Collateral constraints, capital specificity and the distribution of production: the role of real and financial frictions in aggregate fluctuations

Aubhik Khan
The Ohio State University

Julia K. Thomas
The Ohio State University and NBER

September 2009

PRELIMINARY AND INCOMPLETE — Please do not quote without author approval

ABSTRACT

We study the cyclical implications of credit market imperfections in a dynamic, stochastic general equilibrium model wherein firms face persistent shocks to both aggregate and individual productivity. In our model economy, optimal capital reallocation is distorted by two frictions. First, collateralized borrowing constraints limit the investment undertaken by small firms with relatively high productivities. Second, specificity in firm-level capital implies partial investment irreversibilities that lead firms to pursue generalized (S,s) investment rules. This second friction compounds the first in implying that large and relatively unproductive firms carry a disproportionate share of the aggregate capital stock, thereby reducing endogenous aggregate total factor productivity. Moreover, because irreversibilities not only directly induce both downward and upward inertia in firm-level capital adjustment, but also tighten the borrowing limits associated with collateral constraints, they ensure that the negative consequences of a temporary tightening in financial markets are not quickly repaired. In the presence of persistent heterogeneity in both capital and total factor productivity, the effects of a financial shock can be amplified and propagated through large and long-lived disruptions to the distribution of capital that, in turn, imply large and persistent reductions in aggregate total factor productivity. Similarly, the consequences of a negative real shock can be exacerbated and prolonged in the presence of real and financial frictions. This paper seeks to measure the strength of these effects in a calibrated DSGE setting.

KEYWORDS: Financial frictions, capital adjustment frictions, irreversibilities, (S,s) policies, business cycles, total factor productivity, dynamic stochastic general equilibrium.

1 Introduction

Can a large shock to an economy's financial sector produce a large and lasting recession? Can it amplify and propagate the effects of a real shock sufficiently to transform recession into depression? Over the past two years, negative events in the real and financial sectors of the U.S. and other large, developed economies have become increasingly difficult to disentangle. If the current conditions have reawakened interest in business cycle research, they have also made clear how limited are our existing macroeconomic models in their ability to address such topics.

In this paper, we develop a quantitative, dynamic, stochastic general equilibrium model that may better inform current and future discussions regarding the interactions of real and financial shocks in determining the size and frequency of aggregate fluctuations. In our model, firms experience persistent shocks to both aggregate and individual productivity, while credit market frictions interact with real frictions to yield persistent disruptions to the efficient allocation of capital across them, and thus persistent reductions in endogenous aggregate productivity. Calibrating our model to aggregate and firm-level data, we use it as a laboratory in which to obtain quantitatively disciplined answers to the questions raised above.

Capital reallocation is distorted by two frictions in our model, one financial and one real. First, collateralized borrowing constraints limit the investment undertaken by small firms with relatively high productivities. Second, specificity in capital implies partial investment irreversibilities that lead firms to pursue generalized (S,s) rules with respect to their capital adjustments. The second friction compounds the first, further tilting the distribution of production towards larger, less productive firms, and thus reducing endogenous aggregate total factor productivity, and it also exacerbates the direct effects of collateral constraints by reducing the collateral value ascribed to each unit of installed capital. This added element of realism in our setting relative to existing DSGE financial frictions models may be quite important to the transmission and propagation of a financial shock, as we discuss below.

Because specificity in capital induces both downward and upward inertia in firm-level investment activities, and because it tightens the borrowing limits implied by collateralized lending, it ensures that the negative consequences of a temporary tightening in financial markets cannot be quickly reversed. In the presence of persistent heterogeneity in both capital and total factor productivity, the effects of financial frictions are amplified and propagated through large and long-lived disruptions to the distribution of capital that, in turn, imply large and persistent reductions in aggregate productivity. For example, in the presence of only a 5 percent capital irreversibility, we find that steady state output falls by nearly 9 percent when collateralized borrowing limits are introduced. This suggests the potential for large output losses in our model economy following a financial shock, or following a real shock accompanied by a financial one, since the long-run GDP reduction in response to a change in borrowing constraints fails to capture the sharp transitional reductions associated with reallocation following the shock.

As indicated above, we use our model to measure the extent to which a financial shock can spill into the real side of a calibrated economy to produce large and persistent reductions in aggregate employment and GDP on its own, as well as the extent to which it can amplify and prolong the effects of a modest-sized real shock. From the outset, understanding that investment is a small fraction of GDP, it is clear that the reductions in aggregate capital implied by a temporary reduction in available credit are unlikely to deliver sizeable or long-lived aggregate real effects. However, we also know from disaggregated data that there is substantial heterogeneity among firms in their individual productivity levels, and there are real frictions limiting the reallocation of capital across them.[1] As such, the mechanism we explore here focuses on the economy's *effective* capital stock and endogenous total factor productivity.

Our primary question in this study is whether a temporary crisis in financial markets can generate a large and persistent drop in aggregate productivity by disrupting the distribution of capital away from that implied by firms' relative productivities, and thereby distorting the distribution of production. Of course, we are not the first to emphasize reductions in measured TFP arising from a misallocation of resources across firms. Restuccia and Rogerson (2008) show that this channel can be quite important in explaining cross-country per-capita GDP differences. However, we are to our knowledge the first to explore this channel in a quantitative DSGE setting where real frictions slow the reallocation of capital across firms, and where that reallocation is essential in determining the marginal product of the aggregate stock.

The remainder of the paper is organized as follows. Section 2 briefly summarizes the literature most closely related to our work. Next, in section 3, we present our model economy. Section 4 provides some analysis useful in developing a numerical algorithm capable of its solution. In

[1] For direct evidence of large and increasing heterogeneity in firm-level productivity, see Comin and Philippon (2005) and the empirical studies cited therein. Elsewhere, Cooper and Haltiwanger (2006) find it is impossible to reproduce microeconomic investment patterns without both large idiosyncratic shocks and adjustment costs limiting capital reallocation.

section 5, we describe our calibration to moments drawn from postwar U.S. aggregate and firm-level data. Section 6 explores the mechanics of our model in its deterministic steady state and draws some comparisons to the mechanics in a reference model with capital specificity but no financial frictions. Section 7 presents dynamic results. There, we compare our economy's business cycle moments to those in the reference model, and we examine aggregate responses following a persistent shock to exogenous aggregate productivity. Next, we consider the response to an unanticipated reduction in the availability of credit and explore how this response is influenced by both the degree of capital specificity and the expected duration of the financial shock. Finally, section 8 provides brief concluding remarks.

2 Related literature

To be clear, there is a vast existing literature considering the implications of financial market imperfections. For example, Kiyotaki and Moore (1997) study a model of credit cycles and show that collateral constraints can have a role in amplifying and propagating shocks to the value of collateral. More recent studies challenge the finding, however, as one arising from an overly stylized environment. Cordoba and Ripoll (2004) argue that the effects are actually quite small in a more plausibly calibrated model. The explanation for this may be best articulated in a short article by Kocherlakota (2000). However, a common, and likely critical, element across these papers is the abstraction from any additional source of heterogeneity across firms. One notable exception is the recent paper by Buera and Shin (2007). While Buera and Shin emphasize development concerns, their primary finding that financial frictions can have a large and persistent impact on the aggregate transition to a steady state, particularly when capital is initially misallocated, is certainly an informative one for our study. It suggests that our allowance for real capital frictions alongside the financial friction they consider may be quite relevant in magnifying and propagating business cycle fluctuations.

Elsewhere in the investment literature, various empirical and theoretical studies have together mounted a strong case that real frictions limiting the reallocation of capital are essential in explaining microeconomic investment data. (See, for instance, Cooper and Haltiwanger (2006) or Caballero and Engel (1999).) Moreover, these frictions have been shown to add persistence to an economy's aggregate response to shocks (Bertola and Caballero (1994)). Thus, the fact that the financial frictions literature has largely ignored real frictions may be costly along both empirical and theoretical margins. Of course, the same could be said of the investment literature's abstraction from financial frictions, as this abstraction may be critical in the repeated finding that nonconvex capital adjustment costs, as well as investment irreversibilities, have essentially no importance for the aggregate business cycle of a DSGE model economy (e.g., Thomas (2003) and Veracierto (2002)).

To be completed.

There is one existing study that does simultaneously consider real and financial frictions in a dynamic, stochastic setting. Caggese (2007) provides a careful exploration of precisely how collateralized borrowing constraints can interact with investment irreversibility to exacerbate aggregate fluctuations. There are two critical differences in our analysis. The first is our assumption that capital investments are only partly irreversible. The second is general equilibrium.

3 Model

In our model economy, firms face both partial capital fixity and collateralized borrowing limits, which together compound the effects of persistent differences in their total factor productivities to yield substantial heterogeneity in production. We begin our description of the economy with an initial look at the optimization problem facing each firm, then follow with a brief discussion of households and equilibrium. Next, using a simple implication of equilibrium alongside some immediate observations about firms' optimal allocation of profits across dividends and retained earnings, we characterize the capital adjustment decisions of our firms as a variant of the two-sided generalized (S, s) policy that would arise in the presence of investment irreversibilities alone, absent any credit market imperfections. This analysis will show how it is possible for us to derive a convenient, computationally tractable algorithm to solve for equilibrium allocations in our model, despite its three-dimensional heterogeneity in production.

3.1 Production, credit and capital adjustment

We assume a large number of firms, each producing a homogenous output using predetermined capital stock k and labor n, via an increasing and concave production function, $y = z\varepsilon F(k, n)$. Here, z represents exogenous stochastic total factor productivity common across firms, while ε is a firm-specific counterpart. For convenience, we assume that ε is a Markov chain, $\varepsilon \in \mathbf{E} \equiv$

$\{\varepsilon_1, \ldots, \varepsilon_{N_\varepsilon}\}$, where $\Pr(\varepsilon' = \varepsilon_j \mid \varepsilon = \varepsilon_i) \equiv \pi_{ij} \geq 0$, and $\sum_{j=1}^{N_\varepsilon} \pi_{ij} = 1$ for each $i = 1, \ldots, N_\varepsilon$. Similarly, we assume that $z \in \{z_1, \ldots, z_{N_z}\}$, where $\Pr(z' = z_m \mid z = z_l) \equiv \pi^z_{lm} \geq 0$, and $\sum_{m=1}^{N_\varepsilon} \pi^z_{lm} = 1$ for each $l = 1, \ldots, N_z$.

Because our interest is in understanding how financial constraints interact with the specificity of capital in shaping the investment decisions taken by firms in our economy, we must prevent firms from growing so large that none will never again experience a binding borrowing limit. To ensure this does not occur, we impose exogenous exit and entry in the model. In particular, we assume that each firm faces a fixed probability, $\pi_d \in (0, 1)$, that it will be forced to exit the economy following production in any given period. Within a period, prior to investment, firms learn whether they will survive to produce in the next period. Exiting firms are replaced by an equal number of new firms whose initial state will be described below.

At the beginning of each period, a firm is defined by its predetermined stock of capital, $k \in \mathbf{K} \subset \mathbf{R}_+$, by the level of one-period debt it incurred in the previous period, $b \in \mathbf{B} \subset \mathbf{R}$, and by its current exogenous idiosyncratic productivity level, $\varepsilon \in \{\varepsilon_1, \ldots, \varepsilon_{N_e}\}$. Immediately thereafter, the firm learns whether it will survive to produce in the next period. Given this individual state, and having observed the economy's current aggregate state, the firm then takes a series of actions designed to maximize the expected discounted value of the current and future dividends returned to its shareholders, the households in our economy. First, it chooses its current level of employment, undertakes production, and pays its wage bill. Thereafter, it repays its existing debt and, conditional on survival, it chooses its investment, i, current dividends, and the level of debt with which it will enter into the next period, b'. For each unit of debt it incurs for the next period, a firm receives q units of output that it can use toward paying current dividends or investing in its future capital. The relative price q^{-1} reflecting the interest rate at which firms can borrow and lend is, of course, a function of the economy's aggregate state, as is the wage rate ω paid to workers. For expositional convenience, we suppress the arguments of these equilibrium price functions until we have described the model further.

In contrast to the typical setting with firm-level capital adjustment frictions, and unlike a typical environment with financial frictions, real and financial frictions are allowed to interact in our setting. Our firms' borrowing and investment decisions are necessarily inter-related, because each firm faces a collateralized borrowing constraint inside of any period. This constraint takes the form: $b' \leq \Theta k$, where Θ represents the fraction of the firm's capital stock that can be successfully uninstalled and returned to lenders next period in the event of default.

Two external forces together determine what fraction of its capital stock a firm can borrow against - the degree of specificity in capital and the enforceability of financial arrangements. Here, we simply impose both, deferring the question of their foundations for a future study. In particular, we assume that $\Theta = \theta_b \theta_k$, where $\theta_k \in [0, 1]$ is a parameter determining what fraction of a firm's capital stock survives when it is uninstalled and moved to another firm, and $\theta_b \in \mathbf{R}_+$ is the fraction of uninstalled capital that creditors are confident they will be able to repossess should default occur.[2]

If firm chooses to undertake any nonnegative level of investment, then its capital stock at the start of the next period is determined by a familiar accumulation equation,

$$k' = (1 - \delta) k + i \quad \text{for } i \geq 0,$$

where $\delta \in (0, 1)$ is the rate of capital depreciation, and primes indicate one-period-ahead values. Because there is some degree of specificity in capital, the same equation does not apply when the firm undertakes negative investment. In this case, the effective relative price of investment is θ_k rather than 1, so the accumulation equation is instead:

$$\theta_k k' = \theta_k (1 - \delta) k + i \quad \text{for } i < 0.$$

In the analysis section to follow, we will show how the asymmetry that firms face in the cost of capital adjustment naturally gives rise to two-sided (S, s) investment decision rules. For the moment, we simply point out that, in contrast to a nonconvexity in the capital adjustment technology, this type of adjustment friction implies not only investment inaction among firms within their (S, s) adjustment bands, but also some inertia among firms outside of their (S, s) bands. Because there are no increasing returns in the adjustment technology, but instead a linear penalty for negative adjustments, a firm finding itself with an intolerably high capital stock (given its current productivity), will reduce its stock only to the upper bound of its (S, s) inactivity range. Similarly, a firm with too little capital recognizes that it will incur a linear penalty should it later need to shed capital, so it invests only to the lower bound of its inactivity range.

[2] Throughout our numerical exercises in section 6, we assume that the degree of capital irreversibility, $1 - \theta_k$, is a fixed technological parameter. Except where otherwise noted, θ_b is also a fixed parameter. We will allow for an unanticipated change in θ_b in our latter exercises, where we consider the aggregate implications of a financial shock that raises or lowers the confidence of lenders.

It should be clear from the discussion above that, alongside its current productivity draw, a firm's capital adjustment may also be influenced by its ability to borrow (now and in the future), which is in turn affected by the capital (collateral) it currently holds. Note also that the firm's current investment decision may influence the level of debt that it carries into the next period. These observations imply that we must keep track of the distinguishing features of firms along three dimensions: their capital, k, their debt, b, and their idiosyncratic productivity, ε.

We summarize the distribution of firms over (k, b, ε) using the probability measure μ defined on the Borel algebra, $\mathcal{S}$, for the product space $\mathbf{S} = \mathbf{K} \times \mathbf{B} \times \mathbf{E}$. The aggregate state of the economy is then described by (z, μ), and the distribution of firms evolves over time according to a mapping, Γ, from the current aggregate state; $\mu' = \Gamma(z, \mu)$. The evolution of the firm distribution is determined in part by the actions of continuing firms and in part by entry and exit. Following production in each period, fraction π_d of existing firms exit the economy. These firms invest negatively to shed their remaining capital, and are replaced by the same number of new firms. Each new firm has zero debt and productivity $\varepsilon_0 \in \mathbf{E}$ drawn from an initial distribution $H(\varepsilon_0)$, and each enters with an initial capital stock $k_0 \in \mathbf{K}$.[3]

We are now in a position to set out the optimization problem solved by each firm in our economy. Let $v_0(k, b, \varepsilon_i; z_l, \mu)$ represent the expected discounted value of a firm that enters the period with (k, b) and firm-specific productivity ε_i, when the aggregate state of the economy is (z_l, μ), just before it learns whether it will survive into the next period. We state the firm's dynamic optimization problem using a functional equation defined by (1) - (4) below.

$$\begin{aligned} v_0(k, b, \varepsilon_i; z_l, \mu) = \pi_d \max_n [z_l \varepsilon_i F(k, n) - \omega(z_l, \mu) n + \theta_k (1 - \delta) k - b] \\ + (1 - \pi_d) v(k, b, \varepsilon_i; z_l, \mu) \end{aligned} \tag{1}$$

After the start of the period, the firm knows which line of (1) will prevail. If it is not continuing beyond the period, the firm simply chooses labor to maximize its current dividend payment to shareholders. Because it will carry no capital or debt into the future, an exiting firm's dividends are its output, less wage payments and debt repayment, together with the remaining capital it can successfully uninstall at the end of the period. The problem conditional on continuation is more involved, because a continuing firm must choose its current labor and dividends alongside its future

[3] We select k_0 below so that each entrant's capital is χ fraction of the typical stock held across all firms in the long-run of our economy; that is, $k_0 = \chi \int k \widetilde{\mu}(d[k \times b \times \varepsilon])$, where $\widetilde{\mu}$ represents the steady-state distribution.

capital and debt. For expositional convenience, given the partial irreversibility in investment, we begin to describe this problem by defining the firm's value as the result of a binary choice between upward versus downward capital adjustment in (2), then proceed to identify the value associated with each option in (3) and (4).[4]

$$v\left(k,b,\varepsilon_i;z_l,\mu\right)=\max\Big\{v^u\left(k,b,\varepsilon_i;z_l,\mu\right),\ v^d\left(k,b,\varepsilon_i;z_l,\mu\right)\Big\} \tag{2}$$

Assume that $d_m\left(z_l,\mu\right)$ is the discount factor applied by firms to their next-period expected value if aggregate productivity at that time is z_m and the current aggregate state is (z_l,μ). Taking as given the evolution of ε and z according to the transition probabilities specified above, and taking as given the the evolution of the firm distribution, $\mu'=\Gamma\left(z,\mu\right)$, the firm solves the following two optimization problems to determine its values conditional on (weakly) positive and negative capital adjustment. (Here forward, except where necessary for clarity, we suppress the indices for current aggregate and firm productivity.) In each case, the firm selects its current employment and production, alongside the debt and capital with which it will enter into next period and its current dividends, D, to maximize its expected discounted dividends. As above, dividends are determined by the firm's budget constraint as the residual of its current production and borrowing after its wage bill and debt repayment have been covered, net of its investment expenditures.

Conditional on an upward capital adjustment, the firm solves the following problem constrained, respectively, by (i) the fact that investment must be non-negative, (ii)-(iii) the requirements that dividends be non-negative and satisfy the firm's budget constraint and (iv) a borrowing limit determined by its collateral.

$$v^u\left(k,b,\varepsilon_i;z_l,\mu\right)=\max_{n,k',b',D}\Big[D+\sum_{m=1}^{N_z}\pi^z_{lm}d_m\left(z,\mu\right)\sum_{j=1}^{N_\varepsilon}\pi_{ij}v_0\left(k',b',\varepsilon_j;z_m,\mu'\right)\Big] \tag{3}$$

subject to: $$k'\geq(1-\delta)\,k$$

$$0\leq D\leq z\varepsilon F\left(k,n\right)-\omega\left(z,\mu\right)n+q\left(z,\mu\right)b'-b-\left[k'-(1-\delta)\,k\right]$$

$$b'\leq\Theta k$$

The downward adjustment problem differs from that above only in that investment must be

[4] We could instead describe the firm's problem without the binary max operator by adopting an indicator function determining the relative price of capital as 1 in the event of $k'\geq(1-\delta)k$ and θ_k otherwise. Here, for sake of clarity, we opt for the less concise representation, though we will abandon it at some points below.

non-positive and, thus, its relative price is θ_k.

$$v^d(k, b, \varepsilon_i; z_l, \mu) = \max_{n,k',b',D} \Big[D + \sum_{m=1}^{N_z} \pi^z_{lm} d_m(z,\mu) \sum_{j=1}^{N_\varepsilon} \pi_{ij} v_0(k', b', \varepsilon_j; z_m, \mu') \Big] \tag{4}$$

subject to: $$k' \leq (1-\delta)k$$

$$0 \leq D \leq z\varepsilon F(k,n) - \omega(z,\mu)n + q(z,\mu)b' - b - \theta_k[k' - (1-\delta)k]$$

$$b' \leq \Theta k.$$

We will simplify the firm's problem and isolate its decision rules in section 4 below. For now, notice that there is no friction associated with the firm's employment choice, since the firm pays its current wage bill after production takes place, and its capital choice for next period also has no implications for current production. Thus, irrespective of their current debt or their continuation into the next period, all firms sharing in common the same (k, ε) combination select the same employment, which we will denote by $N(k, \varepsilon; z, \mu)$, and hence common production, $y(k, \varepsilon; z, \mu)$. The same cannot be said for the intertemporal decisions of continuing firms, given the presence of both borrowing limits and irreversibilities. Thus, $K(k, b, \varepsilon; z, \mu)$ and $B(k, b, \varepsilon; z, \mu)$ represent the choices of next-period capital and debt, respectively, made by firms sharing in common a complete individual type (k, b, ε).

3.2 Households

The economy is populated by a unit measure of identical households. Household wealth is held as one-period shares in firms, which we denote using the measure λ.[5] Given the prices they receive for their current shares, $\rho_0(k, b, \varepsilon; z, \mu)$, and the real wage they receive for their labor effort, $\omega(z, \mu)$, households determine their current consumption, c, hours worked, n^h, as well as the numbers of new shares, $\lambda'(k', b', \varepsilon')$, to purchase at prices $\rho_1(k', b', \varepsilon'; z, \mu)$. The lifetime

[5] Households also have access to a complete set of state-contingent claims. However, as there is no heterogeneity across households, these assets are in zero net supply in equilibrium. Thus, for sake of brevity, we do not explicitly model them here.

expected utility maximization problem facing each of them is listed below.

$$V^h(\lambda; z, \mu) = \max_{c, n^h, \lambda'} \left[U\left(c, 1 - n^h\right) + \beta \sum_{m=1}^{N_z} \pi^z_{lm} V^h\left(\lambda'; z_m, \mu'\right) \right] \tag{5}$$

subject to

$$c + \int_{\mathbf{S}} \rho_1\left(k', b', \varepsilon'; z, \mu\right) \lambda'\left(d\left[k' \times b' \times \varepsilon'\right]\right) \leq \omega(z, \mu) n^h + \int_{\mathbf{S}} \rho_0(k, b, \varepsilon; z, \mu) \lambda(d[\varepsilon \times k]).$$

Let $C^h(\lambda; z, \mu)$ describe the household choice of current consumption, and let $N^h(\lambda; z, \mu)$ be the allocation of current available to working. Finally, let $\Lambda^h(k', b', \varepsilon', \lambda; z, \mu)$ be the quantity of shares purchased in firms that will begin the next period with k' units of capital, b' units of debt, and idiosyncratic productivity ε'.

3.3 Recursive equilibrium

A *recursive competitive equilibrium* is a set of functions,

$$\left(\omega, q, (d_j)_{j=1}^{N_z}, \rho_0, \rho_1, v_0, N, K, B, V^h, C^h, N^h, \Lambda^h\right),$$

that solve firm and household problems and clear the markets for assets, labor and output, as described by the following conditions.

(i) v_0 solves (1) - (4), N is the associated policy function for exiting firms, and (N, K, B) are the associated policy functions for continuing firms

(ii) V^h solves (5), and $\left(C^h, N^h, \Lambda^h\right)$ are the associated policy functions for households

(iii) $\Lambda^h(k', b', \varepsilon_j, \mu; z, \mu) = \mu'(k', b', \varepsilon_j)$, for each $(k', b', \varepsilon_j) \in \mathbf{S}$

(iv) $N^h(\mu; z, \mu) = \int_{\mathbf{S}} \left[N(k, \varepsilon; z, \mu)\right] \mu(d[k \times b \times \varepsilon])$

(v) $C^h(\mu; z, \mu) = \int_{\mathbf{S}} \Big[z\varepsilon F(k, N(\varepsilon, k; z, \mu)) - (1 - \pi_d)\mathcal{J}\Big(K(k, b, \varepsilon; z, \mu) - (1 - \delta)k\Big)\Big(K(k, b, \varepsilon; z, \mu) - (1 - \delta)k\Big) + \pi_d \theta_k (1 - \delta) k \Big] \mu(d[k \times b \times \varepsilon])$, where $\mathcal{J}(x) = \begin{cases} 1 \text{ if } x \geq 0 \\ \theta_k \text{ if } x < 0 \end{cases}$

(vi) $\mu'(D, \varepsilon_j) = (1 - \pi_d) \int_{\{(k, b, \varepsilon_i) \mid (K(k, b, \varepsilon_i; z, \mu), B(k, b, \varepsilon_i; z, \mu)) \in D\}} \pi_{ij} \mu(d[k \times b \times \varepsilon_i]) + \pi_d \chi(k_0) H(\varepsilon_j)$, for all $(D, \varepsilon_j) \in \mathcal{S}$, defines Γ, where $\chi(k_0) = \{1 \text{ if } (k_0, 0) \in D; 0 \text{ otherwise}\}$

Using C and N to describe the market-clearing values of household consumption and hours worked satisfying conditions (iv) and (v) above, it is straightforward to show that market-clearing requires that (a) the real wage equal the household marginal rate of substitution between leisure and consumption, $\omega(z,\mu) = D_2U(C, 1-N)/D_1U(C, 1-N)$, that (b) the bond price, q^{-1}, equal the expected gross real interest rate, $q(z,\mu) = \beta \sum_{m=1}^{N_z} \pi^z_{lm} D_1U(C'_m, 1-N'_m)/D_1U(C, 1-N)$, and that (c) firms' state-contingent discount factors agree with the household discounted marginal utility of consumption across states $d_j(z,\mu) = \beta D_1U\left(C'_j, 1-N'_j\right)/D_1U(C, 1-N)$. Given these results, we may compute equilibrium by solving a single Bellman equation that combines the firm-level profit maximization problem with these equilibrium implications of household utility maximization, effectively subsuming the implications of households' decisions into the problems faced by firms.

Without loss of generality, we assign $p(z,\mu)$ as an output price at which firms value current dividends and payments and correspondingly assume that firms discount their future values by the household subjective discount factor. Given this alternative means of expressing firms' discounting, the following three conditions ensure all markets clear in our economy.

$$p(z,\mu) = D_1U(C, 1-N) \tag{6}$$

$$\omega(z,\mu) = D_2U(C, 1-N)/p(z,\mu) \tag{7}$$

$$q(z,\mu) = \beta \sum_{m=1}^{N_z} \pi^z_{lm} p\left(z_m, \mu'\right)/p(z,\mu) \tag{8}$$

A reformulation of (1) - (4) then yields an equivalent description of a firm's dynamic problem where each firm's value is measured in units of marginal utility, rather than output, with no change in the resulting decision rules. Suppressing the arguments of the price functions, exploiting the fact that the choice of n is independent of the k' and b' choices, and using the indicator function $\mathcal{J}(x) = \{1 \text{ if } x \geq 0 \text{ ; } \theta_k \text{ if } x < 0\}$ to distinguish the relative price of nonnegative versus negative investment, we have:

$$V_0(k, b, \varepsilon_i; z_l, \mu) = \pi_d \max_n p\Big[z_l\varepsilon_i F(k,n) - \omega n + \theta_k(1-\delta)k - b\Big] + (1-\pi_d)V(k, b, \varepsilon_i; z_l, \mu), \tag{9}$$

$$\text{where } V(k, b, \varepsilon_i; z_l, \mu) = \max_{n,k',b',D}\Big[pD + \beta \sum_{m=1}^{N_z}\sum_{j=1}^{N_\varepsilon} \pi^z_{lm}\pi_{ij} V_0\left(k', b', \varepsilon_j; z_m, \mu'\right)\Big] \tag{10}$$

$$\text{subject to } 0 \leq D \leq z\varepsilon F(k,n) - \omega n + qb' - b - \mathcal{J}\Big(k' - (1-\delta)k\Big)[k' - (1-\delta)k], \tag{11}$$

$$\text{and subject to } b' \leq \Theta. \tag{12}$$

4 Analysis

The problem listed in equations (9) - (10) forms the basis for solving equilibrium allocations in our economy, so long as the prices p, ω and q taken as given by our firms satisfy the restrictions in (6) - (8) above. From here, we begin to characterize the decision rules arising from this problem.[6] Each firm chooses its labor $n = N(k, \varepsilon; z, \mu)$ to solve $z\varepsilon D_2 F(k, n; z, \mu) = \omega(z, \mu)$, which immediately returns its current production, $y(k, \varepsilon) = z\varepsilon F(k, N(k, \varepsilon; z, \mu))$, so that any firm of type (k, b, ε) will achieve current profit flows $\pi(k, b, \varepsilon)$ defined below irrespective of its capital adjustment or borrowing decision.

$$\pi(k, b, \varepsilon) \equiv \pi(k, b, \varepsilon; z, \mu) = z\varepsilon F(k, N(k, \varepsilon; z, \mu)) - \omega(z, \mu) N(k, \varepsilon; z, \mu) - b$$

The challenging objects to determine are D, k' and b' for continuing firms. Turning to these, we will use a simple observation about the implications of borrowing constraints for the value a firm places on retained earnings versus dividends. As long as the firm places non-zero probability weight on encountering a future state in which its borrowing constraint will bind, the shadow value of retained earnings (the discounted sequence of reductions in the multipliers on future borrowing constraints) will necessarily exceed the shadow value of current dividends, p.[7] This means that, as long as the firm may face a binding borrowing limit in the future, it will set $D = 0$. In this case, equation 11 establishes that the firm's choice of k' directly implies its b', the level of debt with which it will enter into the next period. We refer to any such firm as a *constrained* firm, and list the resulting univariate problem it solves after deciding it will pay no dividends in the current period.

$$\begin{aligned} V^C(k, b, \varepsilon_i; z_l, \mu) &= \max_{k' \geq 0} \beta \sum_{m=1}^{N_z} \sum_{j=1}^{N_\varepsilon} \pi^z_{lm} \pi_{ij} V_0\left(k', b', \varepsilon_j; z_m, \mu'\right) \quad \text{subject to:} \qquad (13) \\ b' &= \frac{1}{q}\Big[-\pi(k, b, \varepsilon) + \mathcal{J}\Big(k' - (1-\delta)k\Big)[k' - (1-\delta)k]\Big] \\ \text{and } b' &\leq \Theta \end{aligned}$$

We can make a related observation about the value a firm will place on retained earnings versus dividends if it has accumulated sufficient wealth (via $k > 0$ or $b < 0$) such that collateral constraints will never again affect its investment activities, now or in the future. In this case, the

[6] While firm values and decisions are obviously influenced by the economy's aggregate state, we will suppress the z, μ arguments of the firm-level state vector (and in prices) to reduce notation in many instances below.

[7] This is easily proved using a sequence approach with explicit multipliers on each constraint; see Caggese (2007).

sequence of multipliers on future borrowing constraints are all zero, leaving the firm indifferent between saving its profits internally versus returning them to households. We can exploit this indifference by choosing for the firm which means of saving and borrowing it will adopt. Here, without loss of generality, we assume that the firm returns all earnings to its shareholders and sets its debt at zero in every period here and in future, $b' = 0$. In this case, a second look to equation 11 makes clear that the choice of k' directly implies the level of dividends the firm pays this period, D, which again delivers us a univariate optimization problem. We refer to any such firm as an *unconstrained* firm and denote its value by W.

$$W\left(k, b, \varepsilon_i; z_l, \mu\right) = \max_{k' \geq 0} \Bigg[p \Big[\pi\left(k, b, \varepsilon\right) - \mathcal{J}\Big(k' - \left(1 - \delta\right) k\Big) \left[k' - \left(1 - \delta\right) k\right] \Big] + \beta \sum_{m=1}^{N_z} \sum_{j=1}^{N_\varepsilon} \pi^z_{lm} \pi_{ij} W_0\left(k', 0, \varepsilon_j; z_m, \mu'\right) \Bigg], \text{ where} \tag{14}$$

$$W_0\left(k, b, \varepsilon_i; z_l, \mu\right) = \pi_d p \Big[\pi\left(k, b, \varepsilon\right) + \theta_k \left(1 - \delta\right) k \Big] + (1 - \pi_d) W\left(k, b, \varepsilon_i; z_l, \mu\right)$$

Notice that, if a firm has just become unconstrained, having entered into the period with some lingering debt (savings) $b \neq 0$, its value is linearly reduced (raised) by the associated reduction (rise) in current dividends, which are valued by p. Thus, we can alternatively express the value of any unconstrained firm of type (k, b, ε) as $w\left(k, \varepsilon\right) - pb$, where $w\left(k, \varepsilon\right) \equiv W\left(k, 0, \varepsilon\right)$. Moreover, the firm's beginning-of-period expected value inherits the same property; $W_0\left(k, b, \varepsilon_i; z_l, \mu\right) = w_0\left(k, \varepsilon\right) - pb$, where $w_0\left(k, \varepsilon\right) \equiv W_0\left(k, 0, \varepsilon\right)$.

4.1 (S,s) decisions among unconstrained firms

To characterize the decision rules of an unconstrained firm, it is expositionally useful to adopt the following less concise means of representing the problem in (14).

$$W\left(k, b, \varepsilon_i; z_l, \mu\right) = \max\{W^u\left(k, b, \varepsilon_i; z_l, \mu\right), W^d\left(k, b, \varepsilon_i; z_l, \mu\right)\}, \text{ where:}$$

$$W^u\left(k, b, \varepsilon_i; z_l, \mu\right) = p\pi(k, b, \varepsilon) + p(1 - \delta)k + \max_{k' \geq (1-\delta)k} \Big[-pk' + \beta \sum_{m=1}^{N_z} \sum_{j=1}^{N_\varepsilon} \pi^z_{lm} \pi_{ij} w_0\left(k', \varepsilon_j; z_m, \mu'\right) \Big] \tag{15}$$

$$W^d\left(k, b, \varepsilon_i; z_l, \mu\right) = p\pi(k, b, \varepsilon) + p\theta_k(1 - \delta)k + \max_{k' \leq (1-\delta)k} \Big[-p\theta_k k' + \beta \sum_{m=1}^{N_z} \sum_{j=1}^{N_\varepsilon} \pi^z_{lm} \pi_{ij} w_0\left(k', \varepsilon_j; z_m, \mu'\right) \Big] \tag{16}$$

In the above, W^u and W^d are both strictly increasing in k. This in turn implies that W and W_0 are strictly increasing functions of the unconstrained firm's capital, as are the w and w_0 functions defined above.

We may characterize the capital decision rule for an unconstrained firm by reference to two target capital stocks, the upward and downward adjustment targets that would solve the problems in (15) and (16), respectively, were there no sign restrictions on investment. Define the upward target, k_u^*, as the capital a firm would choose given a unit relative price of investment, and define the downward target, k_d^*, as the capital a firm would choose given a relative price at θ_k.

$$k_u^*(\varepsilon_i) = \arg\max_{k'} \Big[-pk' + \beta \sum_{m=1}^{N_z} \sum_{j=1}^{N_\varepsilon} \pi_{lm}^z \pi_{ij} w_0 \left(k', \varepsilon_j; z_m, \mu'\right) \Big] \tag{17}$$

$$k_d^*(\varepsilon_i) = \arg\max_{k'} \Big[-p\theta_k k' + \beta \sum_{m=1}^{N_z} \sum_{j=1}^{N_\varepsilon} \pi_{lm}^z \pi_{ij} w_0 \left(k', \varepsilon_j; z_m, \mu'\right) \Big] \tag{18}$$

Notice that each target is independent of current capital and depends only on the aggregate state and the firm's current ε. As such, all unconstrained firms that share in common the same current productivity ε have the same upward and downward target capitals. Note also that, because $\theta_k < 1$ (and because the value function w_0 is strictly increasing in k), the upward adjustment target necessarily lies below the downward target: $k_u^* < k_d^*$.

We are now in a convenient position to retrieve the unconstrained firm's capital decision rule. Given a constant price associated with raising (lowering) its capital stock, and because w_0 is increasing in k, the firm selects a future capital as close to the upward (downward) target as its constraint set allows. Thus, the firm's decision rules conditional on upward adjustment and downward adjustment are as follow.

$$k_u(\varepsilon) = \max\{(1-\delta)k, k_u^*(\varepsilon)\} \quad \text{and} \quad k_d(\varepsilon) = \min\{(1-\delta)k, k_d^*(\varepsilon)\}$$

Given these conditional adjustment rules, we know that an unconstrained firm of type (k, b, ε) selects one of three future capital levels, $k' \in \{k_u^*(\varepsilon), k_d^*(\varepsilon), (1-\delta)k\}$. Which one it selects depends only on where its current capital lies in relation to its two targets.

Recalling that $k_u^*(\varepsilon) < k_d^*(\varepsilon)$, if $k \in \left[\frac{k_u^*(\varepsilon)}{1-\delta}, \frac{k_d^*(\varepsilon)}{1-\delta}\right]$ then $k_u(\varepsilon) = (1-\delta)k = k_d(\varepsilon)$, so the firm makes no adjustment to its capital. If, instead, the firm's capital is sufficiently low that its implied stock for next period under no adjustment lies below the upward target, $k < \frac{k_u^*(\varepsilon)}{1-\delta}$, then $k_u(\varepsilon) = k_u^*(\varepsilon)$, while $k_d(\varepsilon) = (1-\delta)k$. In this case, the firm selects $k_u^*(\varepsilon)$, since $(1-\delta)k$ is in the constraint set for upward capital adjustment. Finally, if the firm's implied capital for next period under no adjustment lies above the downward target, $k > \frac{k_d^*(\varepsilon)}{1-\delta}$, then $k_d(\varepsilon) = k_d^*(\varepsilon)$, while $k_u(\varepsilon) = (1-\delta)k$. In this case, the firm selects $k_d^*(\varepsilon)$, since $(1-\delta)k$ is in the constraint set for a downward adjustment. Collecting these observations, we have the following (S,s) capital decision rule for an unconstrained firm.

$$K^W(k,\varepsilon) = \begin{cases} k_u^*(\varepsilon; z, \mu) & \text{if } k < \frac{k_u^*(\varepsilon;z,\mu)}{1-\delta} \\ (1-\delta)k & \text{if } k \in \left[\frac{k_u^*(\varepsilon;z,\mu)}{1-\delta}, \frac{k_d^*(\varepsilon;z,\mu)}{1-\delta}\right] \\ k_d^*(\varepsilon; z, \mu) & \text{if } k > \frac{k_d^*(\varepsilon;z,\mu)}{1-\delta} \end{cases} \tag{19}$$

Given the capital rule, and recalling that $B^W(k,b,\varepsilon) = 0$, we can directly retrieve the firm's dividends, and thus its value.[8]

$$D^W(k,b,\varepsilon) = \pi(k,b,\varepsilon) - \mathcal{J}\Big(K^W(k,\varepsilon_i) - (1-\delta)k\Big)[K^W(k,\varepsilon) - (1-\delta)k] \tag{20}$$

$$W(k,b,\varepsilon_i;z_l,\mu) = pD^W(k,b,\varepsilon) + \beta\sum_{m=1}^{N_z}\sum_{j=1}^{N_\varepsilon}\pi_{lm}^z\pi_{ij}W_0\left(K^W(k,\varepsilon), 0, \varepsilon_j; z_m, \mu'\right), \tag{21}$$

$$\text{where } W_0(k,b,\varepsilon_i;z_l,\mu) = \pi_d p\Big[\pi(k,b,\varepsilon) + \theta_k(1-\delta)k\Big] + (1-\pi_d)W(k,b,\varepsilon_i;z_l,\mu)$$

4.2 Decisions among constrained firms

We now consider the decisions made by a firm that has not previously attained sufficient wealth to become unconstrained. The first essential step is to establish whether or not the firm has crossed the relevant wealth threshold to become unconstrained. If it has, the decision rules isolated above apply. If it has not, the collateralized borrowing constraint will continue to influence its investment decisions, so that the capital and debt decisions remain intertwined.

To ascertain whether a firm of type (k,b,ε) has become unconstrained, we need only consider whether it is feasible for the firm to adopt the capital rule $K^W(k,\varepsilon)$ and pay no dividends in the current period. Recall that, if a firm is unconstrained, it is indifferent between dividend payments versus retained earnings; hence an alternative set of decision rules achieving equal value to those

[8]Equation (20) suggests a potential inconsistency with the restriction on our model that firms may not pay negative dividends. Here, however, the Modigliani-Miller theorem applies. Thus, whenever D^W prescribes the payment of negative dividends, the firm is indifferent and actually assumes a debt of $-D^W/q$, rather than the 0 we record here.

above is $k' = K^W(k,\varepsilon)$, $D = 0$ and $b' = -D^W(k,b,\varepsilon)/q$. We use these alternative rules to verify whether our potentially constrained firm has become unconstrained, simply considering whether their adoption implies a level of debt that is permissible given the firm's ability to borrow. If the firm (k,b,ε) has sufficient wealth that $-D^W(k,b,\varepsilon)/q \leq \Theta k$, it adopts the decision rules (19) - (20), sets $b' = 0$, achieves value $W(k,b,\varepsilon;z,\mu)$ from (21), and exits the period indistinguishable from any other unconstrained firm that entered it with (k,ε).

$$V(k,b,\varepsilon_i;z_l,\mu) = W(k,b,\varepsilon;z,\mu) \text{ iff } -D^W(k,b,\varepsilon)/q \leq \Theta k$$

If the inequality above is not satisfied, the firm remains constrained; for any such firm surviving beyond the current period, $V(k,b,\varepsilon_i;z_l,\mu) = V^C(k,b,\varepsilon_i;z_l,\mu)$. To isolate the decisions made by a continuing constrained firm facing the problem in (13), we again find it useful to adopt a less concise representation and again suppress notation for the aggregate state where convenient.

$$V^C(k,b,\varepsilon_i;z_l,\mu) = \max\{V^u(k,b,\varepsilon_i;z_l,\mu), V^d(k,b,\varepsilon_i;z_l,\mu)\}, \quad \text{where:} \tag{22}$$

$$\begin{aligned} V^u(k,b,\varepsilon_i;z_l,\mu) &= \max_{k' \geq (1-\delta)k} \beta \sum_{m=1}^{N_z} \sum_{j=1}^{N_\varepsilon} \pi^z_{lm} \pi_{ij} V_0\left(k', b'_u(k'), \varepsilon_j; z_m, \mu'\right), \text{ with} \\ b'_u(k') &\equiv \frac{1}{q}\Big(-\pi(k,b,\varepsilon_i) + [k' - (1-\delta)k]\Big) \\ \text{subject to:} \quad & b'_u(k') \leq \Theta k \end{aligned} \tag{23}$$

$$\begin{aligned} V^d(k,b,\varepsilon_i;z_l,\mu) &= \max_{k' \leq (1-\delta)k} \beta \sum_{m=1}^{N_z} \sum_{j=1}^{N_\varepsilon} \pi^z_{lm} \pi_{ij} V_0\left(k', b'_d(k'), \varepsilon_j; z_m, \mu'\right), \text{ with} \\ b'_d(k') &\equiv \frac{1}{q}\Big(-\pi(k,b,\varepsilon_i) + \theta_k[k' - (1-\delta)k]\Big) \\ \text{subject to:} \quad & b'_d(k') \leq \Theta k \end{aligned} \tag{24}$$

We approach the constrained firm's problem as follows. First, given its (k,ε), we isolate a cutoff debt level under which (23) is a feasible option. The lowest choice of k' permitted by the non-negativity constraint on investment is $(1-\delta)k$. If this choice is not affordable given the firm's borrowing constraint, it cannot undertake even a trivial upward capital adjustment. Recalling the definition of $\pi(k,b,\varepsilon)$, this is the case if $\frac{1}{q}[b + \omega N(k,\varepsilon) - z\varepsilon F(k, N(k,\varepsilon))] > \Theta k$. Thus, among any group of firms sharing a common (k,ε), only those with debt not exceeding $b^T(k,\varepsilon)$ can consider an upward adjustment, where the threshold debt level is:

$$b^T(k,\varepsilon) \equiv q\theta_b\theta_k k + z\varepsilon F(k, N(k,\varepsilon)) - \omega N(k,\varepsilon).$$

Firms with $b > b^T(k,\varepsilon)$ do not solve (23); for them, $V^C(k,b,\varepsilon;z,\mu) = V^d(k,b,\varepsilon;z,\mu)$.

To solve the problems (23) - (24), we identify the maximum capitals permitted by the borrowing constraint under upward versus downward capital adjustment, and then impose the relevant sign restrictions on investment to arrive at the constraint sets associated with each option.

$$\begin{aligned}
\overline{k}_u(k,b,\varepsilon) &\equiv (1-\delta)k + \Big[q\theta_b\theta_k k + \pi(k,b,\varepsilon)\Big] \\
\overline{k}_d(k,b,\varepsilon) &\equiv (1-\delta)k + \frac{1}{\theta_k}\Big[q\theta_b\theta_k k + \pi(k,b,\varepsilon)\Big] \\
\Lambda^u(k,b,\varepsilon) &= [(1-\delta)k, \overline{k}_u(k,b,\varepsilon)] \\
\Lambda^d(k,b,\varepsilon) &= [0, \min\{(1-\delta)k, \overline{k}_d(k,b,\varepsilon)]
\end{aligned}$$

Substituting in the debt implied by each capital choice and making use of our findings above, we may express the constrained firm's value as follows.

$$V^C(k,b,\varepsilon_i;\cdot) = \max\{V^u(k,b,\varepsilon_i;z_l,\mu), V^d(k,b,\varepsilon_i;z_l,\mu)\}, \quad \text{where:} \tag{25}$$

$$V^u(k,b,\varepsilon_i;\cdot) = \max_{k'\in\Lambda^u(k,b,\varepsilon)} \beta \sum_{m=1}^{N_z}\sum_{j=1}^{N_\varepsilon} \pi^z_{lm}\pi_{ij} V_0\left(k', \frac{k'-(1-\delta)k-\pi(k,b,\varepsilon_i)}{q}), \varepsilon_j; z_m, \mu'\right),$$

$$V^d(k,b,\varepsilon_i;\cdot) = \max_{k'\in\Lambda^d(k,b,\varepsilon)} \beta \sum_{m=1}^{N_z}\sum_{j=1}^{N_\varepsilon} \pi^z_{lm}\pi_{ij} V_0\left(k', \frac{\theta_k[k'-(1-\delta)k]-\pi(k,b,\varepsilon_i)}{q}, \varepsilon_j; z_m, \mu'\right),$$

and where:

$$\begin{aligned}
V_0(k,b,\varepsilon_i;\cdot) &= \pi_d p[\pi(k,b,\varepsilon) + \theta_k(1-\delta)k] + (1-\pi_d)V(k,b,\varepsilon_i;z_l,\mu), \\
V(k,b,\varepsilon_i;\cdot) &= \begin{cases} W(k,b,\varepsilon;z,\mu) \ \text{if } b \le -\mathcal{J}\Big(K^W(k,\varepsilon_i)-(1-\delta)k\Big)[K^W(k,\varepsilon_i)-(1-\delta)k] \\ \qquad +q\theta_b\theta_k k + y(k,\varepsilon) - wN(k,\varepsilon;z,\mu) \\ V^C(k,b,\varepsilon_i;z_l,\mu) \ \text{otherwise} \end{cases}
\end{aligned}$$

Denoting the capitals that solve the conditional adjustment problems above by $\widehat{k}^u(k,b,\varepsilon_i;\cdot)$ and $\widehat{k}^d(k,b,\varepsilon_i;\cdot)$, and recalling that $D^C(k,b,\varepsilon_i;\cdot) = 0$, we obtain the following decision rules for capital and debt.

$$K^C(k,b,\varepsilon_i;\cdot) = \begin{cases} \widehat{k}^u(k,b,\varepsilon_i;\cdot) & \text{if } V^C(k,b,\varepsilon_i;\cdot) = V^u(k,b,\varepsilon_i;z_l,\mu) \\ \\ \widehat{k}^d(k,b,\varepsilon_i;\cdot) & \text{if } V^C(k,b,\varepsilon_i;\cdot) = V^d(k,b,\varepsilon_i;z_l,\mu) \end{cases} \tag{26}$$

$$B^C(k,b,\varepsilon_i;\cdot) = \frac{1}{q}\Big[\mathcal{J}\Big(K^C(k,b,\varepsilon_i;\cdot)-(1-\delta)k\Big)[K^C(k,b,\varepsilon_i;\cdot)-(1-\delta)k] - \pi(k,b,\varepsilon_i)\Big] \tag{27}$$

The numerical algorithm we use to solve our model is an extension of that described in Khan and Thomas (2003, 2008) using the analysis above. More specifically, our solution involves repeated application of the contraction mapping implied by (25) to solve the constrained firm value function V^C, given the price functions $p(z,\mu)$, $\omega(z,\mu)$ and $q(z,\mu)$ and the laws of motion implied by $\Gamma(z,\mu)$, (π_{ij}) and (π^z_{lm}). In each instance, the starting point is solving (21) to isolate the unconstrained firm value function W, which serves as an input for V^C.

5 Calibration

In the sections to follow, we will consider how the mechanics of our (full) model with real and financial frictions compare to those in two relevant reference models - one where there are no borrowing limits ($\theta_b \to \infty$) and one where there is no specificity in capital ($\theta_k = 1$). These two reference models will help us to isolate how much the *interaction* between credit constraints and micro-level capital rigidities influences our economy's aggregate dynamics. Aside from the values of θ_b and θ_k, all three models share a common parameter set that is selected in our full model to best match moments drawn from postwar U.S. aggregate and firm-level data. To be clear, we do not re-calibrate the reference models; thus, the average capital/output ratio, hours worked, and other important aspects of these economies are allowed to vary as each friction is eliminated.

5.1 Functional forms

Across our model economies, we assume that the representative household's period utility is the result of indivisible labor (Rogerson (1988)): $u(c,L) = \log c + \varphi L$. The firm-level production function is Cobb-Douglas: $z\varepsilon F(k,n) = z\varepsilon k^{\alpha} n^{\nu}$. The initial capital stock of each entering firm is a fixed χ fraction of the typical stock held across all firms in the long-run of our full economy; that is, $k_0 = \chi \int k\widetilde{\mu}(d[k \times b \times \varepsilon])$, where $\widetilde{\mu}$ represents the steady-state distribution therein.

In specifying our exogenous stochastic process for aggregate productivity, we begin by assuming a continuous shock following a mean zero AR(1) process in logs: $\log z' = \rho_z \log z + \eta'_z$ with $\eta'_z \sim N\left(0, \sigma^2_{\eta_z}\right)$. Next, we estimate the values of ρ_z and σ_{η_z} from Solow residuals measured using NIPA data on US real GDP and private capital, together with the total employment hours series constructed by Prescott, Ueberfeldt, and Cociuba (2005) from CPS household survey data, over the years 1959-2002, and we discretize the resulting productivity process using a grid with 3 shock realizations ($N_z = 3$) to obtain (z_l) and (π^z_{lm}). We determine the firm-specific productivity shocks (ε_i) and the Markov Chain governing their evolution (π_{ij}) similarly by discretizing a log-normal process, $\log \varepsilon' = \rho_\varepsilon \log \varepsilon + \eta'$ using 9 values ($N_\varepsilon = 9$).

5.2 Aggregate targets

We set the length of a period to correspond to one year, and we determine the values of β, ν, δ, α, φ and θ_b using moments from the aggregate data as follows. First, we set the household discount factor, β, to imply an average real interest rate of 4 percent, consistent with recent findings by Gomme, Ravikumar and Rupert (2008). Next, the production parameter ν is set to yield an average labor share of income at 0.60 (Cooley and Prescott (1995)). The depreciation rate, δ, is taken to imply an average investment-to-capital ratio of roughly 0.069, which corresponds to the average value for the private capital stock between 1954 and 2002 in the U.S. Fixed Asset Tables, controlling for growth. Given this value, we determine capital's share, α, so that our model matches the average private capital-to-output ratio over the same period, at 2.3, and we set the parameter governing the preference for leisure, φ, to imply an average of one-third of available time is spent in market work. Finally, we select the parameter governing the extent of financial frictions in our model, θ_b, to imply an average debt-to-capital ratio matching that of nonfarm nonfinancial businesses over 1952-05 in the Flow of Funds, 0.366.

5.3 Firm-level targets

The parameters we determine using moments drawn from firm-level data are the exit rate, π_d, the fraction of the steady-state aggregate capital stock held by each entering firm, χ, the extent of reversibilty in capital, θ_k, and the persistence and variability of the firm-specific productivity shocks, ρ_ε and σ_η. We set the exit rate at 10 percent, so that each entering firm expects to remain in the economy for roughly 10 years. Next, we set $\chi = 0.10$ so that entering firms are, on average, one-tenth the size of the typical firm in our economy (Davis and Haltiwanger (1992)).

Finally, we choose θ_k, ρ_ε and σ_η jointly to reproduce three aspects of establishment-level investment data documented by Cooper and Haltiwanger (2006) based on a 17-year sample drawn from the Longitudinal Research Database. These targets are (i) the average standard deviation of investment rates (i/k) across establishments: $\sigma_{i/k} = 0.337$, (ii) the average serial correlation of investment rates: $\rho_{i/k} = 0.058$, and (iii) the fraction of establishment-year observations wherein a positive investment spike ($i/k > 0.20$) occurs: 0.186. While our models has life-cycle aspects affecting firms' investments, the Cooper and Haltiwanger (2006) dataset includes only large manufacturing establishments that remain in operation throughout their sample period, Thus, in undertaking this part of our calibration, we must select an appropriate model-generated sample for comparability with their sample. This we do by simulating a large number of firms for 30 years, retaining only those firms that survive throughout, and then restricting the dates over which investment rates are measured to years 11 through 30 to eliminate life-cycle effects.

5.4 Resulting parameters

The table below lists the parameter set we obtain from our calibration.

β	ν	δ	α	φ	ρ_z	σ_{η_z}	θ_b	π_d	χ	θ_k	ρ_ε	σ_η
0.96	0.60	0.06	0.25	2.15	0.852	0.014	0.56	0.10	0.10	0.95	0.613	0.16

Note that these parameters imply minimal real frictions in our model economy, with only a 5 percent loss incurred in uninstalling capital, and considerable financial frictions, with firms able to take on debt only up to 53 percent of the value of their capital. Also note that firm-level shocks are far more volatile and less persistent than aggregate ones. Given these aspects of the calibration, our model gives rise to a stationary distribution of firms over (k, b, ε) wherein roughly 50 percent of firms are constrained (using the definition from section 4 above) and the remaining 50 percent are unconstrained.

6 Steady state

We begin by considering the implications of borrowing limits and irreversibilities for the typical decisions made in our economy. Figure 1 overviews the stationary distribution of firms in the baseline case of our full model, presenting three slices of the full distribution. In the top panel, we see the distribution of firms over capital and debt levels at the lowest firm-level productivity, while the middle and bottom present the counterparts at the median and highest levels of productivity.

Note that each panel of Figure 1 appears to have two essentially disjoint distributions. The smoother, more connected distribution where capital is relatively high (above 0.6 in the middle panel) and recorded debt is zero corresponds to older, wealthier firms that are unconstrained.

Elsewhere, roughly 10 percent of all firms newly entering the economy each period are scattered across each productivity level according to its ergodic distribution. These firms enter with zero debt and very low initial capital (roughly 0.12), as shown by the right-most $\mu(k, b)$ spike in each panel.

After its first date in production, each new firm begins to take on debt in effort to build up its capital. In the absence of the collateralized borrowing limits, young firms would immediately take on a large, temporary debt that would allow them to jump to the capital stock selected by unconstrained firms with the same current productivity level. Instead, however, the borrowing limits gradualize their adjustments, so we see ripples of these entering firms slowly moving into higher ranges of k and b. Over time, those firms that survive long enough eventually reach a level of capital (and current productivity) such that, having accumulated sufficient collateral, they are no longer affected by borrowing limits. As this occurs, the firm in question leaves the constrained group, jumping to the unconstrained distribution. As would be expected, the mean capital among unconstrained firms rises with firm-level productivity. The same is true for constrained firms, though this is somewhat harder to see in the current figure.

Figure 2 is the no-financial frictions counterpart to Figure 1. It shows the stationary distribution of firms at the same three productivity slices in an otherwise identical economy where no firm is ever constrained (with firms able to borrow far more than the collateral value of their capital). As in the previous figure, new firms enter the economy with low capital. Here, however, these young firms immediately borrow all that they require to reach their unconstrained capital targets in the second period of life, so they join the mass of unconstrained firms with no delay. As such, there are no life-cycle aspects to investment in this economy and, for the most part, firms operate at a scale that is appropriate to their productivity. The quantitative impact of this more efficient allocation of production is that steady state output rises by 8.8 percent relative to our full economy depicted in Figure 1.

Returning to the economy with both frictions in place, Figure 3 illustrates the pure effects of the irreversibility in cases where it does not interact with the financial friction in our economy. Here, we summarize the capital choices made by unconstrained firms entering the period with various levels of capital (measured on the x-axis) and debt (measured on the y-axis), conditional on a current productivity draw. The top panel depicts firms entering with a low productivity value, the middle panel shows those with the median value, and the bottom panel shows those with a high productivity. The z-axis in each panel reports an indicator variable that takes on a value of 1 for unconstrained firms that invest positively to the upward target capital consistent with their current productivity, a value of 2 for those investing negatively to the relevant downward target, and a value of 5 for those that remain inactive with respect to their capital, setting investment to zero. (Areas along the floor of each panel are combinations of (k, b) where firms are financially constrained.)

The region of (k, b) where firms invest to their upward target expands into higher current capital levels as one looks from the top panel downward, since rises in current productivity predict higher marginal product of capital schedules next period. Looking leftward from these regions are the areas with zero investment induced by the irreversibility in capital. While the loss associated with uninstalling capital in our economy is only 5 percent, it nonetheless makes some firms quite reluctant to shed capital. Those with higher current productivities are more so, given the persistence in ε alongside depreciation. As such, the inactivity region expands to higher capital levels as productivity rises, while the region associated with downward investment shrinks, finally disappearing by the bottom panel.

Note that Figure 3 is largely an expositional device. It depicts the directional capital choice adopted by unconstrained firms at each *potential* firm-level state rather than at states actually populated in the economy's stationary distribution. Restricting consideration to those states, the actual fraction of all firms that are unconstrained and adjust to the upward target consistent with their productivity is 18 percent, the fraction remaining inactive is 26 percent, and the fraction undertaking negative investment is 6 percent.

Figure 4 is analogous to figure 3. Again conditional on currently productivity, it illustrates the capital decisions taken by firms, this time considering those that are affected by both the real friction in our economy and the financial one. Such firms are located in regions of the (k, b) space to the right and back where capital is low and/or debt is high. (Areas along the floor of each panel are combinations of (k, b) where firms are unconstrained.) In the steady state, constrained firms' capital decisions are largely determined by life-cycle aspects imposed by collateralized borrowing, so only only two values of the adjustment-type indicator are relevant in this figure. Firms that invest positively to the maximum capital permitted by their ability to borrow are reflected by a value of 3 on the z-axis. Among constrained firms, these are the ones with higher current productivity, comparatively high capital, and comparatively low debt. There are more such firms at higher values of ε, as higher current production expands funds available for investment. To their right and back in each panel are firms with low capital and high debt combinations that are forced to shed capital to repay their existing debt. These firms (with an adjustment-type indicator 4) represent a larger fraction of constrained firms at lower values of ε, as it becomes harder to repay debt when current production is low. Here again, however, the figure is expositional. In the actual stationary distribution, no such firms exist.

Finally, Figure 5 illustrates the levels of output produced across the full range of capital and debt levels. As one would expect, the level of production at any given (k, b) combination rises with the level of productivity, and, examining any single current productivity, production rises with the firm's capital stock. However, the level of debt has no influence here, since we do not require our firms to pay their wage bills until after current production is done. Thus, the economy with no financial frictions has output figures corresponding to this one are essentially unchanged. Of course, as we have seen above, this does not imply that steady state output, productivity and the distribution of production are unchanged, as each of these is influenced by the resulting stationary distribution of firms over the firm-level state space.

7 Results

We begin to explore the dynamics by first considering the effect each friction in our economy has on its typical business cycle. Table 1 presents some commonly reported business cycle statistics derived from an HP-filtered 1000 period simulation of our full model economy under the assumption that aggregate productivity shocks are the only source of aggregate fluctuations, Table 2 presents the corresponding moments when we eliminate financial frictions, and Table 3 is the same economy with neither collateral constraints nor capital specificity. As expected, each friction acts to reduce the average levels of output, capital, and consumption over our simulation. Most notably, average output rises by roughly 8 percent when financial frictions are stripped away, then another 2 percent when the irreversibility is also eliminated.

Moving to consider second moments, there are some small differences across the three tables. Output volatility rises between our full economy and the counterpart model without limits to borrowing, and rises again between that model and the one with no frictions. Despite this, as each friction is lifted, the representative household grows more effective in smoothing its consumption. As the contemporaneous correlation between consumption and production is slightly weakened from one table to the next, consumption's standard deviation (raw and relative) falls. Elsewhere, the volatility of hours worked rises steadily, and is marginally more correlated with output as each friction is eliminated. The same monotone pattern does not follow for investment expenditures, however. There, the relative standard deviation falls from 4.25 percent to 4.05 percent as the financial friction is stripped away, allowing the inertia associated with irreversibility more prominence, while it rises to 4.40 percent when the irreversibility is eliminated.

While we have mentioned some minor differences in the business cycle moments across tables 1 through 3, two points are surely more important. The first is that the business cycle moments drawn from our full model in Table 1 are similar to those of a typical real business cycle model without its complications (for instance, Hansen (1985)). Output volatility is roughly 1.7 percent, consumption is about half as volatile as output, and investment roughly four times as volatile as output. We also see the customary strong positive contemporaneous correlations with output in consumption, investment, hours and wages. While the usual difficulties of excessive investment volatility and weak hours volatility are a bit more pronounced here relative to most representative firm real business cycle models, these come from our differing returns to scale in production rather than either friction we mean to study; the same features are present in Table 3 with both removed.

This brings us to our second point. Despite the differences raised above, the second moments across all three tables are actually quite similar on the whole. Comparing Table 1 to Table 2, in particular, it appears that the typical business cycle in our economy is relatively impervious to some ordinary, ongoing degree of financial frictions, although it implies roughly half of firms in the economy are financially constrained. That noted, however, this is not at all the question we have set out to answer. Our main interest here is to understand what happens when the extent of financial frictions suddenly and unexpectedly grows far more severe than is normal. We explore this question via a series of impulse response figures to which we turn now..

Figure 6 depicts our economy's response to a financial crisis, absent any technology shock. More specifically, it is the response to a 30 percentage point drop in the value of firms' collateral, as generated by a reduction in θ_b, which implies a halving of new debt issuance. In designing this exercise, we assume that firms predict a return to normal financial conditions will ultimately occur. Each period, they place 40 percent probability weight on a full financial recovery in the subsequent period. Thus, when the shock occurs in period 1, they expect that it will persist for 2.5 years.

Although the distribution of capital is predetermined, aggregate production falls by about 0.3 percent (relative to its simulated mean in normal financial times) when the financial shock hits in year 1. This is, of course, a direct consequence of the 0.5 percent fall in the labor input, which is, in turn, a reaction to the reduced expected return to investment. With the sudden reduction in credit, there is a sharp rise in the fraction of firms that are financially constrained (from roughly 50 percent to roughly 67 percent). Underlying this rise, young firms are now far more curtailed in their investment activities relative to the pre-shock economy, and thus will take considerably longer to outgrow the financial frictions and begin producing at a scale consistent with their productivities. Moreover, some mature firms that had been unconstrained in the pre-shock economy now find their collateral insufficient to prevent financial frictions once again influencing their investment plans. These larger constrained firms initially exhibit life-cycle investment similar to that in their youth, accumulating capital in effort to outgrow the new financial friction irrespective of their productivities.

Notice that consumption does not immediately fall when the financial shock hits our economy. Anticipating a sharp distortion to the distribution of production over coming years, and thus unusually low endogenous total factor productivity, the representative household in our economy expects a lowered return to saving. This actually encourages a slight rise in consumption at the impact of the shock, and thus a rise in leisure. This indirect endogenous TFP influence on hours worked is compounded by the fact that the initial aggregate capital stock is roughly 9 percent above that consistent with the tighter borrowing conditions , which further encourages consumption. The fall in investment does not support consumption for long, however; it falls below pre-shock average by year 3, and steadily declines for roughly 10 years before it levels off.

Elsewhere, labor falls much more quickly in the early periods following the shock, with more than a 1 percent drop between dates 1 and 2 alone. Given the severe misallocation of capital at the start of date 2, alongside reductions in the total capital stock, the marginal product of labor drops sharply over early periods, further discouraging its supply relative to date 1. Over the first 5 periods of the impulse, it falls roughly 2.5 percent below its pre-shock level, and thereafter it takes nearly 15 periods to rise back to the level consistent with the new financial setting. This long adjustment period is a reflection of the time that it takes for the capital distribution to settle.

On balance, we take the following observation from Figure 6. A tightening of collateral constraints alone, a purely financial shock, is capable of large and persistent real effects in our model economy. In the example we have shown here, the misallocation of capital arising from tight financial conditions is compounded by the reductions in aggregate capital, productivity, and labor that it causes. As a result, there are protracted adjustments in aggregate quantities lasting a decade or more, and GDP is ultimately reduced by roughly 4 percent, while aggregate consumption falls by 3 percent.

We next consider what implications the prolonged financial crisis from above can have if its onset is shortly followed by a 1 percent negative technology shock. As seen in the top panel of Figure 7, exogenous TFP falls two years after the financial shock hits, and thereafter gradually reverts to its mean. Were credit markets functioning as normal when this TFP shock appeared, output would fall less than 1.5 percent, labor would fall roughly 0.5 percent, and the half-life of the responses would be roughly 5 years. In this case, however, tight credit markets compound the effects of the technology shock to yield a larger and more lasting recession.

With employment and production already contracting due to the increased inefficiency in capital allocation, labor drops almost 3 percent below its average at the impact of the productivity shock, while GDP drops to roughly 3.8 percent below average. Thereafter, although exogenous TFP is smoothly rising back to trend, the financial crisis continues to drag real quantities downward for several more years. Indeed, total production is still about 4 percent below trend a decade later.

8 Concluding remarks

To be added.

References

[1] Bertola, G. and R. J. Caballero (1994) "Irreversibility and Aggregate Investment," *Review of Economic Studies* 61, 223-246.

[2] Buera, F. and Shin, Y. (2007) "Financial Frictions and the Persistence of History: A Quantitative Exploration," manuscript, Northwestern University.

[3] Caballero, R. J. (1999) "Aggregate Investment," chapter 12 in M. Woodford and J. Taylor (eds.) *Handbook of Macroeconomics,* vol. IB. Amsterdam: North Holland.

[4] Caballero, R. J and E. M. R. A. Engel (1999) "Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S, s) Approach," *Econometrica* 67, 783-826.

[5] Caggese, A. (2007) "Financing constraints, irreversibility, and investment dynamics" *Journal of Monetary Economics* 54, 2102–2130.

[6] Comin, D. and T. Philippon (2005) "The Rise in Firm-Level Volatility: Causes and Consequences," NBER Working Paper No. 11388.

[7] Cooley, T. F., Prescott, E. C., 1995. "Economic Growth and Business Cycles." in *Frontiers of Business Cycle Research*, Cooley, T. F. Editor, Princeton University Press.

[8] Cooper, R. W. and J. C. Haltiwanger (2006) "On the Nature of Capital Adjustment Costs," *Review of Economic Studies* 73, 611-633.

[9] Cordoba, J.C. and M. Ripoll (2004) "Credit Cycles Redux" *International Economic Review* 45, 1011-46.

[10] Steven J. Davis, S.J. and J.C. Haltiwanger (1992) "Gross Job Creation, Gross Job Destruction, and Employment Reallocation," *Quarterly Journal of Economics* 107(3), 19-63.

[11] Gomme, P., Ravikumar, B. and P. Rupert (2008), "The Return to Capital and the Business Cycle," Concordia University Working Paper 08-002.

[12] Hansen, G. D. (1985) "Indivisible Labor and the Business Cycle," *Journal of Monetary Economics* 16, 309-327.

[13] Johnson, S. A. (1989) *Spline approximation in discrete dynamic programming with application to stochastic multi-reservoir systems* Unpublished dissertation (Cornell, Ithaca, NY).

[14] Khan, A. and J. K. Thomas (2003) "Nonconvex Factor Adjustments in Equilibrium Business Cycle Models: Do Nonlinearities Matter?" *Journal of Monetary Economics* 50, 331-360.

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

[16] Kiyotaki, N. and J. Moore (1997) "Credit Cycles" *Journal of Political Economy* 105, 211-248.

[17] Kocherlakota, N.R. (2000) "Creating Business Cycles Through Credit Constraints" *Federal Reserve Bank of Minneapolis Quarterly Review* 24, 2-10.

[18] Krusell, P. and A.A. Smith Jr. (1997) "Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns," *Macroeconomic Dynamics* 1, 387-422.

[19] Modigliani, F. and M. Miller (1958). "The Cost of Capital, Corporation Finance, and the Theory of Investment," *American Economic Review* 48, 261–297.

[20] Prescott, E. C. (1986) "Theory Ahead of Business Cycle Measurement," *Federal Reserve Bank of Minneapolis Quarterly Review* 10, 9-22.

[21] Restuccia D. and R. Rogerson (2008) "Policy distortions and aggregate productivity with heterogeneous establishments" *Review of Economic Dynamics* 11, 707-720.

[22] Rogerson, R. (1988) "Indivisible Labor, Lotteries and Equilibrium," *Journal of Monetary Economics* 21, 3-16.

[23] Veracierto, M. L. (2002) "Firm-Level Irreversible Investment and Equilibrium Business Cycles," *American Economic Review* 92, 181-197.

FIGURE 1. Distribution of firms with lowest productivity

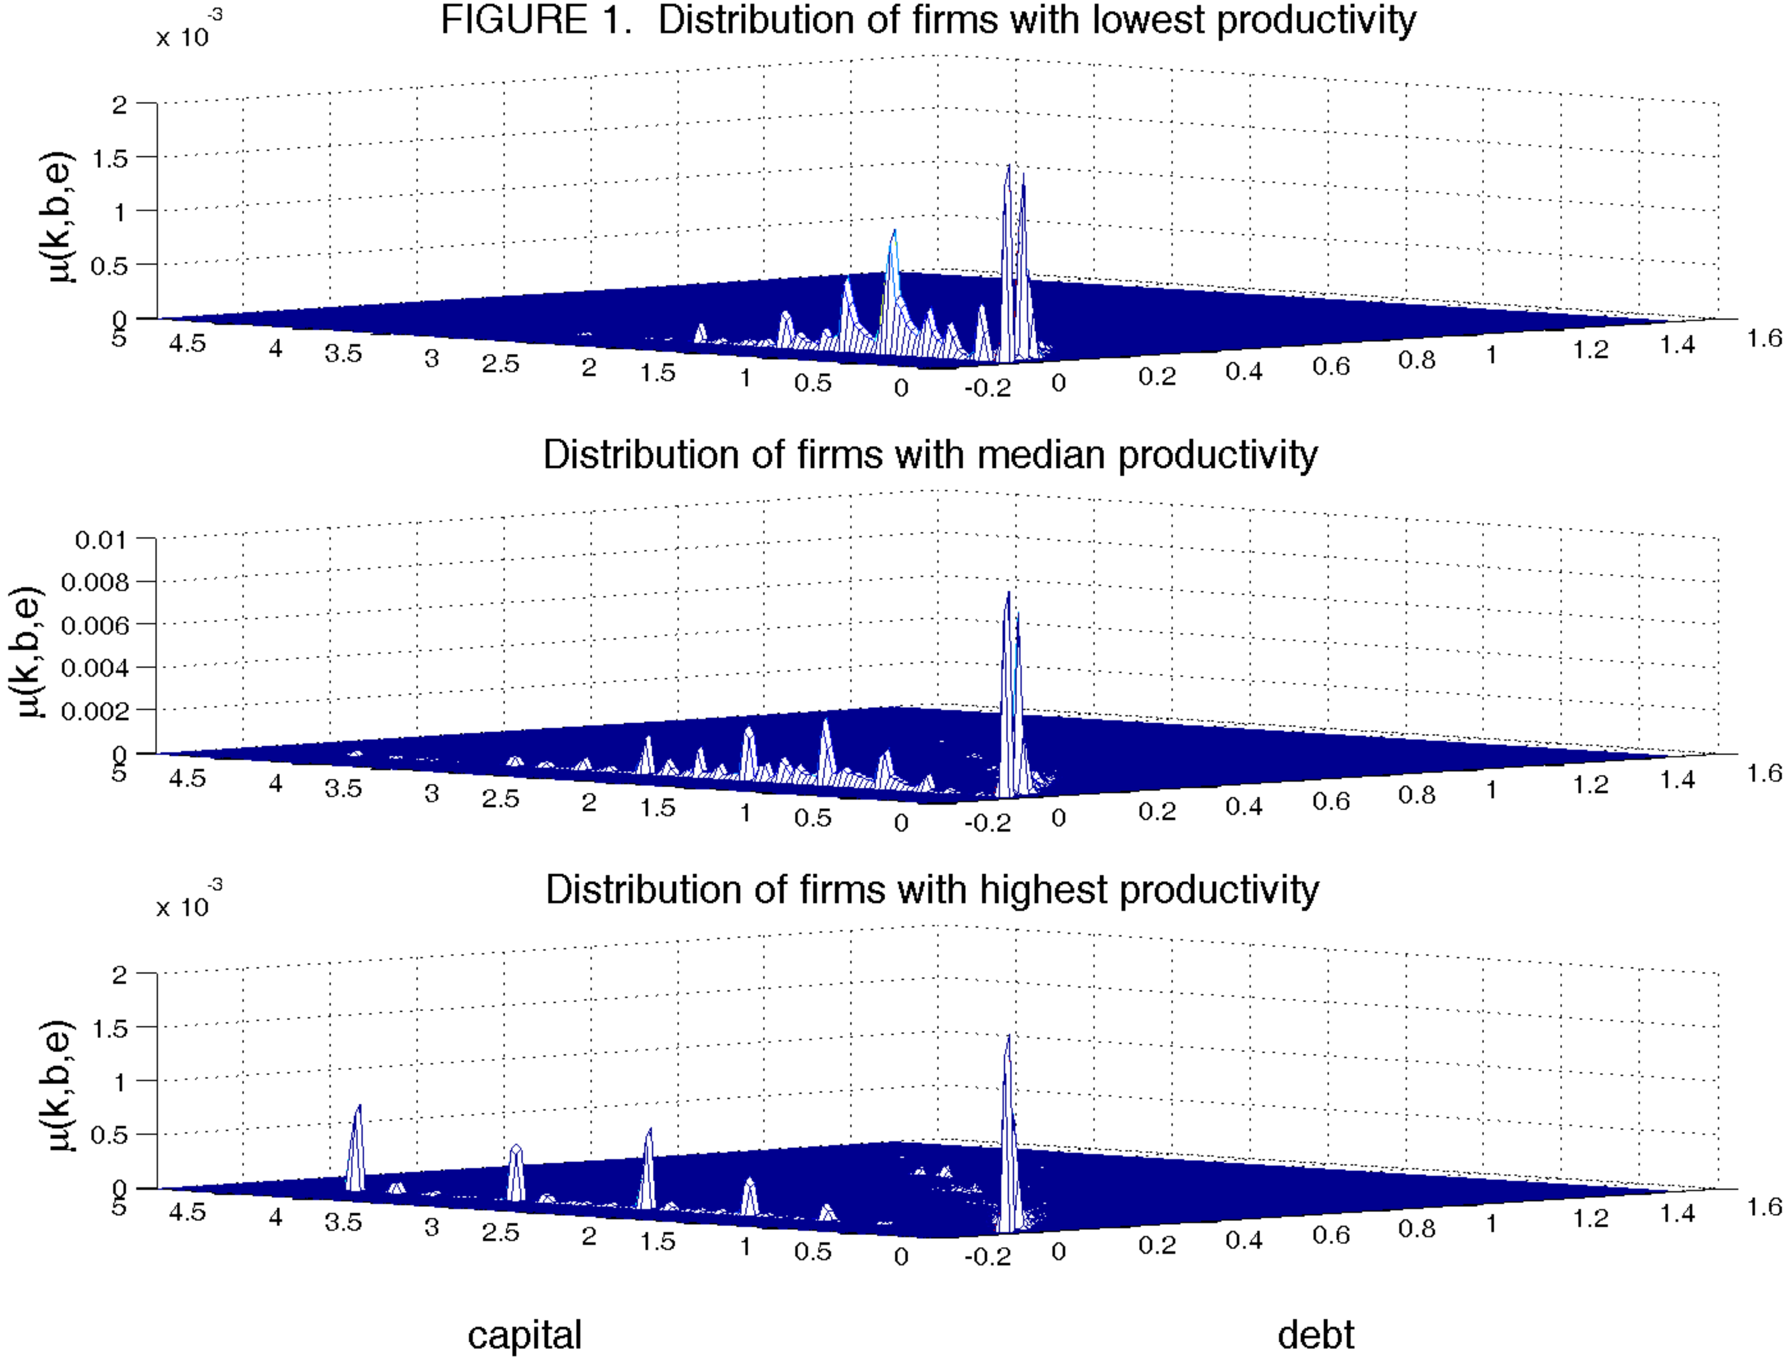

FIGURE 2. Distribution of firms with lowest productivity

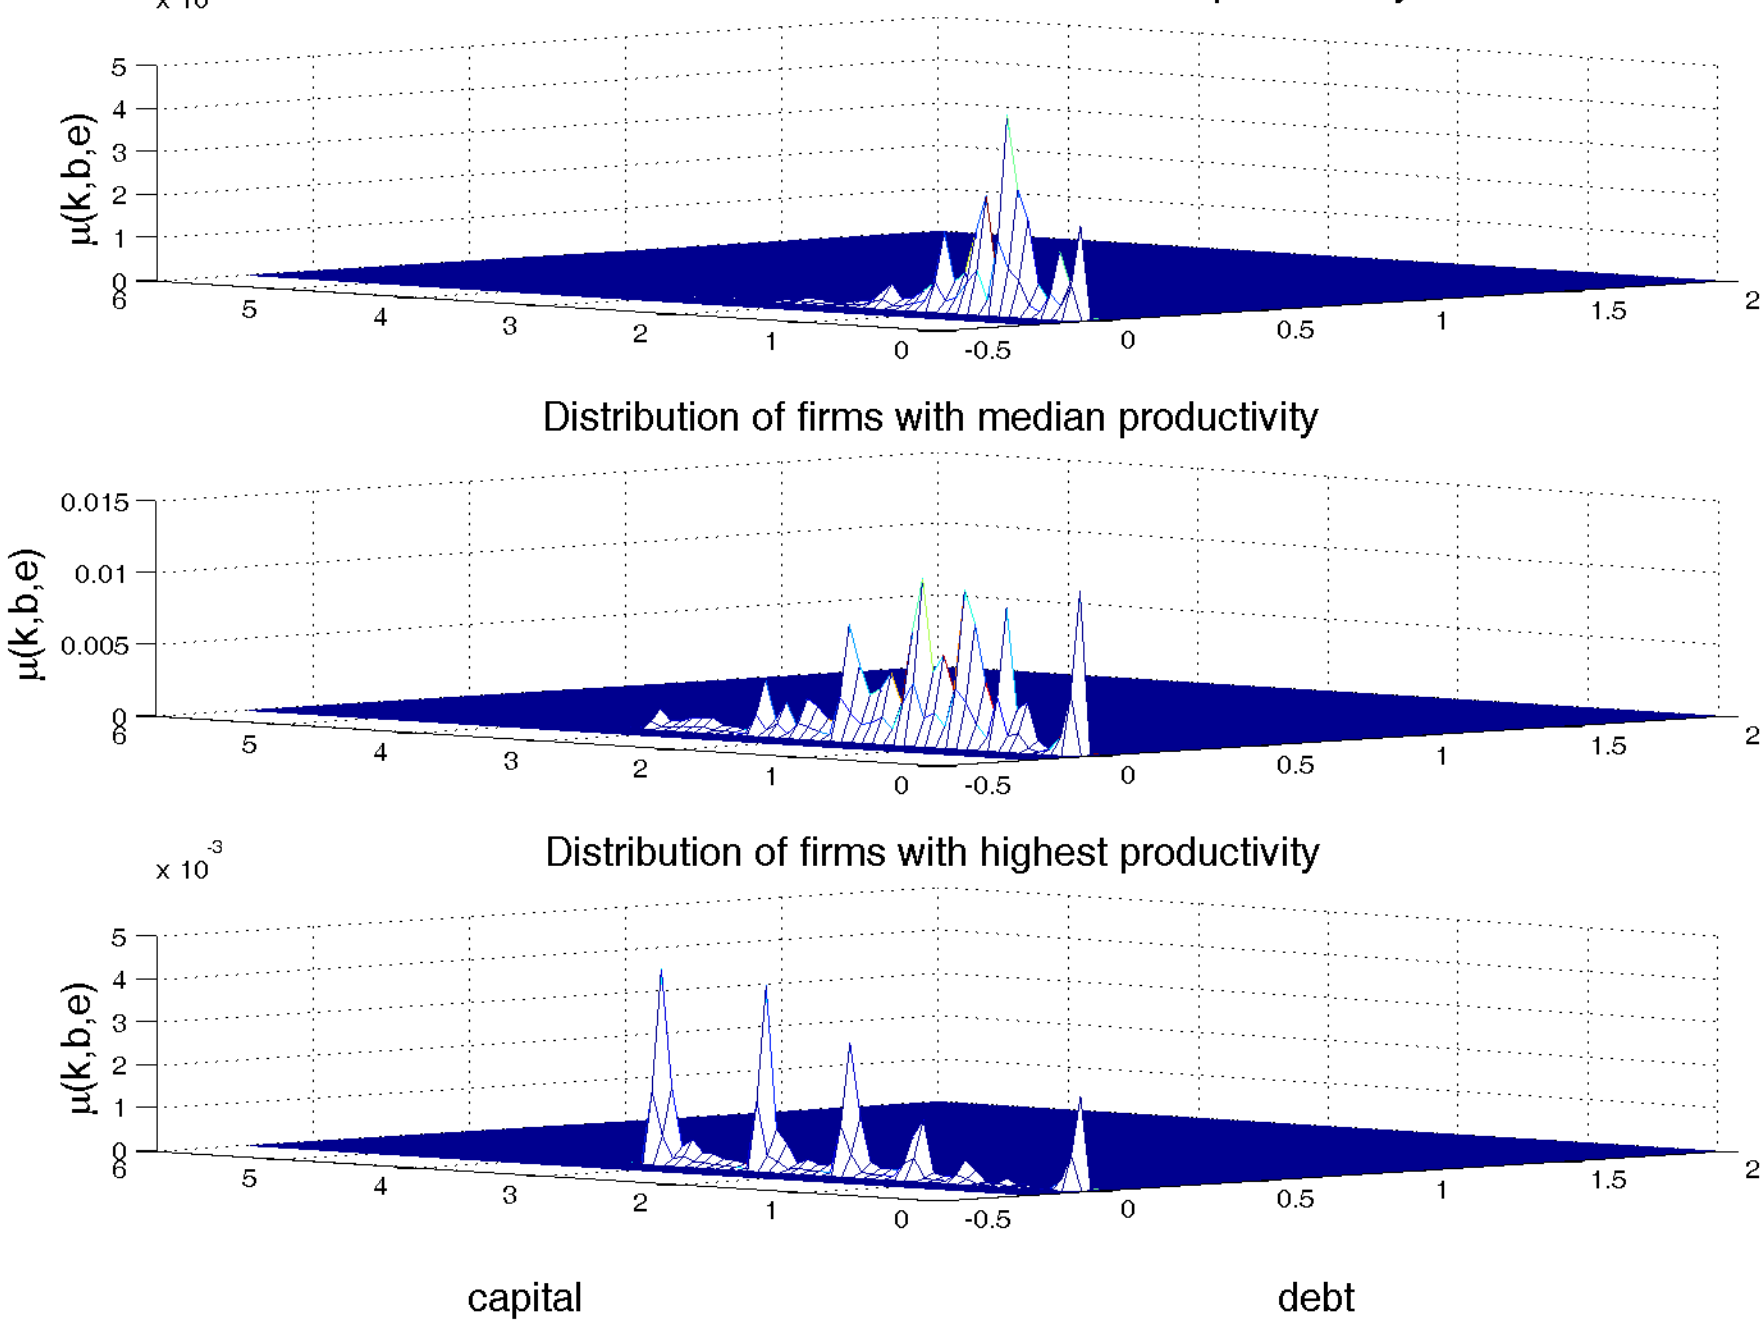

FIGURE 3. Capital choices among unconstrained firms with lowest productivity

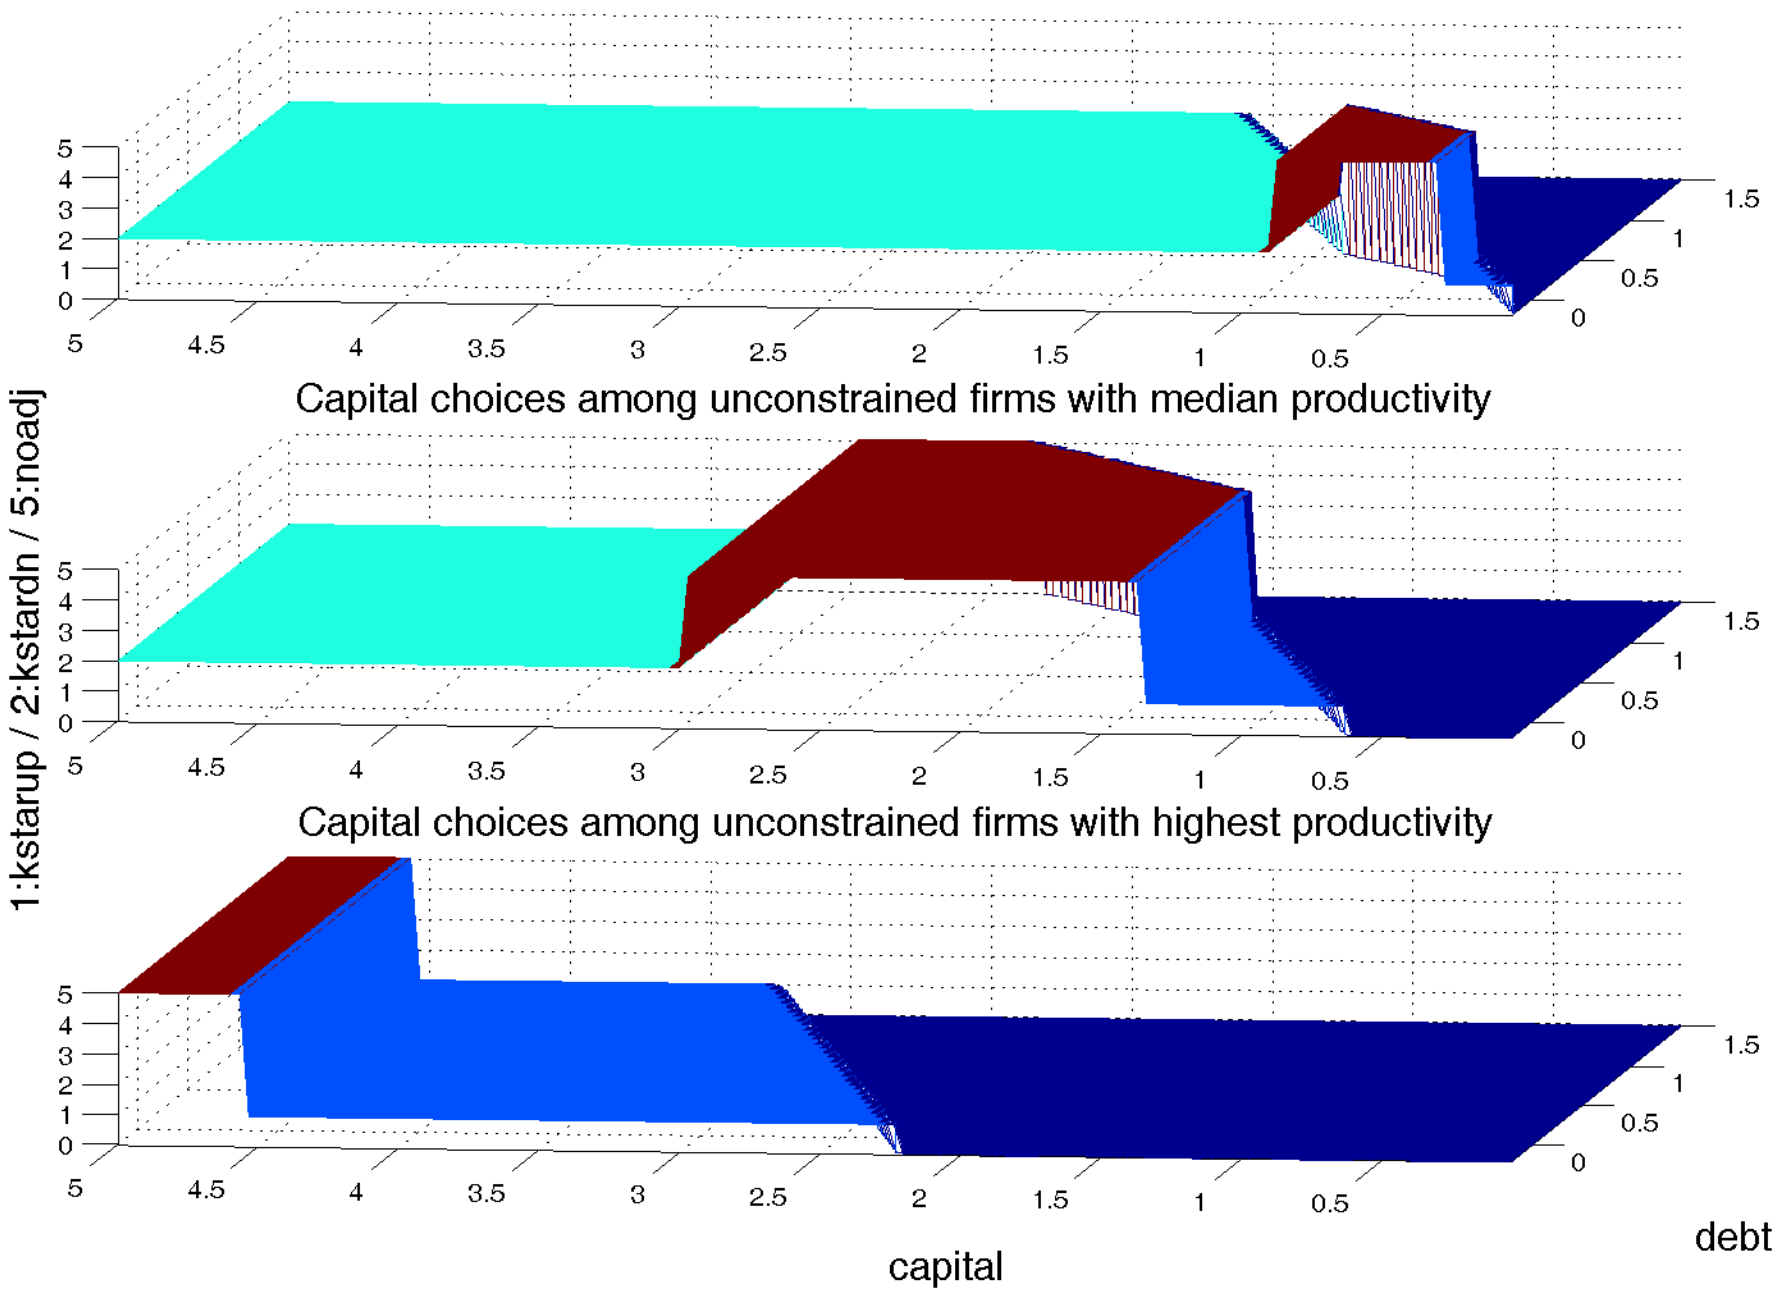

FIGURE 4. Capital choices among constrained firms with lowest productivity

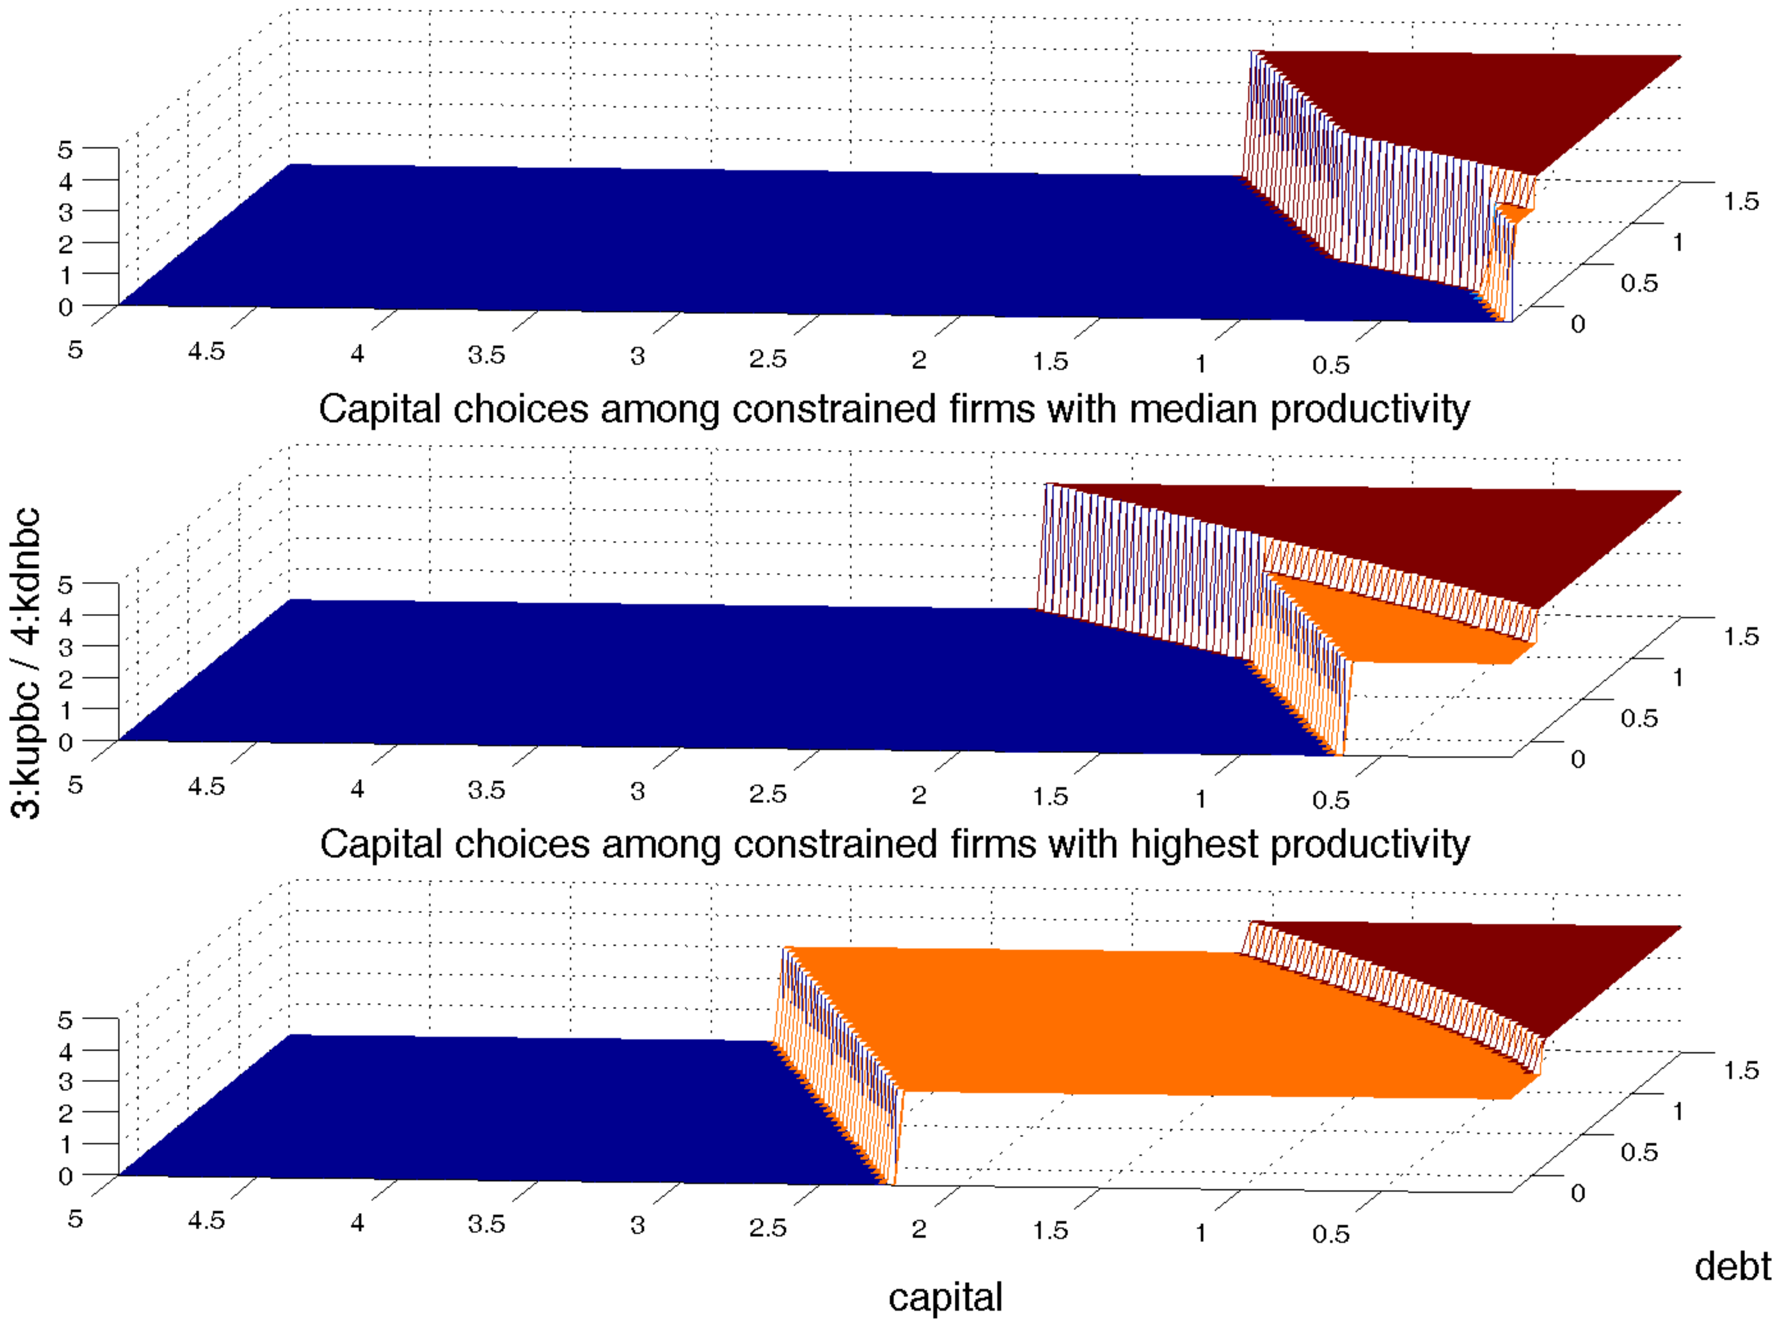

FIGURE 5. Output among firms with lowest productivity

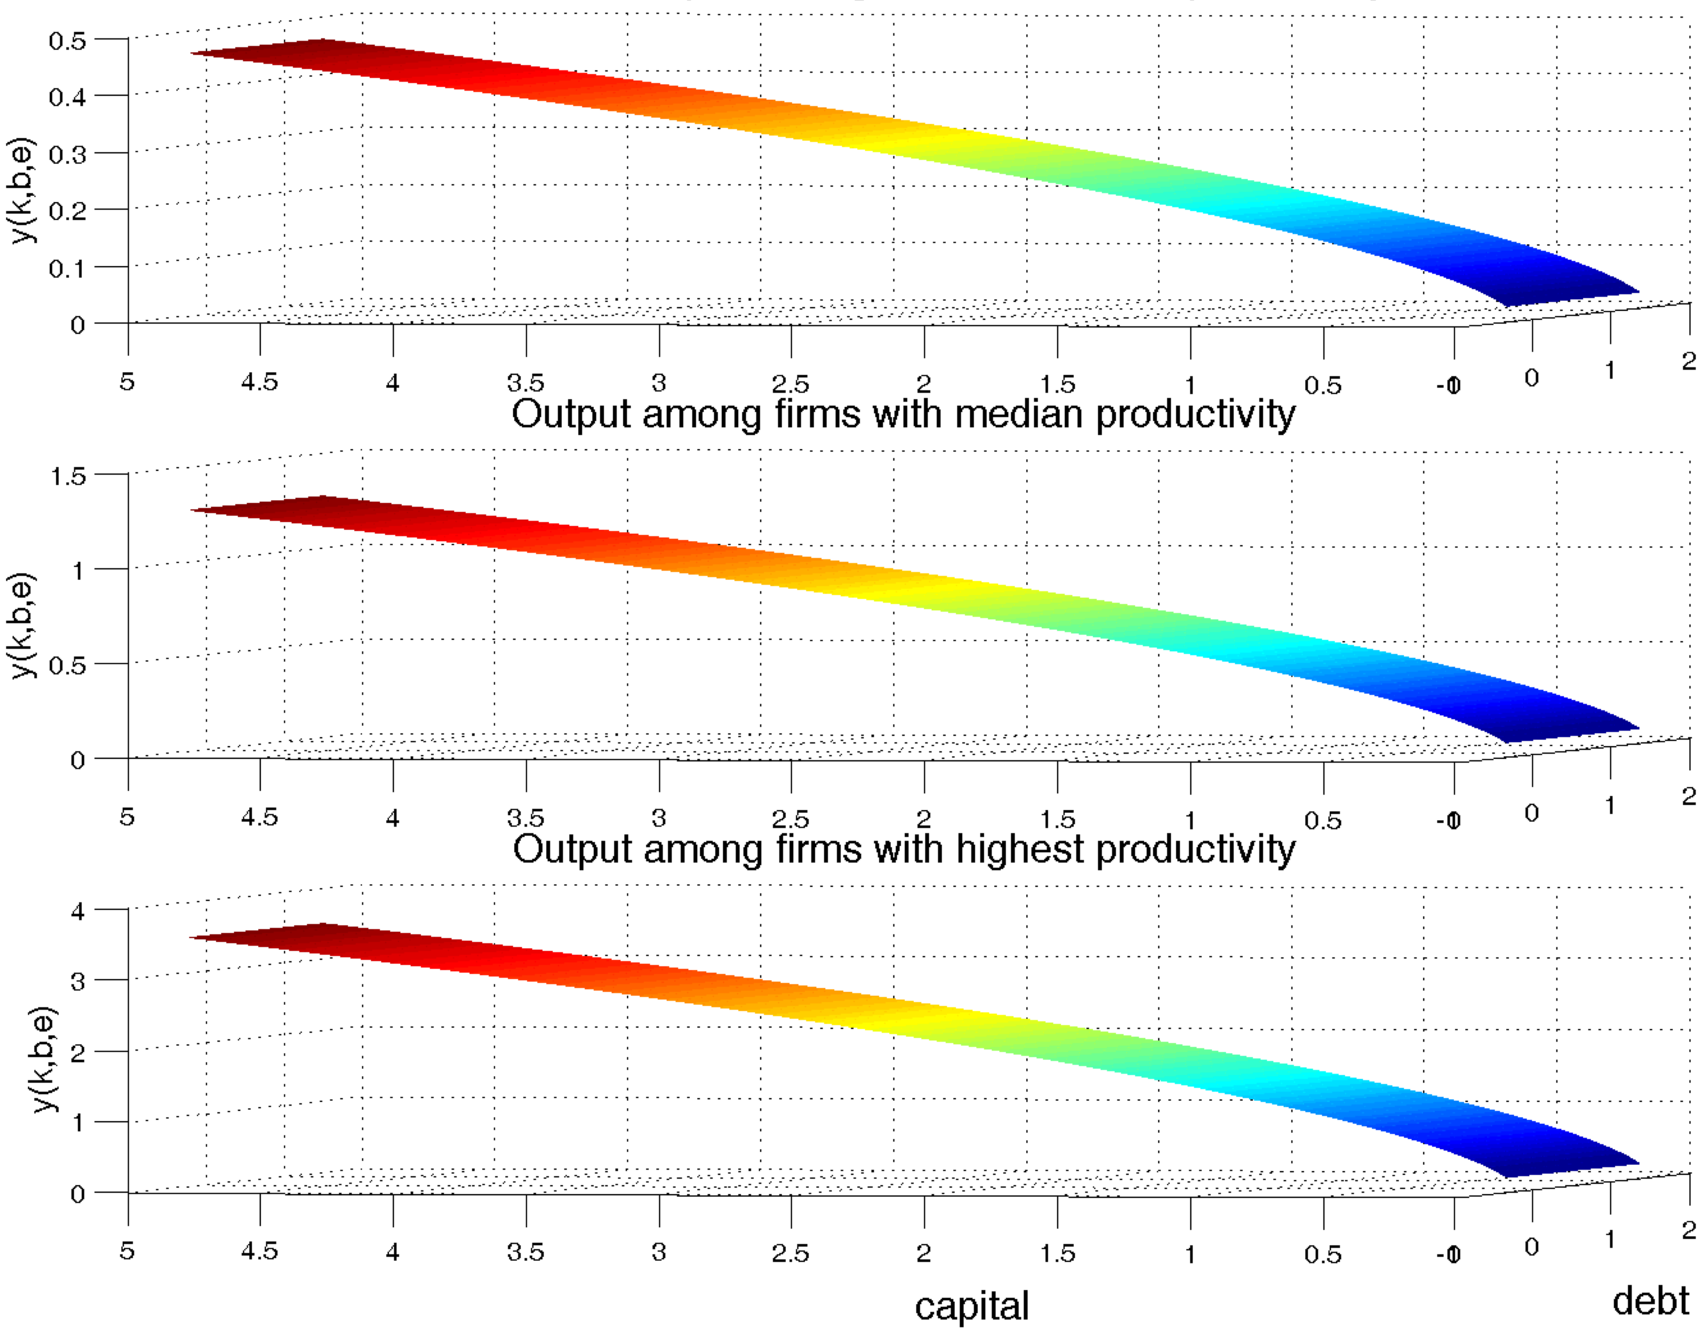

FIGURE 6. Response to Financial Crisis

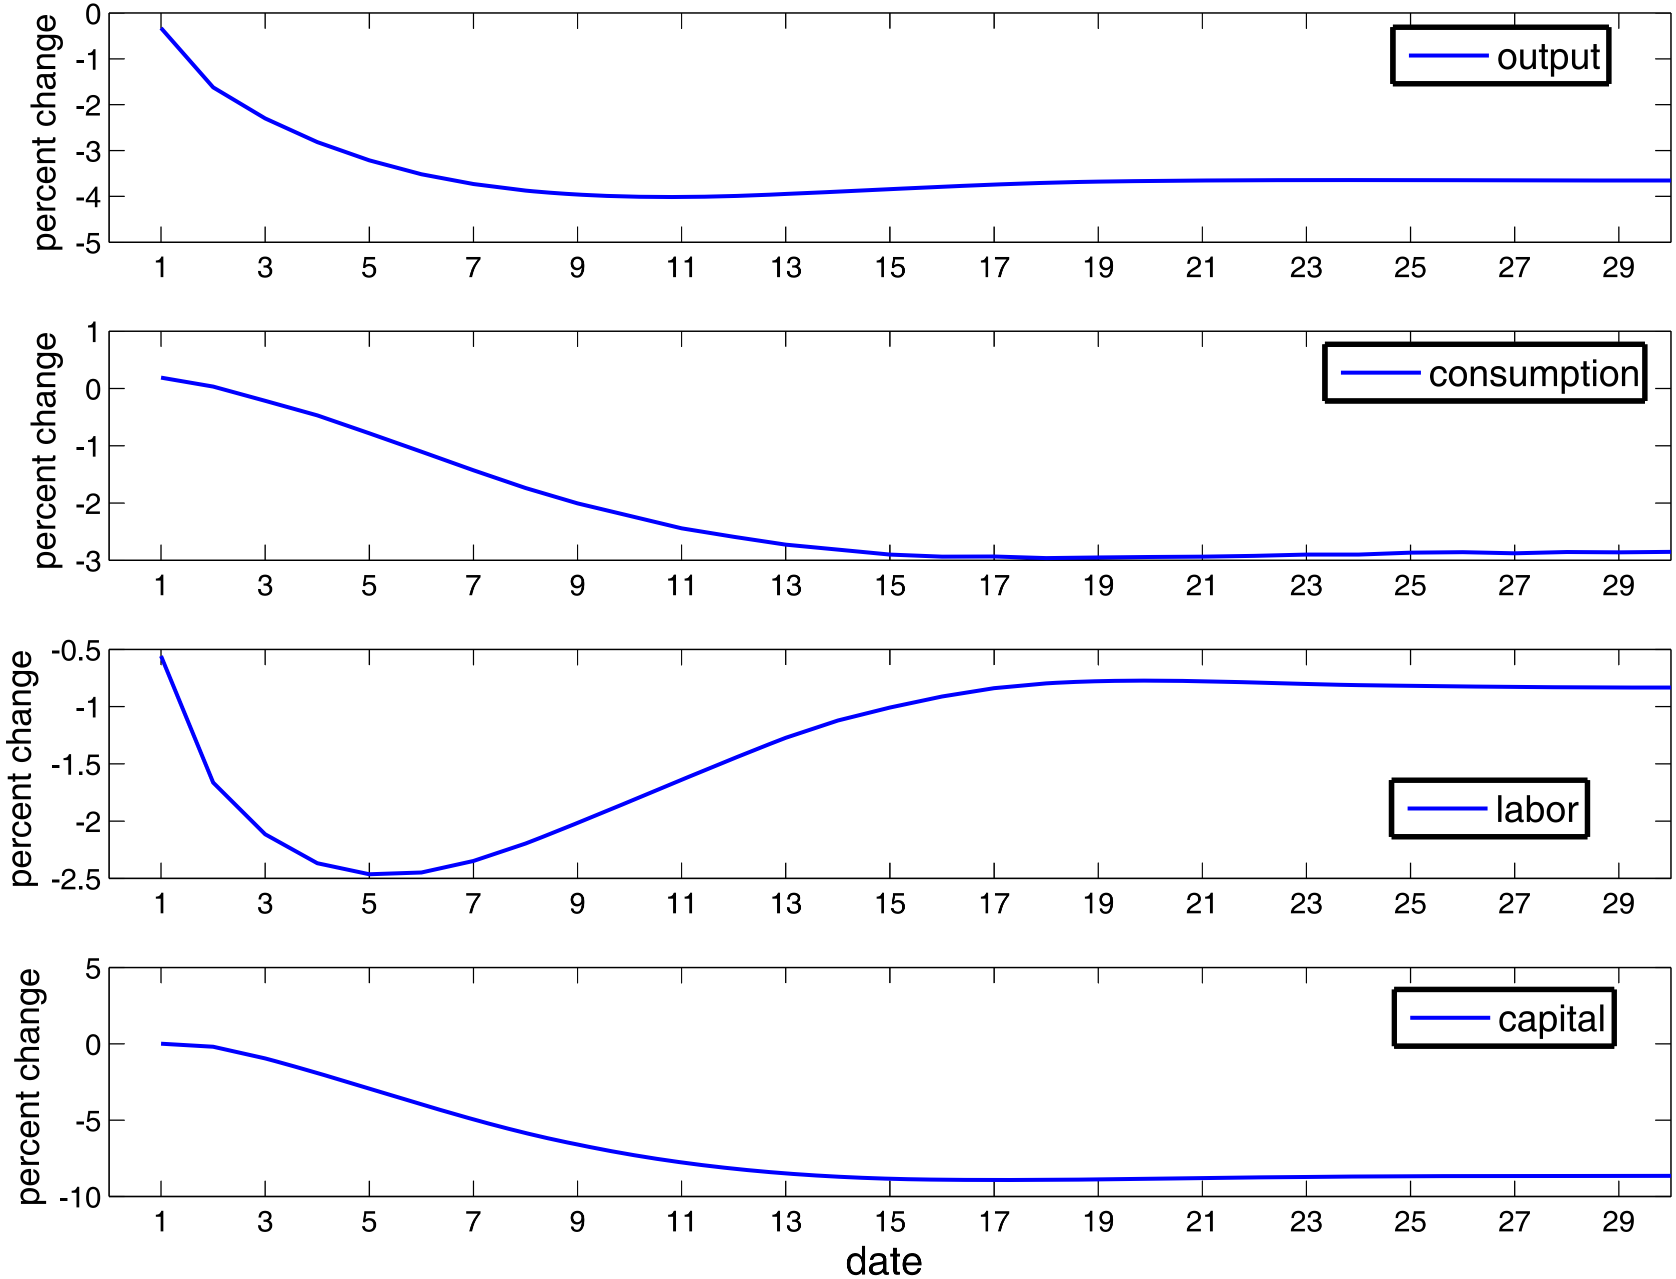

FIGURE 7. Financial Crisis with a Technology Shock

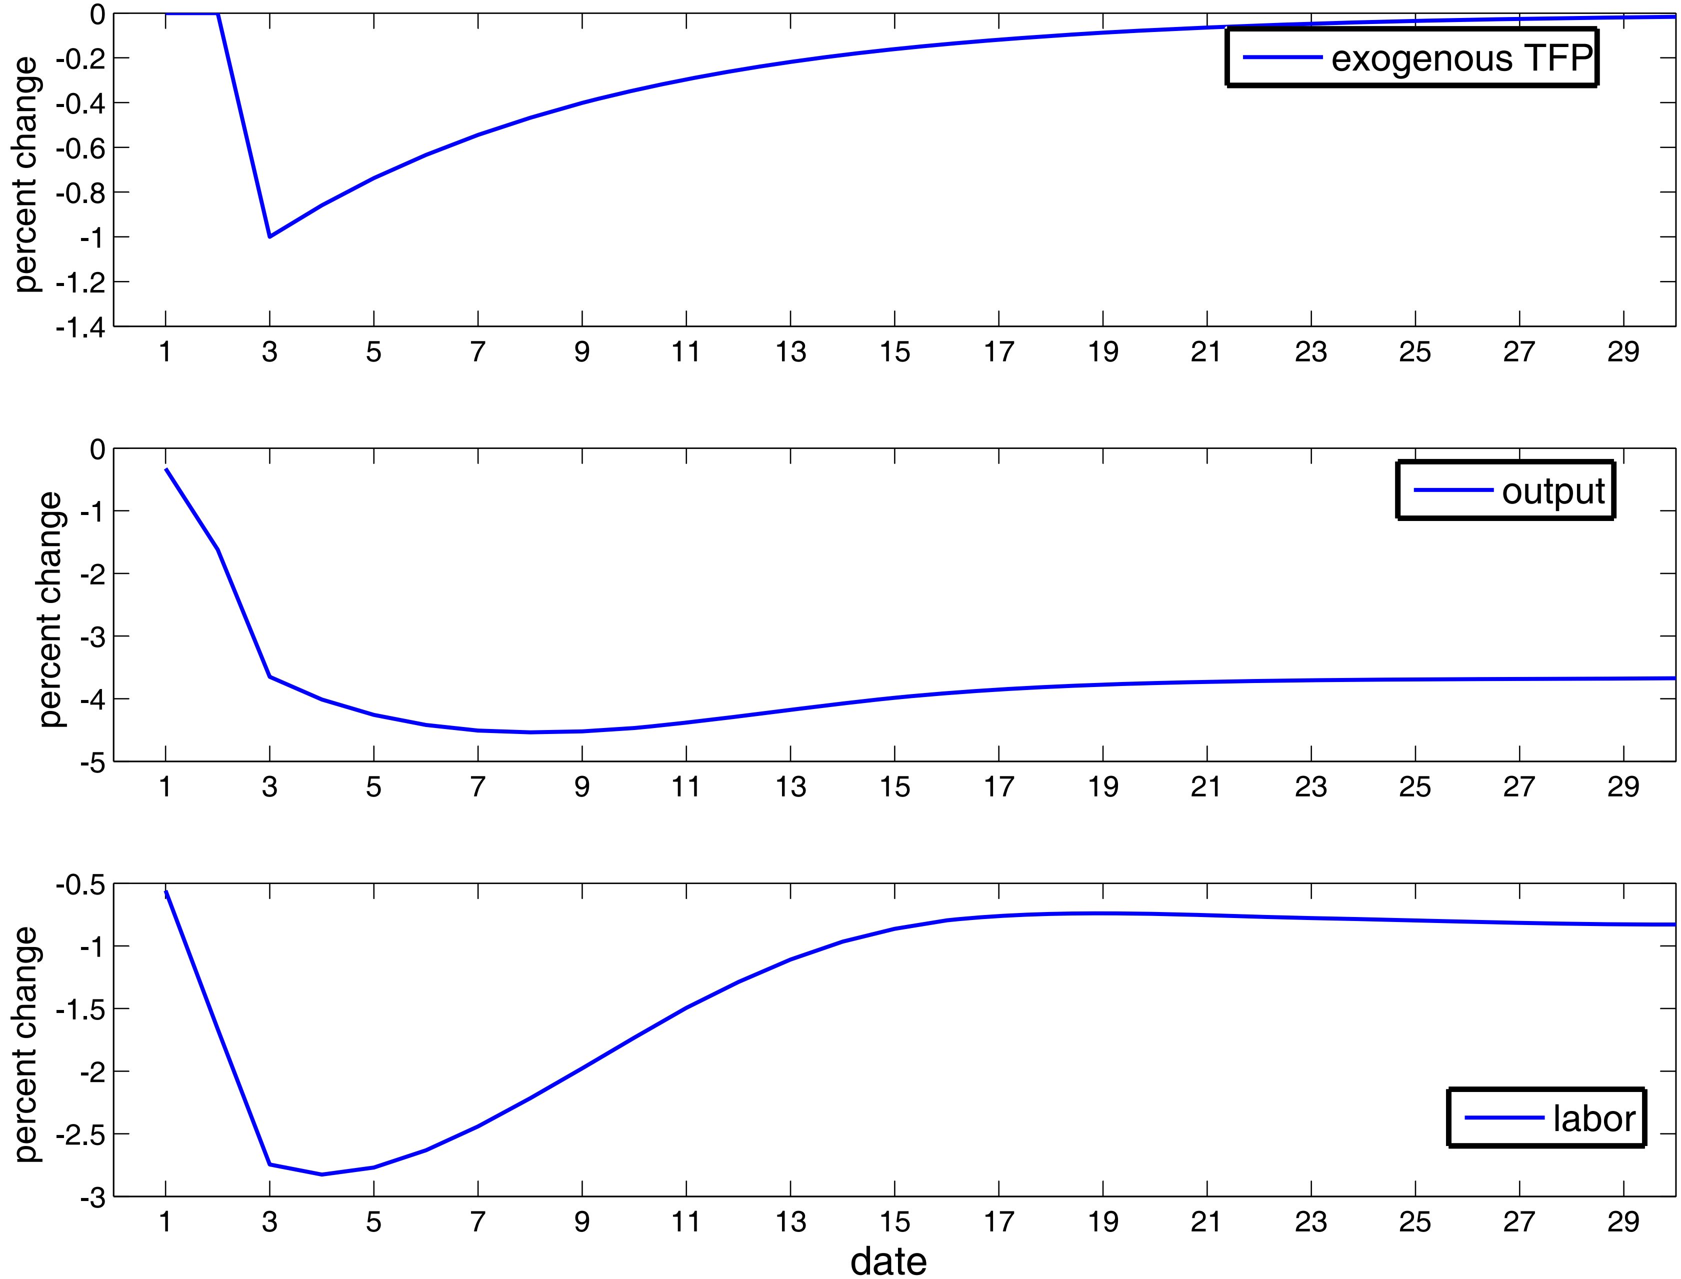

TABLE 1. Business Cycles in the Full Economy

$x =$	Y	C	I	N	K	r
mean(x)	0.535	0.463	0.073	0.323	1.101	0.042
σ_x/σ_Y	(1.695)	0.547	4.254	0.511	0.507	0.516
$corr(x, Y)$	1.000	0.950	0.963	0.943	0.094	0.642

TABLE 2. Business Cycles Without Financial Frictions

$x =$	Y	C	I	N	K	r
mean(x)	0.581	0.492	0.088	0.329	1.329	0.042
σ_x/σ_Y	(1.737)	0.516	4.046	0.546	0.476	0.480
$corr(x, Y)$	1.000	0.940	0.967	0.947	0.073	0.663

TABLE 3. Business Cycles Without Financial or Real Frictions

$x =$	Y	C	I	N	K	r
mean(x)	0.593	0.507	0.086	0.326	1.438	0.042
σ_x/σ_Y	(1.760)	0.492	4.397	0.571	0.458	0.459
$corr(x, Y)$	1.000	0.930	0.965	0.951	0.058	0.677