What Gives? A Study of Firms' Reactions to Cash Shortfalls

Tor-Erik Bakke
University of Wisconsin, Madison

Toni M. Whited*
University of Wisconsin, Madison

November 6, 2008

*Corresponding author. Wisconsin School of Business, University of Wisconsin, Madison, WI 53706-1323. 608-262-6508. toni.whited@simon.rochester.edu. Tim Werner and Heqing Zhu provided excellent research assistance. We are grateful for helpful comments from Harry DeAngelo, Nathalie Moyen, Michael Roberts, and seminar participants at Claremont-McKenna College, Columbia University, Stockholm School of Economics, Norwegian School of Management, Norwegian School of Economics and Business Administration, and Wirtschaftsuniversität Wien.

What Gives? A Study of Firms' Reactions to Cash Shortfalls

Abstract

This paper examines whether firms react to cash shortfalls by cutting investment. We use a regression discontinuity design in which the discontinuity is the point of violation of underfunding of corporate defined benefit pension plans. We reexamine the puzzling evidence in Rauh (2006) that mandatory pension contributions cause sharp investment declines, finding that these results are likely due to the endogeneity that this study is trying to avoid. We also compare firm-year observations in which the firm's pension assets are just barely less than its pension liabilities to observations in which assets are just greater than liabilities. In this quasi-experimental setting, we find little evidence that firms cut back on investment. Instead, they mostly use a variety of financial tools, such as receivables factoring and payout cuts, to fund their pension liabilities.

1. Introduction

A long-standing open empirical question in finance and economics is whether and how financial market frictions affect real economic activity, in particular, corporate investment. Dating back to the influential work of Fazzari, Hubbard, and Petersen (1988), researchers have used the sensitivity of investment to cash flow as a metric for gauging the severity of finance constraints. The intuition behind this test is straightforward. If a firm cannot obtain outside finance, then its investment should respond strongly to movements in its cash flow. However, because cash flow is correlated with investment opportunities, an observed response of investment to cash flow can simply be due to movements in investment opportunities. Therefore, the challenge in identifying the effects of financial frictions is disentangling these effects from those of investment opportunities.

Recently, Rauh (2006) has dealt with this problem by finding a shock to firm cash flow that is arguably exogenous to investment opportunities: the contributions firms must make to their defined-benefit pension plans if the assets backing these plans fall below the estimated liabilities. Although these mandatory contributions themselves are endogenously determined with other real and financial firm decisions, the contributions jump up discretely at the point where the firm's pension assets fall below its pension liabilities. One of the important contributions in Rauh is to exploit this discontinuity to deal with the endogeneity problem. In so doing, Rauh finds that firms cut their capital expenditures 70 cents for every dollar of mandatory pension contributions. This finding is perplexing because it seems odd that firms would prefer to cut productive assets in response to a possibly temporary cash shortfall, instead of using other, less draconian measures.

We argue that this puzzle can be explained. We build on Rauh's ingenious insight of using pension contributions in a regression discontinuity design, but we correct several important flaws in his implementation of this technique and, accordingly, find strikingly different results. We find little evidence that firms respond to mandatory contributions by cutting investment. In particular, we find no evidence of a drop in investment at the point where firms' pensions become under-funded. This finding is important because it is precisely this point of discontinuity that provides for identification of Rauh's effect and justifies the use of the regression discontinuity technique.

Despite the lack of discontinuity in investment, Rauh still finds a significant investment response in his regression analysis. However, we show that Rauh's estimated investment effect is an artifact of a misspecified regression, that is, the effect comes from firms that are deeply underfunded. These troubled firms are systematically different from the rest of the sample. These observed differences suggest that their large pension contributions are endogenously determined with investment opportunities. For example, relative to overfunded firms, severely underfunded firms have less than one sixth the earnings of their overfunded counteparts. Consequently, it is not surprising that these firms invest less. Next, we also show that Rauh's conclusions follow from a misinterpretation the results. Consistent with our finding of no discontinuity in investment, a proper interpretation of Rauh's results reveals no economically or statistically significant effect of pension funding shortfalls on capital expenditures.

Finally, we use the subsample of firms near funding discontinuity to examine whether shortfalls impact corporate behavior along other margins. Firms have a strong tendency to fund shortfalls by factoring receivables and cutting shareholder payout. We find limited evidence of finance constraints along two dimensions. First, firms without bond ratings cut employment growth when they have to fund their pensions. Second, firms likely to be financially constrained have a stronger tendency to avoid the point of a pension funding violation than do firms likely to be unconstrained. These results shed light on a question much broader than that of investment-cash flow sensitivities. They provide evidence on the relative magnitude of real versus financial frictions inasmuch as one would expect firms to adjust on the least costly margin when hit with a shock. One can conclude that for the large fraction of firms that locate close to the point of a pension funding violation, financial frictions are of second order.

To better understand these results, it is important to elaborate on both the econometrics and on the workings of defined benefit (DB) pension plans. We start with the latter. Companies with DB pension plans promise employees a prespecified monthly benefit at retirement. Typically, this benefit is calculated as a function of employee salary and service, and government rules state that firms must contribute to their pension plans to be able to meet these future pension liabilities. If the market value of these contributions (pension assets) exceeds the expected future pension liabilities, the pension plan is considered to be overfunded. Conversely, if the value of the pension assets is less than the pension liabilities, the plan is considered to be underfunded. If the pension plan is overfunded, the firm is free either to contribute or not to contribute to the plan, although contributions above a specified ceiling receive unfavorable tax treatment. If the firm withdraws funds from its pension plans, it can be subject to severe excise taxes. On the other hand, if the plan is underfunded, the firm is required to contribute more funds to its pension plan. These mandatory pension contributions (MPCs) are determined according to a government rule that is a discontinuous function of the funding gap (the difference between pension assets and liabilities).

Fluctuations in the funding gap are for the most part driven by plan contributions and by the performance of the invested pension assets, but not by the value of the pension liabilities. First, if firms experience funding gaps, they do not have much leeway to adjust pension liabilities because of restrictive Internal Revenue Service rules that govern the actuarial assumptions allowed for the valuation of pension liabilities. Second, companies rarely change the amount of benefits promised to their employees. Third, although the present value of pension liabilities depends on the current market interest rate, these changes are dwarfed by the changes in the market value of pension assets that occur with fluctuations in the value of the investment portfolio.

Because firms choose how much to fund their pension plans, the decision whether to risk paying a future MPC is made jointly with investment and financing choices. Consequently, although fluctuations in the value of pension assets are beyond the firm's control, changes in MPCs are endogenous. Put differently, firms optimize subject to the existence of these discontinuities and endogenously choose whether they want to be close to point of a funding violation. That firms actively manage and anticipate future MPCs is vividly illustrated in this quote from Ford Motor Company's 2005 10K form:

> "In 2005, we made \$2.5 billion of cash contributions to our funded pension plans. During 2006, we expect to contribute \$1.5 billion to our worldwide pension plans. . .Based on current assumptions and regulations, we do not expect to have a legal requirement to fund our major U.S. pension plans in 2006."

Clearly, Ford is planning its contributions in 2005 and 2006 to avoid having to face unanticipated financial shocks in 2006.

However, this obvious endogeneity problem can be avoided because the function that links underfunding to MPCs has a discontinuity. Rauh argues that "the sharp nonlinearities of pension funding requirements, particularly around the threshold of underfunding, allow for the identification of an effect of required contributions on investment that is purged of this endogeneity problem." This statement is only partially correct. As explained in detail below, nonlinearities in the funding rules do not provide identification; only the discontinuity at the point of a funding violation does. To understand this issue, it is useful to consider an ideal experimental setting in which one would flip a coin to assign a pension funding violation to a group of firms at random and then compare treated and control groups. This sort of experiment is infeasible, but one can obtain a quasi-experimental setting because the firms that have barely violated the pension funding rules are not much different from those that have barely escaped a violation. Therefore, the near-escapees and near-violators can be thought of as close-to-randomly assigned to a violation, and comparing their investment behavior therefore reveals a causal effect of the funding violation.

This type of regression discontinuity design is valid only for firms that are very close to the point of a funding violation. However, it is also interesting to examine the entire sample, especially if those observations close to the discontinuity are different from those far from the discontinuity. In order to use the whole sample, one can make assumptions, discussed in Van der Klaauw (2002) and delineated below, under which an analysis of the whole sample produces the same results as an analysis only of those observations close to the discontinuity. The two most important are as follows. First, one must control for the distance to the point of violation. Second, one must control for any systematic differences between the firms that are near and far from the cutoff point. Rauh's whole-sample experimental design meets the first condition, but, as explained below, it violates the second. Our proper application of regression discontinuity demonstrates that this problem has a large impact on what we learn from this sort of exercise.

Our results also put the main contribution in Rauh in perspective. Sensitivity of investment to an exogenous cash flow shock at most indicates that external finance is more costly than cutting capital expenditures. By extension, external finance must then be more costly than internal, and this fact is trivially true because of the presence of issuance costs and taxes. Therefore, Rauh's main contribution is not necessarily a demonstration that finance constraints impact investment—a result previously established by many earlier works, as surveyed in Stein (2003). It is instead an introduction of regression discontinuity to the financial economics literature. We build upon this important contribution by following closely the recommendations proposed by numerous econometric studies of regression discontinuity, such as Hahn, Todd, and Van der Klaauw (2001), Van der Klaauw (2002), Imbens and Lemieux (2008), and McCrary (2008). As such, our study not only sheds further light on the ways in which firms respond to cash shortfalls, but also provides practical guidelines for using the regression discontinuity design in a financial economics application.

The current paper naturally falls into the enormous literature on investment and finance constraints, the modern treatment of which starts with Fazzari, Hubbard, and Petersen (1988). This literature is nicely surveyed in Stein (2003), and more recent developments are surveyed in Hennessy and Whited (2007). More specifically, our paper fits into the prior literature that has tried to understand the relation between finance and investment by studying how firms respond to arguably exogenous shocks to cash flow. Obviously, Rauh fits into this category. In addition, Blanchard, Lopez-de-Silanes, and Shleifer (1994) study legal settlements, and Lamont (1997) studies the reaction of the non-oil subsidiaries of oil firms to the dramatic drop in oil prices in the mid-1980s. One difference between this paper and these other three is their exclusive focus on the sensitivity of investment to cash flow. In contrast, we seek to look at a much broader question: the relative magnitude of financial and real frictions. Finally, our paper is most closely related to Chava and Roberts (2008), who also use both a full-sample and a local discontinuity analysis to examine the effects of bond covenant violations on corporate investment.

This paper proceeds as follows. Section 2 describes the data; Section 3 introduces the econometric methodology; Section 4 presents the results, and Section 5 concludes.

2. Data

The empirical analysis in this paper employs an unbalanced panel of Compustat firms from the 2007 Standard and Poor's Compustat industrial files. Definitions of the Compustat variables we use are in the appendix. We restrict our attention to the subsample of Compustat firms that file an IRS 5500 form with the Department of Labor (DOL) and that sponsor defined benefit pension plans. We need to impose the first restriction because our methodology relies on accurately calculating MPCs. All pension plans over a certain size must file an Internal Revenue Service (IRS) form 5500 yearly. We impose the second restriction because only firms with DB pension plans must pay MPCs.

Compustat does provide pension data based on firm SEC filings; however, we do not use this data. As explained in much more detail in Rauh, Compustat pension data is aggregated to the firm level. It is therefore inadequate for our purposes because mandatory pension contributions (MPCs) are determined at the pension plan level—not the firm level. Approximately one third of the firms in our sample have more than one pension plan, and relying on Compustat data, therefore, would lead to inaccurate computation of firm MPCs. Furthermore, firms have significantly more accounting discretion when submitting SEC filings than when filing form IRS 5500. SEC pension data also include both international and domestic pension plans, whereas only domestic plans are required to pay MPCs. Finally, the methods for computing pension liabilities and costs for SEC filings are different from those that are required for computing MPCs.

The sample period runs from 1990 to 1998. The sample starts in 1990 because IRS 5500 forms are first available in this year in standardized form from the DOL. We match pension plans to Compustat firms primarily using CUSIPs. This method confers the advantage that all plans from 1990 to 1998 were required to report the CUSIP of the plan's sponsor. Although we match most of the data by CUSIPs, we occasionally need to match either by employee identification numbers (EIN) or by exact firm names. If all of these methods fail, we match by hand. EINs and firm names only match plans that directly pertain to the Compustat firm. For instance, if the firm's subsidiary sponsors a plan, this plan cannot be matched without access to the CUSIP of the parent firm. Although we attempt to mitigate this problem by using subsidiary names from the Compustat Business Information File, without the CUSIP of the parent of the subsidiary sponsoring the plan, we face potential sample selection issues. For this reason the sample ends in 1998 when a change in reporting requirements no longer required pension plans to report the CUSIP pertaining to the plan . After this matching process we end up with 7905 firm-year observations, a number somewhat smaller than that in Rauh.

For each plan year we extract the following data variables from the plan's IRS 5500 filing: normal cost, pension assets, pension liabilities, and total contributions. Normal cost is the increase in pension liabilities in the current year. Pension assets are the present value of assets in the plan at the beginning of the year. Pension liabilities are the accumulated benefit obligation from Schedule B of the IRS 5500. From 1991 to 1994 the calculation of pension liabilities follows the regulations in the Omnibus Budget Reconciliation Act of 1987. After 1994 these calculations follow the regulations in the Retirement Protection Act of 1994.

How are MPCs determined? Over the sample period from 1990-1998, MPCs are composed of two components: the minimum funding contribution (MFC) and the deficit reduction contribution (DRC). Firms must contribute the larger of the MFC or the DRC. The MFC is defined as the sum of the normal cost (the present value of the pension benefits that accrue during the current year) and an installment payment on unfunded liabilities. It is the normal cost component that provides the necessary discontinuity at the point of underfunding. Given that some discretion is allowed in calculating this installment, we follow Munnell and Soto (2003) and Rauh and estimate it as 10% of the funding gap. Therefore, the MFC is given by

$$MFC = (\text{Normal Cost}) + 0.1\,(\text{funding gap})\,.$$

The DRC is given as a fraction of funding status:[1]

$$DRC = \min\left[0.30, (0.30 - 0.25\,(\text{funding gap} - 0.35))\right].$$

The MPC is then given by the following function:

$$MPC = \max\left[MFC, DRC\right],$$

if the funding gap is negative.

Although the minimum and maximum in the above definitions create kinks in the MPC function, these kinks are not useful for identification because nothing strictly discontinuous happens at these

[1] The formula below reflects the DRC requirement before 1994. Minor changes were made to this function in 1994, and our data analysis below takes these changes into account.

points. As demonstrated in the following section, a point of nondifferentiability does not provide identification. Only a strict discontinuity does. We therefore restrict our analysis to the exact point of underfunding, which is where the MFCs kick in, and which does consists of a point of discontinuity both before and after 1994, which only changed the calculation of the DRC.

This discussion emphasizes the importance of measuring the funding gap. Measuring a funding gap for a firm with one plan is straightforward. Measuring a funding gap for a firm with multiple plans is more complicated. Rauh sums the gaps across all pension plans within a firm. This definition, which we refer to as the average funding status, does not correctly identify the point of discontinuity if some plans are in the red and others are in the black. As an alternative, we define the funding gap as the smallest gap if all are plans have positive funding status, and the sum of the negative gaps if any plans have negative funding status. This definition captures the notion of the distance to the point of violation. As an alternative we have also used a measure of the funding gap defined as the pension-liability-weighted average of the gaps for the plans if all have positive funding status, and the pensions-liability-weighted gaps for the negative plans if any of them have negative funding status. Because these two definitions provide similar results in what follows, we focus on the first for brevity, and we refer to it below as the " pension gap" or "gap."

3. Empirical Strategy

We now outline the basic features of regression discontinuity design. We start with the intuitive description in the introduction that firms that are barely underfunded and those that are barely overfunded can be thought of as close-to-randomly assigned to a pension funding violation. Therefore, by calculating the average differences between the investment (or any other variable of interest) of these two groups of firms, one can estimate the causal effect of the funding violation on investment. This effect is called a local average treatment effect, or LATE.

Regression discontinuity is originally from Thistlethwaite and Campbell (1960), and our presentation follows closely that in Van der Klaauw (2001). Let y_i be a variable of interest, such as investment, employment, liquid assets, or external financing. Let w_i be a violation indicator, and let s_i be the funding surplus as a fraction of pension liabilities. We are interested in estimating the regression

$$y_i = \beta + \alpha w_i + u_i \tag{1}$$
$$w_i = w(s_i) = 1\{s_i \leq 0\},$$

in which α is the average treatment effect from "treatment" with a funding violation. If we try to estimate (1) on a sample of firms with wide variation in funding surpluses and deficits, assignment is not random; so $E(u_i \mid w_i) \neq 0$, and OLS produces biased coefficients.

As we have argued informally, however, we can use a restricted sample to estimate a LATE, which we define formally as

$$LATE = \lim_{s \downarrow 0} E(y \mid s) - \lim_{s \uparrow 0} E(y \mid s). \tag{2}$$

Why does this expression identify the treatment effect, α? To see why, note from (1) that

$$\begin{aligned} \lim_{s \downarrow 0} E(y \mid s) - \lim_{s \uparrow 0} E(y \mid s) &= \alpha \left(\lim_{s \downarrow 0} E(w \mid s) - \lim_{s \uparrow 0} E(w \mid s) \right) + \lim_{s \downarrow 0} E(u \mid s) - \lim_{s \uparrow 0} E(u \mid s) \\ &= \alpha (1 - 0) + \lim_{s \downarrow 0} E(u \mid s) - \lim_{s \uparrow 0} E(u \mid s) \end{aligned} \tag{3}$$

If we assume that $E(u \mid s)$ is continuous in s, then the last term goes to zero and we have

$$\alpha = \lim_{s \downarrow 0} E(y \mid s) - \lim_{s \uparrow 0} E(y \mid s).$$

The assumption that $E(u \mid s)$ is continuous in s is crucial, and it is therefore important to understand what it means in economic terms. If one takes the regression (1) seriously, it implies that the only variable that should determine firm investment, or employment, or external financing, or any other variable we consider is whether a firm's pension assets are greater than its pension liabilities. This interpretation is, of course, absurd, but it points out that many determinants of our variables of interest are omitted from (1) and are therefore implicitly contained in the error term, u_i. The continuity assumption then implies that none of these variables exhibits a discontinuity at the exact point of a pension funding violation. This assumption is from an intuitive standpoint likely to hold because it is hard to imagine that a variable such as Tobin's q would jump down at the point of a pension funding violation.

One difficulty with estimating a LATE is that one cannot necessarily extrapolate one's inferences to the rest of the sample, unless the sample that is close the point of a funding violation is representative of the entire sample. Under certain assumptions, nonetheless, it can be possible to examine the entire sample by using the concept of a control function from Heckman and Robb (1985). Suppose that the only determinant of a pension funding violation is the difference between pension assets and liabilities. Then one can write the regression error, u_i, as

$$u_i = E(u_i \mid s_i) + e_i, \tag{4}$$

in which e_i is, by definition, orthogonal to $w_i = w(s_i)$. Substituting (4) into (1) then gives

$$\begin{aligned} y_i &= \beta + \alpha w(s_i) + E(u_i \mid s_i) + e_i \\ &= \beta + \alpha w(s_i) + k(s_i) + e_i \end{aligned} \tag{5}$$

in which $k(s_i) \equiv E(u_i \mid s_i)$. In general, $k(s_i)$ will be a smooth function of s_i, although it will only be linear if $u_i \mid s_i$ is normally distributed, which is an implausible assumption in this instance. For example, investment is highly skewed. Nonetheless, if we are willing to swallow the assumption that s_i is the only determinant of w_i, we can estimate this regression by including smooth functions of the distance between pension assets and pension liabilities in the regression.

It is hard to swallow the assumption that pension funding status is the only determinant of any firm decisions, but thinking about the assumption points out the key difficulty with estimating (5) on a sample with wide variation in pension funding status. The regression must be very well specified in order for this technique to work. In other words, all other relevant control variables must be included in order to prevent a different source of error-regressor correlation. If not, then if $w(s_i)$ is correlated with anything that is left out of the regression, its coefficient will be biased. Van der Klauuw (2002) puts the point slightly differently by noting that estimating (5) requires strong assumptions to achieve identification. In particular, one has to assume that the effects of s_i (the pension funding gap) on y_i are adequately controlled for by other variables in the regression.

This condition may be violated for a variety of reasons, all of which stem from a disconnect between economic theory and the regressions used to test the theory. For example, if y_i is the rate of investment, then under certain assumptions delineated in, for example, Erickson and Whited

(2000), the q-theory of investment can produce a linear regression of investment on Tobin's q. However, as pointed out in Erickson and Whited (2000) and Whited (2001), the usual measures of Tobin's q only capture from twenty to forty percent of the variation in true investment opportunities. Further, even if one corrects for measurement error, reduced form investment regressions explain at most half of the variation in investment. Consider another example of a decision we examine: hiring. In this case firm-level data on average wages are unavailable in Compustat. Therefore, we lack sufficient data to run a well-specified employment demand equation. In the cases of all of the other decisions we examine, economic theory provides little guidance as to the exact form of a regression-based empirical model. Economic theory can tell us about some of the forces that affect these decisions, but it does not give researchers a specific regression. Therefore, attempting to control for all relevant variables becomes even more challenging. Given these difficulties, we examine only investment using the entire sample, and conduct only a more limited local analysis of all of the other firm decisions we consider.

4. Results

4.1. Summary Statistics

To begin our data analysis we present two simple kernel regressions of investment on the pension gap, one using only observations with a positive gap, and the other using only observations with a negative gap. We scale investment by total assets and the pension gap by pension liabilities (for ease of readability of the graph). We have removed firm fixed effects from investment and then added back in the full sample mean. Figure 1 plots these kernel regressions. The solid line is the fitted regression line, and the dotted lines are 5% confidence bands. The plot also shows the cutoff points for the fifth and tenth percentiles of the distribution of the pension gap. This type of informal analysis is recommended in Imbens and Lemieux (2008), who comment that formal statistical analysis is essentially just a sophisticated version of this sort of basic plot. The intent is to find a discontinuity in the outcome of interest, which in our case is investment. If no discontinuity can be observed, there is little chance that formal analysis will lead to estimates with economically and statistically significant magnitudes.

Four features of Figure 1 stand out. First, the two lines almost meet at a zero pension gap, that is, the point of discontinuity. Further, their difference lies well between the standard error bands. There appears to be no evidence of a jump in the conditional mean of investment at the cutoff point. Second, the regression using only positive gap observations is mostly flat; third, the regression using only negative gap observations is flat around the cutoff but then dips down away from the cutoff; and fourth, this drop off occurs entirely in the bottom 10% of the distribution of the pension gap, that is, for those firms that are deeply underfunded. These last three patterns imply that any differences in investment between the firms that are in violation and those that are not come primarily from the firms that are far away from the cutoff point. This finding questions strongly the use of RD in studying the effect of funding violations on investment. In other words, much of Rauh's identifying data variation is coming from the lower left tail of the distribution of the pension gap, which implies that the endogeneity problems he seeks to avoid are likely driving his results.[2]

Imbens and Lemieux (20080 also recommend plotting relevant covariates in a similar manner. Using the variables in the central regression in Rauh, we also examine analogous kernel regressions of Tobin's q, the ratio of nonpension cash flow to assets, and the ratio of mandatory contributions to assets. These plots are in Figure 2. As in the case of investment, we observe little movement in either cash flow or Tobin's q around the point of discontinuity, and we only observe drop-offs in these variables for the left tail of the distribution of the pension gap. In contrast, we see an almost constant increase in mandatory contributions as the pension gap becomes increasingly negative. Taken together, Figures 1 and 2 show almost no correlation between investment and mandatory contributions, except for deeply underfunded firms. Further, the drop in investment and the rise in mandatory contributions for these firms is accompanied by drops in both Tobin's q and cash flow, thereby indicating that that mandatory contributions are unlikely to be a truly exogenous cash flow shock. Put differently, the unobserved shock that drives firms far below the cutoff and triggers MPCs also affects their investment opportunities, cash flows, and capital expenditures

[2] It is possible that we have omitted other driving variables from our analysis and that a discontinuity might be evident if we control for them. We therefore do the same kernel regressions using the residual from a regresion of investment on Tobin's q and investment. We find similar results.

simultaneously. The analyses in Figures 1 and 2 are, of course, only informal, so we now reinforce our basic finding with a deeper look into the data.

Table 1 presents several summary statistics for our sample of firms, and each column contains results for a particular subsample. All of the variables are scaled by firm assets, except employment, which is measured in logs. For the full sample, the average size of the firm is quite large. For example, the mean value of real assets for all of Compustat between the years 1990 and 1998 is 1,670 million dollars, whereas the mean for this sample is 3,447 million dollars. This result is not surprising in that older, more mature firms tend to be those with DB pension plans. The large average size of the firms in this sample is also evident in the high incidence of firms that have bond ratings relative to the rest of the Compustat universe. Approximately 40% of our firms have bond ratings, whereas only about 25% of publicly listed U.S. firms have S&P bond ratings. The other striking feature of these data is that we observe mandatory contributions in just over 35% of the firm-year observations. This figure is interesting because it indicates that funding violations are economically relevant phenomena to study. In the next two columns we see that this 35% figure changes little for subsamples of firms with and without bond ratings. Another important pattern to notice is the small size of mandatory contributions relative to investment: the latter is on average over 40 times larger than the former. This result will be useful later in interpreting our regressions.

The next two columns examine groups of firms with and without bond ratings. This sample split is of interest inasmuch as firms without bond ratings are more likely to have difficulties raising external finance than firms with bond ratings.[3] These two groups differ in ways that one would predict. For example, the firms without bond ratings tend to be smaller than the firms with bond ratings, they pay fewer dividends, they conduct fewer share repurchases, and they issue more equity. One striking difference between the two groups is that the firms with bond ratings invest more than the firms without bond ratings. This feature of the data is not shared by the rest of Compustat, and this comparison reinforces the idea that this sample is neither random nor typical.

The next two columns of the table investigate the differences between the firms who have positive and negative average funding status. On a variety of levels these two groups of firms are surprisingly

[3]We have also conducted all of our analysis on groups of small and large firms, with largely similar results, which we omit for brevity.

similar. They have almost the same average size, share repurchases, equity issuance, incidence of bond ratings, cash holdings, leverage, long-term debt issuance, advertising expenditures, and R&D expenditures. We see important differences as well: the firms that are not in violation have slightly higher levels of Tobin's q, invest slightly more, and have slightly higher cash flow. One figure that immediately stands out in the third line of the table is that almost 17% of the firms classified as not violating their average funding status across all of their pension plans have at least one plan that is in the red. This finding is important because it means that average funding status is not a good indicator of the point of discontinuity.

The next two columns examine the five percent of firm/year observations that are closest to the point of a zero pension gap and the complement of this set. Although RD allows estimation of the effect of a funding violation close to the point of violation with minimal assumptions, extrapolating those results to the rest of the sample is problematic if the observations close to the cutoff differ systematically from those far from the cutoff. In our sample the firm/year observations close to the cutoff appear broadly similar to those far from the cutoff, except along four important dimensions. The firms close to the cutoff are much larger, they have almost twice the incidence of bond ratings, they issue much more short-term debt, and they do much more advertising. This evidence is consistent with the results in Whited (2009) that firms that face financial constraints tend to avoid known points at which they would experience cash shortfalls. It appears that large firms with easier access to financial markets are those that choose to locate near the cutoff because they are the ones that can easily tap external resources in order to avoid paying a mandatory contribution. The firms that are close to the cutoff also appear to be raising short-term debt, reducing cash balances (dissaving) and doing more advertising, all of which are actions consistent with boosting cash inflows to avoid a mandatory contribution. Although these differences appear to be consistent with the existence of finance constraints, they also indicate that any analysis done close to the cutoff may not have external validity, that is, be relevant for the rest of the sample. We keep this issue in mind as we proceed.

The final column examines firms that are in the bottom fifth percentile of the distribution of the pension gap. This group of firms is interesting because it is primarily from these firms that one observes low investment. These firms are systematically different from the rest of the sample. Not only do they invest less, but they have lower values of Tobin's q, and they have cash flow that is nearly half the sample average. Further, they pay half the dividends that the average firm pays, they are almost half as likely to have a bond rating, they issue less debt, and they increase their cash balances at a much lower rate. Most important, they have less than one sixth the earnings of their overfunded counterparts; their Altman's Z-scores are well below any commonly used cutoffs for an increased probability of bankruptcy, and they are shrinking their workforces at a rate of over 2% per year. All of these characteristics are what one would expect from distressed firms. These differences in observables also hint at potentially large differences in unobservables, and they lend credence to the hypothesis that these differences have as much a role to play in the determination of investment as do mandatory contributions.

We next examine the distributions of average funding status and the pension gap. As explained in McCrary (2008), it is important to inspect whether there is a discontinuity in the distribution of the forcing variable, which is in our case the pension gap, and which is average funding status in Rauh's case. If there is such a discontinuity, it raises the question whether the value of the pension gap is typically manipulated by the individual firms, thus invalidating the design. Figure 3 plots a histogram of the ratio of average funding status to total firm assets, with a bin width of 0.001. Recall that this variable is of interest because it defines whether a firm is in violation in Rauh. We observe a great deal of clustering around the break point, but no bunching of observations on one side or the other. Two explanations are possible. First, firms lose tax benefits from contributing to their plans if the plans are too overfunded. This provision gives them an incentive not to build too much of a cushion against a possible funding shortfall. Second, the Pension Benefit Guarantee Corporation insures the pension plans, which gives firms more of an incentive to underfund.

We next turn to the pension gap, whose histogram is plotted in Figure 4, once again with a bin width of 0.001. As in the case of average funding status we see many firm-year observations clustered around the point at which plan assets equal plan liabilities. To be certain that we do not have any observations bunched on one side of this cutoff, we plot a second histogram in Figure 5, with a narrower bin width of 0.00001. Here we find little visual evidence of bunching. A closer look at the histogram reveals 155 observations in the five bins to the left of the cutoff and 142 observations to the right. This evidence is reassuring in that although many firms tend to locate close to the cutoff, they do not systematically choose to locate on one side or the other. One wonders why firms do not crowd just above the cutoff point if they can indeed manipulate the pension gap. Two explanations are possible. First, they might not care about making a mandatory contribution if they make regular contributions in the first place. This explanation is plausible in light of the evidence in Table 1 that firms close to the cutoff (on both sides) make total contributions that are 15 times larger than their mandatory contributions. Second, even if the firms do care about being above the gap, random fluctuations in the value of pension assets can undo their attempts to avoid a funding violation. Both cases support the idea that the location of firms close to one side or the other is approximately random.

4.2. Regression Discontinuity Analysis

We start our investigation into the ways in which firms respond to cash shortfalls by reexamining the results in Rauh. In particular, we estimate the following regression which is from Rauh:

$$\frac{I_{it}}{A_{i,t-1}} = \alpha_i + \alpha_t + \beta_1 Q_{i,t-1} + \beta_2 \frac{CF_{it}}{A_{i,t-1}} + \beta_3 \frac{FS_{it}}{A_{i,t-1}} + \beta_4 \frac{MC_{it}}{A_{i,t-1}} + u_{it}, \tag{6}$$

in which I_{it} is capital expenditures for firm i at time t, $A_{i,t-1}$ is beginning-of-period book assets, $Q_{i,t-1}$ is the beginning-of-period market to book ratio, CF_{it} is nonpension cash flow, MC_{it} is the mandatory contribution, and FS_{it} is the average funding status of all plans in a firm. All of these regressions contain fixed firm and year effects, and the standard errors are robust to heteroskedasticity and clustering at the firm level.

Before turning to the results, it is important to note that (6) is not an exact application of regression discontinuity because it does not contain an indicator variable for whether the firm has a funding violation or not. Instead, it contains mandatory contributions. Like an indicator variable, MPCs are discontinuous at the point of a funding violation. However, an examination of (3) shows that the effect of a funding violation on investment is not given by the coefficient on mandatory contributions. To see this point heuristically, we simplify (6) by omitting all regressors except mandatory contributions and dropping the $A_{i,t-1}$ deflator.

$$I_{it} = \beta_4 MC_{it} + u_{it}. \tag{7}$$

Then as in (3) we take the left and right limits of both sides of (7) as FS_{it} goes to zero

$$\lim_{FS_{it}\downarrow 0} E\left(I_{it}\left|FS_{it}\right.\right) - \lim_{FS_{it}\uparrow 0} E\left(I_{it}\left|FS_{it}\right.\right) = \tag{8}$$

$$\beta_3\left(\lim_{FS_{it}\downarrow 0} E\left(MC_{it}\left|FS_{it}\right.\right) - \lim_{FS_{it}\uparrow 0} E\left(MC_{it}\left|FS_{it}\right.\right)\right) + \lim_{FS_{it}\downarrow 0} E\left(u_{it}\left|FS_{it}\right.\right) - \lim_{FS_{it}\uparrow 0} E\left(u_{it}\left|FS_{it}\right.\right)$$

$$= \beta_4\left(\lim_{FS_{it}\downarrow 0} E\left(MC_{it}\left|FS_{it}\right.\right) - \lim_{FS_{it}\uparrow 0} E\left(MC_{it}\left|FS_{it}\right.\right)\right), \tag{9}$$

in which the second equality follows from the assumed continuity of u_{it}. The local average treatment effect is then not simply the coefficient β_4, but the entire left side of (9). In practice, to calculate the effect of a funding violation on investment, one must multiply the regression coefficient β_4 by the difference in means between the mandatory contributions of violators and the mandatory contributions of nonviolators, the latter of which is, by definition, zero.

The results from estimating (6) are in Table 2, in which the top panel presents the results in which the regression excludes fixed firm effects, and the bottom panel presents results with fixed firm effects

The first column of Table 2 shows the results from regressing $I_{it}/A_{i,t-1}$ only on the mandatory contributions. Not surprisingly, we find a large negative coefficient. Without any other controls in the regression, this variable picks up all of the effects of the negative shocks that caused the firm to have to pay mandatory contributions in the first place. The second and third columns present a reality check in which the regression includes either $Q_{i,t-1}$ by itself or $Q_{i,t-1}$ and $CF_{it}/A_{i,t-1}$. It is reassuring to note that our results look much like those in the literature. We find a small but statistically significant coefficient on $Q_{i,t-1}$ and a larger, significant coefficient on $CF_{it}/A_{i,t-1}$.

Next, we include mandatory contributions and two measures of the distance to the cutoff for a funding violation: average funding status (Rauh's measure) and the pension gap (our measure). Not surprisingly, in either case the coefficient on mandatory contributions falls somewhat, but it still remains negative and significant. These results using average funding status are almost identical to those in Rauh, although our coefficient on mandatory contributions is somewhat smaller than his. To interpret the coefficient on mandatory contributions, we multiply the estimated value from the third column by the average mandatory contribution for payers, which is 0.0063. We get an estimated effect on investment of 0.0027, which in economic terms is a tiny number, approximately four percent of the total investment of the firms in violation. The economic significance of this number is especially small in light of the fact that most of the decrease of investment for the firms in violation comes from the bottom fifth percentile of the distribution of the pension gap. To put this point another way, we calculate the elasticity of investment with respect to mandatory contributions at the sample mean, finding a low elasticity of 0.042. Finally, these last two points are underscored by the tiny increase in the regression R^2 that occurs with the addition to the regression of either mandatory contributions or the bind variable.

We turn to an exact application of regression discontinuity, in which we include an indicator (Bind) for whether the pension gap is positive or negative. This variable is clearly discontinuous and corresponds exactly to the variable $w(s_i)$ in (1). Interestingly, for the model that contains fixed effects, the coefficient on this indicator variable is insignificantly different from zero if we also include mandatory contributions in the regression. Although this last result casts doubt on the conclusion that firms drastically cut capital expenditures in response to funding violations, it is nonetheless possible that we are not giving the Bind variable a fair chance. Because we include $MC_{it}/A_{i,t-1}$ in this regression, it might be soaking up some of the variation attributable to Bind. We therefore drop $MC_{it}/A_{i,t-1}$. Although we find a significant negative coefficient on Bind in the absence of fixed effects, when we include fixed effects this significance disappears.

We next investigate whether the significant negative coefficients on mandatory contributions are an artifact of measurement error in $Q_{i,t-1}$. We use the following simple linear framework

$$Q_{i,t-1} = \chi_{i,t-1} + \varepsilon_{i,t-1}, \tag{10}$$

in which $\chi_{i,t-1}$ is true unobserved investment opportunities and $\varepsilon_{i,t-1}$ is the measurement error. We assume $\varepsilon_{i,t-1}$ is orthogonal to $\chi_{i,t-1}$ and to all of the other variables in (6) besides $Q_{i,t-1}$. We do not use the high-order moment estimators in Erickson and Whited (2002) because our measure of $Q_{i,t-1}$ is insufficiently skewed when we include fixed effects in the regression.[4] Instead we follow

[4]The lack of sufficient skewness causes the Erickson and Whited (2002) estimators to produce large standard

Erickson and Whited (2005) and ask whether estimating (6) with OLS produces the same coefficient signs as an analogous hypothetical regression estimation in which $\chi_{i,t-1}$ has been substituted for $Q_{i,t-1}$. Erickson and Whited (2005) provide a formula for a threshold for the squared correlation between $Q_{i,t-1}$ and $\chi_{i,t-1}$ above which the OLS coefficients have the same signs as the hypothetical true coefficients. This squared correlation is also the R^2 of (10) and is therefore a useful index of measurement quality for $Q_{i,t-1}$, with a value of 1 indicating a perfect proxy and a value of 0 indicating a worthless proxy. A value of 0 indicates that $Q_{i,t-1}$ can have no worth as a proxy for investment opportunities, but that at the same time the sign of the OLS coefficient is the same as the hypothetical true coefficient. A value near 1 indicates that $Q_{i,t-1}$ can be an excellent proxy but that at the same the OLS coefficient does not have the same sign as the hypothetical true coefficient.

The results are in Table 3. For the regressions that do not contain fixed effects, these estimates hover around 0.25 for cash flow, funding status, and mandatory contributions. To get a sense of what these numbers mean, we turn to Whited (2001), who uses the technique in Erickson and Whited (2000, 2002) to estimate the squared correlation between the market-to-book ratio and true investment opportunities to be around 0.2. In other words, our estimated thresholds are not low enough to ensure with any reasonable confidence that our estimated OLS coefficients on mandatory contributions have the correct sign.

The estimated thresholds for the regressions that contain fixed effects exhibit a slightly different pattern. The thresholds for cash flow and Bind hover between 0.25 and 0.35, but the thresholds for mandatory contributions range between 0 and 0.2. Although this last result indicates that the significance of the coefficient on mandatory contributions in Table 2 may not be an artifact of measurement error in $Q_{i,t-1}$, the evidence is not overwhelming. Further, as modeled in Whited (2009), the coefficient on this kind of a cash shortfall need not indicate the presence of financial constraints. Finally, the estimated measurement quality threshold of 0 on mandatory contributions is highly sensitive to the definition of cash flow used. The measure used in Rauh, nonpension cash flow, correctly excludes mandatory contributions, but it also incorrectly excludes the nonmandatory

errors, so we can draw no useful inferences in any direction from this exercise.

part of total contributions, which are a clear cash outflow. When we account for these discretionary contributions in a measure of cash flow, the effect on the estimated coefficient on mandatory contributions is minimal. However, the estimate of the corresponding measurement quality threshold rises to 0.25.

Even if the significance of this coefficient is not an artifact of measurement error, recall that most of the identifying variation in investment comes from firms with large pension gaps. Therefore, estimating the regression (6) is subject to the same sorts of endogeneity problems that also plague garden variety investment regressions. In other words, it difficult to attribute a causal effect of mandatory contributions on investment because it is difficult to rule out the possibility that an unobserved shock to firm productivity or demand caused both a fall in investment and a fall in the pension gap.

Having established that the effect on investment from a funding violation is minimal in our full sample of firms, we turn to an examination of what decision variables "give" in response to the mandatory contributions. We therefore estimate a local average treatment effect of a funding violation on a variety of real and financial variables. The definition of a LATE in (2) provides no guidance for picking a sample that is "close" to the funding violation threshold. We follow Hoxby (2000), who starts with a tiny sample size and then examines the results when the sample size is increased. This strategy allows us to demonstrate the extent to which our results are sensitive to the cutoff point that defines a small sample.

As pointed out in Hahn, Todd, and Van der Klaauw (2001), a simple difference in means between the near violators and near escapees can have more asymptotic bias than a local linear regression of the variable of interest on a violation indicator and the distance to the point of violation. We therefore estimate the LATE using this sort of a local linear regression, in which we also include fixed firm and year effects. The results for our full sample of firms are in Figure 6. The coefficient on the violation dummy is on the vertical axis, and the pension gap is on the horizontal axis. Larger values for the gap correspond to larger sample sizes. For example, a gap of 0.0002 corresponds to a sample size of 245, and a gap of 0.001 corresponds to a sample size of 1842. The solid line represents the coefficient estimate as a function of the gap width or "bandwidth." The dashed lines represent

5% confidence intervals.

The first panel shows the response of investment. Consistent with Figure 1, we observe a small negative response for the smallest sample size. Although the response is significant, this result is fragile because it disappears with a slightly larger bandwidth. Indeed, we estimate mostly insignificant responses for all of our other "real" variables: R&D, advertising and hiring. The next two panels show a marked decrease in total debt issuance but no significant change in short term debt issuance. Recall the result in Table 1 that firms close to the gap on both sides appear to issue a great deal of short term debt. Therefore, an insignificant difference between the near violators and near escapees suggests that former are using short-term debt to make mandatory contributions and the latter are using short-term debt to avoid a violation in the first place.[5] The next panel shows that firms do not issue equity to fund pension gaps, which makes sense because a firm in poor financial health is unlikely to get a high price for a seasoned offering. The last three panels all depict significant responses. Firms appear to cut share repurchases, to cut dividends slightly, and to cut receivables a great deal. Although quite small, the dividend cut is somewhat puzzling, given that common dividends per share are sticky. To answer this puzzle we also analyze dividends per share. We find almost no change in this variable, which implies that the small but significant dividend reaction is coming either from preferred dividends or from changes in the number of shares outstanding. In sum, the bulk of the evidence shows that factoring of receivables and shareholder payout cuts are the most important sources of funds for small pension funding violations. In addition, because we are only examining local samples, we can attribute a causal effect from these cash shortfalls to the estimated responses.[6]

Figure 7 depicts the LATE estimates for firms with bond ratings. The findings are largely similar, although a few notable differences stand out. First, the negative response of repurchases disappears. Second, the response of hiring is negative and significant but of an economically tiny magnitude. Third, R&D responds negatively to the funding violation and advertising responds

[5] Although not reported in Figure 6, a similar pattern emerges for corporate saving. Firms close to the cutoff dissave a great deal, but we find no difference in the behavior of firms on different sides of the cutoff.

[6] We have also examined the level of employment, the level of cash, the change in cash, total working capital, inventories, and the levels of short- and long-term debt. We find no response for any sample size and therefore do not report these results for brevity.

positively. The seemingly anomalous latter result may be an attempt on the part of firms to boost sales in order to fund the pension gaps.

Figure 8 presents the LATE estimates for the set of firms without bond ratings. For this set of firms we observe responses on many different margins, both real and financial. Near violators issue significantly less debt than near escapees, which we speculate may indicate a reluctance of banks to use longer term debt to finance funding gaps. In contrast, we observe a significant rise in short-term debt for the near violators, a strong and significant drop in hiring, receivables, and dividends. In contrast to the case of the full sample, the drop in total dividends is in part a result of a cut in common dividends per share. This last result is makes intuitive sense because the firms close to the cutoff are those paying high total dividends and high dividends per share. They consequently have ample room with which to decrease dividends. Finally, although we also observe a negative response of advertising, this result depends too much on the bandwidth to be credible.

In sum, Figures 6-8 paint a picture of firms that have ample financial means to deal with violations of pension funding status. As shown in Whited (2009), these types of local regression do a good job of identifying variables that have low costs of adjustment. Therefore, our evidence paints a picture of small financial frictions and much larger real adjustment frictions. This result stands in strong contrast to that in Rauh, who finds a strong reaction of investment to mandatory pension contributions. Most of the difference lies in our strong identification that comes from estimating a local treatment effect. In addition, our large sample results that show a minimal response of investment bolster the typical criticism that estimating a LATE has low power because of the small sample size. One word of caution about these results is that they do not have strong external validity. In other words, except for the case of investment, it is hard to extrapolate these inferences beyond the local sample.

5. Conclusion

This paper has sought to understand how firms react to exogenous cash shortfalls. On a purely theoretical basis, one would expect them to adjust on the margins that entail the fewest costs. To answer the question empirically, we use a regression discontinuity design, in which the discontinuity is the point of violation of underfunding of corporate defined benefit pension plans. We reexamine the evidence in Rauh that the mandatory pension contributions triggered by pension funding violations have a causal effect on investment. We find that this evidence is driven by data variation far away from the point of discontinuity and is therefore likely to be an artifact of endogeneity.

We also examine firm-year observations in which the firm's pension assets are just barely less than its pension liabilities. We compare this group to a control group of firm-year observations in which the firm has just barely escaped having to make a mandatory contribution. In this quasi-experimental setting, we find little evidence that firms cut back on investment. Instead, they use a variety of financial tools, such as factoring of receivables and shareholder payout cuts, to fund their pension liabilities. This evidence suggests that in the sample we consider, the costs of financial adjustment are much less substantial than the costs of real adjustment. Further work to examine this question on a broader scale would be interesting.

Finally, we have provided applied financial economics researchers with several practical guidelines for using regression discontinuity. In particular, we use a graphical analysis to determine whether the variable of interest actually exhibits a jump at the point of discontinuity and to determine whether data variation near or far from the breakpoint drives any results. Similarly, we also check whether observations close to the breakpoint differ systematically from those that are far from the breakpoint. As advocated in McCrary (2008), we also check to see whether firms bunch up on one side of the breakpoint or the other. We emphasize that it is difficult to interpret results when one includes a discontinuous variable that is not a zero-one indicator. Finally, we do a sensitivity analysis when analyzing observations around the breakpoint to ensure that any results are not driven by a choice of bandwidth.

Appendix

For our data from Compustat we define book assets as item 6, cash as item 1, inventories as item 3, working capital as item 4 minus item 5, capital expenditures as item 128, dividends as the sum of items 19 and 21, long-term debt as item 9, short term debt as item 34 plus item 44, the number of common shares as item 25, the share price as item 199, balance sheet deferred taxes as item 74, equity issuance as item 108, employment as item 29, the book value of common equity as item 60, and share repurchases as item 115. As in Rauh we define nonpension cash flow as the sum of items 14, 18, and 43, and cash flow as nonpension cash flow minus total pension contributions from IRS form 5500. The numerator of the market-to-book ratio is book assets minus the book value of common equity minus deferred taxes plus the product of the share price and the number of shares outstanding. The denominator is total book assets. A bond rating is the S&P long-term domestic issuer credit rating (SPDRC).

References

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.

Chava, Sudheer and Michael R. Roberts, 2008, How does financing impact investment? The role of debt covenants, Journal of Finance 63, 2085-2121.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, 2005, Nominal rigidities and the dynamic effects of a shock to monetary policy, Journal of Political Economy 113, 1–45.

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, 2005, Proxy quality thresholds: Theory and applications, Finance Research Letters 2, 131–151.

Fazzari, Steven, R. Glenn Hubbard, and Bruce Petersen, 1988, Financing constraints and corporate investment, Brookings Papers on Economic Activity 1, 144-195.

Gomes, Joao, 2001, Financing investment, American Economic Review 91, 1263–1285.

Hahn, J., P. Todd and W. Van der Klaauw, 2001, Identification and estimation of treatment effects with a regression discontinuity design, Econometrica 69, 201-209.

Heckman, James J. and Richard Robb, 1985, Alternative methods for evaluating the impact of interventions, in J. Heckman and B. Singer eds., Longitudinal Analysis of Labor Market Data, Cambridge: Cambridge University Press.

Hennessy, Christopher A. and Toni M. Whited, 2007, How costly is external financing? Evidence from a structural estimation, Journal of Finance 62, 1705–1745

Hoxby, Caroline, 2000, The effect of class size on student achievement: New evidence from population variation, Quarterly Journal of Economics 115, 1239-1285.

Imbens, Guido W. and Thomas Lemieux, 2008, Regression discontinuity designs: A guide to practice, Journal of Econometrics 142, 615-635.

Lamont, Owen, 1997, Cash flow and investment: evidence from internal capital markets, Journal of Finance 52, 83-109.

McCrary, J., 2008, Testing for manipulation of the running variable in the regression discontinuity design, forthcoming, Journal of Econometrics.

Munnell, Alicia, and Mauricio Soto, 2003, The outlook for pension contributions and profits,Working Paper, Fifth Annual Conference of the Retirement Research Consortium.

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.

Silverman, Bernard, 1986, Density Estimation for Statistics and Data Analysis, Monographs on Statistics and Applied Probability Volume: 26, (Chapman and Hall, London).

Stein, Jeremy, 2003, Agency, information and corporate investment, Handbook of the Economics of Finance, G.M. Constantinides, M.Harris and R.Stulz (eds.), Elsevier Science, 109-163.

Stokey, Nancy L. and Robert E. Lucas, 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.

Thistlethwaite, D., and D. Campbell, 1960, Regression-discontinuity analysis: An alternative to the ex post facto experiment, Journal of Educational Psychology 51, 309-317.

Van der Klaauw, Wilbert, 2002, Estimating the effect of financial aid offers on college enrollment: A regression-discontinuity approach, International Economic Review 43, 1249-1287

Whited, Toni M., 2001, Is it inefficient investment that causes the diversification discount? Journal of Finance 56, 1667–92.

Whited, Toni M., 2009, What can cash shortfalls and windfalls tell us about finance constraints? in Giorgio Calcagnini and Enrico Saltari (eds.) *The Economics of Imperfect Markets* Springer, New York, forthcoming.

Table 1: Summary Statistics

	Full Sample	Bond Rating	No Bond Rating	In Violation	Not in Violation	Close	Far	Bottom 5%
Total Assets	3,413	6,331	1,384	2,608	3,638	10,530	3,048	2,674
Average Funding Status Indicator	0.2185	0.1672	0.2541	1.0000	0.0000	0.0675	0.2261	0.9492
Plan Violation Indicator	0.3510	0.3416	0.3576	1.0000	0.1696	0.4026	0.3483	1.0000
Average Funding Status	0.0354	0.0357	0.0352	-0.0155	0.0496	0.0444	0.0350	-0.0461
Gap	0.0151	0.0134	0.0162	-0.0169	0.0240	0.0000	0.0158	-0.0550
Total Contributions	0.0034	0.0027	0.0039	0.0067	0.0025	0.0015	0.0035	0.0146
Pension Assets	0.1427	0.1331	0.1494	0.1197	0.1491	0.1141	0.1442	0.2972
Pension Liabilities	0.1061	0.0974	0.1122	0.1344	0.0982	0.0689	0.1080	0.3411
Mandatory Contributions	0.0016	0.0010	0.0020	0.0063	0.0002	0.0001	0.0016	0.0171
Investment	0.0694	0.0729	0.0670	0.0634	0.0711	0.0689	0.0694	0.0544
Cash Flow	0.0959	0.1084	0.0872	0.0789	0.1006	0.1034	0.0955	0.0562
Nonpension Cash Flow	0.0995	0.1112	0.0914	0.0864	0.1032	0.1049	0.0993	0.0733
Market-to-Book	1.4811	1.5782	1.4136	1.4207	1.4980	1.5531	1.4774	1.3482
R&D	0.0166	0.0156	0.0173	0.0168	0.0165	0.0169	0.0166	0.0156
Advertising	0.0121	0.0141	0.0107	0.0109	0.0125	0.0187	0.0118	0.0079
Debt-to-Assets	0.2488	0.2741	0.2313	0.2750	0.2415	0.2281	0.2499	0.2322
Bond Rating	0.4100	1.0000	0.0000	0.3138	0.4369	0.6675	0.3969	0.2234
Short Term Debt Issuance	0.0055	0.0070	0.0044	0.0046	0.0057	0.0111	0.0052	0.0018
Long Term Debt Issuance	0.0211	0.0300	0.0149	0.0228	0.0206	0.0104	0.0216	0.0110
Saving	0.0040	0.0045	0.0037	0.0034	0.0042	-0.0002	0.0042	0.0028
Cash	0.0701	0.0569	0.0792	0.0702	0.0700	0.0636	0.0704	0.0716
Dividends	0.0187	0.0235	0.0155	0.0114	0.0208	0.0228	0.0185	0.0102
Common Dividends per Share	0.6147	0.8918	0.4221	0.3374	0.6922	0.9454	0.5978	0.3298
Equity Issuance	0.0124	0.0139	0.0113	0.0152	0.0116	0.0088	0.0125	0.0100
Equity Repurchases	0.0116	0.0150	0.0093	0.0100	0.0121	0.0152	0.0115	0.0072
Employment % Change	0.8098	2.2640	-0.2008	0.3395	0.9412	0.3955	0.8311	-2.5530
Earnings	0.0418	0.0556	0.0323	0.0247	0.0466	0.0507	0.0414	0.0073
Z-Score	2.7805	2.7713	2.7864	2.0147	2.9954	2.8747	2.7761	1.1838

The figures presented are means from a sample of unregulated and nonfinancial firms from the 2007 Compustat annual industrial files. The sample period is from 1990 to 1998. Pension data are from IRS Form 5500. Mandatory contributions are federally mandated contributions to a pension plan with assets less than liabilities. Total contributions are mandatory plus voluntary contributions. Plan Violation Indicator equals one if the firm is making mandatory pension contributions, and zero otherwise. Average funding status is the average across plans of pension assets minus pension liabilites, as a fraction of total firm assets. Average funding status indicator is one if average funding status is positive, and zero otherwise. Gap is the smallest funding gap if all plans have positive gaps, and the sum of the negative gaps if any plans have negative gaps. Bond rating is a dummy that takes a value of 1 if a firm has an S&P long-term bond rating, and zero otherwise. The denominator of market-to-book ratio is total book assets. The numerator is book assets minus the book value of equity plus the market value of equity, and the denominator is book assets. The other variables are self explanatory and are all deflated by total book assets.

Table 2: Global Investment Regressions

No Fixed Effects							
Market-to-Book		0.0135	0.0004	0.0003	0.0004	0.0003	0.0004
		(0.0008)	(0.0009)	(0.0009)	(0.0009)	(0.0009)	(0.0009)
Non-Pension Cash Flow			0.2500	0.2486	0.2476	0.2474	0.2490
			(0.0082)	(0.0082)	(0.0082)	(0.0082)	(0.0082)
MPCs	-0.9520			-0.7233	-0.6526	-0.5764	
	(0.1350)			(0.1340)	(0.1410)	(0.1500)	
Funding Status				-0.0391			
				(0.0094)			
Gap					-0.0274	-0.0335	-0.0152
					(0.0147)	(0.0152)	(0.0145)
Bind						-0.0021	-0.0039
						(0.0014)	(0.0013)
R^2	0.0109	0.0369	0.1388	0.1427	0.1412	0.1414	0.1398
Fixed Effects							
Market-to-Book		0.0236	0.0185	0.0185	0.0185	0.0185	0.0185
		(0.0011)	(0.0011)	(0.0011)	(0.0011)	(0.0011)	(0.0011)
Non-Pension Cash Flow			0.1120	0.1131	0.1128	0.1127	0.1118
			(0.0077)	(0.0077)	(0.0077)	(0.0077)	(0.0077)
MPCs	-0.4507			-0.4261	-0.4153	-0.3692	
	(0.1366)			(0.1467)	(0.1491)	(0.1599)	
Funding Status				0.0209			
				(0.0106)			
Gap					0.0325	0.0320	0.0484
					(0.0163)	(0.0163)	(0.0147)
Bind						-0.0010	-0.0020
						(0.0012)	(0.0012)
R^2	0.6445	0.6648	0.6737	0.6747	0.6747	0.6747	0.6745

Estimates are from a sample of unregulated and nonfinancial firms from the 2007 Compustat annual industrial files. The sample period is from 1990 to 1998. Pension data are from IRS Form 5500. MPCs are mandatory pension contributions to underfunded plans, scaled by total assets. Funding status is the average across plans of pension assets minus pension liabilites, as a fraction of total assets. Gap is the the smallest funding gap if all plans have positive gaps, and the sum of the negative gaps if any plans have negative gaps. Non pension cash flow is income plus depreciation plus pension expense, all scaled by total assets. Bind is a dummy variable for whether the variable "gap" is positive or negative. All regressions contain fixed year effects. Standard errors are corrected for heteroskedasticity and clustering at the industry level.

Table 3: Measurement Quality Thresholds

No Fixed Effects						
Market-to-Book	—	0	0	0	0	0
Non-Pension Cash Flow		0.2534 (0.0223)	0.2524 (0.0225)	0.2533 (0.0225)	0.2532 (0.0226)	0.2537 (0.0225)
MPCs	—		0.2487 (0.0139)	0.2496 (0.0147)	0.2495 (0.0142)	
Funding Status			0.2498 (0.0159)			
Gap				0	0	0
Bind					0	0
Fixed Effects						
Market-to-Book	—	0	0	0	0	0
Non-Pension Cash Flow		0.3792 (0.0222)	0.3780 (0.0219)	0.3781 (0.0222)	0.3780 (0.0219)	0.3793 (0.0221)
MPCs	—		0	0	0	
Funding Status			0			
Gap				0	0	0
Bind					0.3437 (0.1469)	0.2422 (0.0990)

Estimates are from a sample of unregulated and nonfinancial firms from the 2007 Compustat annual industrial files. The sample period is from 1990 to 1998. Pension data are from IRS Form 5500. MPCs are mandatory pension contributions to underfunded plans, scaled by total assets. Funding status is the average across plans of pension assets minus pension liabilites, as a fraction of total assets. Gap is the the smallest funding gap if all plans have positive gaps, and the sum of the negative gaps if any plans have negative gaps. Non pension cash flow is income plus depreciation plus pension expense, all scaled by total assets. All regressions contain fixed year effects. Bind is a dummy variable for whether the variable "gap" is positive or negative. The measurement quality threshold is a level that an index of measurement quality for the mismeasured regressor (true investment opportunities) must surpass a threshold in order for a coefficient to retain the sign obtained via OLS. The index is the squared correlation between true unobservable investment opportunities and the observed market-to-book ratio above. The regressions in this table correspond to those in Table 2.

Figure 1: Kernel Regression of Investment on the Pension Gap

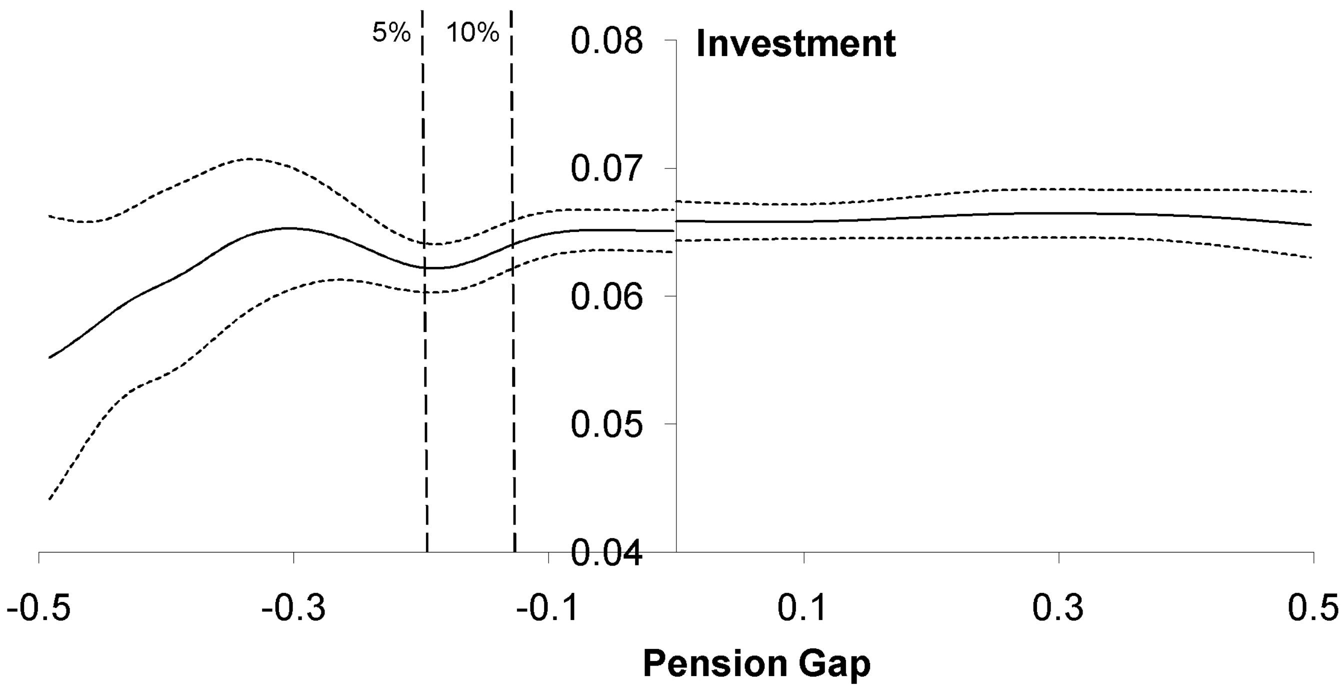

Calculations are based on a sample of firms from Compustat from 1990 to 1998. The graph depicts the univariate relation between the pension gap and investment. The pension gap is defined as the minimum plan surplus if all of the plans in a firm have a surplus and as the sum of the deficits if at least one of a firm's plans has a deficit. Investment is scaled by total firm assets and the pension gap is scaled by pension liabilities. The kernel regression estimation is performed using a Gaussian kernel, and the bandwidth is chosen using cross-validation. Separate regressions are performed for positive and negative pension gaps.

Figure 2: Kernel Regressions of q, Cash Flow and MPCs on the Pension Gap

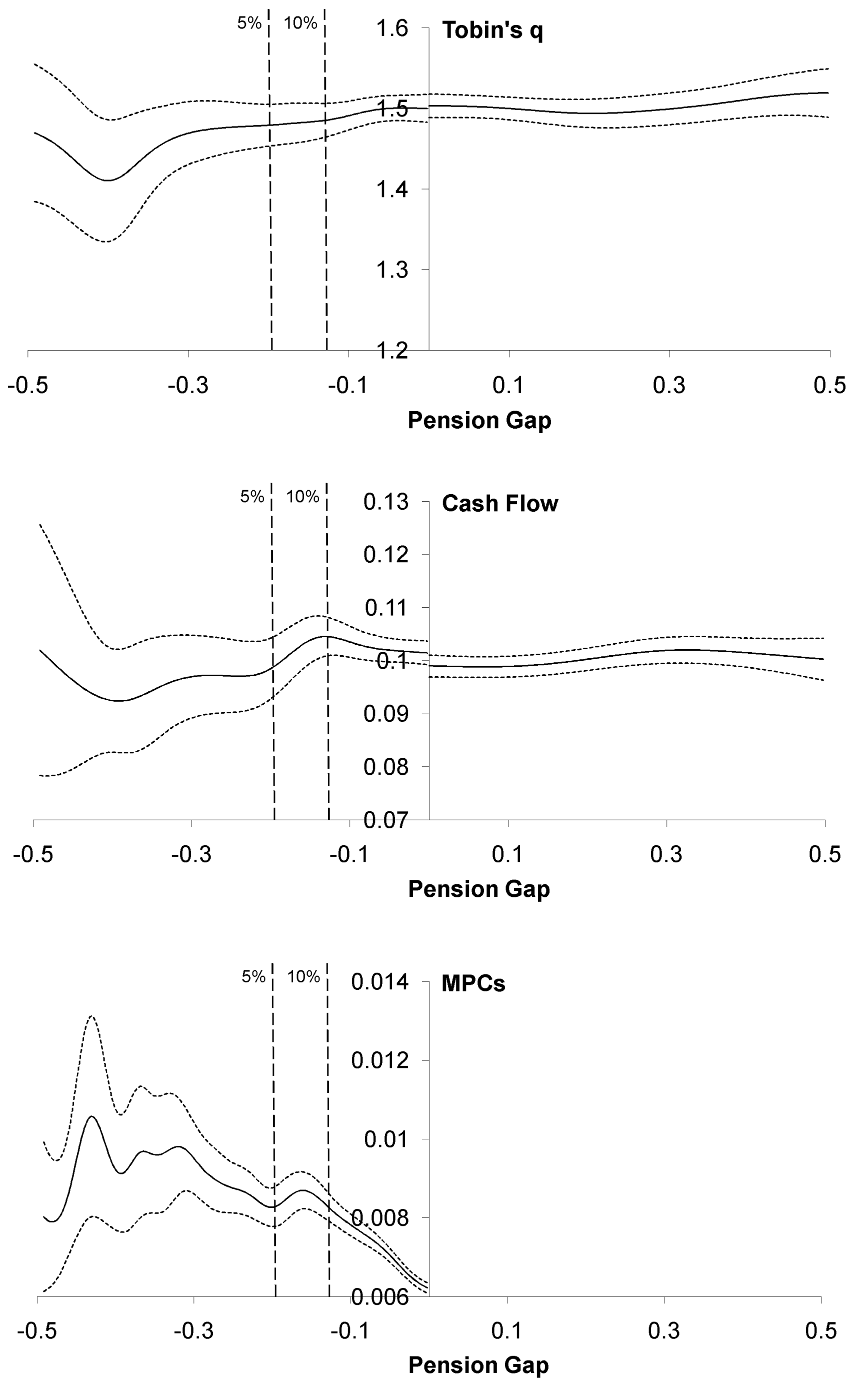

Calculations are based on a sample of firms from Compustat from 1990 to 1998. The graph depicts the univariate relations between the pension gap and Tobin's q, cash flow, and mandatory contributions. The pension gap is defined as the minimum plan surplus if all of the plans in a firm have a surplus and as the sum of the deficits if at least one of a firm's plans has a deficit. Tobin's q is calculated as the market-to-book ratio. Cash flow and mandatory contributions are scaled by total firm assets and the pension gap is scaled by pension liabilities. The kernel regression estimation is performed using a Gaussian kernel, and the bandwidth is chosen using cross-validation. Separate regressions are performed for positive and negative pension gaps.

Figure 3: Histogram of Average Funding Status

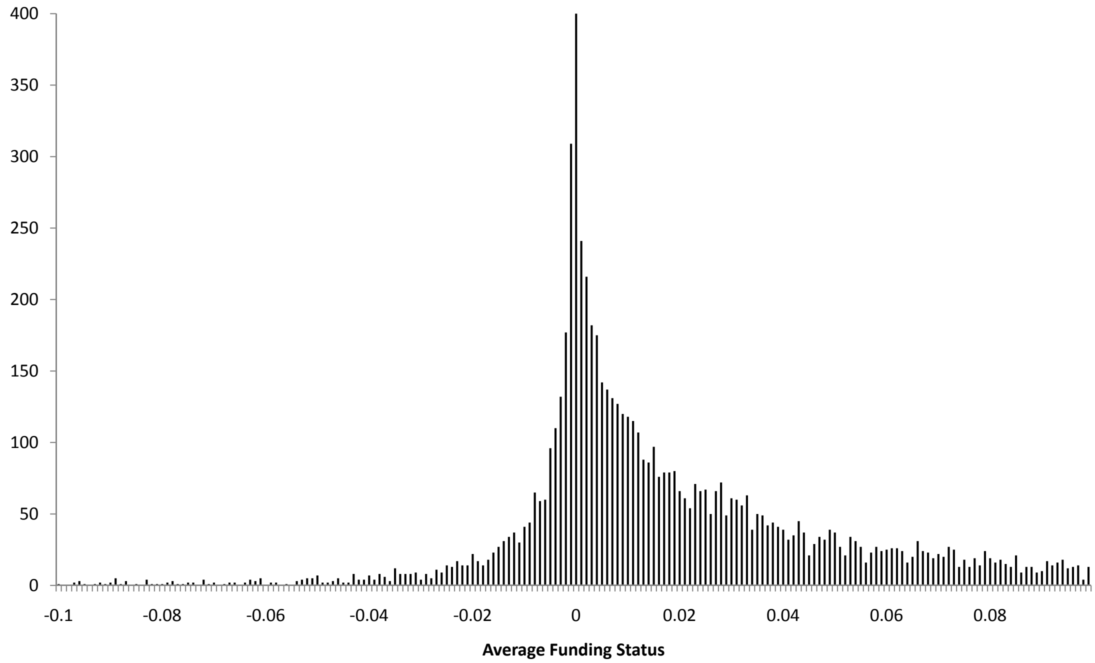

Calculations are based on a sample of firms from Compustat from 1990 to 1998. Funding status is the sum of the pension assets in all of a firm's pension plans less the sum of the corresponding liabilities.

Figure 4: Histogram of Pension Gaps

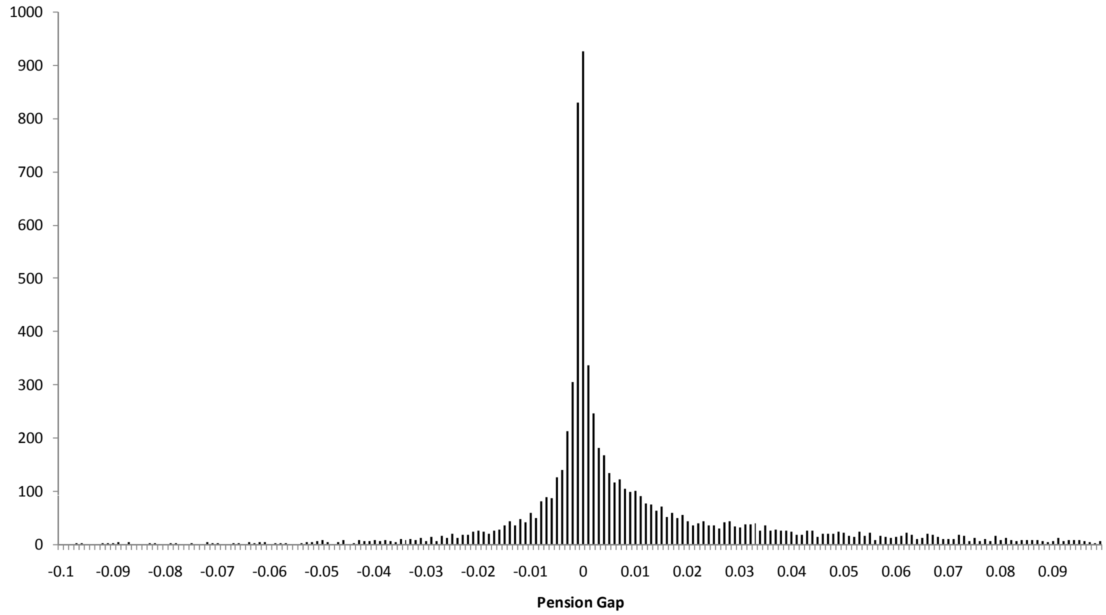

Calculations are based on a sample of firms from Compustat from 1990 to 1998. The pension gap is defined as the minimum plan surplus if all of the plans in a firm have a surplus and as the sum of the deficits if at least one of a firm's plans has a deficit. This figure depicts a histogram of funding gaps as a fraction of total firm assets.

Figure 5: Magnified Histogram of Pension Gaps

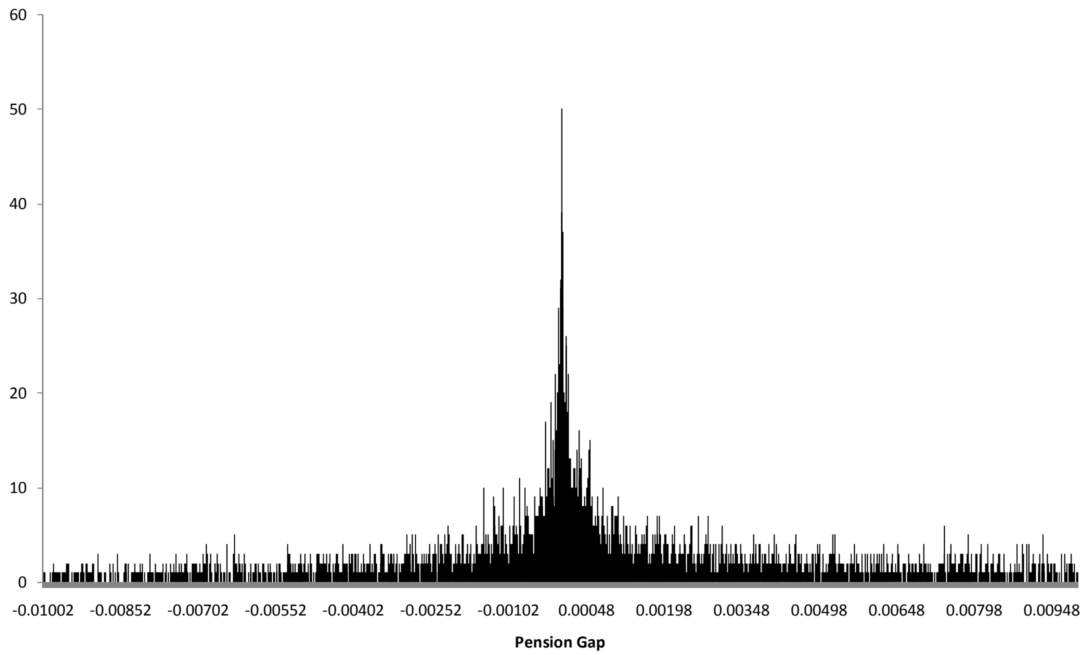

Calculations are based on a sample of firms from Compustat from 1990 to 1998. The pension gap is defined as the minimum plan surplus if all of the plans in a firm have a surplus and as the sum of the deficits if at least one of a firm's plans has a deficit. This figure depicts a histogram of funding gaps as a fraction of total firm assets.

Figure 6: Local Responses to Funding Violations: Full Sample

Calculations are based on a sample of firms from Compustat from 1990 to 1998. This figure depicts the local response of a variety of variables to a pension funding violation. A local response is the coefficient on a pension violation indicator in a regression of a variable of interest on this indicator, firm fixed effects, and fixed year effects. The horizontal axis depicts the sample used for this regression. The sample is defined by the gap window, which is the absolute value of a pension gap as a fraction of firm assets. The pension gap is in turn defined as the minimum plan surplus if all of the plans in a firm have a surplus and as the sum of the deficits if at least one of a firm's plans has a deficit. The sample used for the calculation is all observations for which the absolute value of the pension gaps is less than the specified value on the horizontal axis. All variables except employment and hiring are expressed as a fraction of total assets. Employment is measured as the natural log of millions of employees, and hiring is the change in employment.

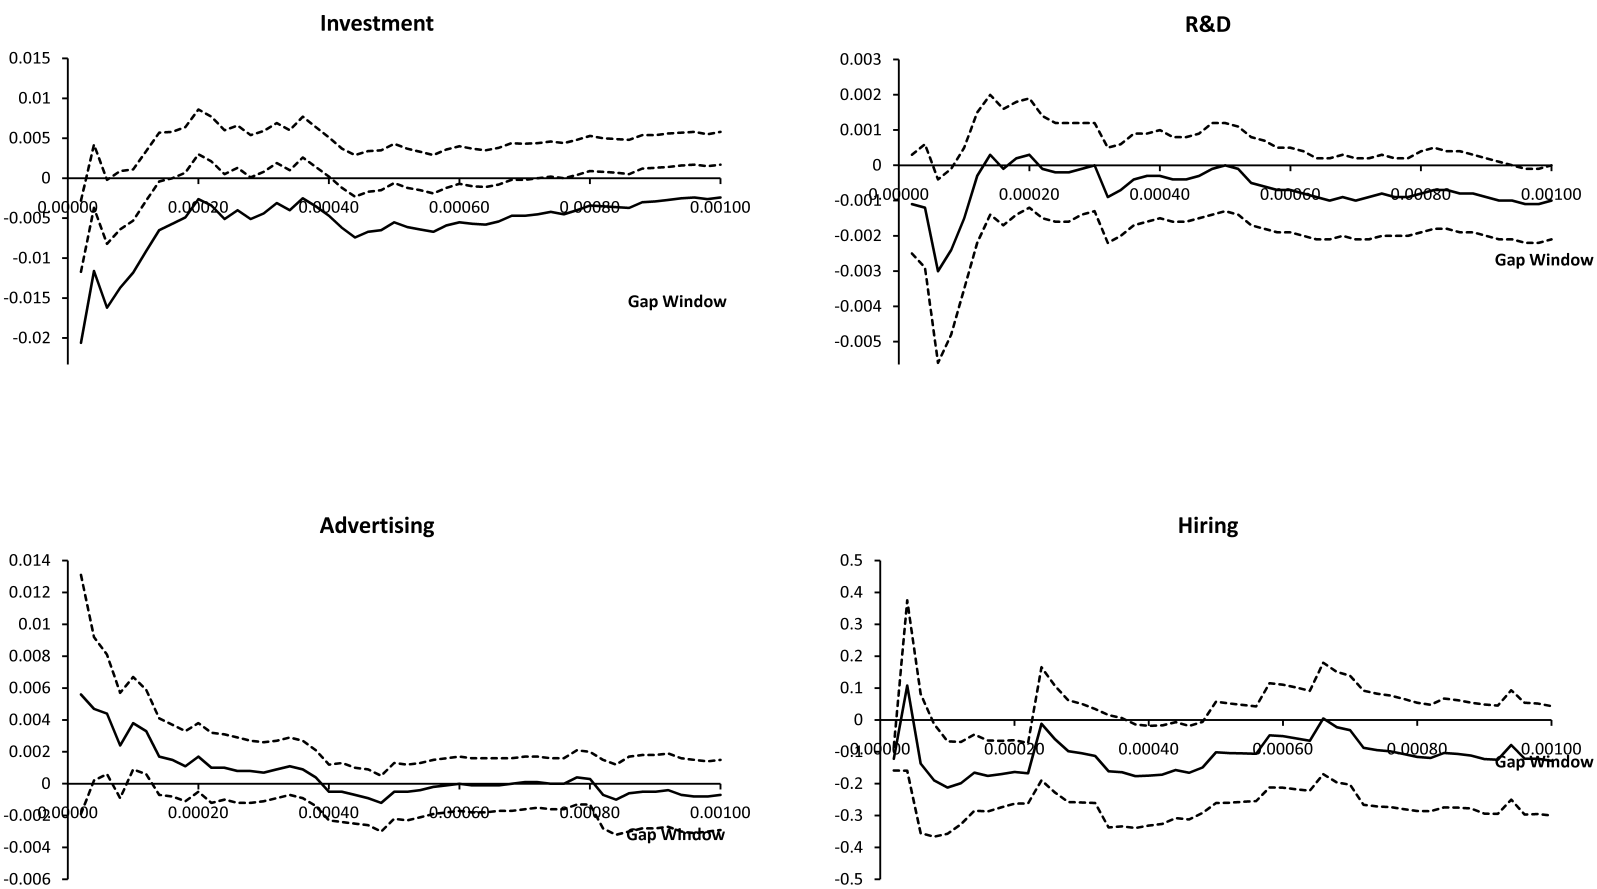

Debt Issuance

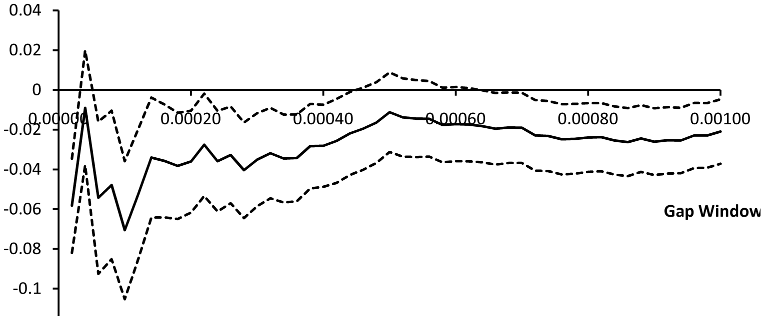

Short Term Debt Issuance

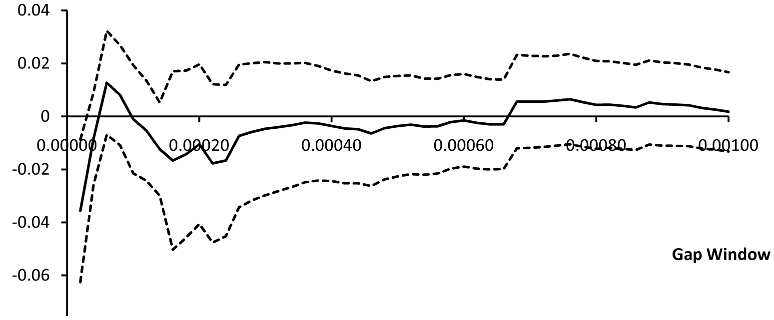

Equity Issuance

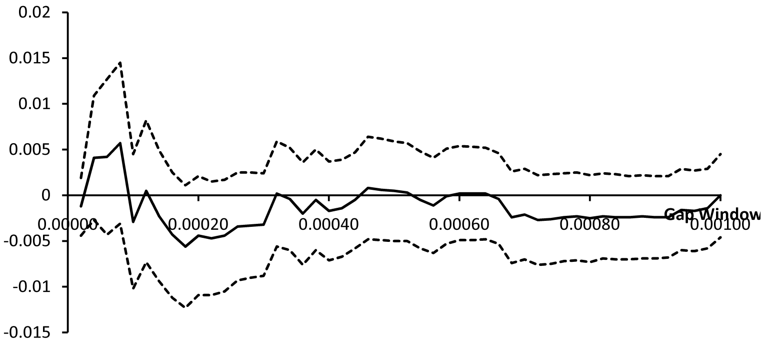

Repurchases

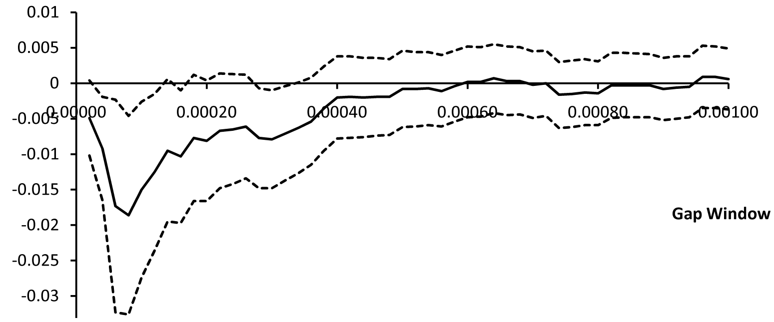

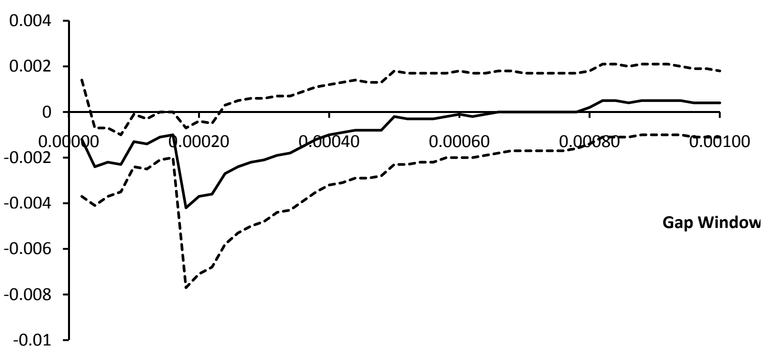

Receivables

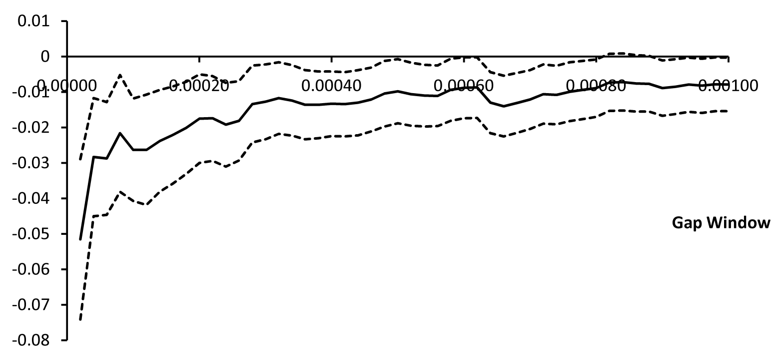

Figure 7: Local Responses to Funding Violations: Firms with Bond Ratings

Calculations are based on a sample of firms from Compustat from 1990 to 1998. This figure depicts the local response of a variety of variables to a pension funding violation. A local response is the coefficient on a pension violation indicator in a regression of a variable of interest on this indicator, firm fixed effects, and fixed year effects. The horizontal axis depicts the sample used for this regression. The sample is defined by the gap window, which is the absolute value of a pension gap as a fraction of firm assets. The pension gap is in turn defined as the minimum plan surplus if all of the plans in a firm have a surplus and as the sum of the deficits if at least one of a firm's plans has a deficit. The sample used for the calculation is all observations for which the absolute value of the pension gaps is less than the specified value on the horizontal axis. All variables except employment and hiring are expressed as a fraction of total assets. Employment is measured as the natural log of millions of employees, and hiring is the change in employment.

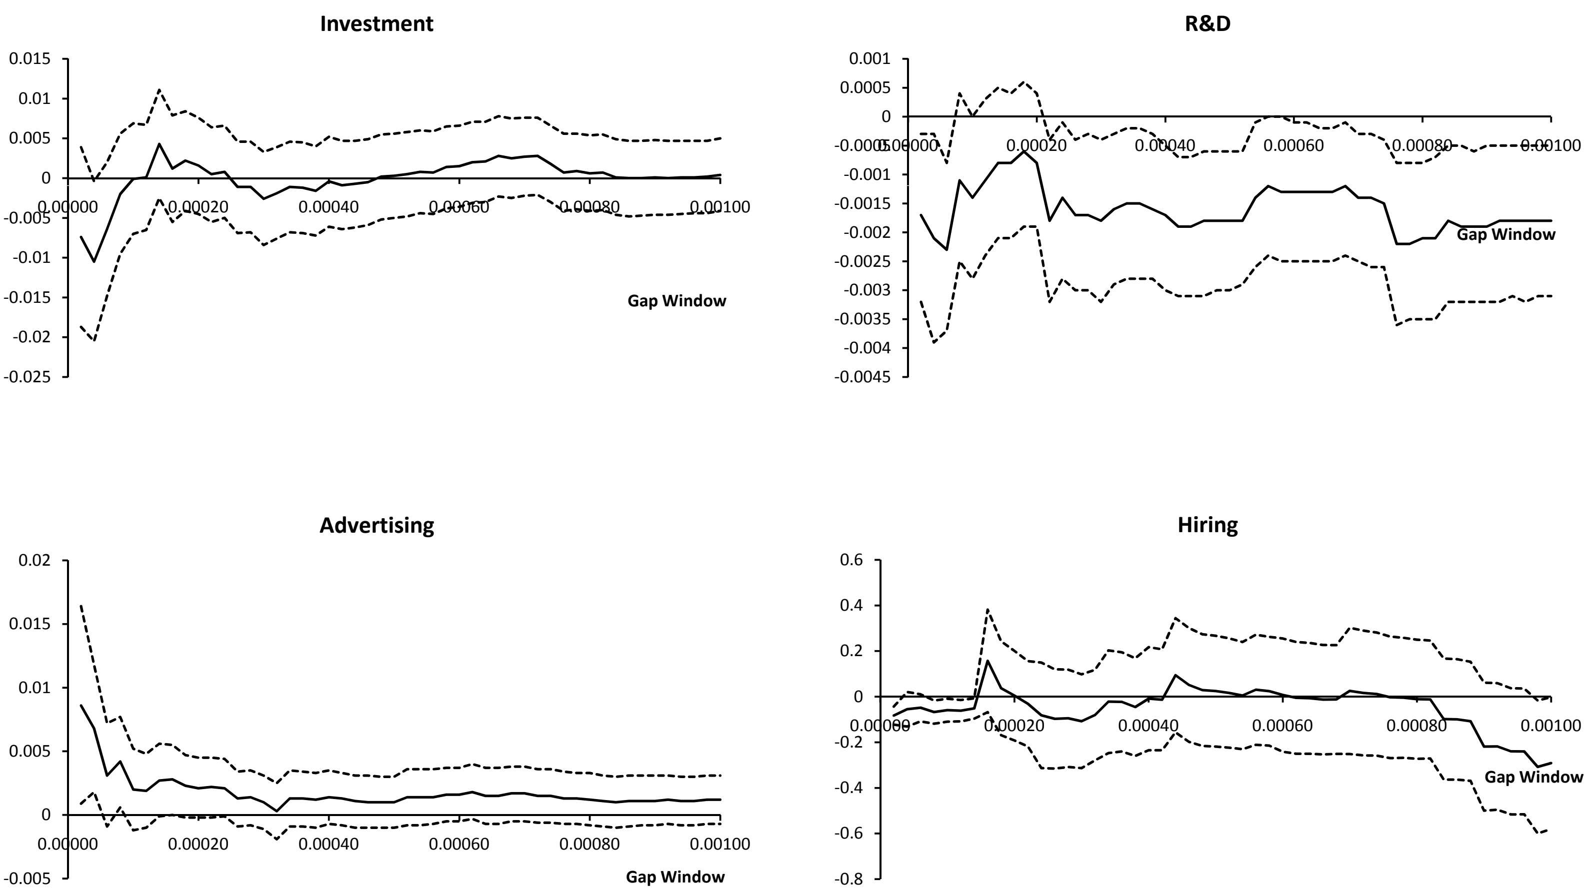

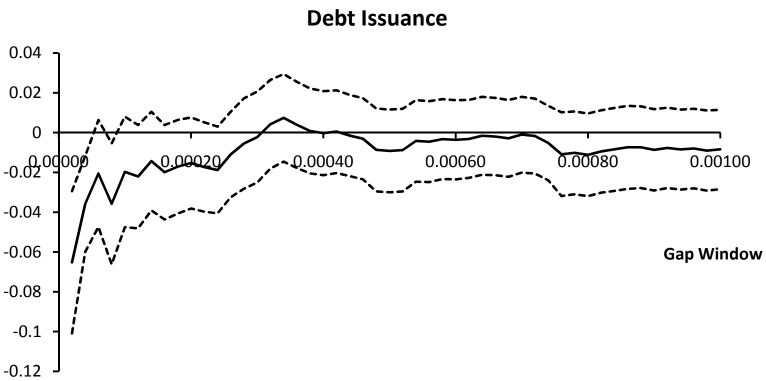
Debt Issuance
0.04
0.02
0
-0.02
-0.04
-0.06
-0.08
-0.1
-0.12
0.00000
0.00020
0.00040
0.00060
0.00080
0.00100
Gap Window

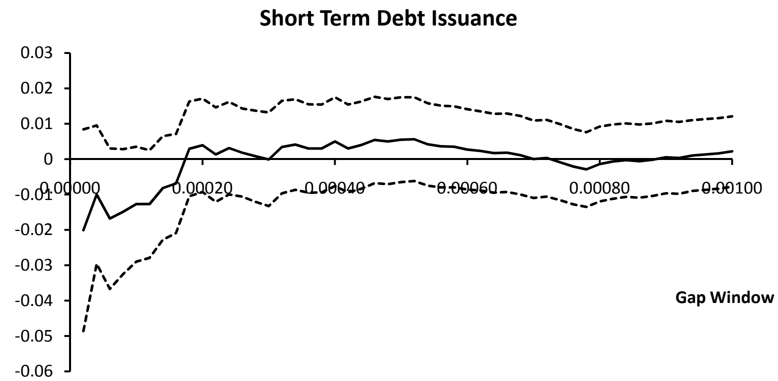
Short Term Debt Issuance
0.03
0.02
0.01
0
-0.01
-0.02
-0.03
-0.04
-0.05
-0.06
0.00000
0.00020
0.00040
0.00060
0.00080
0.00100
Gap Window

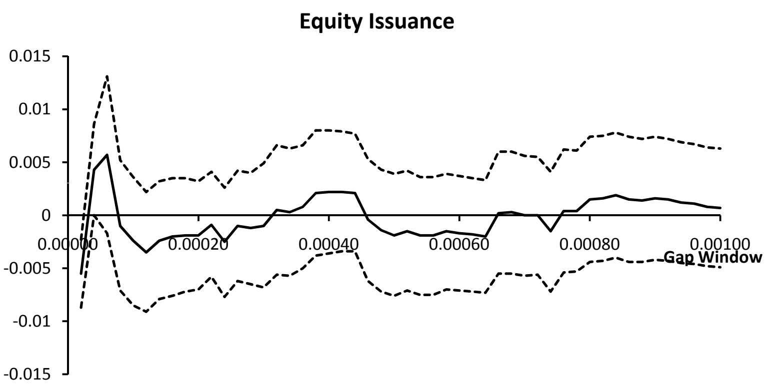
Equity Issuance
0.015
0.01
0.005
0
-0.005
-0.01
-0.015
0.00000
0.00020
0.00040
0.00060
0.00080
0.00100
Gap Window

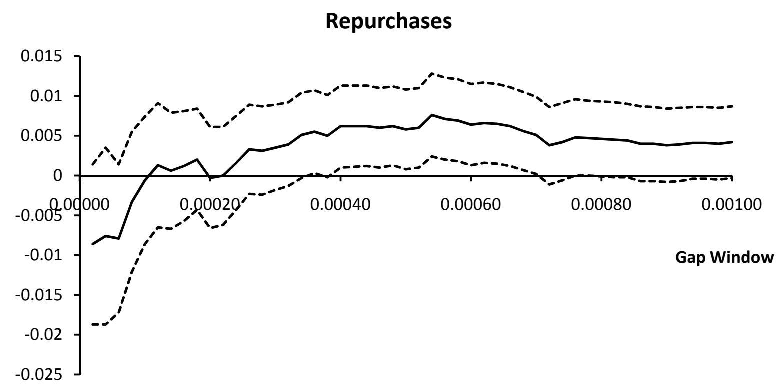
Repurchases
0.015
0.01
0.005
0
-0.005
-0.01
-0.015
-0.02
-0.025
0.00000
0.00020
0.00040
0.00060
0.00080
0.00100
Gap Window

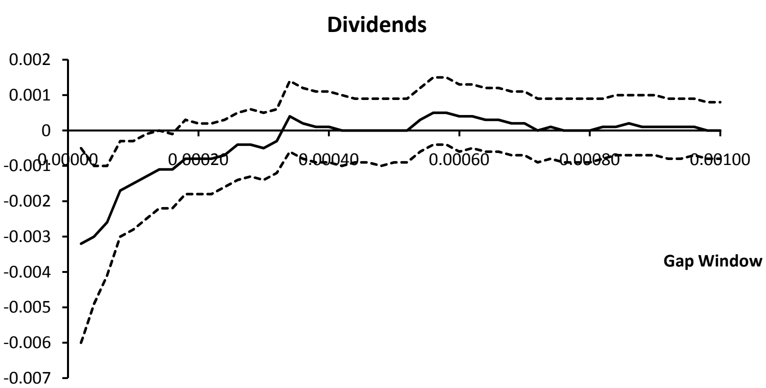
Dividends
0.002
0.001
0
-0.001
-0.002
-0.003
-0.004
-0.005
-0.006
-0.007
0.00000
0.00020
0.00040
0.00060
0.00080
0.00100
Gap Window

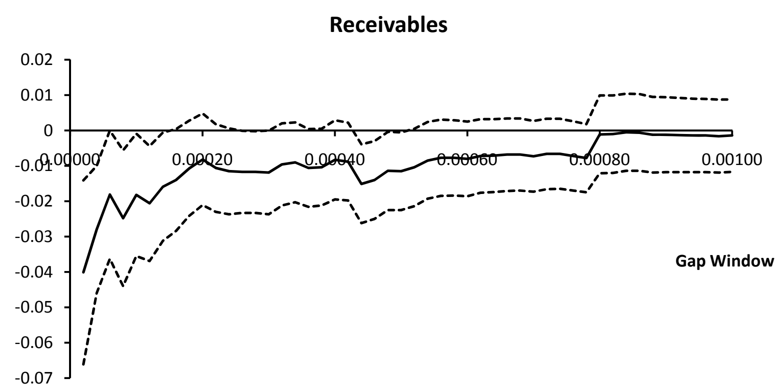
Receivables
0.02
0.01
0
-0.01
-0.02
-0.03
-0.04
-0.05
-0.06
-0.07
0.00000
0.00020
0.00040
0.00060
0.00080
0.00100
Gap Window

Figure 8: Local Responses to Funding Violations: Firms without Bond Ratings

Calculations are based on a sample of firms from Compustat from 1990 to 1998. This figure depicts the local response of a variety of variables to a pension funding violation. A local response is the coefficient on a pension violation indicator in a regression of a variable of interest on this indicator, firm fixed effects, and fixed year effects. The horizontal axis depicts the sample used for this regression. The sample is defined by the gap window, which is the absolute value of a pension gap as a fraction of firm assets. The pension gap is in turn defined as the minimum plan surplus if all of the plans in a firm have a surplus and as the sum of the deficits if at least one of a firm's plans has a deficit. The sample used for the calculation is all observations for which the absolute value of the pension gaps is less than the specified value on the horizontal axis. All variables except employment and hiring are expressed as a fraction of total assets. Employment is measured as the natural log of millions of employees, and hiring is the change in employment.

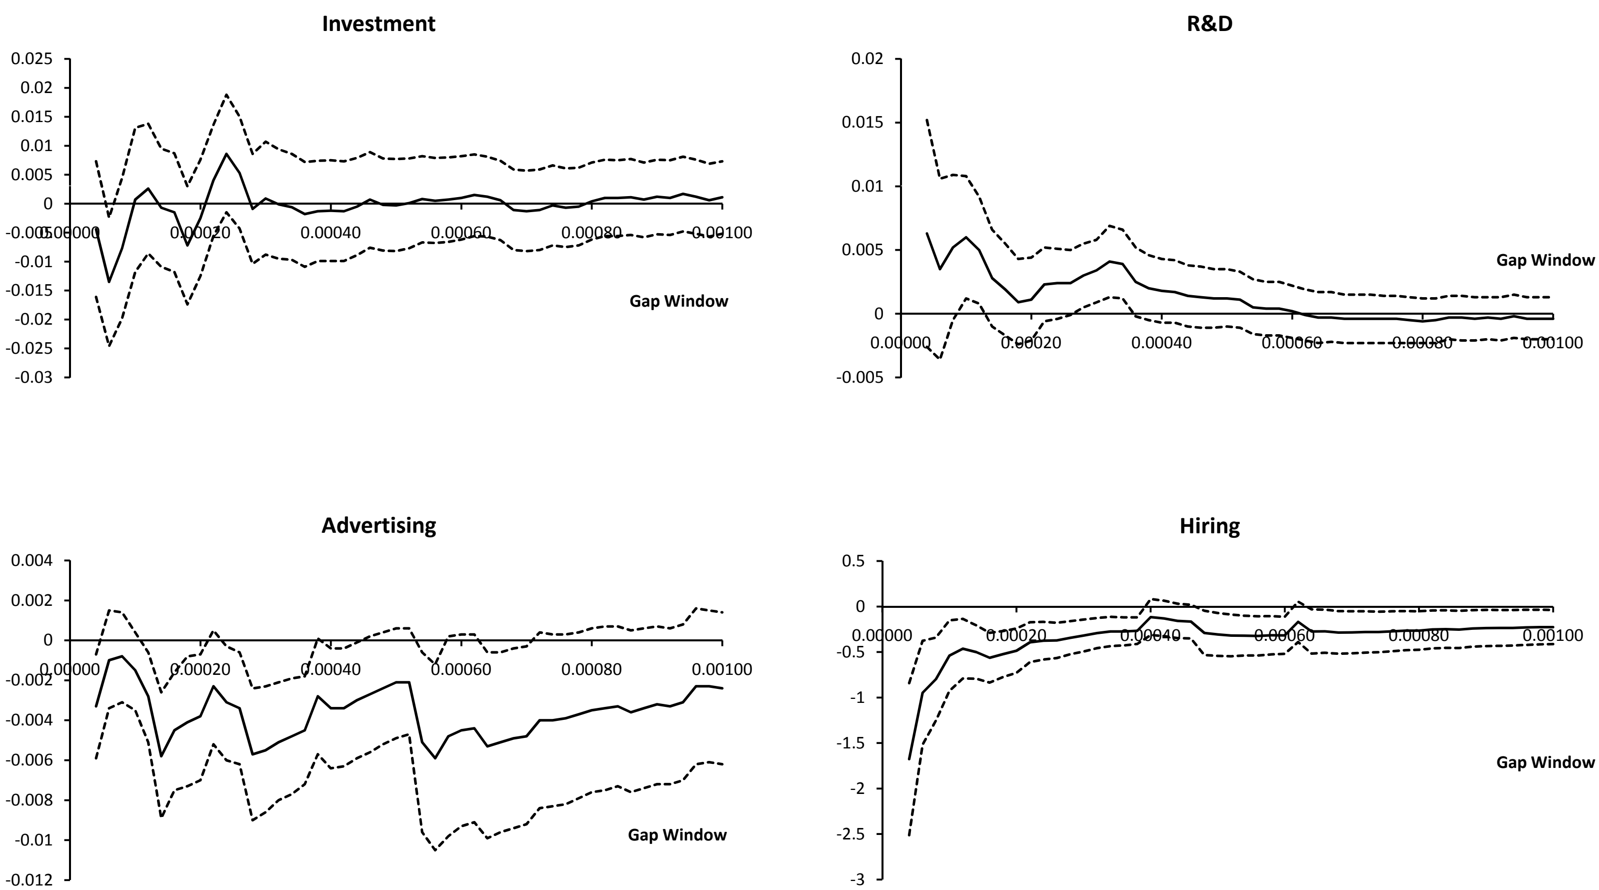

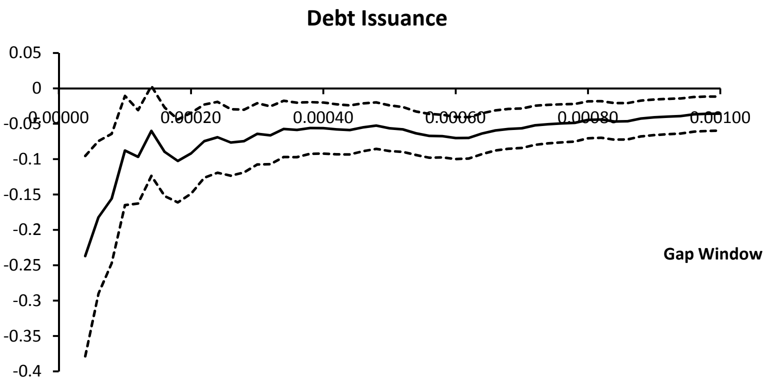
Debt Issuance
0.05
0
-0.05
-0.1
-0.15
-0.2
-0.25
-0.3
-0.35
-0.4
0.00000
0.00020
0.00040
0.00060
0.00080
0.00100
Gap Window

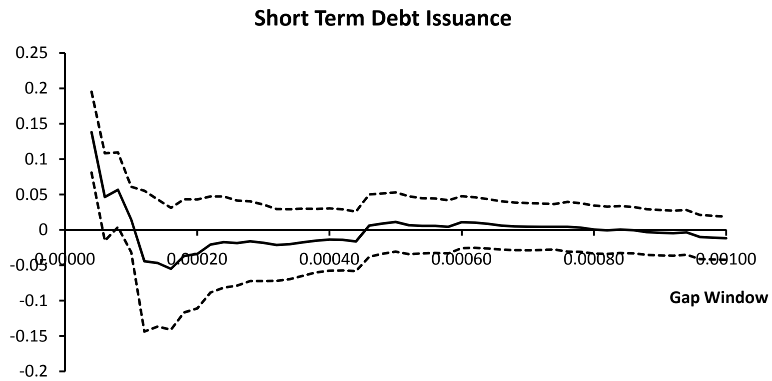
Short Term Debt Issuance
0.25
0.2
0.15
0.1
0.05
0
-0.05
-0.1
-0.15
-0.2
0.00000
0.00020
0.00040
0.00060
0.00080
0.00100
Gap Window

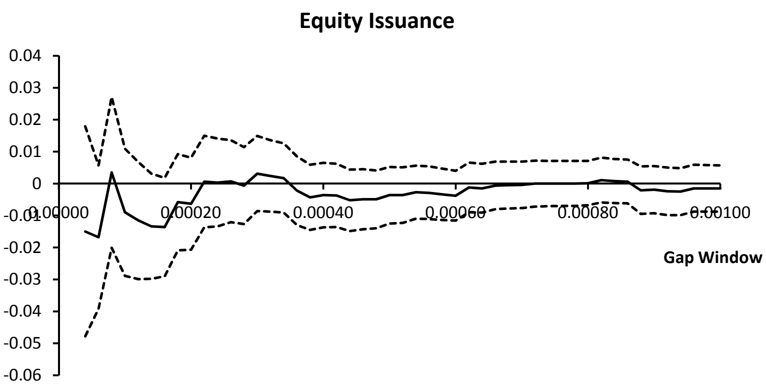
Equity Issuance
0.04
0.03
0.02
0.01
0
-0.01
-0.02
-0.03
-0.04
-0.05
-0.06
0.00000
0.00020
0.00040
0.00060
0.00080
0.00100
Gap Window

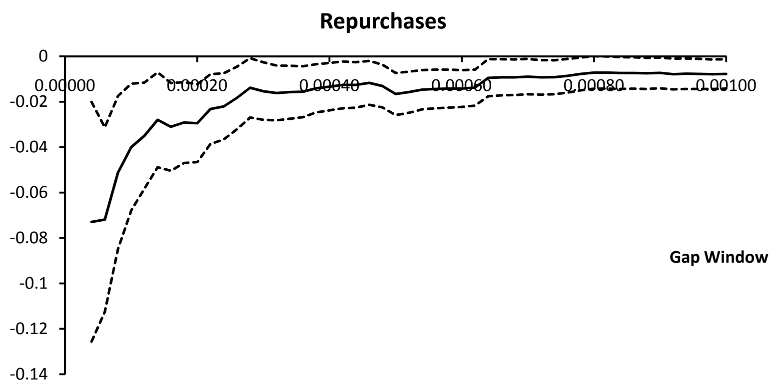
Repurchases
0
-0.02
-0.04
-0.06
-0.08
-0.1
-0.12
-0.14
0.00000
0.00020
0.00040
0.00060
0.00080
0.00100
Gap Window

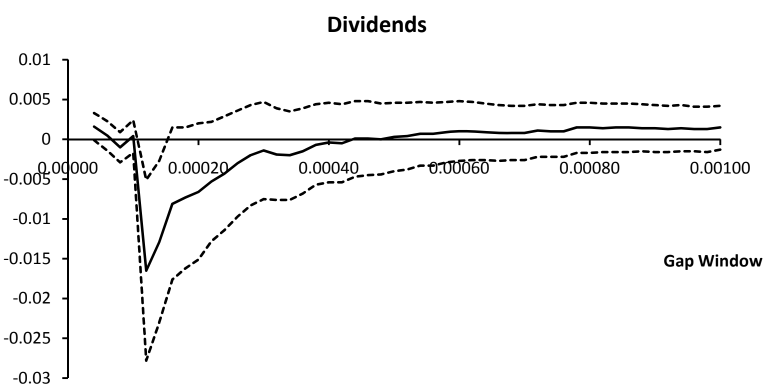
Dividends
0.01
0.005
0
-0.005
-0.01
-0.015
-0.02
-0.025
-0.03
0.00000
0.00020
0.00040
0.00060
0.00080
0.00100
Gap Window

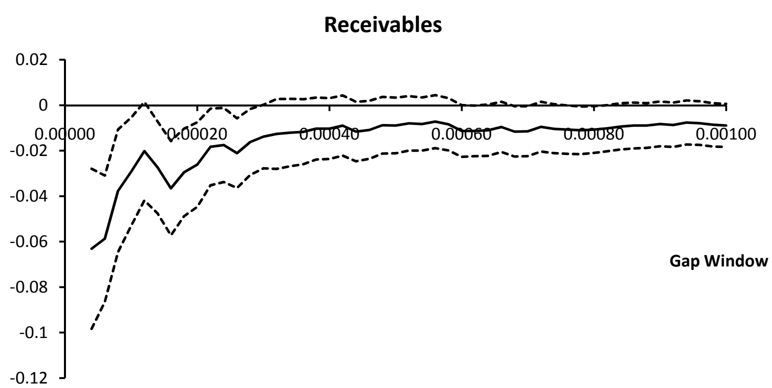
Receivables
0.02
0
-0.02
-0.04
-0.06
-0.08
-0.1
-0.12
0.00000
0.00020
0.00040
0.00060
0.00080
0.00100
Gap Window