Financial Frictions, Investment, and Tobin's q

Dan Cao, Guido Lorenzoni, Karl Walentin

PII: S0304-3932(18)30457-4
DOI: https://doi.org/10.1016/j.jmoneco.2018.08.002
Reference: MONEC 3033

To appear in: *Journal of Monetary Economics*

Received date: 13 February 2017
Revised date: 7 August 2018
Accepted date: 7 August 2018

Please cite this article as: Dan Cao, Guido Lorenzoni, Karl Walentin, Financial Frictions, Investment, and Tobin's q, *Journal of Monetary Economics* (2018), doi: https://doi.org/10.1016/j.jmoneco.2018.08.002

Highlights

- A model of firm-level investment with fully state-contingent financial constraints.
- A characterization of Tobin's q in presence of financial frictions.
- Two forces are driving q: the value of invested capital and future quasi-rents.
- This weakens the correlation between investment and q.
- Implications for investment regressions depend crucially on the nature of shocks.

Financial Frictions, Investment, and Tobin's q

Dan Cao
Georgetown University

Guido Lorenzoni *
Northwestern University and NBER

Karl Walentin
Sveriges Riksbank

June 2018

Abstract

A model of investment with financial constraints is used to study the relation between investment and Tobin's q. A firm is financed by both inside and outside investors. When insiders' wealth is scarce, the firm's value includes a quasi-rent on invested capital. Therefore, two forces drive q: the value of invested capital and future quasi-rents. Relative to a frictionless benchmark, this weakens the relationship between investment and q, generating more realistic correlations between investment, q, and cash flow. The quantitative implications of the model for investment regressions depend crucially on the nature of the shocks hitting the firm.

Keywords: Financial constraints, optimal financial contracts, investment, Tobin's *q*, limited enforcement.

JEL codes: E22, E30, E44, G30.

*Corresponding Author: guido.lorenzoni@northwestern.edu; Address: KGH 3437, 2211 Campus Drive Evanston, Illinois 60208; Telephone: + 1-847-491-8217

1 Introduction

Dynamic models of the firm imply that investment decisions and the value of the firm should both respond to expectations about future profitability of capital. In models with constant returns to scale and convex adjustment costs these relations are especially clean, as investment and the firm's value respond exactly in the same way to new information about future profitability. This is the main prediction of Tobin's q theory, which implies that current investment moves one-for-one with q, the ratio of the firm's financial market value to its capital stock. This prediction, however, is typically rejected in the data, where investment appears to correlate more strongly with current cash flow than with q.

In this paper, we investigate the relation between investment, q, and cash flow in a model with financial frictions. The presence of financial frictions introduces quasi-rents in the market valuation of the firm. These quasi-rents break the one-to-one link between investment and q. We study how the presence of these quasi-rents affects the statistical correlations between investment, q, and cash flow, and ask whether a model with financial frictions can match the correlations in the data.

Our main conclusion is that the presence of financial frictions can bring the model closer to the data, but that the model's implications depend crucially on the shock structure. In a model with financial frictions it is still true that investment and q respond to future profitability, but the two variables now respond differently to information at different horizons. Investment is particularly sensitive to current profitability, which determines current internal financing, and to near-term financial profitability, which determines collateral values. On the other hand, q is relatively more sensitive to profitability farther in the future, which will determine future growth and thus the size of future quasi-rents. Therefore, to break the link between investment and q, requires the presence of both short-lived shocks—which tend to move investment more and have relatively smaller effects on q—and long-lived shocks—which do the opposite.

These points are developed in the context of a stochastic model of investment subject to limited enforcement, with fully state-contingent claims. The ability of borrowers to issue state-contingent claims is limited by the fact that, *ex post*, they can renege on their promises and default. The consequence of default is the loss of a fraction of invested assets. We show that this environment is equivalent to an environment with state-contingent collateral constraints, so the model is essentially a stochastic version of Kiyotaki and Moore (1997) with adjustment costs and state-contingent claims.[1] The model

[1] Related recent stochastic models that combine state-contingent claims with some form of collateral constraint include Lorenzoni (2008), He and Krishnamurthy (2013), Rampini and Viswanathan (2013), Cao and Nie (2017), and Di Tella (2017).

leads to a wedge between average q—which correspond to the q measured from financial market values—and marginal q—which captures the marginal incentive to invest and is related one-to-one to investment.[2] Two versions of the model are analyzed, looking at their implications for investment regressions in which the investment rate is regressed on average q and cash flow.

The first version of the model features no adjustment costs and, under some simplifying assumptions, it can be linearized and studied analytically. When a single persistent shock is introduced, the model has indeterminate predictions regarding investment regression coefficients. This simply follows because in this case q and cash flow are perfectly collinear. With two shocks—a temporary shock and a persistent shock—the one-to-one relation between q and investment breaks down because investment is driven by productivity in periods t and $t+1$ while q responds to all future values of productivity. Finally, a "news shock' is introduced, that allows agents to observe the realization of future productivity shocks J periods in advance. Increasing the length of the horizon J reduces the coefficient on q and increases the coefficient on cash flow in investment regressions. This is due again to the differential responses of investment and q to information on productivity at different horizons.

The model with no adjustment costs, while analytically tractable, is quantitatively unappealing, as it tends to produce too much short-run volatility and too little persistence in investment. Therefore, for a more quantitative evaluation of the model we introduce adjustment costs. The model is calibrated to data moments from Compustat and analyze its implications both in terms of impulse responses and in terms of investment regressions. The baseline calibration is based on the two shocks structure, with temporary and persistent shocks. In this calibration q responds relatively more strongly to the persistent shock while investment responds relatively more strongly to the transitory shock, in line with the intuition from the no-adjustment-cost case. This leads to investment regressions with a smaller coefficient on q and a larger coefficient on cash flow, relative to a model with no financial frictions, thus bringing us closer to empirical coefficients. However, the q coefficient is still larger than in the data and the cash flow coefficient is smaller than in the data. When adding the possibility of news shocks, the disconnect between q and investment increases, leading to further reductions in the q coefficient and increases in the cash flow coefficient.

Fazzari et al. (1988) started a large empirical literature that explores the relation between investment and q using firm-level data. The typical finding in this literature is a

[2] The terminology goes back to Hayashi (1982), who shows that the two are equivalent in a canonical model with convex adjustment costs.

small coefficient on q and a positive and significant coefficient on cash flow.[3] Fazzari et al. (1988), Gilchrist and Himmelberg (1995) and most of the subsequent literature interpret these findings as a symptom of financial frictions at work. More recent work by Gomes (2001) and Cooper and Ejarque (2003) questions this interpretation. The approach taken in these two papers is to look at the statistical implications of simulated data generated by a model to understand the empirical correlations between investment, q and cash flow.[4] In their simulated economies with financial frictions q still explains most of the variability in investment, and cash flow does not provide additional explanatory power. In this paper, we take a similar approach but reach different conclusions. This is due to two main differences. First, Gomes (2001) and Cooper and Ejarque (2003) model financial frictions by introducing a transaction cost which is a function of the flow of outside finance issued each period, while we introduce a contractual imperfection that imposes an upper bound on the stock of outside liabilities as a fraction of total assets. Our approach adds a state variable to the problem, namely the stock of existing liabilities of the firm as a fraction of assets, thus generating slower dynamics in the gap between internal funds and the desired level of investment. Second, we explore a variety of shock structures, which, as argued below, play an important role in our results.

A related strand of recent literature has focused on violations of q theory coming from decreasing returns or market power, leaving aside financial frictions.[5] Our effort is complementary to this literature, since both financial frictions and decreasing returns determine the presence of future rents embedded in the value of the firm. Also in that literature the shock structure plays an important role in the results. For example, Eberly et al. (2008) show that it is easier to obtain realistic implications for investment regressions by assuming a Markov process in which the distribution from which persistent productivity shocks are drawn switches occasionally between two regimes. Abel and Eberly (2011) also show that in models with decreasing returns it is possible to obtain interesting dynamics in q with no adjustment costs, similarly to the results presented in Section 3 for a model with constant returns to scale and financial constraints.

The simplest shock that breaks the link between q and investment in models with financial constraints is a purely temporary shock to cash flow, which does not affect capital's future productivity. Absent financial frictions this shock should have no effect on current investment. This idea is the basis of a strand of empirical literature that tests for financial constraints by identifying some source of purely temporary shocks to cash flow.

[3] See Hubbard (1998) for a survey.

[4] An approach that goes back to Sargent (1980).

[5] See Schiantarelli and Georgoutsos (1990), Alti (2003), Moyen (2004), Eberly et al. (2008), Abel and Eberly (2011), Abel and Eberly (2012).

This is the approach taken by Blanchard et al. (1994) and Rauh (2006), which provide reliable evidence of the presence of financial constraints. Our paper builds on a similar intuition, by showing that in general shocks affecting profitability at different horizons have differential effects on q and investment and asks whether, given a realistic mix of shocks, a model with financial frictions can produce the unconditional correlations observed in the data.

This paper uses the simplest possible model with the features needed: an occasionally binding financial constraint; a dynamic, stochastic structure; adjustment costs that can produce realistic investment dynamics. There is a growing literature that builds richer models that are geared more directly to estimation. In particular, Hennessy and Whited (2007) build a rich structural model of firms' investment with financial frictions, which is estimated by simulated method of moments. They find that the financial constraint plays an important role in explaining observed firms' behavior. In their model, due to the complexity of the estimation task, the financial friction is introduced in a reduced form manner, by assuming transaction costs associated to the issuance of new equity or debt, as in Gomes (2001) or Cooper and Ejarque (2003).[6] This paper takes a complementary route, as it features a more stylized model, but financial constraints coming from an explicitly modeled contractual imperfection.

A growing number of papers uses recursive methods to characterize optimal dynamic financial contracts in environments with different forms of contractual frictions (Atkeson and Cole (2005), Clementi and Hopenhayn (2006), DeMarzo and Sannikov (2006), DeMarzo et al. (2012)). The limited enforcement friction in this paper makes it closer to the models in Albuquerque and Hopenhayn (2004) and Cooley et al. (2004). Within this literature Biais et al. (2007) look more closely at the implications of the theory for asset pricing. In particular, they find a set of securities that implements the optimal contract and then study the stochastic behavior of the prices of these securities. Here, our objective is to examine the model's implication for q theory, therefore we simply focus on the total value of the firm, which includes the value of all the claims held by insiders and outsiders.

Section 2 presents the model. Section 3 contains the case of no adjustment costs. Section 4 contains the model with adjustment costs.

[6] The difference in results, relative to these papers, appears due to the fact that Hennessy and Whited (2007) also match the behavior of a number of financial variables.

2 The model

Consider an infinite horizon economy, in discrete time, populated by a continuum of entrepreneurs who invest in physical capital and raise funds from risk neutral investors.

The entrepreneurs' technology is linear: K_{it} units of capital, installed at time $t-1$ by entrepreneur i, yield profits $A_{it}K_{it}$ at time t. We can think of the linear profit function $A_{it}K_{it}$ as coming from a constant returns to scale production function in capital and other variable inputs which can be costlessly adjusted. Therefore, changes in A_{it} capture both changes in technology and changes in input and output prices. For brevity, we just call A_{it} "productivity". Productivity is a function of the state s_{it}, $A(s_{it})$, where s_{it} is a Markov process with a finite state space $\mathbf{S}$ and transition probability $\pi(s_{it}|s_{it-1})$. There are no aggregate shocks, so the cross sectional distribution of s_{it} across entrepreneurs is constant.

Investment is subject to convex adjustment costs. The cost of changing the installed capital stock from K_{it} to K_{it+1} is $G(K_{it+1}, K_{it}; s_{it})$ units of consumption goods at date t. The function G includes both the cost of purchasing capital goods and the installation cost. G is increasing and convex in its first argument, decreasing in the second argument, and displays constant returns to scale. For numerical results, we use the quadratic functional form

$$G(K_{it+1}, K_{it}; s_{it}) = \phi(s_{it})\left(K_{it+1} - (1-\delta(s_{it}))K_{it}\right) + \frac{\xi}{2}\frac{(K_{it+1}-K_{it})^2}{K_{it}}, \quad (1)$$

in which the state s_{it} can affect both the depreciation rate $\delta(s_{it})$ and the price of capital goods $\phi(s_{it})$.

All agents in the model are risk neutral. The entrepreneurs' discount factor is β and the investors' discount factor is $\hat{\beta}$ and entrepreneurs are more impatient: $\beta < \hat{\beta}$. Investors have a large enough endowment of the consumption good each period so that in equilibrium the interest rate is $1 + r_t = 1/\hat{\beta}$. Each period an entrepreneur retires with probability γ and is replaced by a new entrepreneur with an endowment of 1 unit of capital. When an entrepreneur retires, productivity A_{it} is zero from next period on.The retirement shock is embedded in the process s_{it} by assuming that there is an absorbing state s^r with $A(s^r) = 0$ and the probability of transitioning to s^r from any other state is γ.

Each period, entrepreneur i issues one-period state contingent liabilities, subject to limited enforcement. The entrepreneur controls the firm's capital K_{it} and, at the beginning of each period, can default on his liabilities and divert a fraction $1-\theta$ of the firm's capital. If he does so, he re-enters the financial market as a new entrepreneur, with capital $(1-\theta)K_{it}$ and no liabilities. That is, the punishment for a defaulting entrepreneur is the loss of a fraction θ of the firm's assets.

2.1 Optimal investment

Let us formulate the optimization problem of the individual entrepreneur in recursive form, dropping the subscripts i and t. Let $V(K,B,s)$ be the expected utility of an entrepreneur in state s, who enters the period with capital stock K and current liabilities B. For now, simply assume that the problem's parameters are such that the entrepreneur's optimization problem is well defined. In the following sections, we provide conditions that ensure that this is the case.[7] The function V satisfies the Bellman equation

$$V(K,B,s) = \max_{C\geq 0, K'\geq 0, \{B'(s')\}} C + \beta \mathbb{E}\left[V\left(K', B'\left(s'\right), s'\right) |s\right], \tag{2}$$

subject to

$$C + G\left(K', K; s\right) \leq A(s)K - B + \hat{\beta}\mathbb{E}\left[B'\left(s'\right)|s\right], \tag{3}$$

$$V\left(K', B'\left(s'\right), s'\right) \geq V\left(\left(1-\theta\right)K', 0, s'\right), \forall s', \tag{4}$$

where C is current consumption, K' is next period's capital stock , and $B'(s')$ are next period's liabilities contingent on s'. Constraint (3) is the budget constraint and $\hat{\beta}\mathbb{E}\left[B'(s')|s\right]$ are the funds raised by selling the state contingent claims $\{B'(s')\}$ to the investors. Constraint (4) is the enforcement constraint that requires the continuation value under repayment to be greater than or equal to the continuation value under default.

The assumption of constant returns to scale implies that the value function takes the form $V(K,B,s) = v(b,s)K$ for some function v, where $b = B/K$ is the ratio of current liabilities to the capital stock. The Bellman equation then becomes, using the notation $c = C/K$ and $k' = K'/K$,

$$v(b,s) = \max_{\substack{c\geq 0, k'\geq 0,\\ \{b'(s')\}}} c + \beta\mathbb{E}\left[v\left(b'\left(s'\right), s'\right)|s\right]k', \tag{5}$$

subject to

$$c + G\left(k', 1; s\right) \leq A(s) - b + \hat{\beta}\mathbb{E}\left[b'(s')|s\right]k', \tag{6}$$

$$v\left(b'\left(s'\right), s'\right) \geq \left(1-\theta\right)v\left(0, s'\right), \forall s'. \tag{7}$$

It is easy to show that v is strictly decreasing in b. We can then find state-contingent

[7] In the online appendix we provide a general existence result.

borrowing limits $\overline{b}(s')$ such that the enforcement constraint is equivalent to

$$b'(s') \leq \overline{b}(s'), \forall s'. \tag{8}$$

So the enforcement constraint is equivalent to a state contingent upper bound on the ratio of the firm's liabilities to capital. Relative to existing models with collateral constraints, two distinguishing features of this model are the presence of state-contingent claims and the fact that state-contingent bounds are derived endogenously from limited enforcement.[8]

2.2 Average and marginal q

To characterize the solution to the entrepreneur's problem let us start from the first order condition for k':

$$\lambda G_1(k', 1; s) = \lambda \hat{\beta} \mathbb{E}[b'|s] + \beta \mathbb{E}[v'|s], \tag{9}$$

where λ is the Lagrange multiplier on the budget constraint (6), or the marginal value of wealth for the entrepreneur. The expressions $\mathbb{E}[b'|s]$ and $\mathbb{E}[v'|s]$ are shorthand for $\mathbb{E}[b'(s')|s]$ and $\mathbb{E}[v(b'(s'), s')|s]$. Optimality for consumption implies that $\lambda \geq 1$ and the non-negativity constraint on consumption is binding if $\lambda > 1$.

To interpret condition (9) rewrite it as:

$$\lambda = \frac{\beta \mathbb{E}[v'|s]}{G_1(k', 1; s) - \hat{\beta} \mathbb{E}[b'|s]} \geq 1. \tag{10}$$

When the inequality is strict the entrepreneur strictly prefers reducing current consumption to invest in new units of capital. If C was positive the entrepreneur could reduce it and use the additional funds to increase the capital stock. The marginal cost of an extra unit of capital is $G_1(k', 1; s)$ but the extra unit of capital increases collateral and allows the entrepreneur to borrow $\hat{\beta}\mathbb{E}[b'|s]$ more from the consumers. So a unit reduction in consumption leads to a levered increase in capital invested of $1/(G_1 - \hat{\beta}\mathbb{E}[b'|s])$. Since capital tomorrow increases future utility by $\beta E[v'|s]$, we obtain (10).

Condition (9) can be used to derive our main result on average and marginal q. The value of all the claims on the firm's future earnings, held by investors and by the en-

[8]Other recent models that allow for state-contingent claims include He and Krishnamurthy (2013) and Rampini and Viswanathan (2013). Cao (2018) develops a general model with an explicit stochastic structure that studies collateral constraints with non-state-contingent debt.

trepreneur at the end of the period, is

$$\hat{\beta}\mathbb{E}\left[B'\left(s'\right)|s\right]+\beta\mathbb{E}\left[V\left(K',B'\left(s'\right),s'\right)|s\right].$$

Dividing by total capital invested gives us average q:

$$q^a \equiv \hat{\beta}\mathbb{E}\left[b'|s\right]+\beta\mathbb{E}\left[v'|s\right].$$

Marginal q, on the other hand, is just the marginal cost of one unit of new capital, $q^m \equiv G_1\left(k',1;s\right)$. Rearrange equation (9) and express it in terms of q^a and q^m to get:

$$q^a = q^m + \frac{\lambda-1}{\lambda}\beta\mathbb{E}\left[v'|s\right]. \tag{11}$$

Since $\lambda > 1$ if and only if the non-negativity constraint on consumption is binding, we have proved the following result.

Proposition 1. *Average q is greater than or equal to marginal q, with strict equality if and only if the non-negativity constraint on consumption is binding.*

Equation (11) also shows that the difference between average and marginal q is increasing in the Lagrange multiplier λ and in the future value of entrepreneurial equity $\mathbb{E}\left[v'|s\right]$ (if $\lambda > 1$). As we shall see in the numerical part of the paper, an increase in indebtedness b increases λ but reduces the future value of entrepreneurial equity, so in general the relation between b and $q^a - q^m$ can be non-monotone. There is a cutoff for b such that $\lambda = 1$ below the cutoff and $\lambda > 1$ above the cutoff, so the relation must be increasing in some region.

The first order condition for b' can be written as

$$\hat{\beta}\lambda + \beta v_b\left(b'\left(s'\right),s'\right) = \mu(s'),$$

where $\pi(s'|s)\mu(s')k'$ is the Lagrange multiplier on the debt constraint (8). Using the envelope condition for b to substitute for v_b and using time subscripts, write

$$\lambda_t = \frac{\beta}{\hat{\beta}}\lambda_{t+1} + \frac{1}{\hat{\beta}}\mu_{t+1}. \tag{12}$$

This condition shows that λ_t is a forward looking variable determined by current and future values of μ_{t+1}. Positive values of this Lagrange multiplier in the future induce the entrepreneur to reduce consumption today to increase internal funds available. The forward looking nature of λ_t will be useful to interpret some of our numerical results about news shocks.

Now one can see the role of our assumption $\beta < \hat{\beta}$. If we had $\hat{\beta} = \beta$, condition (12) would imply that if, at some date t, the entrepreneur's consumption is positive and $\lambda_t = 1$, then the non-negativity constraint and the collateral constraint can not be binding at any future date. In other words, once the entrepreneur is unconstrained he can never go back to being constrained. This is due to the assumption of complete state contingent markets. Assuming $\beta < \hat{\beta}$ ensures that entrepreneurs alternate between periods in which they are constrained and periods in which they are unconstrained.

To conclude this section, let us introduce some asset pricing relations that characterize the equilibrium. The notation $G_{1,t}$ and $G_{2,t}$ is shorthand for $G_1 (K_{t+1}, K_t; s_t)$ and $G_2 (K_{t+1}, K_t; s_t)$.

Proposition 2. *The following conditions hold in equilibrium*

$$\lambda_t = \beta \mathbb{E}_t \left[\lambda_{t+1} \frac{A_{t+1} - G_{2,t+1} - b_{t+1}}{G_{1,t} - \hat{\beta} \mathbb{E}_t b_{t+1}} \right], \tag{13}$$

and

$$\hat{\beta} \mathbb{E}_t \left[\frac{A_{t+1} - G_{2,t+1}}{G_{1,t}} \right] \geq 1 \geq \mathbb{E}_t \left[\frac{\beta \lambda_{t+1}}{\lambda_t} \frac{A_{t+1} - G_{2,t+1}}{G_{1,t}} \right]. \tag{14}$$

The last two conditions hold with strict inequality if the collateral constraint is binding with positive probability.

The ratio

$$\frac{A_{t+1} - G_{2,t+1} - b_{t+1}}{G_{1,t} - \hat{\beta} \mathbb{E}_t b_{t+1}}$$

represents the levered rate of return on capital. Condition (13) further illustrates the forward-looking nature of λ_t. In particular, it shows that λ_t is a geometric cumulate of all future levered returns on capital. Condition (13) can also be interpreted as a standard asset pricing condition, dividing both sides by λ_t and observing that $\beta \lambda_{t+1} / \lambda_t$ is the stochastic discount factor of the entrepreneur.

The expression

$$\frac{A_{t+1} - G_{2,t+1}}{G_{1,t}}$$

is the unlevered return on capital. When the collateral constraint is binding the first inequality in (14) is strict and this implies that the expected rate of return on capital is higher than the interest rate $1 + r$. This implies that the levered return on capital is higher than the unlevered return. The entrepreneurs will borrow up to the point at which the discounted levered rate of return is 1, by condition (13). At that point the discounted unlevered return will be smaller than 1, by the second inequality in (14).

Define the finance premium as the difference between the expected return on entrepreneurial capital and the interest rate:

$$fp_t \equiv \mathbb{E}_t \left[\frac{A_{t+1} - G_{2,t+1}}{G_{1,t}} \right] - (1 + r) \,. \tag{15}$$

The first inequality in (14) shows that the finance premium is positive whenever the collateral constraint is binding. This definition of the finance premium is used in an alternative calibration in the online appendix.

3 No adjustment costs

This section considers the case of zero adjustment costs, that is $\xi = 0$ in equation (1). In this case, analytical results can be derived that map directly the shock structure into the coefficients of the investment regression.

With zero adjustment costs, the value function is linear

$$V(K, B, s) = \Lambda(s) \left[R(s) K - B \right], \tag{16}$$

where R is the gross return on capital defined by

$$R(s) \equiv A(s) + \phi(s)(1 - \delta(s)).$$

With a linear value function the borrowing limits are simply

$$\bar{b}(s) = \theta R(s) \,, \tag{17}$$

and they have a natural interpretation: the entrepreneur can pledge a fraction θ of the firm's gross returns.

We now make assumptions that ensure that the problem is well defined and that the collateral constraint is always binding in equilibrium. Assume the following three inequalities hold for all s:

$$\beta\mathbb{E}\left[R\left(s'\right)|s\right] > 1, \tag{18}$$

$$\theta\hat{\beta}\mathbb{E}\left[R\left(s'\right)|s\right] < 1, \tag{19}$$

$$\frac{(1-\gamma)(1-\theta)\beta\mathbb{E}\left[R\left(s'\right)|s, s' \neq s^r\right]}{\phi(s) - \theta\hat{\beta}\mathbb{E}\left[R\left(s'\right)|s\right]} \leq \zeta, \tag{20}$$

for some scalar $\zeta < 1$. Condition (18) implies that the expected rate of return on capital, discounted using entrepreneur's discount factor, is greater than 1, so entrepreneurs prefer investment to consumption. Condition (19) implies that pledgeable returns are insufficient to finance the purchase of one unit of capital, i.e., investment cannot be fully financed with outside funds. This condition ensures that investment is finite. Finally, condition (20) ensures that the entrepreneurs' utility is bounded. The last condition allows us to use the contraction mapping theorem to characterize the equilibrium marginal value of wealth $\Lambda(s)$ in the following proposition. The proof of this lemma and of the following results in this section are in the online appendix.

Lemma 1. *If conditions (18)-(20) hold there is a unique function $\Lambda : \mathbf{S} \to [1, \infty)$ that satisfies the recursion*

$$\Lambda(s) = \frac{\beta(1-\theta)\,\mathbb{E}\left[\Lambda\left(s'\right)R\left(s'\right)|s\right]}{\phi(s) - \theta\hat{\beta}\mathbb{E}\left[R\left(s'\right)|s\right]}, \text{ for all } s \neq s^r, \tag{21}$$

and $\Lambda(s) = 1$ for $s = s^r$.

Condition (21) is a special case of condition (13), in which the constraint is always binding. The following proposition characterizes an equilibrium.

Proposition 3. *If conditions (18)-(20) hold and $\Lambda(s)$ satisfies*

$$\Lambda(s) > \frac{\beta}{\hat{\beta}}\Lambda\left(s'\right), \tag{22}$$

for all $s, s' \in \mathbf{S}$, then the collateral constraint is binding in all states, consumption is zero until the retirement shock, investment in all periods before retirement is given by

$$\frac{K' - (1-\delta(s))K}{K} = \frac{(1-\theta)R(s)}{\phi(s) - \theta\hat{\beta}\sum_{s'}\pi(s'|s)R\left(s'\right)} - (1-\delta(s)), \tag{23}$$

and average q is

$$q^a = \mathbb{E}\left[\left((1-\theta)\beta\Lambda\left(s'\right) + \theta\hat{\beta}\right)R\left(s'\right)|s\right]. \tag{24}$$

Condition (22) ensures that entrepreneurs never delay investment. Namely, it implies that they always prefer to invest in physical capital today rather than buying a state-contingent security that pays in some future state.

The entrepreneur's problem can be analyzed under weaker versions of (18)-(22), but then the constraint will be non-binding in some states. It is useful to remark that we could embed our model in a general equilibrium environment with a constant returns to scale production function in capital and labor and a fixed supply of labor. In this general equilibrium model $A(s)$ is replaced by the endogenous value of the marginal product of capital. It is then possible to derive conditions (18)-(22) endogenously if shocks are small and the non-stochastic steady state features a binding collateral constraint.

Assume now that conditions (18)-(22) hold and let us analyze the model by linearizing the equilibrium conditions (21), (23) and (24) around the non-stochastic steady state. Steady state values are denoted by a bar. A tilde denotes deviations from the steady state, in levels or logs depending on the variable. Namely, level deviations are used for the following variables that are already expressed as ratios: q_t^a, A_t (profits to assets), and the investment rate, or investment to assets ratio,

$$IK_t = \frac{K_{t+1} - (1 - \delta_t) K_t}{K_t}.$$

So, for example, $\tilde{q}_t^a = q_t^a - \bar{q}^a$. Log deviations are used for the variables $\Lambda_t, \delta_t, \phi_t$. So for example, $\tilde{\Lambda}_t = \log \Lambda_t - \log \bar{\Lambda}$. Finally, for R_t, the approximation is

$$\tilde{R}_t = \tilde{A}_t + \tilde{\phi}_t(1 - \bar{\delta}) - \bar{\delta}\tilde{\delta}_t.$$

The steady state price of capital is normalized to $\bar{\phi} = 1$.

The following proposition characterizes the dynamics of investment and average q around the steady state.

Proposition 4. *If the economy satisfies (18)-(22) a linear approximation gives the following expressions for investment, average q and the marginal utility of entrepreneurial wealth* Λ_t*:*

$$\tilde{IK}_t = \frac{1-\theta}{1-\theta\hat{\beta}\bar{R}}\tilde{R}_t + \frac{(1-\theta)\bar{R}}{1-\theta\hat{\beta}\bar{R}} \frac{\mathbb{E}_t\left[\theta\hat{\beta}\tilde{R}_{t+1}\right] - \tilde{\phi}_t}{1-\theta\hat{\beta}\bar{R}} + \bar{\delta}\tilde{\delta}_t, \quad (25)$$

$$\tilde{q}_t^a = \mathbb{E}_t\left[(1-\theta)\,\beta\tilde{\Lambda}_{t+1}\bar{\Lambda}\bar{R} + (1-\theta)\beta((1-\gamma)\bar{\Lambda} + \gamma)\tilde{R}_{t+1} + \theta\hat{\beta}\tilde{R}_{t+1}\right], \quad (26)$$

$$\tilde{\Lambda}_t = \frac{1}{1-\theta\hat{\beta}\bar{R}} \sum_{j=0}^{\infty} \left(\frac{(1-\gamma)\,\bar{\Lambda}}{\gamma + (1-\gamma)\,\bar{\Lambda}} \right)^j \mathbb{E}_t \left[\tilde{R}_{t+j+1}/\bar{R} \ - \tilde{\phi}_{t+j} \right]. \tag{27}$$

Equations (25)-(26) express investment and average q in terms of current and future expected values of productivity. Given assumptions about the exogenous processes for A_t, ϕ_t, δ_t, equations (25) and (26) give us all the information about the variance-covariance matrix of $(\tilde{IK}_t, \tilde{q}_t^a, \tilde{A}_t)$ and thus about investment regression coefficients. In particular, we are interested in the implications of the model for the investment regression

$$IK_{it} = a_{i0} + a_1 q_{it}^a + a_2 CFK_{it} + e_{it}, \tag{28}$$

where CFK_{it} is the ratio of cash flow to assets, which is identified with A_{it} in our model.

We now turn to a battery of examples that show how different shock structures lead to different implications for the variance-covariance matrix of investment, average q and cash flow and thus for investment regressions.

3.1 Examples: productivity shocks

We begin with examples that only include productivity shocks.

Example 1. Productivity $\tilde{A}_t$ follows the AR(1) process:

$$\tilde{A}_t = \rho \tilde{A}_{t-1} + \varepsilon_t,$$

where ε_t is an i.i.d. shock. There are no shocks to the price of capital and depreciation.

In this example, $\mathbb{E}_t \left[\tilde{A}_{t+j} \right] = \rho^j \tilde{A}_t$ so all future expected values of $\tilde{A}_t$ are proportional to the current value. Substituting in (25)-(26), it is easy to show that both $\tilde{q}_t^a$ and $\tilde{IK}_t$ are linear functions of $\tilde{A}_t$. Therefore, in this case cash flow and average q are both, separately, sufficient statistics for investment. This is true even though there is a financial constraint always binding, simply due to the fact that a single shock is driving both variables.

Example 2. Productivity $\tilde{A}_t$ has a persistent component x_t and a temporary component η_t:

$$\tilde{A}_t = x_t + \eta_t,$$
$$x_t = \rho x_{t-1} + \varepsilon_t.$$

There are no shocks to the price of capital and to depreciation.

In this example, we have $\mathbb{E}_t\left[\tilde{A}_{t+j}\right] = \rho^j x_t$, and substituting in (25)-(26), after some algebra, we obtain

$$\tilde{IK}_t = \frac{(1-\theta)\left(1-(1-\rho)\bar{R}\theta\hat{\beta}\right)}{\left(1-\theta\hat{\beta}\bar{R}\right)^2} x_t + \frac{1-\theta}{1-\theta\hat{\beta}\bar{R}} \eta_t,$$

and

$$\tilde{q}^a_t = \left[\beta\left(1-\theta\right)\left(\gamma+\left(1-\gamma\right)\bar{\Lambda}\right)+\theta\hat{\beta}+\frac{\beta\left(1-\theta\right)\left(1-\gamma\right)\left(\gamma+\left(1-\gamma\right)\bar{\Lambda}\right)}{\left(1-\theta\hat{\beta}\bar{R}\right)\left(\gamma+\left(1-\gamma\right)\left(1-\rho\right)\bar{\Lambda}\right)}\bar{\Lambda}\rho\right]\rho x_t.$$

If we run a regression of investment on average q and cash flow, cash flow is the only variable that can capture variations in η_t, so the coefficient on cash flow will be positive and equal to

$$\frac{1-\theta}{1-\theta\hat{\beta}\bar{R}},$$

and cash flow improves the explanatory power of the investment regression. The crucial observation is that average q is affected by the marginal value of entrepreneurial net worth, which is a forward looking variable that reflects expectations about all future excess returns on entrepreneurial capital.[9] Through this channel, average q responds to information about future values of A_t at all horizons. At the same time, investment is only driven by the current and next period value of A_t. The current value determines internal funds, the next period value determines collateral values. Putting these facts together implies that shocks that affect profitability differentially at different horizons can break the link between average q and investment.

To get a quantitative sense of the model implications, let us use the parameter values in the calibrated model of next section, summarized in the first two lines of Table 1 below. However, unlike in that parametrization, let us set the parameter $\xi = 0$ to zero (no adjustment costs) and calibrate the parameters $\bar{\delta}$ and γ to target the average values of q and of the investment rate specified in the next section, which requires setting

$$\bar{\delta} = 0.092 \quad \text{and} \quad \gamma = 0.095.$$

The linearization above yields the following coefficients on Q and cash flow in the investment regression:

$$a_1 = 0.0561 \quad \text{and} \quad a_2 = 1.0273.$$

[9] See the discussion following Proposition 2.

If we use as references the coefficients in Gilchrist and Himmelberg (1995) (0.033 and 0.24), the coefficient on q is close to the empirical counter-part while the coefficient on cash flow is too high. With two shocks and two regressors, the R^2 of the regression is exactly 1.

Notice that in this example, investment, q and cash flow are fully determined by the two random variables x_t and η_t and the coefficients are independent of the variance parameters. This implies that, given all the other parameters, the coefficients of the investment regression are independent of the values of the variances σ_ε^2 and σ_η^2, as long as both are positive. As we shall see, this result does not extend to the general model with adjustment costs.

As an aside, notice that in this example, the coefficient on cash flow is higher for firms with larger values of θ, i.e., for firms that can finance a larger fraction of investment with external funds. These firms respond more because they can lever more any temporary increase in internal funds. This is reminiscent of the observation in Kaplan and Zingales (1997) that the coefficient on cash flow in an investment regression should not be used as measure of the tightness of the financial constraint.

It is useful to remark that our microfoundation of the financial constraint matters for the results derived. In particular, equation (25) makes it clear that investment in our framework depends only on the future value of the firms' asset values at short horizons (R_{t+1}) and not on their value further in the future. This comes from the way we have formulated the participation constraint in (4), which allows the entrepreneur to re-enter financial markets after a default event. Other formulations may make future values of R_t enter in richer ways in current investment, through essentially a "franchise value" effect. It is possible that these forces could increase the correlation between investment and q, but we leave this investigation to future research.

3.2 Examples: additional shocks

We now add shocks to the price of capital ϕ_t and to the depreciation rate δ_t and look at their quantitative implications for investment regressions. Let us begin with ϕ_t.

Example 3. The productivity process is as in example 2. The price of capital follows the AR(1) process

$$\tilde{\phi}_{t+1} = \rho_\phi \tilde{\phi}_t + \nu_{t+1}.$$

To choose a reasonable parametrization for this process, we borrow from the literature on investment-specific technology shocks and with parameters taken from Justiniano

Figure 1: Effect of other shocks on investment regression

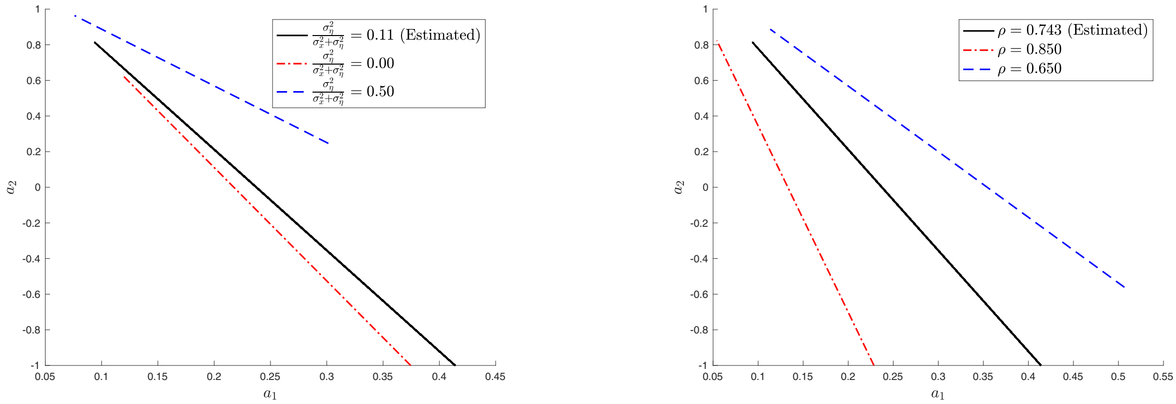

Note: Investment regression coefficients as the persistence and variance of shocks vary. The linear relations are described in Proposition 5.

et al. (2010):

$$\rho_\phi = 0.72 \quad \text{and} \quad \sigma_\nu = 0.063.$$

The coefficients on q and cash flow in the investment regression are now

$$a_1 = 0.1238 \quad \text{and} \quad a_2 = 0.6434.$$

and R^2 is now 0.9525. So adding the price of capital shock reduces the coefficient on cash flow, but increases the coefficient on q. The intuition for this result is that the ϕ_t shock affects investment and q but does not affect cash flow. Therefore it tends to increase the coefficient on q and decrease the coefficient on cash flow in the investment regression.

We have experimented with adding an AR(1) process for the depreciation variable, alone and in combination with ϕ_t shocks, obtaining analogous results. The investment regression coefficients for a variety of parametrizations of the processes of productivity, ϕ_t and δ_t are reported in Figure 3.2. Notice that one fixes the productivity process, there is a linear relation between a_1 and a_2. The processes for ϕ_t and δ_t determine a point in that linear relation, but not the position of the line. The latter is determined by the relative variance of the two productivity shocks (see the right panel) and by the persistence of the persistent component (see the left panel). This is an analytical result which relies on the fact that cash flow is not affected by the other shocks. The result is stated in the following proposition and proved in the online appendix.

Proposition 5. *Suppose productivity follows the process in Example 2. Then the coefficients* a_1, a_2

Figure 2: Effect of news shocks on investment regression

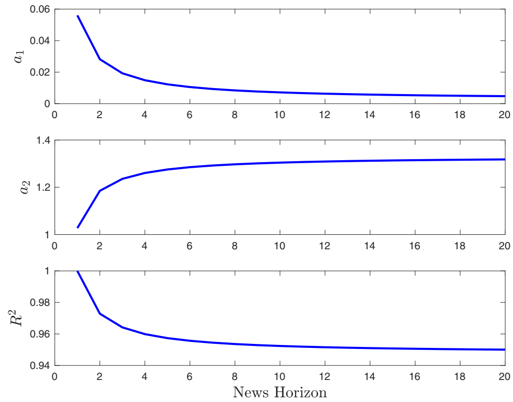

Note: Regression coefficients and R^2 as the news horizon increases

417 *in the investment regression satisfy*

$$\alpha_{q1}\sigma_x^2 a_1 + (\sigma_x^2 + \sigma_\eta^2)a_2 = \alpha_{i1}\sigma_x^2 + \alpha_{i2}\sigma_\eta^2, \tag{29}$$

418 *for some coefficients α_{i1}, α_{i2} and α_{q1} that do not depend on the shock processes for ϕ_t and δ_t.*

3.3 Examples: news shocks

We observed above that the presence of productivity shocks at different horizon alters the relation between q and investment. Building on this observation, we now introduce news shocks, that is shocks that reveal information about future profitability.

Example 4. The productivity process is as in Example 2 but the value of the permanent component x_t is known J periods in advance, with $J \geq 1$.

In the online appendix, we provide derivations for the dynamics of q and investment in this example and prove the following result.

Proposition 6. *In the economy of Example 4, all else equal, increasing the horizon J at which shocks are anticipated decreases the coefficient on average q, increases the coefficient on cash flow, and reduces the R^2 of the investment regression.*

The proof of this result is in the online appendix. Investment, as in the previous example, is just a linear function of productivity at times t and $t+1$, which fully determine current cash flow and collateral values. On the other hand, q is a function of all future values of A_t and, given the presence of news, these values are driven by anticipated future shocks which have no effect on investment. This weakens the relation between q and investment. Moreover, since q is the only source of information about x_{t+1}, and, with news shocks, it becomes a noisier source of information, this also reduces the joint explanatory power of q and cash flow.

Notice that news shocks here are acting very much like measurement error in q, by adding a shock to it that is unrelated to the shocks driving investment. However, financial frictions are essential in introducing this source of error. Absent financial frictions future values of productivity should not affect q, and it is only because q includes future quasi-rents that the relation arises.

To get a sense for the quantitative implications of new shocks, Figure 2 shows how the regression coefficients and R^2 change with the news horizon. Consistently with Proposition 6, as J increases, the coefficient on Q decreases and the coefficient on cash flow increases (starting from a relatively high value when there is no news), and R^2 decreases.

4 Adjustment costs

Let us now turn to the full model with adjustment costs and analyze its implications using numerical simulations. While the no adjustment cost model analyzed above is useful to build intuition, it has unrealistic implications for the responses of investment to shocks. In particular, it produces too large investment volatility for all plausible parametrizations. The model with adjustment costs, on the other hand, can be calibrated to match moments of the observed processes for profits and investment, so we can look at its quantitative implications.

FIrst, our choice of parameters is presented and the equilibrium is characterized in terms of policy functions and impulse responses. We then run investment regressions on the simulated output and explore the model's ability to replicate empirical investment regressions.

In the baseline calibration, there are only productivity shocks. Shocks to the price of capital are added later.

4.1 Calibration

The time period in the model is one year. The baseline parameter values are summarized in Table 1. The first three parameters are pre-set, the remaining parameters are calibrated on Compustat data. We now describe their choice in detail.

The investors' discount factor $\hat{\beta}$ is chosen so that the implied interest rate is 8.7%. As

Table 1: Parameters

Preset	β	$\hat{\beta}$	θ	
	0.90	0.92	0.3	
Calibrated to cash flow moments	$\bar{A}$	ρ	σ_ε	σ_η
	0.246	0.743	0.0713	0.0375
Calibrated to investment and q moments	δ	ξ	γ	
	0.0250	1.75	0.095	

argued by Abel and Eberly (2011) the interest rate used in this type of exercises should correspond to a risk-adjusted expected return. The number chosen is in the range of rates of return used in the literature.[10] The entrepreneurs' discount factor β has effects similar to the parameter γ which governs their exit rate. In particular, both affect the incentives of entrepreneurs to accumulate wealth and become financially unconstrained and both affect the forward looking component of q. Therefore, β is set at a level lower than $\hat{\beta}$ and γ is calibrated.[11] Regarding the fraction of non-divertible assets θ, there is only indirect empirical evidence, and existing simulations in the literature have used a wide range of values. Here $\theta = 0.3$ is chosen in line with evidence in Fazzari et al. (1988) and Nezafat and Slavic (2014). In particular, Fazzari et al. (1988) report that 30% of manufacturing investment is financed externally. Nezafat and Slavic (2014) use US Flow of Funds data for non-financial firms to estimate the ratio of funds raised in the market to finance fixed investment, and find a mean value of 0.284.

The parameters in the second line of Table 1 are calibrated to match moments of the firm-level cash flow time series in Compustat. Profits per unit of capital A_t are the sum of a persistent and a temporary component as in Example 2. Profits per unit of capital in the model, A_{it}, are identified with cash flow per unit of capital in the data, denoted by CFK_{it}.[12] The mean of A_t, denoted by $\bar{A}$, is set equal to average cash flow per unit of capital in the data. The values of ρ, σ_ε and σ_η are chosen to match the first and second order autocorrelation and the standard deviation of cash flow in the data, denoted, respectively, by $\rho_1(CFK)$, $\rho_2(CFK)$ and $\sigma(CFK)$. These moments are estimated using the approach of Arellano and Bond (1991) and Arellano and Bover (1995) and are reported in Table 1.[13]

[10] Abel and Eberly (2011) and DeMarzo et al. (2012) choose numbers near 10%, while Moyen (2004) and Gomes (2001) use $r = 6.5\%$.

[11] Changing the chosen value of β in a reasonable range does not affect the results significantly.

[12] Cash flow is equal to net income before extraordinary items plus depreciation.

[13] We estimate the firm-specific variation in cash-flow by first taking out the aggregate mean for each year and then applying the function xtabond2 in STATA. This implements the GMM approach of Arellano and Bover (1995). This approach avoids estimating individual fixed effects affecting both the dependent variable (cash flow) and one of the independent variables (lagged cash flow), by first-differencing the law of motion for cash flow, and then using both lagged differences and lagged levels as instruments. We use

Table 2: Target moments and model values

Moment	$\rho_1(CFK)$	$\rho_2(CFK)$	$\sigma(CFK)$	$\mu(IK)$	$\sigma(IK)$	$\mu(q^a)$
Target value	0.60	0.41	0.113	0.17	0.111	2.5
Model value	0.60	0.41	0.113	0.23	0.098	2.5

Notice that simply computing raw autocorrelations in the data—as sometimes done in the literature—would lead to biased estimates, given the short sample length.[14] In terms of sample, we use the same sub-sample of Compustat used in Gilchrist and Himmelberg (1995) so that we can compare our simulated regressions to their results.[15]

The next three parameters in Table 1, δ, ξ, and γ, are chosen to match three moments from the Compustat sample: the mean and standard deviation of the investment rate, $\mu(IK)$ and $\sigma(IK)$, and the mean of average q, $\mu(q^a)$. The reason why δ and ξ help determine the level and volatility of the investment rate is intuitive, as these two parameters determine the depreciation rate and the slope of the adjustment cost function. The parameter γ controls the speed at which entrepreneurs exit, so it affects the discounted present value of the quasi-rents they expect to receive in the future and thus average q. However, the three parameters interact, so we choose them jointly—by a grid search—in order to minimizes the average squared percentage deviation between the three model-generated moments and their targets. The target moments from the data and the model generated moments are reported in Table 2.[16]

Notice that there is a tension between hitting the targets for $\mu(IK)$ and $\sigma(IK)$. Increasing any of the parameters, δ, ξ, γ reduces $\mu(IK)$, bringing it closer to its target value, but also decreases $\sigma(IK)$, bringing it farther from its target. Notice also that it is important for our purposes that the model generates a realistic level of volatility in the investment rate, given that IK is the dependent variable in the regressions we will present in Section 4.3 below.

Our calibration also determines the average size of the wedge between average and marginal q. In particular, $\mu(q^a) = 2.5$ is the mean value of average q while ξ and $\mu(IK)$ determine the mean value of marginal q, which is $1 + \xi(\mu(IK) - \delta) = 1.25$. Therefore, the average wedge between average and marginal q is 1.25. Since the presence of the wedge

the first three available (non-autocorrelated) lags in differences as instruments, with lags chosen separately for the 1st and 2nd order autocorrelation estimation. One lagged level is also used as an instrument.

[14]This type of bias was first documented in Nickell (1981). The bias is non-negligible in our sample. For the first-order autocorrelation, the Arellano and Bond (1991) approach gives $\rho_1(CFK) = 0.60$, while the raw autocorrelation in the data is 0.42.

[15]In particular, we restrict attention to the sample period 1978-1989 and use the same 428 listed firms used in their paper.

[16]The target standard deviation $\sigma(IK)$ is a pooled estimate.

Table 3: Calibration of frictionless model

Parameter	δ	ξ	γ	Moment	$\mu(IK)$	$\mu(q^a)$	$\sigma(IK)$
	0.05	1.50	0.125	Target value	0.17	2.5	0.111
				Model value	0.18	1.2	0.116

is what breaks the sufficient statistic property of q it is useful that our calibration imposes some discipline on the wedge's size.

All the simulations assume that entrepreneurs enter the economy with a unit endowment of capital and zero financial wealth (i.e., zero current profits and zero debt). Since the entrepreneurs' problem is invariant to the capital stock and all our empirical targets are normalized by total assets, the choice of the initial capital endowment is just a normalization. We have experimented with different initial conditions for financial wealth, but they have small effects on our results given that—with our parameters—the state variable b converges quickly to its stationary distribution.

It is useful to compare our results to those of a benchmark model with no financial frictions. To make the parametrization of the two models comparable, the parameters δ, ξ and γ are re-calibrated for the frictionless case. The moments and associated parameters are reported in Table 3. Notice that the frictionless model generates a low value of $\mu(q^a)$. For given IK, increasing ξ would increase marginal and average q (which are the same in the frictionless case), but it would reduce the volatility of investment.

4.2 Model dynamics

To describe the model behavior, it helps intuition to use as state variables A_t and n_t rather than A_t and b_t, where $n_t = A_t + \phi_t (1 - \delta_t) - b_t$, is the ratio of net worth (excluding adjustment costs) to assets K_t.

Each row of Figure 3 plots the value function (per unit of capital) v, the optimal investment ratio K'/K, the Lagrange multiplier λ on the entrepreneur's budget constraint, and the wedge between average q and marginal q. Each column corresponds to different values of x_t. In particular, the values reported correspond to the the 20th, 50th and 80th percentile of the unconditional distribution of x. On the horizontal axis there is n, but the domain differs between columns as we plot values between the 10th to 90th percentile of the conditional distribution of n, conditional on the reported value of x.[17]

A higher level of n leads to a higher value v and a higher level of investment K'/K. Moreover, the value function is concave in n. The Lagrange multiplier λ is equal to the

[17] The joint distribution of (n, x) is computed numerically as the invariant joint distribution generated by the optimal policies.

Figure 3: Characterization of equilibrium

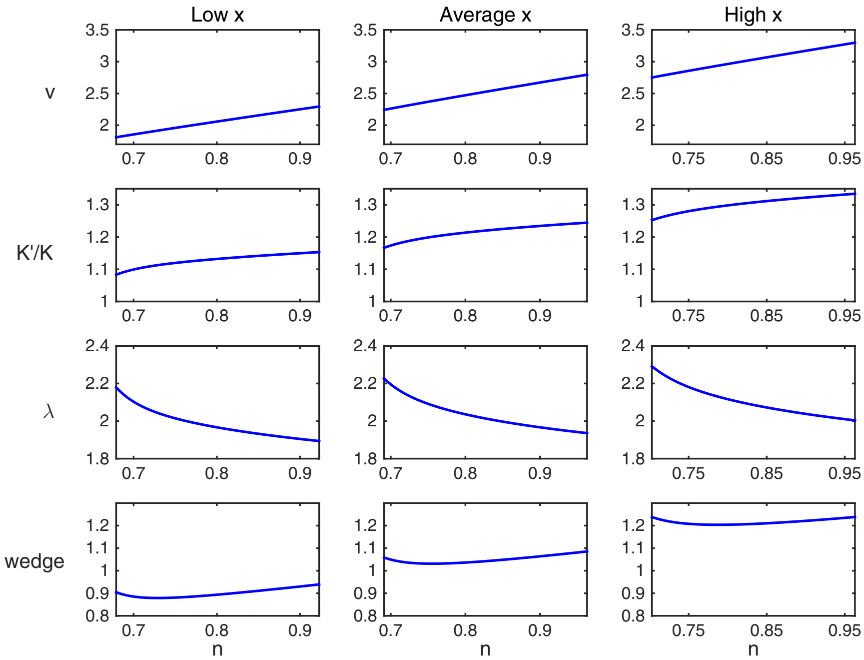

Note: The three columns correspond to the 20th, 50th, and 80th percentile of the persistent component of productivity x. The range for the net worth variable n is between the 10th and 90th percentiles of the distribution of n conditional on x.

541 derivative of the value function and therefore is decreasing in n. The fact that λ is de-
542 creasing in n reflects the fact that a higher ratio of net worth to capital allows firms to
543 invest more, leading to a higher shadow cost of capital G_1 and thus to a lower expected
544 returns on investment. Eventually, for very high values of n we reach $\lambda = 1$. However, as
545 the figures show this does not happen for the range of n values more frequently visited
546 in equilibrium.

547 The bottom row documents how the wedge varies with the level of net worth n and
548 with the persistent component of productivity x. Let us first look at the effect of n. Even
549 though λ is decreasing in n, the wedge, $q^a - q^m$, does not vary much with n for a given
550 value of x. Our analytical derivations in Section 2 help explain this outcome. Recall from
551 equation (11) that the wedge is equal to

$$\frac{\lambda - 1}{\lambda} \beta \mathbb{E} \left[v' | s \right].$$

552 When we reach the unconstrained solution and $\lambda = 1$ the wedge disappears. However,
553 for lower levels of n, for which the constraint is binding, the relation is in general non-

monotone. An increase in n reduces the marginal gain from an extra unit of net worth. However, at the same time it increases the future growth rate of firm's capital stock and so it increases the base to which this marginal quasi-rent is applied. This second effect is captured by the expression $\mathbb{E}[v'|s]$, because the value per unit of capital v' embeds the future growth of the firm and is increasing in n. The plots in the bottom row of Figure 3 show that in the relevant range of n these two effects roughly cancel.

On the other hand, comparing the values of the wedge across columns, shows that persistent component of productivity x has large effects on the wedge and that the wedge is increasing in x. The reason is that higher values of x lead both to higher values of λ, as the marginal benefits of extra internal funds increase with productivity, and to higher values of K'/K and v, because higher productivity allows the firm to raise more external funds and grow faster. Therefore both elements of the wedge increase with higher values of x.

We now present impulse response functions following the two shocks. To construct these impulse response functions, we start at the median values of the state variables n and x. We then introduce a shock, simulate 10^6 paths following the shock, and report the difference between the average simulated paths, with and without the initial shock. Given the non-linearity of the model, the initial conditions for n and x in general affect the responses. However, in our simulations these non-linear effects are relatively small, so the plots below are representative.

The top panel of Figure 4 plots responses to a 1-standard-deviation persistent shock ε.[18] Following a persistent shock all variables increase and return gradually to trend. The response of average q is larger than that of marginal q, thus producing an increase in the wedge. The bottom panel plots responses to the temporary shock η. Also in this case all three variables respond positively, but the response is more short-lived. Moreover, now the response of average q is slightly smaller than the response of marginal q, so the wedge shows a small decrease after the shock.

Average q is a forward-looking variable that incorporates the quasi-rents that the entrepreneur is expected to receive in the future. These quasi-rents are only marginally affected by a temporary shock. In the model with no adjustment costs, the effect is zero. Here, because of adjustment costs, there is a positive effect, due to the fact that the investment response displays a small but positive degree of persistence and high investment in the future increases the future value of installed capital. But the effect is small. In the case of a persistent shock, instead, future quasi-rents are directly affected by higher future

[18]The response of investment K'/K is always proportional to the response of marginal q and is thus omitted.

Figure 4: Impulse response functions

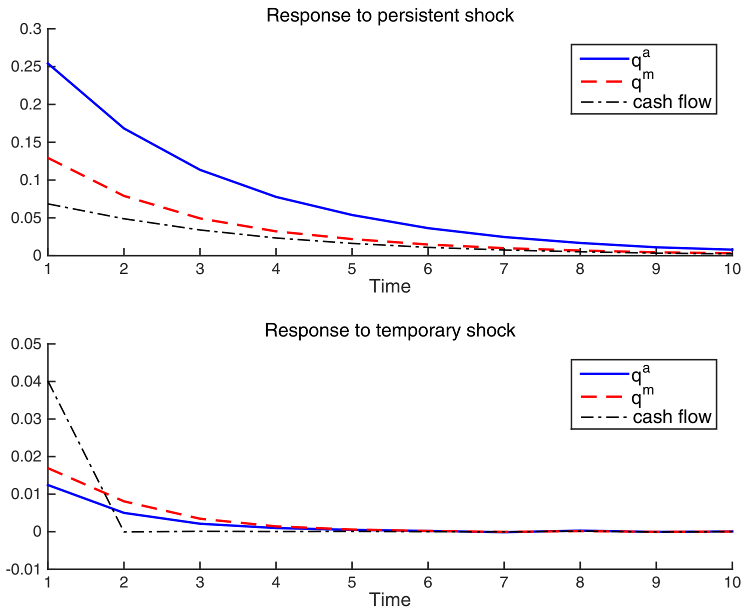

Note: Average paths following a shock at time 1, in (level) deviations from average paths following no shock. Cash flow is cash flow per unit of capital.

productivity, which will lead to faster growth (as shown in Figure 3), thus explaining the large increase in q^a in the top panel of Figure 4.

The discussion following Figure 3, helps to explain the response of the wedge $q^a - q^m$. A temporary shock leads to a pure increase in net worth per unit of capital. The effect of such an increase on the wedge is in general ambiguous and, with our parameter choices, close to zero. In the case of a persistent shock, instead, the effect is unambiguously to increase the wedge.

4.3 Investment regressions

We now turn to investment regressions, and ask whether the model can replicate the coefficients on q and cash flow observed in the data. In particular, we ask to what extent does the presence of a financial friction help to obtain a smaller coefficient on q and a positive and large coefficient on cash flow. To answer this question, simulated data are generated from our model and they are used to run the investment regression (28). In line with the empirical literature, we generate a balanced panel of 500 firms for 20 periods,

Table 4: Investment regressions

	a_1	a_2	R^2	Univariate q^a coefficient	R^2	Univariate CFK coefficient	R^2
Baseline model	0.22	0.14	0.98	0.26	0.98	0.81	0.89
Frictionless model	0.67	0.00	1.00	0.67	1.00	0.95	0.86
GH (1995)	0.033	0.24		0.05			

and allow for firm-level fixed effects in the regression.[19] All reported results are the mean values for 50 simulated panels.

The regression coefficients for the baseline model are presented in the first row of Table 4. As reference points, the second row reports the coefficients that arise in the model without financial frictions and in the last row the empirical estimates in Gilchrist and Himmelberg (1995), which are representative of the orders of magnitude obtained in empirical studies.[20] The table also reports coefficients of univariate regressions of investment on average q and cash-flow separately.

The results for the frictionless benchmark are reported in the second line of Table 4. In this case, average q is a sufficient statistic for investment, the coefficient on cash flow is zero and the coefficient on q is equal to the inverse of the adjustment cost coefficient ξ, which is calibrated to 1.5. This line shows the standard empirical failure of the benchmark adjustment cost model.

Adding financial frictions helps to get a smaller coefficient on q and a positive coefficient on cash flow. The effect is sizable, although the coefficient on q is still large compared to the very small numbers found in empirical regressions. Notice also that the R^2 of the regression is very close to 1. This is not surprising given the simple two-shock structure and the presence of two explanatory variables.[21] Given that the model is non-linear, the R^2 can in general be smaller than 1. However, by experimenting with impulse responses for different initial values of the state variables we have confirmed that, given our parameter values, the model is close to linear in its responses to the two shocks, which helps to explain the high R^2 in Table 4.

The presence of the wedge breaks the one-to-one relation between q and investment and allows for cash flow to have explanatory power in the the investment regression. In particular, as can be seen in Figure 4 the wedge responds in opposite directions to the two shocks, while q^m respond positively to both. So the wedge plays a role somewhat similar

[19]The model features random exit, so to generate a balanced panel we only keep firms for which exit does not occur for 20 periods.

[20]We do not report standard errors, but they are small (less than 0.04) for both coefficients in our simulated data. They are also small in the empirical estimates of Gilchrist and Himmelberg (1995).

[21]For the same reason, in the linear model of Example 2, Section 3, the R^2 is 1.

Table 5: Investment regressions: changing shock variances

σ_ε	σ_η	σ_η^2/σ_A^2	a_1	a_2	R^2	Univariate q^a coefficient	R^2	Univariate CFK coefficient	R^2
0.113	0.000	0.00	0.38	-0.48	0.99	0.25	0.99	0.90	0.98
0.071	0.037	0.11	0.22	0.15	0.98	0.27	0.98	0.81	0.89
0.033	0.080	0.50	0.28	0.11	0.96	0.32	0.95	0.48	0.56
0.006	0.107	0.90	0.34	0.10	0.84	0.38	0.75	0.18	0.32
0.000	0.113	1.00	2.47	0.01	0.92	2.53	0.92	0.11	0.37

to measurement error in dampening the regression coefficient. Notice however that the model still features a strong positive relation between q^a and investment, as documented by the fifth and sixth columns of Table 4, which show that a univariate regression between investment and average q produces a large coefficient and a large R^2 in simulated data (unlike in actual data). In the rest of the paper we investigate shock structures that can potentially weaken this relation.

It is useful to look at how the shock structure affects investment regressions. Table 5 reports regression coefficients and R^2 for different combinations of σ_ε and σ_η, keeping constant the total volatility of A_t. The second row corresponds to the baseline case of Table 4. The third column reports the fraction of variance due to the temporary shock. Here all remaining parameters are kept at their baseline level, in order to focus on how variance parameters affect the result.

The first row of Table 5 shows an extreme case with no temporary shocks. In this case, the coefficient on q is larger than in our baseline and the coefficient on cash flow is actually negative. The last row of the table shows the opposite extreme, with only temporary shocks. Interestingly, also this row displays a larger coefficient on q. The coefficient on cash flow in this case is close to zero. So going to a one-shock model, worsens the model performance in terms of replicating investment regressions. In this case q and investment tend to comove simply because they are driven by the same shock. In these cases, we get close to the sufficient statistic result obtained in the one-shock linear model of Example 1. Example 1 has indeterminate implications for the coefficients, due to the perfect collinearity of q and cash flow. Here, the perfect collinearity result does not hold for two reasons: first, the model displays inertia so past values of x_t determine investment and q, which complicates the correlation structure of investment, q and cash flow; second, the model is non-linear. For these reasons, the bivariate coefficients are determinate even with a single shock, and, in particular, the model prefers to assign a large coefficient on q.[22]

[22] The results in this table may help reconcile our results with the results of Gomes (2001). In particular, although Gomes (2001) uses a different model of financial frictions, it is possible that his result—that q is

Table 6: Investment regressions: shocks to the price of capital

σ_v	a_1	a_2	R^2	Univariate q^a coefficient	R^2	Univariate CFK coefficient	R^2
0.00	0.2212	0.1433	0.9837	0.2622	0.9799	0.8098	0.8883
0.01	0.2238	0.1351	0.9826	0.2622	0.9793	0.8132	0.8846
0.02	0.2309	0.1109	0.9798	0.2624	0.9775	0.8146	0.8740
0.03	0.2421	0.0741	0.9753	0.2626	0.9743	0.8166	0.8562
0.04	0.2561	0.0241	0.9695	0.2627	0.9694	0.8192	0.8314
0.05	0.2715	-0.0337	0.9632	0.2625	0.9630	0.8219	0.8002

The remaining rows of Table 5 illustrate intermediate cases in which both shocks are present. As argued above, both shocks increase investment but they have opposite effects on the wedge and that is what reduces the predictive power of q. So there is some intermediate mix of shocks that adds maximum noise to the information contained in average q and reduces the overall explanatory power of the investment regression. In the table this is visible in the non-monotone relation between the ratio σ_η^2/σ_A^2 and the R^2 of the regression.

While it is intuitive that mixing the two shocks reduces the total explanatory power of investment regressions and reduces R^2, the quantitative effects on the two coefficients a_1 and a_2 are more complex to interpret, as they also depend on the magnitudes of the responses of investment, cash flow, and q to the underlying shocks. In particular, persistent shocks tend to affect more, in relative terms, q than investment, due to the forward looking nature of q and the presence of the financial constraint which dampens the response of investment (see Figure 4).[23] Persistent shocks lead to a smaller response of investment for a given response of q, when compared to temporary shocks. This is immediately visible in the monotone increase in the univariate coefficient with σ_η^2/σ_A^2. The effect on the bivariate coefficient a_1 is more complex as, at the same time, the presence of temporary shocks increases the coefficient on cash flow. Therefore, the relation between each of the coefficients a_1 and a_2 and the variance ratio σ_η^2/σ_A^2 is non-monotone.

The overall take out from Table 5 is that, given all other model parameters, the relative variance of temporary and persistent shocks matter for both the explanatory power and for the individual coefficients in investment regressions.

We can now add to the model additional shocks, as done in the case of no adjustment costs in Section 3. In particular, we add the same AR1 process for the price of capital, with

almost a sufficient statistic for investment—could be driven by his one-shock structure.

[23]The same two reasons identified above (inertia and non-linearity) for one-shock models, explain why in the two-shock model the relative size of the two variances matter for the regression coefficients, unlike in the simple linearized model with no adjustment costs of Section 3, Example 2.

Table 7: News shocks: calibration

	Parameters			Moments			
	δ	ξ	γ	$\mu(IK)$	$\frac{\sigma(IK)}{\sigma(CFK)}$	$\mu(q^a)$	$\sigma(q^a)$
Targets				0.17	0.98	2.5	0.97
No news (7 states)	0.0250	2.00	0.09	0.22	0.79	2.49	0.27
No news (2 states)	0.0200	2.00	0.10	0.22	0.94	2.24	0.33
$J = 1$	0.0275	3.00	0.08	0.21	0.86	2.39	0.42
$J = 2$	0.0225	3.50	0.08	0.20	0.85	2.24	0.45
$J = 3$	0.0225	3.50	0.08	0.19	0.91	2.67	0.59
$J = 4$	0.0275	3.50	0.08	0.19	0.95	2.48	0.59
$J = 5$	0.0300	3.50	0.08	0.19	0.97	2.50	0.63

parameters from Justiniano et al. (2010) as in Section 3.2. Table 6 reports the regression results for different values of the variance of the price of capital shocks. As in the case of no adjustment costs the coefficient on q increases and the coefficient on cash flow decreases. Quantitatively, the slope of the relation between a_1 and a_2 is of a similar order of magnitude, but the relation is a bit flatter (i.e., the negative effect on the cash flow coefficient is relatively smaller than the positive effect on the q coefficient) in the calibration considered here. The underlying intuition is the same. Shocks to the cost of capital affect investment and q but do not affect cash flow, so they weaken the relation between cash flow and investment.

We now turn to news shocks. Example 4 in Section 3 shows that in the case of no adjustment costs news shocks introduce additional noise in average q, thus reducing its predictive power. Here we want to investigate whether the same forces are at work in our full model with adjustment costs and see their quantitative implications.

Introducing news shocks increases the number of state variables, since we need to keep track of anticipated values of x_t. Therefore, to simplify computations, we employ a coarser description of the permanent component of the productivity process, using a two-state Markov process for x_t. The stochastic process for A_t is specified and calibrated as in our baseline but we assume agents observe x_t J periods in advance as in Example 4 in Section 3. We experiment with $J = 1, 2, ..., 5$, re-calibrating the parameters δ, ξ and γ for each value of J. Table 7 reports the calibrated parameters for each value of J. The table also reports our baseline calibration (no news, 7 states) and a calibration with no news and a 2 states Markov chain, which help to evaluate the effect of news on our results.

Table 7 shows that introducing news shocks improves the model's ability to match the empirical level of the investment rate, reducing the value of $\mu(IK)$, while producing similar values for $\sigma(IK)$ and $\mu(q^a)$. The table also reports the volatility of q^a (which is not

Table 8: Investment regressions: news shocks

	a_1	a_2	R^2
No news (7 states)	0.2047	0.1530	0.984
No news (2 states)	0.2434	0.0829	0.985
$J = 1$	0.1920	−0.0121	0.982
$J = 2$	0.1774	0.0161	0.974
$J = 3$	0.1417	0.0502	0.978
$J = 4$	0.1467	0.0628	0.976
$J = 5$	0.1394	0.0824	0.971

used as a target for our calibration), and the table shows that introducing news improves the model's realism in this dimension. The analytical derivations in Section 3 (Example 4) suggest a reason for this: anticipated shocks seem to introduce an additional source of volatility in q^a.

Turning to investment regressions, Table 8 shows regression coefficients and R^2 for different values of J. The coefficient on q^a and the R^2 behave in a similar way as suggested by Example 4: increasing the horizon adds noise in q^a thus reducing the coefficient and the overall R^2. The cash flow coefficient goes down when going from no news to 1 period anticipation, and then increases monotonically in J.

Comparing the cases of no news and the case $J = 5$, the overall take away from Tables 7 and 8 is that news shocks improve the model's ability to match the observed behavior of investment, q and cash flow, both in terms of levels and volatility and in terms of the cross-correlations captured by investment regressions. The central intuition is that news shocks introduce a source of variation in q due to anticipated future shocks, which have little bearing on the contemporaneous movements in investment.

Due to the use of a 2 state Markov chain, the model with news does worse than the baseline in terms of the cash flow coefficient, so it is an open question for future work whether increasing the state space and possibly using alternative models of anticipated news that economize on state variables can further improve the model's empirical performance.[24]

5 Conclusions

The paper shows that financial frictions can help dynamic investment models move closer to the correlations observed in the data. The model in this paper is stylized, but the main conclusions on the role of different shocks are likely to extend to more complex

[24] See for example the information structure in Blanchard et al. (2013).

models. In particular, a promising avenue seems to be to build models where a substantial fraction of the volatility in q is associated to news about profitability relatively far in the future and where these news have relatively small effects on current investment decisions. By assuming risk neutrality, we have omitted an important source of volatility in asset prices, namely volatility in discount factors and risk premia. It is an important open question how these additional sources of volatility affect the correlations investigated here, especially because these factors are likely to correlate with the stringency of financial constraints for individual firms.

References

Abel, A. B. and Eberly, J. C. (2011). How q and cash flow affect investment without frictions: An analytic explanation. *The Review of Economic Studies*, 78(4):1179–1200.

Abel, A. B. and Eberly, J. C. (2012). Investment, valuation, and growth options. *Quarterly Journal of Finance*, 2.

Albuquerque, R. and Hopenhayn, H. A. (2004). Optimal lending contracts and firm dynamics. *The Review of Economic Studies*, 71(2):285–315.

Alti, A. (2003). How sensitive is investment to cash flow when financing is frictionless? *Journal of Finance*, 58:707–722.

Arellano, M. and Bond, S. (1991). Some tests of specification for panel data: Monte carlo evidence and an application to employment equations. *The Review of Economic Studies*, 58(2):277–297.

Arellano, M. and Bover, O. (1995). Another look at the instrumental variable estimation of error-components models. *Journal of Econometrics*, 68(1):29 – 51.

Atkeson, A. and Cole, H. (2005). A dynamic theory of optimal capital structure and executive compensation. Working Paper 11083, National Bureau of Economic Research.

Biais, B., Mariotti, T., Plantin, G., and Rochet, J.-C. (2007). Dynamic security design: Convergence to continuous time and asset pricing implications. *The Review of Economic Studies*, 74(2):345–390.

Blanchard, O. J., de Silanes, F. L., and Shleifer, A. (1994). What do firms do with cash windfalls? *Journal of Financial Economics*, 36(3):337 – 360.

Blanchard, O. J., L'Huillier, J.-P., and Lorenzoni, G. (2013). News, noise, and fluctuations: An empirical exploration. *American Economic Review*, 103(7):3045–70.

Cao, D. (2018). Speculation and financial wealth distribution under belief heterogeneity. *The Economic Journal*. Forthcoming.

Cao, D. and Nie, G. (2017). Amplification and asymmetric effects without collateral constraints. *American Economic Journal: Macroeconomics*, 9(3):222–66.

Clementi, G. L. and Hopenhayn, H. A. (2006). A theory of financing constraints and firm dynamics. *The Quarterly Journal of Economics*, 121(1):229–265.

Cooley, T., Marimon, R., and Quadrini, V. (2004). Aggregate consequences of limited contract enforceability. *Journal of Political Economy*, 112(4):817–847.

Cooper, R. and Ejarque, J. (2003). Financial frictions and investment: requiem in q. *Review of Economic Dynamics*, 6(4):710 – 728.

DeMarzo, P. M., Fishman, M. J., He, Z., and Wang, N. (2012). Dynamic agency and the q theory of investment. *The Journal of Finance*, 67(6):2295–2340.

DeMarzo, P. M. and Sannikov, Y. (2006). Optimal security design and dynamic capital structure in a continuous-time agency model. *The Journal of Finance*, 61(6):2681–2724.

Di Tella, S. (2017). Uncertainty shocks and balance sheet recessions. *Journal of Political Economy*, 125(6):2038–2081.

Eberly, J., Rebelo, S., and Vincent, N. (2008). Investment and value: A neoclassical benchmark. Working Paper 13866, National Bureau of Economic Research.

Fazzari, S. M., Hubbard, R. G., and Petersen, B. C. (1988). Financing constraints and corporate investment. *Brookings Papers on Economic Activity*, 1988(1):141–206.

Gilchrist, S. and Himmelberg, C. P. (1995). Evidence on the role of cash flow for investment. *Journal of Monetary Economics*, 36(3):541 – 572.

Gomes, J. F. (2001). Financing investment. *American Economic Review*, 91(5):1263–1285.

Hayashi, F. (1982). Tobin's marginal q and average q: A neoclassical interpretation. *Econometrica*, 50(1):pp. 213–224.

He, Z. and Krishnamurthy, A. (2013). Intermediary asset pricing. *American Economic Review*, 103(2):732–70.

Hennessy, C. A. and Whited, T. M. (2007). How costly is external financing? evidence from a structural estimation. *The Journal of Finance*, 62(4):1705–1745.

Hubbard, R. G. (1998). Capital-market imperfections and investment. *Journal of Economic Literature*, 36(1):193–225.

Justiniano, A., Primiceri, G. E., and Tambalotti, A. (2010). Investment shocks and business cycles. *Journal of Monetary Economics*, 57(2):132 – 145.

Kaplan, S. N. and Zingales, L. (1997). Do investment-cash flow sensitivities provide useful measures of financing constraints? *The Quarterly Journal of Economics*, 112(1):169–215.

Kiyotaki, N. and Moore, J. (1997). Credit cycles. *Journal of Political Economy*, 105(2):211–248.

Lorenzoni, G. (2008). Inefficient credit booms. *The Review of Economic Studies*, 75(3):809–833.

Moyen, N. (2004). Investment-cash flow sensitivities: Constrained versus unconstrained firms. *The Journal of Finance*, 59(5):2061–2092.

Nezafat, P. and Slavic, C. (2014). Asset prices and business cycles with financial shocks. Michigan State University Working Paper.

Nickell, S. (1981). Biases in dynamic models with fixed effects. *Econometrica*, 49(6):1417–1426.

Rampini, A. A. and Viswanathan, S. (2013). Collateral and capital structure. *Journal of Financial Economics*, 109(2):466 – 492.

Rauh, J. D. (2006). Investment and financing constraints: Evidence from the funding of corporate pension plans. *The Journal of Finance*, 61(1):33–71.

Sargent, T. J. (1980). Tobin's q and the rate of investment in general equilibrium. *Carnegie-Rochester Conference Series on Public Policy*, 12:107 – 154.

Schiantarelli, F. and Georgoutsos, D. (1990). Monopolistic competition and the q theory of investment. *European Economic Review*, 34(5):1061 – 1078.

ACCEPTED MANUSCRIPT