

The Corporate Propensity to Save

Leigh A. Riddick
American University

Toni M. Whited*
University of Wisconsin, Madison

December 5, 2007

*Corresponding author. School of Business, University of Wisconsin, Madison, WI 53706-1323. 608-262-6508. twhted@bus.wisc.edu. We are grateful for helpful comments and suggestions from an anonymous referee, an associate editor, the editors, Cam Harvey and John Graham, as well as Viral Acharya, Tor-Erik Bakke, Morris Davis, Mark Garmaise, Robert Hauswald, Chris Hennessy, Holger Mueller, Emi Nakamura, Michael Roberts, Josef Zechner, and seminar participants at the American University, Boston College, Cornell, Duke, HEC-Lausanne, HEC-Paris, LBS, LSE, Northwestern, the Norwegian School of Management, Rice, Rochester, the University of Vienna, the University of Wisconsin, and Yale. Bobby Hart of Thompson Financial was helpful in data acquisition, and Alex Boquist and Michael R. Sullivan provided research assistance. Leigh Riddick also wishes to acknowledge research support from the Kogod School of Business and American University.

The Corporate Propensity to Save

Abstract

Why and how do corporations accumulate liquid assets? We show theoretically that intertemporal trade-offs between interest income taxation and the cost of external finance determine optimal savings. We find the striking result that, conditional on Tobin's q , saving and cash flow are negatively related because firms lower cash reserves to invest after receiving positive cash-flow shocks, and vice versa. Consistent with theory, we estimate negative propensities to save out of cash flow. We also find that income uncertainty affects saving more than do external finance constraints. Therefore, contrary to previous evidence, saving propensities reflect too many individual forces to be used to measure external finance constraints.

Why do corporations save? In other words, why do they funnel their cash flow into liquid asset holdings rather than into physical capital or into shareholder distributions? This question is challenging because such financial decisions cannot be understood in isolation of the real decisions a corporation makes. Not only is the question challenging, but it is also economically interesting in light of the tendency in recent years of both U.S. and European firms to accumulate unusually high levels of cash. More generally, savings policies matter for all firms because managers must evaluate the trade-offs between using internal and external funds to finance current and future investment. Indeed, several recent studies—for example, Almeida, Campello and Weisbach (2004) and Khurana, Pereira, and Martin (2006)—have used firms’ saving behavior to gauge the cost of external finance.

The goal of this paper is therefore twofold. Although we do not tackle directly the issue of the high level of corporate cash holdings, we shed light on this phenomenon by delving into the economics of the process whereby firms accumulate this cash. We also wish to determine whether and when corporate saving behavior can be a useful indicator of the extent to which firms face external finance constraints. We examine these issues both empirically and theoretically, and we focus on two specific determinants of saving: income uncertainty and the cost of external finance.

On the theoretical side our model provides several insights. In our infinite-horizon framework firms invest, save, produce, raise external finance, and make distributions in the face of uncertainty, physical adjustment costs, taxation, and costly equity issuance. Because interest on cash balances is taxed, the firm faces a dynamic trade-off between this tax penalty and the reduction in expected future financing costs conferred by holding cash. Therefore, the firm’s optimal saving policy depends not only on the cost of external finance, but also on the firm’s expected future financing needs, which, in turn, depend on the firm’s technology and especially on the uncertainty it faces.

In this setting we find that firms hold higher precautionary cash balances when external finance is costly or income uncertainty is high. Firms also hold more cash if their optimal investment policy is lumpy because large investments typically entail costly financing. This connection between cash holding and large investments is consistent with the observed high corporate cash levels that have coincided with the recent wave of mergers and acquisitions,¹ which are a form of lumpy investment.

Our most interesting predictions concern saving, that is, the change in cash. In particular, we find

¹See Berman, Dennis K., Deals Boom Fizzles As Cheap Credit Fades, Wall Street Journal, September 7, 2007.

that, conditional on Tobin's q , saving and cash flow are negatively correlated. Although the result of a negative conditional correlation is, at first, surprising, the intuition is both economically interesting and straightforward. For example, if the firm faces positively serially correlated productivity shocks, conditional on a high shock, the firm's cash flow rises, its capital becomes more productive, and this productivity reverts to its mean slowly. The firm therefore shifts some of its financial asset holdings into physical capital; that is, it invests and dissaves. The amount of this dissaving is typically less than the firm's cash flow. Conversely, the firm accumulates more liquid assets in times when capital productivity is low. This substitution between physical and financial assets manifests itself in a negative conditional sensitivity of saving to cash flow. It is important to condition on Tobin's q because, roughly speaking, this variable capitalizes the value to the firm of holding cash. Indeed, in our model the *unconditional* correlation between saving and cash flow, although small, is positive because firms often do use their cash flow to invest in financial assets.

Naturally, the negative conditional sensitivity of saving to cash flow increases in absolute value with the serial correlation of productivity shocks. It also falls in absolute value as the shocks become more variable because the firm does not react strongly to the small amount of information in high-variance shocks. Finally, the sensitivity rises in absolute value with the cost of external finance because the firm's optimal *level* of cash increases with the cost of external finance. In comparison with a low-cost firm, a high-cost firm therefore has more slack with which to respond to profit shocks, and it saves or dissaves more aggressively to counteract part of the effects of these shocks.

This last result is particularly interesting because it points out that although the levels and changes in cash for a firm are clearly related, a high cash level does not necessarily imply a high, positive sensitivity of the change in cash (saving) to cash flow; nor does a low cash level imply a low sensitivity. This distinction between levels and sensitivities would be impossible to uncover, for example, in a model with a one-period saving decision because in such a setting the change in cash is indistinguishable from the level. A dynamic model such as ours is therefore essential to understanding the trade-offs that affect corporate saving.

Our empirical work is closely tied to our model. To generate exact testable predictions, we solve the model numerically and use the solution to generate a panel of simulated data. We then use these data to run a linear regression, from Almeida, et al. (2004), of the change in cash levels (saving) on

Tobin's q and cash flow. The coefficient on cash flow measures the saving sensitivity, which we also dub the propensity to save.

We then run the same regression on real data, primarily from the United States, but also from Canada, France, Germany, Japan, and the United Kingdom. Although we find positive and significant OLS coefficients on cash flow in all six countries, these coefficients become negative and significant when we correct econometrically for the substantial measurement error in Tobin's q documented in Erickson and Whited (2000, 2006) and Whited (2001). The difference between the two sets of results makes sense given the severe coefficient bias that measurement error can cause. As predicted by our model, we find that saving propensities are less negative in samples of firms with high income variability and more negative in samples with high serial correlation of income. However, we reject our model's prediction that firms typically categorized as financially constrained have more negative saving propensities than their unconstrained counterparts. This result occurs because constrained firms also have highly variable income shocks, which reduce their saving propensities. In sum, from our model and empirical evidence we conclude that the variability and autocorrelation of income shocks are at least as important as the cost of external finance in determining corporate saving. Accordingly, although the sensitivity of saving to cash flow contains information about external finance constraints, too many factors influence this one correlation for it to be used as a summary measure of the cost of external finance.

Our paper fits into both the theoretical and empirical literatures on corporate saving. The theoretical model in this paper is most closely related to that in Whited (2006), in which a firm invests and saves in the face of costly external equity finance and fixed costs of capital adjustment. We extend the model by including a corporate income tax and convex adjustment costs, and we examine empirically the model's implications for saving rather than for investment, as in Whited (2006). Our model is also closely related to the one in Eisfeldt and Rampini (2006), which characterizes the business-cycle properties of aggregate liquidity. They calibrate a general-equilibrium model with a rich specification of uncertainty. Although many of the same economic mechanisms at work in their model also operate in ours, the focus of the two papers is quite different in that we are interested in directly testing the implications of the model at the firm level, instead of calibration at the aggregate level. Another closely related theoretical paper is Gamba and Triantis (2006).

Their model is quite general, allowing for cash holding as well as separate debt and equity finance, although, unlike us, they omit physical adjustment costs. Their main contribution is an explanation of how debt flotation costs can lead to simultaneous cash and debt holdings.

Our paper is most closely related to Almeida et al. (2004). We extend their work along both theoretical and empirical dimensions. Theoretically, their model predicts a positive propensity to save—a result that occurs for two reasons. First, in their model an increase in cash flow is not accompanied by higher capital productivity. Therefore, the firm has no incentive to transform liquid assets into physical assets, as in our model, and increased cash flow produces a pure positive income effect on saving, whereby part of the additional cash flow is used to elevate both the current and future capital stock. Second, physical assets depreciate completely between periods and cannot be used to transfer resources through time. This model feature increases the value of cash holding relative to a model in which capital depreciates. Our multi-period model, in contrast, allows for cash flow shocks that may or may not be tied to productivity, as well as for a variety of capital depreciation rates. Interestingly, our model almost always predicts negative saving propensities.

Empirically, Almeida et al. (2004) find a positive sensitivity in the data and we primarily find a negative sensitivity. The difference lies in our correction for measurement error in Tobin's q . This result is puzzling in light of the argument in Almeida, et al. (2004) that using the sensitivity of cash saving (instead of physical investment) to cash flow as a measure of financing constraints is immune to the measurement error issue. They explain correctly that under the null hypothesis of no financing frictions, saving should not depend on either cash flow or Tobin's q . Therefore, saving can only be sensitive to cash flow in the presence of financial frictions. We elaborate on this argument by using measurement-error consistent estimators to examine the sign and magnitude of the cash-flow effect. As explained in detail below, and also as noted in Greene (1997, p. 440), measurement error in Tobin's q can affect the cash flow coefficient because measurement error in one regressor affects all of the coefficients in a regression if the regressors are correlated with one another. In this case they are correlated because the information about future investment opportunities contained in cash flow leads naturally to a positive correlation between Tobin's q and cash flow. Putting saving instead of investment on the left side of the regression does not eliminate this problem.

Our empirical work is related not only to Almeida, et al. (2004), but also to Khurana, Pereira,

and Martin (2006), who replicate the results in Almeida, et al. (2004) on data from several countries; Acharya, Almeida, and Campello (2007), who examine both the propensity to save out of cash flow and the propensity to issue debt; and Sufi (2006), who uses saving propensities as an explicit metric for gauging the severity of the cost of external finance. Other papers that consider saving propensities include Costa and Paz (2004), Ferrando and Pal (2006), Lin (2007), and Tang (2007).²

The paper is organized as follows. Section 1 presents the model. Section 2 describes the model simulation and its results. Section 3 describes the data, Section 4 presents the estimation procedure and results, and Section 5 concludes. Appendix A contains details concerning data construction, and Appendix B contains Monte Carlo simulations to assess the estimators' finite-sample performance.

I. A Model of Cash Holding

To motivate our empirical work, we consider a discrete-time, infinite-horizon, partial-equilibrium model of investment and saving. First we describe technology, financing, and taxation. Then we move on to a description of optimal financing policies.

A. Technology and Financing

A risk-neutral, firm uses capital, k , and variable factors of production, l , to produce output, and it faces a combination demand and productivity shock, z . Because the variable factors are costlessly adjustable, the firm's per period profit function is given by $\pi(k, z)$, in which the variable factors have already been maximized out of the problem. The profit function $\pi(k, z)$ is continuous, with $\pi(0, z) = 0$, $\pi_z(k, z) > 0$, $\pi_k(k, z) > 0$, $\pi_{kk}(k, z) < 0$, and $\lim_{k \rightarrow \infty} \pi_k(k, z) = 0$. Concavity of $\pi(k, z)$ results from decreasing returns in production, a downward sloping demand curve, or both. The shock z is observed by the producer before he makes his current period decisions. It takes values in $[\underline{z}, \bar{z}]$ and follows a first-order Markov process with transition probability $g(z', z)$, in which a prime indicates a variable in the next period; $g(z', z)$ has the Feller property.

Without loss of generality, k lies in a compact set. As in Gomes (2001), define \bar{k} as

$$(1 - \tau_c) \pi(\bar{k}, \bar{z}) - d\bar{k} \equiv 0, \tag{1}$$

²Our work is also somewhat related to recent empirical work on the determinants of the level of (as opposed to the change in) corporate cash holdings, such as Kim, Mauer, and Sherman (1998), Opler, Pinkowitz, Stulz, and Williamson (1999), Pinkowitz and Williamson (2001), Dittmar and Mahrt-Smitt (2007), Faulkender and Wang (2006), and Foley, Hartzell, Titman, and Twite (2006).

in which d is the capital depreciation rate, $0 < d < 1$, and τ_c is the corporate income tax rate. Concavity of $\pi(k, z)$ and $\lim_{k \rightarrow \infty} \pi_k(k, z) = 0$ ensure that \bar{k} is well-defined. Because $k > \bar{k}$ is not economically profitable, k lies in the interval $[0, \bar{k}]$. Compactness of the state space and continuity of $\pi(k, z)$ ensure that $\pi(k, z)$ is bounded.

Investment, I , is defined as

$$I \equiv k' - (1 - d)k. \quad (2)$$

The firm purchases and sells capital at a price of 1 and incurs adjustment costs that are given by

$$A(k, k') = ck\Phi_i + \frac{a}{2} \left(\frac{k' - (1 - d)k}{k} \right)^2 k. \quad (3)$$

The functional form of (3) is standard in the empirical investment literature, and it encompasses both fixed and smooth adjustment costs. See, for example, Cooper and Haltiwanger (2006). The first term captures the fixed component, $ck\Phi_i$, in which c is a constant, and Φ_i equals 1 if investment is nonzero, and 0 otherwise. The fixed cost is proportional to the capital stock so that the firm has no incentive to grow out of the fixed cost.³ The smooth component is captured by the second term, in which a is a constant. Although curvature of the profit function acts to smooth investment over time in the same way that the quadratic component of (3) does, we include the quadratic component to isolate the effects of smooth adjustment costs. In contrast, curvature of the profit function not only affects investment smoothing but also the relation between firm value and profit.

We now discuss financing. The firm can hold cash, p , via a riskless one-period discount bond that earns taxable interest at a rate $r(1 - \tau_c)$. For simplicity, we do not model personal interest and dividend taxes. What is important for our model is the existence of a tax penalty for saving, which is consistent with recent U.S. tax code. See Hennessy and Whited (2005). To make the choice set compact, we assume an arbitrarily high upper bound on liquid assets, \bar{p} . This upper bound is imposed without loss of generality because our taxation assumptions ensure bounded saving.

All external finance takes the form of equity. This simplification allows us to highlight the interaction between technology, finance constraints, and cash holdings. Also, having a single source of external finance does not affect the qualitative outcome of the model simulations, which only rely on a difference in the costs of external and internal funds. Differentiation among the types

³Replacing ck with a fixed number, F , changes the analysis little because the capital stock is bounded.

of external financing, although important for the study of capital structure, is not required. To preserve tractability, we do not model costs of external equity as the outcome of an asymmetric information problem. Instead, we capture adverse selection costs and underwriting fees in a reduced-form fashion. Accordingly, we define gross equity issuance/distributions as

$$e(k, k', p, p', z) \equiv (1 - \tau_c) \pi(k, z) + p - \frac{p'}{(1 + r(1 - \tau_c))} - (k' - (1 - d)k) - A(k, k'). \quad (4)$$

If $e(k, k', p, p', z) > 0$, the firm is making distributions to shareholders, and if $e(k, k', p, p', z) < 0$, the firm is issuing equity. The external equity-cost function is linear-quadratic and weakly convex:

$$\begin{aligned} \phi(e(k, k', p, p', z)) &\equiv \Phi_e \left(-\lambda_0 + \lambda_1 e(k, k', p, p', z) - \frac{1}{2} \lambda_2 e(k, k', p, p', z)^2 \right) \\ \lambda_i &\geq 0, \quad i = 0, 1, 2, \end{aligned}$$

in which Φ_e equals 1 if $e(k, p, k', p', z) < 0$, and 0 otherwise. Convexity of $\phi(e(k, p, k', p', z))$ is consistent with the evidence on underwriting fees in Altinkilic and Hansen (2000).

The firm chooses (k', p') each period to maximize the value of expected future cash flows, discounting at the opportunity cost of funds, r . The Bellman equation for the problem is

$$V(k, p, z) = \max_{k', p'} \left\{ e(k, k', p, p', z) + \phi(e(k, k', p, p', z)) + \frac{1}{1 + r} \int V(k', p', z') dg(z', z) \right\}. \quad (5)$$

The first two terms represent the excess of cash inflows over cash outflows and the third term represents the continuation value of the firm. The model satisfies the conditions for Theorem 9.6 in Stokey and Lucas (1989), which guarantees a solution for (5). Theorem 9.8 in Stokey and Lucas (1989) ensures a unique optimal policy function, $\{k', p'\} = h(k, p, z)$, if $e(k, k', p, p', z) + \phi(e(k, k', p, p', z))$ is weakly concave in its first and third arguments. This requirement puts easily verified restrictions on $\phi(\cdot)$ that are satisfied by the functional forms chosen below.

B. Optimal Financial Policies

This subsection develops the intuition behind the model by examining its optimality conditions. To simplify the exposition of optimal policies, we assume in this subsection that V is concave and once differentiable. These assumptions are not necessary for the existence of a solution to (5) or of an optimal policy function. We present optimal financial policies, heuristically, in two steps. First, we determine optimal financing under the assumption that the manager ignores the fixed costs of

external equity; that is, he treats $\lambda_0 = 0$. Second, we determine whether the intra-marginal benefits of equity issuance justify the fixed cost.

The optimal interior financial policy, obtained by solving the optimization problem (5), satisfies

$$1 + (\lambda_1 - \lambda_2 e) \Phi_e = \frac{1 + r(1 - \tau_c)}{1 + r} \int V_2(k', p', z') dg(z', z). \quad (6)$$

The right side represents the shadow value of cash balances, and the left side represents the marginal cost of external equity finance. To develop the intuition behind the optimal policy, we use the envelope condition to rewrite (6) as:

$$1 + (\lambda_1 - \lambda_2 e) \Phi_e = \frac{1 + r(1 - \tau_c)}{1 + r} \int (1 + (\lambda_1 - \lambda_2 e') \Phi'_e) dg(z', z). \quad (7)$$

Rewriting (6) as (7) makes it clear that without costly external finance, equation (7) holds as an inequality. In this case the tax penalty for saving implies that the firm never saves; i.e., $p = p' = 0$. In contrast, in the face of costly external finance, if a firm saves a dollar today, it reduces the probability of having to issue new equity tomorrow. It continues to save just to the point where the gain from reducing future equity costs outweighs the tax penalty on saving. Inspection of (6) also reveals that optimal saving policy and optimal investment policy are clearly intertwined.

In some instances the fixed costs of external equity will be larger than the intra-marginal gains from equity issuance. In these cases the firm is in a region of financial inertia in which it neither issues equity nor distributes funds to shareholders. Internal funds are the marginal source of funds and the firm saves any excess cash flows not used for positive NPV projects.

This discussion of the intuition for the value of cash in our model reveals a fundamental difference between our model and the model in Almeida, et al. (2004) that goes beyond the distinction between the finite horizon in their model and the infinite horizon in ours. In their three-period model capital depreciates completely between periods. Therefore cash is the only way to transfer resources between periods. If a firm has insufficient resources to obtain the first-best investment policy, then if it gets a small increment to its income it must save some in order to equate the marginal products of capital across time. This purely mechanical effect does not operate in our model. On a priori grounds a model with complete capital depreciation is unlikely to describe the data, given the small depreciation rates for different types of capital estimated by the Bureau of Economic Analysis. See Fraumeni (1997). We therefore retain the feature of partial capital depreciation each period.

II. Simulations

We solve the model numerically and investigate its implications for reduced-form regressions via simulation. We first describe the parameterization of our baseline simulation and explain the properties of optimal firm behavior. We then explain the experiments we perform on the model and the results of these experiments. We conclude by considering the empirical predictions given by the model and by examining the robustness of the model to our various simplifying assumptions.

A. Model Calibration

The profit function is given by $\pi(k, z) = zk^\theta$, in which we calibrate θ from the estimates of labor shares and mark-ups in Rotemberg and Woodford (1992, 1999). Their estimates, along with the assumptions of a Cobb-Douglas production function and a constant-elasticity demand function, imply that $\theta \approx 0.75$. To specify a stochastic process for the shock z , we follow Gomes (2001) and assume that z follows an $AR(1)$ in logs,

$$\ln(z') = \rho \ln(z) + v', \quad (8)$$

in which $v' \sim N(0, \sigma_v^2)$. Our baseline parameter choices for ρ and σ_v are the averages of the estimates of these two parameters in Hennessy and Whited (2007): the serial correlation of the shock, ρ , is set at 0.66 and the standard deviation of the shock, σ_v , is set at 0.121.

We again follow Hennessy and Whited (2007) to parameterize the financing function, setting $\lambda_0 = 0.389$, $\lambda_1 = 0.053$, and $\lambda_2 = 0.0002$. These settings are from their estimates of the costs of external equity finance for large firms and are therefore conservative, lying only slightly above the figures for underwriting costs in Altinkilic and Hansen (2000). The number 0.389 in the data used by Hennessy and Whited (2007) implies a fee of \$50,332 for the first million dollars of gross equity proceeds. We set the interest rate, r , equal to 4%, which lies between the values chosen by Hennessy and Whited (2007) and Gomes (2001).

To find values for the adjustment cost parameters, c and a , we turn to Cooper and Haltiwanger (2006), who find that both convex and fixed costs of adjustment affect investment. From their estimates we set $c = 0.039$ and $a = 0.049$. We set the depreciation rate equal to 0.15, a figure approximately equal to the average in our data of the ratio of depreciation to the net capital stock.

Finally, to find a numerical solution we need to specify a finite state space for the three state variables. We let the capital stock lie on the points

$$\left[\bar{k}(1-d)^{40}, \dots, \bar{k}(1-d)^{1/2}, \bar{k} \right].$$

We let the productivity shock have 25 points of support, transforming (8) into a discrete-state Markov chain using the method in Tauchen (1986). We let p have 40 equally spaced points in the interval $[0, \bar{p}]$, in which \bar{p} is set to $\bar{k}/2$. The optimal choice of p never hits this upper bound.

We solve the model via iteration on the Bellman equation, which produces the value function $V(k, p, z)$ and the policy function $\{k', p'\} = h(k, p, z)$. In the subsequent model simulation, the space for z is expanded to include 100 points, with interpolation used to find corresponding values of V , k , and p . The model simulation proceeds by taking a random draw from distribution of z' (conditional on z), and then computing $V(k, p, z)$ and $h(k, p, z)$. We use these computations to generate an artificial panel of firms by simulating the model for 10,000 identical firms over 200 time periods, keeping only the last 20 observations for each firm.

B. Simulated Policy Functions

Before presenting our simulation results, we examine the economics behind the model by exploring the simulated policy function, $\{k', p'\} = h(k, p, z)$. We do so by plotting optimal cash flow, investment (net of adjustment costs), saving, and distributions/equity issuance (net of issuance costs) as a function of z for three different (k, p) pairs: low k /medium p , medium k /medium p , and high k /medium p . By high, medium, and low we mean the maximum, median, and minimum values that k and p take in the baseline simulation. We focus on the medium p case for brevity, but we discuss the high and low p cases briefly below. Cash flow is defined precisely as $(1 - \tau_c) \pi(k, z) / k^*$, investment as $((k' - (1 - d)k) - A(k, k')) / k^*$, saving as $(p' / (1 + r(1 - \tau_c)) - p) / k^*$, and net distributions/equity issuance as $(e(k, k', p, p', z) + \phi(e(k, k', p, p', z))) / k^*$, in which k^* is the steady-state level of the capital stock. We deflate our variables of interest by k^* for the three differently sized firms to facilitate comparisons between them.

Figure 1 contains these plots, with the three panels depicting, respectively, a small, a medium, and a large firm. In all three panels cash flow naturally rises with the z shock. These cash flows are, however, distributed differently depending on the size of the firm.

For the small firm investment rises smoothly with cash flow. Despite the presence of adjustment costs, the capital stock is so low and the marginal product of capital so high that a higher value of z almost always means more investment.

In contrast, the behavior of saving is non-monotonic. Although the small firm always saves, saving initially rises with z and then falls. This hump-shaped pattern reflects income and substitution effects. To define these effects, we note that when z rises, both capital productivity and cash flow rise; that is, the firm's value function both becomes steeper and shifts upward. We define the income effect as $\partial p/\partial z$ along the surface where $V_1(k, p, z)$ and $V_2(k, p, z)$ are unchanged. In words, the income effect captures the upward shift and thereby isolates the effect of the extra income generated by z , holding current and future productivity of capital and cash constant. We define the substitution effect as $\partial p/\partial z$ minus the income effect. The substitution effect captures the changed steepness of the value function and thereby isolates the effects of current and future changes in the marginal productivities of cash and capital.

We now use these definitions to explain the hump-shaped pattern. As z rises, capital productivity rises and the firm expects productivity to revert to its mean slowly because of the positive serial correlation in z . The substitution effect implies that the firm saves less because it wants to shift some of its liquid assets into physical assets that have become relatively more productive. The income effect implies that the firm saves more as it increases both physical and liquid assets, both of which continue to have value as z rises. The income effect dominates for low levels of z , but its strength is limited, and the substitution effect dominates for high levels of z .

In the model distributions/equity issuance are a residual. For a small firm the marginal product of capital is sufficiently high that it is optimal for the firm to issue equity and to pay issuance costs regardless of the level of the productivity shock.

The medium-sized firm behaves quite differently. First, the optimal investment rule displays substantial inertia. For low levels of z the firm sells capital, but for intermediate and high levels of z the firm invests. Investment initially rises with z , but then flattens out, rising once again when z is high. Physical adjustment costs cause the stagnation, and the fixed equity issuance cost causes the jump, as the firm goes from making distributions to raising equity. Saving also behaves differently in the medium-sized firm. Saving is always negative and always decreases with z and with cash

flow because the substitution effect always dominates the income effect. This negative correlation between cash flow and saving is crucial for understanding the saving sensitivity results that follow. Finally, for low levels of z , the firm finds it optimal to distribute excess funds to shareholders because the benefits of investing do not outweigh the costs of issuance. However, if z rises to a sufficiently high level, the benefits from investing start to outweigh issuance costs, and the firm issues equity.

The large firm, not surprisingly, sells capital for low to intermediate levels of z because the marginal product of capital is low. Although investment eventually becomes positive as z rises, the presence of adjustment costs combined with the low marginal product of capital cause the rate of investment to level out for very high levels of z . Saving initially declines with z because of the substitution effect, which operates even though the firm is disinvesting because the marginal product of capital always rises with z regardless of optimal investment policy. However, dissaving flattens for high levels of z because the model does not allow for negative cash (i.e. debt).

For brevity we have not plotted the optimal saving and investment rules for firms with either low or high cash balances. In general, the same effects operate, although two further patterns are of interest. First, the income effect on saving dominates for firms with low cash balances. This result makes sense because the value of cash is especially high when cash balances are low. Therefore, the substitution effect that arises from the relative increase in capital productivity is small. Second, and conversely, the substitution effect on saving dominates for firms with larger cash balances.

C. Experiments

With the model intuition in hand we now turn to our simulation results. We investigate two ways in which the model’s parameters affect the firm’s cash and saving policies. We first consider how the parameters affect the level of cash as a fraction of assets, which is defined in our model as the average of p/k over all of the observations in the simulated panel. We then examine how the parameters affect a measure of saving behavior that first appears in Almeida, et al. (2004). Dubbed “the cash-flow sensitivity of cash,” this measure is defined in our model as the regression coefficient, α_1 , in the following regression:

$$\frac{p' - p}{k} = \alpha_0 + \beta \frac{V(k, p, z)}{k} + \alpha_1 \frac{\pi(k, z)}{k} + \alpha_2 \ln(k) + u, \quad (9)$$

in which α_0 , α_1 , α_2 , and β are regression coefficients and u is a regression disturbance, which in our simulations is, by definition, orthogonal to the regressors.⁴ This regression comes directly from Almeida, et al. (2004), and we estimate it with all of the observations in the simulated panel. The intent is to understand better the economics behind this reduced form regression.

In thinking about the results that follow, it is crucial to separate cash levels (p) and saving ($p' - p$) from the saving sensitivity (α_1), which is just the partial correlation between cash flow and saving, holding q and k constant. For example, a negative sensitivity (partial correlation) does not imply that $p' - p < 0$, that is, that the firm always dissaves; nor does $\alpha_1 < 0$ imply that the simple correlation between cash flow and saving is negative. Indeed, in most of our simulations, average saving as a fraction of k is small and positive, and the simple correlation between cash flow and saving is positive, even though $\alpha_1 < 0$.

We examine the sensitivity of our two gauges of cash policy (p/k and α_1) to eight key model parameters: the standard deviation and serial correlation of log profits, σ_v^2 and ρ ; the three equity-cost parameters, λ_0 , λ_1 , and λ_2 ; the curvature of the profit function, θ ; and the fixed and quadratic adjustment cost parameters, c and a . In each of the following experiments, we set all but one of the parameters equal to their baseline levels, allowing the free parameter to range within a given interval. We allow θ to range from 0.6 to 0.9, ρ from -0.8 to 0.8, σ_v from 0.05 to 0.2, λ_0 from 0 to 0.8, λ_1 from 0 to 0.1, λ_2 from 0 to 0.0004, c from 0 to 0.8, and a from 0 to 0.1.

Figure 2 illustrates the sensitivity of the optimal choice of p/k to the model parameters. We first examine the parameters that govern the stochastic shock process. The first panel shows a u-shaped relation between the serial correlation of income, ρ , and cash holdings. For both highly positively and highly negatively correlated shocks, the firm holds high cash balances, choosing lower balances if the shocks are less highly correlated. Two separate effects explain this result. First, as ρ increases, the firm invests in larger amounts because a positive productivity shock signals not only that capital is productive today, but also that it will continue to be productive. The firm therefore wants higher cash balances to lower the probability of needing external finance when it makes these large investments. Second, the higher the serial correlation of an $AR(1)$ process, the higher its variance. If the firm faces an uncertain environment, it expects to tap external finance more often,

⁴Deflating the variables in (9) by $(k + p)$ instead of by k changes the results little.

and it holds higher cash balances. Both effects operate in the same direction for high levels of ρ , but they to offset each other for levels near zero. For levels of ρ far below zero, the second effect dominates.⁵ The intuition about the effect of uncertainty is also evident in the second panel, which depicts a positive relation between cash holdings and σ_v , the standard deviation of the innovations to $\ln(z)$. The increase in cash accompanying an increase in σ_v also has a real options interpretation in which a higher variance leads to a higher option value of cash balances.

The third through fifth panels illustrate the effects of each of the external finance parameters on cash holdings. Not surprisingly, the third and fourth panels show that cash increases with the fixed and linear components of the external finance function, λ_0 and λ_1 , because the value of financial flexibility increases as external finance becomes more costly. However, the relation shown in the fifth panel between the quadratic component, λ_2 , and cash holdings is flat. With λ_0 and λ_1 set to their baseline levels, the effect of λ_2 is second-order.⁶ These results mirror those in the three-period model of Almeida, et al. (2004), which produces a partial derivative of cash with respect to internal funds that is positive for a financially constrained firm, and zero otherwise.

Finally, the sixth through eighth panels display the effects of technology. The sixth panel reveals a hump-shaped relation between production function curvature (θ) and cash. Two different economic forces create this pattern. First, as θ rises, the production function becomes flatter, and the average size of desired investments rises. The firm holds more cash because large investments imply a greater likelihood of needing external finance. Second, as θ rises, the firm is less likely to have to tap external finance because a higher θ implies that a given capital stock can create more internal revenue, and the firm therefore needs to hold less cash. The first effect is stronger for lower levels of θ , and the second effect is stronger for higher levels of θ . The seventh panel shows that cash holding increases with the fixed cost of adjustment. This effect occurs because higher fixed adjustment costs lead to larger investments that occur less frequently. The firm then uses episodes of inaction to accumulate cash, which acts to lower the probability of the firm having to tap external finance when it does invest. Finally, the eighth panel shows that convex adjustment costs have the

⁵Simulations in which ρ increases but the variance of the process is held constant produces a similar result, except that the rise in cash holdings for very low ρ flattens out.

⁶This last result is not an artifact of the presence of quadratic physical adjustment in the model because turning off these costs has no effect on this result.

opposite effect on cash holding. As a increases, the firm makes smaller investment more often, is therefore less likely to have to tap external finance, and holds less cash.

These results on the level of cash balances reassuringly confirm those in Gamba and Triantis (2006), in particular their results on the effects of uncertainty and the cost of external finance. Our results on cash levels are also useful in providing intuition for the main focus of this paper, which is not cash levels, per se, but the propensity to save.

Figure 3 is analogous to Figure 2, except that it depicts how the model parameters affect our second measure of cash policy—the coefficient α_1 in (9); that is, the conditional sensitivity of saving to cash flow. A quick glance at the figure reveals that for almost all model parameterizations this sensitivity is negative. Because this result is the opposite of that produced by the model in Almeida et al. (2004), it is worthwhile to examine the reason for the difference. The answer partly lies in the difference between income and substitution effects. In our model, when a firm receives a positive income shock, its cash flow rises. In addition, if the shock is not too transitory, both the current and future productivity of capital rise. A substitution effect implies that the firm wants to transform its financial assets into relatively more productive physical capital, that is, dissave. An income effect implies that the firm wants to use some of its increased cash flow to increase cash balances, thereby lowering the probability that it will have to use costly external funds to finance future investment. In a regression of saving on q and cash flow, both q and cash flow contain information about capital productivity. Because cash flow, $\pi(k, z)$, is driven in large part by z and because, as demonstrated in Figure 1, z drives the substitution effect, the negative coefficient on cash flow reflects the negative substitution effect. The coefficient on q gives the partial effect on savings of changes in q , holding $\pi(k, z)$ and k constant. It must, therefore, be driven by changes in the only remaining state variable, which is the current cash level, p . In particular, q captures the value of both physical assets and cash, whereas cash flow only captures the value of physical assets. We therefore conjecture that q reflects the increased value of cash following a positive productivity shock and therefore captures the income effect. In contrast, in the Almeida, et al. (2004) model the sensitivity of saving to cash flow is defined as the response of cash holding to an exogenous increase in the firm’s endowment, which does not affect capital productivity. Therefore, their model produces a positive sensitivity in part because it only captures an income effect. As mentioned above, their model also produces a

positive sensitivity because capital depreciates completely between periods.

We now turn to a more detailed discussion of Figure 3. The first panel shows our most interesting simulation result, which is the effect of ρ on α_1 . If the serial correlation of the shock process is highly negative, this sensitivity is large and positive; if ρ is highly positive, it is large and negative; and if ρ is near zero, the sensitivity is near zero. This pattern arises out of the firm's expectation about future needs to tap external finance and about the current and future productivity of capital. If profits are negatively serially correlated, then a positive shock implies an expected productivity decline, which in turn implies a low need for external finance. This income effect promotes dissaving. A stronger substitution effect, however, promotes saving because the expected productivity decline prompts the firm to funnel cash flow into liquid assets and distributions rather than into investment. On the other hand, if the firm faces highly positively correlated income shocks, then the income and substitution effects operate in the opposite direction; and, conditional on q , the firm dissaves when it experiences a positive shock.

The second panel illustrates the effect of the shock variance. The saving sensitivity is always negative, but becomes less so as σ_v increases. As the firm's environment becomes more uncertain, its level of cash increases, but it also becomes more reluctant to change its cash holdings aggressively in response to shocks, which convey little information in an uncertain environment. Finally, because α_1 depends on the variability and autocorrelation of z , any observed cross-sectional variability in σ_v and especially in ρ renders α_1 a poor financial constraint indicator.

The third through fifth panels examine the effects of the cost of external finance. The patterns evident in these panels mirror those in the corresponding panels in Figure 2. In all cases, as the cost of external finance increases, the level of cash increases, and the sensitivity of saving to cash flow becomes more negative. Because a firm with a high cost of external finance optimally holds a high level of cash, it can respond to shocks more aggressively by changing its cash balances. For example, if a positive profit shock hits a firm with a high level of cash, it will dissave a great deal in order to invest. An otherwise identical firm with a low level of cash cannot dissave as much.

Finally, we examine technology. The saving sensitivity becomes more negative as θ increases; that is, as the production function becomes flatter. With a flat production function positive shocks induce large desired increases in the capital stock, and the firm dissaves to fund these investments.

The saving sensitivity also becomes more negative as fixed adjustment costs increase because the firm optimally invests in large amounts. Similarly, the saving sensitivity becomes less negative with the quadratic adjustment cost because the firm optimally invests in small amounts.

The preceding arguments are valid at points in time in which the firm is actively adjusting its capital stock. During periods of inaction, the sensitivity of saving to cash flow is positive because the firm funnels at least part of its cash flow into cash holdings in order to avoid tapping external finance in the future. Under almost all parameterizations of this model the firm adjusts more often than it remains inactive. The observations in which the saving sensitivity is negative therefore outweigh those in which it is positive, and average sensitivity is negative.

The frequent adjustment in our model sets it apart from models of dynamic capital structure with adjustment costs. For example, in Fischer, Heinkel, and Zechner (1989), the firm adjusts its asset and liability composition infrequently. As pointed out in Strebulaev (2006), empirical predictions from this sort of model cannot be based on firm behavior at points in time at which the firm is inactive. The frequent adjustment in our model allows us to sidestep this critique. Frequent adjustment also explains why corporate propensities to save can be negative even though personal propensities to save are typically positive. Although consumers dissave when they purchase durables, these events are infrequent, and because consumers save out of income at other times, average personal saving propensities are therefore positive.

Because we are interested in the effect of measurement error in observed Tobin's q in our data, we conduct a further simulation in which we introduce an additive *i.i.d.* measurement error to $V(k, p, z)/k$. Measurement error biases the positive coefficient on $V(k, p, z)/k$ downward, but it biases the coefficient on $\pi(k, z)/k$ upward because of the strong positive correlation between $\pi(k, z)/k$ and $V(k, p, z)/k$. We find that the baseline simulation requires a great deal of measurement error to reverse the initially negative sign of the coefficient on cash flow. The error variance needs to be at least eight times as large as the variance of $V(k, p, z)/k$. This result is not out of line with the empirical results that follow inasmuch as we find that the measurement quality of observed Tobin's q is extremely low, or equivalently, that the measurement error variance is high.

In sum, these experiments highlight three important pieces of economic intuition. First, corporate saving depends not only on the firm's financial environment, but also on its technology.

Second, variation in capital productivity is critical for our results, because a model cannot capture the firm’s desire to substitute capital for cash if the marginal product of capital is constant. Third, although the levels and changes in cash are related, a high cash level does not necessarily imply a high positive sensitivity of saving to cash flow; nor does a low cash level imply a low sensitivity. We emphasize again that this distinction is impossible to uncover in a model with a one-period saving decision because the change in cash cannot be distinguished from the level of cash.

D. Empirical Predictions

The simulations in Figure 3 delineate the four central empirical predictions we test. First, the sign of α_1 in the regression (9) should be negative. Second, α_1 should increase in absolute value with the cost of external finance. Third, α_1 should decrease in absolute value with σ_v . Fourth, α_1 should increase in absolute value with ρ . We consider the relation between production function curvature and α_1 only in a robustness check because available proxies for curvature are weak, and we do not directly test any predictions concerning the relation between α_1 and adjustment costs because adjustment costs are unobservable. The simulations that examine curvature and adjustment costs do, however, provide intuition that assists with the interpretation of some of our results.

Testing these predictions in this manner forms a strong link between the theory and its test, because the form of the real-data test is identical to the form of the simulated-data theoretical prediction. Further, because the prediction encompasses the entire regression specification, and because the error term in a linear regression (linear projection) is by definition orthogonal to the right-hand-side variables, testing whether $\alpha_1 < 0$ in this manner avoids the usual simultaneity problems that plague regressions in corporate finance. This type of prediction also has precedents in Whited (2006) and Caggese (2006). Finally, this type of a link between theory and its tests is fundamentally different from testing model predictions that take the form of the sign of a partial derivative, because partial derivatives provide incomplete guidance on the entire regression specification.

E. Model Robustness

The model is intentionally sparse to highlight intuition. To assuage concerns that our results are artifacts of the model’s simplicity, in this section we add several more realistic features to the model to examine the robustness of our result of a negative saving propensity.

Our first set of robustness checks focuses on specific features omitted from our baseline model. First, in the baseline model the firm does not have access to a credit line. When we add riskless short-term debt that is secured by the capital stock, as in Hennessy and Whited (2005), the saving propensity of -0.40 in the baseline simulation drops in absolute value to -0.22. Our results are attenuated but not erased because the upper limit to the credit line causes cash to retain its value as a tool to avoid costly external finance. Second, the baseline firm does not smooth distributions to shareholders. To address this possibility, we penalize the firm by the amount of the linear equity issuance cost for every dollar that its distributions fall below the average level of distributions in the baseline simulation. This model feature produces increased cash hoarding because the firm wants to avoid missing a distribution. This higher cash cushion leads to a more negative propensity to save of -0.54. Third, the baseline firm has no fixed costs of production, which could, for example, represent the tendency of young firms to burn profits. We add a cost of production equal to $0.4k$, in which 0.4 is the approximate ratio of selling, general, and administrative expenses to assets in our U.S. sample. This addition to the model produces less cash holding relative to our baseline model because the firm has smaller profits to funnel into liquid assets. Accordingly, the saving propensity drops in absolute value to -0.12. Fourth, we allow the firm to issue risky debt, which we model exactly as in Hennessy and Whited (2007). This approach necessitates the exclusion of any physical adjustment costs. In this case the firm hoards more cash than in the baseline simulation in order to avoid default, and the saving propensity rises in absolute value to -0.59. We choose to exclude these features from our baseline model because they do not change any qualitative simulation outcomes, and a simple structure is important for lending intuition to our empirical work.

Finally, because our model contains only one source of uncertainty, productivity shocks and cash flow are almost perfectly correlated. Therefore, the dissaving that occurs with a positive productivity shock is necessarily accompanied by a rise in cash flow. To ascertain whether our finding of a negative saving propensity is hard-wired by this feature of our model, we allow the net revenue function to take the form $zk^\theta - \eta k^*$, in which k^* is the steady state capital stock for a frictionless version of this model and η is a normally distributed, zero-mean, *i.i.d.* random variable with a variance equal to that of the z shock. This new cost shock takes four points of support, and its transition matrix is given by the method in Tauchen (1986). Not surprisingly, decoupling

cash flows from the productivity shock produces a much smaller (in absolute value) propensity to save. However, unless the variance of the cost shocks is twice the variance of the productivity shock, the saving propensity remains negative. Further, even with a positive saving propensity, the serial correlation and variance of both types of income shocks still have important effects on the saving propensity. Finally, Smets and Wouters (2007) estimate on a macroeconomic level that the variances of cost and productivity shocks are comparable in magnitude. Although this result indicates that a low correlation between cash flow and productivity is unlikely to be a pervasive phenomenon, cross-sectional heterogeneity in this correlation remains an empirical issue that we explore below.

The second set of robustness checks relates to the connection between estimating (9) with simulated data and with real data. The cross sections generated by the model contain 10,000 identical firms over 20 time periods. This simulated cross-section asymptotically generates the same results as a single time series with 200,000 observations. In contrast, our real data contain heterogeneous firms. Therefore, unless a simulated panel of heterogeneous firms can generate a negative sensitivity, the connection between the theory and its tests becomes tenuous.

Adding heterogeneity to the simulated sample should produce a positive sensitivity only when most of the simulated firms come from an environment that generates a positive sensitivity. The predominantly negative sensitivities seen in Figure 3 indicate that a positive sensitivity is only likely to arise in a panel that is heterogeneous along the lines of the serial correlation parameter, ρ . Indeed, we find negative sensitivities when we add heterogeneity by varying the shock standard deviation (σ), issuance costs ($\lambda_0, \lambda_1, \lambda_2$), returns to scale (θ), and fixed and smooth adjustment costs (c, a). In contrast, when we divide the cross section into 10 groups with values of ρ equally spaced between -0.8 and 0.8, we find a small positive sensitivity of 0.008. This result begs the question of the cross-sectional distribution of the serial correlation of income in our real data. To answer the question, we use our U.S. data to estimate a first-order autoregression of operating income either firm-by-firm or industry-by-industry, in which an industry is defined at the three-digit level. In the firm-by-firm autoregressions we find that only 7.1% of our firms have negatively serially correlated income. In the industry-by-industry autoregressions we never find negatively serially correlated income. To add heterogeneity in serial correlation that approximates the situation in our data set, we rerun our simulation with the cross section divided into 10 groups with values of ρ that approximate the

firm-level cross-sectional distribution of income serial correlation in our real data. In this case we do find a negative sensitivity.

It is also interesting to see whether we can generate a positive sensitivity when we add heterogeneity to the sample by varying parameters not examined in Figure 3. We look at the discount factor (β), the rate of capital depreciation (d), and the drift of the shock process, the latter of which we model by adding an intercept to (8). In none of these cases are we able to generate a positive sensitivity. Decreasing the depreciation rate or increasing the discount rate reduces the average size of investments, the need for external finance, and cash levels. The saving propensity remains negative but decreases in absolute value. Increasing either the depreciation rate or the discount factor produces the opposite effects. Changing the drift of the z process alters the average size of investments but leaves the saving propensity negative. As a final note, we can generate a positive sensitivity by turning off the smooth adjustment costs and multiplying the fixed adjustment costs by a factor of 10. In this simulation the firm invests sporadically in large spikes, accumulating cash during the periods of inactivity.

In sum, our result of a negative propensity to save is remarkably robust. We only find three instances in which we can generate a positive saving propensity: with negatively serially correlated productivity shocks, with high-variance cash-flow shocks that are decoupled from productivity shocks, and with high fixed costs of adjustment. We have provided evidence that the first situation is not relevant empirically, and we examine the empirical relevance of the other two situations below.

III. Data and Summary Statistics

We obtain data on U.S. nonfinancial firms from the 2007 Standard and Poor's Compustat industrial files. These data constitute an unbalanced panel that covers 1972 to 2006. We also draw data from Standard and Poor's Compustat Global Issue and Industrial/Commercial for five more countries: Canada, France, Germany, Japan, and the United Kingdom. These data also constitute an unbalanced panel but only cover 1994-2005 because Global Vantage does not report data on all firms for 2006. A description of the regression variables is in Appendix A.

Summary statistics are in Table 1. We see large differences in most instances between the means and medians of Tobin's q (market to book). As explained below, this skewness is essential

for identifying our econometric model. The average ratio of cash to assets ranges from 0.09 to 0.16. These averages hides a sharp increase in this ratio in the last 10 years. For example, in the United States the median has doubled over this time period. The average change in the cash stock is negative in the United Kingdom, Japan, France and Germany, but it is positive in the United States and Canada. The average number of observations per year is in the last column. Both the United States and Japan have on average more than 2,500 observations per year, whereas Canada, France, and Germany each have fewer than 400. Because our estimation method requires a great deal of data, the variation in sample size across countries manifests itself in the precision of the estimates we obtain from each country.

IV. Estimation

In our model we can perfectly observe $V(k, p, z)/k$. In contrast, in our data we use Tobin's q as a proxy for $V(k, p, z)/k$. As explained in Erickson and Whited (2000), Tobin's q is an imperfect proxy, and we must therefore treat the ensuing measurement error. This section outlines our the method we use. It then presents the results from applying this technique to the data.

A. Methodology

We use the estimators in Erickson and Whited (2000, 2002), which employ the structure of the classical errors-in-variables model. Applied to a single cross section, this model can be written as

$$y_i = w_i\alpha + \chi_i\beta + u_i, \tag{10}$$

$$x_i = \gamma + \chi_i + \varepsilon_i. \tag{11}$$

In our application y_i is the ratio of the change in cash to assets, χ_i is the true q of firm i , x_i is an estimate of its true q , and w_i is a row vector of perfectly measured regressors, whose first entry is 1, whose second entry is the ratio of cash flow to assets, and whose third entry is the natural log of total assets. The regression error, u_i , and the measurement error, ε_i , are assumed to be independent of each other and of (w_i, χ_i) , and the observations, $(\varepsilon_i, u_i, w_i, \chi_i)$, $i = 1, \dots, n$, are *i.i.d.* The intercept, γ , in (11) allows for systematic bias in the measurement of χ_i . We do not require any assumptions about the temporal dependence or independence of $(\chi_i, w_i, u_i, \varepsilon_i)$.

To derive a set of tractable moment conditions to be estimated by GMM, we reexpress (10) and (11) in terms of the residuals from the regressions of y_i , x_i , and χ_i on w_i . Let $(\dot{y}_i, \dot{x}_i, \dot{\chi}_i)$ be the residuals from the linear projection of (y_i, x_i, χ_i) on w_i . Then (10) and (11) can be written as

$$\dot{y}_i = \beta \dot{\chi}_i + u_i \quad (12)$$

$$\dot{x}_i = \dot{\chi}_i + \varepsilon_i. \quad (13)$$

If we square (12), multiply the result by (13), and take expectations of both sides, we obtain

$$E(\dot{y}_i^2 \dot{x}_i) = \beta^2 E(\dot{\chi}_i^3). \quad (14)$$

Analogously, if we square (13), multiply the result by (12), and take expectations of both sides, we obtain

$$E(\dot{y}_i \dot{x}_i^2) = \beta E(\dot{\chi}_i^3). \quad (15)$$

As shown in Geary (1942), if $\beta \neq 0$ and $E(\dot{\chi}_i^3) \neq 0$, dividing (14) by (15) produces a consistent estimator for β that equals $E(\dot{y}_i^2 \dot{x}_i) / E(\dot{y}_i \dot{x}_i^2)$. The assumptions, $\beta \neq 0$ and $E(\dot{\chi}_i^3) \neq 0$, are necessary for identification because one cannot divide by zero. These assumptions can be tested via the null hypothesis that $E(\dot{y}_i^2 \dot{x}_i) = 0$ and $E(\dot{y}_i \dot{x}_i^2) = 0$. We refer to this test hereafter as an identification test. It is a useful regression diagnostic inasmuch as it provides information as to whether the coefficient estimates are reliable. For example, if $\dot{\chi}_i$ is near normally distributed, the identification test will not produce a rejection and the coefficient standard errors will be large. It is worth noting that we do not use this test to select our samples.

This estimator is a third-order moment estimator. The innovation in Erickson and Whited (2000, 2002) consists of combining the information in moment equations of orders 2 through 7 via GMM to obtain a more efficient estimator for β . It is possible to estimate many interesting quantities besides β . For example, the coefficient of determination (R^2) of (11), denoted as τ^2 , is given by:

$$\tau^2 = \frac{\mu'_x \text{var}(w_i) \mu_x + E(\dot{\chi}_i^2)}{\mu'_x \text{var}(w_i) \mu_x + E(\dot{\chi}_i^2) + E(\varepsilon_i^2)}. \quad (16)$$

It is a useful index of measurement quality for our proxy for unobservable $V(k, p, z) / k$. A value close to one indicates a nearly perfect proxy, and a value close to zero indicates a nearly worthless proxy. An exactly analogous formula provides the measurement-error consistent estimate of the R^2

of (10). We can also estimate the coefficient vector α , which can be recovered by the identity

$$\alpha = \mu_y - \beta\mu_x, \tag{17}$$

in which (μ_y, μ_x) are the vectors of coefficients in the population projection of (y_i, x_i) on w_i . This identity is useful for understanding why measurement error in Tobin's q biases the cash-flow coefficient even when saving is on the left side of the regression. To simplify the explanation, we isolate α_1 , the cash-flow coefficient, and rewrite the second element of the vector equation (17) as

$$\alpha_1 \equiv \mu_{1y} - \beta\mu_{1x}. \tag{18}$$

The first term in (18), μ_{1y} , is the coefficient on cash flow obtained by regressing saving on only cash flow and firm size. The second term represents the extent to which μ_{1y} changes when one controls for true unobservable Tobin's q . This term clarifies how measurement error in Tobin's q biases the cash flow coefficient. It is well-known that measurement error biases β downward, and the amount of bias is approximately proportional to τ^2 . If $\mu_{1x} = 0$, then this downward bias has no effect on the coefficient on cash flow. However, if $\mu_{1x} \neq 0$, then a downward biased β can affect the cash-flow coefficient, α_1 . Recall that μ_{1x} is the slope coefficient on cash flow one obtains from regressing observable Tobin's q on cash flow and size. Because Tobin's q and cash flow are positively correlated, and because the variance of Tobin's q is much greater than the variance of cash flow, μ_{1x} can be large. In our application it ranges from 1 to 5. Therefore, a small downward bias in β can cause a large upward bias in the OLS estimate of the cash flow coefficient. Finally, it is crucial to note that the bias in the coefficient α_1 does not depend on the correlation between the left-hand-side variable and the right-hand-side variables. Instead, what matters is both the measurement quality of Tobin's q and the covariance matrix of the regressors. Therefore, measurement error in Tobin's q can bias other regression coefficients regardless of the left-hand-side variable.

Appendix B presents Monte Carlo simulations to assess the finite sample performance of these estimators on data closely resembling our own. Of particular interest in these Monte Carlos are the tests of the null hypothesis that the coefficient on cash flow equals its true value. The actual sizes of many of these tests are tiny relative to their nominal sizes. This result indicates that the finite-sample distribution of the GMM estimates of the t-statistics has thin tails.

Because these estimators can only be applied to samples that are arguably *i.i.d.*, we estimate (10) and (11) for each cross section of our unbalanced panel and then pool the yearly estimates via the procedure in Fama and MacBeth (1973). We do not include firm fixed effects in our regressions for four reasons. First, when we compare the results from running fixed-effects OLS to those from using OLS with the Fama-MacBeth approach, we find almost identical coefficient estimates. Second, when we do a standard Hausman test to determine whether a potential fixed effect is correlated with the regressors, we cannot reject the null of no correlation. This result suggests that the within-firm variation in saving and q mirrors the cross-sectional variation. Third, the resulting model almost never passes the GMM identification test because removing fixed effects removes a great deal of data variation and therefore a great deal of the skewness and kurtosis that identify the slope coefficients. For example, the inclusion of fixed effects reduces the coefficient of skewness by more than half. Fourth, the GMM estimates are qualitatively similar, but more unstable, when we include fixed effects. In sum, although it is somewhat unconventional not to control for fixed effects when one has panel data available, we do not because it affects our results little and because we do not wish to use up valuable degrees of freedom or to remove interesting data variation.⁷

Recently, Petersen (2007) has reemphasized that Fama-MacBeth standard errors can produce inflated t-statistics in panel data. Further, our lack of restrictions on the time series properties of $(\chi_i, u_i, \varepsilon_i)$ opens the door for the finite-sample critical values the Fama-MacBeth t-statistics to be much higher than the nominal critical values. We deal with this issue by using the bootstrap in Hall and Horowitz (1996) to calculate the finite-sample distribution of the t-statistics produced with the Fama-MacBeth standard errors. The unit of observation for resampling is the firm. Interestingly, we find that many of these bootstrapped critical values are only slightly higher than the asymptotic critical values, although in several instances we do find bootstrapped critical values as high as 5 for a nominal 5% two-sided t-test, especially in the case of the GMM estimates of the coefficient on χ_i .

B. Results

Table 2 presents the Fama-MacBeth results from estimating (10) via OLS and from estimating (10) and (11) via GMM for each of our six countries. The left panel shows the OLS results, and the

⁷To treat possible within-industry dependence, we have also included two-digit industry dummies, with no qualitative difference in our results.

right panel shows the GMM results. We report the OLS estimate of the regression R^2 in column 3, and the measurement-error consistent GMM estimate of the regression R^2 in column 6. Column 7 contains the estimate of our index of measurement quality: τ^2 . Asymptotic standard errors are in parentheses below each parameter estimate. Asterisks and daggers mark the parameter estimates whose t-statistics exceed, respectively, the 5% bootstrapped and 5% asymptotic critical values.

For each country we test our first prediction that the coefficient on cash flow is negative in a regression of saving on q , cash flow, and size. Our OLS results corroborate earlier findings that the coefficients on both Tobin's q and cash flow are positive for all countries, as in, e.g., Almeida, et al. (2004) and Khurana, et al. (2006).⁸ However, when we apply the Erickson and Whited estimators to correct for measurement error in q , the results change. We find negative coefficients on cash flow in all six countries, and half of the cash-flow coefficients are statistically significant according to our bootstrapped critical values. These results correspond to our model simulation results for firms that have positively serially correlated income processes. The effect of treating measurement error can also be seen in the GMM coefficients on q , which are from seven to fifteen times as high as their OLS counterparts in the different countries. This result can be explained by the attenuation bias in the classical errors-in-variables model, which in this case is very large because of the low estimates for τ^2 ; that is, low estimated measurement quality of observed Tobin's q . Finally, correcting for measurement error increases the regression R^2 substantially because measurement error obscures the contribution of true Tobin's q to the variation in saving.

The flip in the sign of the cash-flow coefficient can be understood as follows. On average in our regressions μ_{1y} is about 0.15, and μ_{1x} is about 2. The biased OLS estimates of β , which hover around 0.02, therefore produce a positive cash-flow coefficient when plugged into the identity, $\alpha_1 \equiv \mu_{1y} - \beta\mu_{1x}$, with these values for μ_{1y} and μ_{1x} . In contrast, the consistent GMM estimate of β , which hovers around 0.2, produces a negative cash-flow coefficient when plugged into this identity.

What is the intuition behind this econometric result? First, because μ_y is approximately the simple correlation between saving and cash flow, and because the OLS estimate of β is biased downward severely, the OLS estimate of α_1 only picks up this simple correlation. A positive OLS

⁸Our OLS estimates of the coefficients on q and cash flow are larger than those in Almeida, et al. (2004) because we use trimmed data and they do not. When we do not trim out data, we get OLS results that are similar to theirs.

coefficient makes sense in that on average companies should save part of cash flow shocks and then invest the rest or return it to shareholders. Conversely, our negative GMM coefficients do not imply that firms dissave more than the amount of a cash flow shock, because these coefficients only pick up the partial correlation between saving and cash flow, Tobin's q and size held constant.

Our simulations provide further insight into the GMM results. Recall that in the regressions on simulated data, Tobin's q in part reflects the value of cash to the firm and therefore is likely to pick up the income effect. If one uses a noisy proxy for true Tobin's q and therefore controls only for part of its variation, cash flow ends up picking up this income effect. Although the model provides some intuition for the result, it does not provide a complete explanation, because it is difficult to believe that the income effect on saving is extremely strong. However, Tobin's q picks up not only the income effect but also other reasons for saving, such as changes in agency issues, that are capitalized by the stock market. Clearly, not controlling for these motivations for saving causes cash flow to pick them up and contributes to making its OLS coefficient positive.

We now turn to the reliability of these regression results. Table 3 presents four summary statistics from the yearly regressions that underlie the Fama-MacBeth estimates in Table 2: the fraction of the years with negative cash-flow coefficients, with cash-flow coefficients that are significantly negative at the 5% level, in which the overidentifying restrictions of the model are rejected at the 5% level, and in which the null of no model identification is rejected at the 5% level. In nearly all years for all countries we find negative cash flow coefficients, and in more than 60% of the years the coefficients are significant. Sample size affects our ability to find statistical significance when using the GMM estimator because a great deal of data is required to estimate high order moments with precision. As is common when using international data, in some countries—particularly Canada, France, and Germany—we have smaller samples than we would like. Another difficulty that leaves us with insignificant coefficients is our frequent failure to reject the null of an unidentified model in three of the six countries: Canada, the United Kingdom, and Germany. Both of these problems tend to manifest themselves in high standard errors. Given these difficulties, our finding of a high incidence of significant coefficients is all the more striking.

Finally, Table 3 shows that we fail to reject the overidentifying restrictions from our yearly GMM estimates in all countries for most years. This result mitigates concerns about possible model

misspecification. For example, although u_i is by construction uncorrelated with (χ_i, w_i) , it may not be independent of (χ_i, w_i) . Similarly, ε_i may not be independent of (χ_i, w_i) , the true form of the regression (10) may be nonlinear, or the sample may not be *i.i.d.* Nonetheless, even though the classical errors-in-variables model is not a perfect representation of the relationship between saving, cash flow, and Tobin's q , our tests indicate that we have a useful approximation. Our failure to reject is even more interesting in light of the Monte Carlo result in Appendix B that this test tends to over-reject slightly in finite samples. More important, Appendix B shows that this test has good finite-sample power to detect modest amounts of model misspecification.

Next we test our second prediction that the cash-flow coefficient is more negative for firms with more costly external finance. In contrast to the case of our simulations, in which we know exactly which firms are constrained, no perfect measures of the severity of external finance constraints exist. Nonetheless, we use two commonly used measures: firm size and the existence of a bond rating.⁹ Although our Japanese sample is large enough to perform sample splits, we omit these results for brevity because they are similar to, but slightly weaker than those obtained using U.S. data.

We define a firm as large if the book value of its assets lies above the 67th percentile and small if its assets lie below the 33rd percentile. In the literature on finance constraints, size is often used as an indicator of the cost of raising external funds, and Hennessy and Whited (2007) demonstrate that it is an accurate indicator. Firm size confers an important advantage over other indicators such as dividend payout. Size can be considered exogenous, because it is not a choice variable for the manager in the short run and because our estimates exploit cross sectional data variation instead of time series variation. The intuition behind using the existence of a bond rating is that a firm with a bond rating has undergone a great deal of public scrutiny and is less likely to encounter the asymmetric information problems that lead to finance constraints. We regard bond ratings as exogenous, because agencies that provide bond ratings tend to base their judgments more on a consistent history of good financial and operating performance than on current operating decisions.

The first half of Table 4 presents the results from these two sample splits. Our OLS results confirm those in Almeida, et al. (2004) that small firms and firms without bond ratings have a

⁹We omit other, commonly used measures of finance constraints such as dividend payout and the KZ index (from Kaplan and Zingales, 1997) because they are endogenously determined with investment. We do not consider commercial paper ratings because too few firms have them.

stronger response of saving to cash flow than their unconstrained counterparts. Our GMM estimates, however, paint a different picture. All of the estimates are negative, and the unconstrained groups of firms have cash-flow coefficients that are significantly more negative than those of the constrained firms. This result does not support the model's prediction that firms with more costly external finance have more negative cash-flow coefficients. We conjecture that the effect of uncertainty on saving dwarfs the effect of the cost of external finance, which should be present in light of the evidence in Hennessy and Whited (2007) that small firms face more costly external finance than large firms. When we examine differences between the large and small firms, one characteristic that stands out is the marked difference in the degree of uncertainty that they face. We estimate a first-order panel autoregression of the ratio of operating income to assets for both groups of firms, using the technique in Holtz-Eakin, Newey, and Rosen (1988). The standard deviations of the error term for the small firms and the firms without bond ratings are 0.149 and 0.107, respectively. In contrast, the standard deviations for the large firms and the firms with bond ratings are only half as large at 0.075 and 0.067. As demonstrated in the model simulations, firms that face a great deal of uncertainty do not make large changes in their cash holdings in response to income shocks.

It is natural to examine at this point the conjecture that uncertainty matters more than finance constraints by testing our third and fourth predictions, which concern the relation between saving propensities and the serial correlation and variance of income. To this end we estimate an $AR(1)$ for operating income (scaled by total assets) firm by firm, only using firms with at least six consecutive observations. We then sort our sample by the estimates of serial correlation and residual standard deviation, throw out the middle third, and compare saving propensities across the top and bottom thirds. Because firm-by-firm autoregressions produce noisy estimates, we discard the middle third to minimize the possibility that we incorrectly classify individual firms.¹⁰

The next four lines of Table 4 contain the results for firms grouped by the standard deviation of the residual of the $AR(1)$ process for operating income. As predicted by the model, the low standard deviation group has a large, significantly negative cash-flow coefficient, whereas the high standard deviation group has a cash-flow coefficient that, while statistically different from zero, is also statistically different from the coefficient from the high standard deviation group. Not

¹⁰We have also estimated these autoregressions on an industry level with almost identical results.

surprisingly, given the results on large and small U.S. firms, the low standard deviation group contains firms considerably larger than those in the high standard deviation group. The mean level of assets for the former is 4,053 million 1997 dollars, and the mean level of assets for the latter is 63 million 1997 dollars. This result is of further interest because it demonstrates that the degree of uncertainty is not proxying for finance constraints. Otherwise we should have seen a more negative cash-flow coefficient on the high-uncertainty firms, but we do not.

The results from examining samples with low and high serial correlation are in the final four lines of Table 4. The GMM estimates of the cash-flow coefficient are significantly higher for the high serial correlation firms than for the low serial correlation firms. In the first group the median serial correlation is near 0.8, and in the second it is near zero. For the low serial correlation group the cash-flow coefficient is only significant if one considers the asymptotic 5% critical value instead of the bootstrapped 5% critical value. These results support the simulation that shows a low saving sensitivity if the serial correlation of the income process is low, but a large negative saving sensitivity if the serial correlation is high.

Just as we do for our full-sample results, we also present in Table 5 summary statistics describing the yearly regressions underlying the Fama-MacBeth estimates. We once again find a very low incidence of overidentifying restriction rejections and a fairly high incidence of identification test rejections. Both of these pieces of evidence support the reliability of our split-sample results.

We next test directly whether uncertainty matters as much as finance constraints by running a regression of saving on Tobin's q , the log of assets, cash flow, a constraint dummy, a low uncertainty dummy, the interaction of each of these dummies with cash flow, the interaction of the two dummies with each other, and the triple interaction of both of these dummies with cash flow. We use each of our measures of finance constraints separately. The results are in Table 6, which for brevity reports only the cash-flow coefficient and the coefficients on the interaction terms. For both size and the existence of a bond rating, we find a negative coefficient on cash flow and a positive coefficient on the interaction of cash flow and the constraint dummy. These coefficient estimates confirm the differential sensitivity results in Table 4.

Two pieces of evidence are of particular interest in Table 6. Both concern the figure in the last column, which is the sum of the coefficients on the three interaction terms. It measures the net effect

of being constrained and having low uncertainty. When we use size to proxy for finance constraints, this coefficient is positive, but insignificantly different from zero if we use our bootstrapped critical values. When we use the bond-rating dummy, this coefficient is negative and significant. In the first case stripping away the high-uncertainty firms from the constrained group leaves no differential sensitivity between this smaller constrained group and the rest of the sample. In the second case stripping away the high-uncertainty firms allows the predicted negative relation between saving propensities and finance constraints to be apparent. Nonetheless, the main message of this table is that both uncertainty and finance constraints affect the cash flow coefficient and that whatever its sign, this coefficient cannot be used as a summary measure of finance constraints.

Although we have demonstrated in the appendix the accuracy and good finite sample properties of the estimators we use, it would be reassuring to find a different measurement-error consistent estimator that produced similar results. To this end, we construct the proxy for Tobin's q in Cummins, Hassett, and Oliner (2006) and Bond and Cummins (2001), which is built on the idea that Tobin's q should be the expected discounted profits from using capital. This measure uses analysts' estimates of earnings and growth rates from I/B/E/S/ to proxy for expected future profits. Although it is difficult to argue that this proxy is superior or inferior to the usual market-based proxy, these papers do argue convincingly that one can use lagged observable variables as instruments, because these proxies are based on forecasts that are, by definition, orthogonal to information known at the time of the forecast. As argued in Erickson and Whited (2000), such is not the case for the usual market-based proxy for Tobin's q because the measurement errors in this proxy are highly serially correlated. When we use the proxy based on analysts' forecasts, and use the same instrument sets as in Cummins, Hassett, and Oliner (2006) and Bond and Cummins (2001), we replicate their result that the coefficient on cash flow in an investment regression is insignificantly different from zero. Using this proxy in a saving regression produces a coefficient of -0.310 with a heteroskedasticity and autocorrelation robust standard error of 0.057. Although this result confirms the rest of our evidence, we use this technique only as a robustness check because the data required to construct this alternate proxy for Tobin's q are available for less than half of our original sample.

It is interesting to compare our results with those in Erickson and Whited (2000), who find cash-flow coefficients near zero when applying measurement-error consistent estimators to a regression of

investment on q and cash flow. This result may seem puzzling because the investment and saving regressions have the same right-hand-side variables, because the magnitudes of the OLS estimates of the cash-flow coefficients are similar in the two regressions, and because measurement error bias propagates through the covariance matrix of the regressors. The difference between our results and theirs lies in their use of a different proxy for investment opportunities, whose estimate of τ^2 is about 0.4—twice as large as the estimate for our proxy. Therefore, even though both studies find positive OLS estimates of cash-flow coefficients, in the saving regression these estimates are more severely biased upward than they are in the investment regression. This differential bias implies that the true coefficient in the saving regression is negative and the true coefficient in the investment regression is zero. Not surprisingly, in results not reported for brevity, we are unable to find positive OLS cash-flow coefficients in our saving regressions when using the proxy from Erickson and Whited (2000); and, as is the case here, we find negative GMM estimates of the cash-flow coefficients.

We next investigate the two interesting instances in which our model predicts a positive coefficient on cash flow: cash flow shocks independent of productivity and intermittent investment. A careful examination of the first situation is beyond the scope of this paper inasmuch as empirical examples of exogenous movements in cash flow are very difficult to find.¹¹ Nonetheless, we examine this issue in an informal way by looking at the two industries in which Tobin's q and cash flow are negatively correlated: SIC 28 (chemicals) and SIC 38 (measurement instruments). These two industries account for approximately 10% of our observations. Our measurement error consistent estimators produce mixed results. In approximately half of the years in our sample we find positive coefficients and some of them are significantly different from zero. Although none of the negative coefficients are significant, the average of the coefficients over the sample period is negative. Although this result tells us nothing about a causal relationship between the importance of exogenous cash-flow shocks and the sign of the cash flow effect, it is interesting in that it is consistent with the prediction of our model that profit shocks unrelated to productivity should produce a positive sensitivity. To understand the source of our result, we examine the dimensions on which these firms differ from the rest of the sample. They are remarkably similar in terms of size, the autocorrelation and variance

¹¹We are only aware of three such studies in the empirical investment literature: Blanchard, Lopez-de-Silanes, and Schleifer (1994), Lamont (1997), and Rauh (2006).

of profits, the incidence of bond ratings, and leverage. They do invest less and hold more cash than the rest of the sample, but these differences are small. The similarity between the two subsamples makes it difficult to attribute the positive cash flow coefficients to the presence of finance constraints. Instead, we conjecture that technological or competitive factors are driving our result.

To examine the second situation we look at firms in the lowest size decile in our sample. As shown in Whited (2006), these micro firms are the only ones in Compustat whose investment is intermittent. Indeed, in our sample, these firms have an investment rate less than 1% in 12% of the firm-year observations. The corresponding figure for the rest of the sample is 4%. In accordance with the prediction of our model that intermittent investment produces a positive sensitivity of saving to cash flow, we do find a positive and significant saving sensitivity in most years. However, as above, it is difficult to attribute this result to finance constraints because these firms also have highly variable income with low serial correlation. It is also worth noting that we obtain negative sensitivities in the other nine size deciles.¹² In sum, although we can isolate some samples for which the sign of the cash flow coefficient is positive, the bulk of our results point to a negative coefficient. This result should not be surprising in light of the evidence in Smets and Wouters (2007) and Whited (2006) that the conditions required for a positive sensitivity are unlikely to be pervasive.

V. Conclusion

The issue of corporate saving has recently received much attention, in large part because of the build-up of liquid assets in recent years by both U.S. and European firms. Prior empirical research, including papers by Opler, et al. (1999), Faulkender and Wang (2006), Almeida, et al. (2004), and Khurana, et al. (2006), has addressed two related issues: why firms hold cash and why firms save, that is, change their cash holdings. We have, for the most part, addressed the second issue. In so doing, we take care to model a firm's saving, financing, and real investment decisions simultaneously in a stochastic, dynamic framework, to form a strong link between our theory and our empirical tests, and to account for measurement error in our empirical work.

¹²We have also examined firms in highly concentrated industries, firms in competitive industries, firms with low and high return volatility, and firms in industries with low and high depreciation rates. These sample splits capture, respectively, curvature of the profit function (which is a function of demand elasticity), uncertainty about both productivity and cost shocks, and depreciation rates. None of these subsamples produce a positive cash flow coefficient.

This approach leads us to conclusions richer than previous theoretical and empirical results. Our dynamic model predicts that, conditional on Tobin's q , the firm counteracts movements in cash flow with opposite movements in saving. This negative propensity to save occurs because a positive productivity shock causes both cash flow and the marginal product of capital increase. The firm then wants to decrease its cash stock to buy capital goods that have become relatively more productive, that is, dissave and invest. In contrast, the sensitivity of saving to cash flow in earlier models with one-period saving decisions, such as the one in Almeida, et al. (2004), is typically positive because good cash flow news is modeled as an increase in the firm's endowment. The marginal product of capital is unaffected, and the firm has no incentive to transform liquid assets into productive assets. Instead, the firm wants to hold high levels of cash to avoid costly external finance. This income effect operates in our model as well, but our model allows for the existence of a substitution effect, which is stronger. This distinction between income and substitution effects requires an endogenous investment decision and variation in capital productivity, which earlier work did not consider explicitly. This work could not, therefore, reach our results.

We find strong empirical support for our negative-sensitivity result in data from six countries. When we estimate this sensitivity using OLS, we find the standard result in the literature that the sensitivity of saving to cash flow is positive. However, when we correct econometrically for measurement error in Tobin's q , we find the opposite result, except in two small subsamples of firms. We also find that the saving sensitivity increases in absolute value with the serial correlation of income and decreases with the variance of income shocks. Interestingly, the effect of uncertainty on the propensity to save out of cash flow is empirically at least as strong as the effect of finance constraints. Consequently, propensities to save cannot be used as summary measures of either the cost of external finance or of any of the other multitude of factors that affect these propensities.

Taken together, our model and evidence demonstrate that a wide variety of both real and financial factors determine corporate saving. Our results also reemphasize that any model estimation involving Tobin's q can suffer similarly from significant mismeasurement bias. The variable q is important in many contexts, and its mismeasurement has the potential to bias results in any context in which q is correlated with other variables in the model. Because q is a broad measure of firm health, it is likely that this correlation issue will be important in many situations, as it is here.

Appendix A. Data Definitions

We select the sample by first deleting any firm-year observations with missing data. Next, we delete any observations for which total assets, the gross capital stock, or sales are either zero or negative. Then for each firm we select the longest consecutive times series of data and exclude firms with only one observation. Finally, we omit all firms whose primary SIC classification is between 4900 and 4999, between 6000 and 6999, or greater than 9000, because our model is inappropriate for regulated, financial, or quasi-public firms.

We define data variables from Global Industrial/Commercial as follows: Book assets is Item 89; investment is Item 193; operating income is Item 14; cash flow is the sum of Items 11 and 32; and cash is Item 60. The numerator of the market-to-book ratio is the sum of the market value of equity (Item 3 times Item 13 in Global Issue) and total book assets minus the book value of equity (Item 105+Item 135), and the denominator is book assets. For our U.S. data from Compustat we define book assets as item 6, operating income as item 13, cash flow is the sum of items 14 and 18, cash as item 1, the number of common shares as item 25, and the share price as item 199. In our regressions we scale both saving and cash flow by total assets. We delete the top and bottom 1% of our regression variables.

Appendix B. Monte Carlo Experiments

In order to allay skepticism of empirical results that have been produced by unusual estimators on fairly small samples, in Table 7 we report the results of a Monte Carlo simulation using artificial data similar to our real data, both in terms of sample size and observable moments. These simulations are of particular interest because these estimators have most commonly been used on investment regressions instead of saving regressions, and because saving and investment have different statistical properties. Most importantly, the distribution of investment is highly skewed, whereas the distribution of saving is much more symmetric.

We do three experiments. For each we generate 10,000 simulated cross sections from (12) and (13). The first has a sample size of 3,000, the second a sample size of 1,200, and the third a sample size of 200. These numbers correspond to the size of the largest and smallest cross sections in our data set, as well as to an intermediate size. For each simulation we set the parameters β , α , r^2 , and τ^2 approximately equal to the averages of the corresponding GMM estimates from Tables 2 through 7. Each observation is of the form (y_i, x_i, w_i) , generated according to (10)-(11) so that (y_i, x_i, w_i) has, on average over the simulation samples, first and second moments equal to, and higher-order

moments comparable with the corresponding average sample moments from our real data.

For the third-, fourth-, fifth-, and sixth-order GMM estimators, Table 14 reports the mean value of the estimator of our parameter of interest, α_1 . It also reports its mean absolute deviation (MAD), the probability that an estimate is within 20% of its true value, and the actual size of a nominal 5% two-sided test of the null hypothesis that α_1 equals its true value. For the small and intermediate sample sizes Table 14 shows that the fourth-order GMM estimator (GMM4) gives the best estimates in terms of expected value, MAD, and probability concentration. For the large sample size the GMM6 estimator performs best. Because the performance of the GMM4 and GMM6 estimators is similar for the large sample size, we therefore use the GMM4 estimator for our empirical work. Also of interest in this table are the tiny actual sizes of the test of the null hypothesis that α_1 equals its true value for the intermediate and large sample sizes.¹³

Table 8 explores the power of the J-test to detect misspecification. We examine four likely types of departures from the linear errors-in-variables model. Each is obtained by introducing one type of misspecification into the correctly specified baseline simulation described above. First, we make y_i depend nonlinearly on χ_i ; second, we mismeasure the capital stock by multiplying (y_i, x_i, w_i) from the baseline sample by an *i.i.d.* lognormal variable; third, we introduce a correlation between u_i and χ_i ; and fourth, we violate the *i.i.d.* assumption by allowing different quintiles of our simulated observations to be generated by different distributions. We limit the degree of each misspecification so that the absolute biases in the GMM estimates of α_1 do not exceed 0.3, which is approximately the absolute value of the cash flow coefficient estimated in our real data. For the first three types of misspecification we find that the fourth, fifth, and sixth order GMM *J*-tests exhibits usefully large power for the largest sample size, ranging from 0.403 to 0.995. The test is more powerful for larger sample sizes, and all of the power figures are larger than the fractions of rejections we obtain in our empirical work. For the fourth type we find that the coefficient estimates are affected little, even though the J-test has lower power to detect non-*i.i.d.* samples. Finally, we did not combine misspecifications, which we suspect would further increase test power.

¹³This result is the opposite of that found in Erickson and Whited (2000) for *investment* regressions.

References

- Acharya Viral, Heitor Almeida, and Murillo Campello, 2007, Is cash negative debt? A hedging perspective on corporate financial policies, *Journal of Financial Intermediation* 16, 515-554.
- Almeida, Heitor, Murillo Campello, and Michael S. Weisbach, 2004, The cash flow sensitivity of cash, *Journal of Finance* 59, 1777–1804.
- Altinkilic, Oya and Robert S. Hansen, 2000, Are there economies of scale in underwriting fees? Evidence of rising external financing costs, *Review of Financial Studies* 13, 191–218.
- Blanchard, Olivier Jean, Florencio Lopez-de-Silanes, and Andrei Shleifer, 1994, What do firms do with cash windfalls, *Journal of Financial Economics* 36, 337-360.
- Bond, Stephen P and Jason G. Cummins, 2001, Noisy share prices and the q model of investment. Unpublished working paper. Institute for Fiscal Studies.
- Caggese, Andrea, 2006, Testing financing constraints on firm investment using variable capital, *Journal of Financial Economics*, forthcoming.
- Cooper, Russell and John Haltiwanger, 2006, On the nature of capital adjustment costs, *Review of Economic Studies* 73, 611–633.
- Costa, Cristiano and Lourenco Paz, 2004, Are Brazilian firms savings sensitive to cash windfalls?, Working Paper, University of Maryland.
- Cummins, Jason, Kevin Hassett, and Stephen S. Oliner, 2006, Investment behavior, observable expectations, and internal funds, *American Economic Review* 96, 796-810.
- Dittmar, Amy and Jan Mahrt-Smith, 2007, Corporate governance and the value of cash holdings, *Journal of Financial Economics* 83, 599-634.
- Eisfeldt, Andrea L. and Adriano A. Rampini, 2006, Financing shortfalls and the value of aggregate liquidity, Working paper, Northwestern University.
- Erickson, Timothy and Toni M. Whited, 2000, Measurement error and the relationship between investment and q , *Journal of Political Economy* 108, 1027–57.
- Erickson, Timothy and Toni M. Whited, 2002, Two-step GMM estimation of the errors-in-variables model using high-order moments, *Econometric Theory* 18, 776–799.
- Erickson, Timothy and Toni M. Whited, 2006, On the accuracy of different measures of Q , *Financial Management* 35, 5-33.
- Fama, Eugene and James MacBeth, 1973, Risk, return, and equilibrium: empirical tests, *Journal of Political Economy* 81, 607–36.
- Faulkender, Michael and Rong Wang, 2006, Corporate financial policy and cash holdings, *Journal of Finance* 61, 1957–1990.
- Ferrando, Annalisa and Rozalia Pal, 2006, Financial constraints and firm’s cash policy in the Euro area, Working Paper, European Central Bank.
- Fischer, Edwin O., Robert Heinkel, and Josef Zechner, 1989, Dynamic capital structure choice: theory and tests, *Journal of Finance* 44, 19–40.
- Foley, C. Fritz, Jay C. Hartzell, Sheridan Titman, and Garry Twite, 2006, Why do firms hold so much cash? A tax-based explanation, *Journal of Financial Economics*, forthcoming.
- Fraumeni, Barbara, 1997, The measurement of depreciation in the U.S. national income and product accounts, *Survey of Current Business* 77, 7-23.
- Gamba, Andrea and Alexander Triantis, 2006, The value of financial flexibility, *Journal of Finance*, forthcoming.

- Geary, R. C., 1942, Inherent relations between random variables, *Proceedings of the Royal Irish Academy A* 47, 63–76.
- Gomes, Joao, 2001, Financing investment, *American Economic Review* 91, 1263–1285.
- Greene, William H., 1997, *Econometric Analysis*, Prentice Hall, Upper Saddle River, NJ.
- Hall, Peter, and Horowitz, Joel L., 1996, Bootstrap critical values for tests based on generalized-method-of-moment estimators, *Econometrica* 64, 891–916.
- Hennessy, Christopher A. and Toni M. Whited, 2005, Debt dynamics. *Journal of Finance* 60, 1129–1165.
- Hennessy, Christopher A. and Toni M. Whited, 2007, How costly is external financing? Evidence from a structural estimation, *Journal of Finance* 62, 1705–1745
- Holtz-Eakin, Douglas, Whitney K. Newey, and Harvey Rosen, 1988, Estimating vector autoregressions with panel data, *Econometrica* 56, 1371-1395.
- Kaplan, Steven N., and Luigi Zingales, 1997, Do investment-cash flow sensitivities provide useful measures of financing constraints? *Quarterly Journal of Economics* 112, 169-215.
- Khurana, Inder K., Xiumin Martin, and Raynolde Pereira, 2006, Financial development and the cash flow sensitivity of cash, *Journal of Financial and Quantitative Analysis*, forthcoming.
- Kim, Chang-Soo, David C. Mauer, and Ann E. Sherman, 1998, The determinants of corporate liquidity: Theory and evidence, *Journal of Financial and Quantitative Analysis* 33, 335–359.
- Lamont, Owen, 1997, Cash flow and investment: evidence from internal capital markets, *Journal of Finance* 52, 83-109.
- Lin, Yi-Chen, 2007, The cash flow sensitivity of cash: Evidence from Taiwan, *Applied Financial Economics*, 1, 1–12.
- Opler, Timothy, Lee Pinkowitz, René Stulz, and Rohan Williamson, 1999, The determinants and implications of corporate cash holdings, *Journal of Financial Economics* 52, 3–46.
- Petersen, Mitchell A., 2007, Estimating standard errors in finance panel data sets: comparing approaches, *Review of Financial Studies*, forthcoming.
- Pinkowitz, Lee, and Rohan Williamson, 2001, Bank power and cash holdings: Evidence from Japan, *Review of Financial Studies* 14, 1059–1082.
- Rauh, Joshua, 2006, Investment and financing constraints: evidence from the funding of corporate pension plans, *Journal of Finance* 61, 33-71
- Rotemberg, Julio J. and Michael Woodford, 1992, Oligopolistic pricing and the effects of aggregate demand on economic activity, *Journal of Political Economy* 100, 1153–1207.
- Rotemberg, Julio J. and Michael Woodford, 1999, The cyclical behavior of prices and costs. In: Taylor, J. B., Woodford, M. (Eds.), *Handbook of Macroeconomics*, Vol. 1B. North Holland, Amsterdam, pp. 1051-1135.
- Smets, Frank and Rafael Wouters, 2007, Shocks and frictions in U.S. business cycles: A Bayesian DSGE approach, *American Economic Review* 97, 586-606.
- Stokey, Nancy L. and Robert E. Lucas, 1989, *Recursive Methods in Economic Dynamics*. Harvard University Press, Cambridge, Mass. and London.
- Strebulaev, Ilya, 2006, Do tests of capital structure mean what they say? *Journal of Finance*, forthcoming.
- Sufi, Amir, 2006, The real effects of debt certification: evidence from the introduction of bank loan ratings, *Review of Financial Studies*, forthcoming.

- Tang, Tony T., 2007, Information asymmetry and firms' credit market access: Evidence from Moody's credit rating format refinement, Working paper, University of Chicago.
- Tauchen, George, 1986, Finite state Markov-chain approximations to univariate and vector autoregressions, *Economics Letters* 20, 177–181.
- Whited, Toni M., 2001, Is it inefficient investment that causes the diversification discount?" *Journal of Finance* 56, 1667–92.
- Whited, Toni M., 2006, External finance constraints and the intertemporal pattern of intermittent investment, *Journal of Financial Economics* 81, 467–502.

Table 1: Summary Statistics

	Investment	Market to Book Ratio	Cash Flow	Change in Cash Stock	Cash Stock	Total Assets	Average Obs. per Year
United States							
Mean	0.0967	1.4311	0.1341	0.0031	0.1119	1522	2,258
Median	0.0621	1.1831	0.1405	0.0003	0.0587	119	
Canada							
Mean	0.1212	1.3345	0.0772	0.0042	0.0919	1037	341
Median	0.0694	1.1312	0.0874	0.0000	0.0295	247	
United Kingdom							
Mean	0.0755	1.5043	0.0809	0.0025	0.1174	1054	796
Median	0.0442	1.3134	0.0968	0.0001	0.0710	115	
Japan							
Mean	0.0282	1.1940	0.0326	-0.0027	0.1647	461	2,070
Median	0.0182	1.0978	0.0307	-0.0021	0.1389	230	
France							
Mean	0.0107	1.2511	0.0615	-0.0078	0.1279	1116	353
Median	0.0072	1.1102	0.0686	-0.0010	0.0913	84	
Germany							
Mean	0.0519	1.2489	0.0704	-0.0102	0.1075	1418	332
Median	0.0384	1.1498	0.0793	-0.0023	0.0586	122	

Calculations are based on a sample of U.S. nonfinancial firms from Compustat from 1972 to 2006 and a sample of international firms from Global Vantage from 1994 to 2005. Investment, cash flow, the cash stock, and the change in the cash stock are all deflated by total assets. Total assets are expressed in millions of 1997 U.S. dollars.

Table 2: Pooled Saving Regressions: All Countries

Country	OLS			GMM4			
	q	CF	R^2	q	CF	R^2	τ^2
United States	0.029* [†] (0.003)	0.103* [†] (0.009)	0.112* [†] (0.010)	0.283* [†] (0.016)	-0.397* [†] (0.060)	0.440* [†] (0.031)	0.255* [†] (0.014)
Canada	0.045* [†] (0.006)	0.053* [†] (0.025)	0.144* [†] (0.026)	0.213* [†] (0.018)	-0.076* [†] (0.022)	0.495* [†] (0.051)	0.323* [†] (0.041)
United Kingdom	0.009 [†] (0.002)	0.103* [†] (0.016)	0.047* [†] (0.013)	0.427 (0.076)	-0.485* [†] (0.168)	0.356* [†] (0.042)	0.137* [†] (0.026)
Japan	0.019* [†] (0.002)	0.141* [†] (0.019)	0.049* [†] (0.005)	0.318* [†] (0.040)	-0.162* [†] (0.037)	0.255* [†] (0.020)	0.113* [†] (0.015)
France	0.021* [†] (0.003)	0.126* [†] (0.033)	0.084* [†] (0.013)	0.263* [†] (0.084)	-0.304* [†] (0.097)	0.303* [†] (0.060)	0.226* [†] (0.060)
Germany	0.018 [†] (0.004)	0.078* [†] (0.020)	0.082* [†] (0.018)	0.310* [†] (0.073)	-0.200* [†] (0.087)	0.354* [†] (0.069)	0.122* [†] (0.025)

Calculations are based on a sample of U.S. nonfinancial firms from Compustat from 1972 to 2006 and a sample of international firms from Global Vantage from 1994 to 2005. GMM estimates are from the fourth-order estimator in Erickson and Whited (2000). The dependent variable is the change in the stock of cash divided by total assets. CF stands for cash flow divided by total assets; q stands for the market-to-book ratio; and τ^2 is the coefficient of determination of the measurement equation. Fama-MacBeth standard errors are below the average estimates in parentheses. An asterisk indicates that the t-statistic exceeds the 5% bootstrapped critical value. A dagger indicates that the t-statistic exceeds the 5% asymptotic critical value.

Table 3: Yearly Saving Regressions Summary: All Countries

	Fraction of Negative Cash Flow Coefficients	Fraction of Significant Negative Cash Flow Coefficients	Fraction of Overidentifying Restriction Rejections	Fraction of Identification Test Rejections
United States	0.886	0.714	0.085	0.800
Canada	0.833	0.250	0.083	0.583
United Kingdom	0.833	0.333	0.250	0.417
Japan	1.000	0.333	0.250	0.750
France	0.750	0.333	0.083	0.167
Germany	0.667	0.250	0.000	0.500

Calculations are based on a sample of U.S. firms from Compustat from 1972 to 2006 and a sample of international firms from Global Vantage from 1994 to 2005. The first column contains the fraction of the yearly estimates that are negative. The second column contains the fraction of the yearly estimates that are significantly negative at the 5% level, using asymptotic critical values. The third column contains the fraction of the yearly tests of overidentifying restrictions that produce rejections at the 5% level. The fourth column contains the fraction of the yearly identification tests that produce rejections at the 5% level.

Table 4: Split Sample Regressions: United States

Subsample	OLS			GMM4			
	q	CF	R^2	q	CF	R^2	τ^2
Small	0.045*† (0.004)	0.134*† (0.011)	0.166*† (0.015)	0.265*† (0.019)	-0.147*† (0.071)	0.522*† (0.034)	0.300*† (0.020)
Large	0.006*† (0.001)	0.083*† (0.008)	0.046*† (0.006)	0.281*† (0.054)	-0.856*† (0.172)	0.183*† (0.027)	0.342*† (0.031)
No Bond Rating	0.032*† (0.003)	0.110*† (0.010)	0.122*† (0.012)	0.244*† (0.023)	-0.247*† (0.068)	0.444*† (0.036)	0.291*† (0.038)
Bond Rating	0.016*† (0.003)	0.046† (0.015)	0.070*† (0.009)	0.219*† (0.030)	-0.815*† (0.130)	0.254*† (0.034)	0.417 (0.022)
High Standard Deviation	0.037*† (0.004)	0.128*† (0.008)	0.150*† (0.013)	0.315*† (0.062)	-0.274*† (0.090)	0.517*† (0.037)	0.264*† (0.018)
Low Standard Deviation	0.014*† (0.002)	0.081*† (0.009)	0.058*† (0.008)	0.299*† (0.098)	-0.836*† (0.280)	0.322*† (0.033)	0.366*† (0.025)
High Serial Correlation	0.023*† (0.003)	0.088*† (0.009)	0.102*† (0.008)	0.248*† (0.025)	-0.579*† (0.072)	0.380*† (0.033)	0.344*† (0.017)
Low Serial Correlation	0.033*† (0.004)	0.122*† (0.009)	0.122*† (0.011)	0.213*† (0.025)	-0.074 (0.045)	0.416*† (0.037)	0.266*† (0.018)

Calculations are based on a sample of U.S. nonfinancial firms from Compustat from 1972 to 2006. GMM estimates are from the fourth-order estimator in Erickson and Whited (2000). The dependent variable is the change in the stock of cash divided by total assets. CF stands for cash flow divided by total assets; q stands for the market-to-book ratio; and τ^2 is the coefficient of determination of the measurement equation. Serial correlation is the first-order autoregressive coefficient on the ratio of operating income to assets, and standard deviation is the standard deviation of the residual from this regression. Fama-MacBeth standard errors are below the average estimates in parentheses. An asterisk indicates that the t-statistic exceeds the 5% bootstrapped critical value. A dagger indicates that the t-statistic exceeds the 5% asymptotic critical value.

Table 5: Yearly Saving Regressions Summary: Split Samples

	Fraction of Negative Cash Flow Coefficients	Fraction of Significant Negative Cash Flow Coefficients	Fraction of Overidentifying Restriction Rejections	Fraction of Identification Test Rejections
Small	0.571	0.371	0.057	0.629
Large	0.886	0.429	0.171	0.514
No Bond Rating	0.742	0.600	0.114	0.800
Bond Rating	0.857	0.600	0.029	0.571
High Standard Deviation	0.686	0.514	0.086	0.457
Low Standard Deviation	0.800	0.600	0.200	0.457
High Serial Correlation	0.971	0.657	0.086	0.457
Low Serial Correlation	0.571	0.257	0.171	0.571

Calculations are based on a sample of U.S. nonfinancial firms from Compustat from 1972 to 2004. The first column contains the fraction of the yearly estimates that are negative. The second column contains the fraction of the yearly estimates that are significantly negative at the 5% level, using asymptotic critical values. The third column contains the fraction of the yearly tests of overidentifying restrictions that produce rejections at the 5% level. The fourth column contains the fraction of the yearly identification tests that produce rejections at the 5% level.

Table 6: Uncertainty versus Finance Constraints

Constraint Indicator	CF	$CF \times D_C$	$CF \times D_L$	$CF \times D_C \times D_L$	Sum
Size	-0.384*† (0.070)	0.347*† (0.076)	-0.653*† (0.082)	0.443*† (0.075)	0.163† (0.072)
Bond Rating	-0.533*† (0.070)	0.318*† (0.062)	-0.972*† (0.097)	0.478*† (0.087)	-0.177*† (0.052)

Calculations are based on a sample of U.S. nonfinancial firms from Compustat from 1972 to 2004. GMM estimates are from the fourth-order estimator in Erickson and Whited (2000). The dependent variable is the change in the stock of cash divided by total assets. CF stands for cash flow divided by total assets; D_C is a dummy variable that takes a value of 1 if a firm is classified as constrained, and 0 otherwise. D_L is a dummy variable that takes a value of 1 if a firm is classified as having low uncertainty, and 0 otherwise. “Sum” refers to the sum of the coefficients in columns 3–5. Size and bond rating are the two constraint indicators. Fama-MacBeth standard errors are below the average estimates in parentheses. An asterisk indicates that the t-statistic exceeds the 5% bootstrapped critical value. A dagger indicates that the t-statistic exceeds the 5% asymptotic critical value.

Table 7: Monte Carlo Performance of GMM and OLS Estimators

	OLS	GMM3	GMM4	GMM5	GMM6
Sample Size = 200					
$E(\hat{\alpha}_1)$	0.180	-0.234	-0.223	-0.168	-0.201
$MAD(\hat{\alpha}_1)$	0.481	0.409	0.299	0.326	0.342
$P(\hat{\alpha}_1 - \alpha_1 \leq 0.2 \alpha_1)$	0.001	0.085	0.130	0.150	0.130
T-test Size		0.032	0.062	0.074	0.090
J-test Size			0.235	0.329	0.529
Sample Size = 1200					
$E(\hat{\alpha}_1)$	0.183	-0.349	-0.329	-0.275	-0.288
$MAD(\hat{\alpha}_1)$	0.483	0.277	0.134	0.094	0.087
$P(\hat{\alpha}_1 - \alpha_1 \leq 0.2 \alpha_1)$	0.000	0.151	0.357	0.486	0.565
T-Test Size		0.001	0.008	0.007	0.011
J-test Size			0.136	0.323	0.394
Sample Size = 3000					
$E(\hat{\alpha}_1)$	0.184	-0.348	-0.331	-0.292	-0.300
$MAD(\hat{\alpha}_1)$	0.484	0.201	0.078	0.049	0.043
$P(\hat{\alpha}_1 - \alpha_1 \leq 0.2 \alpha_1)$	0.000	0.221	0.507	0.704	0.796
T-Test Size		0.000	0.002	0.001	0.001
J-test Size			0.096	0.317	0.418

Indicated expectations and probabilities are estimates based on 10,000 Monte Carlo samples. The samples are generated by

$$\begin{aligned} y_i &= q_i \beta + w_i \alpha + u_i \\ x_i &= \gamma + \chi_i + \varepsilon_i, \end{aligned}$$

in which χ_i and ε_i are distributed as a chi-squared variables. u_i is distributed as a negative chi-squared variable. GMM n denotes the GMM estimator based on moments up to order $M = n$. OLS denotes estimates obtained by regressing y_i on x_i and w_i . MAD denotes mean absolute deviation. ‘‘T-Test Size’’ refers to the actual size of a nominal 5% test of the null hypothesis that α_1 equals its true value. ‘‘J-Test Size’’ refers to the actual size of a nominal 5% test of the overidentifying restrictions.

True Value: $\alpha_1 = -0.3$.

Table 8: Monte Carlo Performance of J-Test


	GMM4	GMM5	GMM6
Sample Size = 200			
Nonlinear Regression	0.275	0.289	0.661
Mismeasured Denominator	0.246	0.362	0.636
Correlated Error Regressor	0.410	0.425	0.671
Non- <i>i.i.d.</i> Sample	0.185	0.387	0.485
Sample Size = 1200			
Nonlinear Regression	0.424	0.558	0.848
Mismeasured Denominator	0.325	0.403	0.600
Correlated Error Regressor	0.542	0.743	0.918
Non- <i>i.i.d.</i> Sample	0.163	0.394	0.533
Sample Size = 3000			
Nonlinear Regression	0.504	0.725	0.960
Mismeasured Denominator	0.403	0.466	0.651
Correlated Error Regressor	0.676	0.936	0.995
Non- <i>i.i.d.</i> Sample	0.272	0.382	0.592

The table reports the fraction of J-test rejections at a 5% nominal critical value. The samples are generated by

$$\begin{aligned}
 y_i &= q_i\beta + w_i\alpha + u_i \\
 x_i &= \gamma + \chi_i + \varepsilon_i,
 \end{aligned}$$


in which χ_i and ε_i are distributed as a chi-squared variables. u_i is distributed as a negative chi-squared variable. GMM n denotes the GMM estimator based on moments up to order $M = n$.

Figure 1: Policy Functions



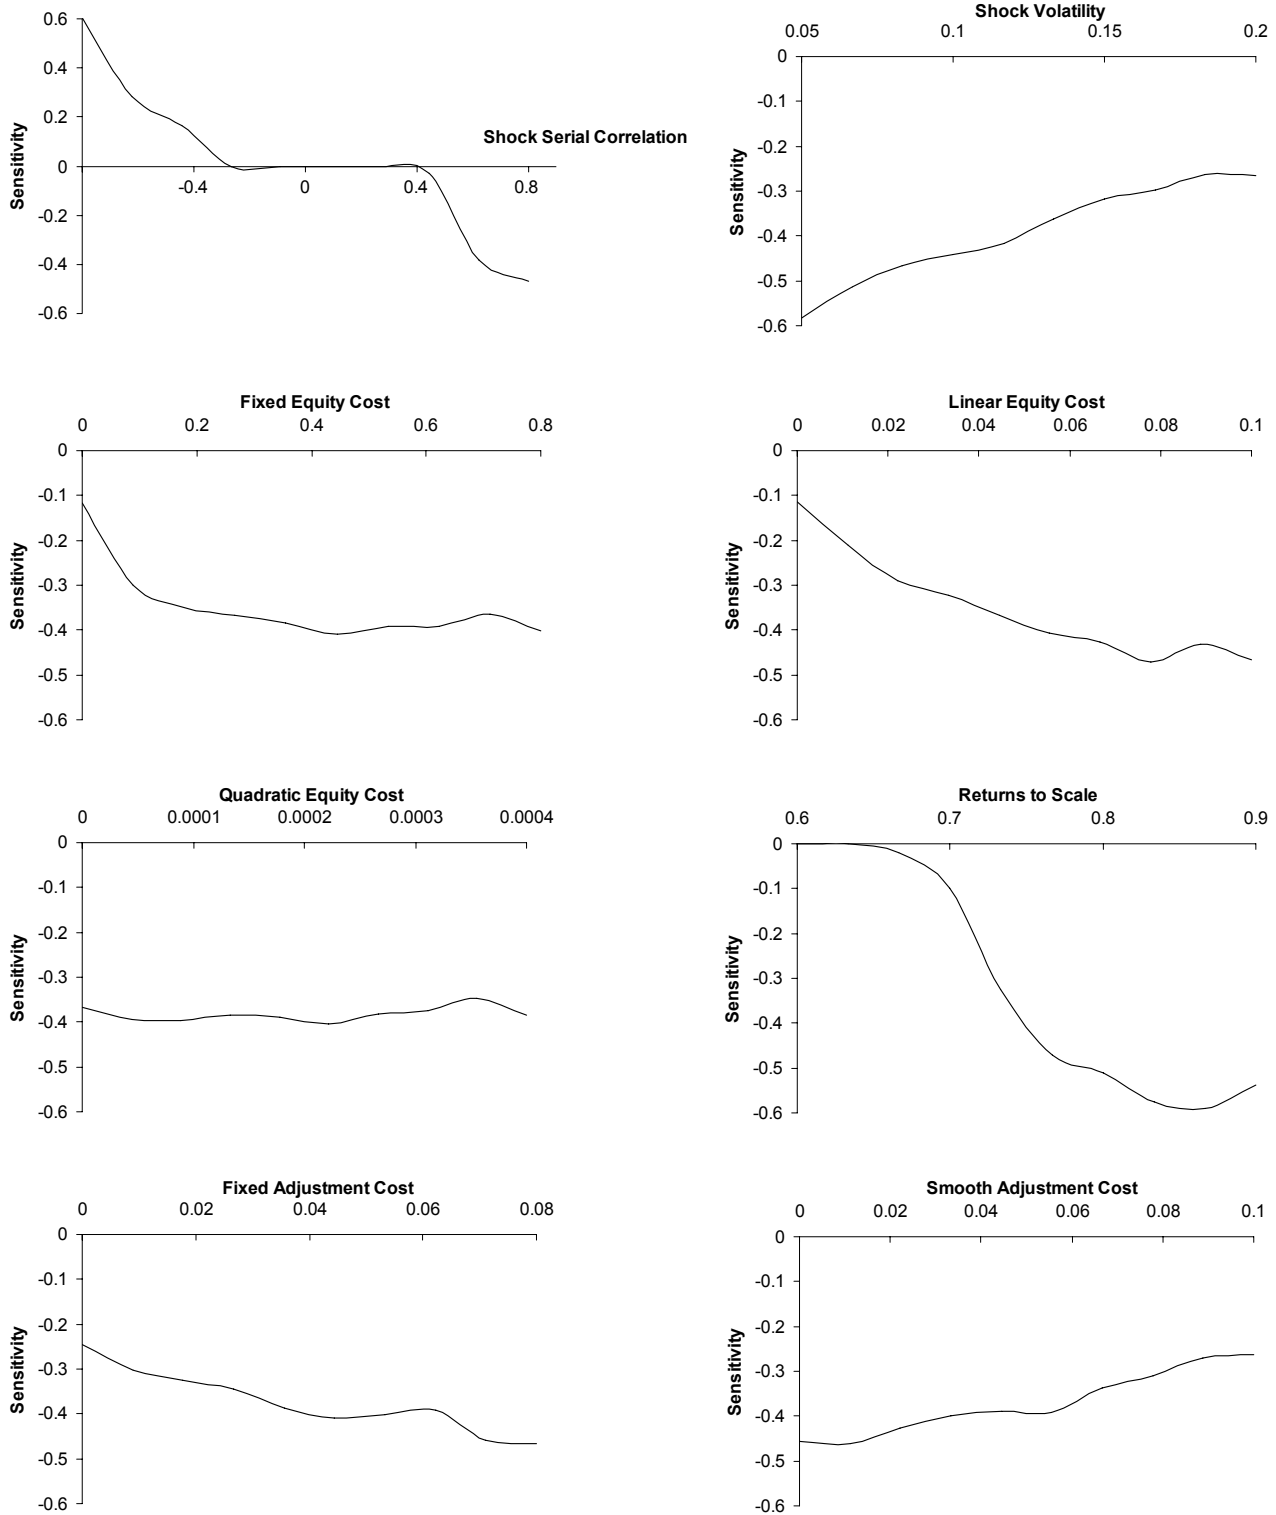
This figure depicts the optimal response of investment, saving, cash flow, and distributions/equity issuance in response to the productivity shock, z , in the revenue function zk^θ . The first panel depicts the response of the smallest firm in the simulated sample, the second panel depicts the response of the median firm in the simulated sample, and the third depicts the response of the largest firm in the simulated sample.

Figure 2: The Ratio of Cash to Assets



This figure depicts the relation between various model parameters and the average optimal ratio of the stock of cash to assets. Each graph is constructed by varying the model parameter on the horizontal axis. For each of these parameter values the model is solved and simulated, and the average ratio of cash to assets is taken over all of the simulated data points.

Figure 3: The Cash-Flow Sensitivity of Cash



This figure depicts the relation between various model parameters and the sensitivity of the change in the cash stock to cash flow, holding constant Tobin's q and the capital stock. Each graph is constructed by varying the model parameter on the horizontal axis. For each of these parameter values the model is solved and simulated, and the average sensitivity is calculated over all of the simulated data points.