

Calibrating Credit Risk Dynamics in Private Infrastructure Debt

FRÉDÉRIC BLANC-BRUDE, MAJID HASAN,
AND TIMOTHY WHITTAKER

**FRÉDÉRIC
BLANC-BRUDE**

is the director of EDHEC Infrastructure Institute and EDHEC Asia Pacific at EDHEC Business School in Singapore.
frederic.blanc-brude@edhec.edu

MAJID HASAN

is an associate research director of EDHEC Infrastructure Institute at EDHEC Business School in Singapore.
majid.hasan@edhec.edu

TIMOTHY WHITTAKER

is an associate research director of EDHEC Infrastructure Institute at EDHEC Business School in Singapore.
tim.whittaker@edhec-risk.com

This article proposes an empirical evaluation of credit risk in private infrastructure project finance (PF) debt. Despite its importance in the practice of banking, PF has received little attention in the finance literature. But a growing interest among asset owners in search of yield and diversification in private debt in general and infrastructure debt in particular suggests that documenting the risk profile of private debt instruments has become an increasingly relevant empirical question.

Existing empirical research on private infrastructure project credit risk relies on the possibility of observing actual events of default among a population in order to measure default frequencies, that is, so-called reduced form models that represent credit risk as an exogenous random variable or hazard rate, which have been produced by credit rating agencies (see, e.g., Moody's Investors Service [2017]).¹

These studies are built using data pooled from banks and can include hundreds or thousands of loans originated by project finance lenders over the past two decades or more. As such, they usually claim to be representative of the underlying population of originated loans (Moody's [2017, p. 2]). While this may be the case (assuming no reporting biases among lenders), the related claim that such datasets can provide an

unbiased view of the credit risk of private infrastructure loans is not self-evident.

In making these claims, rating agencies assume that a representative dataset of credit instruments allows the observation of enough credit events to derive unbiased default frequencies.

However, this need not be the case if the population of instruments being sampled offers only limited evidence of credit events and what might cause them. Hence, there exists a fallacy of composition—inferring that the observation of a representative sample of credit instruments necessarily allows the observation of an equally representative set of credit events.

If credit events are sufficiently frequent within a population of loans, this assumption can be reasonable. If, however, credit events are sufficiently rare, or are clustered in time and space, then historical events of default may not provide unbiased estimates of future credit risk to creditors.

This point is well illustrated by the near-zero, 10-year marginal probability of default for project finance loans reported in the most recent Moody's study (Moody's Investors Service [2017, p. 22]) In the very large sample available to Moody's, there is close to zero reported (first) default events in the 10th year after financial close. By this measure, conditional on no default during the first 10 years, reported default risk in

private project finance debt is equivalent to that of a AAA rated bond, yet it typically carries a credit spread in the range of 80–250 bps, and when they are rated, infrastructure project loans tend to remain in the BBB/BBB–range even after 10 years, directly contradicting rating agencies’ own reported default probabilities.

Insurance statisticians are familiar with these issues. It can be difficult to derive meaningful statistics about the likelihood of events that are seldom observable, let alone in sufficient numbers to control for the multiple variables that might have determined them. The likelihood of an earthquake in a given location or that of the mid-air collision of two commercial airplanes are typical examples. Likewise, modeling a hazard rate for investors in private infrastructure project debt is made more difficult by the difficulty in observing large samples of the phenomenon of interest.

Moreover, infrastructure investment is characterized by the large size of each investment and a highly illiquid private market, meaning that individual investors are usually exposed to a handful of assets and not the mean asset available at that time. Hence, even if credit events were frequent enough to model default frequencies accurately for a representative investor holding a large sample of available loans, the immense majority of creditors would not be better informed about their own risk exposure.

In this article, we propose a simple but powerful approach to deal with the dearth of information available about credit risk in private infrastructure debt: we extend the structural model of credit risk in private illiquid debt we put forward in Blanc-Brude and Hasan [2016], using Bayesian inference to extract robust credit risk estimates from observable data on cash flow ratio collected at the individual borrower level. We also provide a calibration of the model using a unique dataset of project finance cash flow covering 14 European countries and going back two decades.

Our objective is to calibrate a model of *distance to default* (Kealhofer [2003]) in infrastructure project finance. Hence, the absence of observable default events does not limit our ability to model and predict default, since we can measure default risk *before* default events occur and even if they never occur. As a result, our approach also allows introducing control variables that are unavailable to reduced form models relying on observing actual default events.

This approach can be applied to a representative sample of loans, as is the case in existing studies, or to an individual portfolio or even a single loan, thus removing the requirement to hold a representative portfolio to derive meaningful information for an investor in private infrastructure debt.

The rest of this article is organized thus: first, we provide a very brief overview of the credit risk framework that underpins our empirical investigation. Next, we present our data and initial descriptive results, and the third section proposes a model to estimate the dynamics of private borrowers’ credit ratios. The fourth section describes our results, compares them with those from prior studies, and briefly discusses industrial applications.

FRAMEWORK

The only empirical research available on the credit risk of infrastructure projects is provided by credit rating agencies that have pooled information about realized default events. This data is privately held and has not been available for independent academic research. Existing academic research on credit risk in project finance has focused on the pricing of credit risk, that is, the determinants of credit spreads in bank loans and bonds used to finance infrastructure projects (see, e.g., Blanc-Brude and Strange [2007]; Esty [2001, 2002]; Esty and Megginson [2003]; Kleimeier and Megginson [2000]; and Sorge and Gadanecz [2008]). This literature typically finds that project finance credit spreads do not have the same determinants as plain vanilla corporate debt but offers little insight into how risky such instruments actually are.

Blanc-Brude and Hasan [2016, henceforth “BBH”] put forward a technical framework to measure credit risk in private illiquid debt using a structural approach—that is, modeling the mechanism by which defaults occur as opposed to the occurrence of default events. The BBH framework requires the knowledge of debt service coverage ratio (DSCR) dynamics of individual borrowers as its primary input. The DSCR is

$$DSCR_t = \frac{CFADS_t}{DS_t^{BC}} \quad (1)$$

in each period $t = 1, 2, \dots, T$ for a project financing of maturity T ; DS^{BC} is the debt service in the “base case,” that is, in the current debt contract.

Because infrastructure project companies are contractually prevented from raising additional funds (see Gatti [2013, Section 7.2.3.11.2, on negative covenants]), the DSCR provides an unambiguous definition of default thresholds. The hard default threshold is easily defined as $DSCR_t$ (that is, Moody's definition of default (Moody's Investors Service [2017])). Likewise, the DSCR can also be used to unambiguously define "technical" default thresholds (Yescombe [2002]), that is, covenant breaches defined in terms of the DSCR level ($DSCR_t$ (1.x) that qualify as credit events and provide additional control or step-in rights to lenders.

It follows that the statistical distribution of $DSCR_t$ provides a direct measure of the probability of default at time t . A firm can be considered in default if its DSCR falls below 1.x, and its probability of "hard" default is simply the probability that the DSCR falls below this threshold when $x = 0$.

BBH also show that knowledge of DSCR dynamics allows computing the *distance to default* or DD_t measure typically used in structural models of credit risk à la Merton [1974], since

$$DD_t = \frac{1}{\sigma_{DSCR_t}} \frac{DS_{t-1}^{BC}}{DS_t^{BC}} \left(1 - \frac{1}{DSCR_t} \right), \quad (2)$$

where σ_{DSCR_t} is the standard deviation of the annual percentage change in the $DSCR$ value.

Robust estimation of $DSCR_t$ dynamics in infrastructure projects thus allows implementation of a powerful model of credit risk that can address the main shortcomings of existing reduced form results; it can predict default even when there is no default event to be observed.

DATA AND PRELIMINARIES

Implementing the BBH approach to modeling credit risk requires collecting cash flow data, computing DSCRs, and modeling their dynamics. In what follows, we describe our dataset and initial empirical findings. DSCR data are found to be noisy and their dynamics nonlinear, which motivates the modeling approach taken in the third section of this article, "Modeling DSCR Dynamics."

Dataset

Our dataset of DSCR is built using manually collected and verified data from the audited statements

of accounts of individual project companies. For the purpose of describing DSCR dynamics, the dataset also includes individual project size, leverage, industrial sectors, countries, business model (whether the projects are contracted or merchant), and date schedules, including financial close and construction completion dates. The data obtained include information reported in companies' balance sheets, and in income and cash flow statements, and allows DSCRs, leverage, and other financial variables to be computed in each period.

In the literature, DSCRs are typically computed using an operating income (Ciochetti et al. [2003]; Harris and Raviv [1990]), but this can under/overestimate the cash flow available for debt service in practice. For instance, if a project is drawing down additional debt to make its debt payment, then the cash available for debt service will exceed operating income. Similarly, if the project is investing capital in physical assets, then the cash flow available for debt service will be less than its operating income.

In an effort to create a meaningful proxy for credit risk, we compute what we call an "economic" DSCR, thus:

$$DSCR_t = \frac{C_{bank,t} + C_{op,t} + C_{IA,t} + C_{dd,t} - C_{inv,t}}{DS_{senior,t}} \quad (3)$$

where $C_{bank,t}$, $C_{op,t}$, $C_{IA,t}$, $C_{dd,t}$, and $C_{inv,t}$ denote cash at bank, cash from operating activities, cash withdrawal from investment account, cash from debt drawdowns, and cash invested in physical investments, at time t . Thus, our computation includes cash at bank, including any debt service reserve account or other cash, which are typically not included in the DSCR certificates reported to project finance lenders. This formula provides us with a DSCR that is directly linked to the default threshold described in the BBH framework, which is also the Moody's definition of default: debt repayments cannot be made.

Thus, we compute realized DSCR observations across a range of infrastructure projects spanning more than 17 years, representing the largest such sample available for research to date, and conduct a series of statistical tests and analyses to establish the most adequate approach to modeling and predicting future DSCR levels and volatility.

Our dataset includes 267 projects spanning two revenue risk families ("contracted" and "merchant"),

EXHIBIT 1

DSCR Summary Statistics

	Nb firms	Without Outliers				With Outliers			
		Nb obs	50th-Q	Mean	SD	Nb obs	50th-Q	Mean	SD
Contracted	215	1615	2.21	2.68	1.71	1774	2.21	5.44	22.06
Merchant	52	428	1.92	2.68	1.92	475	1.92	3.79	10.62
Total	267	2043	2.16	2.68	1.75	2249	2.16	5.09	20.20

Note: Statistics without outliers are computed using DSCR observations between the 5th and 95th quantiles.

in seven sectors², from the 2000 to 2016, from all major European markets.³

Exhibit 1 presents summary statistics. Raw DSCR data contain some large outliers, especially on the upside. During the early years of a project's life, especially when bonds are used, CFADS can be very high in that period since debt proceeds are invested in physical capital in subsequent years. Other instances of high DSCRs occur in the last years of certain loans, when very little outstanding debt remains to be paid relative to the firm's free cash flow. These outliers impact reported mean and standard deviation. Exhibit 1 also reports descriptive statistics after truncating the data at the 90th quantiles of the empirical distribution.

We note that contracted projects tend to have higher DSCR than merchant ones (using our formula) which may come as a surprise. They tend to be hoarding more cash than their merchant counterparts, we discuss why in the section "Model Calibration Results."

Initial Findings

Further examination of the data reveals the following stylized facts:

- 1. Contracted and merchant are different:** The DSCRs of contracted and merchant infrastructure have different dynamics at the 1% level of statistical significance, that is, they belong to distributions with different means and volatility;⁴
- 2. Truncated DSCRs follow a Gaussian process:** While the raw DSCR data do not follow any clear statistical distribution, due to high outliers, we find that truncating the DSCR sample between [0, 5] for contracted and merchant families, achieves a reasonable goodness-of-fit with the log-normal density function.⁵

3. Sectors matter less than business models:

Testing for statistically significant differences of DSCR mean and volatility across industrial sectors yields no results once the business model families identified above are controlled for. These results are not reported here.

- 4. Project size and leverage matter:** We proxy size by the total value of the project's assets, and leverage is computed as the ratio of total outstanding senior liabilities to total assets. While project size has roughly similar distributions for the two families, project leverage is very different, and Contracted projects have a higher leverage compared to Merchant projects. For both contracted and merchant infrastructure, higher leverage is associated with lower DSCR levels and standard deviation. The effect of size on realized DSCRs is less strong, but higher size tends to be associated with higher DSCR levels and volatility for both Contracted and Merchant projects.

- 5. Linear models are inadequate:** Using ordinary least square (OLS) models, we further investigate the explanatory power of business models, sectors, regions, leverages, size, financial ratios and project and calendar years. OLS models yield weak results with low adjusted R^2 because DSCRs have dynamic profiles in project time, and in the cross-section, may be impacted by project-specific and macro-economic factors. OLS models require data with constant variance and limited or no serial correlation, but as Exhibit A3 shows, DSCRs exhibit non-constant variance (the Breush-Pagan test rejects the null hypothesis of homoskedasticity), significant autocorrelation (the Durbin-Watson test rejects the null hypothesis of no autocorrelation in the regression residuals) and non-normal residuals (the Shapiro test rejects the null of Gaussian

residuals), all of which suggests that linear models are ill-suited to model our DSCR data.

6. **A dynamic model is necessary:** To account for the dynamic nature of the data, we also fit a panel regression model using fixed effects for project time, while calendar years and sectors are controlled for in the cross-section.⁶ The most interesting specification controls for initial investment calendar years (in the cross-section) while taking into account fixed project year effects (in time). These models highlight the existence of multiple-year effects in DSCR dynamics, some driven by the project life-cycle and some by the state of the project finance sector at the time of financial close. Stronger trends can also be detected in merchant projects, which tend to de-leverage faster than contracted one.⁷

We conclude that descriptive statistics and linear regression models fail to capture DSCR dynamics in full. In the next section, we discuss an approach to build a more powerful model of DSCR dynamics.

MODELING DSCR DYNAMICS

To model the dynamics of DSCRs in infrastructure projects, we make the following assumptions inspired by the initial results described earlier:

- The DSCR follows a stochastic process that is a simple translation of the firm's free cash flow process, as per Equation (1);
- This process can be in two "states": *risky* or *safe*. In the risky state, $DSCR_t$ follows a latent process characterized by a log-normal distribution of *unobservable* mean and variance parameters, which determine the firm's credit risk, as per Equation 2;
- In the safe state, $DSCR_t$ is relatively high and its dynamics have no impact on credit risk, which can be assumed to be negligible, that is, $Pr(DSCR_t = 0) \sim 0$ conditional on being in the safe state.
- The DSCR process can transition between states, with some probability at each period in time.

In this setup, understanding credit risk in infrastructure projects consists of estimating the Gaussian distribution parameters of the $DSCR_t$ in the risky state, as well as the probability of $DSCR_t$ being in the risky state, at each point in time.

In what follows, we proceed in three steps: we first represent the transitions between risky and safe states as a standard Markov chain model. We then show how estimating state transition probabilities as well as the DSCR distribution parameters in the risky state can be done by applying simple Bayesian inference techniques, requiring only basic calculus as long as the assumption of log-normality of $DSCR_t$ in the risky state can be maintained. This assumption is deemed reasonable in light of the descriptive results above. Bayesian inference then allows for calibrating a $DSCR_t$ model even when limited information is available. For instance, state transitions can be rare, especially after the first few years of a loans' life.

DSCR States as a Markov Process

As well as having time-varying dynamics, we noted earlier that DSCRs could take high values, suggesting virtually zero probability of default at that time. Moreover, these high values were found to curtail the goodness-of-fit of a log-normal density function when applied to the data, whereas both Contracted and Merchant DSCRs could effectively be considered log-normal if the sample was truncated at a threshold $\overline{DSCR} = 5$.

These findings suggest a simple state-transition model of the DSCR process with two distinct states:

1. A safe (i.e., risk-free) state "s" in which $DSCR_t > \overline{DSCR}$, and
2. A risky state "r" in which $DSCR_t < \overline{DSCR}$.


An example path from state to state followed by an individual project is illustrated by Exhibit 2.

In the safe state, the realized DSCR is so high that no matter how volatile the process might be, from a senior creditor perspective, the probability of default is not significantly different from zero. The debt is (conditionally) risk-free. As before, in expectation at time t , an infrastructure project may transit in and out of the safe state at each point in the future, with some probability (π_{st}). In the safe state, estimating the parameters of the DSCR distribution, in particular estimating its variance, is also irrelevant.

In the risky state, a project's DSCR may take values between 0 and some higher threshold $\overline{DSCR}$. From this state, it may either stay in the *risky* state at the next period, or transit into the safe state "s," described above.

EXHIBIT 2

Illustration of the DSCR Path between States



Within the risky state, the project may default if the DSCR falls below 1.

In the risky state, we know from our results above that if the upper threshold is set at $\overline{DSCR} = 5$, the data follow a log-normal process, the parameters of which (position and scale) can be estimated using the particle filtering approach described in the section “Particle Filtering.”

Estimating state transition probabilities amounts to estimating the components of a matrix describing a Markov process. As suggested above, the DSCR process can take one of two states S_t at time t : a risky state defined as $DSCR_t \geq 0$ and denoted by $S_t = r$, or a safe state such that $DSCR_t > \overline{DSCR}$ and denoted by $S_t = s$. The probability of being in the risky state is defined as $Pr(DSCR_t \geq 0) = Pr(S_t = r) = p_t$, while $Pr(DSCR_t > \overline{DSCR}) = Pr(S_t = s) = q_t = 1 - p_t$.

In a Markov process, future DSCR states are a function of the current state. Denoting time $i = t - 1$, let $\pi_{rs} = Pr(S_{t+\tau} = s | S_{t+i} = r)$ be the state transition probabilities between states r and s , with the one-step transition probability matrix given by

$$S_{t+i} = \begin{pmatrix} \pi_{rr} & \pi_{rs} \\ \pi_{sr} & \pi_{ss} \end{pmatrix} \quad (4)$$

Here, π_{rr} is the probability of being in the risky DSCR state at time $t + \tau$ conditional on having been in the same state at time $t + i$, and π_{rs} is the probability of transitioning to the safe state at time $t + \tau$ conditional on having been in the risky state at time $t + i$.

The probability of being in the risky state at $t + \tau$ conditional on the realized state at $t + i$ is thus

$$p_{t+\tau} = p_{t+i}\pi_{rr} + (1 - p_{t+i})(1 - \pi_{ss}) \quad (5)$$

and in the matrix notation

$$\begin{bmatrix} p_{t+\tau} \\ q_{t+\tau} \end{bmatrix} = S_{t+i} \cdot \begin{bmatrix} p_{t+i} \\ q_{t+i} \end{bmatrix} \quad (6)$$

That is, the probabilities of being in the risky (safe) state in period $t + \tau$ are determined by the product of the transition matrix with the probabilities of being in the risky (safe) state in the previous period $t + i$.

Hence, starting from any point in time, for which we know which state the DSCR is in (i.e., $DSCR_t$ is either strictly greater than 1 or not), we can compute the probabilities of being in the risky and safe states at future periods by successively applying the transition matrix.

According to Equation (6), we can know the conditional probabilities of being in the risky or safe state in each future period $t + \tau$ by estimating S_{t+i} across the project life cycle for $i = 0, \dots, (T - 1)$, as well as initial state conditions.

Estimating Transition Probabilities

In Markov switching models, the transition probability from a state i to a state j is estimated by counting the observed number of transitions from state i to state j and dividing by the total number of transitions from state i . That is,

$$\hat{\pi}_{i,j} = \frac{n_{i,j}}{\sum_{k=1}^N n_{i,k}}, \quad (7)$$

where $\hat{\pi}_{i,j}$ denotes the estimated transition probability from state i to j , $n_{i,k}$ denotes number of transitions from state i to state k , and N denotes the total number of states.

In Bayesian Markov switching models, we start with forming a prior belief about transition probabilities between different states, based on information available otherwise, which is then updated as one observes actual transitions between states.

By definition, the values of any $\mathbf{S}_{t+\tau}$ are such that each line of the state transition matrix must add up to one, that is, $\pi_{rr} + \pi_{rs} = 1$. Hence,

$$\mathbf{S}_t = \begin{pmatrix} \pi_{rr} & 1 - \pi_{rr} \\ \pi_{rs} & 1 - \pi_{rs} \end{pmatrix},$$

that is, each row of $\mathbf{S}_{t+\tau}$ matrices is equivalent to an independent Bernoulli draw of parameter π_{rr} or π_{rs} , and we only need to estimate π_{rr} and π_{rs} to know the entire transition matrix at time $t + \tau$.

Say we can observe a population of N projects at time t , with n of “successes” (realized DSCR transitions between two given states), these data (call it Y) follow a binomial distribution (the outcome is binary) with the likelihood

$$\mathcal{L}(Y|\pi) = \binom{N}{n} \pi^n (1 - \pi)^{N-n}$$

where $\binom{N}{n} = \frac{N!}{n!(N-n)!}$ is the binomial coefficient.

According to Bayes’ Law

$$p(\pi|Y) \propto p(\pi)\mathcal{L}(Y|\pi)$$

that is, the posterior (distribution) is proportional to the prior (distribution) times the likelihood.

We can give a *beta* prior density to $Pr(\pi)$, such that

$$p(\pi; \alpha, \beta) = \frac{\Gamma(\alpha + \beta)}{\Gamma(\alpha)\Gamma(\beta)} \pi^{\alpha-1} (1 - \pi)^{\beta-1}$$

The beta distribution has a domain on $[0, 1]$ which can usefully represent a *probability* and can take *any* shape on its domain. The beta distribution is also *conjugate* with respect to the binomial likelihood, so that the product of the prior (beta) and the likelihood (binomial) is another beta distribution, which incorporates the information obtained from observing the data.

$$\begin{aligned} p(\pi|Y) &\propto p(\pi)\mathcal{L}(Y|\pi) \\ &\propto \pi^n (1 - \pi)^{N-n} \times \pi^{\alpha-1} (1 - \pi)^{\beta-1} \\ &\propto \text{Beta}(\alpha + n, \beta + N - n) \end{aligned} \quad (8)$$

... to a normalizing constant which does not depend on π .

Hence, the sufficient statistics to update the prior distribution of π are N and n , which we know to be observable, that is, with an observable population of N projects, we can count the number n of draws corresponding n realised DSCRs strictly greater than one at time $t + \tau$, given that we also observed DSCRs strictly greater than one at the previous period.

In other words, by assuming that the true value of π_{rr} is the mean of a beta distribution of parameters (α, β) (the meta-parameters), given that the likelihood function of the data follows a binomial distribution of parameter π_{rr} with N data points, we can update the values of the meta-parameters each time we observe n transitions (in this case projects staying in the risky state from one period to the next) among N new data points.

The posterior distribution of π_{rr} summarizes the state of our knowledge by combining information from newly available data expressed through the likelihood function, with *ex ante* information expressed through the prior distribution.

The posterior distribution of $Pr(\pi_{rr})$ then becomes a *new prior* each time new empirical observations become available. Bayesian inference thus allows sequential learning about the expected state transitions of projects’ DSCRs.

The same process is used to estimate π_{rs} , after observing projects transiting from the risky to the safe state.

Bayesian Parameter Estimation in the Risky State

We know from the results above that $DSCR_t$ is serially correlated and can change its risk profile during the investment life cycle of infrastructure projects. In other words, the expected value $E(DSCR_t)$ and the volatility $\sigma_{DSCR_t}^2$ are partly determined by the lagged values of the same quantities at time $t - 1$, and partly by innovations (incremental changes) or shocks happening at time t . However, the true mean and volatility of the $DSCR_t$ process are unobservable and can only be imperfectly measured by observing realized mean and variance of observed DSCRs.

Moreover, as argued earlier, there may not be sufficient realized DSCRs available to form accurate estimates of the latent DSCR mean and volatility of existing projects, partly because of a lack of data, and partly because the existing projects that were financed and

structured several decades ago may not be representative of projects that are being financed and structured today.

Therefore, we aim to build a model of the mean and volatility of $DSCR_t$ that is capable of integrating information contained in realized DSCRs, without making static assumptions about the underlying process, and learning from new observations as they become available.

Numerous models exist that aim to determine the position of a dynamic system and, based on the latest round of observations (measurements), to predict where it will be positioned in future periods. Such systems are frequently applied in robotics, aero-spatial tracking, and chemistry. Here, we apply Bayesian particle filtering to estimate the position of a given infrastructure project in a mean/volatility DSCR “plane” at a given point in time, and to predict its position—its DSCR mean and variance “coordinates” so to speak—in the following periods.

As before, Bayesian inference allows the parameters of the distribution of interest (here the DSCR at time t) to be treated as stochastic quantities, thus reflecting the limits of our current knowledge of these parameters, or their stochastic nature. Thus, each parameter of the DSCR distribution is given a distribution of its own and the variance of, for example, the parameter representing the volatility of DSCRs, represents our current uncertainty about the true value of this parameter.

In this setup, we first build a *prior* distribution of the DSCR process, given the current state of knowledge about infrastructure debt investments. In each period for which DSCR data become observable, this *prior* knowledge is updated to derive a more precise *posterior* probability distribution of $DSCR_t$.

In the risky state, $DSCR_t$ follows a log-normal process,

$$\log(DSCR_t) \sim N(m_t, p_t), \quad (9)$$

where m_t is the location parameter of the distribution, and p_t its precision, which is defined as the inverse of its variance, or $p_t = 1/\sigma^2$.

Hence, the latent state of the $DSCR_t$ process is $X_t = (m_t, p_t)$, and the state equation is

$$X_t = X_{t-1} + \phi_{t-1}U_{t-1} + W_t, \quad (10)$$

where U_t denotes various factors that may affect the mean and volatility of a project's DSCR, and ϕ_t is the project's exposure to these factors. In other words,

the parameters m_t and p_t of the $DSCR_t$ process are assumed to follow a stochastic autoregressive process with one lag, which is affected by k known factors, with the random innovation or disturbance W_t . For the set of factors, U_t , that may affect the DSCR distribution, we consider the project's business model, sector, region, and time to maturity. Thus, a representative contracted project in the UK will have an exposure of $U_{k,t} = 1$ for $k \in \{\text{Contracted}, \text{GBR}\}$, while it may have an exposure anywhere in the range of $[0, 1]$ to various sectors, determined by the number of U.K. Contracted projects in various sectors.

Unknown parameters (whether they are stochastic or not) are given a probability distribution. Here, m_t , the mean of the log-normal $DSCR_t$ process follows a normal distribution of meta-parameters μ_t and δ_t , and the precision p_t of the $DSCR_t$ process follows a gamma distribution of meta-parameters α and β . That is,

$$m_t \sim \phi_{m,t}U_t + N(\mu_t, \delta_t) \quad (11)$$

$$p_t \sim \phi_{p,t}U_t + \Gamma(\alpha_t, \beta_t) \quad (12)$$

The state vector X_t is written $X_t = ((m_t|\mu_t, \delta_t), (p_t|\alpha_t, \beta_t))$.

For ease of interpretation, we define $\phi_{m,t}$ coefficients such that they are zero for the representative project in every year. In the case of m_t , this can be done by rewriting the above Equation for m_t as follows:

$$m_t \sim \mu_t + \phi_{m,t}w_t + \phi_{m,t}(U_t - w_t) + N(0, \delta_t), \quad (13)$$

where w_t is the weight of the representative project on each factor, $U_{k,t}$, which is determined by the fraction of the number of DSCR observations with $U_{k,t} = 1$ relative to the total number of observations. Now, we can redefine μ_t as

$$\mu_t \equiv \mu_t + \phi_{m,t}w_t, \quad (14)$$

and write m_t as

$$m_t \sim N(\mu_t, \delta_t) + \phi_{m,t}(U_t - w_t) \quad (15)$$

Thus, for the representative project, the exposure to each factor, U_t , is equal to the weight, w_t , and the mean m_t is distributed as $N(\mu_t, \delta_t)$. However, for a project k in a specific family ($U_{k,t} = 1$), the mean, $m_{k,t}$ is given by the mean of the representative project, $N(\mu_t, \delta_t)$,

plus a contribution due the fact that the weight of the project for k th factor, given by $U_{k,t}$, deviates from the representative project's exposure to the k th factor, given by $w_{k,t}$.

Similarly, p_t can be rewritten relative to the representative project as

$$p_t \sim \Gamma(\alpha_t, \beta_t) + \phi_{p,t}(U_t - w_t) \quad (16)$$

As is well documented in the literature, under such parameterization the prior distribution of the parameters is conjugate (has the same functional form) to the likelihood of the data, which makes implementing Bayes rule straightforward and computationally easy.

Indeed, the conjugate prior of a log-normal process is a gamma-normal distribution (Fink [1997]), that is, as a function of m and p , the likelihood function is proportional to the product of a gamma distribution of p (with parameters a and b) with a normal distribution (with mean μ and precision δ) of m conditional on p .

Next, if the realized DSCR data Y follow a log-normal process of mean m and precision p , their likelihood function is given by

$$\mathcal{L}(m, p|Y) \propto p^{N/2} e^{\left(-\frac{p}{2} \sum_{n=1}^N (\log Y_n - m)^2\right)} \quad (17)$$

where N is the number of observations.

This relationship relates observations Y to the latent state X .

Following (Fink [1997]), the sufficient statistics (required data) to update a prior distribution are the number of observations N , $E[Y] = \frac{\sum_{n=1}^N \ln(Y_n)}{N}$, and SS

the sum of squared deviation of the log data about m ; and the joint posterior distribution $Pr(m^+, p^+)$ is given by the meta-parameters

$$\begin{aligned} \alpha^+ &= \alpha + \frac{N}{2} \\ \beta^+ &= \left(\frac{1}{\beta} + \frac{SS}{2} + \frac{\delta N (\bar{Y} - \mu)^2}{2(\delta + N)} \right)^{-1} \\ \mu^+ &= \frac{\delta \mu + N \bar{Y}}{\delta + N} \\ \delta^+ &= \delta + N \end{aligned}$$

Thus, each time a new set of DSCR data is observed, we know N , $\bar{Y}$ and SS , and the posterior values of α^+ , β^+ , μ^+ , and δ^+ can be computed according to Equation 18, and the posterior parameters m^+ and p^+ of the distribution of $DSCR_t$ derived, incorporating prior knowledge and the new information. Finally, coefficients for various factors can be computed as

$$\begin{aligned} \phi_{m,k,t} &= \frac{\bar{Y}_k - E[m^+]}{1 - w_{k,t}} \\ \phi_{p,k,t} &= \left(\frac{N_k}{(Y_k - \bar{Y}_k)^2} - E[p^+] \right) \frac{1}{1 - w_{k,t}} \end{aligned} \quad (18)$$

where $\bar{Y}_k = E[Y|U_{k,t} = 1]$ is the mean conditional on $U_{k,t} = 1$, N_k is the number of observations for $U_{k,t} = 1$, and $w_{k,t} = \frac{N_k}{N}$ is the relative number of observations for $U_{k,t} = 1$.

Particle Filtering

Filtering models are a form of signal processing and aim to arrive at some best-estimate of the value of a system, given some limited and possibly noisy measurements of that system's behavior.

Given any initial belief about the mean and variance of $DSCR_{t_0}$ —drawn for example from the historical project family mean and variance—and assuming a log-normal DSCR process, deriving the prior values of the state vector meta-parameters $X_{t_0} = ((\mu_{t_0}^-, \delta_{t_0}^- | m_{t_0}^-), (\alpha_{t_0}^-, \beta_{t_0}^- | p_{t_0}^-))$, is a matter of simple arithmetic.


Next, given the prior distributions of m_{t_0} and p_{t_0} , we make 1,000 draws for each parameter to generate 1,000 “particles”, that is, each particle i is a pair (m, p) , that is, a possible occurrence of the DSCR state X_{t_0} given the meta-parameters.

We then observe the data (realised $DSCR_{t_i}$) in the first investment period and compute the likelihood L_i for each one of the 1,000 particles given the data, as per Equation (17) on page 17.

Normalized likelihood scores w_i^8 are then used to rank individual particles, which are then resampled by weight, that is, each particle is duplicated $1000 \times w_i$ times and only the first 1,000 particles by rank are kept in the sample. Thus, the resampled particles are updated according to how likely they are to be the true mean and

EXHIBIT 3

Generating Particles Using Prior Knowledge to Estimate $DSCR_t$ Mean and Variance Using a Particle Filter



variance of the DSCR given all the DSCR observations. And the distribution of DSCR mean, m , and precision, p , is updated accordingly.

The resulting posterior parameters of the DSCR distribution at time t_1 then become the prior estimates of the DSCR process at time t_2 , before any observations are made at that time, and the filtering and updating process starts again.


Exhibits 3 and 4 show this procedure schematically. In Exhibit 3, particles are generated based on the prior distribution of m and p . These particles are then resampled using the likelihood of observed DSCR in Exhibit 4, and the particles with higher likelihood of explaining the observations get higher weight. The resampled particles then provide the updated estimate of DSCR mean and variance.

Hence, whether we are observing realized DSCRs for a whole sample of projects or for a single one, we can estimate the current and future trajectory of the DSCR process in a mean/volatility plane.

In summary, our approach consists of filtering the parameters of the DSCR distribution in the “risky” state in which we can reasonably assume that it follows a log-normal process, as well as the transition probabilities in and out of that state at each point in the project life cycle. In the next section, we implement this approach to our dataset.

EXHIBIT 4

Updating $DSCR_t$ Mean and Variance Estimates Using Resampled Particles



MODEL CALIBRATION RESULTS

In this section, we implement our approach to estimate $DSCR_t$ dynamics in project finance debt and calibrate the BBH model of credit risk.

DSCR State Transitions

The safe state threshold is set at $\overline{DSCR} = 5$, that is, high enough to justify the assumption of a zero-conditional probability of default and that of a log-normal DSCR process in the risky state.

We start with a uniform prior for the probability of staying within the same state. That is, p_{rr} and p_{ss} are all assumed to follow beta distribution with $\alpha = \beta = 1$. In the next period, we observe transitions between states, and update our prior estimates of α and β according to

$$\alpha_{ii} = \alpha_{ii} + n_{ii} \quad (19)$$


$$\beta_{ii} = \beta_{ii} + N_i - n_{ii} \quad (20)$$

where $i \in \{r, s\}$, and N_i denotes total number of transitions from state i , and n_{ii} denotes total number of transitions from state i to state i . The average probability of staying within the same state, p_{ii} , is then given by the definition of the mean of beta-distributed variables:

$$p_{ii} = \frac{\alpha_{ii}}{\alpha_{ii} + \beta_{ii}} \quad (21)$$

EXHIBIT 5

Transition Probabilities from the Safe State for Contracted and Merchant Projects in Operation Time



These updated (posterior) estimate of α_{ii} and β_{ii} are then used as prior estimates for transition probabilities for the next period. This evolution of transition probabilities captures both the effects of time variation in true underlying transition probabilities as well as the effect of learning about these true probabilities. As we move forward in time, our prior becomes more and more informed, and estimated transition probabilities become more stable.

Exhibit 5 shows transition probabilities in each period between safe and risky states. We see that the transition probabilities from the safe state to the risky (safe) state go down (up) in project time for both Contracted and Merchant projects, while transition probabilities from the risky state (not shown here) stay largely unchanged. This suggests that projects de-risk over time, as the projects that are in the safe state become more likely to remain in the safe state, over time. This de-risking trend is stronger for Merchant projects for which the probability of staying in the safe state conditional on being in the safe state goes up to 60% in operation year 17.

Generally speaking, the risky DSCR state is highly persistent, that is, the probability of staying in this state once the process is in it, is high, for both project families,

EXHIBIT 6

Average State Transition Probabilities (%) between Risky and Safe States


	All		Contracted		Merchant	
	R	S	R	S	R	S
R	96.71	3.28	97.16	2.84	95.08	4.91
S	69.39	30.60	67.85	32.14	75.04	24.96

indicating that the projects are likely to stay in the risky state once they are in the risky state. The RR transition probability is close 100% (not shown here), is much higher than the SS transition probability, which is close 50%, which suggests that the projects that are in the risky state tend to stay in the risky state, while projects that are in the safe state may transition into the risky state. This is also illustrated in Exhibit 6, which shows average per-period transition probabilities for the two states.

Finally, Exhibit 7 shows the probabilities of being in one of two states for Contracted and Merchant projects. For both types of projects, the probability of being in the risky state decreases over time, while the probability of being in the safe state increases over time.

EXHIBIT 7

Probability of Being One of the Two DSCR States for Contracted and Merchant Projects



Filtered Values of the DSCR Distribution

To calibrate the risky state dynamics, we estimate the DSCR log-normal distribution parameters m and p at each point in time as described above.

As before, prior values are assumed for the four meta-parameters— μ , δ , α , β —of the DSCR distribution, and updated sequentially using a particle filter. The values of these meta-parameters (μ , δ , α , β) are shown in Exhibit A4. Meta-parameters give us the DSCR distribution parameters of m and σ in Exhibit A4. Due to the relatively large number of DSCR observations in the cross-section in initial years, the effect of initial prior values fades away within the first period in this case, and the variance of estimated distribution parameters (Δm , $\Delta \sigma$) is very low, indicating that the mean values of the distribution parameters ($\bar{m}$, $\bar{\sigma}$) are estimated with a high confidence. Once the distribution of parameters of the DSCR distribution is known, we can take the mean values of these parameters ($\bar{m}$, $\bar{\sigma}$) as our best guess for the parameters of the DSCR distribution.

Exhibit 8 shows the estimated mean and SD of $DSCR_t$ in the risky state for various project groups. Conditional on being in the risky state, different groups can exhibit various levels of mean and standard deviation of $DSCR_t$. These levels may seem counterintuitive since

different business models are found to have similar volatility at some points in time; however, unconditional levels of credit risk are also determined by the DSCR level and the probability of the risky state for each group (we return to this in the next section).

The effect of different project characteristics on the mean and SD of DSCR can also be seen in Exhibit A5, which shows the sensitivity (beta) of DSCR mean and SD estimates to the change in a specific project characteristic: the DSCR SD is typically less sensitive to changes in project characteristics than the DSCR mean, and among various project characteristics, ESP (Spain) region has the largest effect on DSCR levels.


CREDIT RISK IMPLICATIONS

Together the results described in the previous section provide a complete picture of the credit risk of private infrastructure project debt over the sample period.

The probability of the DSCR being in the risky state, times the probability of $DSCR_t$ falling below a certain threshold in the risky state is the equivalent of the marginal default frequencies (PD) reported by rating agencies. Cumulative default rates, default risk over calendar (as opposed to project) time and distance to default can also be reported.

EXHIBIT 8

Conditional DSCR Mean and Volatility for Different Project Groups (i.e., conditional on being in the “risky” state when $DSCR_t$ [5])



Panel A of Exhibit 9 shows the unconditional probability that $DSCR_t$ falls below the hard default threshold at each point in project time for the full sample period. Only the first default of payment is counted, which is consistent with Moody’s definition of default (missing one payment) and computation of marginal PDs in Moody’s [2017]. Consistent with rating agencies results, we find that project finance borrowers tend to de-risk over time. However, we find a range of default risk profiles depending on the control variables used. At the aggregate level (black line), we find higher levels of marginal PD than in Moody’s [2017]. This is not surprising since we measure all potential defaults, as opposed to a biased sample of actual credit events. We note that after 10 years, conditional on no default until that time, PD in project finance senior debt is typically not equal to zero, in contrast with reported PDs in Moody’s [2017], but more in line with typical credit ratings.


The ability to introduce control variables in our model allows for differentiating between different project business models, sectors, or countries and suggests very different risk profiles. For instance, projects financed in Spain exhibit, at the beginning of their lives, very high

credit risk but also the most dynamic change in their risk profile over time. Merchant projects are markedly riskier than contracted ones and projects in the UK, especially those created under the Private Finance Initiative (PFI) are found to be considerably less risky, primarily because they exhibit high levels of expected DSCR when cash at bank is included in the computation of $CFADS_t$.

Panel B of Exhibit 9 shows the same result excluding projects that actually did default, that is, it is the credit risk of the surviving population of borrowers 16 years after origination. Looking at Spain again, the bump of PD around year 10 in Panel A, corresponding to the collapse of the toll road sector and series of defaults/bankruptcies in 2012–2013, has disappeared in Panel B, since only surviving borrowers (that did not default) are included. We note that credit risk was still very high in early years for Spanish projects that survived, because they tended to be highly leveraged and generated limited free cash flow relative to their debt repayments. Over time, they were de-risked either because they gradually repaid their debt or were restructured before a hard default occurred, thus reducing credit risk.

EXHIBIT 9

Probabilities of Hard Default for Various Project Families with and without Defaulted Borrowers
(computed as the probabilities of $DSCR_t$ falling below 1:0)



Note: Panel A comprises all projects, including the ones that defaulted, and Panel B includes only projects that never defaulted.

While Panel B of Exhibit 9 is a biased result (survivorship bias), it shows that under the BBH framework, credit risk can be measured even when no default can be observed. Thus, a very prudent or lucky lender only picking borrowers that did not default over the past 16 years can still be considered to be exposed to significant credit risk.

These results are also presented in terms of cumulative default risk in project time (Panel A) and marginal risk in calendar time (Panel B) in Exhibit 10. As before, we report higher cumulative PDs in project time than rating agencies: across all sectors and countries 10-year cumulative PD is close to 12%, almost double the figure reported in Moody's [2017]. We note that the 10-year cumulative probability of default in Spain is close to 50%.

In calendar time (Panel B of Exhibit 10), the aggregate PD follows the business cycle, decreasing from the early 2000s until the 2008 financial crisis and increasing again from 2009 onwards, but not quite as much. This is consistent with previously reported evidence and is driven by the much larger number of

contracted projects relative to merchant ones in the underlying population of infrastructure projects in Europe since 2008. In effect, merchant project PD taken in isolation tends to peak much more markedly at bad times in the economic cycle, especially Spanish projects, but also any merchant project as illustrated by the blue line in Exhibit 10.

Finally, Exhibit 11 reports DD_t (distance to default) as per Equation 2. As before, projects tend to de-risk over time as their mean $DSCR_t$ increases and their volatility decreases. PFI projects in particular are found to be 2.5 standard deviations away from a hard default after less than a decade of operation, which confers them a very low credit risk. Conversely, merchant projects exhibit a stable distance to hard default over time at 1.5 standard deviations from the default threshold. DD_t in Spain is the lowest in early project year but shifts upward closer to the sample mean as numerous projects fail to survive the first 10 years and leave the reporting sample, having experienced hard defaults.

In this article, we have conducted the first empirical study of $DSCR$ dynamics in infrastructure project

EXHIBIT 10

Cumulative Probabilities of Hard Default and Probabilities of Hard Default by Year





EXHIBIT 11

Distance to Hard and Soft Default at Thresholds of 1:0 and 1:1



finance using a new and unique dataset of 15 years of realized DSCR data for 267 projects in Europe and covering two broad categories of projects: those receiving a contracted income and those exposed to merchant or commercial risk.

We conclude that relying on reduced form models that require observing numerous default events is ill-suited in a sector characterized by a limited number of observable credit events. Instead, the structural approach proposed by BBH can be calibrated using observable

DSCR data and Bayesian inference. Our results suggest that credit risk in project finance is heterogeneous and driven by project business models, but also regional and sector specific, resulting from the procurement choices and types of creditors active in various markets. At two ends of the spectrum, contracted projects procured in the UK (primarily PFI projects) exhibit very low if any credit risk, whereas merchant projects procured and financed in Spain represent a high-risk form of credit.

APPENDIX A

EXHIBIT A1

Non-Parametric Test Results for the Difference in Mean and Variance between the Contracted and Merchant DSCR Subsamples

Wilcoxon Test		Bartlett Test		Kolmogorov–Smirnov Test	
Test Stat	p-Value	Test Stat	p-Value	Test Stat	p-Value
494183.00	0.03	18.14	0.00	0.14	0.00

Note: Mann–Whitney, Bartlett, and Kolmogorov–Smirnov tests are non-parametric tests for the equality of means, variances, and distributions of two samples, respectively.

EXHIBIT A2

Estimated DSCR Distribution Parameters and the Corresponding p-Values for the Cramér–von Mises Goodness-of-Fit Test

	μ	$\Delta\mu$	σ	$\Delta\sigma$	Mean	SD	Stat	p-Value
Contracted	0.70	0.01	0.38	0.01	2.17	0.86	0.38	0.11
Merchant	0.60	0.02	0.44	0.02	2.01	0.92	0.45	0.11

EXHIBIT A3

Test Results for Heteroskedasticity, Autocorrelation and Normality of Residuals for Log(Dscr) Ordinary Least Squares (OLS) Regression

	Test	Statistic	p-Value
Heteroskedasticity	Breusch–Pagan	806.00	0.00
Autocorrelation	Durbin–Watson	1.12	0.00
Normality	Shapiro	0.80	0.00

EXHIBIT A 4

Meta-Parameters of the DSCR Distribution for Contracted and Merchant Families

Project Year	Contracted				Merchant			
	μ	δ	α	β	μ	δ	α	β
1	0.70	1.00	1.00	1.00	0.70	1.00	1.00	1.00
2	0.71	18.50	9.75	28.53	0.67	7.50	4.25	11.37
3	0.73	48.33	24.67	75.27	0.66	17.00	9.00	24.85
4	0.73	87.50	44.25	139.29	0.65	28.75	14.88	41.60
5	0.72	136.00	68.50	214.79	0.64	41.80	21.40	64.92
6	0.71	228.00	114.50	370.98	0.61	66.60	33.80	113.15
7	0.71	337.40	169.20	565.12	0.61	95.60	48.30	161.92
8	0.70	456.60	228.80	793.15	0.60	126.40	63.70	210.73
9	0.70	576.00	288.50	1034.12	0.60	158.00	79.50	261.52
10	0.72	690.00	345.50	1253.45	0.58	189.00	95.00	307.12
11	0.74	794.20	397.60	1445.28	0.59	218.00	109.50	343.08
12	0.76	887.80	444.40	1635.29	0.62	243.60	122.30	386.84
13	0.76	969.20	485.10	1816.55	0.63	266.20	133.60	422.55
14	0.78	1041.80	521.40	1948.40	0.67	285.20	143.10	452.95
15	0.79	1103.40	552.20	2094.02	0.68	301.00	151.00	480.84

EXHIBIT A 5

DSCR Mean and SD Coefficients (Beta) for Various Project Characteristics

Project Year	Contracted		Merchant		GBR		ESP	
	Mean	SD	Mean	SD	Mean	SD	Mean	SD
1	0.01	0.00	-0.10	0.67*	0.05	0.00	-0.25*	0.00
2	0.01	-0.00	-0.09*	0.34**	0.06	-0.00	-0.16**	0.02
3	0.02	-0.00	-0.08*	0.24*	0.07	-0.02	-0.14**	0.03
4	0.01	-0.00	-0.07	0.19*	0.09	-0.02	-0.14*	0.03
5	0.01	-0.00	-0.08	0.15*	0.09*	-0.02	-0.14*	0.01
6	0.02	-0.01	-0.08	0.01	0.11*	-0.03	-0.12*	0.01
7	0.02	-0.00	-0.08	0.01	0.12*	-0.03*	-0.14*	0.01
8	0.02	-0.01	-0.10*	0.01	0.11*	-0.03*	-0.14*	0.01
9	0.03	-0.01	-0.12*	0.01	0.10*	-0.02*	-0.15**	0.01
10	0.03*	-0.01	-0.13**	0.02	0.08	-0.02	-0.16**	0.03
11	0.04*	-0.01	-0.11*	0.03	0.07	-0.02	-0.15**	0.03*
12	0.03	-0.02	-0.11*	0.03	0.06	-0.02	-0.15*	0.02
13	0.03	-0.02	-0.09	0.03	0.04	-0.01	-0.13*	0.01
14	0.03	-0.02	-0.08	0.03	0.04	-0.01	-0.10*	0.00
15	0.02	-0.02	-0.08	0.03	0.04	-0.01	-0.08	-0.00

*, **, and *** indicate statistical significance at the 10%, 5%, and 1% level, respectively.

ENDNOTES

¹In each period, a sample of live loans can be observed, a certain number of which go into default during that period. The hazard rate is simply the ratio of the number of reported defaults to that of the number of live loans at the beginning of each period (Moody's [2017 p. 17]).

²Transportation, Telecoms, Oil & Gas, Industrial, Government Services, Environmental Services, and Energy.

³UK, Spain, France, Italy, Portugal, Germany, Norway, Sweden, Ireland, the Netherlands, Poland, Slovakia, and Austria.

⁴Exhibit A1 provides the results of non-parametric tests of the null hypothesis that both contracted and merchant

infrastructure have the same mean and variance. The resulting Mann–Whitney, Bartlett, and Kolmogorov–Smirnov test statistics lead to the conclusion that the null hypothesis can be rejected with a high degree confidence (low p-value). As a result, we conclude that contracted and merchant projects exhibit different DSCR distributions.

⁵Increasing the upper cutoff level for the DSCR up to 10 does not significantly affect the goodness-of-fit in most operation years, but the goodness-of-fit deteriorates in some operation years. Exhibit A2 shows that within this DSCR range of [0;5], we cannot reject the null hypothesis that DSCRs follow a truncated log-normal distribution. This result is also true at each point in time t in the investment life.

⁶That is, the model fits a different intercept for each investment year, as opposed to each firm, and controls for calendar-year specific effects through calendar-year dummies.

⁷We also test for the impact of profit margins, asset turnover (revenue/total assets ratio), cash return on total assets (operating cash/total assets), capex coverage (operating cash/capital expenditures), and capital expenditures to revenue ratio on realized DSCR levels. Only asset turnover is significant for both business model families, with a positive coefficient, which is intuitive.

$$w_i = \frac{L_i - \min}{\max - \min}.$$

REFERENCES

- Blanc-Brude, F., and M. Hasan. “A Structural Credit Risk Model for Illiquid Debt.” *The Journal of Fixed Income*, Vol. 26, No. 1 (2016), pp. 6–19.
- Blanc-Brude, F., and R. Strange. “How Banks Price Loans to Public–Private Partnerships: Evidence from the European Markets.” *Journal of Applied Corporate Finance*, 19 (2007), pp. 94–106.
- Ciochetti, B.A., Y. Deng, G. Lee, J.D. Shilling, and R. Yao. “A Proportional Hazards Model of Commercial Mortgage Default with Originator Bias.” *The Journal of Real Estate Finance and Economics*, 27 (2003), pp. 5–23.
- Esty, B.C. “Structuring Loan Syndicates: A Case Study of the Hong Kong Disneyland Project Loan.” *Journal of Applied Corporate Finance*, 14 (2001), pp. 80–95. <http://doi.wiley.com/10.1111/j.1745-6622.2001.tb00440.x>.
- . “Returns on Project Financed Investment.” *Journal of Applied Corporate Finance*, 15 (2002), pp. 71–86.
- Esty, B.C., and W.L. Megginson. “Creditor Rights, Enforcement, and Debt Ownership Structure: Evidence from the Global Syndicated Loan Market.” *The Journal of Financial and Quantitative Analysis*, 38 (2003), pp. 37–59.
- Fink, D. “A Compendium of Conjugate Priors.” *Montana The Magazine of Western History*, 1997, pp. 1–47.
- Gatti, S. *Project Finance in Theory and Practice*. Academic Press, 2013.
- Harris, M., and A. Raviv. “Capital Structure and the Informational Role of Debt.” *The Journal of Finance*, Vol. 45, No. 2 (1990), pp. 321–349.
- Kealhofer, S. “Quantifying Credit Risk I: Default Prediction.” *Financial Analysts Journal*, Vol. 59, No. 3 (2003), pp. 30–44.
- Kleimeier, S., and W.L. Megginson. “Are Project Finance Loans Different from Other Syndicated Credits?” *Journal of Applied Corporate Finance*, 13 (2000), pp. 75–87.
- Merton, R.C. “On the Pricing of Corporate Debt: The Risk Structure of Interest Rates.” *The Journal of Finance*, 29 (1974), pp. 449–470.
- Moody’s Investors Service. *Default and Recovery Rates for Project Finance Bank Loans, 1983–2015*. Technical Report, 2017.
- Sorge, M., and B. Gadanecz. “The Term Structure of Credit Spreads in Project Finance.” *International Journal of Finance & Economics*, 13 (2008), pp. 68–81.
- Yescombe, E.R. *Principles of Project Finance*. Academic Press, 2002.
- To order reprints of this article, please contact David Rowe at drowe@iijournals.com or 212-224-3045.