

NBER WORKING PAPER SERIES

CORPORATE DEBT MATURITY AND THE REAL EFFECTS OF THE 2007 CREDIT
CRISIS

Heitor Almeida
Murillo Campello
Bruno Laranjeira
Scott Weisbenner

Working Paper 14990
<http://www.nber.org/papers/w14990>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2009

We thank Jaehoon Lee for excellent research assistance. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2009 by Heitor Almeida, Murillo Campello, Bruno Laranjeira, and Scott Weisbenner. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Corporate Debt Maturity and the Real Effects of the 2007 Credit Crisis
Heitor Almeida, Murillo Campello, Bruno Laranjeira, and Scott Weisbenner
NBER Working Paper No. 14990
May 2009
JEL No. E22,E32,G31,G32

ABSTRACT

We use the 2007 credit crisis to assess the effect of financial contracting on real corporate behavior. We identify heterogeneity in financial contracting at the onset of the crisis by exploring ex-ante variation in long-term debt maturity. Our empirical methodology uses an experiment-like design in which we control for observed and unobserved firm heterogeneity via a differences-in-differences matching estimator. We study whether firms with large portions of their long-term debt maturing right at the time of the crisis observe more pronounced outcomes than otherwise similar firms that need not refinance their debt during the crisis. Firms whose long-term debt was largely maturing right after the third quarter of 2007 reduced investment by 2.5% more (on a quarterly basis) than otherwise similar firms whose debt was scheduled to mature well after 2008. This relative decline in investment is statistically significant and economically large, representing approximately one-third of pre-crisis investment levels. A number of falsification and placebo tests confirm our inferences about the effect of credit supply shocks on corporate policies. For example, in the absence of a credit shock ("normal times"), the maturity composition of long-term debt has no effect on investment outcomes. Likewise, maturity composition has no impact on investment when long-term debt is not a major source of funding for the firm.

Heitor Almeida
University of Illinois at Urbana-Champaign
515 East Gregory Drive, 4037 BIF
Champaign, IL, 61820
and NBER
halmeida@illinois.edu

Murillo Campello
University of Illinois at Urbana-Champaign
4039 BIF
515 East Gregory Drive, MC- 520
Champaign, IL 61820
and NBER
campello@illinois.edu

Bruno Laranjeira
University of Illinois at Urbana-Champaign
1206 South Sixth St.
Champaign IL 61820
laranjei@illinois.edu

Scott Weisbenner
University of Illinois at Urbana-Champaign
Department of Finance
340 Wohlers Hall, MC-706
1206 South Sixth Street
Champaign, IL 61820
and NBER
weisbenn@illinois.edu

1 Introduction

Do financial contracts have real implications? How do firms respond to shifts in the supply of credit? The endogeneity of financing and investment decisions makes it difficult to answer these questions. To complicate matters, credit supply shocks often confound financial and economic factors that affect firm behavior. One common-place approach to studying the effect of credit shocks on firm behavior is to look at financing activity (e.g., loans or equity issues) that takes place over the credit cycle. Unfortunately, this approach is compromised by the fact that observed transactions may reflect both a shift in the supply of credit (e.g., lower supply of loans in a monetary contraction) as well as a shift in the demand for credit (firms demand less loans if economic conditions adversely affect their investment). Likewise, it is difficult to identify a causal link going from firm financing to firm investment during a credit contraction because economic considerations (the firm’s business fundamentals) may drive both ex-ante financial contracting and ex-post real outcomes.

We develop a novel strategy to gauge the effect of financial contracting on real corporate outcomes following a shift in the supply of credit. We do so by using the credit crisis (or “panic”) of August 2007. This event is unique among other credit shortage episodes in that it originated from problems arising from non-corporate assets: housing mortgages.¹ In other words, the 2007 episode provides for a shock to the supply of external financing that was not caused by the weakening of firm business fundamentals. We note, however, that simply exploring an event of this type is insufficient to identify a causal link between financial contracting and corporate outcomes. In particular, while general credit conditions may exacerbate the relation between variables such as financial leverage and investment, one cannot ascertain whether financial contracting *causes* firms to behave in a particular way. To establish that channel, one needs to identify a feature of financial contracts whose variation can be considered exogenous at the time of the credit shock.

We identify heterogeneity in financial contracting at the onset of the 2007 crisis by exploring ex-ante variation in *long-term debt maturity*. In a nutshell, we examine whether firms with large fractions of their long-term debt maturing at the time of the crisis are “forced” to adjust their behavior (e.g., by cutting capital expenditures) in ways that are more pronounced than otherwise similar firms that need not refinance their long-term obligations during the crisis. Long-term debt is typically publicly held, and difficult to renegotiate on short notice during crises. If financial contracting affects real corporate outcomes, we would expect a firm’s inability to refinance its obligations to

¹Gorton (2008) provides a detailed analysis of the various forces leading to a sharp reduction in liquidity (further aggravated by collateral calls) that affected financial institutions dealing with subprime-based derivatives starting in August 2007. The liquidity crisis had an immediate effect on instruments that are traded among financial institutions as well as those used to fund credit facilities used by corporations.

affect the firm’s post-crisis policies. To the extent that these effects are large, they constitute true costs of financial distress and may provide evidence that the terms of financial contracting — in this case, contract maturity — can cause real corporate outcomes.

It is important that we discuss how our focus on long-term debt maturity works as an identification tool. The literature on the determinants of debt structure has shown that the choice between *short- versus long-term* debt is correlated with firm characteristics such as size, profitability, and credit ratings (see, e.g., Barclay and Smith (1995) and Guedes and Opler (1996)). As such, in general, the determination of debt maturity creates difficulties for the identification of causal effects of financial contracting on real outcomes. For example, as shown by the literature, small firms are more likely to use short-term debt. Since small firms are also more likely be financially constrained (e.g., Rauh (2006)), it may not be surprising to find that firms that borrow on the short end are likely to suffer the most in the event of a credit contraction. In other words, the same firm characteristics that predict short-term borrowing may also predict how a firm will fare during a crisis. Perhaps because of these difficulties, previous literature has not examined the real consequences of maturity structure.

Rather than contrasting short- and long-term debt, we look at the *proportion of long-term debt* that matures right after August 2007 to assess how firms are affected by credit contractions. Because multiple, cumulative decisions affecting the maturity of a firm’s long-term debt were made several years prior the 2007 crisis, whether the firm was scheduled to refinance a large fraction of its long-term debt *right around the crisis* is plausibly exogenous to the firm’s performance following the crisis. Of course, while the firm could in principle renegotiate its long-term debt right at the time of the crisis, in practice, it is rather difficult for a firm to instantaneously readjust the profile of its long-term debt, especially when the firm anticipates liquidity problems (see Bolton and Scharfstein (1996)).²

While we argue that cross-firm variation in the proportion of long-term debt that comes due right after August of 2007 is exogenous to firm outcomes over the crisis, one might wonder if other sources of firm heterogeneity could underlie the relations we might observe. To tackle this concern, we use a differences-in-differences matching estimator approach that incorporates observable firm characteristics and accounts for unobservable firm effects. The goal of our empirical methodology is to replicate an experiment-like setting in which firm financial status can be seen as a “treatment” that is plausibly exogenous to corporate outcomes. To minimize concerns about selection, we match firms that we expect to be more susceptible to the negative effects of financial distress (i.e., those firms that happened to have a non-trivial fraction of their long-term debt coming due when the crisis hits) with “control” firms that we do not expect to be susceptible to distress (i.e., those firms that

²We empirically verify that firms do not change the profile of their long-term debt maturity in the period leading up to the 2007 crisis.

happened to have most of their long-term debt coming due many years after the shock). We match these two groups of firms on the basis of their asset size, industry classification, credit ratings, Q , long-term leverage ratio, cash flows, and cash holdings. This matching is meant to assure that we are comparing otherwise similar firms, with the only salient difference between the two groups being the composition of their long-term debt maturity. The tests we perform further account for time-invariant firm heterogeneity by comparing within-firm changes in the outcome variables of interest from the period that precedes the credit shock to the period that follows the shock. We perform these tests using the Abadie and Imbens (2002) matching estimator.

Importantly, we consider a number of alternatives to our basic empirical design. These alternative tests provide checks for the logic of our empirical design and further minimize concerns about “hard-wiring” in our results. For example, we perform a battery of falsification tests that replicate our matching estimator procedure in non-crisis periods. In principle, a firm whose debt matures at a time in which credit is easily available should not display a distressed-type behavior that can be linked to long-term debt maturity. It is only the interaction of firm debt maturity and a credit crisis that should affect investment. In addition, we re-define our treatment and control groups based on the degree to which long-term debt is an important component of firm financing. According to the logic of our strategy, for those firms for which long-term debt is only a small fraction of firm financing, we should not see a strong link between investment spending and the fact that some long-term debt is maturing during the crisis. Finally, to further assure that the assignment of firms into treatment and control groups is exogenous to the post-2007 crisis outcomes, we also perform tests in which we measure maturity structure several years prior to the credit crisis. Using a pre-determined fraction of long-term debt maturity as our sorting variable allows us to rule out more subtle unobserved heterogeneity stories, such as “smart CEOs” anticipating the August 2007 panic and refinancing (prior to the crisis) the part of their firms’ long-term debt that is scheduled to mature in 2008.

Our findings are as follows. We first document the existence of pronounced cross-sectional variation in the maturity structure of long-term debt at the onset of the 2007 credit crisis — variation that stems from contracting decisions made *several years prior* to that credit shock. The cross-sectional variation in long-term debt maturity is persistent across time, and we find no evidence that it changed in the years leading up to the 2007 crisis. These results are interesting in their own right and suggest that future researchers may use long-term debt maturity structure to gauge a plausibly unconfounded source of heterogeneity in firm financial status. Indeed, after matching on numerous characteristics, we are able to isolate firms with a large fraction of long-term debt maturing right after the crisis (the treated firms) that are virtually identical to other firms whose debt happens

to mature in later years (the control group). These groups of firms are identical across *all aspects* we consider other than maturity structure.³ For example, the two groups of firms display similar investment rates in the three quarters immediately leading up to the crisis (7.8% of capital on a quarterly basis for the treatment group, and 7.3% for the control group).

We then show that a firm's debt maturity structure has real consequences for post-crisis outcomes.⁴ For firms in the treatment group, quarterly investment rates decrease to 5.7% of capital on average — a fall of 2.1%. In contrast, firms in the control group hardly change their investment — quarterly investment over capital actually increases by 0.1% of capital. The estimate of the differences-in-differences in investment behavior increases from 2.2% to 2.5% in our baseline experiment after employing the Abadie and Imbens (2002) matching estimator. This drop in investment is economically substantive as it represents a decline of approximately one-third of pre-crisis investment levels. Confirming the logic of our strategy, the relation between maturity structure and investment disappears when we use firms with insignificant amounts of long-term debt in the experiment. On the flip side, that relation strengthens when we focus on firms for which long-term debt is a more important source of financing (in this case, the drop in investment is 3.4%).

We also find that the effect of maturity structure on investment is robust to variations in the definitions of treatment and control groups. Moreover, it holds *only* for the 2007 crisis. To show this latter point, we replicate our base tests for all years between 2000 and 2006. More precisely, over that window, we compare the changes in investment for firms with a lot of long-term debt maturing within one year with firms for which most of the outstanding long-term debt matures well over one year. We find that maturity structure is *uncorrelated* with changes in investment for these placebo periods. This finding suggests that the negative effect of debt maturity on investment that we uncover is indeed due to firms' inability to refinance the maturing portion of their long-term debt in the one period in recent years when there was a pronounced credit squeeze.

Adding to our robustness checks, we find results that are similar to those coming from our baseline experiment when we use 2005 maturity structure information to predict which firms are likely to have a large fraction of their long-term debt maturing in 2008 (i.e., firms for which a large fraction of long-term debt was scheduled to mature in three years). These firms reduce investment in

³To be precise, there are no statistically significant differences between treatment and control firms across the dimensions of size, industry, credit ratings, Q , long-term leverage ratio, cash flows, and cash holdings. It is worth emphasizing that using the matching estimator per se *does not ensure* the matching accuracy we report. The quality of our matching stems from the choice of a likely unconfounded source of firm heterogeneity (maturity structure).

⁴Anticipating the details of the baseline experiment, the pre-crisis period is defined as the first three quarters of 2007, and the post-crisis period is defined as the first three quarters of 2008 to control for seasonality effects. The treatment group contains firms for which the fraction of long-term debt maturing within one year is greater than 20%, while the control group contains firms for which that fraction is lower than 20%. The maturity variables are calculated at the end of 2007, and the experiment only includes firms whose long-term debt is greater than 5% of assets.

the aftermath of the crisis relative to comparable firms that are predicted (based on 2005 data) to have most of their debt maturing well beyond 2008.⁵ This pre-determined maturity test helps us to ascertain that the assignment of firms into treatment and control groups is exogenous to the 2007 crisis outcomes. In particular, it leaves little room for a story in which some firms anticipated the credit panic and selectively refinanced their long-term debt before it happened.

Naturally, the large effect of maturing debt on investment in 2008 raises the question of whether firms adjusted along other margins to accommodate the joint effects of the credit crisis and the need to repay a lot of debt in the short run. In particular, firms may have adjusted other real and financial policies, such as drawing down cash balances, reducing inventory stocks, repurchasing less shares, and cutting dividends. To provide some evidence on this point, in the last part of the paper, we perform a simple, “back-of-the-envelope” analysis of how the treated firms responded to the crisis. Our calculations suggest that the firms that were burdened with large amounts of maturing debt in 2008 tapped their “least costly” sources of funds. Notably, consistent with Almeida, Campello, and Weisbach (2004), the brunt of the shock to external funding was primarily absorbed by firms’ cash balances. Reductions in inventory were also pronounced across treated firms (consistent with Fazzari and Petersen (1993)). Interestingly, these firms on average seemed to be more willing to cut investment than to cut dividends, at least during the first three quarters of 2008.⁶

There are only a handful of empirical papers looking at the dispersion of corporate debt maturity. These papers consider issues other than the effect of supply shocks on real corporate policies. Examples are Barclay and Smith (1995), Stohs and Mauer (1996), and Guedes and Opler (1996). Barclay and Smith, for example, report that firms that have few growth options and are large have more long-term debt in their capital structures. Stohs and Mauer report that asset maturity is positively related to debt maturity. Finally, Guedes and Opler show that large firms with investment-grade credit ratings typically borrow on the short end and on the long end of the maturity spectrum, while firms with speculative-grade credit ratings borrow in the middle of the spectrum. While these papers look at determinants of debt maturity, they do not examine the real consequences of maturity structure.

Theory has also looked at the determinants of maturity structure (e.g., Diamond (1991, 1993, 2004) and Flannery (1986)).⁷ The theoretical literature often suggests that both low- and high-credit quality firms are likely to borrow short-term, but for different reasons. High-credit quality firms borrow short-term to signal that they are not concerned with the possibility of a liquidity shock that

⁵The differences-in-differences in investment across the two groups of firms is -1.4% (significant at the 5% level).

⁶Overall, firms did start to cut dividends at a record pace in 2009, as the economy worsened and corporate survival became a concern. However, these dividends were not reduced or eliminated immediately after the August 2007 credit crisis hit, with investment seeming to have been cut first in the early stages of the financial crisis.

⁷See also Greenwood, Hanson, and Stein (2008), and Brunnermeier and Yogo (2009).

might trigger refinancing. Low-credit quality firms use short-term bank financing because they have no alternative options other than debt financing with restrictive covenants and frequent renegotiations. The existing theory highlights the identification problem that we tackle in this paper. For instance, firms that use short-term bank financing are naturally more likely to be affected by a credit supply shock. As a result, one cannot measure the effect of maturity structure on real outcomes simply by relating the amount of short-term debt to post-crisis outcomes.

Finally, our paper is related to recent studies on the effects of credit supply shocks. Chava and Purnanandam (2008) examine the effects of the 1998 Brazil-Russia-LTCM crisis on corporate valuation. The authors find a larger valuation impact of that crisis upon bank-dependent firms whose main banks had greater exposure to Russia. Lemmon and Roberts (2008) examine the effects of a contraction in the supply of risky credit (junk bonds) caused by changes in regulation and the collapse of Drexel Burnham Lambert. Their evidence suggests that risky firms' leverage remained constant while their investment declined as a result of changes in the junk-bond market landscape. Our study differs from these papers in that our strategy dispenses with the need to focus on (more marginal) bank-dependent or risky firms to assess the impact of credit supply shocks. Similarly to our paper, Duchin, Ozbas, and Sensoy (2009) focus on the impact of the current credit crisis on corporate investment. However, their identification strategy uses firms' cash and debt positions to sort firms into treatment and control groups. As discussed above, long-term debt maturity structure provides for a much cleaner identification strategy than the more basic dimensions of financial policy that those authors adopt.

The remainder of our paper is organized as follows. We discuss our empirical strategy in Section 2. Our baseline result that the financial contracting (debt maturity structure) affects real corporate outcomes (investment) is presented in Section 3. In Section 4, we conduct a number of additional tests designed to check the robustness of our results. Section 5 concludes the paper.

2 Empirical Design

We start this section by describing our basic experimental design and the matching estimator methodology we employ in the paper. We then describe the data used in our tests.

2.1 The “Experiment”

Our basic insight is that of exploring variation in long-term debt maturity at the onset of the 2007 financial crisis as a way to identify the effect of credit supply shocks on corporate policies. Of course, the relevant question is how would the composition of long-term debt maturity affect real corporate

policies. In a Modigliani-Miller world of frictionless capital markets, debt maturity is irrelevant because firms can always refinance and recontract their way around the potential effects of a balloon debt payment. What is special about credit crises is that financial markets are arguably less than frictionless during those times. The 2007 crisis, in particular, shut down traditional modes of corporate financing, such as commercial paper, bond placements, bank loans, and secondary equity issuance. In such an environment, soon-to-mature debt can effectively reduce corporate investment, as firms find it difficult to substitute across alternative funding sources, while at the same time trying to avoid defaulting on their debt payments. As a result, firms that are “unlucky” to have large chunks of debt maturing around the 2007 crisis may be expected to face tighter financing constraints than firms that do not have to finance balloon debt payments during that same period.

To gauge the magnitude of the 2007 credit squeeze, one can consider the time series of the 3-month LIBOR and commercial paper spreads over treasuries. These series are plotted in Figure 1.⁸ There is a sharp, large shock to all of these spreads around August 2007. Spreads go up from levels lower than 0.5% between 2001 and the summer of 2007, to levels between 1% and 2% following August 2007. In particular, in July 2007 the average 3-month LIBOR spread was 0.53%. This spread jumped to 1.31% in the month of August, and stayed above 1% in the subsequent months. This event, which stemmed from disruptions in the market for mortgage-back securities, represented a noticeable contraction in the supply of credit following a long period — multiple consecutive years — of easy credit. And importantly, it generated fears in the financial markets that additional hikes in the cost of funds were about to come (see, e.g., Gorton (2008)). It is fair to argue that panic of August 2007 provides us with a unique opportunity to identify the effects of supply contractions on corporate policies.

FIGURE 1 ABOUT HERE

Our identification strategy requires two conditions to be met. First, and most simply, there has to be enough variation in debt maturity to allow for comparisons across firms. In particular, there must exist a significant group of firms that have a spike in maturity in the period right after the crisis. Naturally, one could expect firms to have well-diversified maturity structures, so that they are never forced to repay or refinance significant amounts of debt in any particular year. If that was the case, it would be harder for us to implement our proposed strategy.

Figure 2 depicts the distribution of debt maturities for the sample of firms that we use in this paper (the data are described in detail in Section 2.3), calculated at the onset of the financial crisis (fiscal-year end 2007). For each firm, we have information on the amount of long-term debt that

⁸These data are available from <http://www.federalreserve.gov/datadownload/>.

matures in each of the following five years: 2008, 2009, 2010, 2011, and 2012.⁹ Figure 2 reports these amounts, as a fraction of total long-term debt. If maturity structure was well diversified, we would expect this distribution to have a large mass around a specific value (for example, if firms tend to regularly issue 10-year bonds we would expect to see a mass at the value of 10%). Clearly, there is significant variation in maturity structures. Consider, for example, the fraction of long-term debt that is due in one year (i.e., in 2008). Figure 2 suggests that there exists a significant number of firms whose long-term debt maturity concentrates in the year of 2008. By the same token, many firms do not have any significant amount of long-term debt maturing in 2008. Similar variation in maturities obtains for the other individual years. For example, there are many firms with maturity spikes occurring in 2012. These firms are similar to the ones with concentrated maturity in 2008, in that they, too, allow their debt maturity to concentrate in a particular year; however, their maturity is concentrated in a future year that lies far beyond the 2007 crisis.

FIGURE 2 ABOUT HERE

Two other features of the distribution of debt maturity measured at the end of 2007 are noteworthy (and useful for our “experimental” design). First, the distributions of long-term debt maturing in the individual years beyond 2008 (2009 through 2012) look fairly similar to the distribution of long-term debt maturing in 2008. This suggests that firms may not always try to renegotiate in advance and elongate maturities of debts that are soon to come due. Second, as depicted in Figure 3, the distributions of the long-term debt maturity of firms for years prior to 2007 are strikingly similar to that of 2007. In other words, there is no evidence of changes in long-term debt maturity structure in the years leading up to the 2007 crisis.

FIGURE 3 ABOUT HERE

The second condition that must be satisfied to validate our identification strategy is that the variation in debt maturity at the onset of the crisis needs to be exogenous to observed post-crisis outcomes. In particular, one might worry that the same variables that determine the pre-crisis distribution of long-term debt maturity are also related to post-crisis corporate investment. For example, suppose that firms that have high growth opportunities tend to issue debt of longer maturities. Then, it might not be surprising to observe that these firms invest more, relative to other firms that have shorter-maturity debt.

⁹We also know the amount of long-term debt that matures in more than five years (starting in 2013), though we do not have year-by-year information beyond five years.

Our empirical strategy addresses this issue in several complementary ways. First, we focus on variation in long-term debt maturity (only) rather than in short- versus long-term debt to sort firms into treatment and control groups. This choice ensures that we are not simply comparing low-quality firms that must issue short-term debt to firms that can — and do — issue long-term debt. Second, we use matching estimator techniques (described in greater detail in Section 2.2) that minimize concerns about selection. This enables us to account for the effect of observables by matching firms with long-term debt maturing right after the credit crisis to “control firms” of similar size, industry, credit ratings, Q , long-term leverage ratio, cash flows, and cash holdings, with the key difference between the treatment and control firms being when their long-term debt, which was contracted long ago, happens to mature. Third, we consider a number of variations to our baseline empirical design that allow us to further distinguish among alternative explanations for the results we obtain.

Among other things, we perform a battery of falsification tests that replicate our matching estimator procedure in non-crisis periods. To see the logic behind these tests, consider Figure 3. The 2006 distribution of long-term debt maturity shows a significant fraction of firms with a lot of long-term debt maturing within one year’s time (in this case, 2007). These firms are similar to firms at the end of 2007 that have a lot of long-term debt maturing within one year’s time (in this case, 2008) in that both sets of firms have fast-approaching spikes in the maturity of their long-term debt. However, the former group of firms *should not* be affected by the credit crisis because their debt matures *before* the crisis. Our empirical tests will directly compare these two groups of firms. Figure 3 also shows a group of firms that have, as of fiscal-year end 2005, a large fraction of their long-term debt maturing in 2008 (i.e., the third year after 2005). In one of our robustness tests, we will use this *pre-determined* long-term debt maturity distribution to sort firms into treatment and control groups. This test allows us to verify whether some managers were better able to anticipate the effects of the credit crisis and refinance their long-term debt before the credit crisis hit, we can thus check whether unobserved managerial quality could explain the post-crisis differences in investment behavior.

2.2 Matching Estimators

The use of matching estimators is a crucial feature of our test strategy and it is important to explain how we employ this technique. Recall, we want to test whether firms that need to refinance their long-term financial obligations at the time of a credit crisis alter decisions related to real-side variables. In particular, we want to determine whether re-financing constraints affect real firm outcomes. Our goal is to develop an identification strategy that is akin to a “random” experiment: the firm’s long-term debt maturity structure and developments in the financial markets coincide such that the firm

is in need of refinancing a large fraction of its debt in the midst of a sharp credit contraction. If one believed that a firm's debt maturity was truly randomly assigned across firms, then it would suffice to compare the ex-post outcomes of firms that had significant debt maturing around the time of the crisis with those whose debt happened to mature at a later date. Our analysis, however, needs to allow for the fact that we are not using a true laboratory experiment, but instead relying on observational data.

Short of running a pure-randomizing experiment with the firms' financing constraints, the econometric challenge is then to gauge firms' likely outcomes had they *not* been caught between a credit crisis and the need to refinance their debt. Since we are interested on the impact of firm financing on real outcomes, we need to carefully identify a group of firms that also face the credit crisis and are virtually similar to those whose debt matures during the crisis period *except* for the fact that their debt is not maturing in the crisis. That is, we need to pin down the counterfactual firm policy in 2008 — what would the firms with debt maturing in 2008 have done *if* their debt had matured at a later date.

Traditionally, researchers have dealt with this problem in the context of ordinary least squares (OLS) estimators, where the group of interest is differentiated from other observations via an indicator variable. Under this standard parametric approach, the impact of the variable of interest on observed outcomes is measured by the coefficient for the selected group's indicator variable. The regression specification is determined according to a set of theoretical priors about the endogenous variable. Such models are often simpler, linear representations of the theory. In corporate finance research, controls such as firm size, profitability, and leverage are customarily added to a specification to capture additional sources of firm heterogeneity. If left unmodeled, that sort of variation could jeopardize the OLS estimator as it could explain both a firm's selection into the group of interest and its observed outcomes.

While the traditional OLS approach is perfectly valid, there are a few concerns with its implementation. First, the inclusion of control variables in the specification does nothing to address the fact that the groups being compared may have different distributions of the control variables (for example, comparison groups may have vastly different size or profitability distributions). Unfortunately, OLS estimates will not alert the researcher that a poor distributional overlap might yield an ineffective control set. Second, and relatedly, the OLS approach will allow for extreme outliers in the estimation, outliers that can bias the estimates of interest substantially — OLS models are notoriously weak in dealing with outliers. Finally, the OLS approach may place undue importance on linear model parametrization in the estimation process. Depending on the application, one can improve the estimation of group differences by allowing for non-linear modeling of the outcomes of interest as well as by way of non-parametric methods.

The estimation strategy that we use in this paper is less parametric and more closely related to the notion of a randomized experiment. We use matching estimators in all of the tests performed in this paper. The idea behind this family of estimators is that of isolating *treated* observations (in our application, firms with debt maturing during the crisis) and then, from the population of non-treated observations, look for *control* observations that best “match” the treated ones in a number of dimensions (*covariates*). In this estimation framework, the set of counterfactuals are restricted to the matched controls. In other words, it is assumed that the treated group would have behaved as the control group actually did if the treated group had not been treated. The matches are carefully made so as to ensure that treated and control observations have identical distributions along the covariates chosen (dimensions such as firm size, profitability, long-term leverage, etc.). Inferences about the treatment of interest (re-financing constraints) are based on comparisons of the ex-post outcomes of treatment and control groups (outcomes such as investment spending).¹⁰ Although a number of matching estimators are available, we employ the Abadie and Imbens (2002) estimator.¹¹ Their non-parametric procedure most naturally fits our application.

The Abadie-Imbens estimator allows one to match a treated firm with a control firm, with matching being made with respect to both categorical and continuous variables. The estimator aims at producing “exact” matches on categorical variables. Naturally, the matches on continuous variables will not be exact (though they should be close). The procedure recognizes this difficulty and applies a “bias-correction” component to the estimates of interest.¹² In matching estimations, the specification used is less centered around the idea of representing a model that fully explains the endogenous variable. Instead, the focus is in ensuring that variables that might *both* influence the selection into treatment and observed outcomes are appropriately accounted for in the estimation. For example, the outcome that we are most interested in is investment spending.¹³ While there are numerous theories on the determinants of corporate investment, we only include in our estimations covariates for which one could make a reasonable case for simultaneity in the treatment–outcome relation. Among the list of categorical variables we include in our matching estimations are the firm’s industrial classification (2-digit SICs) and the rating of its public bonds (either speculative grade, investment grade, or unrated). Our non-categorical variables include the firm’s market-to-book ratio (or “*Q*”), cash flow, size, and the ratio of long-term debt to total assets. Although our original quasi-randomization

¹⁰In the treatment evaluation literature this difference is referred to as the average treatment effect for the treated, or ATT (see Imbens (2004) for a review).

¹¹We refer the reader to Abadie and Imbens (2002) for a detailed discussion of their matching estimator. Here we apply the bias-corrected, heteroskedasticity-consistent estimator implemented in Abadie et al. (2004).

¹²We will illustrate the effect of this bias correction in our paper’s central table (Table 3).

¹³In the last section of the paper, we look at other outcomes, such as cash savings, inventories, and dividend payout.

approach already makes it hard to tell a story under which the covariates we consider would still predict both treatment and outcomes, it is commonly accepted that those covariates capture a lot of otherwise unobserved firm heterogeneity.

Lastly, we note that we model the outcomes in our experiments in a differenced form, that is, we perform differences-in-differences estimations. To be precise, rather than only comparing the *levels* of investment of the treatment and control groups after the experiment, we compare the *changes* in investment across the groups. In other words, we compare the difference in behavior across the two groups of firms after the credit crisis with the difference in their behavior before the crisis. We do so because the investment levels of the treated and controls could be different prior to the event defining the experiment, and continue to be different after that event, in which case our inferences could be potentially biased by these uncontrolled differences (in actuality, this turns out *not* to be the case, adding further reassurance to our experiment-like design).

2.3 Data Description

To conduct our empirical tests, we use data from COMPUSTAT’s North America Fundamentals Annual, Fundamentals Quarterly, and Ratings files. Our sampling disregards observations from financial institutions (SICs 6000–6999), not-for-profit organizations and governmental enterprises (SICs greater than 8000), as well as ADRs. We require firms to provide valid information on their total assets, capital expenditures, property, plant and equipment, long-term debt maturing within one year, total long-term debt maturing in more than one year, cash holdings, and operating income.

Our data selection criteria and variable construction approach follows that of Almeida, Campello, and Weisbach (2004), who study the effect of financing constraints on the management of internal funds, and that of Frank and Goyal (2003), who look at external financing decisions. Similar to Almeida, Campello, and Weisbach (2004), we discard from the raw data those firm-years for which the value of total assets is less than \$10 million, and those displaying asset growth exceeding 100%. We further require that firm’s quarterly sales be positive and that the log of sales growth does not exceed 100%. We also eliminate firms for which liabilities such as notes payables, bank overdrafts, and loans payable to officers and stockholders are greater than 1% of total assets. In addition, we require firms to have long-term debt maturing beyond one year that represents more than 5 percent of total assets. This screen ensures that firms are capable of issuing debt with longer maturities. These debt-related restrictions help assure that the results in our paper do not come from comparisons between “low-quality” firms that need to rely on very short-term obligations with “high-quality” firms that can issue long-term debt.

We focus on firms that have 2007 fiscal year-end months in September, October, November, December, or January. The sample of firms with these fiscal year-end months corresponds to more than 80% of the universe of firms in fiscal year 2007. This restriction is due to the timing of the credit shock, which happened around August of 2007. For our benchmark specification, we want to avoid firms that filed their 2007 annual report before the crisis. These firms could have used the time period between filing the annual report and the credit crisis to rebalance their debt maturity. The variables that detail the amount of long-term debt maturing within one, two, three, four, and five years from the date of the report are only available in the annual COMPUSTAT file (they are not available on a quarterly basis). Thus, for a December fiscal-year-end firm, we cannot use the third quarter report to obtain a breakdown of timing of the debt maturity composition as of 9/30/2007, we instead use the firm’s 2007 annual report to obtain the debt maturity breakdown as of 12/31/2007. Our final (2007) sample consists of 1,067 individual firms.¹⁴

We use the Abadie and Imbens (2002) matching estimator (described in Section 2.2) to test our hypothesis. The outcome variable used in the matching estimator is the average change in investment during the first 3 fiscal quarters of the year relative to the first three fiscal quarters of the previous year (expressed on a quarterly basis). In our basic experiment, we are examining the change in the average quarterly investment over the first three quarters of 2008 relative to the first three quarters of 2007. Investment is defined as the ratio of quarterly capital expenditures (COMPUSTAT’s *capxy*) to the lag of quarterly property, plant and equipment (*ppentq*). As discussed earlier, we match firms based on Q , cash flow, size, cash holdings, and long-term leverage. Q is defined as the ratio of total assets plus market capitalization minus common equity minus deferred taxes and investment tax credit ($atq + prccq \times cshoq - ceqq - txditcq$) to total assets (atq). Cash flow is defined as the ratio of net income plus depreciation and amortization ($ibq + dpq$) to the lag of quarterly property, plant and equipment. Size is defined as the log of total assets. Cash holdings are defined as the ratio of cash and short-term investments (*cheq*) to total assets. Long-term leverage is defined as the ratio of total long-term debt ($ddl + dltt$) to total assets. For the matching procedure, we use averages in the first 3 quarters of 2007 for all of these variables as covariates.

We also match firms both on industry and credit ratings categories. Industry categories are given by firms’ two-digit SIC codes. Our credit ratings categories follow the index system used by S&P and are defined as: investment grade rating (COMPUSTAT’s *splticrm* from AAA to BBB–), spec-

¹⁴Complete data on accounting and financial variables is currently available from COMPUSTAT only up to the third quarter of 2008. Importantly, we note that the time window we currently explore naturally fits our proposed experiment. Among other considerations, it is difficult to ascertain whether data from the fourth quarter of 2007 belongs to the treatment window, and whether the realization of real-side variables (outcomes) in the fourth quarter of 2008 confound the credit shock with the effects of the fiscal and monetary policies put in place to address the original crisis.

ulative rating (*splticrm* from SD to BB+), and unrated (*splticrm* is missing). Matching treatment and control firms within the same industry *and* within the same debt ratings categories assures that differences in firms’ underlying business activity (e.g., product demand) and credit quality may not explain our results.

We construct treatment and control groups based on firms’ long-term debt maturity schedule. COMPUSTAT’s items *dd1*, *dd2*, *dd3*, *dd4*, and *dd5* represent, respectively, the dollar amount of long-term debt maturing during the first year after the annual report (long-term debt maturing in 2008 for firms with a December 2007 fiscal year-end), during the second year after the report (long-term debt maturing in 2009 for firms with a December 2007 fiscal year-end), during the third year after the report, and so on. We note that COMPUSTAT’s *dltt* represents the dollar amount of long-term debt that matures in more than one year. Accordingly, a firm’s total long-term debt is calculated as $dd1+dltt$. In our benchmark specification, the treatment variable is defined by the ratio of long-term debt maturing within one year (*dd1*) to total long-term debt ($dd1+dltt$). Firms for which this ratio is greater than 20% (alternatively, less than or equal to 20%) are candidates for treatment (alternatively, control) assignments. In Section 4.4, we experiment with alternative definitions of treatment and control groups.

3 Results

We start by providing summary statistics on our samples of treated, control, and non-treated firms. Our initial goal is to show that our procedure does a good job of matching treatment to control firms along observable dimensions. We then present our baseline empirical results.

3.1 Summary Statistics

Table 1 reports the (pre-crisis) median values for our matching variables (the covariates), and also for investment normalized by capital. Our data selection procedure yields a group of 86 treated firms. Panel A of Table 1 compares the characteristics of the treated sample with those of the remaining 981 firms that do not make it into the treated group. Firms in the treatment group have similar pre-crisis investment levels to non-treated firms. However, the median values for the other variables differ significantly across these two groups of firms. The treated firms have higher median Q , cash flows, and cash holdings. Treated firms are also smaller and have a lower median leverage ratios. As discussed above, these sample differences are expected, given that we are relying on observational data rather than running a true laboratory experiment. The goal of matching estimator techniques is to precisely control for these distributional differences that can affect both the selection into the

treatment group and the post-crisis outcomes.

TABLE 1 ABOUT HERE

Panel B compares median values for treated and control firms. The Abadie-Imbens estimator identifies a match firm for each firm in the treatment group (thus, we have 86 firms in both the treated and control groups). Notably, there are *no* statistical differences in the median values of the covariates we consider across treated and control firms.¹⁵

Table 2 compares the entire distributions — rather than just the median — of the various matching covariates across the various groups. The results mirror those reported in Table 1. Panel A shows that treated firms differ significantly from non-treated firms. In particular, the Kolmogorov-Smirnov tests of differences in distribution produce statistically significant differences in distribution for all of the matching covariates (again, there is no difference in the pre-crisis level of investment across the two groups of firms). As in Table 1, these differences vanish when we compare the treated firms to the sample of closely-matched control firms — Panel B of Table 2 highlights that there are no statistical differences in the *distributions* of the various matching covariates across the treated and control firms. This evidence supports the assertion that the matching estimator moves our “experiment” much closer to a true randomization in which treatment and control groups differ only with respect to when their long-term debt happens to mature.

TABLE 2 ABOUT HERE

3.2 The 2007 Credit Crisis

We now examine the investment behavior of our treated and control firms around the 2007 credit crisis. Panel A of Table 3 shows that prior to the crisis, both the treated and control firms were investing at similar levels. The average investment-to-capital ratio in the three first quarters of 2007 (the pre-crisis period) is 7.8% for the treated firms, and 7.3% for the control firms. The difference is statistically insignificant, as indicated in the last row of the table. The fact that both groups of firms have statistically indistinguishable investment levels in the pre-crisis period is consistent with the statistics reported in Tables 1 and 2. It provides additional evidence that our matching procedure produces groups of firms that are virtually identical in terms of covariates that could be associated with firm investment behavior.

TABLE 3 ABOUT HERE

¹⁵We use the Continuity Corrected Pearson χ^2 statistic to test for differences in the medians of the covariates across the two groups.

Panel A of Table 3 also shows the investment levels in the first three quarters of 2008 (the post-crisis period). Notice that the investment policies of the treated and control firms are no longer similar. For firms in the treatment group, the average investment drops to 5.7% of capital (a fall of 2.1%). In contrast, for firms in the control group, investment remains virtually unchanged (increasing by a statistically insignificant 0.1%). These figures suggest that investment decreased by 2.2% more for firms that happened to have a lot of long-term debt maturing right after the credit crisis hit, relative to otherwise similar firms whose long-term debt did not come due so soon.

Panel A also reports the differential change in investment that is produced by the Abadie-Imbens matching estimator (the average treatment effect, or ATT). The ATT difference is equal to 2.5%. That estimate would equal 2.2% (the simple average difference effect) if it were not for the “bias-correction” that is embedded in the estimator that helps dealing with the issue of matching on continuous variables (see Section 2.2). The fact that the estimated difference increases from 2.2% to 2.5% after the bias-correction suggests that the inability to match firms perfectly in terms of continuous variables creates a small downward bias in the observed cross-group differences in investment growth.

The results in Panel A support the notion that the financing constraints that were generated by firms’ debt maturity structures created frictions in the corporate sector following the 2007 crisis, frictions that manifested themselves in the form of lower corporate investment rates. Our estimates imply that investment for the treated firms during the first three quarters of 2008 fell by about one-third of their pre-crisis investment levels. Given the similarity between firms in the treatment and control groups, the evidence presented is already highly suggestive of a true effect of debt maturity on investment. In order to further strengthen the interpretation of the results, however, we replicate exactly the same “experiment” that we run for the crisis period around a *placebo period* one-year earlier. That is, we use 2006 maturity information to sort firms into treatment and non-treated groups and 2006 covariates to produce a matched group of firms. Then, we examine firms’ investment behavior during the first three quarters of 2007. This placebo test can help us rule out alternative explanations for the results reported in Panel A. For example, there could be unobservable characteristics that predict both decreases in investment and a short-maturity profile for long-term debt, and which are not captured by the matching estimator procedure described in Section 2.2. However, if this is the case, then maturity structure and investment should be correlated in 2006 as well, and not just in the crisis period.

The results from this placebo test are shown in Panel B of Table 3. As in Panel A, treated and control firms have virtually identical investment behavior in 2006. Firms with more than 20% of their long-term debt maturing in 2007 (the treatment group) have investment representing 7.3% of

capital, on average, in the first three quarters of 2006, while their control group counterparts invest 7.2%. Strikingly, there is *no difference* in investment behavior across these two groups of firms in 2007, despite the different maturity profiles of long-term debt (both invest 6.9% on average in the first three quarters of 2007). The average treatment effect (ATT) in this case is virtually zero, and statistically insignificant. The differences between the ATTs in Panels A and B provide further evidence that maturing long-term debt had a *causal* negative effect on investment in the aftermath of the 2007 credit crisis.

4 Extensions and Robustness Tests

In this section, we test additional implications of our basic argument, provide evidence that the benchmark results are robust to variations in the empirical specification we use, and show that the 2007 crisis results (reported in Table 3) do not obtain in non-crisis periods. We also provide a “back-of-the-envelope” calculation that shows how firms with balloon debt payments in 2008 — the treated firms — responded to the credit crisis along other dimensions besides investment policy.

4.1 Evidence from Non-Crisis Periods

Our identification strategy relies on the assumption that firms with maturing long-term debt find it difficult to refinance their obligations by tapping other financing sources (e.g., long-term debt or equity). The 2007 credit crisis provides us with an ideal setting in which this assumption is likely to hold. By the same token, the assumption is unlikely to hold in periods of easier credit. If our identification strategy is correct, we would then expect *not to find* the same effects of maturity structure on investment during non-crisis periods. Panel B of Table 3 verifies whether this is true for the year of 2006 (one year before the August 2007 credit event). Here, we extend these placebo tests for years prior to 2006, reporting results on a year-by-year basis as well as pooled over the pre-crisis 2000-2006 period.¹⁶ To replicate our testing strategy for years prior to 2006, we sort firms into treatment and non-treatment groups considering maturity structures measured in 2000 to 2005, using previous years’ covariates to construct control groups. We then examine the differential change in investment for treated and control firms. We perform this test for each individual non-crisis year, using the same sampling criteria and definitions of treatment and control groups that we used for the credit crisis period.

The results are depicted in Table 4, which also reports the results for 2006 and 2007 for quick

¹⁶We start in 2000 because it is difficult to classify the late 1990’s as a non-crisis period in light of episodes such as the LTCM debacle and the Asian crisis. In addition, information on debt maturity in COMPUSTAT thins out before 2000.

reference. The estimated difference in investment changes across treatment and control groups is economically small and statistically insignificant for *all years* between 2000 and 2006. The pooled estimate between 2000 and 2006 is 0.3%, which is insignificantly different from zero. These results strongly support our assertion that debt maturity affects investment through a (re-)financing constraint channel in the aftermath of the financial crisis.

TABLE 4 ABOUT HERE

4.2 Changing Long-Term Leverage Cutoffs

Long-term debt maturity should matter only for firms that have significant amounts of long-term debt in their capital structures. Accordingly, in our benchmark specification we considered only those firms for which the ratio of long-term debt maturing in more than one year to total assets was higher than 5%. This screen also ensures that firms are capable of issuing debt with longer maturities.

In Table 5, we experiment with different inclusion rules that are designed to check the logic behind our strategy. Provided that the sample remains large enough, increasing the cutoff for the fraction of long-term debt in firms' capital structures should result in larger post-crisis effects of maturity on investment. By the same token, including firms that do not have significant long-term debt should weaken our estimated effects.

Table 5 shows evidence that is consistent with these hypotheses. In the first column, we report the changes in investment for treatment and control groups (from 2007 to 2008) that are constructed using the entire sample of firms; that is, also including those firms whose long-term debt is less than or equal to 5% of assets. Consistent with expectations, the estimated differences between treatment and control groups *disappears* after this change. The simple differences-in-differences estimate is 0.0% while the ATT is now positive at 0.2% (both are statistically insignificant).¹⁷ This contrasts with our benchmark result, which is reported in the second column of the table. In the third column, we perform an alternative experiment that only includes firms whose long-term debt is greater than 10% of assets. Now, the differential change in investment across treatment and control groups *increases* to 3.4% of capital. This evidence is consistent with the hypothesis that treated firms found it difficult to refinance their maturing long-term debt in the post-crisis period, and thus were forced to substantially cut their investment spending.

TABLE 5 ABOUT HERE

¹⁷Notice that in this case investment decreases by 1.1% for both samples (treated and control). In particular, the fall in investment for the control group is consistent with the notion that the credit crisis had an overall negative effect on investment, possibly due to a combination of credit supply and investment demand effects. Duchin, Ozbas, and Sensoy (2009) also report a reduction in 2008 investment.

4.3 Pre-Determined Maturity Tests

Our baseline experiment uses maturity variables measured at the end of 2007, just a few months following the August credit panic. As explained in Section 2.3, we made this choice to make sure that we capture the extent to which firms are constrained by debt maturity in the aftermath of the crisis. This requirement should increase the power of our tests. However, it raises the concern that measured variation in maturity reflects the anticipated effects of the crisis. A particularly problematic alternative explanation is the following. Suppose that higher quality managers were more likely to anticipate the effects of the credit crisis either in 2006 or early in 2007. Then, it is possible that unobservable managerial quality could explain both longer maturity profile and superior firm performance in the aftermath of the crisis. In other words, such refinancing in anticipation of the financial crisis by “smart CEOs,” would leave only the “dumb CEOs” with long-term debt maturing in 2008, and these “dumb CEOs” may be forced to cut investment for non-maturity-related reasons after the credit crisis hits. The placebo tests of Section 4.1 do not address this challenge, because this is a crisis-specific story.

A simple way to ensure that the anticipation of the crisis by “smart CEOs” does not drive our results is to use maturity variables measured in years prior to the end of 2006 and the end of 2007. For example, we can examine firms’ maturity profiles at the end of 2005 — over a full year before the crisis — and identify a group of firms that had a large fraction of their long-term debt maturing in three years (i.e., in 2008). Since it is unlikely that even the best manager could have anticipated the 2007 credit crisis back in 2005, such a modification of our basic specification can address the unobservable managerial quality story. For robustness, we also experiment with using a maturity profile measured an additional two years earlier, fiscal-year end 2003, which is the earliest we can go back given that COMPUSTAT only provides the maturity structure of long-term debt on a yearly basis for the next five years. Naturally, as we go back to earlier years to measure maturity, the effect of this pre-determined maturity structure on 2008 investment should decrease in magnitude (since the maturity information might become stale as we go back further in time). For both earlier snapshots (2003 and 2005), the treatment group again includes firms that have more than 20% of their long-term debt at the time maturing in 2008. Indeed, other than using pre-determined maturity profiles to assign treatment and non-treatment groups, all other components of the experiment remain unchanged. Accordingly, the outcome variables are defined identically to those in Panel A of Table 3, that is, changes in investment between the first three quarters of 2008 and the first three quarters of 2007.

The results (untabulated) suggest that the pre-determined maturity profiles also help predict changes in investment around the credit crisis. As expected, the effects of maturity on investment (−1.4% when using the 2005 maturity structure and −0.6% when using the 2003 maturity structure)

are somewhat smaller than those estimated in Table 3, nonetheless, they are still economically substantive.¹⁸ These results suggest that the managerial quality hypothesis cannot explain the relation between debt maturity and investment that we report in Table 3.

4.4 Different Specifications for the Matching Estimator

We have also experimented with several variations in our procedure to construct treatment and control groups, and with the set of matching covariates. To illustrate the robustness of our results, we report two of these exercises in this section.

Our benchmark specification defines the treatment group as all firms for which the ratio of long-term debt maturing within one year to total long-term debt is greater than 20%. The non-treated group contains all the other firms that satisfy the sampling restrictions (in particular, a minimum level of long-term debt over assets as discussed in Section 4.2). We also considered a control group that includes only firms that have more than 20% of their long-term debt maturing in exactly *five years* (that is, in 2012). These firms are similar to those in the treatment group in that they allow their maturity structures to be poorly diversified across maturities. However, they happen to have concentrated their maturity in a time period that lies far in the future.¹⁹ The estimated difference in investment changes (the matching estimator ATT) remains negative, equal to -1.6% , and statistically significant (standard error of 0.9) after this change in definition.

Our benchmark specification includes matching covariates that are potentially associated with both investment and maturity structure. We have also experimented with including the 2007 investment level among the set of matching covariates to ensure that we are comparing firms that were at the same starting point of investment before the crisis. The matching estimator's average treatment effect is virtually unchanged after this modification in the set of covariates; point estimate of -2.3% , with a standard error of 0.9.

4.5 How Did the Treated Firms Respond to the Credit Crisis?

The evidence so far suggests that firms with large amounts of debt maturing in 2008 (the treatment group) were forced to decrease investment in order to be able to repay their maturing debt. However, investment is not the only policy variable that these firms could have adjusted in the aftermath of the crisis. Here, we examine post-crisis changes in other policy variables that the treated firms could have used to absorb the effect of the credit squeeze. Even if it was difficult or impossible for firms

¹⁸The difference in investment using the end-of-year 2005 debt maturity is significant at the 5 percent level. The difference in investment using the end-of-year 2003 debt maturity is statistically insignificant (t -statistic equal to 1.0).

¹⁹We choose five years because this is the maximum one-year fraction that is available in COMPUSTAT. As mentioned earlier, for debt maturing in more than five years, we have no information on the specific year of maturity.

to respond to the crisis by issuing additional external finance, they can potentially make up for the debt payment by adjusting other variables, such as drawing down cash reserves, reducing stocks of inventory, repurchasing less shares, and/or cutting dividends. If the treated firms found it necessary to cut investment (which is a costly measure), we would also expect them to adjust, for example, the amount of share repurchase activities that they undertake in the aftermath of the crisis.²⁰ In addition, we would expect firms to draw down on their cash balances and reduce inventories. The literature suggests that cash balances are held in part to hedge against negative shocks such as the 2007 crisis (see, e.g., Almeida, Campello, and Weisbach (2004)). In addition, there is evidence that firms use inventories to smooth out the effects of fluctuations in the availability of internal funds (Fazzari and Petersen (1993)). Finally, firms can cut dividends to help finance debt repayments and mitigate investment reductions.

To provide some evidence on these additional policies, we provide a crude, “back-of-the-envelope” analysis of how the treated firms in our experiment responded to the credit crisis. Across our treated firms, we calculated the average amount of long-term debt due in 2008, as well as the amount of “cuts” conducted elsewhere to help pay off this debt (besides investment reductions) — inventories, share repurchases, dividends, and cash holdings. These variables were present for 77 of our 86 treated firms. We thus use these firms in the subsequent analysis.

For this sample of 77 firms, we compute the average changes in all of the policy variables above, between the first three quarters of 2007 and the first three quarters of 2008. For our two stock variables (cash holdings and inventories), we just take the differences in the average value of their levels in the first three quarters of 2008 relative to the first three quarters of 2007. For the quarterly flow variables (investment, share repurchases, and dividends), we convert the differences in the average quarterly flow to an annual flow basis for ease of comparison with the stock variables. For example, the quarterly reduction of investment (normalized by capital) of 2.1% for the treated firms reported in the first row of Panel A Table 3, represents an annual decline of 8.4%. To facilitate comparisons with our estimate of the fall in investment, we normalize all the other variables by the value of the capital stock as well. We then take the averages across all 77 of our treated firms to see how much they drew down their cash reserves, cut dividends, etc. We then compare that to the average amount of debt they had coming due in 2008.

Figure 4 provides a visual illustration of the treated firms’ response to the credit crisis. In this figure, we report the average changes in the policy variables as a fraction of the total amount of long-term debt maturing in 2008 for the treated firms. The decline in investment spending in 2008

²⁰The survey evidence in Brav, Graham, Harvey, and Michaely (2005) suggests that share repurchases are the residual after the investment and dividend decisions have been made.

represents about one-eighth of the amount of long-term debt these firms had coming due in 2008. By comparison, the treated firms drew down their cash reserves in amounts that represented about two-fifths of the amount of debt due in 2008. These firms reduced share repurchases (relative to 2007 levels) by an amount representing about one-tenth of the debt due. And reductions in their inventories accounted for another 7% of the 2008-maturing debt. Given executives' strong aversion to cutting dividends (see Brav, Graham, Harvey, and Michaely (2005)), it is perhaps not surprising that dividend cuts during 2008 accounted for only 1% of the amount of debt due for the treated firms, with the remaining 30% to be explained by other factors (such as reductions in R&D, labor costs, and asset sales).²¹

While admittedly done solely for purposes of providing a crude calculation for how the treated firms responded overall to the financial crisis, the set of numbers depicted in Figure 4 fits our economic intuition very well. In particular, the figure suggests that firms that were burdened with large amounts of maturing debt in 2008 tapped their least costly sources of funds (such as cash holdings) in order to mitigate the effects of maturing debt, but had to ultimately cut back on real activities, such as investment spending.

FIGURE 4 ABOUT HERE

5 Concluding Remarks

We use the August 2007 credit panic to assess the effect of financial contracting on real corporate policies. In particular, we test whether firms with large fractions of their long-term debt maturing at the time of the crisis observe more pronounced negative outcomes than otherwise similar firms whose debt structure is such that they do not need to refinance their obligations during the crisis. Our empirical methodology aims at replicating an experiment-like design in which we control for observed and unobserved firm heterogeneity via a differences-in-differences matching estimator.

We find evidence that long-term financial contracting has a sizeable effect on firms' real and financial policies when the firms face a credit supply shock. Firms whose long-term debt was largely maturing right after the third quarter of 2007 reduce investment (normalized by capital and measured on a quarterly basis) by 2.5% more than otherwise similar firms whose debt matures well after the crisis. This relative decrease in investment for firms with short maturity structures is statistically significant and economically large (approximately one-third of the pre-crisis level of investment for these firms). A number of falsification and placebo tests confirm our inferences about the effect of

²¹Campello, Graham, and Harvey (2009) survey 574 U.S. CFOs at the end of 2008. These managers report cuts of 11% in their firms' R&D expenditures and another 4% in their work force. Moreover, nearly 50% of the CFOs surveyed say that they sold assets in 2008 to cope with the credit squeeze.

credit supply shocks on corporate policies.

Our results contribute to the literature in a number of ways. First, our unique identification strategy reveals a novel link between debt maturity structure and corporate investment. In particular, our results point out to the importance of maturity structure for corporate financial flexibility. Second, our results provide evidence that the 2007 credit crisis had real effects on corporate behavior in 2008. Third, our evidence suggests that debt maturity structure is an important variable in understanding how credit supply shocks spread through the corporate sector. Undoubtedly, understanding the effects of credit cycles (and credit crises in particular) is not only of interest for corporate finance researchers, but also for economic policymakers. More broadly, our findings — obtained from an experiment-like framework — provide new evidence that financial contracting has causal effects on real corporate outcomes.

References

- Abadie, A., and G. Imbens, 2002, "Simple and Bias-Corrected Matching Estimators for Average Treatment Effects," NBER Technical Working Paper #0283.
- Abadie, A., D. Drukker, J. Herr, and G. Imbens, 2004, "Implementing Matching Estimators for Average Treatment Effects in Stata," *Stata Journal* 4, 290-311.
- Almeida, H., M. Campello, and M. Weisbach, 2004, "The Cash Flow Sensitivity of Cash." *Journal of Finance* 59, 1777-1804.
- Barclay, M., and C. Smith Jr., 1995, "The Maturity Structure of Corporate Debt," *Journal of Finance* 50, 609-631.
- Bolton, P., and D. Scharfstein, 1996, "Optimal Debt Structure and the Number of Creditors," *Journal of Political Economy* 104, 1-25.
- Brav, A., J. Graham, C. Harvey, and R. Michaely, 2005, "Payout Policy In The 21st Century," *Journal of Financial Economics* 77, 483-527.
- Brunnermeier, M., and M. Yogo, 2009, "A Note on Liquidity Risk Management," NBER Working Paper #14727.
- Campello, M., J. Graham, and C. Harvey, 2009, "The Real Effects of Financial Constraints: Evidence from a Financial Crisis," Working Paper, University of Illinois and Duke University.
- Chava, S., and A. Purnanandam, 2008, "The Effects of Banking Crisis on Bank-Dependent Borrowers," Working Paper, University of Michigan.
- Diamond, D., 1991, "Debt Maturity Structure and Liquidity Risk," *Quarterly Journal of Economics* 106, 709-737.
- Diamond, D., 1993, "Seniority and Maturity of Debt Contracts," *Journal of Financial Economics* 33, 341-368.
- Diamond, D., 2004, "Presidential Address: Committing to Commit Short-Term Debt When Enforcement is Costly," *Journal of Finance* 59, 1447-1479.
- Duchin, R., O. Ozbas, and B. Sensoy, 2009, "Costly External Finance, Corporate Investment, and the Subprime Mortgage Credit Crisis," Working Paper, University of Southern California.
- Fazzari, S., and B. Petersen, 1993, "Working Capital and Fixed Investment: New Evidence on Financing Constrains", *RAND Journal of Economics* 24, 328-342.
- Flannery, M., 1986, "Asymmetric Information and Risky Debt Maturity Choice," *Journal of Finance* 41, 19-37.
- Frank, M., and V. Goyal, 2003, "Testing the Pecking Order Theory of Capital Structure," *Journal of Financial Economics* 67, 217-248.

- Gorton, G., 2008, "The Panic of 2007," NBER Working Paper #14358.
- Greenwood, R., S. Hanson, and J. Stein, 2008, "A Gap-Filling Theory of Corporate Debt Maturity Choice," Working Paper, Harvard University.
- Guedes, J., and T. Opler, 1996, "The Determinants of the Maturity of Corporate Debt Issues," *Journal of Finance* 51, 1809-1833.
- Imbens, G., 2004, "Nonparametric Estimation of Average Treatment Effects under Exogeneity: A Review," *Review of Economics and Statistics* 86, 4-29.
- Lemmon, M., and M. Roberts, 2008, "The Response of Corporate Financing and Investment to Changes in the Supply of Credit," forthcoming, *Journal of Financial and Quantitative Analysis*.
- Rauh, J., 2006, "Investment and Financing Constraints: Evidence from the Funding of Corporate Pension Plans," *Journal of Finance* 61, 33-71.
- Stohs, M., and D. Mauer, 1996, "The Determinants of Corporate Debt Maturity Structure," *Journal of Business* 69, 279-312.

Table 1: Characteristics of Treated, Non-Treated, and Control Firms at the end of 2007: Median Test

The sample of firms is defined as those that have a long-term leverage ratio greater than 0.05 at the end of 2007 (i.e., long-term debt represents more than 5 percent of total assets). The 1,067 sample firms are split into treated and non-treated groups. The treated firms are defined as those for which the percentage of long-term debt maturing within one year (i.e., 2008) is greater than 20 percent and non-treated firms are defined as those for which the percentage of long-term debt maturing within one year is less than or equal to 20 percent. Control firms are a subset of the non-treated firms selected as the closest match to the treated firms based on a set of firm characteristics: Q , cash flow, size, cash holdings, long-term debt normalized by assets, 2-digit SIC industry, and credit ratings. There are 86 treated firms and 86 control firms. The medians of Q , cash flow, size, cash holdings, and long-term leverage are displayed for the three samples of firms (treated, non-treated, and controls). The average quarterly investment-to-capital ratio over the first three quarters of 2007 is also displayed. See text for further variable definitions. The test for a difference in the medians of a firm characteristic across two groups is conducted by calculating the Continuity Corrected Pearson χ^2 statistic, with the p -values of this test reported at the bottom row of each panel.

	Q	Cash Flow	Size	Cash	LT Leverage	Investment
<i>Panel A: Medians for Treated and Non-Treated Firms in 2007</i>						
Treated	1.728	0.076	5.870	0.080	0.244	0.047
Non-Treated	1.499	0.056	6.784	0.045	0.294	0.047
Difference	0.229	0.020	-0.914	0.035	-0.050	0.000
Median Test p -value	0.005	0.009	0.005	0.005	0.093	0.918
<i>Panel B: Medians for Treated and Control Firms in 2007</i>						
Treated	1.728	0.076	5.870	0.080	0.244	0.047
Control	1.599	0.070	6.266	0.063	0.233	0.051
Difference	0.129	0.006	-0.396	0.017	0.011	-0.003
Median Test p -value	0.286	0.446	0.286	0.879	0.647	0.879

***, **, * indicate significance at the 1, 5, and 10 percent levels, respectively.

Table 2: Distributional Tests of Treated, Non-Treated, and Control Firms at the end of 2007

The sample of firms is defined as those that have a long-term leverage ratio greater than 0.05 at the end of 2007 (i.e., long-term debt represents more than 5 percent of total assets). The 1,067 sample firms are split into treated and non-treated groups. The treated firms are defined as those for which the percentage of long-term debt maturing within one year (i.e., 2008) is greater than 20 percent and non-treated firms are defined as those for which the percentage of long-term debt maturing within one year is less than or equal to 20 percent. Control firms are a subset of the non-treated firms selected as the closest match to the treated firms based on a set of firm characteristics: Q , cash flow, size, cash holdings, long-term debt normalized by assets, 2-digit SIC industry, and credit ratings. There are 86 treated firms and 86 control firms. The medians of Q , cash flow, size, cash holdings, and long-term leverage are displayed for the three samples of firms (treated, non-treated, and controls). The average quarterly investment-to-capital ratio over the first three quarters of 2007 is also displayed. See text for further variable definitions. The 25th percentile, median, and 75th percentile are reported for each firm characteristic. The test for a difference in the distribution of a firm characteristic across two groups is conducted by calculating the Corrected Kolmogorov-Smirnov D statistic, with the p -values of this test reported in the rightmost column.

		25th %	Median	75th %	Kolmogorov-Smirnov Test p -value
<i>Panel A: Characteristics of Treated vs. Non-Treated Firms in 2007</i>					
Q	Treated	1.341	1.728	2.305	0.006
	Non-Treated	1.185	1.499	2.081	
Cash Flow	Treated	0.033	0.076	0.150	0.013
	Non-Treated	0.026	0.056	0.116	
Size	Treated	4.320	5.870	7.640	0.000
	Non-Treated	5.730	6.784	7.883	
Cash	Treated	0.021	0.080	0.184	0.005
	Non-Treated	0.017	0.045	0.126	
LT Leverage	Treated	0.159	0.244	0.356	0.096
	Non-Treated	0.186	0.294	0.427	
Investment	Treated	0.027	0.047	0.095	0.365
	Non-Treated	0.027	0.047	0.082	
<i>Panel B: Characteristics of Treated vs. Control Firms in 2007</i>					
Q	Treated	1.341	1.728	2.305	0.160
	Control	1.263	1.599	2.063	
Cash Flow	Treated	0.033	0.076	0.150	0.676
	Control	0.043	0.070	0.124	
Size	Treated	4.320	5.870	7.640	0.676
	Control	4.549	6.266	7.237	
Cash	Treated	0.021	0.080	0.184	0.416
	Control	0.019	0.063	0.161	
LT Leverage	Treated	0.159	0.244	0.356	0.977
	Control	0.154	0.233	0.341	
Investment	Treated	0.027	0.047	0.095	0.915
	Control	0.028	0.051	0.091	

Table 3: Differences-in-Differences of Firm Investment Before and After the Fall 2007 Credit Crisis with a Placebo Test Conducted a Year Before the Credit Crisis

Panel A of this table presents an estimate of the change in average quarterly investment rates from the first three quarters of 2007 to the first three quarters of 2008 (before and after the fall 2007 credit crisis). Panel B presents an estimate of the change in investment from the first three quarters of 2006 to the first three quarters of 2007 (a placebo test conducted before the credit crisis). In Panel A, the average of quarterly investment during the first three quarters of 2008 and the first three quarters of 2007 is calculated for the treated firms and control firms, as well as the difference in the difference between the two groups of firms over the two years. The average quarterly investment is normalized by the capital stock at the preceding quarter; that is, by lagged property, plant, and equipment. The treated firms are defined as those for which the percentage of long-term debt maturing within one year (i.e., 2008) is greater than 20 percent and control firms are defined as those for which the percentage of long-term debt maturing within one year is less than or equal to 20 percent. Control firms are a subset of the non-treated firms selected as the closest match to the treated firms based on a set of firm characteristics: Q , cash flow, size, cash holdings, long-term debt normalized by assets, 2-digit SIC industry, and credit ratings. There are 86 treated firms and 86 control firms in Panel A. Panel B is constructed analogously, but the tests are conducted one year earlier (before the credit crisis). There are 113 treated firms and 113 control firms in Panel B. ATT is the Abadie-Imbens bias-corrected average treated effect matching estimator (Matching Estimator). Heteroskedasticity-consistent standard errors are in parentheses.

Average Quarterly Investment / Capital Stock (in percentage points)			
	2008	2007	Difference
<i>Panel A: Investment Before and After the Fall 2007 Credit Crisis</i>			
Investment in 2008 (Q1 to Q3) vs. Investment in 2007 (Q1 to Q3)			
Treated Firms	5.7*** (0.5)	7.8*** (0.9)	-2.1** (0.8)
Control Firms	7.3*** (0.6)	7.3*** (0.7)	0.1 (0.7)
Difference	-1.6*** (0.6)	0.6 (1.0)	-2.2** (1.0)
Matching Estimator (ATT)			-2.5** (1.1)
<i>Panel B: The Placebo Test</i>			
Investment in 2007 (Q1 to Q3) vs. Investment in 2006 (Q1 to Q3)			
Treated Firms	6.9*** (0.7)	7.3*** (0.6)	-0.4 (0.7)
Control Firms	6.9*** (0.7)	7.2*** (0.8)	-0.3 (0.8)
Difference	0.0 (0.8)	0.1 (0.8)	-0.1 (1.0)
Matching Estimator (ATT)			0.0 (1.1)

***, **, * indicate significance at the 1, 5, and 10 percent levels, respectively.

Table 4: Differences-in-Differences of Firm Investment from One Year to the Next: 2000 through 2007

This table presents an estimate of the change in investment from the first three quarters of a given year to the first three quarters of the next year. The first row replicates the Differences-in-Differences and Matching Estimator (ATT) from Panel A of Table 3 and the second row replicates the Differences-in-Differences and Matching Estimator (ATT) from Panel B of Table 3. Analogous results are then presented for the other years. The treated firms are defined as those for which the percentage of long-term debt maturing within one year (i.e., 2008) is greater than 20 percent and control firms are defined as those for which the percentage of long-term debt maturing within one year is less than or equal to 20 percent. Control firms are the closest matches to the treated firms based on a set of firm characteristics (please see the description in Table 3 for details). ATT is the Abadie-Imbens bias-corrected average treated effect matching estimator (Matching Estimator). Heteroskedasticity-consistent standard errors are in parentheses.

Investment Change	Difference in the change in investment between treated and control firms	Matching Estimator (ATT)
2008 – 2007	–2.2** (1.0)	–2.5** (1.1)
2007 – 2006	–0.1 (1.0)	0.0 (1.1)
2006 – 2005	0.2 (1.0)	0.2 (1.0)
2005 – 2004	–0.7 (0.5)	–0.5 (0.5)
2004 – 2003	0.3 (0.5)	0.2 (0.5)
2003 – 2002	0.2 (0.5)	0.3 (0.5)
2002 – 2001	0.2 (0.9)	0.6 (0.9)
2001 – 2000	0.7 (0.9)	1.3 (1.0)
Pooled Analysis: All Years Before Fall 2007 Credit Crisis	0.1 (0.3)	0.2 (0.3)

***, **, * indicate significance at the 1, 5, and 10 percent levels, respectively.

Table 5: Differences-in-Differences of Firm Investment Before and After the Fall 2007 Credit Crisis: Different Leverage Cutoffs

This table presents an estimate of the change in investment from the first three quarters of 2007 to the first three quarters of 2008 (before and after the fall 2007 credit crisis) for various long-term leverage cutoffs: leverage ratio of more than 0%, leverage ratio of more than 5%, and leverage ratio of more than 10%) The leverage ratio cutoff of 5% presented in the middle column (i.e., long-term debt represents more than 5 percent of assets) reproduces the results presented in Panel A of Table 3 for ease of comparison. The treated firms are defined as those for which the percentage of long-term debt maturing within one year (i.e., 2008) is greater than 20 percent of total assets and control firms are defined as those for which the percentage of long-term debt maturing within one year is less than or equal to 20 percent of total assets. Control firms are the closest matches to the treated firms based on a set of firm characteristics (please see the description in Table 3 for details). ATT is the Abadie-Imbens bias-corrected average treated effect matching estimator (Matching Estimator). Heteroskedasticity-consistent standard errors are in parentheses.

	Long-Term Leverage > 0%	Long-Term Leverage > 5%	Long-Term Leverage > 10%
Change in Investment for Treated Firms	-1.1* (0.6)	-2.1** (0.8)	-2.7** (1.2)
Change in Investment for Control Firms	-1.1* (0.5)	0.1 (0.7)	-0.5 (1.0)
Difference	0.0 (0.7)	-2.2** (1.0)	-2.2 (1.5)
Matching Estimator (ATT)	0.2 (0.8)	-2.5** (1.1)	-3.4** (1.3)

***, **, * indicate significance at the 1, 5, and 10 percent levels, respectively.

Figure 1: Historical LIBOR and Commercial Paper Spreads

This figure displays the 3-month LIBOR and commercial paper (CP) spreads over treasuries, for the period of January 2001 to June 2008. The data is from <http://www.federalreserve.gov/datadownload/>.



Figure 2: Composition of Long-Term Debt Maturity at the end of 2007

This figure displays the amount of long-term debt maturing in the years of 2008 to 2012, as a fraction of total long-term debt, for the sample of firms described in Section 2.3. Maturity structure is measured at the end of the 2007 fiscal year.




Figure 3: Composition of Long-Term Debt Maturity: 1999 to 2006

This figure displays the amount of long-term debt maturing in one to five years away from an initial year t , as a fraction of total long-term debt, for the sample of firms described in Section 2.3. Maturity structure is measured at the end of fiscal year t , with t varying from 1999 to 2006.




Figure 4: How did Treated Firms Pay Off Their Debt?

This figure displays changes in policy variables from the first three quarters of 2007 to the first three quarters of 2008, as a fraction of the amount of long-term debt maturing in 2008, for the sample of 77 treated firms for which we have complete data on investment, cash holdings, cash dividends, inventories, and share repurchases. Treated firms are those which have more than 20% of their long-term debt maturing in 2008.

