EFEL 7/22/02
WHITED 3:15PM

External Finance Constraints and the Intertemporal Pattern of Intermittent Investment

Toni M. Whited*
University of Iowa
Department of Finance
Henry B. Tippie College of Business
Iowa City, IA 52242
(319) 335-0882
toni-whited@uiowa.edu

First version: April 30, 2001
This revision: March 9, 2002

*This paper previously circulated under thet itle "Internal Capital Markets, External Finance Constraints, and Lumpy Investment." I would like to thank Tim Erickson, Tom George, Jon Garfinkel, John Nasir,a nd seminar participants at the London Business School, the University of British Columbia, the University of Iowa, the University of Utah, Iowa State University, and the University of North Carolina for insightful comments. Gonul Colak provided excellent research assistance.

External Finance Constraints and the Intertemporal Pattern of Intermittent Investment

Abstract

Do external finance constraints affect the timing of large, indivisible investment projects? Simulations of a model with fixed capital-stock adjustment costs establish the hypothesis that external finance constraints lower a firm's investment hazard: the probability of undertaking a large project today as a function of the time since the last project. Hazard model estimation that controls for productivity and adjustment costs supports this hypothesis. Small firms that do not pay dividends have significantly lower hazards than small firms that do; small stand-alone firms have significantly lower hazards than small segments of conglomerates, and accumulation of liquid assets raises hazards.

The past fifteen years have seen a flood of empirical studies of the effects of external finance constraints on corporate investment. Such an outpouring of research stems from two important sources of interest. First, external finance constraints have implications for monetary policy transmission and tax policy. Second, the topic is of interest to financial economists, since imperfections in financial markets have implications for capital structure, capital budgeting decisions, and the operation of internal capital markets.

The intuition behind the connection between finance and investment starts with theoretical models of imperfect information. These models show that information asymmetry leads to a divergence between the costs of internal and external funds or, at the extreme, to a rationing of external funds. However,s uch models provide little guidance for the direction of empirical work, since few have both endogenous investment and finance decisions, and since few are couched in terms of observable variables. Therefore, empirical studies have turned to two loose arguments to motivate tests of the connection between finance and investment. First, finance constraints may cause an excess sensitivity of investment to internal funds; and second, they may affect the firm's incremental intertemporal investment allocation. This paper tackles this topic from a new, unexploited angle; briefly, one that examines the effects of finance constraints on the timing of large indivisible investment projects.

To understand the contribution of this idea, it is useful to discuss progenitors and relatives of the present paper. Most of thee mpirical research in this area has followed the methods first outlined in Fazzari, Hubbard, and Petersen (1988), who argued and found that if a firm cannot obtain outside finance, its investment will respond strongly to movements in cash flow, holding investment opportunities constant. More recently, however, Kaplan and Zingales (1997) and Cleary (1999) have provided evidence that cash-flow sensitivity need not identify *a priori* liquidity constrained firms. Further, the simulations in Gomes (2001) and the empirical work in Erickson and Whited (2000), Bond and Cummins (2001), and Cooper and Ejarque (2001) have demonstrated that the results from these earlier papers are due to measurement error in the usual proxy for investment opportunities: Tobin's q. In sum, these recent papers question whether the large body of research on cash-flow sensitivity has taught us much about the way in which external finance constraints affect investment. Thus, they have re-opened the door to understanding the mechanism whereby finance and investment interact.

A different line of research estimates directly theE uler equation of an intertemporal investment model using generalized method of moments. For example, Whited (1992) and Bond and Meghir (1994) show that augmentations of the Euler equation that account for financial constraints improve its fit. This approach has the advantage of avoiding the difficult problem of measuring q. Euler equation studies havep rovided convincing evidence that external finance constraints affect the rate of intertemporal substitution between investment today and investment tomorrow. However, these papers examine only marginal decisions, since they are based on models with convex capital-stock adjustment costs.

In contrast, common intuition suggests that finance constraints are at least as likely to alter a firm's decision about undertaking a large project or not; that is, they ought to have lumpy in addition to smooth effects. Loosely speaking, although finance constraints could affect a firm's decision to spread out theb uilding of a new plant over an extra month (a "marginal" intertemporal decision), they are more likely to affect a firm's decision to delay thee ntiren ew-plant project (a "lumpy" intertemporal decision). Further motivation for studying finance constraints in the context of lumpy investment comes from studies that have found a great deal of lumpy adjustment in plant-level data (Doms and Dunne, 1998; Cooper, Haltiwanger, and Power, 1999). For example, Doms and Dunne (1998) find that from 25% to 40% of an average plant's cumulative investment over 17 years is concentrated in a single year. If, as this evidence suggests, investment decisions are lumpy, then external finance constraints are quite likely to have lumpy effects.

Lumpy adjustment occurs when firms facen onconvex adjustment costs.[1] Two strategies have dominated thee mpirical literature on estimating and testing investment models with nonconvexities. First, as illustrated, for example, in Caballero, Engel,a nd Haltiwanger

[1]Other sources of lumpy adjustment, such as irreversibility, indivisible capital goods, and different puchase and sale prices for capital, can be though of as examples or extreme cases of nonconvex adjustment costs.

(1995) and Caballero and Engel (1999), one can construct a measure of the "gap" between the firm's actual and desired capital stock, where the latter typically comes from a theoretical frictionless model. Testable hypotheses emerge from this characterization because the reaction of investment to the gap depends on the nature of adjustment costs. However, as pointed out in Cooper and Willis (2001), because specifying an optimal capital stock requires a specific structural model and because an optimal capital stock needs to be defined in terms of a model with frictions, it is easy to mismeasure the gap: a problem that can lead to misleading inferences. This problem is analogous to the difficulty of measuring q, and it is also a generic problem with estimation of a structural model, since the resulting inferences can be fragile with respect to the choice of model assumptions.

A second, less structural, method is estimating an adjustment hazard function. In general, a hazard rate is the probability of an event occurring during the next instant, given that it has not occurred for a length of time, t. A hazard function is the relationship between this probability and t. In the context of investment, an adjustment hazard is the relationship between the probability of a large change in the capital stock at a certain point in time and the length of times ince the last large adjustment. Testable hypotheses arise because the shapeo f the hazard function depends on then ature of adjustment costs. Since hazard function estimation does not depend on the assumptions of a structural model, it avoids them easurement issues with the gap approach. Although this approach is not without measurement problems, as discussed at length below, they area rguably much less severe than the problems in measuring q. On the other hand,t he lack of a direct link between any findings and a specific model requires care in interpreting these findings.

I follow the second approach, primarily to minimize measurement problems. Specifically, I examine thei mplications of external finance constraints on a capital stock adjustment hazard. If a firm faces no external finance constraints but does face indivisible investment or nonconvex costs of adjustment, then its investment hazard should beu pward sloping. In other words, the probability of a large change increases as a spell of inaction lengthens. This sort of intuition comes from models with threshold effects, such as those in Caballero and Leahy (1996), Cooper and Haltiwanger (1998), Cooper, Haltiwanger, and Power (1999), and Caballero (1999). In these models large purchases followed by periods of inactivity are optimal.[2] If a firm faces fixed adjustment costs, it only invests when its capital stock is sufficiently far from the optimal level, otherwise preferring to remain inactive to avoid any lump-sum costs. Therefore, if a firm has recently adjusted, its desired capital stock is close to thea ctual, and the probability is low that it will undertake another large adjustment. However, thel onger thet ime since the last major investment, the morel ikely it is that cumulated productivity shocks and depreciation will havec hanged the marginal profit of capital sufficiently to warrant further investment. In other words, the hazard slopes up.[3]

In a world with financial frictions, the shape of the hazard depends both on adjustment costs and access to external finance. With nonconvex costs of physical adjustment, external financial constraints act as an additional cost of adjusting the capital stock, thereby furthering the delays between episodes of intense investment. In other words, an external finance premium increases the exercise price for the option to defer investment, thereby postponing exercise. Although the hazard wills till slope up, it will lie below that of an otherwise identical unconstrained firm. I use model simulations to derive these theoretical results, where the model incorporates both fixed adjustment costs and finance constraints.

I then test thei dea on a sample of firms and segments of firms from COMPUSTAT, classified into groups according to whether they ought to face external finance constraints or not. COMPUSTAT data have the obvious disadvantage that a firm, especially a large firm, is an aggregation of several different decision making units. If these individual units act in unison, then their behavior should resemble the behavior of an individual unit, and investment should occur episodically. However, this scenario is unlikely; and to the extent that these individual units act asynchronously, their aggregated investment will appear to be

[2]Models with fixed costs of adjustment have been used to show that lumpy adjustment and inactivity characterize a wide variety of economic decisions. For a modelo f inventories, see Caplin (1985); for a model of durables consumption, see Eberly (1994); for a model of capital structure, see Fischer, Heinckel, and Zechner (1989); and for a model of portfolio choice, see Vayanos (1998).

[3]The hazard for a firm with convex costs of adjustment is downward sloping because of intertemporal smoothing.

smoothed out over time. This problem could confound any empirical findings. For example, if I compare the investment of a firm composed of many units with the investment of a firm composed of one unit, then the two may have very differently shaped hazards even though both face no external finance constraints.

To mitigate this problem, I limit my sample to small conglomerate segments and small single-segment firms, since small segments or single-segment firms are less likely to be composed of a large number of decision making units. I then compare the investment hazards of *a priori* constrained and unconstrained groups of observations. First, I group small firms according to whether they have a consistent history of paying dividends or not. Because dividend payment is *prima facia* evidence of the availability of internally generated funds, one can assume that a firm that never pays dividends will be more likely to need external finance than one that does. Although dividend and investment decisions are determined simultaneously and although dividend behavior is an imperfect indicator of access to external capital markets, I argue below that neither of these concerns affects the size of my tests.

My second experiment is based on the idea that large conglomerates have better relationships with external capital markets than small single-segment firms, thus allowing the conglomerates egments access to less costly finance. Therefore, I compare the investment behavior of small segments of conglomerates with that of small single-segment firms. As discussed below, the experiment is also tangentially related to the idea in Williamson (1975) that conglomerates operate internal capital markets.

Finally, I examine the effect of the stock of liquid assets on the height of the estimated hazard. Two opposing ideas motivate this experiment. First, firms with high stocks of liquid assets may have access to cheaper external finance, since asset holdings are a characteristic of financial health. On the other hand, firms with low liquid assets may already have access to external finance and not need to carry liquid assets on their balance sheets. The direction of the effect of liquid assets on the hazard should distinguish these two hypotheses.

One contribution of this paper is evidence of episodic, lumpy investment in COMPUSTAT data. Therefore, using a framework of nonconvex adjustment to study external finance constraints has a strong basis in realism. The central contribution, however, is new evidence of the interdependence of finance and investment. I find that small firms that never pay dividends have upward sloping hazards, but that small firms that do pay dividends have significantly higher hazards. Further, although both small segments and small single-segment firms have upward sloping hazards, the firms have significantly lower hazards. These results are consistent with the presence of both nonconvex adjustment costs and finance constraints. Finally, in accord with the idea that good financial health lowers the cost of external finance, I find that the stock of liquid assets raises the estimated hazard.[4]

The rest of the paper is organized as follows. Section 2 presents a simple model that incorporates both fixed costs of capital stock adjustment and external finance constraints. Section 3 presents the results from simulating this model. Section 4 describes the data and a number of summary statistics. Section 5 contains the hazard model results, and section 6 concludes. The details of the simulation are in the appendix.

I. A Simple Model of Lumpy Investment

To motivate the empirical tests used below, and especially to provide structure for the choice of the control variables in my estimation, I consider a discrete-time partial-equilibrium model of a producer that uses current-period capital, K, to produce output. The producer's per period revenue function is given by $\Pi(K, z)$, where $\Pi(0, z) = 0$, $\Pi_z(K, z) > 0$, $\Pi_K(K, z) > 0$, $\Pi_{KK}(K, z) < 0$, and $\lim_{K\to\infty} \Pi_K(K, z) = 0$. z is a productivity shock that is observed by thep roducer before he makes his current period decisions, but not observed by the econometrician. It has support on the interval $(0, \infty)$ and has a stationary Markov transition function $q(z', z)$, where a prime denotes a variable in the subsequent period. $\Pi(K, z)$ can be thought of as a reduced-form production function wherev ariable factors of production have already been maximized out of the problem. The concavity of $\Pi(K, z)$ results from

[4]This paper leaves to further research the failure of previous empirical investment studies to isolatea ny differences between costly external finance and a hard finance constraint. Therefore, throughout the paper I use the terms "external finance constraints" and "costly external finance" interchageably. The word constraint should be interpreted as a surpassable obstacle, rather than an insurpassable obstacle.

decreasing returns in production and/or a downward sloping demand curve.

The firm purchases and sells capital at a price of 1 and incurs a fixed cost, cK, whenever investment is not equal to zero. The fixed cost is proportional to the capital stock so that the firm can never grow out of the fixed cost. This specification abstracts from many sources of nonconvexity in the firm's problem such as irreversibility or differing purchase and sales prices for capital, incorporating nonconvexities only through the fixed cost cK.

The capital stock evolves according to a standard capital stock accounting identity:

$$I \equiv K' - (1-d)K, \tag{1}$$

where d is the constant rate of depreciation, $0 < d < 1$. The producer is risk neutral and maximizes the value of future cash flows, discounting them at a constant factor, β, $0 < \beta < 1$. This model can be thought of either as a partial equilibrium model of a firm or, equivalently, as a model of a general equilibrium economy with production and consumption, where a representative consumer has utility linear in both consumption and leisure.[5]

Thus far the model is fairly standard and says nothing about financing costs. It would be ideal to model external finance costs endogenously. However, for the purpose of understanding the behavior of investment hazards, such an approach becomes analytically intractable. Therefore, I model external finance costs loosely after the idea of the pecking-order theory of capital structure, (Myers, 1984). Following Gomes (2001), I assume that whenever the optimal choice of I remains smaller than revenue, the firm uses internal funds for investment. However, whenever desired investment exceeds revenue, the firm can only proceed if it obtains external funds at a premium. This assumption can be thought of as the outcome of a an information theoretic model of external finance. To quantify the idea I define the excess of desired investment over revenue as $e(K,z) \equiv I - \Pi(K,z)$ and then specify a financing cost function $\phi(e(K,z))$, where $\phi(e(K,z)) = 0$ if the firm faces no external finance constraints or if $e(K,z) \leq 0$. If $e(K,z) > 0$, $\phi(e(K,z)) > 0$ and $\phi_e(e(K,z)) > 0$. Note the discrete difference between the cost of funds when the firm moves from internal to external sources.

[5]Adding risk aversion or decreasing marginal utility of leisuret o the model changes its quantitative but not qualitative predictions.

Although the financing function is uninformative about the source of external funds, it is nonetheless appropriate for a model that focuses on investment behavior.[6]

Let $V(K,z)$ denote current value of the firm and define it as:

$$V(K,z) = \max\left\{V^i(K,z), V^n(K,z)\right\}, \tag{2}$$

where the superscripts "i" and "n" refer to investment and no investment, respectively. The corresponding Bellman equations are

$$V^n(K,z) = \Pi(K,z) + \beta \int V(K(1-d), z')\, dq(z', z) \tag{3}$$

$$V^i(K,z) = \max_I \left\{\Pi(K,z) - I - cK - \phi(I - \Pi(K,z)) + \beta \int V(K', z')\, dq(z', z)\right\}. \tag{4}$$

A unique solution to this maximization problem requires that $\Pi_K(K,z)\left[1 + \phi'(e(K,z))\right]$ be decreasing in the capital stock. The existence of a unique solution to (2) is then guaranteed by Theorem 9.6 in Lucas and Stokey (1989). I characterize the solution to this problem by the value function $V(K,z)$ and the policy function $I = g(K,z)$.

II. Simulations

This model has no closed-form solution unless one makes homogeneity assumptions such as those in Caballero and Leahy (1996). Since such assumptions would eliminate any discrete difference between the costs of internal and external funds, I investigate the implications for the solution to this problem via simulation. In order to do so, I need to choose functional forms for the revenue, adjustment cost, and financing functions and the stochastic process for thep roductivity shocks. I also need to assign values to the fixed cost of adjustment, the discount factor, and thed epreciation rate. Because I am not literally estimating this structural model, the intent of the design is simply to generate qualitative conclusions that are robust to perturbations in the design parameters. Details are in the appendix.

[6] Instead of using a financing function to model costly external finance, one can specify that $I \leq A$, where A is the stock of financial assets, which are governed by the intertemporal budget constraint $A' = A/\beta + \Pi(K,z) - I$. This model produces identical qualitative conclusions.

I solve the model via value function iteration, which yields the policy and value functions. Whether constrained or unconstrained, the firm follows a two-sided (S, s) policy. If the marginal product of capital gets either too large or too small, the firm adjusts it back to a single return point. I then simulate the model for 10,000 time periods to generate the hazard functions for the simulated firms. In this simulation I define an adjustment or "spike" as a rate of net investment that exceeds 20%.[7]

Figure1 presents the hazard functions from thes imulations of thec onstrained and unconstrained firms. In this figure theh orizontal axis measures thea mount of time since the previous adjustment of the capital stock, and the vertical axis measures the adjustment hazard. Noticet hed ifference between theh azards of theu nconstrained and constrained firms. The hazard of theu nconstrained firm slopes upward steeply, which, as noted in thei ntroduction, is a pattern consistent with the presence of fixed costs of adjustment. The hazard of the constrained firm also slopes upward, but it is lower. Because the firm essentially faces an extra fixed cost of adjusting its capital stock, it will do so less frequently. The external finance function affects not only the cost of adjustment but the marginal productivity of capital. On the margin capital not only adds to production, but it alleviates the external finance premium. This addition to the marginal productivity of capital will raise the investment hazard, since a firm should adjust more often if it is more productive. However, because of decreasing returns to scale, the direct negative cost-of-adjustment effect is stronger.

Although the differences in theh azards of constrained and unconstrained firms is an empirically testable implication of my model, a number of other factors affect hazards: factors that need to be accounted for in any tests. One such important issue is aggregation of asynchronous actions within a firm. The effects of aggregation are illustrated in Figure 2, which contains graphs of the hazards from two types of "conglomerate" firms, where each type can be constrained or unconstrained. I construct the conglomerates by assuming

[7]Although I use a variety of thresholds in the following hazard estimations, for expositional brevity I limit myself to one threshold in my simulations.

they are composed of either two or six *i.i.d.* units, each of which is identical to the unit represented in Figure 1. For each type of conglomerate, I allow the individual units to be either all constrained or all unconstrained. Because the conglomerates are composed of *i.i.d.* units, they represent a worst-case scenario of the difficulties induced by aggregation, sincea firm composed of units whose decisions are positively correlated will behave in a manner more like an individual unit. As in Figure 1, the hazard for the constrained "small" conglomerate in Figure 2 lies below the hazard for the unconstrained small conglomerate, though both are lower than those in Figure 1. The pattern exhibited by the pair of hazards for the "large" conglomerates is quite different: both are at the samel ow level. Here, because of thea synchronous actions of the conglomerate units, and becauset he rate of investment contains total conglomerate assets in the denominator, ther ate of investment for the conglomerate as a whole rarely crosses a spike defining threshold, even though the individual units of the conglomerateb ehave exactly as thosed epicted in Figurel . Also, adding costly external finance to the model affects the hazard little, because the effect of aggregation dwarfs the effect of the finance constraint. This simulation result underlines the importance of considering aggregation when trying to uncover the effects of external finance constraints with COMPUSTAT data, which covers many large diversified firms.

Three further factors that could affect the heights of the hazards are productivity, adjustment costs, and depreciation. Starting with the model of an unconstrained firm, to model high productivity, I adjust the mean of the innovation of z to be 0.1 rather than 0; to model high adjustment costs, I double the value of c; and to model high depreciation, I double the value of d. The results from these experiments are in Figure 3, where the hazard of the more productive firm is higher than that of the unconstrained firm, the hazard of the firm with high adjustment costs is lower, and the hazard of the firm with high depreciation is higher. The intuition behind these findings is that firms with high productivity, low adjustment costs, or high depreciation should optimally want to replace capital more often, all elsee qual. This theoretical finding indicates thei mportance of controlling for all of these factors when comparing the hazards of different groups of firms.[8]

III. Data and Summary Statistics

My data are from the non-financial firms in the combined annual, research, and full coverage 2000 Standard and Poor's COMPUSTAT industrial files that are also covered by the ten most recent COMPUSTAT business information files. The latter contain data from 1983 through 1999. In late 1997 SFAS 131 changed the way in which firms define their segments. The concepts of industrial and geographic segments have been replaced by "operating segments" as defined by the company's management. This change renders data from 1998 inconsistent with earlier data. BecauseI want long consistent times eries on the segments, I only use data from 1984 until 1997.

I select the sample by first deleting any firm-year observations with missing data. Next, I delete any observations for which total assets, the gross capital stock, or sales are either zero or negative. Omitting the firm-year observations that report zero sales is important because I use sales growth as a proxy for a component of investment opportunities. Further, I delete any observations if thes um of segment assets deviates by more than 25% from reported total firm assets or if sales growth is greater than two hundred percent or less than one hundred percent. Finally, I include a firm or segment only if it has at least four consecutive years of complete data; and I omit all firm- and segment-level observations whose primary SIC classification is between 4900 and 4999 or between 6000 and 6999, since my model of investment is inappropriate for regulated or financial firms. I end up with between 1302 and 2953 single-segment firms per year, between 1091 and 1463 multiple-segment firms per year, and between 1978 and 2589 segments of multiple-segment firms per year.[9]

---

[8]One further factor that could affect the hazard is the variance of thei nnovation to z. However, changes in this variablea ffect the hazard little under a wide variety of model parameterizations. Intuitively, when z has a high variance, the marginal product of capital is morel ikely to hit oneo f thet hresholds. On the other hand, thefi rm will respond by widening thei naction interval. Theset wo effects appear to cancel each other out. For similar intuiton, see Bertola and Caballero (1990).

[9]The number of segments is close to the number of firms, because even though the firm-level data on the conglomerate may be complete, the segment-level data often is not.

Table 1 provides summary statistics for my three key variables: the ratio of capital expenditures to total assets, sales growth, and total assets. Not surprisingly, the conglomerates are substantially larger than their segments or the single-segment firms. The smaller size of the conglomerate segments and the single-segment firms suggests that aggregation issues may be less important for these two groups. Notice also that the segments of conglomerates tend to be larger than stand-alone firms: the second column of the table shows that the median size of a segment of a conglomerate is more than twice the median size of a stand-alone firm. One final feature of this table stands out: the single-segment firms have much higher sales growth than the conglomerates or the segments of conglomerates, even though their investment is not substantially higher. This pattern is loosely consistent with the idea that external finance constraints affect investment.

Next I examine the extent to which firms engage in large investment projects. Since most firms in COMPUSTAT invest at least a small bit every period, the definition of a large project requires thought. Low observed rates of investment probably occur because of maintenance and because some types of investment may well be subject to convex adjustment costs. However, when a firm undertakes a large project, we ought to observe a much higher than normal rate of investment. To capture this idea I define an investment "spike" in terms of "net" investment—the ratio of investment to total assets minus the firm average ratio of depreciation to total assets. I subtract depreciation in order to concentrate on "above-normal" investment. Also, I use average depreciation to smooth out the high observed annual fluctuations in this variable, since this variation can be attributed more to accounting practices than changes in economic depreciation. I then define a spike as an observation in which net investment is greater than either 0.10, 0.15, or 0.20.[10] I use several spike thresholds in order to check the robustness of my results to the criteria for measuring spikes.

[10] My definitions of spikes have lower cutoff points than the 0.2 level used in Cooper, Haltiwanger, and Power (1999) for two reasons. First, because the segment data do not contain any information on the capital stock, I define the rate of investment as the ratio of capital expenditures to total assets, whereas they define it as the ratio of capital expenditures to the capital stock. This latter ratio is on average twice as large as my measure of the rate of investment. Second, while they define spikes in terms of gross investment, I define them in terms of net investment.

Table 2 shows the fraction of observations for which net investment is less than zero, the fraction for which gross investment is equal to zero, and the fraction consisting of investment spikes. Two types of evidence in this table providep *rima facia* evidence of fixed costs of adjustment. First, the second column shows that for twenty to thirty percent of the observations, investment falls below depreciation. If inaction is defined as a period in which the firm allows its capital stock to depreciate, then this evidence of a great deal of inaction favors models with nonconvex costs. Second, in a world with convex adjustment, I ought to see very few rates of investment greater than any of my spike thresholds. The table shows that this is not the case. Note that the single-segment firms exhibit more investment spikes than either the conglomerates or their segments. This observation is consistent with the idea that larger firms or segments contain more decision making units. In these firms even though individual decision making units may be experiencing investment spikes, when the outside observer measures these spikes as a percentage of total assets, they appear small. Also the percentage of single-segment firms experiencing spikes is not much smaller than the 14 percent figure reported by Cooper and Haltiwanger (1998). This similarity suggests that the sort of lumpy adjustment observed in plants may also be present in firms.

One aspect of investment behavior remains to be examined: the timing of these spikes. As a first pass, I calculate histograms of what I term "inaction spells;" that is, the length of time between spikes. For any given firm or segment, if the final spell is ongoing during the last year of data on the firm, I cut the length of the spell off at that year and then note that the observation is censored. Because in my hazard estimation I truncate any spells longer than seven years, I present only the first seven spell lengths. Figure 4 contains these histograms for single-segment firms, conglomerates, and individual segments of these conglomerates. In each group and for each spike threshold most of the observations are bunched at year one. However, the figures hows that many firms do experience long inaction spells: a strategy that is unlikely to be optimal in the face of convex adjustment costs.

Next, I calculate summary statistics on these inaction spells, which arec ontained in Table 3. For each of my three groups of observations, I present the number and average length of spells corresponding to each of my spike-defining thresholds. The results for each of these groups are broadly similar, though the conglomerates do tend to have slightly longer spells than either thes ingle-segment firms or the segments. As in thes imulations of the previous section, this result is likely duet o the aggregation of asynchronous actions. One interesting feature of the table is the difference between the segments and the single-segment firms. Their mean spell lengths are quite similar, yet the segments have substantially lower sales growth. Given the difference in sales growth, one would expect the single-segment firms to be adjusting more often; so perhaps external finance constraints are hindering adjustment. Providing more specific evidence of this conjecture is the subject of the rest of the paper.

IV. Estimation

A. Methods

My theoretical model predicts that the hazard of an investment spike will, *ceteris paribus*, be lower for a firm that must use costly external finance than for a firm that does not. A number of different techniques exist for estimating hazard functions. The simplest method consists of calculating for each length of an inaction spell and for each year, ther atio of the number of firms that experience spikes to then umber of all firms that have remained inactive for at least as long. These simple empirical hazards could then be compared to the simulated hazards. However, this approach can lead to biased hazard function estimates unless onec ontrols for cross-sectional heterogeneity. To seet his point in the context of investment spikes, suppose we observe a cross section containing two types of firms that face fixed adjustment costs: low cost and high cost. Suppose also that there are twice as many low cost firms as high cost firms. If we could observe a long time series on each firm, all would have upward sloping hazards. However because the low cost firms replace their capital more often than the high cost firms, in a cross section we see more replacements of relatively new capital than of older capital, and a simple empirical hazard will slope downward.

Difficulties such as this can be solved by using a duration model, since it is possible to account for observable time-varying covariates, such as productivity, as well as unobservable heterogeneity across firms. Loosely speaking, an empirical hazard can be thought of as a sort of histogram, whereas the results from estimating a duration model can be thought of as a "conditional" histogram. The most likely candidate for the source of unobservable heterogeneity is thel evel of adjustment costs, since I can control for theo ther important non-financial determinants of investment. Caballero and Engle (1999) emphasize that cross-sectional heterogeneity in adjustment costs is likely to exist, and they find that a structural investment model that allows for heterogeneity.explains aggregate investment better than a model that does not. Using a model that incorporates cross-sectional heterogeneity lowers thep robability that my results are an artifact of an incidental correlation between real adjustment costs and measures of access to external financial markets.

To control for heterogeneity, I uset hee stimation technique in Meyer (1990), which accounts for observable and unobservable heterogeneity, and which allows the shape of the hazard to be estimated nonparametrically. The following brief description of this technique follows Meyer (1990) closely. First, I define theh azard function for an individual firm i, $\lambda_i(t)$. Let T_i be the length of a firm's inaction spell. The hazard is defined as

$$\begin{aligned}\lambda_i(t) &\equiv \lim_{\Delta\to 0^+} \frac{\Pr\left(t+\Delta > T_i \geq t \mid T_i \geq t\right)}{\Delta} \\ &= \frac{f(t)}{1-F(t)},\end{aligned}$$

where $F(t)$ is the distribution of the inaction spells and $f(t)$ is their density. This definition is quite general, and to estimate the hazard it is necessary to impose some structure on its form. I use a proportional hazards specification:

$$\lambda_i(t) = \lambda_0(t)\exp\left(x_i(t)'\beta\right),$$

where$x_i(t)$ is a column vector of covariates, β is the corresponding vector of unknown coefficients, and $\lambda_0(t)$ is called the baseline hazard. The parametric part of this specification is the linear modelling of thec ovariates. Then onparametric part is the baselineh azard, which is not restricted to take any particular shape. Note that the existence of the covariates allows the hazard to shift up and down depending on their values and on β. Equivalently, the existence of the covariates essentially changes the units in which time is measured.

Estimating the parameters of a hazard function is exactly analogous to estimating the parameters of a density. However, since time is measured at discrete intervals, to write down a likelihood function, it is convenient to express the hazard at time t as the product of the hazards for the time intervals leading up to time t. As a first step in this process, note that the probability that an inaction spell will last until time $t+1$, given that it has lasted to time t can be written as a function of $\lambda_i(t)$ as follows:

$$
\begin{aligned}
\Pr(T_i \geq t+1 \mid T_i \geq t) &= \exp\left[-\int_t^{t+1} \lambda_i(s)\, ds\right] \\
&= \exp\left[-\exp\left(x_i(t)'\beta\right)\int_t^{t+1} \lambda_0(s)\, ds\right].
\end{aligned} \tag{5}
$$

Define

$$
\gamma(t) \equiv \ln\left(\int_t^{t+1} \lambda_0(s)\, ds\right).
$$

Then (5) can be written as

$$
\Pr(T_i \geq t+1 \mid T_i \geq t) = \exp\left[-\exp\left(x_i(t)'\beta + \gamma(t)\right)\right].
$$

I can use this specification to write down the likelihood function. First, define C_i as the censoring time for an individual inaction spell. For example, if a firm experiences a spike in 1994, if it never experiences another, and if the data on the firm end in 1997, the censoring time is three. I also censor any spell lengths longer than seven years, where I have chosen this number because it is the largest for which I can estimate all of the elements of $\gamma(t)$ for all of my samples. Depending on the spike threshold and sample, this seven-year rule affects from 0.6 to 5 percent of the observations. The likelihood function for a sample of N *i.i.d.* individual spells that accounts for this sort of right censoring can be written as

$$
\mathcal{L}(\gamma,\beta) = \prod_{i=1}^{N}\left\{\left[1-\exp\left(-\exp\left(x_i(h_i)'\beta+\gamma(h_i)\right)\right)\right]^{\delta_i} \times \prod_{t=1}^{h_i-1}\exp\left(-\exp\left(x_i(t)'\beta+\gamma(t)\right)\right)\right\},
$$

where $\gamma \equiv [\gamma(0), \gamma(1), \ldots \gamma(T-1)]'$, $\delta_i = 1$ if $T_i \leq C_i$ and 0 otherwise, and $h_i = \min(T_i, C_i)$. Thefi rst term in squareb rackets is 1 if thei naction spell is censored, and the second term is just the probability of a spell lasting at least until h_i. The corresponding log-likelihood can be written as

$$L(\gamma,\beta)=\sum_{i=1}^{N}\left[\delta_i \ln\left[1-\exp\left(-\exp\left(x_i(h_i)'\beta+\gamma(h_i)\right)\right)\right]-\sum_{t=1}^{h_i-1}\exp\left(x_i(t)'\beta+\gamma(t)\right)\right]. \tag{6}$$

As this point, although thee mpirical model can account for observable heterogeneity via the inclusion of $x_i(t)$, it still does not account for unobserved heterogeneity. To address this issue, I once again follow Meyer (1990) and assume that unobserved heterogeneity takes a multiplicative form:

$$\lambda_i(t)=\omega_i\lambda_0(t)\exp\left(x_i(t)'\beta\right).$$

Here ω_i is a random variable that is assumed to be independent of $x_i(t)$. To construct a tractable log-likelihood function, one usually assumes a parametric functional form for the distribution of ω_i. A commonly-used distribution is a gamma with a mean of one. In this case the log-likelihood is

$$\begin{aligned} L(\gamma,\beta) \;=\; & \sum_{i=1}^{N}\ln\left\{\left[1+\sigma^2\sum_{t=1}^{h_i-1}\exp\left(x_i(t)'\beta+\gamma(t)\right)\right]^{-\left(1/\sigma^2\right)}\right. \\ & \left.-\delta_i\left[1+\sigma^2\sum_{t=1}^{h_i}\exp\left(x_i(t)'\beta+\gamma(t)\right)\right]^{-\left(1/\sigma^2\right)}\right\}, \end{aligned} \tag{7}$$

where σ is the variance of the gamma distribution. The estimation procedure chooses the shape of the hazard to maximize the likelihood of observing the inaction spells in the sample.

My specification of the model allows $x_i(t)$ to contain sales growth, two-digit industry dummies,a nd year dummies. The year dummies capture the effects of interest rates and the business cycle on individual investment decisions. The industry dummies capture several important factors that could affect the hazard. First, differences in competitiveness across industries could have a strategic affect on a firm's decision to invest. Second, differences in the types of capital used across industries could affect depreciation rates, differences in returns to scale, and adjustment costs.

I also need somec ontrol for "investment opportunities." Given adjustment costs, the adjustment hazard depends only on thep robability of thefi rm reaching an adjustment trigger, which in turn depends primarily on the expected rate of increase in the marginal product of capital. Therefore, in this model "investment opportunities" is not equivalent to marginal q, but to this speed. Indeed, Caballero and Leahy (1996) show that in the presence of fixed costs of adjustment that occur with each project (as opposed to per unit of time), the relationship between investment and marginal q is not a function but a correspondence. The non-financial determinants of the rate of reaching an adjustment trigger are the mean of the innovation to the z shock, the depreciation rate, and returns to scale. Because technology determines the second two factors, and becauset echnology is likely to vary morea cross industries than across firms within an industry, the industry dummies control for much of the variation in these components of investment opportunities. I include sales growth as an admittedly imperfect proxy for the mean of the innovation to the z shock.

Although the proxy is imperfect, the model *can* shed some light on its quality,a nd sales growth is, in itself, an easily-measured variable. It is straightforward to show in my model that if the firm has a homogeneous revenue function and if the firm invests just enough to replace depreciated capital, expected sales growth is equal to the population mean of the innovation, and observed sales growth equals the innovation. Since optimal behavior implies that firms either remain completely inactive or invest in large spikes, average sales growth will bea n imperfect proxy for them ean of thei nnovation. However, if in a simulation mean sales growth and the innovation mean are highly correlated, the proxy should at least theoretically be of high quality. Using the model of the previous section I simulate 1000 firms, where for each the innovation mean is drawn from a truncated lognormal distribution. The squared cross-sectional correlation between the innovation mean and average sales growth is 0.823. In other words variation in average sales growth accounts for 82.3 percent of the variation in the innovation mean.[11] Because this figure is quite high, I have confidence that the qualitative nature of the following results is unlikely to be an artifact of measurement error. In contrast, thec orresponding figure for Tobin's q found by Erickson and Whited

[11]Using actual instead of average sales growth in the estimation produces identical results, since the estimation procedure averages observed sales growth.

(2000) is only approximately 40 percent.

B. Results

I use this model to compare the estimated hazards of firms and segments categorized along three dimensions: size, diversification, and access to external financial markets. First, however, I examine the importance of controlling for cross-sectional heterogeneity. For brevity, I use only small single-segment firms to examine this issue, primarily because this group is involved in all of my tests. I do, however, note below any departures from the results for this group. I define an observation from a small single-segment firm as one from a firm whose real assets in the year the inaction spell ends lie below the stand-alone firm median in that year. This year-by-year definition of a "small" firm controls for real firm growth.

Table 4 presents ther esults from estimating two models with data from small single-segment firms: one model controls for cross-sectional heterogeneity and the other does not. Each column contains the model estimates for each of the different spike thresholds—0.10, 0.15, and 0.20.[12] Note first that the log-likelihoods from the models on the right, which do not control for heterogeneity, are much lower than thel og-likelihoods from the models on the left, which do control for heterogeneity. Indeed, standard likelihood ratio tests produce rejections of all of the no-heterogeneity models. Consistent with this result is the significance of the estimates of the heterogeneity variance, which arel abeled σ^2. Another contrast in the model estimates can bef ound in thec oefficients on sales growth. As expected, they are positive and significant for the model with heterogeneity: higher sales growth increases the probability of ending an inaction spell. However, they are negative and insignificant for the model without heterogeneity. This latter anomalous result adds further support to the hypothesis that the no-heterogeneity model is not the right one.

The most interesting aspect of the table is the comparison of the baseline hazards, which are indicated by the entries labeled "$\exp(\gamma_i)$." For example, the estimate of $\exp(\gamma_3)$ is the probability of ending an inaction spell, conditional on the spell lasting at least three years.

[12] The estimated coefficients on the year and industry dummies are omitted. In all of the results that follow, most of these estimates are significantly different from zero.

Note that all of the estimated hazards from the heterogeneity model slope upwards; that is, the estimates of $\exp(\gamma_i)$ increase with i. This result is broadly consistent with one of the predictions from the theoretical model: fixed costs of adjustment result in upward sloping hazards. In contrast, the estimated hazards from the no-heterogeneity model slope downwards, mirroring the unconditional histograms presented in Figure 4. The stark differences in the slopes of the baseline hazards from the two models can be seen in Figure 5, which presents the estimates from the twenty percent threshold columns. The difference in slope indicates that for the small single-segment firms, the cross-sectional pattern of investment spikes does not match the pattern for any particular firm. The result can be understood as follows: the inclusion of heterogeneity in this hazard model acts loosely as the inclusion of fixed effects in a panel-data model in that both allow the behavior of an individual to stand out. Given the superior performance of the heterogeneity model, all results that follow will be from this specification.

Next I examine the hazards from groups of firms categorized by size and diversification. As explained in the previous section, investment by large firms or segments is more likely to be thep roduct of aggregated asynchronous decisions, and hazard estimates from these groups are unlikely to be upward sloping. Table 5 contains hazard model estimates for three groups of observations: those from large single-segment firms, small conglomerates, and large conglomerates. In order to render my size classifications comparable to those for the single-segment firms,I define "small" using the year-by-year medians of the single-segment firms. The most important result in Table 5 is the low flat baseline hazards for all three groups of firms. This finding is consistent with the prediction of the theoretical model that aggregation lowers hazards. A related difference is the insignificance of the estimates of the heterogeneity variance. Roughly speaking, for these firms, the cross-sectional distribution of spells is close to the flat individual time-series distribution. Interestingly enough, the aggregation effect even appears in thes mall conglomerates. Not only is small size a necessary condition for eliminating aggregation effects, but low diversification is as well. Finally, this evidence has an important implication for the empirical investment literature that uses firm size as a proxy for access to external financial markets. Any differences in the behavior of small versus large firms may be due to aggregation issues rather than finance constraints.[13] The evidence in Table 5 does admit, however, a clear alternative interpretation: aggregation issues are unimportant, conglomerates and large single-segment firms face finance constraints, and small single-segment firms do not. Although implausible, this possibility begs for an attempt to disentangle finance effects from aggregation effects.

To do so, I concentrate my analysis on small single-segment firms and small segments of conglomerates, focusing on my central categorization criterion—access to external financial markets. By excluding large stand-alone firms, large segments, and diversified firms from the analysis, I hope to mitigate the aggregation effects that could contaminate my results. The null hypothesis for the two tests that follow is that theb aseline hazard is thes ame across groups of observations classified by indicators of access to external finance. Note that the rest of the hazard is allowed to vary across groups. The alternative hypothesis is that the baseline hazard for a "financially constrained" group of firms is lower than the baseline hazard for an "unconstrained" group. Structuring the null and alternativeh ypotheses in this way is important, because neither of my sample-splitting variables will produce perfect sorting of observations into constrained and unconstrained groups, and because imperfect sorting will only lower the power of the tests but will not affect the size.

The first experiment along this line is a comparison of small single-segment firms who differ in their dividend policies. As explained in the introduction, the "constrained" group will consist of observations from firms with a consistent history of paying no dividends before the end of an inaction spell. The "unconstrained" group will consist of all other observations. Using lagged dividend behavior as a classification variable mitigates the simultaneity problem that arises because dividends and investment are joint decisions. In other words, this sample splitting variablei s at the very least predetermined, if not exogenous. More importantly,

[13]Studies that use size as a proxy for financing constraints include Gilchrist and Himmelberg (1995), Kaplan and Zingales (1997), and Erickson and Whited (2000). The first claims that small firms appear more constrained than large firms; thes econd claims that small firms appear less constrained than large firms; and the third finds no difference.

using a consistent history of no dividends as a splitting criterion increases the likelihood that the dividend and investment decisions are independent. Table 6 shows that the estimated baseline hazards for both groups are upward sloping and that those for the dividend group are higher than those for the no dividend group. These differences are significant at the five percent level in years two through seven for all thresholds. This difference is illustrated in Figure 6, which shows the baseline hazards for the 20 percent threshold.

As a second test of the idea that external finance constraints lower investment hazards, I run separate hazard models on my group of small single-segment firms and a group of same-sized segments of conglomerates. Estimates from these models are in Table 7, which shows that the firms have lower hazard rates than the segments. Indeed, two-sided t-tests of the null that the hazard rates are the same across groups produce rejections at the five percent level in all but four of the twenty-one instances. The difference in the hazard rates for the 0.20 threshold is illustrated in Figure 7. To the extent that belonging to a conglomerate is an indicator of easy access to finance, these results also are consistent with the idea that external finance constraints can affect investment. The results are not, furthermore, an artifact of the segments belonging to different industries than the firms. Ninety-one percent of the two-digit industries represented by the sample of small firms are also represented by the sample of small segments, and eighty-seven percent of the two-digit industries represented by the sample of small segments are also represented by the sample of small firms. However, the result says littlea bout thee fficiency of internal capital markets, since any test of efficiency needs to compare the behavior of all the segments within a firm. At the very least, it appears that internal capital markets are not inefficient *enough* to render the behavior of small segments the same as the behavior of small single-segment firms.

As a final test I include as an explanatory variable lagged net liquid assets: net working capital less inventories, where this quantity is expressed as a fraction of total assets. The intent is to see if accumulation of liquid assets precedes an investment spike. I do not split the sample in this case, since liquid assets are a continuous variable and have no clear break point. As noted in the introduction, liquid assets could have two opposing effects on the height of the hazard. First, small firms with low liquid asset positions may have limited access to debt markets, presumably because they lack the collateral to back their borrowing. Therefore, liquid assets should have a positive coefficient in the hazard model. Further, an accumulation of liquid assets could indicate the presence of financial constraints if the firm needs to save thef unds for a large project rather than obtaining them externally. This behavior is also consistent with a positive coefficient. In contrast, Opler, Pinkowitz, Stulz, and Williamson (1999) note that firms with access to external financial markets do not *need* to keep stocks of liquid assets on hand. In this case the coefficient on liquid assets should be negative. An insignificant coefficient could mean one of two things: the above two effects offset one another, or financea nd investment are independent. Ther esults from a model that includes liquid assets arei n Table8 . As in the other hazard models for the small single-segment firms, the estimated hazards slope upwards. The number of observations is lower here, because some of the firms do not report information on liquid assets. Note the positive estimates of the coefficients on liquid assets, which are significant for the 10 and 15 percent thresholds. This result supports the role of liquid assets as a sign of financial health or saving behavior, and it is consistent with the presence of external finance constraints.

One possible alternative explanation for all of these results is based on the idea that firms respond to financec onstraints by undertaking reduced-size projects more often. In that case constrained firms should have higher hazards, and the interpretations of all of the above results should be reversed. However, because the results are robust to the choice of several different spike thresholds, I view this scenario as unlikely.

V. Conclusion

This paper has tackled the question of the interaction between finance and investment from a new angle—one that examines the timing of large indivisible investment projects. One contribution of this different approach is its basis in a realistic view of firm investment decisions. Instead of relying on predictions from models with smooth costs of adjustment, the paper operates on the assumption that the most important costs of adjusting the capital stock are fixed. This choice stems from the intuitive observation that external finance constraints are more likely to affect large investment projects than incremental additions to the capital stock. A second advantage is of this approach deals with measurement issues. I have argued that because my model offers guidance in finding a simple, easily measured control for productivity, the measurement issues are not as severe as those facing regressions of investment on qa nd cash flow. Finally, for researchers interested in the interaction between finance and investment, a new anglea ppears necessary, given thec ontradictory and inconclusive evidence from a decade and a half of cash-flow sensitivity tests.

Within this framework I use a simple theoretical model to show that, *ceteris paribus*, costly external finance lowers the hazard function for investment spikes. In other words, given that a firm has not undertaken a large investment project for a certain length of time, it is less likely to undertake another if it faces costly external finance than if it does not. I also demonstrate that the aggregation of decisions in large firms can mask this result.

I take this idea to data by using a hazard model in which I control for firm size, industry, macroeconomic effects, and a proxy for productivity. First, I find evidence of lumpy investment in firm-level data, which adds credence to the idea of testing for financial constraints in the context of fixed costs of physical adjustment. Second, I find that aggregation of asynchronous decisions affects investment hazards. As predicted by my fixed-costs model, small single-segment firms have upward sloping hazards; and large single-segment firms and both largea nd small conglomerates have lower hazards than small single-segment firms. This result casts doubt on the common practice of using size as a proxy for access to external capital markets. The most important result is evidence that access to cheap finance lowers investment hazards. Small single-segment firms that pay dividends have significantly higher hazards than small single-segment firms that do not. In addition, small single-segment firms have significantly lower hazards than small segments of conglomerates. Finally, stocks of liquid assets raise hazards.

In sum, the paper has provided a new type of evidence that access to external finance does indeed influence firms' real investment decisions. Because looking for evidence of finance constraints in the context of models with real nonconvexities appears to be fruitful, future research could indeed explore other ways to exploit thesem odels. Once avenuec onsists of looking at plant-level data. Other, more methodological avenues include simulation and structural estimation. One challenget o structural estimation is the lack of closed-form solutions for many models with nonconvexities—a challenge possibly solved with simulation estimators.

Appendix

I start by assuming that production takes place according to

$$\Pi(K,z) = zK^{\theta}. \tag{8}$$

Ideally, I would like to estimateθ with my data from COMPUSTAT. However, because these data do not contain sufficient information on payments to variable factors to estimate a production function, I follow Cooper and Haltiwanger (1998) and set $\theta = 0.51$.

Next I consider the financing function, whose shape requires considerably more thought. External finance may be more costly than internal finance for several reasons. First, information asymmetries may external investors to require a "lemons" premium. Similarly, external investors may require premia because external equity exacerbates manager-shareholder conflicts, and because debt can cause underinvestment problems. Second, monitoring costs are important for bank loans, and transactions costs are important for seasoned debt and equity offerings, as well as bank loans. Because little research has been done to quantify the first type of costs, I follow Gomes (2001) and focus only on transactions costs. This strategy will provide a very conservative estimate of the costs of external finance. To quantify these costs I use the estimates in Altinkilic and Hansen (2000) for seasoned equity issues. (See their Table 2.) Their regression results imply an external finance function of the form

$$\phi(e) = 0.0404 + 0.0264(e), \tag{9}$$

where e is a dummy variable for the gross amount of financing as a percentage of firm assets. To find a value for the fixed cost, c, I turn to Cooper, Haltiwanger, and Power (1999) and set $c = 0.2$. Finally, I set the discount ratee qual to 6%, which implies a discount factor $\beta = 0.943$; and I set the depreciation rate equal to the average in my data of depreciation divided by total assets: 0.047. Fifty percent changes in thea bove parameters result in identical qualitative conclusions.

Next, I specify a stochasticp rocess for the shock, z. Following Caballero and Leahy

(1996), I assume that z follows a random walk in logs,

$$\ln(z') = \ln(z) + u', \tag{10}$$

where $u' \sim N(0, 0.11^2)$.

Finally, to find a numerical solution I need to specify a finite st ates pacef or thet wo state variables. I let the capital stock lie on the points

$$\left[k^*(1-d)^{30}, \ldots, k^*(1-d)^2, k^*(1-d), k^*, k^*/(1-d), \ldots, k^*/(1-d)^{30}\right],$$

where k^* is the steady-statec apital stock of a model without any adjustment costs. I let the productivity shock have nine points of support, transforming (10) into a discrete-state Markov chain using the method in Tauchen (1986).

References

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.

Bertola, Giuseppe and Ricardo J. Caballero, 1990, Kinked adjustment costs and aggregate dynamics, *N.B.E.R. Macroeconomics Annual,* 237–288.

Bond,S tephen and Costas Meghir, 1994, Dynamic investment models and the firm's financial policy, *Review of Economic Studies* 61, 197–222.

Caballero, Ricardo J., 1999, Aggregate investment, in John B. Taylor and Michael Woodford, eds.: *Handbook of Macroeconomics, Volume 1A.* (North Holland, Amsterdam).

Caballero, Ricardo J. and Eduardo M. R. A. Engel, Explaining investment dynamics in U.S. manufacturing: a generalized (S,s) approach, *Econometrica 67*, 783–826.

Caballero, Ricardo J., Eduardo M. R. A. Engel, and John C. Haltiwanger, 1995, Plant level adjustment and aggregate investment dynamics. *Brookings Papers on Economic Activity* 2, 1-39.

Caballero, Ricardo and John V. Leahy, 1996, Fixed costs: the demise of marginal q, N.B.E.R. Working Paper 5508.

Caplin, Andrew S., 1985, The variability of aggregate demand with (S,s) inventory policies, *Econometrica 53*, 1395-1409.

Cleary, Sean, 1999, The relationship between firm investment and financial status, *Journal of Finance* 54, 673–692.

Cooper, Russell and John Haltiwanger, 1998, On the nature of capital adjustment costs, Boston University Working paper.

Cooper, Russell, John Haltiwanger, and Laura Power, 1999. Machine replacement and the business cycle: lumps and bumps, *American Economic Review* 89, 921–46.

Cooper, Russel and Jonathan Willis, 2000, The economics of labor adjustment: minding the gap, Boston University Working paper.

Doms, Mark and Timothy Dunne, 1998. Capital adjustment patterns in manufacturing plants, *The Review of Economic Dynamics* 1, 409-429.

Eberly, Janice C., 1994, Adjustment of consumers' durables stocks: evidence from automobile purchases. *Journal of Political Economy* 102, 403-36.

Erickson, Timothy and Toni M. Whited, 2000, Measurement error and the relationship between investment and q, *Journal of Political Economy* 108, 1027-57.

Fazzari, Steven, R., Glenn Hubbard, and Bruce C. Petersen, 1988, Financing constraints and corporate investment, *Brookings Papers on Economic Activity* 1, 141–195.

Fischer, Edwin O., Robert Heinkel, and Josef Zechner, 1989, Optimal dynamic capital structure choice: theory and tests, *Journal of Finance* 44, 19-40.

Gilchrist, Simon and Charles P. Himmelberg, 1995, Evidence on the role of cash flow for investment, *Journal of Monetary Economics* 36, 541-572.

Gomes, Joao,2 001, Financing investment, *American Economic Review*, in press.

Kaplan, Stephen, and Luigi Zingales, 1997, Do financing constraints explain why investment is correlated with cash flow? *Quarterly Journal of Economics* 112, 169–216.

Meyer, Bruce D., 1990, Unemployment insurance and unemployment spells, *Econometrica* 58, 757–782.

Myers, Stewart C., 1984, The capital structure puzzle, *Journal of Finance* 39, 575–592.

Opler, Tim, Lee Pinkowitz, Rene Stulz, and Rohan Williamson, 1999, The determinates of corporate cash holdings, *Journal of Financial Economics* 52 3-49.

Stokey, Nancy L. and Robert E. Lucas with Edward C. Prescott, 1989, *Recursive Methods in Economic Dynamics.* (Harvard University Press, Cambridge, Mass. and London)

Tauchen, George, 1986, Finite state Markov-chain approximations to univariate and vector autoregressions, *Economics Letters* 20, 177-181.

Vayanos, Dimitrios, 1998, Transaction costs and asset prices: a dynamic equilibrium model, *Review of Financial Studies* 11, 1-58.

Whited, Toni M., 1992, Debt, liquidity constraints, and corporate investment: Evidence from panel data, *Journal of Finance* 47, 1425-1460.

Williamson, Oliver E., 1975, Markets and hierarchies: Analysis and antitrust implications. (Collier McMillan Publishers, Inc., New York).

Figure 1
Theoretical Adjustment Hazards
Single Segment Firms

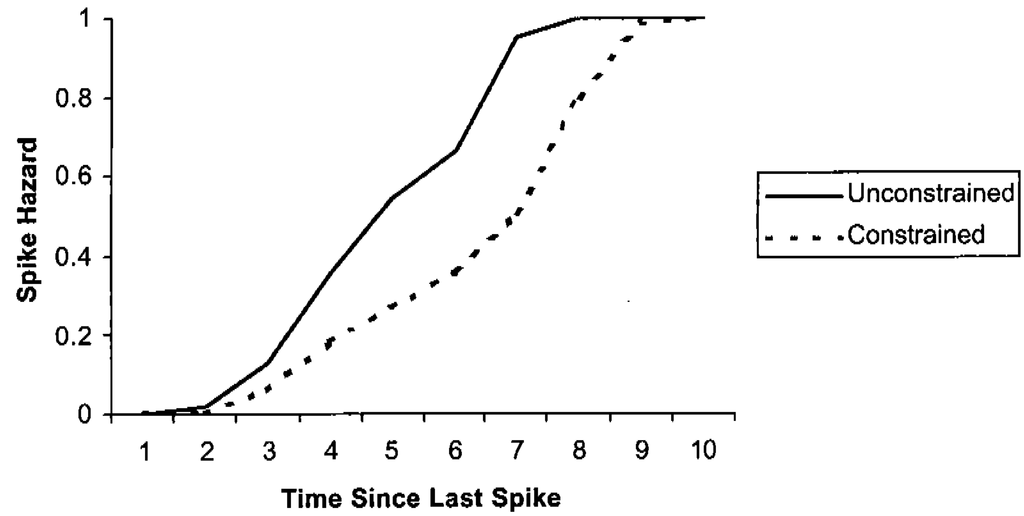

The hazard functions are simulated from the investment model in Section 2. The horizonatal axis measures the length of time since the last investment spike, and the vertical axis measures the probability of a spike, given that the firm has remained inactiveu p to that time. "Baseline" refers to a firm without costly external finance, and "constrained" refers to a firm with costly external finance.

Figure 2
Theoretical Adjustment Hazards
Conglomerates

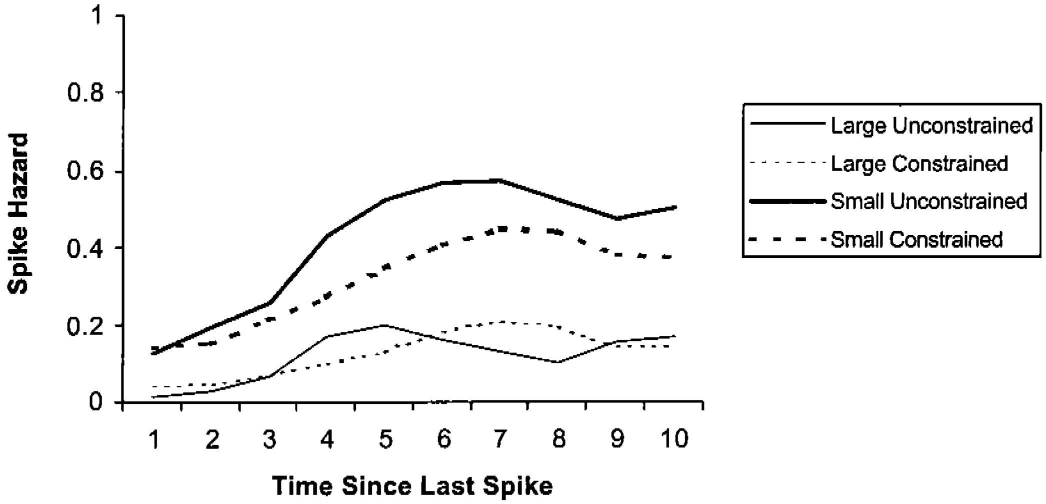

The hazard functions are simulated from the investment model in Section 2. The horizonatal axis measures the length of time since the last investment spike, and the vertical axis measures the probability of a spike, given that the firm has remained inactiveu p to that time. "Baseline" refers to a firm without costly external finance, and "constrained" refers to a firm with costly external finance. A "large" conglomerate contains six units, and a "small" conglomerate contains two units.

Figure 3
Theoretical Adjustment Hazards
Single Segment Firms

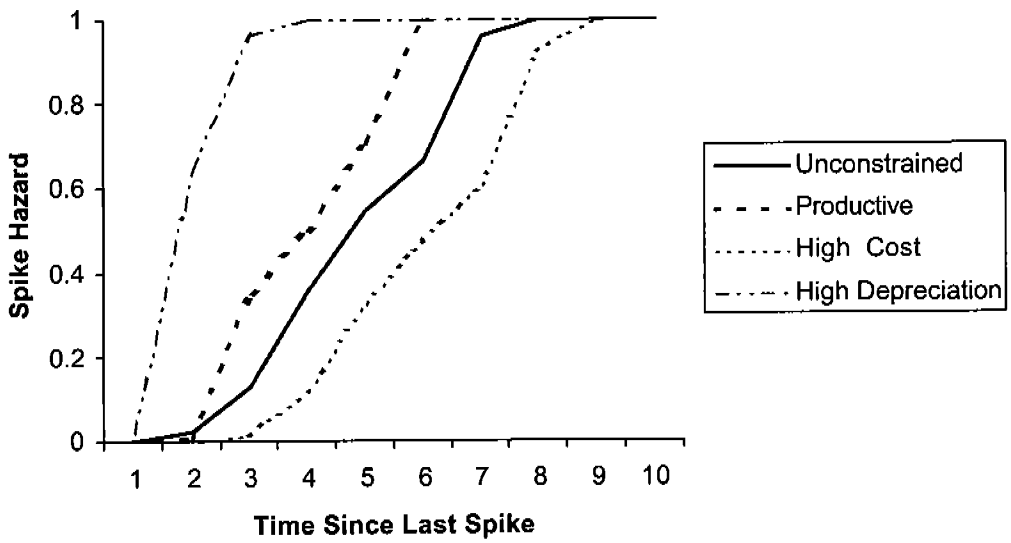

The hazard functions are simulated from the investment model in Section 2. The horizonatal axis measures the length of time since the last investment spike, and the vertical axis measures the probability of a spike, given that the firm has remained inactiveu p to that time. "Baseline" refers to a firm without costly external finance. All other hazards result from perturbations of the basic model. "Productive" refers to a firm with a shock to total factor productivity with a mean of 0.1 instead of 0. "High cost" refers to a firm with a doubled fixed cost of adjustment, and "high depreciation" refers to a firm with a doubled depreciation rate.

Figure 4
Inaction Spell Histograms

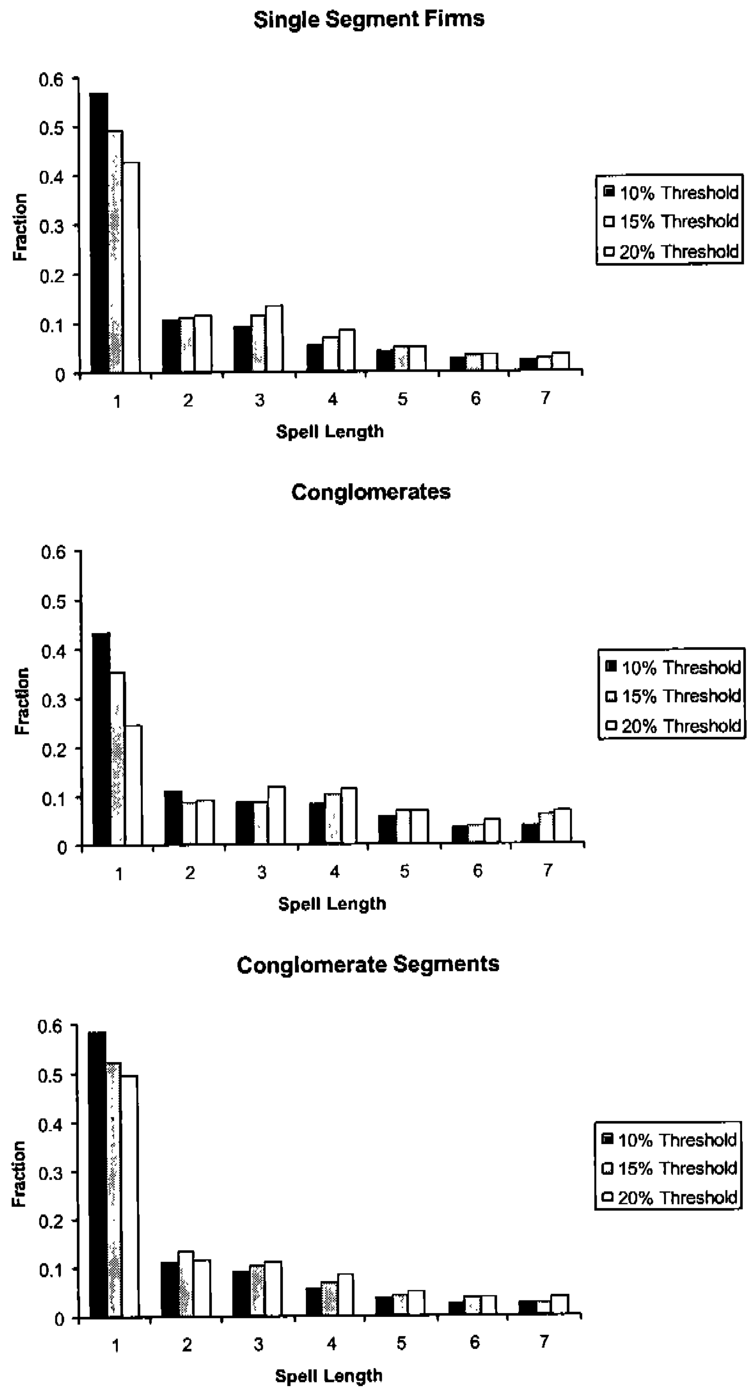

Calculations areb ased on a sampleo f non-financial firms and segments of these firms from the combined annual and full coverage 2000 Standard and Poor's COMPUSTAT industrial files that are also covered by COMPUSTAT'S Business Information File. The sample period is 1984 through 1997.

Figure 5
Estimated Hazards
Heterogeneity versus No Heterogeneity

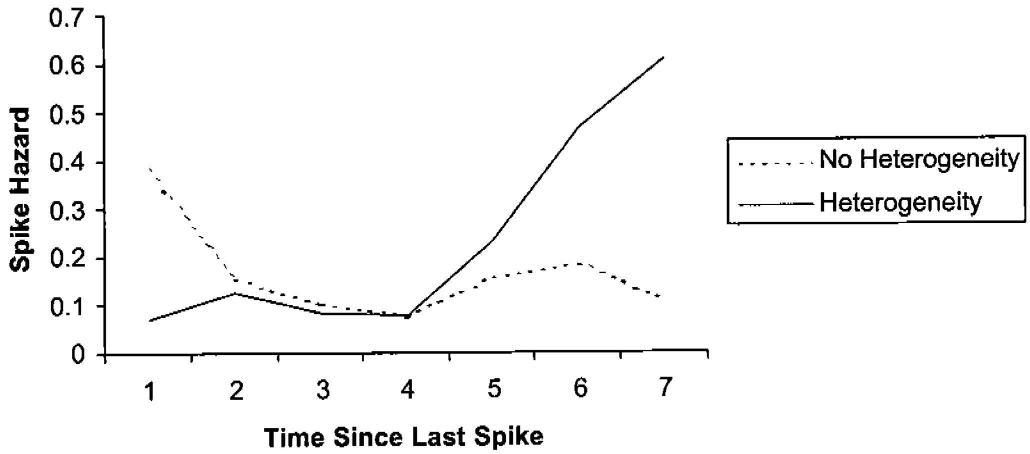

Estimates are from the twenty percent threshold columns of Table 4. "Heterogeneity" refers to estimates from a model that allows for unobservable cross-sectional heterogeneity. "No Heterogeneity" refers to estimates from a model that does not. The horizonatal axis measures the length of time since the last investment spike, and the vertical axis measures the probability of a spike, given that the firm has remained inactive up to that time. Calculations are based on a sampleo f single-segment non-financial firms from thec ombined annual and full coverage2 000 Standard and Poor's COMPUSTAT industrial files that are also covered by COMPUSTAT'S Business Information File. The sample period is 1984 through 1997.

Figure 6
Estimated Hazards
Dividends versus No Dividends

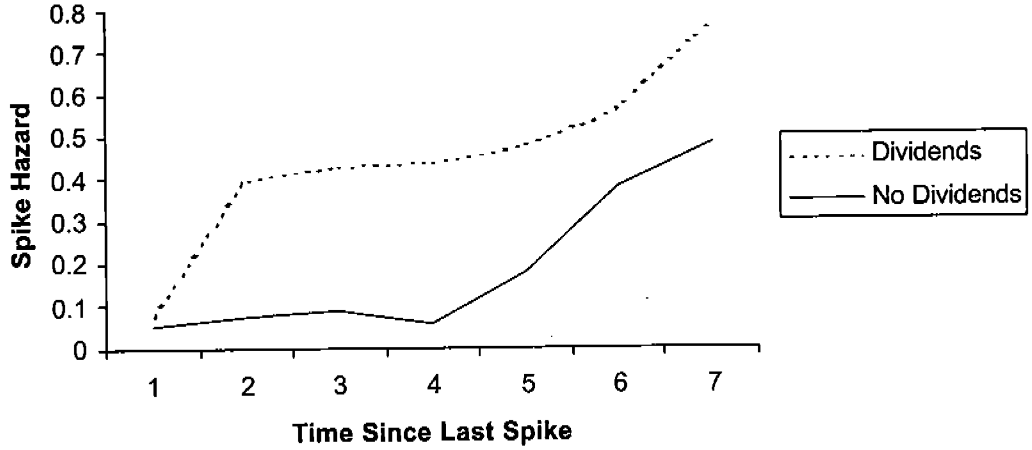

Estimates aref rom the twenty percent threshold columns of Table 5. "No Dividends" refers to estimates from a sample of small single-segment firms that consistently pay no dividends. "Dividends" refers to estimates from a sample of small single-segment firms that does pay dividends. Theh orizonatal axis measures the length of time since the last investment spike, and the vertical axis measures the probability of a spike, given that the firm has remained inactive up to that time. Calculations are based on a sample of single-segment non-financial firms from the combined annual and full coverage 2000 Standard and Poor's COMPUSTAT industrial files that are also covered by COMPUSTAT'S Business Information File. The sample period is 1984 through 1997.

Figure 7
Estimated Hazards
Small Firms versus Small Segments

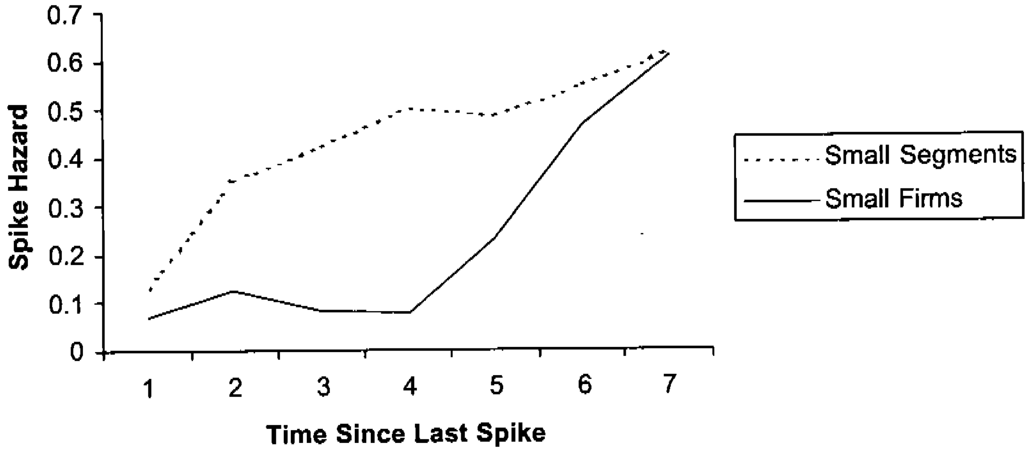

Estimates are from the twenty percent threshold columns of Table 6. "Small Firms" refers to estimates from a sample of small single-segment firms. "Small Segments" refers to estimates from a sample of small segments of conglomerates. The horizonatal axis measures the length of time since the last investment spike, and thev ertical axis measures the probability of a spike, given that thefi rm has remained inactive up to that time. Calculations areb ased on a sample of single-segment non-financial firms from the combined annual and full coverage 2000 Standard and Poor's COMPUSTAT industrial files that are also covered by COMPUSTAT'S Business Information File. The sample period is 1984 through 1997.

Table 1: Summary Statistics

Calculations are based on a sample of non-financial firms and segments of firms from the combined annual and full coverage 2000 Standard and Poor's COMPUSTAT industrial files that are also covered by COMPUSTAT's Business Information File. The sample period is 1984 through 1997. Assets are expressed in millions of 1992 dollars.

	Mean	Median	Standard Deviation
Single-Segment Firms			
Investment/Assets	0.082	0.055	0.089
Sales Growth	0.178	0.081	0.489
Depreciation/Assets	0.053	0.042	0.076
Assets	608	61	3460
Multiple-Segment Firms			
Investment/Assets	0.073	0.057	0.067
Sales Growth	0.092	0.042	0.387
Depreciation/Assets	0.049	0.043	0.046
Assets	2262	244	6803
Segments			
Investment/Assets	0.078	0.058	0.123
Sales Growth	0.081	0.035	0.368
Depreciation/Assets	0.059	0.048	0.087
Assets	981	141	3506

Table 2: Inaction and Investment Spikes

Calculations are based on a sample of non-financial firms and segments of firms from the combined annual and full coverage 2000 Standard and Poor's COMPUSTAT industrial files that are also covered by COMPUSTAT's Business Information File. The sample period is 1984 through 1997. I is gross investment, $\bar{D}$ is a firm's average reported depreciation, and A is total assets.

	Fraction of Observations				
	$\frac{I}{A}-\frac{\bar{D}}{A}>0.1$	$\frac{I}{A}-\frac{\bar{D}}{A}>0.15$	$\frac{I}{A}-\frac{\bar{D}}{A}>0.2$	$\frac{I}{A}-\frac{\bar{D}}{A}<0$	$\frac{I}{A}=0$
Single-Segment Firms	0.134	0.081	0.050	0.252	0.010
Multiple-Segment Firms	0.093	0.046	0.026	0.223	0.004
Segments	0.090	0.050	0.030	0.313	0.018

Table 3: Inaction Spells

Calculations are based on a sample of non-financial firms and segments of firms from the combined annual and full coverage 2000 Standard and Poor's COMPUSTAT industrial files that are also covered by COMPUSTAT's Business Information File. The sample period is 1984 through 1997.

Threshold	0.10	0.15	0.20
Single-Segment Firms			
Number	3279	1724	1037
Average Length	2.581	2.978	3.295
Fraction Censored	0.298	0.379	0.445
Average Length Censored	5.099	5.240	5.334
Fraction Uncensored	0.702	0.621	0.555
Average Length Uncensored	1.510	1.595	1.663
Conglomerates			
Number	1371	640	333
Average Length	3.521	4.220	4.757
Fraction Censored	0.390	0.519	0.622
Average Length Censored	6.275	6.497	6.502
Fraction Uncensored	0.610	0.481	0.378
Average Length Uncensored	1.758	1.766	1.889
Segments			
Number	2232	1106	620
Average Length	2.411	2.571	2.745
Fraction Censored	0.319	0.382	0.405
Average Length Censored	4.439	4.363	4.637
Fraction Uncensored	0.681	0.618	0.595
Average Length Uncensored	1.460	1.465	1.458

Table 4: Semiparametric Hazard Model Estimates: SmallS ingle-Segment Firms

Calculations are based on a sample of non-financial firms from the combined annuala nd full coverage 2000 Standard and Poor's COMPUSTAT industrial files that are also covered by COMPUSTAT's Business Information File. The sample period is 1984 through 1997. The rows labeled $\exp(\gamma_i)$ contain estimates of the baseline hazard, where the subscript refers to the number of years since the last spike. The row labeled σ^2 contain estimates of the variance of cross-sectional heterogeneity of the hazards. Standard errors are in parentheses under the parameter estimates.

	Heterogeneity			No Heterogeneity		
Threshold	0.10	0.15	0.20	0.10	0.15	0.20
Sales Growth	0.1169	0.2091	0.2237	-1.3156	-1.5084	-1.5599
	(0.0599)	(0.1221)	(0.1270)	(1.8957)	(1.8753)	(2.4180)
$\exp(\gamma_1)$	0.1080	0.0719	0.0725	0.3907	0.3632	0.3792
	(0.0088)	(0.0083)	(0.0114)	(0.0161)	(0.0220)	(0.0322)
$\exp(\gamma_2)$	0.1153	0.0860	0.1248	0.1620	0.1617	0.1559
	(0.0106)	(0.0115)	(0.0214)	(0.0116)	(0.0144)	(0.0123)
$\exp(\gamma_3)$	0.1462	0.1397	0.0798	0.1339	0.1453	0.1012
	(0.0161)	(0.0163)	(0.0120)	(0.0059)	(0.0053)	(0.0059)
$\exp(\gamma_4)$	0.1281	0.0875	0.0770	0.1256	0.0773	0.0741
	(0.0122)	(0.0080)	(0.0072)	(0.0051)	(0.0084)	(0.0084)
$\exp(\gamma_5)$	0.1896	0.1616	0.2294	0.1403	0.1522	0.1522
	(0.0153)	(0.0226)	(0.0219)	(0.0073)	(0.0075)	(0.0077)
$\exp(\gamma_6)$	0.3511	0.2153	0.4659	0.1865	0.1504	0.1851
	(0.0204)	(0.0128)	(0.0342)	(0.0069)	(0.0098)	(0.0075)
$\exp(\gamma_7)$	0.5230	0.5498	0.6051	0.1842	0.1308	0.1101
	(0.0245)	(0.0687)	(0.0763)	(0.0045)	(0.0104)	(0.0065)
σ^2	1.4389	1.5144	1.8615			
	(0.0560)	(0.0784)	(0.0814)			
Log-likelihood	-79.5269	-25.4187	-10.7285	-129.9532	-41.8881	-18.0428
Sample Size	1243	700	467	1243	700	467

Table 5: Semiparametric Hazard Model Estimates: Firms Grouped by Size and Diversification

Calculations are based on a sample of non-financial firms from the combined annual and full coverage 2000 Standard and Poor's COMPUSTAT industrial files that are also covered by COMPUSTAT's Business Information File. The sample period is 1984 through 1997. The rows labeled $\exp(\gamma_i)$ contain estimates of the baseline hazard, where the subscript refers to the number of years since the last spike. The row labeled σ^2 contain estimates of the variance of cross-sectional heterogeneity of the hazards. Standard errors are in parentheses under the parameter estimates.

	Large Single-Segment Firms			Small Conglomerates			Large Conglomerates		
Threshold	0.10	0.15	0.20	0.10	0.15	0.20	0.10	0.15	0.20
Sales Growth	0.4849	0.3184	0.5152	0.1629	0.1905	0.4240	0.3901	0.5016	0.7231
	(0.0179)	(0.0765)	(0.0924)	(0.1152)	(0.1653)	(0.4785)	(0.0250)	(0.0090)	(0.1687)
$\exp(\gamma_1)$	0.1210	0.1273	0.0814	0.1027	0.0887	0.0585	0.0888	0.0814	0.0439
	(0.0030)	(0.0086)	(0.0130)	(0.0113)	(0.0156)	(0.0240)	(0.0030)	(0.0007)	(0.0100)
$\exp(\gamma_2)$	0.0426	0.0478	0.0390	0.0369	0.0350	0.1121	0.0507	0.0355	0.0354
	(0.0010)	(0.0048)	(0.0072)	(0.0069)	(0.0096)	(0.0525)	(0.0019)	(0.0003)	(0.0092)
$\exp(\gamma_3)$	0.0418	0.0448	0.0378	0.0576	0.0281	0.1325	0.0375	0.0262	0.0542
	(0.0005)	(0.0057)	(0.0080)	(0.0110)	(0.0096)	(0.0788)	(0.0020)	(0.0002)	(0.0149)
$\exp(\gamma_4)$	0.0334	0.0469	0.0339	0.0462	0.0467	0.5859	0.0602	0.0355	0.0184
	(0.0017)	(0.0051)	(0.0072)	(0.0103)	(0.0164)	(0.3243)	(0.0019)	(0.0003)	(0.0079)
$\exp(\gamma_5)$	0.0377	0.0343	0.0274	0.0527	0.0197	0.0169	0.0382	0.0336	0.0192
	(0.0006)	(0.0021)	(0.0035)	(0.0069)	(0.0029)	(0.0069)	(0.0005)	(0.0003)	(0.0048)
$\exp(\gamma_6)$	0.0232	0.0272	0.0206	0.0572	0.0167	0.6099	0.0165	0.0299	0.0204
	(0.0002)	(0.0018)	(0.0037)	(0.0055)	(0.0020)	(0.4665)	(0.0034)	(0.0003)	(0.0039)
$\exp(\gamma_7)$	0.0341	0.0977	0.0391	0.0672	0.0479	0.3329	0.0343	0.0305	0.0223
	(0.0027)	(0.0173)	(0.0656)	(0.0330)	(0.0083)	(0.0030)	(0.0051)	(0.0057)	(0.0297)
σ^2	0.0016	0.0126	0.1044	0.0060	0.0953	2.7403	0.0068	0.0009	0.0779
	(0.0012)	(0.2299)	(0.1484)	(0.2908)	(0.2167)	(0.2539)	(0.0692)	(0.0090)	(0.4647)
Log-likelihood	-311.2549	-78.0375	-23.7617	-11.2987	-2.8107	-0.8391	-78.8649	-13.8878	-3.0490
Sample Size	2036	1024	570	370	195	111	1001	445	222

Table 6: Semiparametric Hazard Model Estimates: Small Single-Segment Firms Grouped by Lagged Dividend Payout

Calculations are based on a sample of non-financial firms from the combined annuala nd full coverage 2000 Standard and Poor's COMPUSTAT industrial files that are also covered by COMPUSTAT's Business Information File. The sample period is 1984 through 1997. The rows labeled $\exp(\gamma_i)$ contain estimates of the baseline hazard, where the subscript refers to the number of years since the last spike. The row labeled σ^2 contain estimates of the variance of cross-sectional heterogeneity of the hazards. Standard errors are in parentheses under the parameter estimates.

	No Dividends			Dividends		
Threshold	0.10	0.15	0.20	0.10	0.15	0.20
Sales Growth	0.0530	0.1292	0.4427	0.7845	0.7028	0.5819
	(0.1194)	(0.1397)	(0.1258)	(0.1111)	(0.1904)	(0.2608)
$\exp(\gamma_1)$	0.0823	0.0578	0.0539	0.0922	0.0593	0.0747
	(0.0071)	(0.0066)	(0.0082)	(0.0153)	(0.0141)	(0.0192)
$\exp(\gamma_2)$	0.0788	0.0579	0.0755	0.3525	0.4166	0.3932
	(0.0088)	(0.0080)	(0.0094)	(0.0264)	(0.0172)	(0.0175)
$\exp(\gamma_3)$	0.1218	0.1070	0.0865	0.4103	0.3858	0.4271
	(0.0115)	(0.0152)	(0.0149)	(0.0114)	(0.0056)	(0.0157)
$\exp(\gamma_4)$	0.1142	0.0673	0.0554	0.4607	0.4881	0.4369
	(0.0083)	(0.0049)	(0.0062)	(0.0088)	(0.0158)	(0.0090)
$\exp(\gamma_5)$	0.1431	0.1557	0.1747	0.4979	0.5356	0.4783
	(0.0082)	(0.0140)	(0.0127)	(0.0066)	(0.0167)	(0.0090)
$\exp(\gamma_6)$	0.2264	0.1565	0.3811	0.6407	0.5467	0.5573
	(0.0089)	(0.0089)	(0.0247)	(0.0157)	(0.0103)	(0.0159)
$\exp(\gamma_7)$	0.3012	0.3991	0.4808	0.6380	0.7694	0.7731
	(0.0027)	(0.0527)	(0.0721)	(0.0105)	(0.0417)	(0.0550)
σ^2	1.3706	1.3654	1.7185	2.1365	2.5983	2.9094
	(0.0602)	(0.0863)	(0.1020)	(0.1013)	(0.1759)	(0.2527)
Log-likelihood	-72.9731	-22.2039	-9.8178	-8.8570	-3.4872	-1.3726
Sample Size	943	515	349	300	185	118

Table 7: Semiparametric Hazard Model Estimates: Small Single-Segment Firms versus Small Segments

Calculations are based on a sample of non-financial firms from the combined annuala nd full coverage 2000 Standard and Poor's COMPUSTAT industrial files that are also covered by COMPUSTAT's Business Information File. The sample period is 1984 through 1997. The rows labeled $\exp(\gamma_i)$ contain estimates of the baseline hazard, where the subscript refers to the number of years since the last spike. The row labeled σ^2 contain estimates of the variance of cross-sectional heterogeneity of the hazards. Standard errors are in parentheses under the parameter estimates.

	Single-Segment Firms			Segments		
Threshold	0.10	0.15	0.20	0.10	0.15	0.20
Sales Growth	0.1169	0.2091	0.2237	0.6525	0.0533	0.2199
	(0.0599)	(0.1221)	(0.1270)	(0.0530)	(0.1369)	(0.1652)
$\exp(\gamma_1)$	0.1080	0.0719	0.0725	0.1131	0.0849	0.1290
	(0.0088)	(0.0083)	(0.0114)	(0.0101)	(0.0079)	(0.0083)
$\exp(\gamma_2)$	0.1153	0.0860	0.1248	0.3939	0.3123	0.3517
	(0.0106)	(0.0115)	(0.0214)	(0.0520)	(0.0674)	(0.0750)
$\exp(\gamma_3)$	0.1462	0.1397	0.0798	0.5038	0.4242	0.4203
	(0.0161)	(0.0163)	(0.0120)	(0.1719)	(0.1453)	(0.1444)
$\exp(\gamma_4)$	0.1281	0.0875	0.0770	0.5544	0.5770	0.5020
	(0.0122)	(0.0080)	(0.0072)	(0.1071)	(0.1119)	(0.1065)
$\exp(\gamma_5)$	0.1896	0.1616	0.2294	0.4724	0.4837	0.4820
	(0.0153)	(0.0226)	(0.0219)	(0.0097)	(0.0106)	(0.0117)
$\exp(\gamma_6)$	0.3511	0.2153	0.4659	0.5874	0.5869	0.5460
	(0.0204)	(0.0128)	(0.0342)	(0.0572)	(0.0619)	(0.0517)
$\exp(\gamma_7)$	0.5230	0.5498	0.6051	0.7756	0.6600	0.6126
	(0.0245)	(0.0687)	(0.0763)	(0.1112)	(0.1071)	(0.1458)
σ^2	1.4389	1.5144	1.8615	1.9361	2.2395	2.5750
	(0.0560)	(0.0784)	(0.0814)	(0.0541)	(0.0814)	(0.0895)
Log-likelihood	-79.5269	-25.4187	-10.7285	-94.7680	-31.1166	-10.6828
Sample Size	1243	700	467	1002	571	345

Table 8: Semiparametric Hazard Model Estimates: Liquid Asset Model

Calculations are based on a sample of non-financial firms and segments of conglomerates from the combined annual and full coverage 2000 Standard and Poor's COMPUSTAT industrial files that are also covered by COMPUSTAT's Business Information File. The sample period is 1984 through 1997. The rows labeled $\exp(\gamma_i)$ contain estimates of the baseline hazard, where the subscript refers to the number of years since the last spike. The row labeled σ^2 contain estimates of the variance of cross-sectional heterogeneity of the hazards. Standard errors are in parentheses under the parameter estimates.

Threshold	0.10	0.15	0.20
Sales Growth	0.0591	0.2008	1.1434
	(0.1104)	(0.1001)	(0.0671)
Liquid Assets	1.6160	1.6521	0.1776
	(0.4678)	(0.5372)	(0.1691)
$\exp(\gamma_1)$	0.1021	0.0632	0.0701
	(0.0087)	(0.0076)	(0.0092)
$\exp(\gamma_2)$	0.1268	0.0941	0.1513
	(0.0120)	(0.0128)	(0.0159)
$\exp(\gamma_3)$	0.1487	0.1622	0.2545
	(0.0163)	(0.0191)	(0.0096)
$\exp(\gamma_4)$	0.1422	0.0920	0.1768
	(0.0131)	(0.0064)	(0.0050)
$\exp(\gamma_5)$	0.2247	0.1998	0.3039
	(0.0183)	(0.0306)	(0.0040)
$\exp(\gamma_6)$	0.3968	0.2350	0.4814
	(0.0228)	(0.0110)	(0.0180)
$\exp(\gamma_7)$	0.5553	0.4425	0.5606
	(0.0296)	(0.0625)	(0.0237)
σ^2	1.4927	1.6156	2.2541
	(0.0583)	(0.0771)	(0.1072)
Log-likelihood	-113.2989	-35.8652	-17.3519
Sample Size	1164	651	438