A Market-Based Study of the Cost of Default*

Sergei A. Davydenko   Ilya A. Strebulaev   Xiaofei Zhao

October 2011

Abstract

Although the cost of financial distress is a central issue in capital structure and credit risk studies, reliable estimates of its size are difficult to come by. This paper proposes a novel method of extracting the cost of default from the change in the market value of a firm's assets upon default. Using a large sample of firms with observed prices of debt and equity that defaulted over 14 years, we estimate the cost of default for an average defaulting firm to be 21.7% of the market value of assets. The costs vary from 14.7% for bond renegotiations to 30.5% for bankruptcies, and are substantially higher for investment-grade firms (28.8%) than for highly-levered bond issuers (20.2%), which extant estimates are based on exclusively.

Keywords: Default; Bankruptcy; Renegotiation; Costs of financial distress

JEL Classification Numbers: G21, G30, G33

*Sergei Davydenko (davydenko@rotman.utoronto.ca) and Xiaofei Zhao (xiaofei.zhao08@rotman.utoronto.ca) are at Joseph L. Rotman School of Management, University of Toronto. Ilya Strebulaev (istrebulaev@stanford.edu) is at the Graduate School of Business, Stanford University, and NBER. We thank Peter Christoffersen, Jeff Coles, Pascal François, Jerome Taillard, Jan Mahrt-Smith, Alan White, and seminar participants at Arizona State University, HEC Montreal, SUNY Binghamton, the University of Toronto, and the European Finance Association Stockholm meetings for comments and suggestions. Financial support from SSHRC is gratefully acknowledged.

1. Introduction

The cost of financial distress is among the most important factors thought to affect corporate financing decisions. It is a crucial parameter both in studies of capital structure and in models of corporate securities pricing. However, there is no consensus in the literature regarding the size of the costs arising from financial distress and default, in particular because separating their effect on firm value from the effect of contemporaneous economic distress is notoriously difficult. Our paper proposes a novel approach to estimating the cost of default, based on the idea that the change in the combined market value of the firm's debt and equity upon default reflects the total cost of default, as well as the degree to which it is a surprise for investors. Using a large sample of bond defaults from 1997 to 2010, we find the cost for an average defaulting firm to be 21.7% of the market value of assets. Default costs vary from 14.7% for bond renegotiations to 30.5% for bankruptcies, and are substantially higher for investment-grade firms (28.8%) than for highly-levered firms (20.2%), which extant estimates are based on exclusively.

The costs of financial distress include both direct and indirect components. Direct costs, such as lawyers' fees in bankruptcy, are relatively straightforward to estimate, but they do not exceed a few percent of firm value. Indirect costs of financial distress are much more difficult to measure, but also potentially much larger than direct costs.[1] As financial distress (the inability to meet required debt payments) typically occurs simultaneously with economic distress (deteriorating economic fundamentals), the effect of "pure" financial distress on firm value is difficult to identify empirically. As a result, quantitative estimates of the value loss due to financial distress are so rare that currently, most if not all studies that make use of such estimates have to rely on just one systematic study of distress costs, that of Andrade and Kaplan (1998), henceforth referred to as AK.

Based on 30 highly-leveraged transactions (HLT) that became distressed between 1987 and 1992, AK conclude that the total cost of financial distress for HLTs is likely to be in the 10% to 20% range. Their numerical estimates have been applied in studies of capital structure (Graham (2000); Molina (2005); Almeida and Philippon (2007); Elkamhi, Ericsson, and Parsons (2011)),

[1] Estimates of direct bankruptcy costs range from 3.1% (Weiss (1990)) to 5.3% (Warner (1977)) to 6% (Altman (1984)) of firm value (see also Bris et al. (2006)). Indirect costs of distress may arise, for example, due to managerial distraction, distortions in the customer-supplier relationship (Titman (1984)), losses from asset fire sales (Shleifer and Vishny (1992)), and agency costs of debt due to asset substitution (Jensen and Meckling (1976)) and debt overhang (Myers (1977)).

implementations of structural bond-pricing models (Eom, Helwege, and Huang (2004); Huang and Huang (2003)), calibrations of dynamic models of the levered firm (Miao (2005)), studies of the effect of macroeconomic variables on asset prices and capital structure (Bhamra, Kuehn, and Strebulaev (2010a, 2010b)), and other settings. Yet, as AK point out, the HLTs that their sample consists of may have chosen to become highly levered precisely because their distress costs were unusually low. As a result, for a typical firm AK's estimates may be biased downward, and applying them to non-HLTs may be problematic.

The primary goal of this paper is to provide new estimates of the total costs associated with default and bankruptcy based on the market prices of defaulted firms not limited to HLTs. We do this by combining new data on the market values of debt and equity for defaulted firms with a novel estimation procedure that extends the event study methodology to events such as default announcements, which may be partially anticipated by investors long before they occur.

At the heart of our approach is the idea that investors anticipate default only partially, so that the announcement of default contains an element of surprise. As a result, upon the announcement, the total market value of the firm's debt and equity changes, and the size of the change reflects both the cost of default and the degree to which it is unanticipated. Loosely speaking, if $D + E$ is the total market value of the firm's debt and equity, V is the value of assets in the absence of default, and $V - c$ is their "recovery" value in default (where c is the cost of default), then the pre-default value of the firm is $D + E = (1 - q) \times V + q \times (V - c)$, where q is investors' risk-adjusted estimate of the default probability. Upon default, the value of the firm decreases to $V - c$. The total firm-level price reaction to default equals $(1 - q) \times c$, which is the cost of default scaled down due to investors' partial anticipation of default. We formalize this intuition in a dynamic setting, and evaluate the conditional probability of default, q, from historical defaults and debt prices. This allows us to undo the effect of partial anticipation on the price reaction and compute the total cost of default, c.

We apply this approach to a sample of 175 firms, which defaulted between 1997 and 2010, and for which market prices of bonds, bank loans, and equity are observed both just prior to and shortly after default. For an average defaulting firm, we estimate the mean (median) cost of default to be 21.7% (22.1%) of the market value of assets. These estimates are robust to assumptions of the model that we use to back out investors' expectations of default. We also allow for the possibility that investors' information about the value of assets is noisy, so that observed price reactions to default may reflect investors' realization that the asset value was lower than thought. The reduction in our mean estimates due to this learning-from-default effect does not exceed 3 percentage points.

Looking at different types of default, we find the costs of a distressed bond exchange to be 14.7%, compared with the average cost of bankruptcy of 30.5%. Importantly, sample firms with lower default costs have lower ratings when they issue bonds. For highly levered (original-issue junk) firms default costs average 20.2%, which is similar to AK's estimates for HLTs. However, the costs are higher (28.8%) for fallen angels (firms originally rated investment grade), which may be more representative of a typical firm for the purposes of capital structure studies. Nonetheless, we also find that marginal and ex ante expected risk-adjusted default costs remain small in comparison with the likely tax benefits of debt, consistent with Graham (2000), Elkamhi et al. (2011), and other studies based on AK's averages. Interestingly, for 30% of defaults the market value of the firm increases upon the announcement, implying negative default costs. For these firms, the absence of default likely involves value-destroying activities, whereas default may precipitate a value-increasing shake-up (e.g., Hotchkiss and Mooradian (1997), Taillard (2011)). Finally, consistent with Acharya et al. (2007), we find the cost of default to vary with industry conditions, as predicted by Shleifer and Vishny (1992).

Our estimation procedure can be viewed as a generalization of the event study methodology (e.g., Brown and Warner (1985)). Event studies look at the price reaction to various corporate announcements. They typically deal with potential information leakages that can affect market prices before the event by extending the observation window backwards. Unfortunately, this approach cannot be applied to studying the cost of default because the *timing* of default itself is systematically related to the value of the firm. Our approach allows the timing of the event of interest (in our case, default) to depend endogenously on the quantity to be measured (in our case, the value of the firm's assets), and the possibility of the event to affect prices for an arbitrarily long time. The procedure can be applied to any defaulted firm with observed market prices of debt and equity, including those distressed both economically and financially.

A number of existing empirical studies focus on various components of distress costs, such as price discounts in asset fire sales (Pulvino (1998)), risk shifting (Eisdorfer (2008)), the loss of market share in industry downturns (Opler and Titman (1994)), and the debt overhang problem (Franks and Sanzhar (2006)), but do not assess how they interact with each other. With the exception of

Andrade and Kaplan (1998) and a case study by Cutler and Summers (1988), most studies that attempt to estimate total distress costs do so by studying the behavior of non-distressed firms, rather than the realized value loss in default. Based on the time series behavior of market prices of debt and equity, Korteweg (2010) estimates expected distress costs for highly levered firms between 15% and 30%. His estimates reflect total ex post costs as well as the probability of becoming distressed, which are not identified separately. Similar to his estimates, we find that the expected cost of default (that is, the decrease in the value of the firm due to expected default) just prior to default for an average firm is 14%. Using structural estimation, Hennessy and Whited (2007) estimate implied bankruptcy costs to be between 8.4% and 15.1%. For comparison, the mean firm-level price reaction to bankruptcy in our sample is –18.5%. Glover (2011) finds that for observed leverage ratios to be consistent with the trade-off theory, the average cost of default must be as high as 45%. At the same time, because of the self-selection problem, the sample average among defaulting firms is only 24.6%, which is only slightly above what we observe empirically.

Overall, to date, Andrade and Kaplan (1998) remains the main reference for numerical estimates of firm-specific ex post costs of financial distress. AK overcome the need for disentangling the effects of financial and economic distress by focusing on 30 distressed highly levered transactions. Because the firms in their sample have above-average operating performance, AK conclude that they are financially distressed due to their high leverage, but are not economically distressed. AK estimate the total change in the value of the firm between the onset of distress and its resolution, and attribute it to distress costs.

Our estimation procedure offers several potential advantages. First, our sample includes not only original-issue junk firms (i.e., highly levered bond issuers), but also fallen angels that were rated investment grade at the time of bond issuance, and these appear to have systematically higher default costs. Second, our estimates are based on the change in the observed market value of the firm around the announcement of default. By contrast, AK's estimates are based on the change in cash flow margins from before the onset of distress to its resolution, multiplied by the industry median ratio of the firm value to cash flow, plus 2% that they add to account for direct costs of bankruptcy based on other empirical studies. Such an approach may be less accurate when changes in cash flow margins are transitory and as such do not translate into a proportional shift in the market asset value; when cash flow multiples differ across firms within an industry; or when the firm's direct costs of bankruptcy are unusually high or low. Third, most firms in our sample are distressed not only financially but also economically, which is a far more common situation than that of purely financial distress (Asquith, Gertner and Scharfstein (1994)). Because economic distress depletes firms' assets in the run-up to default, estimates of default costs are affected when expressed as a proportion of the value of assets at the time of default. Fourth, as debt pricing data sets become more readily available, our procedure can be applied to larger and more recent samples, reducing noise in the estimates and facilitating cross-sectional analysis.

A potential limitation of our estimates is that they are based on the price reaction to the default event, and as such cannot be directly applied to measure agency costs of financial distress incurred by non-defaulting firms. However, existing estimates indicate that such costs are unlikely to be substantial (Mello and Parsons (1992); Parrino and Weisbach (1999); Moyen (2007)). Moreover, AK do not find the cost of distress for defaulting and non-defaulting firms to be statistically different. Given that even for highly-levered firms our estimates are at the upper bound of AK's range, our results suggest that distress costs may be higher than previously thought.

The remainder of this paper is organized as follows. Section 2 discusses our estimation procedure. Section 3 describes the data. Empirical results are reported in Section 4. Section 5 concludes. The derivation of the equations and the procedure used to estimate the market value of the firm are described in the appendices.

2. Estimating the costs of default

In this section, we describe the approach that we use to estimate the unobservable costs of default from observed market prices of debt and equity before and after default. Our estimation procedure is based on the idea, first introduced in models of risky debt by Duffie and Lando (2001), that the information that investors have about firms' economic fundamentals is noisy and incomplete. As a result, investors generally cannot conclude with certainty whether or not any given firm is so distressed that it is about to default in the next instant. Indeed, if investors had enough information to replicate the timing of managers' decision to default, then an announcement of default would never be a surprise. Hence, by the time the firm defaulted, its debt and equity prices would have gradually converged to their post-default "recovery" values, and upon the announcement of default prices would not move even if default involved deadweight value losses.

Empirically, however, it is well known that upon default firms' assets exhibit large abnormal returns. Clark and Weinsten (1983) and Lang and Stulz (1992) document abnormal stock returns at bankruptcy of around –20% to –30%, whereas Warner (1977) finds that prices of public bonds of bankrupt railroads fall by 9.2% in the month of bankruptcy. These large price reactions to default imply that default is not perfectly anticipated by investors. Duffie and Lando (2001), Jarrow and Protter (2004) and Giesecke (2006) argue that investors are only partially informed about crucial parameters that determine the timing of default. They show that under certain assumptions these information imperfections imply that the assets of the distressed firm can be priced as if, conditional on information available to investors, default were a random event with a hazard rate that is a function of the firm's economic conditions. Observed pre-default debt and equity prices reflect both the "recovery" value that the firm's assets would have in default, and their "continuation" value in the absence of default, with the difference between the two arising because default is costly. By observing market values of firms immediately prior to default and their recovery values immediately after default, and by parameterizing the default hazard, one can solve for the implied continuation value that the firm's assets would have if default were never to occur. The net cost of default can then be found by subtracting the recovery value of the firm from the continuation value of assets.

2.1. A static illustration

To illustrate the key idea of our approach, consider a simple static example of a levered firm that has to make a single (and final) debt payment of B at time T. If the firm does not default on the debt payment, the value of its productive assets, also referred to as their "continuation" value, will be equal to V. If it defaults, the "recovery" value of the assets L will generally be different.[2] The net cost of default, c, is defined as the difference between the value of assets absent the possibility of default and their value in default: $c = V - L$. Default may be costly due to transactions costs of arranging a distressed bond exchange, legal fees in bankruptcy, lost sales due to customers' unwillingness to buy from a defaulted company, opportunity costs of management's time, expected asset fire sale discounts, and other factors. For some firms, default may also be beneficial (its net cost may be negative) if it precipitates a value-increasing shake-up, like a sale of the firm to higher-value users, which self-serving managers may resist in the absence of default.

[2] In structural models of risky debt, such as Merton (1974), V has a natural interpretation as the present value of assets or asset-generated cash flows in the absence of any financing imperfections.

In the base case, we assume that investors observe both V and L,[3] and use them in conjunction with an estimate of the risk-neutral probability of default, q, to calculate the value of the firm's financial claims, such as debt and equity. The value of the firm, M, is the total value of all such claims. If investors believe that default is possible but not certain, M will depend both on V and on L. As econometricians, we observe the value of the firm prior to the scheduled payment, M, and, in case of default, its recovery value, L. Unlike the investors, we do not know the continuation value, V. Our task is to estimate the cost of default, $c = V - L$, from observed prices of debt and equity.

Suppose that just prior to time T investors know both V and L but lack full information regarding some other important economic parameters that affect the firm's ability to make the required debt payment. For example, investors may be unsure if the firm has enough liquid assets to repay the debt, and if not, whether it will be able to raise the required cash from external sources. As a result of these information imperfections, up until the maturity of debt investors can neither be sure that the firm will make the debt payment, nor know with certainty that it will not. They determine the market prices of debt and equity at T_{-} (i.e., just prior to time T) given their assessment of the risk-neutral probability of default q, conditional on the information available to them. Investors' estimate of q may be based, for instance, on the distance-to-default (a volatility-adjusted measure of market leverage based on the Merton (1974) model) and the firm's accounting ratios (e.g., Altman's (1968) z-score).

In this setting, the market value of the firm at time T_{-}, i.e., the total value of its debt and equity, is the probability-weighted average of the continuation and recovery values of its assets:

$$M = V \times (1 - q) + L \times q. \tag{1}$$

Given this relationship, we can compute the cost of default implied by market prices as follows. First, we estimate investors' conditional default probability q, for example, from bond prices or from survival analysis of firms at risk of failure. Second, if the firm does default, we can measure its recovery value L and its pre-default value M. These are, respectively, the total market value of the firm's debt and equity immediately after default, and their value just prior. Third, we solve Equation (1) for the unobserved continuation value of assets V. Finally, we find the cost of default

[3] In robustness checks (Subsection 4.4.3) we use a model in which investors can only observe a noisy signal about V.

as $c = V - L$.

Our approach can be interpreted as adjusting the observed firm price reaction upon default so as to undo the effect that partial anticipation of default has on pre-default debt and equity prices. To see this, notice that Equation (1) can be re-written as:

$$M - L = c \times (1 - q). \tag{2}$$

The left-hand side of this equation is equal to the (negative of the) jump in the firm value upon the announcement of default, i.e., the firm-level price reaction to default. The right-hand side equals the cost of default, times one minus investors' conditional probability of default, which is a measure of the extent to which default is a surprise. As long as default is partially anticipated, so that the conditional probability of default is positive, Equation (2) implies that (the negative of) the change in the firm value upon default is smaller than the cost of default. At the same time, the two are closely related, and the less default is expected by investors, the closer the price reaction is in magnitude to the total default cost. The sign of the cost of default is always opposite to that of the observed firm price reaction.

Our estimation procedure can be viewed as a generalization of the event study methodology (e.g., Brown and Warner (1985)). If the event of interest (in our case, default) is partially anticipated by investors, the observed price reaction at the time of the event is the lower boundary for the total value effect of the event. Event studies deal with partial anticipation by extending the observation window backwards. Unfortunately, this approach cannot be applied to studying the cost of default, because investors may be factoring in the possibility of default for a long time prior to the actual announcement. Moreover, the firm's decision to default may be systematically related to the value of assets, which summarizes the degree of the firm's economic distress. We overcome these difficulties associated with the event study design by explicitly evaluating investors' conditional probability of default, and adjusting pre-default prices accordingly.

2.2. Base-Case Dynamic Model

The static model discussed above ignores the fact that in reality debt payments are spread over time. As a result, if the firm does not default at time t, its value at t_+ still differs from the asset value V, as it is affected by the possibility of default in the future. To account for this effect, one needs to specify investors' expectations about the future dynamics of the asset value and the default process.

In this subsection, we describe the dynamic model that our base-case estimates are based on. The model merges important features of both reduced-form and structural models of credit risk. At the same time, our approach is structured so as to minimize the reliance on a number of debatable assumptions of such models, such as default boundary conditions used in structural models.

2.2.1. The default hazard

The central assumption behind our approach is that due to information imperfections, investors cannot predict the timing of default perfectly. We assume that, as a result, there exists a default hazard rate, which is a function of investors' information.[4] Conditional on this information, default is a realization of a Poisson process stopped at its first jump. This approach to modeling default is common in reduced-form models of risky debt pricing (e.g., Duffie and Singleton (1999), Madan and Unal (1998)). However, most reduced-form models also assume that the default hazard is driven by some latent risk factors, inferred from the time-series behavior of credit spreads. In contrast, we explicitly specify the hazard rate as a function of observed firm characteristics.

To focus on the most important salient information available to investors, we assume that the hazard rate is a function of the firm's asset value and its outstanding debt. Specifically, under the real probability measure $\mathbb{P}$ the default hazard $\lambda_t^{\mathbb{P}}$ is:

$$\lambda_t^{\mathbb{P}} = e^{\beta_0 + \beta_1 \log \frac{V_t}{B}}, \tag{3}$$

where V_t is the market value of assets, B is the face value of debt, and β_0 and β_1 are fixed parameters. The ratio of the market value of assets to the face value of debt measures the firm's economic solvency and captures the degree of economic distress that the firm is in. The assumption that this ratio is a sufficient statistic for default is standard in many structural models of credit risk, starting from Merton (1974) and Black and Cox (1976). This ratio is the main input for computing the

[4] Duffie and Lando (2001) are the first to introduce asymmetric information in a structural model. They show that the default process in their model can be described using a hazard rate. Giesecke (2006) generalizes the conditions under which a hazard rate exists in such models.

distance-to-default and the EDF by Moody's/KMV, both of which are now widely used in academic literature and practical applications as a measure of the firm's default risk (e.g., Berndt et al.(2005)). Empirically, Davydenko (2011) shows that the ratio of the market value of assets to the face value of debt is by far the most powerful variable explaining the timing of default. Its explanatory power exceeds that of most other conventional default predictors (e.g., those entering Altman's (1968) *z*-score) put together, and in regression analysis most such factors become insignificant in its presence.[5] Hence, we are unlikely to lose much in accuracy by following structural models and focusing on the asset-to-debt ratio exclusively.

To relate observed asset prices to investors' expectations about default, we need the mapping between the actual and the risk-neutral probability measures. For Poisson processes, the change of the probability measure affects the intensity of jump arrivals (see, e.g., Shreve (2004), as well as Gorbenko and Strebulaev (2010) for an application to finance). We therefore assume that under the risk-neutral measure $\mathbb{Q}$, default is also a doubly-stochastic process, and that its intensity is a multiple of the real-measure intensity,

$$\lambda_t^{\mathbb{Q}} = \xi \lambda_t^{\mathbb{P}}, \tag{4}$$

where $\xi \geq 1$ is the risk premium associated with default, to be estimated from the data.[6]

At this point, several observations on our specification are in order. Our model combines the tractability of a reduced-form model with the economic intuition of structural models, which predict that default is driven by deteriorating economic fundamentals. At the same time, it does *not* rely on two common structural assumptions that are easiest to challenge on empirical grounds. First, in contrast to structural models such as Black and Cox (1976), Leland (1994), and others, we do not assume that there is a sharp value-based default boundary separating defaulting and nondefaulting firms. Contrary to this assumption, Davydenko (2011) finds that some firms default while their asset value is still relatively high, and others manage to avoid default at very low asset values, so that the reliance on the assumption of a sharp boundary known in advance is not very accurate. Second, we make no assumptions regarding how the firm's assets are divided between creditors and

[5] The only exception is measures of balance sheet liquidity, such as the current ratio, although their explanatory power is an order of magnitude lower than that of V_t/B. We control for liquidity in robustness checks, reported in Subsection 4.4.1, and find its effect to be small.

[6] Bhamra, Kuehn, and Strebulaev (2010a; 2010b) show that this representation is general and that the risk premium is a function of the primitives of the economy, such as preferences.

shareholders in default. While most structural models assume that the absolute priority rule (APR) is enforced in default, empirical studies of distressed reorganizations find that the APR is often violated in practice (e.g., Franks and Torous (1989)). Our approach is based on the aggregate value of the firm, and does not depend on its split between debt and equity.

2.2.2. The pricing equation

To relate the price reaction at default announcement to the cost of default, we proceed by specifying the risk-neutral dynamics of the continuation value of assets V_t and their recovery value, L_t. Following the standard assumptions in credit risk literature, we assume that V_t follows a geometric Brownian motion under the risk-neutral measure $\mathbb{Q}$:

$$dV_t = rV_t dt + \sigma V_t dW_t^{\mathbb{Q}}, \tag{5}$$

where r is the risk-free rate, σ is the volatility of assets, and $dW_t^{\mathbb{Q}}$ is a Brownian motion defined on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{Q}, (\mathcal{F}_t)_{t\geq 0})$. All parameters, as well as the face value of debt, B, are known constants. We also follow the literature (e.g., Leland (1994)) in assuming that the recovery value of the firm is a constant fraction of the asset value:

$$L_t = (1 - \alpha)V_t, \tag{6}$$

where α is the proportional cost of default.

Investors observe both V_t and L_t, and also evaluate the conditional risk-neutral default intensity $\lambda_t^{\mathbb{Q}}$. Then, as shown in the Appendix, at any time t up to default (or maturity T, whichever comes first) the value of the firm can be expressed as:

$$M_t = L_t + (V_t - L_t)\mathbb{E}_t^{\mathbb{Q}} \left[\frac{V_T e^{-r(T-t)}}{V_t} e^{-\int_t^T \lambda_u^{\mathbb{Q}} \, du} \right], \tag{7}$$

where the expectation $\mathbb{E}^{\mathbb{Q}}$ is conditional on all information available to investors.

This equation relates the market value of the firm to the continuation value of its assets and their recovery value. For a firm that defaults at time $t = \tau$, M_τ can be observed as the market value of the firm just prior to the announcement of default, and L_τ as its value immediately after. Thus, we can solve Equation (7) for V_τ, and compute the cost of default as $c = V_\tau - L_\tau$.

It is important to emphasize the economic content of V and its implications for interpreting our empirical results. In our model, V is the value of a copy-cat firm, identical to the firm that we observe, but for which default does not affect the value of the firm (i.e., it is costless). Note that, in general, V may not coincide with the value of unlevered assets of the firm, for two reasons. First, tax benefits of debt are incorporated in V but not in the unlevered firm value.[7] Second, even if default per se does not change the total value of the firm, it can still result in wealth re-distributions among the various stakeholders. This gives rise to agency conflicts, which may affect the value of the firm. For example, the levered firm may underinvest ("debt overhang" of Myers (1977)) or overinvest in risky projects ("asset substitution" of Jensen and Meckling (1976)). To the extent that such activities are costly, V could be below the unlevered value of assets, and default costs could be lower than the full cost of financial distress.

While we cannot ascertain the quantitative impact of such factors, it is unlikely to be substantial. Studies like Mello and Parsons (1992) and Parrino and Weisbach (1999) estimate the effect of agency costs on firm value to be below 2%, which is small in comparison with total distress costs. Andrade and Kaplan (1998), whose estimates of distress costs presumably incorporate agency costs, do not find significant differences between the costs for defaulting and non-defaulting firms. Moreover, our estimates of default costs are similar to or higher than AK's.

Similar to the static case of the previous subsection, Equation (7) implies a relationship between the cost of default $c = V_\tau - L_\tau$ and the firm-level price reaction to the default announcement, $M_\tau - L_\tau$:

$$M_\tau - L_\tau = c \times \mathbb{E}_\tau^{\mathbb{Q}} \left[\frac{V_T e^{-r(T-\tau)}}{V_\tau} e^{-\int_\tau^T \lambda_u^{\mathbb{Q}} \, \mathrm{d}u} \right], \tag{8}$$

where the expectation term on the right-hand side parallels the "surprise" component of the default announcement in the Equation (2) of the static model. Essentially, this term is the probability of no default until maturity, adjusted for the expected growth in the firm's assets between τ and T.

[7]Given that both before and after default most firms are loss-making with high debt levels (Gilson (1997)), and thus the expected present value of income taxes is very low, the effect of taxes in our sample is likely to be small.

2.2.3. Identifying the risk premium

We estimate the default risk premium parameter, ξ, from the market prices of debt before default and observed debt recovery rates. Consider a generic firm financed with a zero-coupon bond with a promised payment of \$1 at maturity. If the firm defaults, creditors receive R dollars at the time of default, where R is the recovery rate, assumed to be constant. As shows in the Appendix, the market value of the bond can be found from:

$$D_t = e^{-r(T-t)}\mathbb{E}_\tau^{\mathbb{Q}}\left[e^{-\xi\int_t^T \lambda^{\mathbb{P}}(V_u)\,\mathrm{d}u}\right] + \xi\int_t^T \mathbb{E}_\tau^{\mathbb{Q}}\left[Re^{-\int_t^\tau [r+\xi\lambda^{\mathbb{P}}(V_u)]\,\mathrm{d}u}\lambda^{\mathbb{P}}(V_\tau)\right]\mathrm{d}\tau. \tag{9}$$

For firms that are about to default, we observe the market value of debt immediately prior to default, as well as debt recovery rates following default. Given an estimate of the value of assets at that time, V_τ, we can evaluate the physical default hazard, $\lambda^{\mathbb{P}}(\cdot)$, from Equation (3), and solve Equation (9) for ξ. The idea is similar to solving Equation (7) for V_t based on the value of the firm before and after default. The main difference is that, to treat the risk premium as a market-wide parameter, in implementing the model we do not solve Equation (9) for each individual firm. Instead, we solve it once for an "average" firm in each year, using as inputs the average values of the parameters over all firms that defaulted in that year. We thus obtain time-varying estimates of the risk premium.[8]

2.3. Implementation

Our estimation procedure involves the following major steps. First, using a sample of defaulting and non-defaulting firm-months, we estimate the parameters of the hazard function under the real measure by means of survival analysis. Second, we compute the risk premium embedded in observed debt prices, and use it to transform the hazard rate to the risk-neutral measure. Third, we solve Equation (7) for V.

A complicating factor is that $\lambda_t^{\mathbb{Q}}$ in Equation (7) is a function of V_t, which we do not know initially. To estimate jointly the parameters of the risk-neutral default hazard function, the default risk premium, and the continuation value of assets, we employ the following iterative procedure.

[8] Several papers estimate the jump-to-default risk premium by comparing risk-neutral default probabilities implied by bond or CDS spreads with physical probabilities implied by credit ratings or the Expected Default Frequency (EDF) from Moody's/KMV (Driessen (2005); Berndt et al.(2008); Hull, Predescu, and White (2005)).

2.3.1. The iterative estimation procedure

Step 1. As an initial approximation for V_t, we choose $V_t^{(1)} = M_t$, i.e., we use the observed firm value as an initial guess for the continuation value of assets.

Step 2. We apply standard tools of parametric survival analysis (see, e.g., Kalbfleisch and Prentice (2002) and Shumway (2001)) to estimate the parameters of the hazard function $\lambda_t^{\mathbb{P}}\left(V_t^{(1)}\right)$ specified in Equation (3), using maximum likelihood for the whole sample of firms, including firm-month observations that do not correspond to default. This yields parameter estimates $\beta_0^{(1)}$ and $\beta_1^{(1)}$.

Step 3. Next, we estimate the default risk premium. For each year in the sample, we select firms that defaulted in that year. Based on the current approximation, we find the average market value of assets for these firms at default, $\overline{V_\tau^{(1)}}$. We also compute the average price of debt just before default, $\overline{D}$, and the average recovery rate, $\overline{R}$, as well as average volatility, debt maturity, and risk-free rate (these parameters do not change from iteration to iteration). We plug these inputs into Equation (9). Using 1,000 value path simulations, we solve the equation for ξ, and obtain an approximation of the risk-neutral hazard function in year i as $\lambda_t^{Q\ (1)} = \xi_i e^{\beta_0^{(1)}}\left(V_t^{(1)}/B\right)^{\beta_1^{(1)}}$. For years in which we have less than 10 defaults, we estimate ξ using the characteristics of an average firm in the sample.

Step 4. For firm-month observations that correspond to default, we solve Equation (7) for V_t using 1,000 value path simulations per firm.[9] In particular, this yields $V_\tau^{(2)}$, where $t = \tau$ is the month of default. The implied proportional default costs for each defaulted firm are thus $\alpha^{(2)} = 1 - L_\tau/V_\tau^{(2)}$.

Step 5. For all other observations we find $V_t^{(2)}$ from a modification of Equation (7) that uses α instead of L_t as an input:

$$M_t = (1-\alpha)V_t + \alpha\mathbb{E}_t^{\mathbb{Q}}\left[V_T e^{-r(T-t)} e^{-\int_t^T \lambda_u^{\mathbb{Q}}\,\mathrm{d}u}\right]. \tag{10}$$

To do so, for non-defaulting firms we assume that the proportional cost of default is equal to the sample average of $\alpha^{(2)}$. For firm-month observations of defaulting firms prior to default, we use the firm-specific estimates of $\alpha^{(2)}$.

[9] The details of our simulation algorithms are available upon request.

Step 6. We return to step 2, and re-estimate the hazard rate coefficients using $V_t^{(2)}$. We repeat steps 2 through 5 until $\beta_0^{(k)} - \beta_0^{(k-1)}$, $\beta_1^{(k)} - \beta_1^{(k-1)}$, and $V_\tau^{(k)} - V_\tau^{(k-1)}$ all become less than $\epsilon = 10^{-5}$.

2.3.2. The choice of the model inputs

The variables that the model uses as inputs are computed as follows. Prior to default, the market value of the firm M_t is estimated as the total value of all bonds, bank debt, and common and preferred equity, as described in Appendix B. Because of data limitations, these estimates are only available on a monthly basis. Hence, the value of the firm at default, denoted M_τ in Equation (8), is approximated by its value at the end of the last calendar month prior to default. Similarly, the recovery value of the firm L_τ is observed at the end of the calendar month of default. To separate the price reaction to default from the general market movement in the month of default, we subtract the market return from the defaulted firm's return and adjust the recovery value of assets accordingly.

We calculate the volatility of the firm's assets σ as the standard deviation of monthly asset returns for the median firm in the industry, as follows. First, we estimate the standard deviation of each firm's monthly returns, as in Choi and Richardson (2008), excluding post-default months and firms with fewer than 10 consecutive monthly firm value observations. Second, we find the median asset volatility in each of Fama and French's 50 industries. The use of industry rather than firm-specific volatility estimates increases the number of usable observations and reduces noise. Moreover, because the median firm in the industry is typically not distressed, its firm and asset values are very close (see Subsection 4.3 below). Therefore, asset volatility can be estimated as the volatility of the firm, which is much easier to measure, as it does not have to be adjusted for the unobserved expected default costs.

Debt maturity, $T - t$, is the weighted average of maturities of all debt instruments, assuming that all bank debt has a maturity of one year. The face value of debt, B, is the total debt outstanding at the end of the previous fiscal quarter, as reported in Compustat. Finally, the risk-free rate r is the five-year constant-maturity Treasury rate, which is chosen to match the average debt maturity of 5.3 for firms at default.

3. Data description

3.1. Data sources and sample selection

Our estimation procedure is based on debt and equity prices before and after default. Accordingly, our sample construction involves merging data on firm defaults and on bond, loan, and equity prices, as well as accounting information and details of firms' debt structure.[10]

According to the definition of default used by the rating agency Moody's, bond defaults comprise bankruptcy filings, distressed bond exchanges, and missed or delayed bond payments. Thus, default events include both bankruptcies and out-of-court renegotiations with bondholders.[11] Covenant violations, failed bond exchange offers, and renegotiations of bank loans do not alter bondholders' cash flow, and hence are not considered events of default. Although it is an interesting question how costly such actions are in distress, in this study we exclude them in order to focus on measuring the cost of bond defaults.

Our main source of information on defaults is the Default Risk Service (DRS) database distributed by Moody's, which includes all defaults on rated bonds between 1970 and 2010. We amend these data in several ways. For distressed exchanges, DRS reports the date of successful completion as the date of default. Yet the price reaction we want to study is realized at the time of the announcement of the exchange, which DRS does not report. For this reason, we collect information on announcement dates for distressed exchange offers from news reports in Factiva. We also use Factiva to determine the outcomes of defaults not available in DRS. Not all defaults in DRS are independent events, both because firms often default together with their wholly owned subsidiaries, and also because DRS often reports multiple default events within a short period of time.[12] We deal with these issues by focusing on defaults by parent companies only, and by looking at the

[10]To increase the precision of estimated hazard function coefficients, we also use non-defaulting firms included in the Merrill Lynch indices in survival analysis (see the previous subsection). For each month, we estimate values of these firms by applying the same procedures as for defaulting firms.

[11]Moody's defines bond default as "any missed or delayed disbursement of interest and/or principal, bankruptcy, receivership, or distressed exchange, where (i) the issuer offered bondholders a new security or package of securities that amount to a diminished financial obligation (such as preferred or common stock, or debt with a lower coupon or par amount), and (ii) the exchange had the apparent purpose of helping the borrower avoid default"(Keenan, Shtogrin, and Sobehart (1999), p. 10). Standard & Poor's adopts a similar definition; the minor differences pertain to grace period defaults and defaults on preferred stock.

[12]Hotchkiss (1995) finds that many firms emerging from Chapter 11 bankruptcy continue to experience difficulties and often default again.

first default event during our sample period. Finally, we classify defaults as 'formal bankruptcies' and 'out-of-court renegotiations' as follows. If the default event is a missed bond payment or a distressed exchange offer not followed by bankruptcy in the same calendar month, we classify it as a (de-facto) renegotiation of the bond contract.[13] Default events involving a bankruptcy filing in the same calendar month are classified as bankruptcy reorganizations.

Our procedure calls for the estimation of the total market value of the firm's debt and equity in the month prior to and in the month following default. To estimate the market value of bonds, we use monthly bond prices from Merrill Lynch. These are available for constituents of the Merrill Lynch U.S. Investment Grade Index and High Yield Master II Index, available since December 1996. Consequently, our sample period extends from January 1997 to December 2010.[14] Bank loan prices are based on quotes from the LSTA/LPC Mark-to-Market Pricing Database. The database includes monthly secondary-market loan quotes, each obtained from several dealers. These data are available for 69% of defaulting firms, but only for 40% of defaults. For the remaining firms that borrow from banks, we compute loan prices as a function of bond prices, as described in Appendix B. Where available, equity prices are obtained from CRSP. However, firms are often delisted from the exchange some time prior to default. For such cases, we search CapitalIQ for OTC equity price quotes. Also, while accounting information is primarily from quarterly Compustat, we use debt structure data from CapitalIQ to compute the proportion of debt that is owed to banks. For descriptive information on bonds, we use Mergent's Fixed Income Securities Database (FISD). We use bond, loan, and equity prices in conjunction with the debt structure data to estimate market values of total debt and equity at the end of the last calendar month preceding default and in the calendar month of default, as described in Appendix B.

We select the sample as follows. The DRS lists 2,675 defaults between January 1997 and December 2010. We first exclude non-US firms, and retain only defaults by industrial, transportation, and utility companies. We also remove dividend omissions and other events other than public bond defaults. After combining repeated defaults and defaults by firms related through parent-subsidiary relationships, we are left with 727 unique defaulting firms. We manually merge these to FISD and

[13] This definition of renegotiation does not preclude the firm from filing for bankruptcy in the next calendar month. However, such cases are infrequent in the sample. Where renegotiations are followed by bankruptcy, the median time period between default and filing in the sample is 130 days.

[14] Schaefer and Strebulaev (2008) describe the Merrill Lynch data in detail.

the Merrill Lynch data, and obtain bond prices for 514 of them.[15] We manually merge these firms with CRSP and also search for equity prices from CapitalIQ. Share prices are available in the month prior to default for only 240 of the 514 firms. They are missing for many defaulting firms, because a large majority of them are original-issue speculative-grade firms (i.e., highly levered bond issuers), and many of those are post-LBO and other private firms, especially in the latter part of the sample. Finally, we require that bond and equity prices be available both in the month prior and in the month following default, which further reduces the sample to 183 firms.

Our estimation procedure assumes that investors observe the value of assets at default. (In robustness checks, we also allow for the possibility that the value of assets can only be observed with noise, but assume that the noise is unbiased.) However, if the information that investors have about the firm is systematically biased due to fraud, then the price reaction upon default may be primarily due to investors' learning about the fraud, rather than due to default costs. As a result, our estimates could also be biased. For this reason, we exclude firms which allegedly were involved in fraud within two years of default, using the list of large corporate fraud cases compiled by Dyck, Morse and Zingales (2010). After removing Enron and 7 other firms with alleged fraud, we are left with our final sample of 175 defaulted firms.[16]

3.2. Descriptive statistics

The composition of the sample by year of default and broad industry group is shown in Table 1. Although the sample spans 14 years, it is dominated by firms that defaulted during the dot-com crash of the early 2000s. As many as 28.6% of sample firms defaulted in 2001, when the rate of default was one of the highest since the Great Depression (Giesecke et al. (2011)). By contrast, there are relatively few defaults from 2008–2010, as a large majority of defaults by non-financials during this period were by private, post-buyout firms with no traded equity. The sample also covers "calm" periods with low default rates, which allows us to study the effect of macroeconomic conditions on the cost of default. Panel B shows the sample composition by industry. Telecommunications, Wholesale and retail trade, and Consumer goods are the best-represented industries in the sample

[15]The Merrill Lynch indices do not include bonds with par amounts of less than $100 million, nor those with remaining maturity below one year. For firms that have such bonds outstanding at the time of default, we approximate their market value based on prices of other bonds of the same issuer, with similar maturity and seniority.

[16]For the 8 fraud cases, both the price reaction at default and the estimated cost of default are about twice the sample average, consistent with bias in investors' perception of the value of assets. When these firms are included in the sample, the average estimate of the cost of default increases by 1 percentage point.

(19.4%, 15.4%, and 13.7%, respectively).

Table 2 reports the number of sample defaults by the type and outcome of default. As Panel A shows, 37.7% of firms default by filing for bankruptcy, 10.9% complete a distressed bond exchange, and 51.4% miss or delay a bond payment. Panel B reports the incidence of bankruptcy in the calendar month of default. It shows that 12.2% of bond payment defaults are quickly followed by a bankruptcy filing. We refer to the rest of the payment delays and omissions, as well as to successful bond exchanges, as "renegotiations" or "workouts".[17] In untabulated analysis, we find that an additional 75.5% of payment defaults and 15.8% of bond exchanges result in bankruptcy within two years. Overall, as many as 84.6% of bond defaults are followed by bankruptcy either immediately, or after some time within two years of the first default event. Finally, panel C shows eventual outcomes of default, with successful emergence from Chapter 11 being the most common outcome by far.

Table 3 reports general descriptive statistics for defaulted firms. As many as 87.1% of them are original-issue junk-bond issuers, meaning that they had a speculative grade rating when they last issued bonds. The remaining 12.9% are "fallen angels", i.e., firms that are rated investment grade at the time of bond issuance, later downgraded to junk. Thus, although our sample is not limited to highly-levered bond issuers by design, it is nonetheless dominated by them, as is any random sample of firms that default on their bonds. The firms in the sample are naturally larger in size than a typical Compustat firm, because all of them issue public bonds. They appear distressed based on measures of leverage, profitability, and liquidity. About 76% of their debt is in bonds. The weighted average debt maturity is 5.3 years, and the median firm has 2 bonds outstanding.

4. Empirical results

4.1. Asset returns at default

As discussed in Section 2, the cost of default is proportional to (the negative of) the change in the market value of the firm upon default. The observed firm-level price reaction to default announcements is at the heart of our estimates of the cost of default. Table 4 summarizes the firm price

[17] Thus, "workouts" include bond exchange offers, which are bona fide bond contract renegotiations, as well as missed or delayed bond payments not followed by a bankruptcy filing within the same month. Payment defaults reduce the value of creditors' cash flow compared to those specified in the bond contract, and as such constitute a de facto out-of-court debt restructuring.

reaction in the month of default, as well as returns on specific asset classes over the same month. The returns are adjusted by subtracting the return on S&P 500 over the same month.

For a typical firm, the announcement of default results in large value losses. The mean (median) firm-level market-adjusted return in the month of default is –12.2% (–9.9%). The value of the firm falls much more for bankruptcies (by 18.5% on average) than for nonbankruptcy bond defaults (7.3%). Such large price reactions to default imply that, although investors might anticipate it to a certain degree, the announcement of default nonetheless contains a significant element of surprise, which in turn means that pre-default security prices are informative about the continuation value of assets.

The magnitude of the return on individual classes of assets (bonds, loans, and equity) in Table 4 is inversely related to the seniority of the asset: For an average firm, the equity return in the month of default is –21.2%, the loan return is only –4.9%, and the return on bonds falls in between, at –16.2%. This ranking is to be expected, given that payoffs in default are increasing with seniority. For very distressed firms that have not yet defaulted, the value of junior claims such as equity comes mostly from the option value on the firm's recovery, which is greatly reduced in default. In contrast, banks usually have a senior claim on the firm's assets in bankruptcy, and hence loan prices do not fall nearly as much.

Table 4 also shows that asset returns at default are highly heterogenous, ranging from –43.3% to +12.4% between the first and the last deciles. Moreover, adjusted for the market return, the value of the firm *increases* upon default for 30% of defaults, including 18% of bankruptcies and as many as 38% of non-bankruptcy bond defaults. Similarly, Andrade and Kaplan's (1998) estimates of distress costs are also negative for 8 out of 30 firms, or 27% of their sample. A positive price reaction at default means that, even though there may be administrative costs of renegotiation and bankruptcy, the net cost of default is nonetheless negative for these firms. In the absence of default and reorganization, the status quo for such firms likely involves value destruction in ongoing operations, which makes default good news for investors. Consistent with this conjecture, Andrade and Kaplan (1998) find that an important component of costs of financial distress is firms' tendency to delay reorganization, and Davydenko and Rahaman (2011) find that a large number of firms that are worth more dead than alive are able to avoid reorganization or delay it for years, while financing ongoing losses by liquidating assets. For such firms, default may increase value, as documented by

Hotchkiss and Mooradian (1997) for reorganizations involving vulture investors, and by Taillard (2011) for asbestos-related bankruptcies.

4.2. Estimates of the cost of default

Table 5 reports the main results of the paper – our estimates of the cost of default. The mean (median) cost for all bond defaults in the sample is 21.7% (22.1%) of the market value of assets. Default costs are highly heterogenous, varying from –22.5% at the first decile to +65.6% at the tenth decile. Figure 1 shows the distribution of the cost estimates.

The total cost of bankruptcy is on average more than twice as large as the cost of a non-bankruptcy default, 30.5% versus 14.7%. Our estimates of bankruptcy costs are much larger than direct costs of bankruptcy such as lawyers' fees, which are typically found to be within several percentage points of the (book) value of the firm (e.g., Altman (1984); Weiss (1990)).[18] The following factors contribute to the substantial size of these estimates. First, default usually occurs at advanced stages of insolvency, so that the market value of assets just prior to default on average is only 66% of the face value of debt (Davydenko (2011)). This implies that the denominator of our estimated cost-to-value ratio is substantially lower than that in AK's study of firms that are not economically distressed. Second, on average, the value of the bankrupt firm falls by 18.5% in the month of bankruptcy alone, which provides a lower bound on total bankruptcy costs. Third, bankruptcy filings are usually at least partially anticipated by investors, which means that firm values prior to default already incorporate some of the bankruptcy costs, so that the price reaction to the bankruptcy announcement is only a fraction of the total cost of bankruptcy. Indeed, our estimates imply that the observed price reaction is only about half of the total costs of default, while the other half is already incorporated in the pre-default firm value. Overall, our evidence suggests that indirect costs of financial distress are substantially larger than direct costs.

Panel B of Table 5 compares default costs for different outcomes of default. Although the eventual outcome is not known with certainty at the time of default, investors' perceptions of the likely scenarios should affect the price reaction to default, and hence we expect our estimates to be

[18] It should be noted that previous studies express bankruptcy costs as a proportion of the book value of assets, whereas our estimates are normalized by market asset values, which at default average only 45.2% of the book value. Nonetheless, even after adjusting for the differences in the denominator, our estimates of total bankruptcy costs far exceed the direct costs found in aforementioned studies.

correlated with the outcome. The most interesting result in this panel is the contrast between default costs for firms that eventually emerge from bankruptcy (23.1%) and those which are eventually liquidated or sold (41.4%). One interpretation of these estimates is that liquidations are substantially costlier than going-concern reorganizations. An alternative possibility is that the way the firm is reorganized in bankruptcy is endogenous, so that firms that are costlier to reorganize end up in liquidation, while those for which reorganization is feasible are preserved as a going concern and subsequently emerge from bankruptcy. Interestingly, the average estimated net cost of a bond exchange that is not followed by bankruptcy within two years is negative, indicating that such renegotiations are value-increasing overall. Finally, Panel C suggests that there is more heterogeneity in default costs within industries than across industries.

An important question concerns the applicability of estimates obtained from observed defaults to other, non-defaulting firms. For example, Andrade and Kaplan (1998) warn that their estimates of the cost of financial distress may be biased downward, because the HTLs that constitute their sample may have chosen to become highly levered precisely due to their lower-than-usual distress costs. Glover (2011) calibrates a structural model under the assumption that firms follow the trade-off theory of capital structure, and shows that average distress cost among defaulting firms could be substantially lower than the population average.

We investigate this issue by splitting the sample based on the rating that the firm had the last time it issued bonds. Under the self-selection hypothesis, we expect firms with lower costs of default to be more willing to issue high-yield bonds. By contrast, high default-cost firms would reduce their expected losses from default by choosing lower leverage. Table 6 presents evidence to this effect. Consistent with self-selection, firms with lower ratings at the time of bond issuance generally have lower default costs. The average cost for original-issue high-yield firms is 20.2%, which is at the higher end of but close to the 10% to 20% range found by Andrade and Kaplan (1998) for distressed HLTs. However, for investment-grade firms (which in our sample later become fallen angels) the mean cost is 28.8%. Our subsample of fallen angels includes a disproportionate number of regulated utilities, which may default for idiosyncratic reasons and may not be representative of other investment-grade issuers. Panel B shows that when we exclude such firms, the mean (median) estimate for investment-grade firms increases to 31% (30.7%). The correlation between the cost of default and the at-issuance firm rating (coded as 1 for AAA, 2 for AA, etc.) is –0.16%, significant at the 5% level. Thus, firms with low default costs do seem to be over-represented among risky firms that are more likely to default.

These findings have important implications for the interpretation of sample statistics obtained from defaulting firms. Consistent with Andrade and Kaplan's (1998) conjecture, their 10% to 20% range found based on HLTs may have to be revised upwards when evaluating the cost of default for a typical investment-grade firm. Similarly, one has to be selective when applying our sample averages, which are also tilted towards original-issue junk firms, which comprise a large majority of firms observed to default. As a rough guide, if our estimates are to be used to compute ex ante expected default costs for non-distressed investment-grade firms, it may be more appropriate to use average costs of about 30% instead of 20%.

4.3. Ex ante and marginal costs of default

So far, our focus has been on estimating default costs at the time of default. Many applications, such as studies of optimal capital structure, are concerned with the expected risk-adjusted present value of default costs prior to default. In this subsection, we estimate the risk-adjusted expected cost of default as the difference between the market value of assets and the value of the firm in each month up to default. Panel A of Table 7 reports the estimates. Just prior to default, the expected cost for the median firm is 10.0%, which is 45% of the total median cost of default. Thus, by the time a firm declares default, its debt and equity prices already incorporate almost half of its total cost, whereas the other half is realized as a price reaction upon the announcement.

For firms that are still far away from default conditional default, probabilities are small, and expected costs of default are much lower. For investment-grade firms, the sample mean is 2.6% of the asset value, and the median is only 1%. Moreover, even these estimates are in all likelihood biased upwards, because our investment-grade subsample is dominated by future defaulters whose ratings are barely above junk. If we assume that a typical investment-grade firm is like the three A-rated firms in our sample, then the expected default cost is only about 0.5% of the firm value. These numbers are similar to those of Elkamhi et al. (2011), who use AK's estimates of total ex post costs and find that ex ante expected costs are generally below 1% of firm value.

Finally, we also calculate the default-driven marginal cost of debt, defined as the decrease in the firm value when the face value of debt increases by $1. The estimates are reported in Panel B of

Table 7. The mean (median) marginal cost of default is 6.5% (5.6%) for investment-grade firms, and slightly higher for high-yield firms.[19] These estimates are much below marginal corporate tax rates estimated as in Graham (1996), which for the same firm-month observations in our sample average 20.6%. Overall, consistent with the conclusions of Elkamhi et al. (2011) and Graham (2000), our findings suggest that ex ante default costs are too small to offset tax advantages of debt.

4.4. Robustness checks and alternative specifications

In this subsection, we investigate the robustness of our results. First, we vary the setup of the base-case hazard model. Second, we look at whether our estimates are likely to be affected by our use of monthly prices. Finally, we study the implications of the possibility that investors do not observe the market value of assets perfectly prior to default.

4.4.1. Robustness to assumptions of hazard model

A potential concern is that, although rooted in default-predicting studies, our specification for the default hazard is necessarily to some degree ad hoc. We therefore investigate whether our estimates are sensitive to various specific assumptions of our hazard model. In our base-case model, the risk-neutral hazard rate is specified as $\lambda_t^{\mathbb{Q}} = \xi e^{\beta_0 + \beta_1 \log \frac{V_t}{B}}$, where ξ is estimated separately for each year. The estimates under alternative hazard rate specifications are reported in Table 8. First, we use a different functional form of the hazard function, assuming that $\lambda_t^{\mathbb{P}} = \xi e^{\alpha_0 + \alpha_1 \frac{V_t}{B}}$. As can be seen from the second row of Table 8, this change in the specification hardly affects our results. Second, time-varying risk premium estimates may be noisy for years in which only a handful of firms default. Therefore, we also use a common value for ξ, estimated from Equation (9) based on the characteristics of an average firm in the whole sample. The third row of the table shows that our estimates are robust to the assumption of a constant risk premium.

An important assumption of the model is that the ratio of the market value of assets to the face value of debt, V_t/B, can be used as a sufficient statistic for investors' conditional default hazard. Davydenko (2011) shows that, although this ratio is far and away the most powerful predictor of default, balance sheet liquidity also has incremental predictive power. As a result, ignoring liquidity may bias the hazard rate and affect the estimates of the cost of default. We assess the

[19] In a recent paper, van Binsbergen, Graham, and Yang (2010) estimate the all-in marginal cost of debt, and find it to be substantially above the marginal cost of default.

robustness of our estimates by adding liquidity as a second factor that can affect the default hazard: $\lambda_t^{\mathbb{Q}} = \xi e^{\gamma_0 + \gamma_1 \log \frac{V_t}{B} + \gamma_2 QR}$, where QR is the quick ratio (cash and accounts receivable over current liabilities). We assume that to price assets at time t, investors expect QR to stay constant at its current level QR_t until the debt matures. With this modification, liquidity affects the baseline hazard for each firm but not its expected future dynamics. The fourth row of Table 8 reports our estimates for this specification, and shows that the results are very similar to the base case. Although not shown in the table, we find that the main effect of modifying the hazard function in this way is on the cross-sectional distribution of estimated costs rather than on sample averages. Overall, Table 8 shows that our reported estimates are not sensitive to the specification of the default hazard.

4.4.2. Robustness to the use of monthly prices

One potential source of bias in our estimates is the use of monthly prices instead of shorter observation windows around default. If values of defaulting firms decrease systematically prior to default, they will be lower immediately before default than at the end of the previous calendar month. As a result, the asset return in the month of default may be systematically higher than the price reaction to the default announcement, leading to an upward bias in our estimates of default costs.

To investigate this possibility, we compare price reactions and default cost estimates for defaults that happen within the first 14 days of the month with those happening in the remainder of the month. If the decline in firm value since the end of the last calendar month systematically increases observed asset returns over and above the price reaction to the default announcement, we would expect this bias to be higher for defaults that happen late in the month. Ceteris paribus, this would result in higher asset returns and higher estimated default costs for such firms. In untabulated tests, we find no evidence of this effect in the data: The average cost of default for the first half of the month is actually slightly higher than that for the second half (23.6% vs. 20.0%), although the difference is not statistically significant. Thus, our use of monthly data is unlikely to bias our estimates significantly.

4.4.3. Robustness to core model assumptions and the effect of learning from default

Our model assumes that investors know with certainty the continuation value of assets, V. But what if the value of assets is not observed perfectly? In this case, the default announcement itself may cause investors to update their beliefs about the value of assets, most likely downwards. For a given size of the cost of default, this learning-from-default effect could increase the observed price reaction. More generally, the total price reaction is the sum of two parts, one due to learning and the other due to default costs. Hence, our estimates of default costs, which assume away the first component, could be biased upwards. Below, we explore the extent of this potential bias.

To explore the implications of information imperfections and learning from default, we use the framework developed by Duffie and Lando (2001). Their model assumes that investors cannot observe the value of assets perfectly, and receive periodic noisy signals about it. Another important assumption is that there exists a default boundary, that is, a threshold value of assets below which the firm defaults. Calibrating the model to debt and equity prices before and after default, we can back out the implied default boundary as well as investors' noisy signal about the value of assets. The idea behind the use of a boundary model is that for firms that are known to default, the continuation value of assets just prior to default must coincide with the default boundary. This feature allows us to infer the true value of assets at default, even though it is not perfectly observed by investors. An added advantage of using an alternative model specification is that it allows us to ascertain that our estimates are driven primarily by the observed price reaction to default rather than by our modeling choices. The downside of using a boundary-based model is that such models often lack accuracy in the cross-section (Eom et al. (2004), Davydenko (2011)). For this reason, we limit the use of this model to robustness checks.

In the interest of space, we outline the model only briefly.[20] Suppose that there exists a default boundary V_b, known to investors. When the value of assets, V, falls below V_b for the first time, managers (who know both V and V_b) declare default, and the value of the firm changes to L. Furthermore, due to information imperfections, investors do not know V with certainty. Instead, at time t they receive a noisy but unbiased signal $\hat{V}_t$ about V_t, such that:

$$V_t = \hat{V}_t(1 + a\psi), \quad \psi \sim G(\psi),\ E[\psi] = 0,\ E[\psi^2] = 1. \tag{11}$$

Here, ψ is random noise, and the parameter a determines the precision of the signal. Although different distributions can be used, to be specific, in empirical applications we consider the case

[20] Full details are available from the authors upon request.

of a normally distributed noise, and the distribution derived in Duffie and Lando (2001), in which investors observe the firm some time prior to receiving a noisy signal at t.[21]

It is important to note that although investors' signal is unbiased at time t_-, it is not conditionally unbiased given default at time t. Indeed, unless investors are certain at t_- that the firm is about to default, they assign a positive probability to V_t being strictly above the default boundary V_b. If the firm defaults at t, they learn that $V_t = V_b$, which is the lowest possible value for V. Hence, they update their estimate of V_t downward. As a result of this learning from default, even if the cost of default were zero, the value of the firm would fall upon the default announcement to V_b. Since V is bounded from below by its at-default value, V_b, the existence of a default boundary implies the strongest possible effect from learning. Therefore, we can use this model to estimate the maximum bias that learning could induce in our estimates.

Using the well-known distributional results from the first-passage time literature, one can write down closed-form expressions for the value of debt and the firm in the absence of noise. Given creditors' recovery rate R and the recovery value of the firm L, the former will be a function of the "distance to default" V/V_b (e.g., see Leland (1994)), whereas the latter will depend on V as well as on V/V_b.[22] In the presence of noise, debt and firm values are obtained by conditioning on the noisy signal, $\hat{V}$. As a result, the value of debt is a function of $\hat{V}/V_b$, and the value of the firm a function of both $\hat{V}/V_b$ and $\hat{V}$:

$$D_t = D\left(\hat{V}_t/V_b, R, F\right), \quad (12)$$

$$M_t = M\left(\hat{V}_t, \hat{V}_t/V_b, L, F\right). \quad (13)$$

As before, for firms that are observed ex post to default at $t = \tau$, we can compute firm and debt values just prior to default (M_τ and D_τ, respectively), as well as their values immediately after (L and R). Moreover, from the fact that these firms defaulted we know that just prior to default the

[21] In the simplest case, Duffie and Lando (2001) assume that the value of assets is observed perfectly at some prior time 0. Subsequently, at time t investors receive a noisy signal $\hat{V}_t$. They derive the distribution of V_t conditional on a) the signal $\hat{V}_t$, b) the value of assets at time 0, V_0, and c) the fact that the firm did not default between 0 and t. In implementing the Duffie-Lando model, we assume that investors last observe the true value of assets one year before default (our conclusions are not dependent on this one-year assumption.) Because we do not know the value of assets at that time, we infer it from the observed value of the firm, assuming constant proportional default costs. The details are available online and from the authors upon request.

[22] The intuition behind this result is the same as in the intensity model: Because the cash flow from debt in the absence of default is known and fixed, the value of debt depends on the value of assets only through the risk-neutral probability of default, which in such models is a function of V/V_b. By contrast, the value of the firm also depends on V directly; in fact, the two would coincide if default were costless.

true value of assets, V_τ, must have been equal to the default boundary, V_b. This allows us to infer the values of interest as follows. First, we solve Equation (12) for $\hat{V}/V_b$. Next, we substitute this ratio into Equation (13) and solve it for $\hat{V}$. Third, we find the default boundary as $\frac{\hat{V}}{\hat{V}/V_b}$. Finally, the cost of default equals $\alpha \equiv L/V_\tau - 1 = L/V_b - 1$.

We estimate the cost of default implied by this model for different levels of the asset noise, a, using the characteristics of an average firm in the sample. The results are presented in Fig.2. When the noise parameter is zero, the value of assets is observed perfectly, which corresponds to our base-case scenario. The cost of default implied by this model in the absence of noise for the average firm is 21.6%. This estimate is virtually identical to the average cost of 21.7%, which we obtain from our hazard-rate model, despite substantial differences in assumptions. This suggests that our main results are insensitive to how default expectations are inferred.

For positive values of a, part of the observed drop in the value of the firm at default is due to investors' learning that $V_t = V_b$. Nonetheless, the quantitative effect on our estimates is limited. Using a structural model, Korteweg and Polson (2010) estimate the standard deviation of noise in the value of assets for junk firms to be about 5–7%.[23] As Fig. 2 shows, for $a = 7\%$, the cost of default obtained from the Duffie-Lando model is 18.9%, whereas under normal noise it is 18.4%. Given that the presence of a default boundary amplifies the effect of learning in this model, we conclude that the bias in our base-case estimates arising due to information imperfections is unlikely to be substantial.

4.5. Regression results

In this section we explore the determinants of default costs, focusing in particular on the effect of economy-wide and industry distress. Shleifer and Vishny (1992) identify asset fire sales as a potentially important source of financial distress costs. They argue that when a financially distressed firm needs to sell assets, other firms in the same industry are likely to be distressed at the same time. As a result, assets may have to be sold at a discount to deep-pocketed industry outsiders who are not their most efficient users, resulting in value losses.

Given the dearth of empirical estimates of distress costs, the effect of industry conditions has been studied either for particular industries, such as airlines (Pulvino (1998)), or indirectly, by

[23] For bankrupt firms, Gilson, Hotchkiss, and Ruback (2000) find that signals about firm value are unbiased, but noisier.

observing debt recovery rates in default (Acharya et al. (2007)). Andrade and Kaplan (1998) regress their distress cost estimates on industry equity returns and find the correlation between the two to be negative as expected, but not statistically significant.

Using our sample of defaulted firms, which both is much larger and also covers a longer time period, we test the Shleifer-Vishny fire-sale hypothesis by regressing our default cost estimates on measures of economy-wide and industry distress. Our macroeconomic proxies include the annual rate of default on rated bonds reported by Moody's (Ou et al. (2011)), the return on the S&P 500 over the previous year, and the rate of GDP growth. Industry-specific proxies include the median profitability and the median equity return over the previous year for firms in the same 3-digit SIC industry. Under the fire-sale hypothesis, these variables are expected to be negatively correlated with default costs, with the exception of the rate of default, which is expected to be positively correlated. We report regressions for the whole sample, as well as separately for bankruptcies, for which we expect the effect of fire-sales to be particularly strong.

The results are presented in Table 9. For bankruptcy filings, both macroeconomic and industry-specific distress measures are significant and have the predicted sign. When we look at all defaults, the macro factors become insignificant, but industry-specific distress measures remain strongly significant. Importantly, our evidence corroborates the findings of Acharya et al. (2007) that industry-level variables dominate economy-wide distress measures: Columns (11) to (13) show that when both types of proxies are included simultaneously, industry variables remain significant but macroeconomic factors do not. Overall, these findings support Shleifer and Vishny's (1992) prediction that industry distress makes financial distress costlier.

In untabulated tests, we also estimate cross-sectional regressions of default costs on various firm-specific factors used by Andrade and Kaplan (1998), including measures of debt structure complexity, such as the number of outstanding bonds; firm size; and proxies for asset tangibility, such as the ratio of fixed to total assets. Similar to AK, few of these variables are significant determinants of the size of our default cost estimates. The only exception is the ratio of bank to total debt, for which the regression coefficient is negative and generally statistically significant at the 5% level for the bankrupt subsample. The negative correlation with the fraction of bank debt may suggest that banks reduce bankruptcy costs by facilitating restructuring of bankrupt firms. An alternative interpretation is that the composition of debt is endogenous, and that bank prefer borrowers with high-quality collateral, for which default costs are low. Yet another possibility is that the presence of banks may be factored in investors' conditional probability of default and thus in the price reaction to default. Our specification of the default hazard may not fully capture this cross-sectional variation in the probability of default, as well as the effect of other firm-specific variables. Such a dependence could make our cross-sectional comparisons unreliable, which is why we do not attempt to conduct a detailed cross-sectional analysis of the estimated default costs.

5. Conclusions

By combining a novel estimation approach with unique data on market values for a large sample of firms that defaulted between 1997 and 2010, this paper obtains market-implied estimates of the total cost of default, which are not limited to highly levered firms. Although the average cost of default for all firms in the sample is in the region of 20%, the cost for investment-grade firms is closer to 30%. This figure may be more applicable to high- and medium-grade firms than estimates obtained on samples limited to highly-levered firms. As debt prices of defaulting firms become more readily available, our procedure can be applied to firms that are rarely seen to default. This would allow for more precise estimation of the ex ante cost of debt, and thus for better understanding of financing choices that a typical firm faces.

Our estimation procedure adjusts the price reaction to default for the effect of partial anticipation on pre-default prices. It would be interesting to extend our model to allow for more precise cross-sectional identification of the probability of default, which could facilitate cross-sectional comparisons and provide new insights into the nature and the determinants of distress costs. The idea behind the approach potentially can also be applied to other settings in which the market reaction to a corporate event may be dampened by its partial anticipation by investors, such as in the context of mergers and acquisitions. This could advance our understanding of the full effect that such events have on the firm value, beyond the capabilities of traditional event studies.

Appendix A: Derivation of the pricing equation

In what follows, all probabilities and expectations are under the risk-neutral measure $\mathbb{Q}$. We assume that:

1. The market value of the firm's productive assets V_t (i.e., the continuation value of the firm) follows a geometric Brownian motion
$$dV_t = rV_t dt + \sigma V_t dW_t. \tag{A1}$$
2. Default follows a doubly stochastic process: conditional on knowing the history of the risk factor, default time is the first jump of a heterogenous Poisson process with conditional risk-neutral intensity λ_t.
3. The "recovery" value of the firm L_t (i.e., its value in a hypothetical default at time t) is a constant fraction of its continuation value:
$$L_t = (1-\alpha)V_t. \tag{A2}$$
4. The firm is financed with a single discount bond which promises to pay \$1 at maturity. If the firm defaults, the bondholders receive a recovery payment of R at the time of default.

If the firm does not default by the maturity date T, its value at maturity equals the all-equity asset value, V_T. If the firm defaults some time prior to maturity, the value of assets at maturity is $L_T = (1-\alpha)V_T$. Conditional on no prior default, the market value of the firm M_t for $t \leq T$ can be expressed as

$$M_t = e^{-r(T-t)}\mathbb{E}_t\left[V_T 1_{\{\tau \geq T\}} + (1-\alpha)V_T 1_{\{\tau < T\}}\right]. \tag{A3}$$

Rearranging the above equation and using the fact that $V_t = e^{-r(T-t)}\mathbb{E}_t[V_T]$ yields:

$$\begin{aligned} M_t &= e^{-r(T-t)}\mathbb{E}_t\left[V_T 1_{\{\tau \geq T\}} + (1-\alpha)V_T(1-1_{\{\tau \geq T\}})\right] \\ &= e^{-r(T-t)}\mathbb{E}_t\left[(1-\alpha)V_T\right] + \mathbb{E}_t\left[\alpha V_T e^{-r(T-t)}1_{\{\tau \geq T\}}\right] \\ &= (1-\alpha)V_t + \alpha\mathbb{E}_t\left[V_T e^{-r(T-t)}\mathbb{E}_t[1_{\{\tau \geq T\}}]\right] \\ &= (1-\alpha)V_t + \alpha\mathbb{E}_t\left[V_T e^{-r(T-t)}e^{-\int_t^T \lambda_u(V_u)\,\mathrm{d}u}\right], \end{aligned} \tag{A4}$$

which is Equation (10) of the main text. The last step uses the fact that at time t we know $\tau > t$ and conditional on the information up to T, the default process is an non-homogeneous Poisson process stopped at its first jump. Hence, $\mathbb{E}_t[1_{\{\tau \geq T\}}]$ is the non-default probability and we have

$$\mathbb{E}_t[1_{\{\tau \geq T\}}] = e^{-\int_t^T \lambda_u(V_u)\,\mathrm{d}u}. \tag{A5}$$

Since $L_t = (1-\alpha)V_t$, Equation (A4) can be re-arranged as

$$M_t = L_t + (V_t - L_t)\mathbb{E}_t\left[\frac{V_T e^{-r(T-t)}}{V_t}e^{-\int_t^T \lambda_u(V_u)\,\mathrm{d}u}\right], \tag{A6}$$

which is Equation (7) of the main text.

Similarly, the value of debt can be found as

$$\begin{aligned} D_t &= \mathbb{E}_t\left[e^{-r(T-t)}1_{\{\tau > T\}} + Re^{-r(\tau-t)}1_{\{\tau \leq T\}}\right] \\ &= e^{-r(T-t)}\mathbb{E}_\tau\left[e^{-\int_t^T \lambda(V_u)\,\mathrm{d}u}\right] + \int_t^T \mathbb{E}_\tau\left[Re^{-\int_t^\tau [r+\lambda(V_u)]\,\mathrm{d}u}\lambda(V_\tau)\right]\mathrm{d}\tau \end{aligned} \tag{A7}$$

where we have used the fact that the conditional distribution of the default time in the doubly-stochastic framework is given by

$$f_\tau(\tau|V_t) = e^{-\int_t^\tau \lambda(V_u)\,\mathrm{d}u}\lambda(V_\tau). \tag{A8}$$

This yields Equation (9) of the main text.

Appendix B: Computing the market value of the firm

For each sample firm, we estimate monthly market values of the firm as the sum of market values of bonds, bank debt, and equity. The firm's bond structure is inferred from the history of outstanding bond amounts in the FISD database for each bond issued by the firm and its wholly owned subsidiaries. The market value of bonds included in the Merrill Lynch indices (MLI) is calculated by multiplying the currently outstanding amount by the bond price. Bonds with remaining maturity of less than one year or face value under $100 million are not included in the MLI. The market value of these bonds is calculated assuming that their yield equals the weighted-average yield of all quoted bonds of the same issuer on each date. If in any given month no bond prices are available for the firm, the firm-month observation is excluded from the sample.

Estimates of bank loan prices are based on quotes provided by the LSTA/LPC Mark-to-Market Pricing service, available from May 1998. On average, for each loan-month, the data base provides a mean price quote from 3 dealers. When there are several loans outstanding for a firm, we use their mean price, resulting in 7.5 dealer quotes per bank debt price on average (the median is 4). LSTA/LPC quotes are available for 69% of the sample firms, but only for 40% of firm-months that correspond to default. For firm-months not included in this database, the market price of bank debt is estimated as a quadratic function of the weighted-average bond price, as follows:

$$\begin{array}{rcccccc} P_{bank} & = & 40.18 & + & 1.045 \times P_{bond} & - & 0.00461 \times P_{bond}^2, \\ & & (14.2) & & (12.9) & & (-8.45) \end{array}$$

where P_{bank} and P_{bond} are average loan and bond prices in cents on the dollar, respectively, and t-statistics adjusted for firm clustering are reported in parentheses. The quadratic term controls for nonlinearities that arise due to different priorities of loans and bonds in bankruptcy. The regression produces an R^2 of 75.5% and is not substantially improved by the inclusion of additional firm-specific or macroeconomic controls.

Preferred equity is rarely important in the sample; its par value is below 5% of the face value of debt for 79% of firms at default. Preferred stock is worth little in default, and thus its par value is likely to vastly overstate its market value in distress. Varma (2003) finds mean recovery rates for preferred stock of 15.3%, compared with 36.1% for senior unsecured bonds (the most common bond type by far). Hence, to approximate the market value of preferred stock, we assume that its price relative to par is equal to the constant fraction 15.3/36.1=0.424 of the firm's current bond price. Sensitivity analysis shows that this approximation has a negligible effect on our estimates.

For the median firm at default, bonds and bank loans together constitute about 98% of total debt. Firms may make use of other types of borrowing, such as commercial paper, mortgages, and project finance debt. Because commercial paper (rare in the sample) has short maturity and is backed by credit lines, and most other debt types are secured, we assume that all such debt obligations are similar to bank debt and have the same price-to-par ratio. These types of debt are not frequently used by risky firms that dominate our sample, so this approximation affects only a small fraction of the firms.

Where available, we use equity prices from CRSP. However, firms are occasionally delisted from the stock exchange and disappear from CRSP some time before default. For these cases, we use OTC equity prices from CapitalIQ. Finally, we also rely CapitalIQ for the details of the firms' debt structure, including the split of debt between bonds and bank loans. The market value of the firm is then computed as the weighted average of the values of common and preferred stock, and all outstanding debt instruments.

References

Acharya, Viral V., Sreedhar T. Bharath, and Anand Srinivasan, 2007, Does industry-wide distress affect defaulted firms? – Evidence from creditor recoveries, *Journal of Financial Economics* 85, 787–821.

Almeida, Heitor, and Thomas Philippon, 2007, The risk-adjusted cost of financial distress, *Journal of Finance* 62, 2557–2586.

Altman, Edward I., 1968, Financial ratios, discriminant Analysis and the prediction of corporate bankruptcy, *Journal of Finance* 23, 589–609.

Altman, Edward I., 1984, A further empirical investigation of the bankruptcy cost question, *Journal of Finance* 39, 1067–1085.

Andrade, Gregor, and Steven N. Kaplan, 1998, How costly is financial (not economic) distress? Evidence from highly leveraged transactions that became distressed, *Journal of Finance* 53, 1443–1493.

Asquith, Paul, Robert Gertner, and David Scharfstein, 1994, Anatomy of financial distress: An examination of junk bond issuers, *Quarterly Journal of Economics* 109, 625–658.

Berndt, Antje, Rohan Douglas, Darrell Duffie, Mark Ferguson, and David Schranz, 2008, Measuring default risk premia from default swap rates and EDFs, Working paper.

Bhamra, Harjoat S., Lars-Alexander Kuehn, and Ilya A. Strebulaev, 2010a, The levered equity risk premium and credit spreads: A unified framework, *Review of Financial Studies* 23, 645–703.

Bhamra, Harjoat S., Lars-Alexander Kuehn, and Ilya A. Strebulaev, 2010b, The aggregate dynamics of capital structure and macroeconomic risk, *Review of Financial Studies* 23, 4187–4241.

van Binsbergen, Jules H., John R. Graham, and Jie Yang, 2010, The cost of debt, *Journal of Finance* 65, 2089–2136.

Black, Fischer, and John C. Cox, 1976, Valuing corporate securities: Some effects of bond indenture provisions, *Journal of Finance* 31, 351–367.

Bris, Arturo, Ivo Welch, and Ning Zhu, 2006, The costs of bankruptcy, *Journal of Finance* 61, 1253–1303.

Brown, Stephen J., and Jerold B. Warner, 1985, Using daily stock returns: The case of event studies, *Journal of Financial Economics* 14, 3-31.

Choi, Jaewon, and Matthew Richardson, 2008, The volatility of the firm's assets, Working paper, New York University.

Clark, Truman A., and Mark I. Weinstein, 1983, The behavior of common stock of bankrupt firms, *Journal of Finance* 38, 489–504.

Cutler, David M., and Lawrence H. Summers, 1988, The costs of conflict resolution and financial distress: evidence from the Texaco-Pennzoil litigation, *RAND Journal of Economics* 19, 157–172.

Davydenko, Sergei A., 2011, When do firms default? A study of the default boundary, Working paper, University of Toronto.

Davydenko, Sergei A., and Muji R. Rahaman, 2011, Excessive continuation and the costs of flexibility in financial distress, Working paper, University of Toronto.

Driessen, Joost, 2005, Is default event risk priced in corporate bonds?, *Review of Financial Studies* 18, 165–195.

Duffie, Darrell, and David Lando, 2001, Term structures of credit spreads with incomplete accounting information, *Econometrica* 69, 633–664.

Duffie, Darrell, and Kenneth J. Singleton, 1999, Modeling term structures of defaultable bonds, *Review of Financial Studies* 12, 687–720.

Dyck, Alexander, Adair Morse, and Luigi Zingales, 2010, Who blows the whistle on corporate fraud?, *Journal of Finance* 65, 2213–2253.

Eisdorfer, Assaf, 2008, Empirical evidence of risk shifting in financially distressed firms, *Journal of Finance* 63, 609–637.

Elkamhi, Redouane, Jan Ericsson, and Christopher A. Parsons, 2011, The costs of financial distress and the timing of default, *Journal of Financial Economics*, forthcoming.

Eom, Young Ho, Jean Helwege, and Jing-zhi Huang, 2004, Structural models of corporate bond pricing: An empirical analysis, *Review of Financial Studies* 17, 499–544.

Franks, Julian R., and Sergey V. Sanzhar, 2006, Evidence on Debt Overhang from Distressed Equity Issues, Working paper, London Business School.

Franks, Julian R., and Walter N. Torous, 1989, An empirical investigation of U.S. firms in reorganization, *Journal of Finance* 44, 747–769.

Giesecke, Kay, 2006, Default and information, *Journal of Economic Dynamics & Control* 30, 2281–2303.

Giesecke, Kay, Francis A. Longstaff, Stephen M. Schaefer, and Ilya A. Strebulaev, 2011, Corporate Bond Default Risk: A 150-Year Perspective, *Journal of Financial Economics* 102, 233–250.

Gilson, Stuart C., 1997, Transactions costs and capital structure choice: Evidence from financially distressed firms, *Journal of Finance* 52, 161–196.

Gilson, Stuart C., Edith S. Hotchkiss, and Richard S. Ruback, 2000, Valuation of bankrupt firms, *Review of Financial Studies* 13, 43–74.

Glover, Brent, 2011, The expected cost of default, Working paper, Carnegie Mellon University.

Gorbenko, Alexander S., and Ilya A. Strebulaev, 2010, Permanent vs. Transitory Shocks: Explaining Corporate Finance Policies, *Review of Financial Studies* 23, 2591–2647.

Graham, John R., 1996, Proxies for the marginal tax rate, *Journal of Financial Economics* 42, 187–221.

Graham, John R., 2000, How big are the tax benefits of debt?, *Journal of Finance* 55, 1901–1941.

Hennessy, Christopher R., and Toni M. Whited, 2007, How Costly is External Financing? Evidence from a Structural Estimation, *Journal of Finance* 62, 1705–1745.

Hotchkiss, Edith S., and Robert M. Mooradian, 1997, Vulture investors and the market for control of distressed firms, *Journal of Financial Economics* 43, 401–432.

Hotchkiss, Edith S., 1995, Postbankruptcy performance and management turnover, *Journal of Finance* 50, 3–21.

Huang, Jing-zhi, and Ming Huang, 2003, How much of the corporate-treasury yield spread is due to credit risk?, Working paper, Pennsylvania State University.

Hull, John, Mirela Predescu, and Alan White, 2005, Bond prices, default probabilities, and risk premiums, *Journal of Credit Risk* 1, 53–60.

Jarrow, Robert A., David Lando, and Stuart M. Turnbull, 1997, A Markov model for the term structure of credit risk spreads, *Review of Financial Studies* 10, 481–523.

Jarrow, Robert A., and Philip Protter, 2004, Structural versus reduced form models: A new information based perspective, *Journal of Investment Management* 2, 1–10.

Jensen, Michael C., and William H. Meckling, 1976, Theory of the firm: Managerial behavior, agency costs and ownership structure, *Journal of Financial Economics* 3, 305–360.

Kalbfleisch, John D., and Ross L. Prentice, 2002, The Statistical Analysis of Failure Time Data, 2nd Edition, Wiley, New York.

Keenan, Sean C., Igor Shtogrin, and Jorge Sobehart, 1999, Historical default rates of corporate bond issuers, 1920–1988, Working paper, Moody's Investors Service.

Korteweg, Arthur G., 2010, The net benefits to leverage, *Journal of Finance*, forthcoming.

Korteweg, Arthur G., and Nicholas Polson, 2010, Corporate credit spreads under uncertainty, Working paper, Stanford University.

Lang, Larry H.P., and Rene M. Stulz, 1992, Contagion and competitive intra-industry effects of bankruptcy announcements, *Journal of Financial Economics* 32, 45–60.

Leland, Hayne E., 1994, Corporate debt value, bond covenants, and optimal capital structure, *Journal of Finance* 49, 1213–1252.

Madan, Dilip B., and Haluk Unal, 1998, Pricing the risks of default, *Review of Derivatives Research* 2, 121–160.

Mello, Antonio S., and John E. Parsons, 1992, Measuring the Agency Cost of Debt, *Journal of Finance 47*, 1887–1904.

Merton, Robert C., 1974, On the pricing of corporate debt: The risk structure of interest rates, *Journal of Finance 29*, 449–470.

Miao, Jianjun, 2005, Optimal capital structure and industry dynamics, *Journal of Finance*, 60, 2621–2660.

Molina, Carlos A., 2005, Are firms underleveraged? An examination of the effect of leverage on default probabilities, *Journal of Finance* 60, 1427–1459.

Moyen, Nathalie, 2007, How big is the debt overhang problem?, *Journal of Economic Dynamics & Control* 31, 433–472.

Myers, Stewart, 1977, Determinants of corporate borrowing, *Journal of Financial Economics* 3, 799–819.

Opler, Tim, and Sheridan Titman, 1994, Financial distress and corporate performance, *Journal of Finance* 49, 1015–1040.

Ou, Sharon, David Chlu, and Albert Metz, 2011, Corporate default and recovery rates, 1920–2010, Special Comment, Moody's Investors Service.

Parrino, Robert, and Michael S. Weisbach, 1999, Measuring investment distortions arising from stockholder-bondholder conflicts, *Journal of Financial Economics* 53, 3–42.

Pulvino, Todd C., 1998, Do asset fire sales exist: An empirical investigation of commercial aircraft sale transactions, *Journal of Finance* 53, 939–978.

Schaefer, Stephen M., and Ilya A. Strebulaev, 2008, Structural models of credit risk are useful: Evidence from hedge ratios on corporate bonds, *Journal of Financial Economics* 90, 1–19.

Shreve, Steven E., 2004, Stochastic Calculus for Finance II: Continuous-Time Models, Springer, New York.

Shleifer, Andrei, and Robert Vishny, 1992, Liquidation values and debt capacity: A market equilibrium approach, *Journal of Finance*, 47, 1343–1366.

Shumway, Tyler, 2001, Forecasting bankruptcy more accurately: A simple hazard model, *Journal of Business* 74, 101–124.

Taillard, Jérôme P., 2011, Thriving in the midst of financial distress? An analysis of firms exposed to asbestos litigation, Working paper, Boston College.

Titman, Sheridan, 1984, The effect of capital structure on a firm's liquidation decision, *Journal of Financial Economics* 13, 137–151.

Varma, Praveen, 2003, Recovery rates on defaulted corporate bonds and preferred stocks, Moody's Corp.

Warner, Jerold B., 1977, Bankruptcy costs: some evidence, *Journal of Finance* 32, 337–347.

Weiss, Lawrence A., 1990, Bankruptcy resolution: direct costs and violation of priority of claims, *Journal of Financial Economics* 27, 285–314.

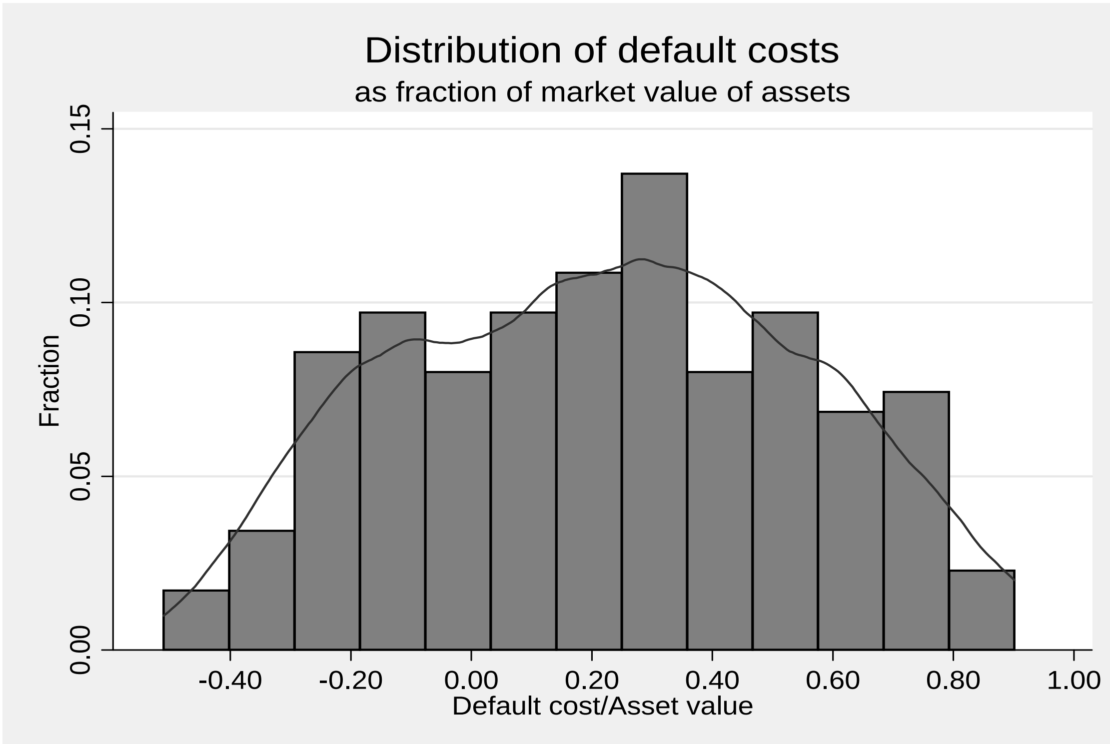

Fig. 1. This graph presents the histogram of default cost estimates in the sample.

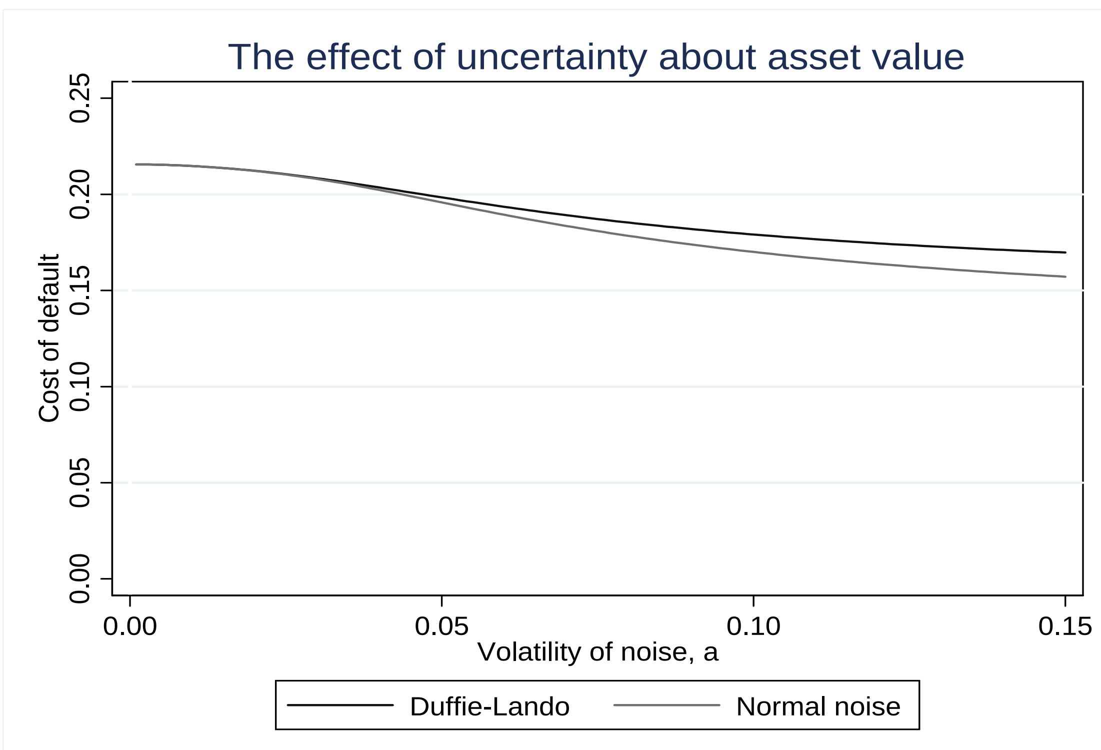

Fig. 2. This graph shows estimates of the cost of default for the average firm implied by the boundary model as a function of the standard deviation of noise.

Table 1: Sample composition

This table reports the number of defaults in the sample, by year of default (Panel A) and by industry (Panel B). The sample consists of non-financial U.S. public firms that defaulted on their public bonds between January 1997 and December 2010. Default events are bond payment omissions (including those rectified within the grace period), distressed bond exchanges, and bankruptcy filings.

	No. of defaults	% of sample
Panel A: Number of defaults by year		
1997	5	2.9%
1998	13	7.4%
1999	24	13.7%
2000	19	10.9%
2001	50	28.6%
2002	28	16.0%
2003	19	10.9%
2004	5	2.9%
2005	7	4.0%
2008	2	1.1%
2009	2	1.1%
2010	1	0.6%
All	175	100%
Panel B: Industry composition		
Consumer goods	24	13.7%
Business equipment	7	4.0%
Steel	9	5.1%
Other manufacturing	21	12.0%
Telecommunications	34	19.4%
Wholesale and retail trade	27	15.4%
Transportation	8	4.6%
Energy & Utilities	14	8.0%
Other industries	31	17.7%
All	175	100%

Table 2: Default events, bankruptcy, and outcomes

This table reports the incidence of bankruptcy filings and the eventual outcome of default for sample firms, by the type of the first default event (bankruptcy filing, payment omission, or distressed bond exchange). Panel A gives the total number of defaults by the first default event. Panel B reports whether, following default, there was a bankruptcy filing in the same calendar month. Panel C reports the eventual outcomes of default.

	First default event			
	Bankruptcy filing	Distressed exchange	Payment default	Total
Panel A: First default events				
All defaults	66	19	90	175
	37.7%	10.9%	51.4%	100%
Panel B: Renegotiations and immediate bankruptcies				
Bankruptcy	66	-	11	77
	100%		12.2%	44.0%
Renegotiation	-	19	79	98
		100%	87.8%	56.0%
Panel C: Eventual outcomes of default				
Creditors paid in full	-	-	7	7
			7.8%	4.0%
Bond exchange completed	-	13	2	15
		68.4%	2.2%	8.6%
Emerged from bankruptcy	55	4	69	128
	83.3%	21.1%	76.7%	73.1%
Acquired or liquidated	10	2	11	23
	15.2%	10.5%	12.2%	13.1%
Other	1	-	1	2
	1.5%		1.1%	1.1%

Table 3: Descriptive statistics

This table reports descriptive statistics for firms at default. *Original-issue junk firms* are those which had a speculative-grade rating at the most recent bond issuance. *Book leverage* is the ratio of total debt to book assets. *Market leverage* is the ratio of the market value of debt to the market value of the firm. *EBIT* is the sum of pretax income and interest expenses. *Interest coverage ratio* is the ratio of EBITDA, calculated as the sum of pretax income, interest expense, and depreciation, to interest expense. *Industry asset volatility* is the annualized median standard deviation of monthly firm returns in the industry, using Fama-French's 50 industries. *Quick ratio* is the sum of cash and accounts receivable divided by current liabilities. *Current ratio* is the ratio of current assets to current liabilities. *Debt maturity* is the weighted average of maturities of all outstanding debt instruments, assuming that all bank debt has a maturity of one year. *Debt interest rate* is interest expense in the last quarter divided by the average outstanding debt in that quarter. Accounting variables are observed at the end of the last fiscal quarter preceding default. All other variables are as of the end of the last calendar month preceding default.

	Mean	Median	Std.dev.	N
% Original-issue junk firms	87.1%			171
Total assets ($ Mil.)	3,271	1,034	7,856	175
Book leverage	0.803	0.747	0.383	175
Market leverage	0.851	0.906	0.161	175
Sales/Book assets	0.240	0.190	0.237	172
EBIT/Total assets	-0.109	-0.024	0.282	163
% Negative net income	90.1%			172
Interest coverage ratio	-3.151	-0.169	7.343	162
Industry asset volatility	0.276	0.279	0.048	175
Quick ratio	0.575	0.397	0.616	164
Current ratio	0.977	0.770	0.830	164
Short-term/Total debt	0.200	0.029	0.309	173
Debt maturity	5.27	4.63	2.90	175
Debt interest rate	8.8%	8.8%	2.3%	172
Bonds/Total debt	0.762	0.842	0.258	175
Number of bond issues	4.63	2.00	10.41	175

Table 4: Asset returns at default and debt recovery rates

This table reports statistics on market-adjusted returns for different asset classes in the month of default, as well as debt recovery rates. *Total return* is the weighted-average return on equity, loans, and bonds, in the calendar month of default, less the return on S&P 500. *Debt return* is the weighted-average return on loans and bonds, calculated similarly. Returns on bonds, bank debt, and equity are also adjusted for the market return. *Debt recovery rate* is the weighted-average market value of all of the firm's outstanding debt instruments at the end of the calendar month of default, expressed as a proportion of the face value of total debt.

	Mean	Median	Std.dev.	10%	90%	Return > 0
Panel A: All defaults, N=175						
Total return	-12.2%	-9.9%	22.4%	-43.3%	12.4%	0.30
Equity return	-21.2%	-24.7%	43.2%	-72.7%	23.7%	0.24
Debt return	-10.8%	-7.8%	22.6%	-39.9%	11.8%	0.30
Bond return	-16.2%	-13.8%	28.6%	-55.0%	14.0%	0.29
Bank debt return	-4.9%	-3.7%	13.1%	-19.8%	7.9%	0.32
Debt recovery rate	43.5%	40.8%	22.7%	14.0%	76.0%	
Panel B: Renegotiations, N=99						
Total return	-7.3%	-6.3%	22.1%	-37.9%	17.4%	0.38
Equity return	-11.5%	-14.7%	43.5%	-57.3%	35.3%	0.29
Debt return	-6.5%	-4.7%	22.9%	-35.6%	17.6%	0.38
Bond return	-8.7%	-6.8%	26.9%	-44.4%	26.9%	0.38
Bank debt return	-4.0%	-3.2%	12.1%	-19.2%	7.9%	0.38
Debt recovery rate	47.4%	44.8%	22.9%	18.4%	79.4%	
Panel C: Bankruptcy filings, N=76						
Total return	-18.5%	-14.5%	21.3%	-51.0%	6.4%	0.18
Equity return	-33.7%	-33.9%	39.7%	-78.0%	12.4%	0.17
Debt return	-16.3%	-12.7%	21.0%	-50.5%	8.0%	0.20
Bond return	-26.0%	-24.6%	28.0%	-64.6%	4.8%	0.16
Bank debt return	-6.0%	-3.9%	14.3%	-23.0%	9.1%	0.24
Debt recovery rate	38.4%	36.7%	21.4%	11.4%	66.6%	

Table 5: Estimates of the costs of default

This table reports estimates of the costs of default, expressed as a proportion of the market value of assets at the end of the last calendar month prior to default. Panel A reports the statistics for all sample firms, and separately for firms that do and do not file for bankruptcy in the calendar month of default. Panels B and C report default costs by industry and by the eventual outcome of default.

	Mean	Median	Std.dev.	10%	90%	N
Panel A: By type of default						
All defaults	21.7%	22.1%	33.0%	-22.5%	65.6%	175
Renegotiations	14.7%	11.4%	33.9%	-27.2%	69.2%	98
Bankruptcy filings	30.5%	30.7%	29.8%	-11.6%	65.6%	77
Panel B: By outcome of default						
Acquired or liquidated	41.4%	43.1%	27.1%	9.4%	76.7%	23
Emerged from bankruptcy	23.1%	25.4%	31.8%	-22.5%	64.3%	128
Bond exchange completed	-7.4%	-10.3%	25.6%	-44.4%	28.8%	15
Creditors paid in full	5.2%	-8.5%	34.9%	-27.2%	74.7%	7
Panel C: By industry						
Consumer goods	23.6%	27.6%	28.7%	-27.2%	50.2%	24
Business equipment	9.7%	-5.9%	51.9%	-51.1%	83.6%	7
Steel	48.5%	44.8%	23.9%	6.9%	76.7%	9
Other manufacturing	24.9%	30.6%	32.3%	-21.8%	59.5%	21
Telecommunications	18.4%	16.5%	35.3%	-17.9%	74.1%	34
Wholesale and retail trade	27.5%	30.7%	35.9%	-28.2%	69.2%	27
Transportation	19.5%	13.4%	22.9%	-5.2%	61.7%	8
Energy & Utilities	25.1%	16.9%	22.1%	8.4%	54.8%	14
Other industries	10.4%	6.3%	32.2%	-22.5%	50.1%	31

Table 6: Default costs by original rating

This table reports estimates of the costs of default, expressed as a proportion of the market value of assets at the end of the last calendar month prior to default, by firm rating as of the date of the most recent bond issuance. Panel A reports statistics for all firms. Panel B excludes Energy & Utility firms. *Fallen angels* and *original-issue junk firms* are firms that were rated investment-grade and speculative grade, respectively, at the time of the last bond sale preceding default.

	Mean	Median	Std.dev.	10%	90%	N
Panel A: All firms						
A	26.0%	26.0%	40.8%	-2.8%	54.8%	2
BBB	29.1%	22.6%	23.6%	-2.0%	57.5%	20
BB	25.9%	27.6%	26.9%	-29.9%	60.5%	18
B	21.4%	25.6%	33.4%	-22.4%	65.0%	100
CCC	15.9%	8.3%	37.1%	-28.8%	76.2%	28
CC	-16.6%	-20.6%	12.3%	-26.3%	-2.8%	3
All fallen angels	28.8%	22.6%	24.2%	-2.8%	55.5%	22
All original-issue junk firms	20.2%	21.9%	33.5%	-24.9%	65.6%	149
Panel B: Excluding Utilities						
BBB	31.0%	30.7%	25.5%	-9.5%	59.5%	15
BB	25.9%	27.6%	26.9%	-29.9%	60.5%	18
B	21.1%	25.7%	34.1%	-22.5%	64.3%	93
CCC	15.9%	8.3%	37.1%	-28.8%	76.2%	28
CC	-16.6%	-20.6%	12.3%	-26.3%	-2.8%	3
All fallen angels	31.0%	30.7%	25.5%	-9.5%	59.5%	15
All original-issue junk firms	19.9%	21.9%	33.9%	-24.9%	64.3%	142

Table 7: Ex ante (expected) and marginal costs of default

This table reports expected and marginal default costs prior to default. Expected default cost is the difference between the continuation value of assets and the market value of the firm, expressed as a proportion of the market value of assets. Marginal default cost is the decrease in the value of the firm when the face value of debt increases by $1. For each firm, the sample consists of all firm-month observations since December 1996 and up to default. The reported statistics are calculated using firm means for each firm-rating combination.

	Mean	Median	Std.dev.	10%	90%	N
Panel A: Expected costs						
A	0.5%	0.6%	0.6%	-0.2%	1.0%	3
BBB	3.0%	1.2%	4.4%	-0.6%	8.9%	18
BB	4.5%	2.0%	8.4%	-1.3%	9.5%	44
B	4.6%	4.6%	9.2%	-5.7%	15.9%	135
CCC	10.3%	9.0%	19.1%	-14.4%	35.2%	122
CC	12.8%	7.3%	22.4%	-11.9%	47.7%	40
C	19.6%	6.0%	28.2%	1.3%	72.8%	6
Firms at default	14.0%	10.0%	20.9%	-10.2%	42.8%	171
All investment grade	2.6%	1.0%	4.1%	-0.2%	7.5%	21
All high yield	7.8%	4.8%	15.8%	-8.1%	28.9%	347
Panel B: Marginal costs						
A	8.4%	8.4%	5.1%	4.8%	12.1%	2
BBB	6.2%	5.6%	5.9%	1.0%	13.6%	17
BB	8.8%	7.9%	5.3%	2.5%	17.1%	43
B	8.7%	9.3%	6.3%	1.0%	15.1%	128
CCC	8.6%	8.9%	6.7%	0.4%	16.3%	116
CC	8.2%	8.5%	5.9%	0.6%	15.9%	39
C	10.2%	10.0%	6.4%	1.7%	18.7%	6
Firms at default	8.7%	8.7%	7.2%	0.5%	16.3%	157
All investment grade	6.5%	5.6%	5.7%	1.0%	13.6%	19
All high yield	8.6%	8.9%	6.3%	1.1%	16.1%	332

Table 8: Robustness to model assumptions

This table reports estimates of the cost of default for different functional forms of the default hazard, $\lambda_t^{\mathbb{P}}$, and different assumptions about the default risk premium, ξ. The base-case uses $\lambda_t^{\mathbb{P}} = e^{\beta_0+\beta_1 \log \frac{V_t}{B}}$ and year-specific risk premium estimates. In the second row, the hazard function is assumed to be exponential: $\lambda_t^{\mathbb{P}} = e^{\alpha_0+\alpha_1 \frac{V_t}{B}}$. The third row uses a constant risk premium, ξ, estimated based on the characteristics of an average firm in the sample. In the fourth row, the hazard function is assumed to depend on balance sheet liquidity, as follows: $\lambda_t^{\mathbb{P}} = e^{\gamma_0+\gamma_1 \log \frac{V_t}{B}+\gamma_2 QR}$, where QR is the quick ratio, defined as cash and receivables over current liabilities, assumed constant for each firm. The number of observations in each row is 175.

	Mean	Median	Std.dev.	10%	90%
Base-case model	**21.7%**	**22.1%**	**33.0%**	**-22.5%**	**65.6%**
Exponential hazard function	21.7%	22.4%	33.2%	-23.0%	68.7%
Constant risk premium	22.2%	23.7%	33.2%	-22.3%	65.7%
Liquidity in default hazard	21.1%	22.2%	33.7%	-26.0%	66.3%

Table 9: The effect of macroeconomic and industry conditions and the cost of default

This table presents OLS regressions of the cost of default. Column (1) through (5) are for all firms, and columns (6) to (13) are for bankruptcies only. *Default rate* is the proportion of rated bonds that default in each calendar year, as reported by Moody's. *Market return* is the return on S&P 500 over the year preceding default. *GDP* is the percentage increase in the GDP in the year of default. *Industry profitability* is the profit margin of the median firm in the same 3-digit SIC industry. *Industry return* is the equity return of the median firm in the same 3-digit SIC industry. Values of t-statistics are reported in parentheses. Coefficients marked ***, **, and * are significant at the 1%, 5%, and 10% significance level, respectively.

	All defaults					Bankruptcies							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)	(12)	(13)
Const.	0.20***	0.22***	0.20***	0.23***	0.22***	0.17**	0.31***	0.42***	0.30***	0.30***	0.20**	0.31***	0.37***
	(3.16)	(8.66)	(3.58)	(9.05)	(8.59)	(2.05)	(9.29)	(5.62)	(9.03)	(9.19)	(2.37)	(9.10)	(4.75)
Default rate	0.005					0.054*					0.042		
	(0.22)					(1.77)					(1.30)		
Market return		0.0003					-0.004**					-0.001	
		(0.26)					(-2.02)					(-0.63)	
GDP growth			0.004					-0.024*					-0.014
			(0.40)					(-1.72)					(-0.96)
Industry profitability				-0.008***					-0.010***		-0.009**		
				(-2.67)					(-2.78)		(-2.18)		
Industry return					-0.002**					-0.0026***		-0.002**	-0.002**
					(-2.41)					(-2.96)		(-2.35)	(-2.64)
N	175	175	175	167	167	77	77	77	75	72	75	72	72
R^2	0.000	0.000	0.001	0.042	0.034	0.040	0.051	0.038	0.096	0.111	0.117	0.116	0.123