibrium dynamics, validating the specification. The goodness

[26]The recovered true of motion refers to the specification that considers up to the seven-period lagged aggregate cash stocks, where R^2 is the highest in Table D.6.

of fitness (R^2) in the time series implied by the recovered law of motion is greater than 0.999 for both G_N and G_p in all shock realizations.

D.1 Proposition and proof of satiation point

Proposition 6 (The existence of the target cash-holding level)**.**
Suppose policy functions are non-trivial: $n'(n_1, z) > 0$ for some $n_1 > 0$ and $d(n_2, z) > 0$ for some $n_2 > 0$, given z. Then, there exists $\overline{n}(z) > 0$ such that $n'(n, z) \leq \overline{n}(z)$ for $\forall n \geq 0$.

Proof. To prove the proposition by contradiction, suppose there is no such $\overline{n}(z)$. That is, $n'(n, z) < n'(n + \epsilon, z)$ for $\forall(n, z)$ and $\forall\epsilon > 0$.

Define cash on hand as $m(n, z) = \pi(z) + n$. Then,

$$d(n, z) + q^n n'(n, z) = m(n, z).$$

$m(n, z)$ strictly increases in n. Due to the monotone preference on greater d and n' and strict monotonicity of m on n, d and n' weakly increase in n. Now consider $\widetilde{n}$ such that $n'(\widetilde{n}, z) > 0$ and $d(\widetilde{n}, z) > 0$. Such $\widetilde{n}$ exists as n' and d weakly increase in n. For example $\widetilde{n} = \max\{n_1, n_2\}$.

Then, for a marginal incremental ϵ in cash, the marginal cost of hoarding cash is 1 (forgone dividend), while the marginal benefit out of hoarding cash is $\frac{q}{q^n}$:[27]

$$\underbrace{1}_{\text{Marginal cost}} > \underbrace{\frac{q}{q^n}}_{\text{Marginal benefit}}.$$

where, $q = \beta$ at the non-stochastic steady state. This implies that for the extra cash, the firm does not have an incentive to hoard it in the cash reserve. Therefore, $d(n + \epsilon, z) =$

[27] In this argument, the non-negativity constraint does not matter, as $n'(\widetilde{n}, z) > 0$.

$d(n,z)+\epsilon$, if $d(n,z)>0$. Then, from a firm's budget constraint,

$$\begin{aligned} q^n n'(\widetilde{n}+\epsilon,z) &= \widetilde{n}\epsilon + \pi(z) - d(\widetilde{n}+\epsilon,z) \\ &= \widetilde{n} + \pi(z) - (d(\widetilde{n}+\epsilon,z)-\epsilon) \\ &= \widetilde{n} + \pi(z) - d(\widetilde{n},z) \\ &= q^n n'(\widetilde{n},z). \end{aligned}$$

Therefore, any extra increase in the current cash stock $\widetilde{n}$ does not change the future cash stock:

$$q^n n'(\widetilde{n},z) = q^n n'(\widetilde{n}+\epsilon,z),$$

which is a contradiction. Therefore, there exists the target cash stock $\overline{n}(z)$. ■

E Social planner's problem

The social planner's problem can be formulated in the following recursive form:

$$V^H(n;A) = \max_{c,n',L^H} \left[\log(c) - \frac{\eta}{1+\frac{1}{\chi}} (L^H)^{1+\frac{1}{\chi}} + \beta \mathbb{E}\left[V^H(n';A') \right] \right] \tag{30}$$

subject to:

$$c + q^n(X)n' = f(L^H) + n$$

$$A' \sim \Gamma(A)$$

The social planner's state variables are the aggregate cash holding n and the exogenous productivity A. The cash return is determined by $q^n(X)$ as in the baseline model.

F Additional figures

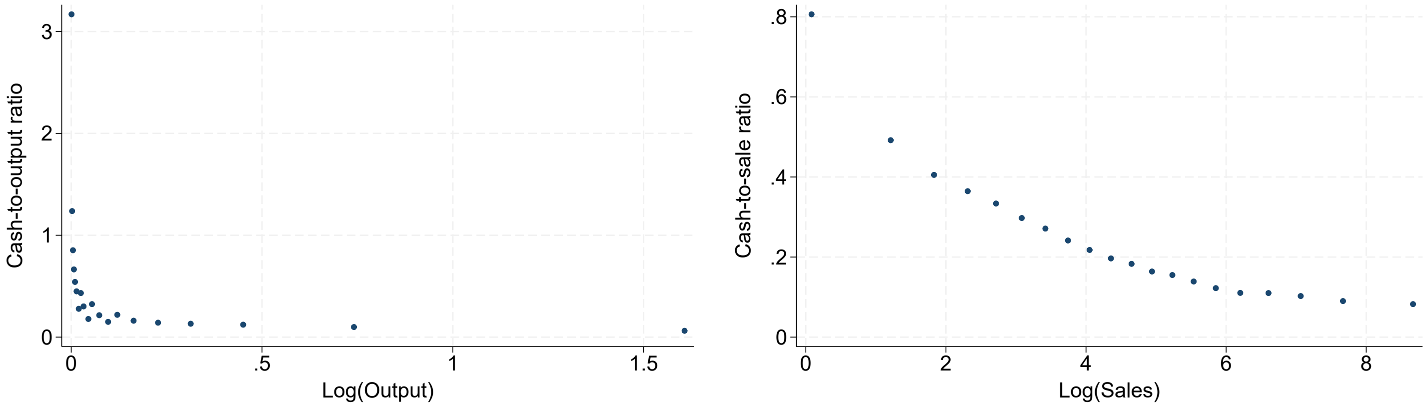

(a) Model (b) Data

Figure F.11: Correlation between cash-to-output ratio and output in the model (panel a) and in the data (panel b).

Notes: On the left panel, the figure plots the the correlation between cash-to-output ratio, on the y-axis, and log of output, on the x-axis, in the baseline model. On the right panel, the figure plots the empirical correlation between cash-to-sales ratio, on the y-axis, and log of sales, on the x-axis.

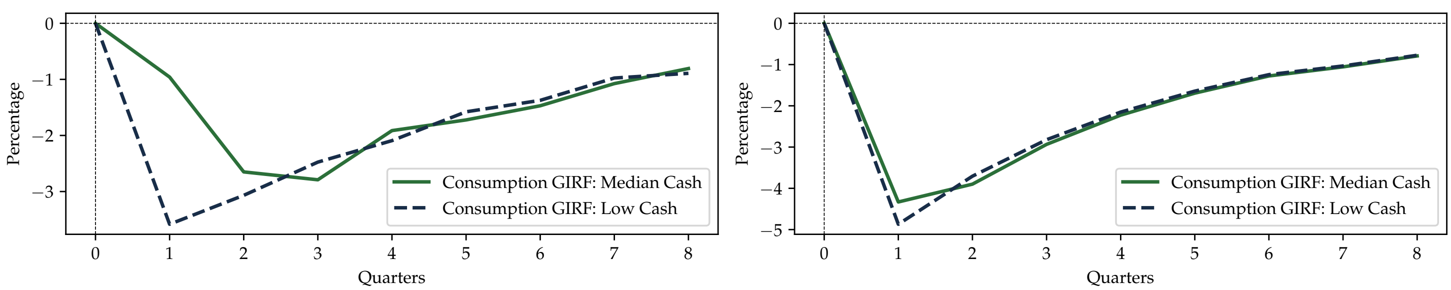

(a) Heterogeneous firms (Baseline) (b) Representative firm

Figure F.12: GIRF: Consumption responses to negative shocks for median and low cash in the baseline model (panel a) and the representative firm model (panel b).

Notes: The figure plots the consumption % change to a negative aggregate TFP shocks when cash is at median level (green solid line) vs lowest level (blue dashed line) over the entire simulation.

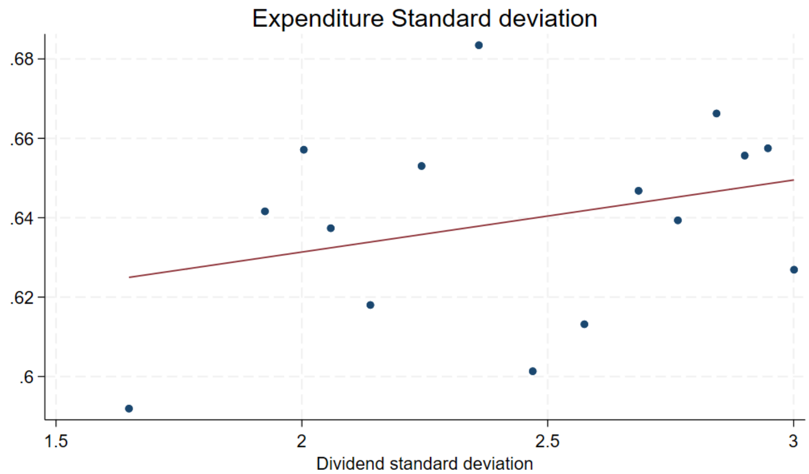

Figure F.13: Standard deviations of percentage deviations from mean for dividend and expenditure.

Notes: The figure plots the within household standard deviation of percentage deviations from mean expenditure as a function of the standard deviation of percentage deviations from mean dividend income. The regression includes, as controls, the standard deviation of labor income percentage deviations from the mean, average cash and standard deviation of percentage deviations from mean cash.

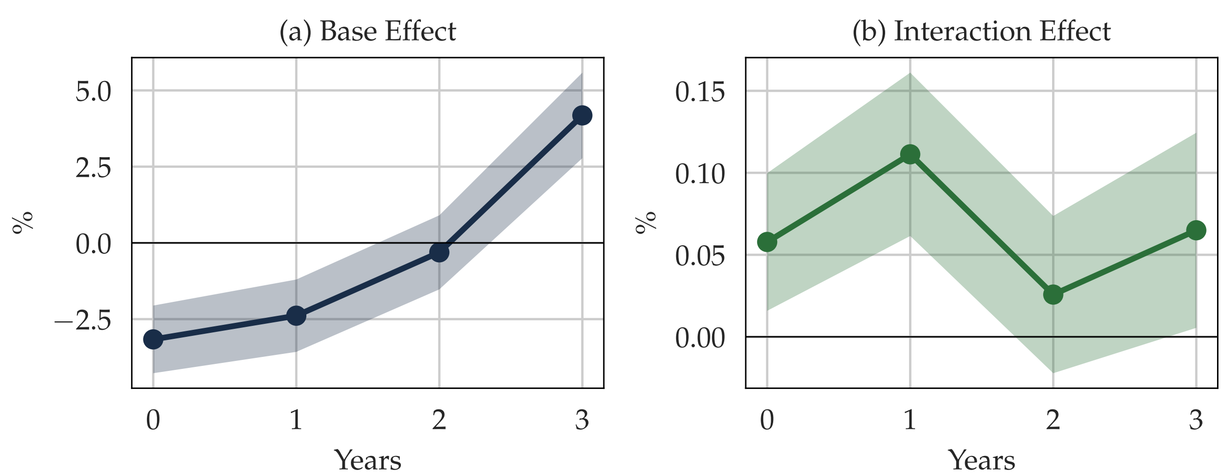

Figure F.14: Dynamic dividend reaction to negative GDP and interaction with cash to assets

Note. The figure plots the impact of recessions on panel (a) up to three years after the recession. Panel (b) plots the coefficient associated with the interaction term between cash to total assets ratio and the recession indicator.Shaded bands represent 90% confidence intervals. Log of one plus real dividends considered.

G Proofs two-period model

Proof.

$$\frac{\partial}{\partial n_0}\left|\frac{\partial log(c_2(A_1-\Delta))}{\partial \Delta}\right| = \frac{\partial}{\partial n_0}\left|\frac{-1}{c_2(A_1-\Delta)}\right| \tag{31}$$

$$= -\left(\frac{1}{c_2(A_1-\Delta)}\right)^2 n_1^{*'}(n_0) \leq 0. \tag{32}$$

■