A Publication of the EDHEC Infrastructure Institute-Singapore

Cash Flow Dynamics of Private Infrastructure Project Debt

Empirical evidence and dynamic modelling

March 2016

with the support of
NATIXIS

EDHECinfra
Singapore Infrastructure Investment Institute

Table of Contents

The author would like to thank Noël Amenc, Lionel Martellini, Anne-Christine Champion, Julien Michel, Marie Monnier-Consigny and Qi Wang for useful comments and suggestions. They would also like to thank the investors and banks that contributed data for this study, including members of the Long-Term Infrastructure Investors Association and the European Investment Bank, as well as the data collection efforts of Bui Nguyen Ngoc Ha, Trang (Christy) Tran, Divya D/O Sangaraju, Cai Qiyu and Zhang Chunyu. Financial support from Natixis is acknowledged. This study presents the author's views and conclusions which are not necessarily those of EDHEC Business School or Natixis.

Foreword

The purpose of the present publication, "Cash Flow Dynamics of Infrastructure Project Debt", which is drawn from the NATIXIS research chair at EDHEC Infrastructure Institute-Singapore, is to empirically validate a number of hypotheses put forward in previous EDHEC-Risk Institute papers on the valuation and risk measurement of private infrastructure debt instruments.

In this paper, the authors rely on a new and unique dataset of infrastructure project cash flows to compute and document the dynamics of debt service cover ratios (DSCRs) in private infrastructure projects, thus allowing the implementation of a fully-fledged structural credit risk model.

This dataset is one of the early outputs of a substantial effort launched at EDHEC to collect private information and document the investment characteristics of infrastructure investments, and create a global database of infrastructure cash flows spanning several decades.

This paper offers a powerful validation of the insights developed in earlier publications of the EDHEC/NATIXS Research chair, including the existence of tractable DSCR dynamics in infrastructure project finance, and the ability to predict credit events and value credit instruments using such metrics.

In a context were data paucity remains a concern and complete time series of cash flows covering the decade-long life of investments remain rare, the authors also propose a novel approach to modelling and predicting the "trajectories" of DSCRs in infrastructure projects. The paper documents the existence of homogenous "families" of cash flow dynamics in infrastructure projects that are best explained by the contractual characteristics of these investments as well as the initial financial structuring choices made jointly by project sponsors and creditors.

Hence, this paper also addresses the question of the path-dependency found in project dynamics and the resulting serial correlation of returns, especially in cases where achieving full diversification might be challenging.

Borrowing from statistical methods usually applied in the physical sciences, the authors show how new information arriving sequentially, which is typical of long-term investment in infrastructure, can be integrated to infer the conditional parameters of cash flow distributions and the related credit risk measures for subgroups or even individual investments.

We are grateful to NATIXIS for their support of this study in the context of this research chair at EDHEC Infrastructure Institute-Singapore.

Noël Amenc

Associate Dean, EDHEC Business School

About the Authors

Frédéric Blanc-Brude is Director of the EDHEC Infrastructure Institute–Singapore and represents EDHEC Business School on the Advisory Board of the Global Infrastructure Facility of the World Bank. He holds a PhD in Finance (King's College London), and degrees from London School of Economics, the Sorbonne and Sciences Po Paris.

Majid Hasan is Research Associate and Head of Asset Pricing at EDHEC Infrastructure Institute-Singapore. He holds a Masters in Applied Mathematics from Western University, Canada, and a Bachelors in Physics from the University of Punjab, Pakistan. He is currently completing his PhD at EDHEC Business School.

Tim Whittaker is a Senior Research Engineer at EDHEC Infrastructure Institute-Singapore and Head of Data Collection. He holds a Master of Business (Financial Management) an, Bachelors of Economics and Commerce from the University of Queensland, and a PhD in Finance from Griffith University. He also has a previous experience as a corporate finance analyst.

Executive Summary

Executive Summary

The objectives of this paper are to document the statistical characteristics of debt service cover ratios (DSCRs) in infrastructure project finance, and to develop and calibrate a model of DSCR dynamics.

For this purpose, we collect a large sample of realised DSCR observations across a range of infrastructure projects spanning more than 15 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 modelling and predicting future DSCR levels and volatility.

Using these results, we build a model of the conditional probability distribution of DSCRs at each point in the life of infrastructure projects.

DSCR Dynamics Incorporate Key Information About the Credit Risk of Infrastructure Debt

In a previous paper (Blanc-Brude et al., 2014), we showed that debt service cover ratios can play a pivotal role in the modelling of credit risk in infrastructure project finance.

This is because the DSCR of an infrastructure project company, which measures the amount of cash available to make the current period's debt service, provides us with

1. An unambiguous definition of the point of hard default (default of payment), i.e. $DSCR = 1$, and
2. An equally unambiguous definition of key technical default covenants i.e. $DSCR = 1.x$, while both types of default events create significant embedded options for creditors following a credit event.
3. Moreover, knowledge of DSCR dynamics is sufficient to estimate the firm's distance to default (DD), which is the workhorse of the so-called Merton or structural credit risk model.
4. DSCR dynamics can also be combined with future debt service to compute the expected value and volatility of the firm's future free cash flow, which is instrumental in measuring enterprise value in the case of infrastructure projects, since they derive their value almost entirely from future operating cash flows.

For these reasons, documenting DSCR dynamics using realised DSCR data is an important part of the objective to create investment benchmarks of private infrastructure debt, as described in the roadmap published by EDHEC-Risk Institute in 2014 (Blanc-Brude, 2014).

A Combination of Empirical Analysis and Statistical Modelling is Necessary

DSCRs in infrastructure project finance are mostly undocumented both in industry and academic empirical literature. While DSCR information is routinely collected by the

Executive Summary

creditors of infrastructure projects, this type of data is typically confidential and not available in large datasets.

From such data paucity, especially in time series, it follows that empirical observations alone are not sufficient to document the expected behaviour of infrastructure project cash flows over their entire investment life, and a combination of *ex ante* modelling and empirical observations is necessary.

Finally, private infrastructure investment tends to be characterised by very large individual investments, almost necessarily leading to poorly diversified portfolios. This suggests that assuming the mean-reversion of investors' infrastructure debt portfolios may not be realistic and that idiosyncratic risk should be taken into account.

In particular, individual infrastructure investments can exhibit significant "path dependency" and investors cannot necessarily take for granted the notion that they are exposed to the "median infrastructure project."

For both sets of reasons (data limitations and the importance of firm-specific risk), an adequate model of the DSCR should be able to capture conditional dynamics and explicitly integrate the different credit "states" that an infrastructure project might go through.

This can help both learning from the data as and when it becomes available, and taking into account the path-dependency of each instrument when estimating future cash flows, instead of assuming a reversion to the population mean.

Current academic and industry literature is static in nature and relies on "reduced form" credit models, which are likely to be biased given the nature of empirical data available and, in the current state of empirical knowledge, can only address a limited number of dimensions of private infrastructure debt investment: the historical frequency of default events, and to some extent, average recovery rates.

For these reasons, in this paper we endeavour to better document the dynamics of DSCRs in infrastructure project finance and build a model of DSCR dynamics using Bayesian inference to describe credit state transitions and to estimate the mean and variance of the DSCR in each state and at each point in an instrument's life. This allows better prediction of defaults, conditional on the actual trajectory of individual investments or groups of projects. The ability to predict cash flows and their volatility is then instrumental in the implementation of the private infrastructure debt valuation model defined in Blanc-Brude et al. (2014).

Dividing Infrastructure Investments into Groups Defined by Their "Business Model"

In Blanc-Brude et al. (2014), we described two generic and intuitive types of infrastructure project companies and called them "contracted" and "merchant".

Executive Summary

This distinction was informed by the casual observation that the financial structure of infrastructure project finance vehicles is often such that it requires, at the onset, either a rising or a flat "base case" DSCR profile.

A rising base case DSCR profile then implies an increasing volatility of $DSCR_t$. That is, creditors would demand a higher DSCR in the future to protect themselves against rising expected volatility of the cash flows available for debt service (CFADS). Such projects would also have longer "tails"[1] and exhibit between 70% and 80% of initial senior leverage. We argued that such structuring decisions signalled infrastructure projects that were exposed to commercial risks, such as a power plants selling electricity at market prices, and referred to these projects as *Merchant* infrastructure.

Conversely, we argued that the decision to structure a project while requiring a lower and flatter base case DSCR profile implied the expectation of a lower and constant conditional volatility of cash flows. We observed that projects with little to no market risk are financed with such a flat DSCR base case and also have shorter tails and a higher level of senior leverage, usually around 90%. Examples of these projects include social infrastructure projects, such as schools or hospitals that receive a fixed payment from the public sector, or energy projects that benefit from a long-term "take-or-pay" purchase agreement. We called these projects *Contracted* infrastructure.

1 - The amount of time between the original loan maturity and the end of the project's life, thus allowing higher recovery rates in the event of restructuring.

In this paper, we endeavour to determine statistically whether **realised DSCR dynamics** fall into categories determined by the distinctions made above between Contracted and Merchant infrastructure, as well as exogenous conditions at the time of financing and when the data is observed. We then use our results to design a model of DSCR dynamics.

The Largest Sample of DSCR Data Available for Research to Date

Our dataset of *realised* DSCRs is built using data manually collected and verified from the audited statements of accounts of several hundred project companies, as well as DSCR data reported by private contributors.

We hand collected 15 years of realised DSCR data for more than 200 projects in Europe and the United States covering our two broad categories of projects (those receiving a contracted income and those exposed to merchant or commercial risks), in seven sectors, from the early 1990s to 2015.

Our initial analysis of the data reveals some important points that confirm our intuition: the average credit risk profile of infrastructure projects can be usefully captured by categorising instruments in broad groups or families of underlying "business models."

The two groups correspond to two distinctive DSCR processes, with statistically different mean and variance parameters and following different project time dynamics. We also find, as intuition

Executive Summary

predicts, that contracted infrastructure DSCRs in the cross section are much less affected by macro-variables or the business cycle than merchant projects.

We confirm our hypotheses that the DSCR profile of an infrastructure project is strongly related to the firm's total business risk, and show that more highly leveraged projects achieve lower levels of realised DSCR volatility i.e. in project finance high leverage signals low asset risk as initially argued by Esty (2003).

But while descriptive statistics and linear regression models provide some insights about the determinants of the DSCRs, they fail to capture DSCR dynamics in full. Indeed, we find that the DSCR profiles of individual projects and families of projects are highly non-linear, auto-regressive and heteroskedastic (variance is not constant).

Hence, a more advanced model that can capture these dynamics is needed.

Tracking the "Coordinates" of the DSCR Distribution in the Mean-Variance State-Space

If the $DSCR_t$ is serially correlated and can change profile during the investment lifecycle of infrastructure projects, the *ex post* trajectory of individual projects could in principle correspond to any combination of high/low expected value $E(DSCR_t)$ and high/low volatility $\sigma^2 DSCR_t$. The DSCR of populations of projects would equally reflect the weighted trajectory of their constituents in a $DSCR_t$ mean/variance "plane".

Numerous models exist that aim to determine the position of a dynamic system and, based on the latest round of observations, to predict where it will be positioned in future periods. Such systems are frequently used in robotics, aero-spatial and chemistry applications. In this paper, we apply such approaches 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.

In the descriptive part of our analysis of the data, we show that realised DSCRs can be fitted to a lognormal process up to their 90^{th} and 85^{th} quantiles for contracted and merchant projects, respectively, at each point in their lifecycle, which allows us to develop an easily tractable model of parameter inference.

Hence, we propose a two-step modelling strategy combining a three-state model corresponding to break up points in the otherwise lognormal dynamics of the DSCR, with a filtering model to infer the values of the Lognormal process parameters (its "coordinates") in the state in which the DSCR is indeed lognormal.

Three-State Transition Probabilities

The DSCR process is assumed to occur in any one of three states at time t: a risky state (R) in which it is indeed an autoregressive lognormal process, a default state

Executive Summary

Figure 1: Illustration of the DSCR path between states

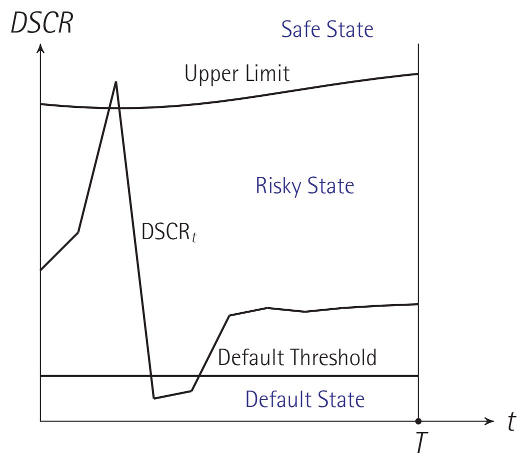

Figure 2: Probability in being one of the three DSCR states for contracted and merchant projects (D: default state; R: risky state; S: safe state).

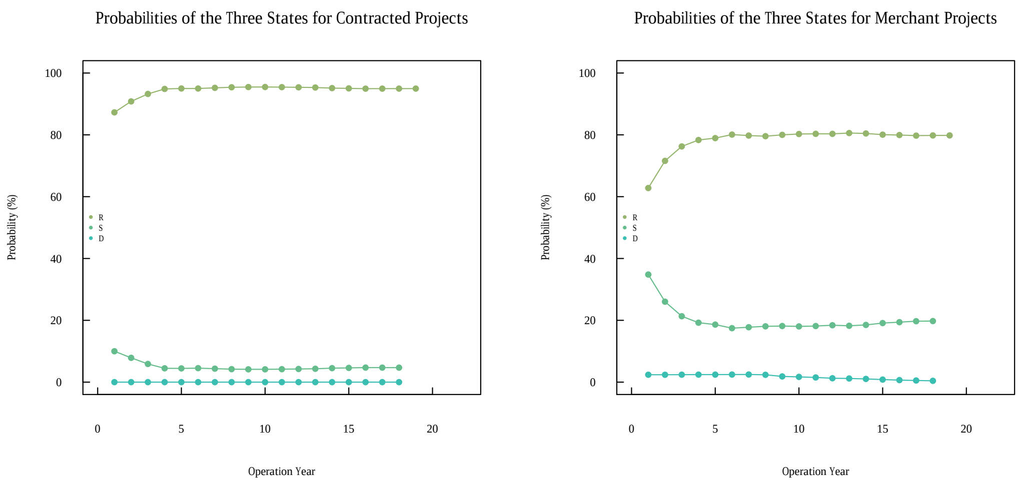

(D) defined by a threshold corresponding to $DSCR_t = 1$ in which the DSCR process stops until it emerges from default; and a safe state (S), corresponding to high realised values above the "good-lognormal-fit" quantile, in which case, as long as the DSCR stays in that state, the project debt is considered risk-free. This illustrated by Figure 1.

Hence, once a project's DSCR breaches the *hard default* threshold represented by $DSCR_t = 1$, it enters the default state, which it may or may not leave after a number of periods. In this state, creditors can take over the firm and optimise the value of exercising this option depending on the size of their exit costs and of restructuring costs. They may decide to waive the event of default or engage in negotiations with project sponsor to restructure the firm and its debt or indeed take over the firm and seek another sponsor (see Blanc-Brude et al., 2014, for a formal model).

Executive Summary

Hence, the firm may transit out of the default state (into the risky state) with some probability (say, π_{dr}) at the next period, or stay in this state and again transit out of default at the next period, etc.

In this state, the DSCR process effectively stops (in most cases, there is no debt service), hence estimating its mean and variance is irrelevant since the project is already in default.

In the safe state, on the contrary, the realised 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 (say, π_{sr}).

In this state, estimating the parameters of the DSCR distribution, in particular estimating its variance, is also irrelevant.

Finally, in between the default and safe states, a project's DSCR may take values between 1 and some higher threshold $\bar{DSCR}$. From this state, it may either stay in the *risky* state at the next period, or transit ouf of it into the state of default "d" or the safe state "s", both described above.

In this state, we know from our empirical results that if the upper threshold is set at the 85th/90th quantile of our DSCR sample, the data follows a lognormal process, the paramaters of which (position and scale) have to be estimated (see below).

Formally, this setup amounts to a relatively simple model of conditional state transition probabilities, which can be set in terms of a series of binomial draws and calibrated using Bayesian inference given some prior knowledge (e.g. we know from credit rating studies that projects tend to stay in default for 2.3 years) and counting the number of projects crossing the different state thresholds, conditional on which state they are in at the previous period.

The combination of the conditional probabilities of switching state at each point in time are then combined into the probability of being in any given state at time t, which is illustrated by figure 2.

For contracted projects the probability of being in the risky state is much higher compared to the probability of being in the other two states i.e. contracted projects are more likely to stay in the "normal" risky state.

For merchant projects, the probability of being in the risky state is lower, while the probabilities of being in the default and safe states are higher compared to the corresponding probabilities for contracted projects. Thus, merchant projects are found to have more diverse DSCR trajectories in state space, and each state is less persistent (stable).

This result confirms that path dependency can be an important dimension of

Executive Summary

infrastructure investment insofar as assets are more or less heterogenous and can be difficult to fully diversify very large and bulky assets. For instance, our results suggest that contracted infrastructure is more homogenous than merchant projects, which are more likely to follow paths that diverge strongly from the mean of the population.

Group and Individual DSCR Trajectories

To determine the value of the lognormal process parameters in the "risky" state discussed above, we propose to use a straightforward implementation of so-called particle filtering models to infer the parameter values of the DSCR's lognormal process in the risky state i.e. the state in which documenting and tracking the volatility of the DSCR really matters, because it is a direct measure of credit risk.

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 behaviour. Given our modelling objectives to accommodate small samples, and to avoid assuming static values for the $DSCR_t$ distribution parameters, we must be able to revise any existing parameter estimates once new data becomes available. This process is best estimated iteratively using Bayesian inference techniques described in detail in the paper.

We show that such a framework allows the dynamics of DSCR to be derived in well defined groups of projects as well as individual projects, including tracking the individual DSCR "path" followed by investments that do not necessarily correspond to the median infrastructure project.

The estimated dynamics of the DSCR process in contracted and merchant projects is shown in Figure 3, which describes the change in density of the DSCR process in investment time, and figure 4, which describes the trajectory of the DSCR state in the mean/standard deviation plane.

From such results, certain credit risk conclusions are immediately available, such as the expected default frequency for hard defaults but also any level of technical default ($DSCR_t = 1.x$) as shown in figure 5.

These results allows us to characterise the behaviour of groups of infrastructure projects which exhibit reasonably homogeneous dynamics, however, we know that highly idiosyncratic trajectories and path dependency should be a point of interest in a context where diversification is difficult to achieve in full.

Hence, we also show that the ability to infer both the expected value and the volatility of the DSCR process allows us to take a much more informed view on the credit risk of projects that substantially deviate from their base case.

For instance, consider an infrastructure project that follows an oft-observed trajectory: while it remains in the risky state throughout its life, it starts off with a relatively high DSCR, implying a merchant-

Executive Summary

Figure 3: DSCR densities for contracted and merchant families.

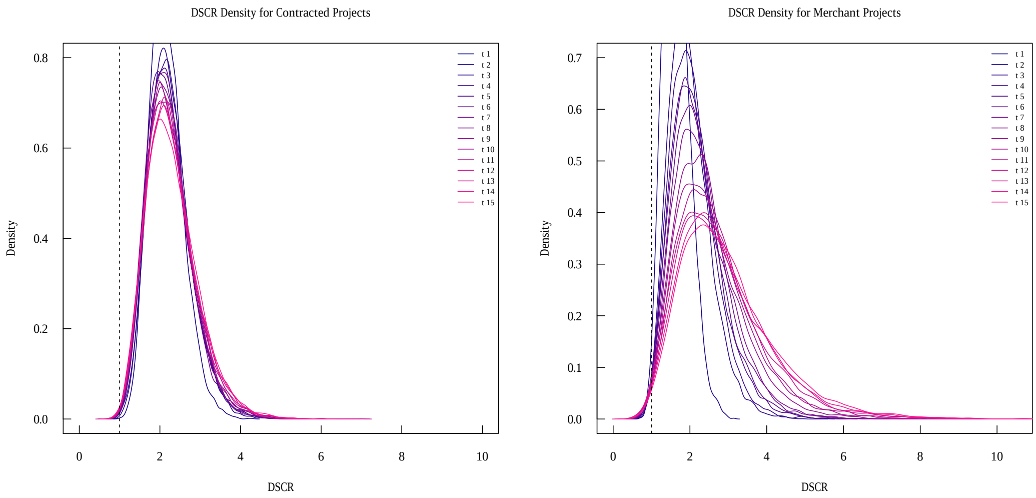

Figure 4: DSCR trajectories in the state (m, σ) plane, for both families.

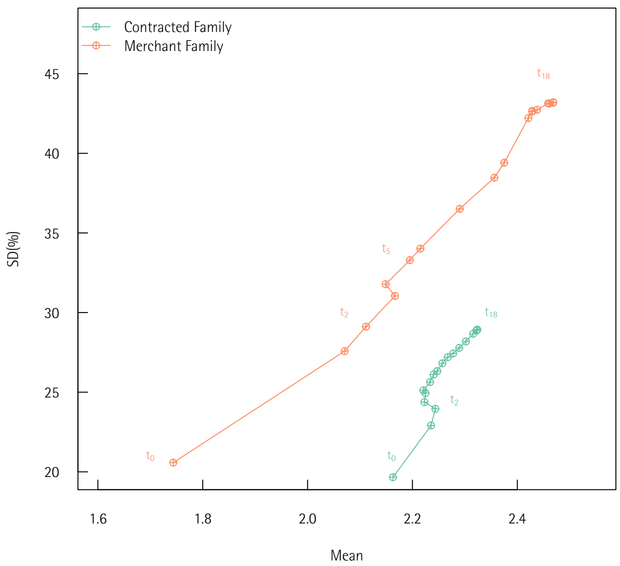

type structure with relatively high DSCR volatility, but later on undergoes a large downward shift in its realised DSCR level, e.g. as the result of a negative demand shock, while its DSCR realised volatility from that point onwards also decreases markedly.

A concrete case of such a trajectory could be a toll road experiencing significant loss of traffic after an economic recession, but for which the residual "baseload" traffic is much less volatile than before the shock, and still sufficiently high to keep the DSCR out of the default state.

Such a project would not be adequately captured by the mean DSCR process of its original family, even though this was the

Figure 5: Probabilities of hard and soft defaults for contracted and merchant families; computed as the probabilities of $DSCR_t$ falling below 1.0 and 1.05, respectively.

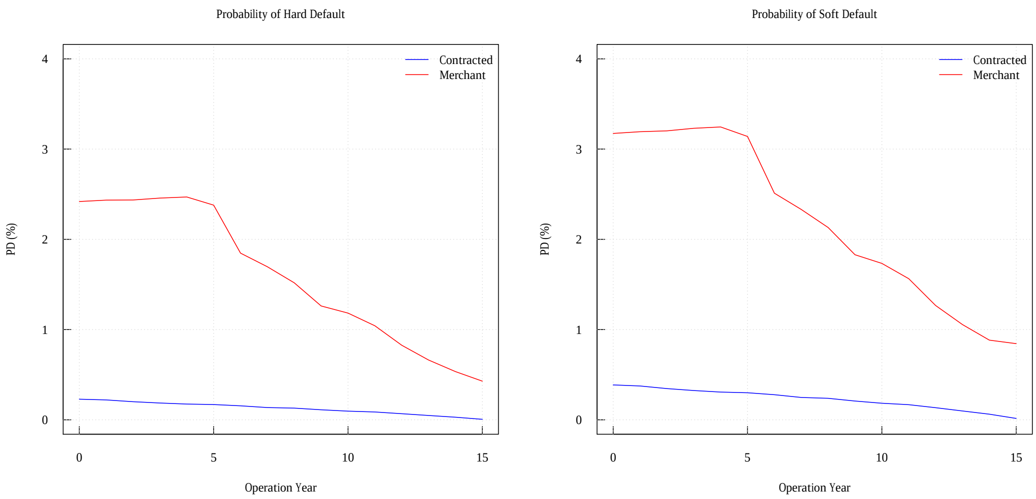

Figure 6: Filtered DSCR quantiles and standard deviation for a single project experiencing a negative shock.

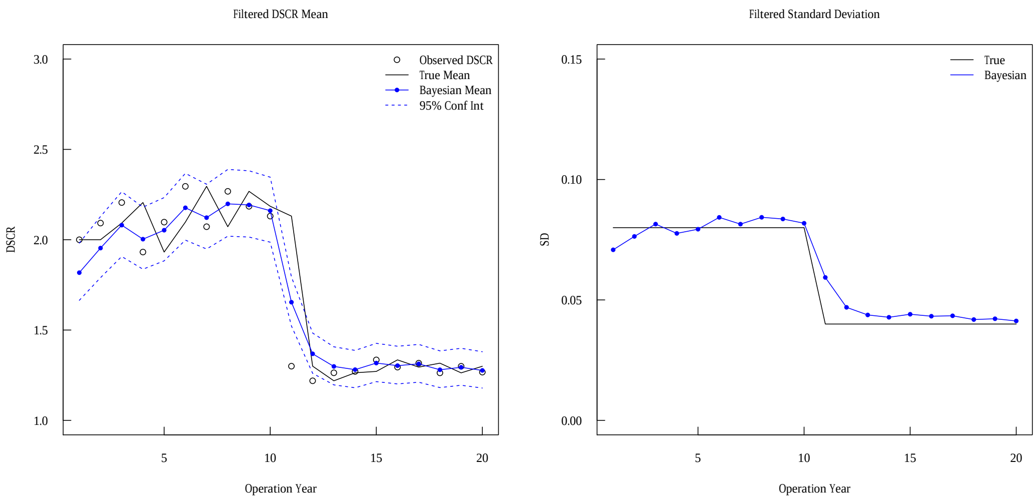

best available starting point to anticipate its behaviour at t_0.

In this illustration, we know the "true" underlying DSCR process that is otherwise unobservable, and how it is impacted by the negative demand shock. The point of the example is to show that as we observe realised information for individual investments, our estimates of the true process can quickly converge to the true value and then track it as it evolves during the life of the investment.

Figure 6 shows the filtered DSCR mean and standard deviation along with the realised DSCR values and the true standard deviation of the project. As soon as the DSCR diverges from its original trajectory the model takes this new information into account, and if the divergence persists, future estimates of the expected value of $DSCR_t$ are updated

Figure 7: Project's distance to lockup, soft, and hard defaults and the corresponding probabilities of lockup, soft, and hard defaults, estimated using filtered DSCR mean and volatility.

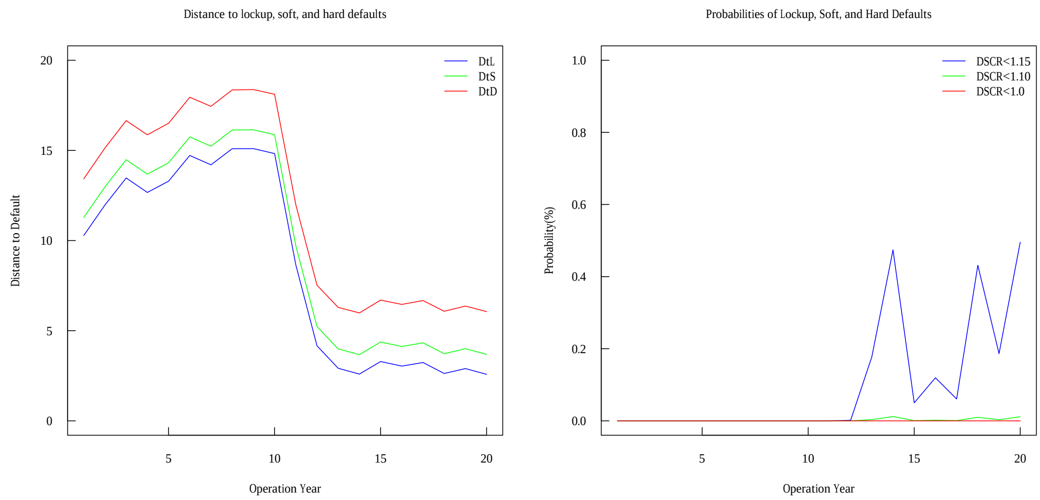

accordingly. Likewise, initial estimate of the volatility of $DSCR_t$ on the right panel of figure 6 are corrected as information about the lower realised volatility becomes integrated into each posterior value.

The ability to revise the DSCR dynamics of individual projects directly leads to the revision of their risk metrics. For example, figure 7 shows the probabilities of dividend lockup, soft default, and hard default, respectively, and suggests that the negative jump in the DSCR, combined with lower realised volatility of DSCR, has no noticeable effect on the project's probability of hard default, a negligible impact on probability of soft default, but a noticeable impact on the probability of a dividend lockup.

Towards Larger Samples of DSCR Data

This paper shows that a powerful statistical model of credit risk relying on DSCR dynamics can be developed, and provides important insights for the valuation of private credit instruments in infrastructure project finance.

It also militates for standardising the data collection and computation of items such as the debt service cover ratio in infrastructure project finance, and for pooling this information in central repositories where it can be used to create the investment metrics that investors need (and regulators require) to be able to invest in large, illiquid assets such as private infrastructure project debt.

Such analyses will be further developed as new data is collected and standardised to improve our ability to track the DSCR path of individual and groups of infrastructure projects, and increase the number of control variables and the robustness of parameter estimates.

EDHEC is committed to the continued development of this research agenda, both in

Executive Summary

terms of data collection and technological development.

1. Introduction

1. Introduction

1.1 Objectives

The objectives of the paper are to document the statistical characteristics of debt service cover ratios (DSCRs) in infrastructure project finance, and to develop and calibrate a model of DSCR dynamics.

For this purpose, we collect a large sample of realised DSCR observations across a range of infrastructure projects spanning more than 15 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 modelling and predicting future DSCR levels and volatility.

Using these results, we build a model of the conditional probability distribution of DSCRs at each point in the life of infrastructure projects.

1.2 Motivation

1.2.1 A Structural Approach to Credit Risk in Project Finance

In a previous paper (Blanc-Brude et al., 2014), we show that debt service cover ratios can play a pivotal role in the modelling of credit risk in infrastructure project finance.

The DSCR of a firm measures the amount of cash available to make the current period's debt service. The higher the DSCR, the more cash the firm has at its disposal to meet its debt obligations in a given year. The DSCR is written:

$$\mathrm{DSCR}_t = \frac{\mathrm{CFADS}_t}{\mathrm{DS}_t^{BC}} \tag{1.1}$$

where CFADS_t is the cash flow available for debt service, and DS_t^{BC} is the "base case" or current debt service at time t.

Because infrastructure project companies are typically contractually barred from raising additional funds (see Gatti, 2013, section 7.2.3.11.2 on negative covenants), the statistical distribution DSCR provides a direct measure of their probability of default.

A firm can be considered in default if its DSCR falls below 1, (as long as it is computed in such a way that it captures the total cash available for debt service i.e. $\mathrm{DSCR}_t = 1$). This provides an unambiguous definition of the default threshold.

For project companies, the DSCR is often also used to specify so-called "technical" default thresholds in debt covenants (Yescombe, 2002), which provide additional control rights to lenders, such as the power to take over the project company, or restructure its debt if its DSCR falls below a pre-defined minimum level. In this case the default threshold is known to be $DSCR = 1.x$

Understanding the dynamics of DSCRs thus allows implementation of the so-called structural approach to credit risk, since DSCR levels correspond to economically significant default thresholds, the breach of which constitutes an event of default by definition. Modelling the default process from a cash flow perspective also improves on the standard Merton approach to credit risk, which relies on

1. Introduction

total asset and liability values to characterise the mechanism leading to default, thus requiring the use of valuation proxies and making the definition of the default point itself a matter of discussion.

Indeed, Blanc-Brude et al. (2014) show that knowledge of DSCR dynamics is sufficient to estimate the firm's distance to default (DD), which is the workhorse of the Merton approach (Kealhofer, 2003).

Thus,

$$DD_t = \frac{1}{\sigma_{DSCR_t}} \frac{\mathrm{DS}^{\mathrm{BC}}_{t-1}}{\mathrm{DS}^{\mathrm{BC}}_{t}} \left(1 - \frac{1}{DSCR_t}\right) \quad (1.2)$$

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

In the same paper, Blanc-Brude et al. (2014) show that DSCR dynamics are useful to value the whole firm since, following equation 1.1, the free cash flow of the firm can be written:

$$\mathrm{CFADS}_t = \mathrm{DSCR}_t \times \mathrm{DS}^{\mathrm{BC}}_t \quad (1.3)$$

The same relationship holds in expectation and with regard to variance since $\mathrm{DS}^{\mathrm{BC}}_t$ is a constant.

Hence, well-documented DSCR dynamics. combined with future debt service obligations (which are a known constant at the time of valuation), allow for the computation of the expected value and the volatility of the firm's free cash flow at each point in the future. The discounted sum of which is the expected firm value and its volatility.

This is true of infrastructure projects because they derive their value almost entirely from future operating cash flows and do not hold assets that have any value outside of their operation such as intellectual property or financial assets i.e. infrastructure projects are highly relationship-specific.

Using these ideas, Blanc-Brude et al. (2014) develop a structural credit risk model for project finance, adapting the Merton model with debt restructuring post credit-event (Black and Cox, 1976) to the cash flow metrics described above. This allows for the derivation of the *full distribution of future creditor losses* and the computation of standard credit risk metrics such as loss-given default, expected shortfall or effective duration.

1.2.2 The Need for an Advanced Model

Unfortunately, we cannot simply observe DSCR values to calibrate the credit risk approach described above. Indeed, the DSCR process in project finance is mostly undocumented in empirical literature. While DSCR information is routinely collected by the creditors of infrastructure projects, this type of data is typically confidential and not available in large datasets.

The task of collecting DSCR data also runs into one of the main limitations of empirical research in long-term, private investments such as infrastructure: for most observable and investable instruments, a substantial proportion of cash flows remains to be observed today (see Blanc-Brude, 2014,

1. Introduction

for a detailed discussion of the limitations of collecting infrastructure data). It follows that empirical observations alone are not sufficient to document the expected behaviour of infrastructure project cash flows over their entire investment life, and a combination of *ex ante* modelling and empirical observations is necessary.

Finally, private infrastructure investment tends to be characterised by very large individual investments, which almost necessarily leads to poorly diversified portfolios. This suggests that assuming the mean-reversion of an investor's infrastructure debt portfolio may not be realistic, and that idiosyncratic risk should be taken into account.

In particular, individual infrastructure investments can exhibit significant "path dependency" and investors cannot necessarily take for granted the notion that they are exposed to the "median infrastructure project."

For both sets of reasons (data limitations and the importance of firm-specific risk), an adequate model of the DSCR should be able to capture conditional dynamics that explicitly integrate the different credit "states" that an infrastructure project might go through.

This can involve both learning from the data as and when it becomes available, and taking into account the path-dependency of each instrument when estimating future cash flows, instead of assuming a reversion to the population mean.

Existing academic work on credit risk in project finance and infrastructure remains limited and static in nature.[2]

Cohort-based studies of default events such as Moody's (2014) and Standard and Poor's (2013) shed some light on the dynamics of credit risk in project finance loans and report observed events of defaults and reported loss given defaults. These studies document decreasing probabilities of default over project life, and higher recovery rates for project finance loans compared to corporate finance loans.

However, they rely on a "reduced form" approach, representing credit risk as an exogenous process that impacts firms randomly, failing to capture conditionality and state-dependence and effectively assuming no persistence in the evolution of an individual instrument's credit risk. The unconditional nature of such studies does not allow a project's credit risk profile to be updated based on its past performance. Credit risk is implicitly expected to mean-revert irrespective of whether individual projects out- or under-perform *ex ante* expectations.

Still, one such study reports an average time to emergence from default of around 2.3 years (Moody's, 2012), indicating that once a project is in default, it is likely to stay in that state for several periods. This implies state dependence in credit risk dynamics, which is not fully captured by reduced form models.

At the opposite end of the credit research spectrum, individual credit ratings do incor-

2 - Esty (2001, 2002); Esty and Megginson (2003); Blanc-Brude and Strange (2007); Sorge and Gadanecz (2008) study the determinants of project finance loan spreads and find that creditor-rights, syndicate-structure, project leverage, loan covenants, and country risk affect credit spreads, while loan maturity and size do not affect project finance loan spreads the way they affect corporate finance loans. These analyses are static, and do not seek to study the dynamic nature of credit risk in project finance.

1. Introduction

porate new information and aim to provide a revised outlook on the credit quality of rated instruments. They however fail to capture any other trends present in a given population of instruments, which could be revealed by a more quantitative approach.

For the reasons above, in this paper we endeavour to better document the dynamics of DSCRs in infrastructure project finance and build a model of DSCR dynamics using Bayesian inference to describe credit state transitions and to estimate the mean and variance of the DSCR in each state and at each point in an instrument's life.

Next, we discuss a series of hypotheses about the distribution of DSCRs, drawn from the economic and financial literature, and that we will rely on to elicit informed priors of DSCR dynamics.

1.3 Hypotheses

1.3.1 DSCRs and the Financial Literature

The literature on the optimal financial structuring of companies suggests a direct relationship between the DSCR in project finance and the underlying risk profile of the firm.

In finance theory, the financial structure of the firm is determined as an equilibrium contract between various stakeholders of the firm rationally acting in their self-interest, under a given set of firm-specific and economic conditions. As a result, the project company and its financial structure can be seen as an outcome of an optimisation process that maximises some combination of the value of different stakeholders' share of the project company.

Starting from the work of Modigliani and Miller (1958) arguing the irrelevance of financial structure for the value of the firm, this literature has evolved to show that financial structure can affect firm value for several reasons, which may include information asymmetries, agency problems between stakeholders, differential tax treatment of debt and equity payouts, bankruptcy costs etc.

Thus, the question of the optimal financial structure of a project company, as well as the choice of project financing itself is the outcome of this optimisation process maximising firm value under a given set of conditions.

Shah and Thakor (1987) demonstrate the optimality of project financing under asymmetric information for high risk ventures, and show that more leverage is optimal for such projects, as it increases the tax-shield benefits in high-income states, which are all the more all likely for higher risk projects. Thus, in contrast to standard corporate finance literature , they suggest a direct relationship between project riskiness and project leverage. Chemmanur and John (1996) argue that project financing can be optimal even in symmetric-information due to control benefits, and highly leveraged project financing is optimal for projects with higher benefits of control.

1. Introduction

John and John (1991) and Flannery et al. (1993) argue that project financing arrangements result from the firm's efforts to minimise the impact of agency problems (which may lead to underinvestment (Myers, 1977) or asset substitution (Jensen and Meckling, 1976)).

This literature provides some possible relations between a project's characteristics and its DSCR profile that can be empirically tested with our data, and may provide explanations for our findings. At least two hypotheses of interest include:

- The existence of a relation between project leverage and its DSCR;
- The existence of a relation between project DSCR and the firm's overall riskiness.

Next, we formulate our main hypothesis about the existence of homogenous families of DSCR processes.

1.3.2 Families of DSCR Processes

In Blanc-Brude et al. (2014), we described two generic and intuitive types of infrastructure project companies and called them "contracted" and "merchant".

This distinction was informed by the casual observation that the financial structure of infrastructure project finance firms is often such that it requires, at the onset, either a rising or a flat "base case" DSCR profile.

A rising DSCR profile would exhibit both a rising mean and an increasing volatility of $DSCR_t$. That is, creditors would demand a higher DSCR in the future to protect themselves against rising expected volatility of CFADS. Such projects would also have longer "tails" [3] and exhibit between 70% and 80% of initial senior leverage.

We argued that such structuring decisions signalled infrastructure projects that were exposed to commercial risks, such as a power plants selling electricity at market prices, and referred to these projects as *Merchant* infrastructure.

Conversely, we argued that the decision to structure a project while requiring a flatter base case DSCR profile implied the expectation of a much lower and constant conditional volatility of cash flows. We observed that projects with little to no market risk are financed with a flat DSCR base case and also have shorter tails and a higher level of senior leverage, usually around 90%.

Contrary to projects with a rising DSCR, which effectively de-leverage as their lifecycle unfolds, projects with a constant base case DSCR thus stay highly leveraged until the end of the debt's life (otherwise their DSCR would rise). Examples of these projects include social infrastructure projects, such as schools or hospitals that receive a fixed payment from the public sector, or energy projects that benefit from a long-term "take-or-pay" purchase agreement. We called these projects as *Contracted* infrastructure.

Such base case DSCR profiles correspond to what creditors require of each project in the face of uncertainty about their ability

3 - The amount of time between the original loan maturity and the end of the project's life, thus allowing higher recovery rates in the event of restructuring.

1. Introduction

to service their debt, hence they embody creditors ' *ex ante* views about the riskiness of the firm and its debt.

In other words, we argued that the combined decision taken by both creditors and equity sponsors to structure infrastructure projects in a certain manner at their onset is a *quasi-price signal* and can be interpreted as revealing most information available to the different parties at the time of financial close, as well as their risk preferences and market conditions at the time.

In this paper, we determine statistically whether **realised DSCR dynamics** fall into categories determined by the distinctions made above between Contracted and Merchant infrastructure, as well as exogenous conditions at the time of financing and when the data is observed. We then use our results to design a model of DSCR dynamics.

1.4 This Paper

The rest of this paper is structured thus: in Chapter 2, we describe the data collection process and provide a number of initial findings based on descriptive statistics, non-parametric tests and regression analyses. We conclude that the nature of the data requires combining a dynamic state transition model with a particle filtering approach to infer the parameters of the DSCR distribution in each state, at each point in time.

Chapter 3 describes our modelling approach, from the filtering of the DSCR distribution parameters to the definition of state transition probabilities between different DSCR regimes. These ideas are developed and implemented using our dataset in Chapters 4 and 5, respectively.

Finally, in Chapter 6, we examine the implications of our findings in terms of conditional credit risk measures in infrastructure debt at the onset of a project and how it may be updated in time as realised DSCR states become known for individual instruments.

Chapter 7 concludes.

2. Data Collection and Analysis

2. Data Collection and Analysis

In this chapter, we describe the process by which our DSCR dataset was created in section 2.1 and provide a series of descriptive statistics, non-parametric and goodness-of-fit test results in section 2.2. Next, section 2.3 describes a preliminary analysis using standard linear regression techniques, while section 2.4 discusses these first results and the most appropriate way to model DSCR dynamics going forward.

2.1 DSCR Sources and Computations

Our dataset of *realised* DSCRs is built using manually collected data and verified from the audited statements of accounts of several hundred project companies, as well as DSCR data reported by private contributors to the EDHEC Infrastructure Investment Database. [4]

4 - See edhec.infrastructure.institute for more details.

The data obtained from company filings includes information reported in companies' balance sheets, income and cash flow statements, and allows DSCRs, leverage and other financial variables to be computed in each period. Privately contributed data comes from creditors and investors who have access to realised DSCR data and information about realised and future debt service.

Approximately 90 per cent of the total sample was thus obtained from audited accounts, while the remainder was sourced from private contributors. For the purpose of describing DSCR dynamics, the dataset also includes individual project size, leverage, industrial sectors, countries, revenue risk family (as discussed in chapter 1) and date schedules including financial close and construction completion dates.

DSCRs are thus either reported directly or computed using a combination of cash at bank, cash from operating activities, cash from investing activities and cash from debt and equity drawdowns in each period. The data is reported annually.

Project companies reporting their annual accounts do not consistently provide cash flow statements, hence cash flows from operations, investment and financing activities are estimated employing the "indirect method", which employs changes in accrual accounts and operating profit to estimate cash flows from operations. Similar approaches were employed for the cash flows from investing and financing activities. Further details of this procedure are provided in section 8 of the Appendix.

Whilst the application of the indirect method is straightforward, issues did arise when applying this estimation procedure. These included changes in accounting standards, annual reports with balance sheets not balancing, notes to financial accounts not reconciling with the balance sheet, and the inconsistent application of accounting standards across similar projects at the same time.

When these and similar issues were encountered, judgement was applied to ensure that the resulting cash flow statements were consistent with an understanding of the

2. Data Collection and Analysis

underlying project's cash flows. To ensure consistency, the cash flows for all projects, even projects that provided a cash flow statement, were estimated in a consistent manner.

Having estimated the cash flows, the variables of interest in this paper include cash flow available for debt service (CFADS), and debt service.

2.1.1 Computing DSCRs

We argued in section 1 that DSCRs provide key insights into the credit risk of infrastructure projects, as they correspond to an economically significant measure of the ability of the firm to service its debt at each point in time. In practice, however, DSCRs may not always be calculated to reflect all the cash available for the debt service, and can fail to provide economically meaningful information.

Some of the commonly used definitions in the literature (Ciochetti et al., 2003; Harris and Raviv, 1990) use a project's operating income to compute the DSCR. However, such a definition can under/over-estimate 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 pinpoint a meaningful proxy of credit risk, we examine several possible definitions of the DSCR of the firm: we compute the DSCR under three different definitions

$$\mathrm{DSCR}_1 = \frac{C_{bank} + C_{op} + C_{IA}}{DS_{senior}} \tag{2.1}$$

$$\mathrm{DSCR}_2 = \frac{C_{bank} + C_{op} + C_{IA} + C_{dd}}{DS_{senior}} \tag{2.2}$$

$$\mathrm{DSCR}_3 = \frac{C_{bank} + C_{op} + C_{IA} + C_{dd} - C_{inv}}{DS_{senior}}, \tag{2.3}$$

where C_{bank}, C_{op}, C_{IA}, C_{dd}, and C_{inv} denote cash at bank, cash from operating activities, cash withdrawal from investment account, cash from debt drawdowns, and cash invested physical investments, respectively.

For all three definitions DS_{senior} denotes the total debt service, and it consists of interest and principal payments of senior loans and bonds. Figure 8 shows histograms of our dataset for each definitions of the DSCR.

The first definition only takes into account the beginning-of-period cash at bank, and any amount withdrawn from the investment account in computing the CFADS. This definition, however, underestimates the DSCR in the first few years of the project, mostly during construction years, when the project company may have negative operating cash, and may be making debt payments by drawing down more debt.

The second definition remedies this shortcoming by adding the debt drawn in that period to the CFADS calculation. This second

2. Data Collection and Analysis

definition, however, may overestimate the CFADS if a significant fraction of the amount of debt drawn is invested into the project.

Therefore, we subtract the project investment from the CFADS calculation in the third definition. This third definition is the most reasonable from an economic standpoint, and we use it as our default definition in what follows. [5]

2.2 Descriptive Statistics

In this section, we report summary statistics for realised DSCR observations under our preferred definition of the DSCR, as discussed above.

Our data sample consists of a total of 207 projects spanning the two revenue risk families (contracted and merchant), in seven sectors [6], from the early 1990s to 2015, as Figure 9 illustrates. However the sample is limited in size until the late 1990s and incomplete for 2015, and so we report results for the 1999-2014 period in the rest of this paper.

Table 1 presents a breakdown of the firms reporting their DSCR observations by revenue risk family, sector and region. All data for contracted infrastructure is sourced from the UK, while DSCR data for merchant infrastructure represents a combination of UK, US and a number of other OECD countries. [7] We note that contracted projects have a lower standard deviation, skewness, and kurtosis of their realised DSCR than merchant ones.

Raw DSCR data contains some large outliers, especially on the upside. These often arise during the early years of the project life when bonds are used to finance projects. The issuing of bonds significantly increases the CFADS in that period. The proceeds of the bond issue are only invested in the 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.

This makes mean and standard deviation calculations less reliable as they are significantly affected by such large outliers.

We also note that for such high values of the DSCR, the ability of the firm to repay its debt in the current, or the next period must, *ceteris paribus*, be very close to certainty. We return to this notion of a "safe" state for project debt in the conclusion of this chapter and use it in the rest of the paper.

For the purpose of describing the data, we compute mean and standard deviation of the DSCR between the 10th and 90th quantiles of the empirical distribution. This allows the effect of outliers to be avoided, while still observing the volatility in the DSCR data around its median level, as illustrated by Figure 10.

Next, Table 2 shows the number of DSCR observations, median and standard deviation in investment time after operations start (post-construction). Again, contracted infrastructure projects exhibit a

5 - We also used the following three definitions of the DSCR: $DSCR_4 = \frac{C_{op}+C_{IA}+C_{dd}-C_{inv}}{DS_{senior}}$, $DSCR_5 = \frac{C_{op}}{DS_{senior}}$, and $DSCR_6 = \frac{C_{bank}+C_{op}}{DS_{senior}}$, but $DSCR_3$ remained the most economically meaningful definition.

6 - Transportation, Telecoms, Oil & Gas, Industrial, Government Services, Environmental Services and Energy

7 - These countries are, Austria, the Virgin Islands (US), France, Finland, Slovakia, Poland, Germany, Ireland, Canada and the Netherlands

Figure 8: DSCR histograms for the three definitions of the DSCR given in Equations 2.1 to 2.3, respectively, over the period 1984-2015.

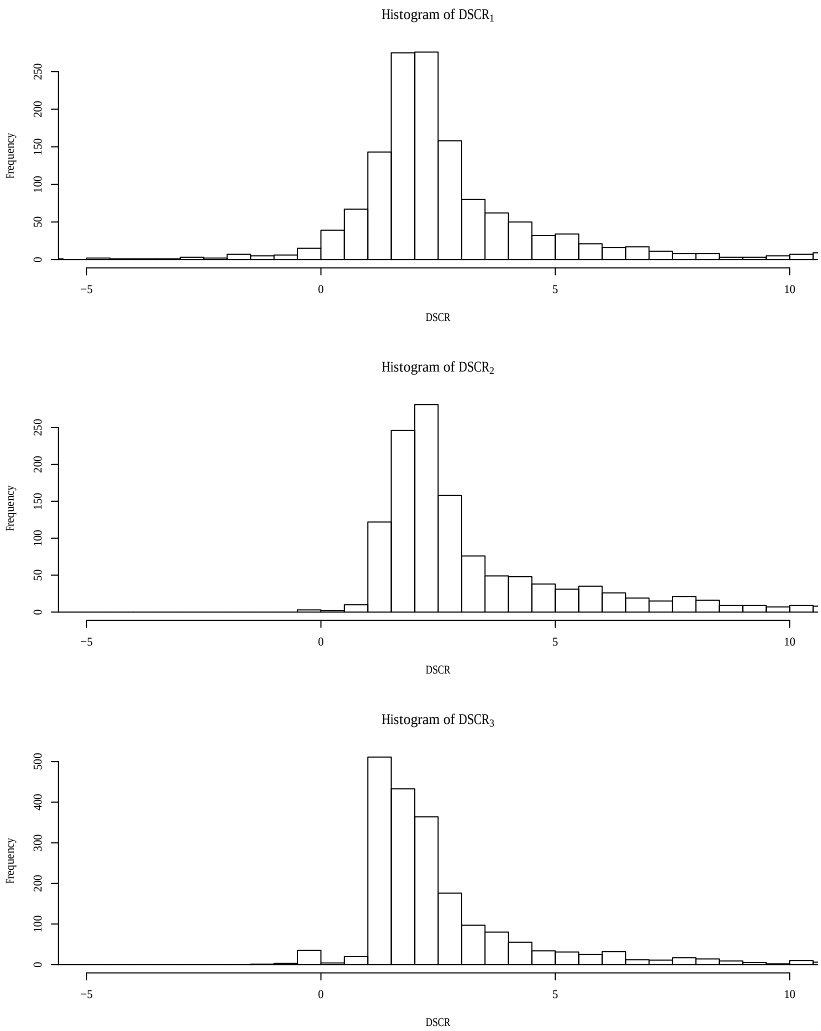

Table 1: DSCR summary statistics by revenue risk family and region

	Obs	1st and 2nd Moments			Higher Moments		Extremes	
	N	*Median*	*Mean*	*SD*	*Skew*	*Kur*	*10th-Q*	*90th-Q*
Contracted	1,001	2.269	2.311	0.851	1.109	4.618	1.406	5.351
Merchant	428	2.015	2.987	3.533	3.423	16.664	1.071	23.824
UK	1,403	2.383	2.692	1.738	1.974	7.131	1.394	10.918
US	121	1.480	1.177	0.218	−0.086	3.051	1.258	2.000
Other	11	1.174	1.898	2.230	2.471	7.114	1.074	10.364
Total	1,429	2.225	2.350	1.128	1.557	5.315	1.304	6.480

more stable median and standard deviation of the DSCR over time than merchant ones.

Tables 17 in the Appendix provide a more detailed breakdown of DSCR statistics by family of revenue risk. DSCR observations for contracted projects not only have a lower standard deviation than merchant ones, they also possess thinner tails, as evidenced by their low values for the 80th and 90th quantiles. The DSCR distribution

2. Data Collection and Analysis

Figure 9: DSCR and cash flow observations by reporting firm (investment start year=financial close)

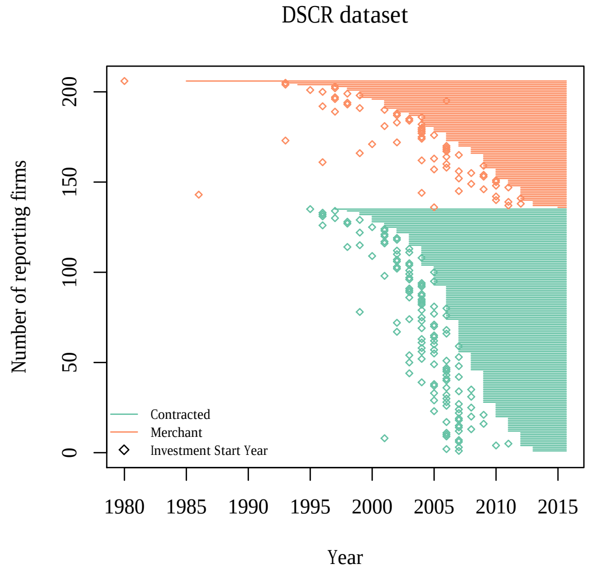

Table 2: Post-construction realised DSCR statistics, 1999-2014.

	Contracted			Merchant		
Year	*N*	*Median*	*SD*	*N*	*Median*	*SD*
1	108	2.397	0.933	45	2.130	2.108
2	121	2.255	0.479	54	1.836	0.897
3	121	2.259	0.380	54	1.828	0.491
4	115	2.058	0.300	48	1.749	0.660
5	112	2.148	0.292	42	1.685	0.517
6	97	2.229	0.250	35	1.994	0.985
7	80	2.248	0.332	30	1.732	0.921
8	72	2.209	0.375	22	2.483	2.215
9	50	2.374	0.382	18	2.139	2.075
10	36	2.289	0.427	14	1.855	0.850
11	25	2.279	0.936	11	1.481	1.112
12	15	3.109	0.846	9	1.830	1.435
13	13	2.498	1.494	5	2.347	0.365
14	10	3.385	1.807	9	2.223	0.605
15	8	5.279	1.626	4	16.550	18.796
16	6	5.227	0.501	5	2.812	0.734

Figure 10: Post-construction realised DSCR for contracted and merchant projects, with (left panel) and without outliers (right panel).

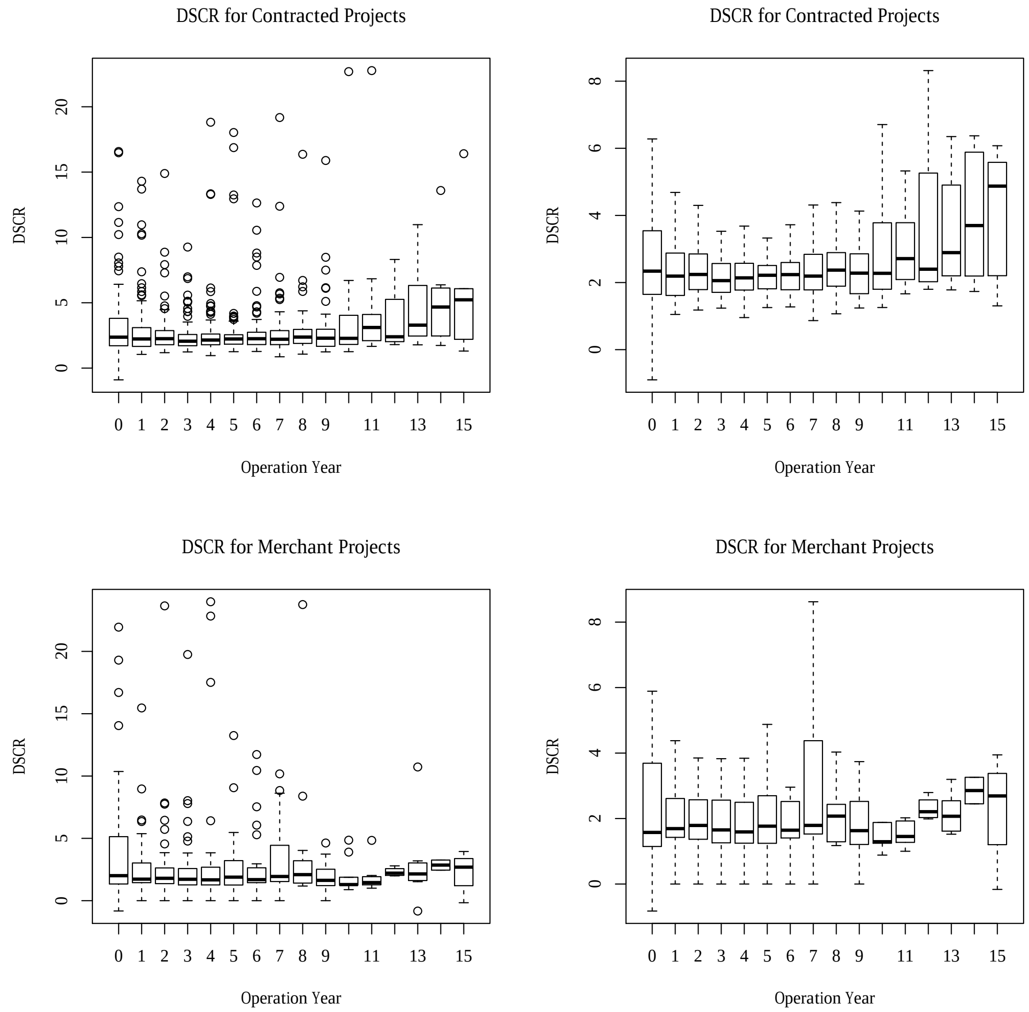

for contracted projects tends to be more concentrated around its median.

In both types of projects, we note that DSCRs tend to increase during the investment life. This effect is partly attributable to the fact that the projects financed in earlier periods tend to have higher realised DSCRs and partly because DSCRs tend to increase during investment life. We return to these two effects in section 2.3 when we conduct panel regressions with fixed time-effects.

2.2.1 Difference of mean and variance between DSCR families

Next, we test whether DSCRs follow similar distributions in contracted and merchant infrastructure. Figure 11, which shows empirical densities of realised DSCRs for the two revenue risk families, suggests that this is not the case.

Table 3 provides the results of non-parametric tests of the null hypothesis that both contracted and merchant infrastructure have the same mean and variance.

2. Data Collection and Analysis

Figure 11: Empirical densities of DSCRs in contracted and merchant infrastructure, 1999-2014.

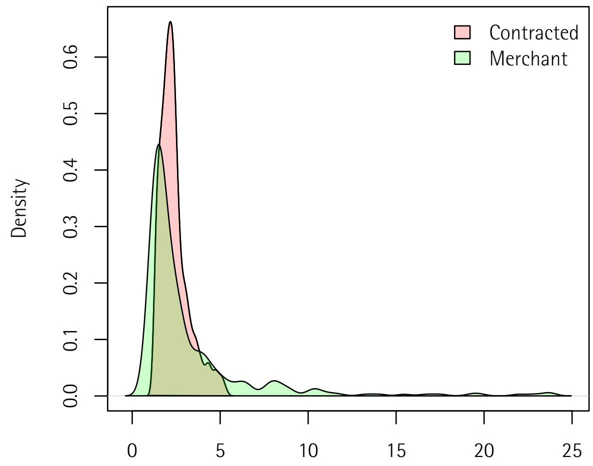

Figure 12: Lognormal fit to the DSCR distribution for contracted projects.

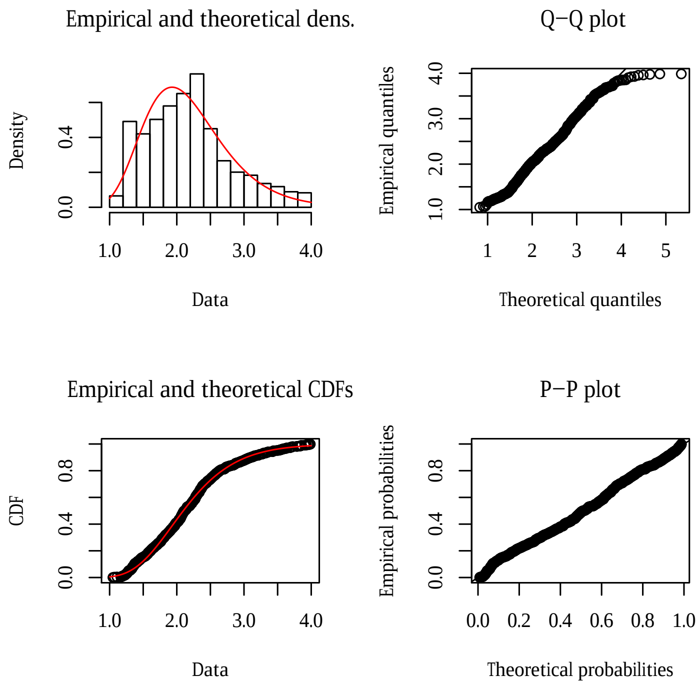

The resulting Mann-Whiteney, Bartlett, and Kolmogorov-Smirnov test statistics lead to the conclusion that the null hypothesis can be rejected with a high degree confidence (low p-value) as a result, we conclude contracted and merchant DSCR processes possess different mean and variance processes.

2.2.2 Goodness-of-fit

Next, we explore the goodness-of-fit of the data with respect to standard probability density functions. The most natural

2. Data Collection and Analysis

Figure 13: Lognormal fit to the DSCR distribution for merchant projects.

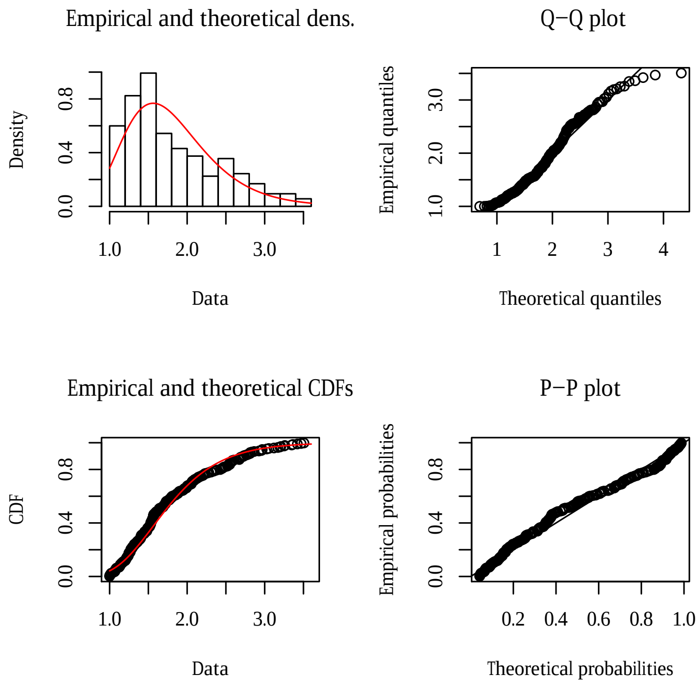

Table 3: Non-parametric test results for the difference in mean and variance between the contracted and merchant DSCR sub-samples. Mann-Whitney, Bartlett, and Kolmogorov-Smirnov tests are non-parametric tests for the equality of means, variances and distributions of two samples, respectively.

	Mann-Whitney		Bartlett		Kolmogorov-Smirnov	
	Statistic	*p-value*	*Statistic*	*p-value*	*Statistic*	*p-value*
Contracted-Merchant	186,065	0.001	1,169	0.000	0.18256	0.000

Table 4: Estimated DSCR distribution parameters and the corresponding p-values for the chi-squared goodness-of-fit test. μ and σ denote the estimated log-mean and log-standard deviation of the lognormal distribution.

	μ		σ				χ^2 fit	
	Estimate	*Error*	*Estimate*	*Error*	*Mean*	*SD*	*Statistic*	*p-value*
Contracted	0.74	0.010	0.289	0.007	2.186	0.644	15.011	0.682
Merchant	0.55	0.019	0.316	0.014	1.815	0.589	9.946	0.448

Table 5: Quantile threshold of observed DSCRs below which the data can be fitted to a lognormal distribution, at each point in investment time

	Contracted				Merchant			
	Quantile	*DSCR*	*Statistic*	*p-value*	*Quantile*	*DSCR*	*Statistic*	*p-value*
1	0.900	8.184	0.064	0.814	0.850	11.838	0.070	0.706
2	0.900	5.610	0.076	0.609	0.850	6.356	0.132	0.162
3	0.900	3.975	0.077	0.589	0.850	3.885	0.068	0.746
4	0.900	3.488	0.060	0.869	0.850	5.117	0.046	0.683
5	0.900	4.260	0.062	0.839	0.850	3.346	0.082	0.514
6	0.900	3.434	0.101	0.255	0.850	5.411	0.038	0.799
7	0.900	4.342	0.132	0.617	0.850	7.018	0.156	0.259
8	0.900	4.302	0.120	0.214	0.850	9.975	0.067	0.753
9	0.900	3.719	0.114	0.246	0.850	7.736	0.066	0.781
10	0.900	6.131	0.059	0.874	0.850	8.163	0.108	0.195
11	0.900	6.296	0.049	0.971	0.850	6.436	0.124	0.395
12	0.900	6.231	0.090	0.386	0.850	10.069	0.066	0.778
13	0.900	7.990	0.064	0.802	0.850	6.178	0.055	0.926
14	0.900	12.833	0.071	0.702	0.850	9.222	0.076	0.606
15	0.900	9.939	0.079	0.562	0.850	10.341	0.154	0.277

2. Data Collection and Analysis

candidate for the DSCR distribution is the exponential family, which includes lognormal, gamma, weibull, and exponential distributions. This is because the DSCR observations mostly lie between $[0, \infty]$, with a uni-modal distribution and with a peak around 2.

However, due to heavy right-hand tails in realised DSCR observations, neither one of these distributions achieves a high goodness-of-fit when fitted over the whole sample. Tables 17 in the appendix confirms that the 10th, 20th and 50th quantiles of our DSCR observations lie much closer to each other than the 50th, 80th, and 90th quantiles.

This is, again, a reflection of the fact that DSCR can take very high values that appear to suggest an altogether different (and lower) level of credit risk in infrastructure projects in some states of the world.

Since our aim is to build a parsimonious model of DSCR dynamics, we test different thresholds below which the data does exhibit a high goodness-of-fit with exponential functions, in particular with the Lognormal function, since its calibration is analytically tractable using Bayesian inference.[8]

8 - The likelihood of the data is conjugate with the prior distribution of its parameters.

Table 4 shows that truncating the sample at the 90th and 85th quantiles for contracted and merchant families, respectively, achieves a very high goodness-of-fit with the lognormal density function i.e. the log of DSCR observations at each point in time follows a Gaussian process. As Table 5 illustrates, below these thresholds, we cannot reject the null hypothesis that DSCRs follow a lognormal distribution at each point in time t in the investment life. Figures 12 and 13 also confirm a high degree of Gaussian fit for logs of realised DSCR values up the relevant thresholds.

It follows that, as originally suggested, but not documented in Blanc-Brude et al. (2014), we can reasonably assume that **DSCRs in infrastructure project finance follow a lognormal process.**

Next, we examine the role of several candidate factors explaining DSCR variability.

2.2.3 DSCR determinants

Sectors

Table 6 breaks down DSCR observations by sectors and revenue risk family. Contracted infrastructure has a lower volatility than merchant projects irrespective of the sector. However, some variation in the level of volatility of realised DSCR exists within families, across different sectors. For instance, the median and standard deviation of the DSCR is higher for the transport sector than the social infrastructure sector, for both contracted and merchant projects.

Overall, industrial sectors seem less relevant than revenue risk families to explain DSCR levels and volatility. We test for the statistical significance of sector variables in section 2.3 using panel regression techniques.

2. Data Collection and Analysis

Table 6: Realised DSCR statistics post-construction by sectors for contracted and merchant infrastructure, 1999-2014

	Contracted			Merchant		
Sector	*N*	*Median*	*SD*	*N*	*Median*	*SD*
Commercial and Industrial	22	2.638	0.279	–	–	–
Energy	24	1.719	0.357	258	1.644	0.579
Environmental Services	6	1.461	0.010	8	2.597	0.854
Government Services	820	2.207	0.349	2	1.393	–
Oil and Gas	7	3.527	0.217	57	3.738	7.168
Telecom	9	2.655	0.861	9	2.968	1.182
Transport	113	3.133	1.423	94	2.771	1.796

Figure 14: Scatter plots of realised DSCRs, project leverage and size

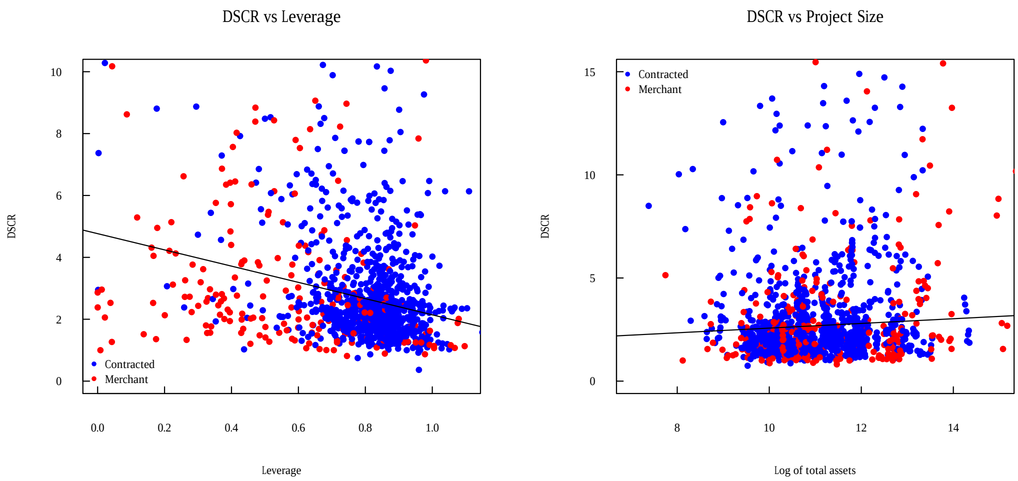

Table 7: Leverage and project size by business model, for all observations over the period of 1999-2014.

	Contracted				Merchant			
	N	*10th-Q*	*Median*	*90th-Q*	*N*	*10th-Q*	*Median*	*90th-Q*
Leverage	1,255	0.659	0.848	0.954	325	0.090	0.581	0.916
Size	1,301	9.532	10.823	12.543	819	7.806	10.445	12.823

Table 8: Realised DSCR statistics post operation start by project leverage and size for Contracted and merchant projects, 1999-2014.

	Contracted			Merchant		
	N	*Median*	*SD*	*N*	*Median*	*SD*
Leverage Below Median	509	2.284	0.521	227.000	2.814	3.136
Leverage Above Median	438	2.278	0.308	37.000	2.382	0.624
Size Below Median	491	2.231	0.475	127.000	2.727	1.280
Size Above Median	456	2.307	0.354	171.000	2.794	2.809

Leverage and size

Table 7 shows the impact of project leverage and size for the two families. Project size is proxied by the maximum of the total value of the project's assets over the project's life, and leverage is computed as the maximum of the ratio of total outstanding senior liabilities to total assets over the project's life.

While project size has roughly similar distributions for the two families, project leverage is very different. Contracted projects not only have a higher median leverage compared to merchant projects,

2. Data Collection and Analysis

their leverage is also more concentrated. This finding is in-line with the hypothesis put forward in section 1, according to which in project finance, high leverage is generally associated with low asset risk, and therefore that contracted infrastructure should, on average, possess a higher leverage than its merchant counterpart.

Table 8 looks at whether realised DSCRs vary systematically by project leverage and size and suggests that in the case of contracted infrastructure, leverage and DSCRs are only related insofar as higher leverage tends to be associated with lower standard deviation of realised DSCR. Median realised DSCRs are otherwise unchanged by the relative change in leverage for contracted infrastructure. The effect of size on realised DSCRs in contracted projects is more ambivalent since larger projects tend to have higher but less volatile realised DSCRs.

This can also be seen in Figure 14, which shows a negative relation between the DSCR and leverage, and a slightly positive relation between the DSCR and total assets. This suggests that DSCRs have a negative relation with project leverage, and that highly leveraged contracted projects tend to have the lowest realised DSCRs.

For merchant infrastructure, Table 8 shows that both higher leverage and size seem to correspond to higher levels of realised DSCRs. Again, high leverage is also associated with much lower volatility of DSCRs, suggesting a direct (and inverse) relationship between leverage and credit risk in project finance, as previously argued in the literature.

Next, we conduct more robust tests using regression analysis.

2.3 Regression analysis

We now use regression analysis to further investigate the explanatory power of different variables in determining the level of realised DSCR. We consider determinants such as revenue risk family, sector, region, leverage and size, financial ratios and project and calendar years.

2.3.1 OLS Model

From prior knowledge and the descriptive statistics presented above, we know that DSCRs have dynamic profiles in project investment time, while they may be impacted in the cross-section by project-specific or macro factors. We first test a simple ordinary-least squares (OLS) regression model using dummy variables for project time "buckets" (each bucket corresponding to a 5-year period) and calendar year dummies to proxy for the business cycle.

Table 24 in the appendix reports the results. We find that in the contracted family, calendar years have no statistically significant impact on realised DSCRs whereas they do in merchant projects, even though this is not true across all sectors, with energy and oil & gas DSCRs apparently the most affected by the business cycle. In these two sectors, realised DSCRs also exhibit statisti-

2. Data Collection and Analysis

cally significant and continuous growth in project time (buckets).

Across the board, leverage is a significant and negative explanatory factor for the level of realised DSCR, which confirms our initial hypothesis that the decision to structure projects with higher leverage, signals lower asset risk.

Indeed, while we argued that the initial required or "base case" DSCR is low (leverage is high) when creditors anticipate low credit risk, a high *realised* DSCR suggests that the investment turns out to be highly capable of repaying its creditors, especially as these high realised DSCRs are associated with low realised volatility of DSCRs, which is the case for contracted infrastructure.

We also report in Table 9 that our DSCR data exhibits non-constant variance (the Breush-Pagan test rejects the null hypothesis of homoskedasticity), significant auto-correlation (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.

2.3.2 Panel regressions

Next, to control for the impact of a project time separately, we fit a panel regression model using fixed effects for project time[9] while calendar years and sectors are controlled for in the cross-section.

9 - i.e. The model fits a different intercept for each investment year, as opposed to each firm

In the most interesting specification, reported in the Appendix in Table 25, we use initial investment calendar years as explanatory variables in the cross-section instead of contemporary year dummies, which greatly improves the power of the model in the case for the merchant family (50% adjusted-R^2), suggesting a significant impact of market conditions at the time of financing in setting average DSCR trajectories through the initial choice of financial structuring. The impact is much more muted (i.e. not statistically significant in most years) in the case of the contracted family.

Fixed project year effects are reported in Table 26 for both contracted and merchant projects. Contracted infrastructure tends to have constant project year effects with an average realised DSCR above 2 ($exp(0.8)$) for the first 15 years of investment, rising towards the end of the period. Merchant projects exhibit a more dynamic and much higher level of average realised DSCR in project time, oscillating between 2.7 and 8 on average.

Finally, we fit the same project time fixed effect panel model replacing calendar year effects with financial ratios extracted from the accounts of the relevant firms. Table 27 in the appendix reports 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 realised DSCR levels, while Table 28 reports the year fixed effects.

2. Data Collection and Analysis

Table 9: Test results for heteroskedasticity, autocorrelation and residuals normality for log(DSCR) ordinary least squares regression.

	Test	Statistic	p-value
Heteroskedasticity	Breusch-Pagan	358.305	0.000
Autocorrelation	Durbin-Watson	0.857	0.000
Normality	Shapiro	0.695	0.000

Profit margin, asset turnover, and cash return on total assets turn out to be highly significant with positive coefficients for the contracted family, and profit margin and asset turnover turn out to be significant for the merchant family. We note the large coefficient for profit margin in contracted infrastructure.

Indeed, the DSCR for contracted projects may be more sensitive to changes in profitability (such as the profit margin) as revenue is largely fixed (contracted), while for merchant projects higher profitability may come at the expense or be the result of lower aggregate revenue leading to cost cutting, leaving the realised DSCR unchanged.

We also note that once we control for annual changes in financial ratios, project year fixed effects cease to be significant except for a rise in realised DSCRs in the later years in the sample of contracted projects and in the earlier years in merchant projects. This is best explained by the fact that the panel is not balanced and suggests that older investments had higher average DSCRs in the contracted family, while more recent merchant projects also tend to have higher DSCRs than the sample average in each investment year.

2.4 Conclusions

In conclusion, the key findings drawn from the descriptive statistics and analysis are as follows:

1. The two DSCR families, contracted and merchant, exhibit different distribution functions, with contracted infrastructure exhibiting a lower DSCR mean and volatility;
2. Realised DSCRs exhibit an investment life dynamic in both families, and tend to go up over the life of the project. This trend is much more pronounced for merchant projects;
3. DSCR realisations up to a fixed quantile can be modelled with a lognormal distribution for each family;
4. Realised DSCRs exhibit statistically significant serial auto-correlation;
5. Regression residuals for realised DSCRs exhibit serial correlations, heteroskedasticity, and non-normality, indicating the need for a time-varying, stochastic model of the DSCR.

Clearly, descriptive statistics and linear regression models provide some insights about the determinants of the DSCRs, but also fail to capture DSCR dynamics in full. This requires taking an auto-regressive and heteroskedastic process into account. Next, we discuss our approach to build a powerful model of DSCR dynamics.

3. Approach and Methodology

3. Approach and Methodology

In this chapter, we describe our proposed approach and methodology to model the DSCR dynamics of infrastructure projects using available realised DSCR observations.

In line with the credit risk framework defined in Blanc-Brude et al. (2014), we aim to characterise the expected value and conditional volatility of DSCRs in infrastructure projects at each point in their investment life. Indeed, as per Equation 1.2, the mean value and the standard deviation of the DSCR at time t are sufficient to compute the distance to default of an infrastructure project.

In what follows, we characterise DSCR dynamics as a "latent process" i.e. an unobservable vector of the DSCR distribution parameters at each point in time, and show that we can estimate the value of these parameters using a filtering approach not unlike the the ones used to determine the physical coordinates of a moving object (Section 3.1). Next, we use a simple state-transition model to represent the full path-dependency of this process, taking into account the possibility of a prolonged state of default as well as a "safe" state, for very high values of the DSCR (Section 3.2).

Next, we detail these two dimensions of our approach before implementing them in turn using our dataset in Chapters 4 and 5, respectively.

3.1 The DSCR as a Latent Process

3.1.1 Intuition

The conclusions of chapter 2 helped confirm our intuition that $DSCR_t$ are serially correlated and can change its risk profile during the investment lifecycle of infrastructure projects. In other words, the expected value $E(DSCR_t)$ and the volatility $\sigma^2 DSCR_t$ are partly determined by the values of the same quantities at time $t-1$, and partly by "innovations" (incremental changes) or shocks happening at time t.

Hence, the *ex post* trajectory of individual projects could correspond to any combination of high/low $E(DSCR_t)$ and high/low $\sigma^2 DSCR_t$, and the DSCR of populations of projects would equally reflect the weighted trajectory of their constituents in a $DSCR_t$ mean/variance plane, as illustrated by figures 15 and 16.

However, the true coordinates (m,σ^2) of the $DSCR_t$ process are unobservable (it is a latent process) and can only be imperfectly measured by observing DSCRs and their realised mean and variance (which are themselves noisy because of measurement errors).

This is compounded by the fact that infrastructure projects financed today, may not follow the exact same future paths as projects financed two decades ago (for instance because choices of financial structure are influenced by market conditions and prudential regulation), and by the paucity of available historical data, so that realised data cannot be assumed to predict the behaviour of new projects very well. In

3. Approach and Methodology

Figure 15: Idealised DSCR trajectory of a broad family in the DSCR mean-variance plane.

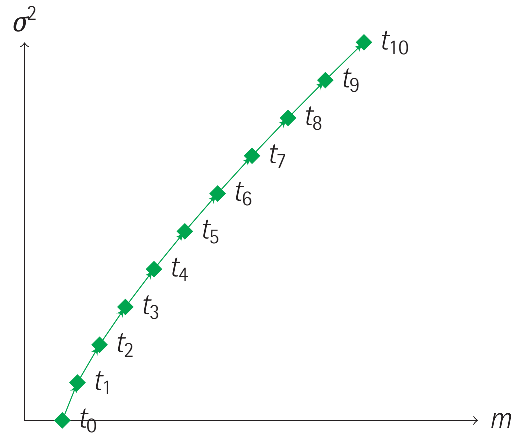

Figure 16: Example DSCR trajectories of individual projects in the DSCR mean-variance plane.

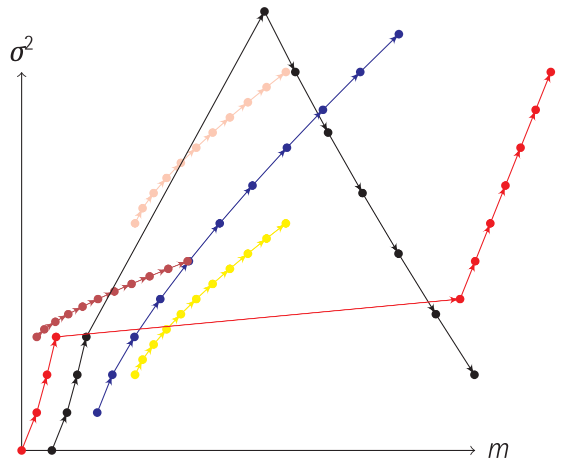

other words, we can only learn so much from realised DSCR data.

Recognising that our knowledge of $DSCR_t$ dynamics remains conditional, we aim to build a model of the future mean and volatility of $DSCR_t$ that is capable of both integrating past information about realised 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 such approaches 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.

Next, we discuss how "particle filtering" models allow us to estimate and predict

3. Approach and Methodology

DSCR dynamics for individual projects or families of projects.

3.1.2 Particle Filtering

We aim to build a time-varying model of the DSCR mean and volatility but can only observe a limited number of realised DSCR observations especially relative to the number of relevant control variables.

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 behaviour.

Latent Process

Say that we can observe certain data Y (here, realised DSCRs) and, as illustrated by figure 17, this "observation process" is related to hidden variables (the latent process X) by some known functional form and the dynamic system describing the evolution of the latent process is known probabilistically.

Here, the latent state of the DSCR system at time t is simply the vector of parameters X_t of its distribution at time t and the latent process can be written:

$$X_t = f(X_{t-1}) + W_t \tag{3.1}$$

This is known as a Markov process and the **state equation** of the system. W_t is a "noise" of the Markov process, indicating potential innovations or shocks in the process outcome, alongside the dynamic imposed by $f(.)$, a known function. Next, the observation process is formally written:

$$Y_t = g(X_t) + V_t \tag{3.2}$$

which is known as the **observation equation**, with V_t, the noise of the observation process and $g(.)$, also a known function.

The filtering problem consists of estimating the values of $X_{t+\tau}$ given the noisy measurement Y at all times $t, t+1, \ldots, t+\tau$, that is, $Pr(X_{t+\tau}|Y_{t+1}, Y_{t+2}, \ldots Y_{t+\tau})$.

In other words, in this very general setting, as long as we can relate our observations in each period to some function of the system's state, as well as the current state estimate to its previous realisation, we can aim to estimate the parameters that define this auto-regressive relationship.

Bayesian Parameter Estimation

Given our modelling objectives are to accommodate small samples, avoid assuming static values for the $DSCR_t$ distribution parameters, and be able to revise any existing parameter estimates once new data becomes available, this process is best estimated iteratively using Bayesian inference techniques.

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, repre-

3. Approach and Methodology

Figure 17: An example of transitions between Markov states. The true state takes values of X_1, X_2, X_3, etc while. Y_1, Y_2, Y_3 are noisy observations of the true state at times t_1, t_2, t_3, etc.

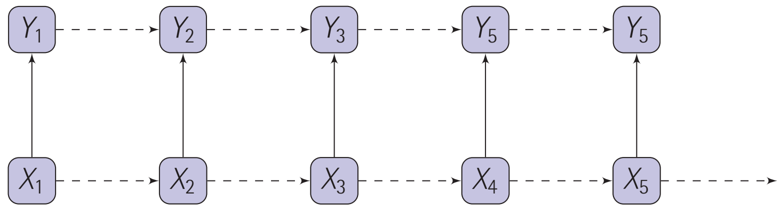

sents our current uncertainty about the true value of this parameter.

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 becomes observable, this *prior* knowledge is updated using Bayesian inference techniques, which we discuss next, to derive a more precise *posterior* probability distribution of $DSCR_t$.

In this case, based on the findings reported in chapter 2, we can assume that the DSCR process follows a lognormal distribution in each period (we address the treatment of outliers in realised DSCRs in the next section).

That is,

$$log(DSCR_t) \sim N(m_t, p_t), \qquad (3.3)$$

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 s**tate equation** is written:

$$X_{t-1} = X_t + W_t \qquad (3.4)$$

In other words, the parameters m_t and p_t of the $DSCR_t$ process are assumed to be autoregressive with one lag and innovation or disturbance W_t.

In a Bayesian setup, unknown paramaters (whether they are stochastic or not) are given a probability distribution. Here, m_t, the mean of the lognormal $DSCR_t$ 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,

$$m_t \sim N(\mu_t, \delta_t) \qquad (3.5)$$

$$p_t \sim \Gamma(\alpha_t, \beta_t) \qquad (3.6)$$

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

As is well documented in the literature, under such parameterisation 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 Lognormal 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

3. Approach and Methodology

parameters a and b) with a Normal distribution (with mean μ and precision δ) of m *conditional on p.*

Next, if the realised DSCR data Y follow a lognormal process of mean m and precision p, its likelihood function is given by:

$$\mathcal{L}(m, p|Y) \propto p^{N/2} e^{\left(-\frac{p}{2}\sum_{n=1}^{N}(\log Y_n - m)^2\right)} \tag{3.7}$$

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, $\bar{Y} = \frac{\sum_{n=1}^{N}\ln(Y_n)}{N}$, 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 :

$$\begin{aligned}
\alpha^+ &= \alpha + \frac{N}{2} \\
\beta^+ &= \left(\frac{1}{\beta} + \frac{SS}{2} + \frac{\delta N(\bar{Y} - \mu))^2}{2(\delta + N)}\right)^{-1} \\
\mu^+ &= \frac{\delta\mu + N\bar{Y}}{\delta + N} \\
\delta^+ &= \delta + N
\end{aligned} \tag{3.8}$$

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 3.8, and the posterior parameters m^+ and p^+ of the distribution of $DSCR_t$ derived, incorporating prior knowledge and the new information.

Particle Filtering

Given any initial belief about the mean and variance of $DSCR_{t_0}$ – drawn for example from the historic project family mean and variance – and assuming a lognormal 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, as described in section 8.4 in the Appendix.

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" i.e. 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 $DSCR_{t_1}$) in the first investment period and compute the likelihood L_i for each on of the 1,000 particles given the data, as per equation 3.7.

Normalised likelihood scores w_i [10] are then used to rank individual particles, which are then resampled by weight i.e. 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 variance of the DSCR given all the DSCR observations. And the distribution of DSCR mean, m, and precision, p, is updated in accordance with 3.8.

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

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

3. Approach and Methodology

Figure 18: Genearting particles using prior knowledge to estimate $DSCR_t$ mean and variance using a particle filter.

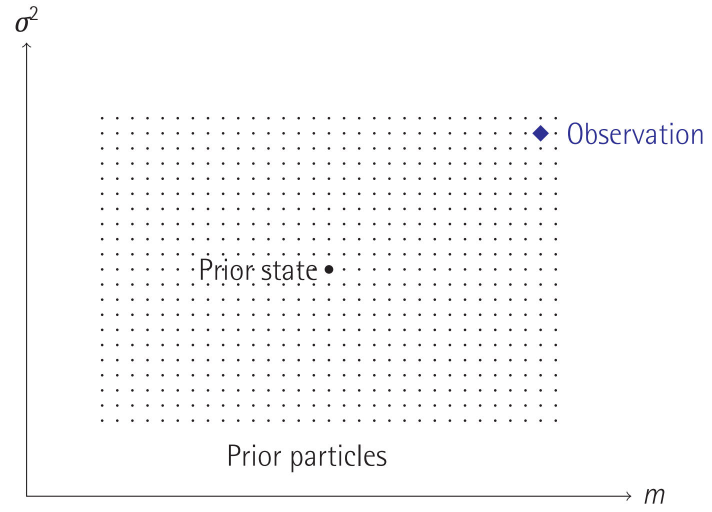

Figure 19: Updating $DSCR_t$ mean and variance estimates using resampled particles.

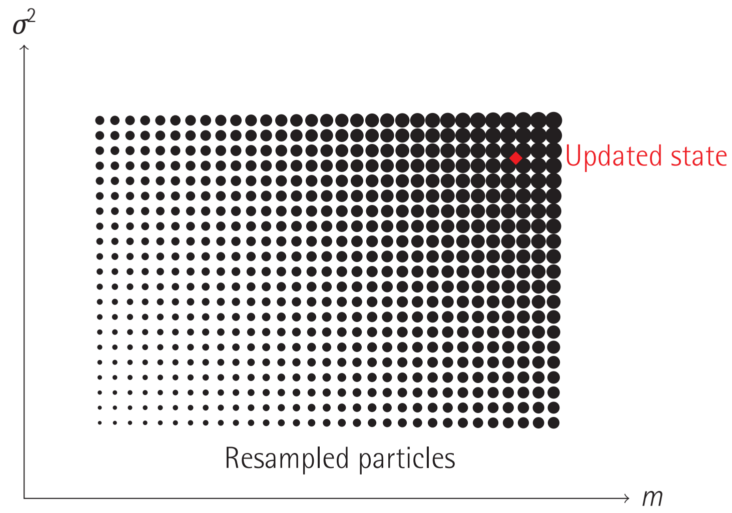

at that time, and the filtering and updating process starts again.

Figures 18 and 19 show this procedure schematically. In Figure 18 particles are generated based on the prior distribution of *m* and *p*. These particles are then resampled using the likelihood of observed DSCR in Figure 19, 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 realised 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. We return to and implement this estimation approach in chapter 4.

3. Approach and Methodology

First, we discuss our proposed treatment of the DSCR process when it ceases to follow a lognormal function.

3.2 DSCR States

3.2.1 Intuition

As well as having time-varying dynamics, we noted in chapters 1 and 2 that projects could stay in default for several periods and that, at the opposite end of the scale, DSCRs could take some very high values, suggesting virtually zero probability of default.

Moreover, these high values were found to curtail the goodness-of-fit of a lognormal density function when applied to the data, whereas both contracted and merchant DSCRs could effectively be considered lognormal if the sample was truncated at a threshold $\overline{DSCR}$ corresponding to the 90th and 85th quantile, respectively.

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

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

An example path from state to state followed by an individual project is illustrated by figure 20.

Hence, once a project's DSCR breaches the *hard default* threshold represented by $DSCR_t = 1$, it enters the default state, which it may or may not leave after a number of periods. In this state, creditors have the option to take over the firm or restructure debt, and they can maximise the value of exercising this option depending on the size of their exit and restructuring costs. They may decide to waive the event of default or engage in negotiations with project sponsor to restructure the firm and its debt or indeed take over the firm and seek another sponsor (see Blanc-Brude et al., 2014, for a formal model).

Hence, the firm may transit out of the default state (into the risky state) with some probability (π_{dr}) at the next period, or stay in this state and again transit out of default at the next period, etc.

In this state, the DSCR process effectively stops (there is no debt service), hence estimating its mean and variance is irrelevant.

In the safe state, on the contrary, the realised 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_{sr}).

In this state, estimating the parameters of the DSCR distribution, in particular estimating its variance, is also irrelevant.

3. Approach and Methodology

Finally, in between the default and safe states, a project's DSCR may take values between 1 and some higher threshold $D\bar{S}CR$. From this state, it may either stay in the *risky* state at the next period, or transit out of it into the state of default "d" or the safe state "s", both described above.

In this state, we know from our results in chapter 2 that if the upper threshold is set at the 85th/90th quantile of our DSCR sample, the data follows a lognormal process, the paramaters of which (position and scale) can be estimated using the particle filtering approach described earlier in section 3.1.

Next, we discuss how state transition probabilities may be estimated and describe the method implemented in chapter 5.

3.2.2 State Transition Probabilities

DSCR States as a Markov Process

Estimating state transition probabilities amounts to estimating the components of a matrix describing another Markov process.

Say for now that the DSCR process can take one of two states S_t at time t: a risky state defined as $DSCR_t \geq 1$ and denoted by $S_t = r$, or a default state such that $DSCR_t < 1$ and denoted by $S_t = d$.

The probability of being in the risky state is defined as $Pr(DSCR_t \geq 1) = Pr(S_t = r) = p_t$ while $Pr(DSCR_t < 1) = Pr(S_t = d) = q_t = 1 - p_t$. q_t is also the probability of default at time t.

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

$$\mathbb{S}_{t+i} = \begin{pmatrix} \pi_{rr} & \pi_{rd} \\ \pi_{dr} & \pi_{dd} \end{pmatrix}$$

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_{rd} is the probability of transiting to the default 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 realised state at $t + i$ is thus written:

$$p_{t+\tau} = p_{t+i}\pi_{rr} + (1 - p_{t+i})(1 - \pi_{dd}) \tag{3.9}$$

And in the matrix notation

$$\begin{bmatrix} p_{t+\tau} \\ q_{t+\tau} \end{bmatrix} = \mathbb{S}_{t+i}. \begin{bmatrix} p_{t+i} \\ q_{t+i} \end{bmatrix} \tag{3.10}$$

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

Hence, staring 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 default states at future periods by successively applying the transition matrix.

3. Approach and Methodology

Figure 20: Illustration of the DSCR path between states

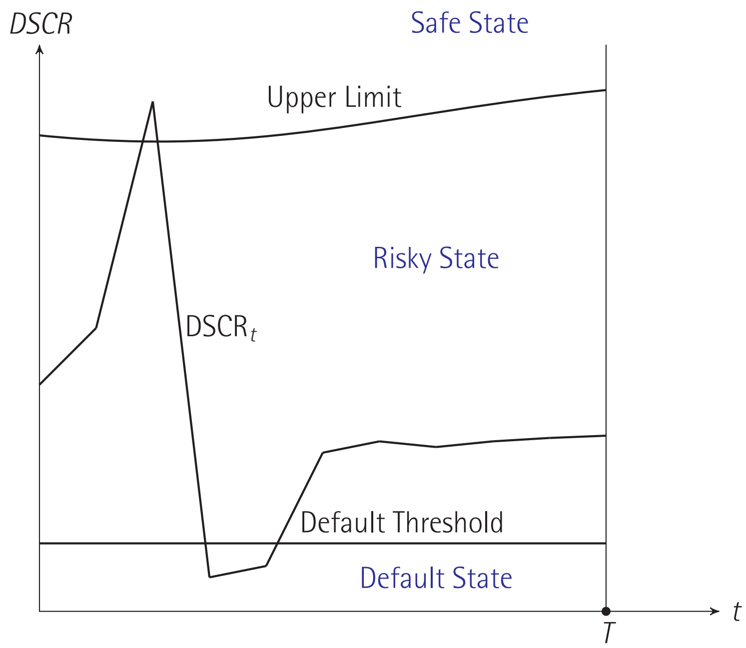

Figure 21: Illustration of the transition probabilities between two Markov states $S_t = d$ and $S_t = r$. The task of the state transition model is to estimate the transition probabilities between the two states.

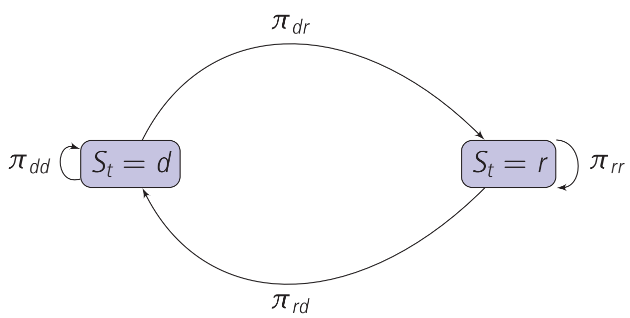

According to equation (3.10), we can know the conditional probabilities of being in the risky or default state in each future period $t + \tau$ by estimating $\mathbb{S}_{t+i}$ across the project lifecycle for $i = 0, \ldots (T - 1)$, as well as initial state conditions.

For example, in the context of a greenfield (new) infrastructure project, **initial conditions** at t_0 are unambiguously set to $\pi_{rr} = p_0 = 1$ and $\pi_{dd} = q_0 = 0$, since a new project cannot start its life in the default state.

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,

$$\hat{\pi}_{i,j} = \frac{n_{i,j}}{\sum_{k=1}^{N} n_{i,k}}, \tag{3.11}$$

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.

3. Approach and Methodology

These models can be used to estimate both time-independent and time-varying transition probabilities. However, since these models rely solely on observed transitions to estimate state transition probabilities, a large number of unbiased time-series observations are required to estimate these probabilities reliably.

In the case of infrastructure projects, as discussed in detail in Blanc-Brude (2014), the frequency of observations at any point in time is too limited, and the number of control variables required to explain the variance of realised DSCRs –from initial to realised market conditions, to project and macro-level characteristics – is too large.

As before, a Bayesian approach to estimating state-transition probabilities is more promising.

In Bayesian Markov switching models, we start with forming a prior belief about transition probabilities between different states, based on information available otherwise (e.g. a project does not start its life in the default state), which is then updated as one observes actual transitions between states.

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

$$\mathbb{S}_t = \begin{pmatrix} \pi_{rr} & 1-\pi_{rr} \\ \pi_{rd} & 1-\pi_{rd} \end{pmatrix}$$

That is, each row of $\mathbb{S}_{t+\tau}$ matrices is equivalent to an independent Bernoulli draw of parameter π_{rr} or π_{rd}, and we only need to estimate π_{rr} and π_{rd} 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" (realised DSCR transitions between two given states), this data (call it Y) follows a binomial distribution (the outcome is binary) with the likelihood:

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

where $\binom{N}{n} = \frac{N!}{n!(N-k)!}$ 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

3. Approach and Methodology

obtained from observing the data.

$$\begin{aligned} p(\pi|Y) &\propto p(\pi)\mathcal{L}(Y|\pi) \\ &\propto \pi^n(1-\pi)^{N-n} \times \pi^{\alpha-1}(1-\pi)^{\beta-1} \\ &\propto Beta(\alpha+n, \beta+N-n) \end{aligned} \tag{3.12}$$

...to a normalising 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 i.e. 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 a 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) amongst N new data points.

The posterior distribution of π_{rr} summarises 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_{rd}, after observing projects transiting from the risky to the default state.

In chapter 5, we implement this approach to estimating state transitions for all three states described above.

3.3 Conclusion

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 lognormal process, as well as the transition probabilities in and out of that state at each point in the project lifecycle.

Moreover, estimating the parameters of the DSCR process outside of the risky state is unnecessary because the other two possible states do not represent any credit risk: the firm is either already in default or its debt is risk-free (in that period).

In the next two chapters, we implement this approach to our dataset.

4. Stochastic DSCR Model

4. Stochastic DSCR Model

In this chapter, we implement our filtering approach of the DSCR distribution parameters discussed in Chapter 3 using data for each family of infrastructure projects.

We report our results in Section 4.1 and some implications in Section 4.2.

4.1 Filtered Values of the DSCR Distribution

As described in Chapter 3, we estimate the DSCR lognormal distribution parameters m and p at each point in time.

We start by assuming prior values for the four meta-parameters — μ, δ, α, β — of the DSCR distribution, and then update them sequentially using the particle filter described earlier.

Our choice of prior values are shown in Figure 22. The mean of the log of DSCR is largely concentrated between 0 and 1 with a mean of about 0.5. This prior distribution for mean of log(DSCR) is consistent with the observed DSCR mean values that lie in the range of 1 to 3. The precision of the log(DSCR) distribution is concentrated between 0.5 and 1.5, which is chosen to be consistent with the observed DSCR standard deviations that generally fall in the range of 0.5 to 2.

Given these choices of prior distributions of meta-parameters, Table 10 shows estimated posterior parameter values, as well as the implied DSCR mean and standard deviation for all periods, for each family. Due to the relatively large number of DSCR observations in the cross section during the first few operational years, the effect of prior values fades away within the first two periods in this case.

Figure 23 shows the estimated mean and 99.5% confidence intervals of $DSCR_t$ for both project families. The parameter estimates follow a similar trend as observed values reported in Chapter 2, but are much smoother, which is a consequence of the increasing precision of parameter estimates. Having filtered out the noise in the observations makes our estimates less sensitive to jumps in the data, while learning from each new observation.

The difference in DSCR dynamics between contracted and merchant projects is also clear in Figure 27, which shows a much more dynamic evolution of both the mean and standard deviation of DSCR in the merchant case. In particular, relative to merchant projects, contracted infrastructure exhibits less DSCR volatility and a constant realised volatility in time, as is also illustrated in Table 11, which shows the per period rate of change of the mean and standard deviation of DSCRs in contracted and merchant infrastructure.

For both families, filtered mean and standard deviation show an increasing trend, but to very different extents. For the contracted family, the DSCR mean increases from about 2.2 to 2.3, and standard deviation increases from about 20 to 30% over 15 years. But for merchant projects, the DSCR mean increases from about 2 to 2.5, while standard deviation increases

4. Stochastic DSCR Model

Figure 22: Prior distributions of meta-parameters of log(DSCR) distribution.

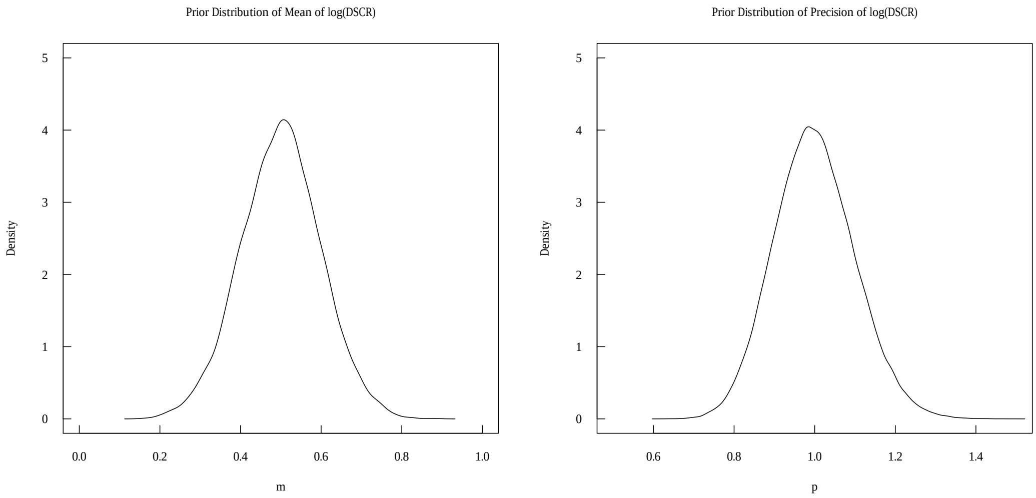

Table 10: Bayesian estimates of DSCR distribution parameters for both families. The initial prior values for μ, δ, α, β are set to be 0.5, 0.1, 500, and 100, and posterior for each period is used as a prior for the next period. M and SD denote mean and standard deviation of the DSCR distribution.

	Contracted						Merchant					
Yr	μ^+	δ^+	α^+	β^+	M	SD	μ^+	δ^+	α^+	β^+	M	SD
1	0.752	194	675	26	2.163	19.663	0.535	111	918	38	1.744	20.578
2	0.779	303	732	37	2.236	22.910	0.691	158	942	69	2.071	27.580
3	0.780	414	789	44	2.244	23.957	0.707	207	966	79	2.111	29.118
4	0.770	519	844	49	2.223	24.377	0.727	250	988	91	2.167	31.045
5	0.770	622	897	54	2.225	24.944	0.717	287	1,006	97	2.149	31.781
6	0.767	712	943	58	2.221	25.116	0.734	319	1,022	107	2.195	33.294
7	0.772	784	981	62	2.234	25.637	0.741	346	1,036	113	2.215	34.017
8	0.774	849	1,015	67	2.240	26.111	0.766	365	1,045	131	2.290	36.515
9	0.776	894	1,038	70	2.247	26.328	0.788	382	1,054	145	2.356	38.473
10	0.779	926	1,055	73	2.257	26.814	0.793	394	1,060	153	2.375	39.412
11	0.783	946	1,067	76	2.267	27.197	0.802	404	1,065	175	2.421	42.212
12	0.787	959	1,074	78	2.278	27.433	0.804	412	1,069	179	2.428	42.663
13	0.791	969	1,080	80	2.289	27.783	0.804	416	1,071	179	2.429	42.624
14	0.796	976	1,084	83	2.302	28.189	0.807	423	1,074	180	2.438	42.734
15	0.800	981	1,088	86	2.316	28.667	0.814	426	1,076	184	2.459	43.140

Figure 23: Filtered DSCR values for contracted and merchant families.

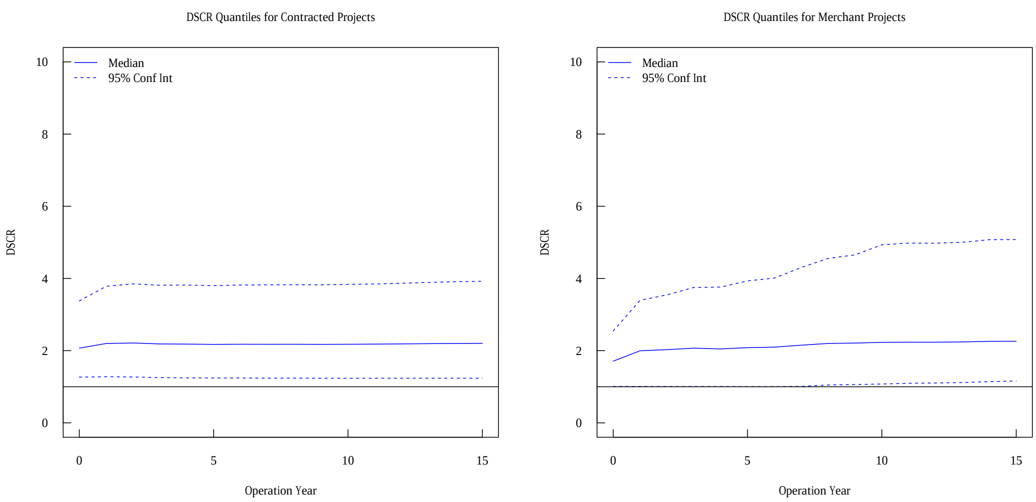

4. Stochastic DSCR Model

Figure 24: *Filtered DSCR mean (left) and volatility (right) for contracted and merchant families.*

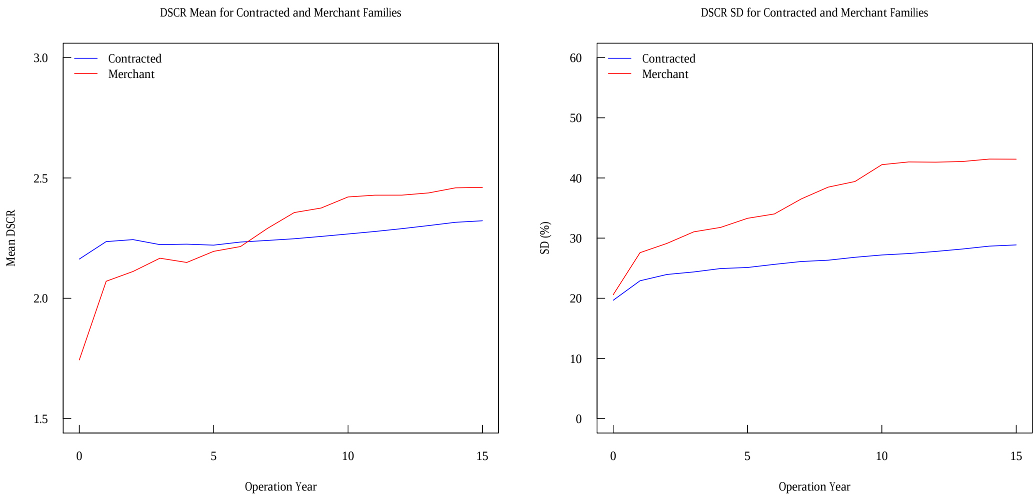

Figure 25: *DSCR densities for contracted and merchant families.*

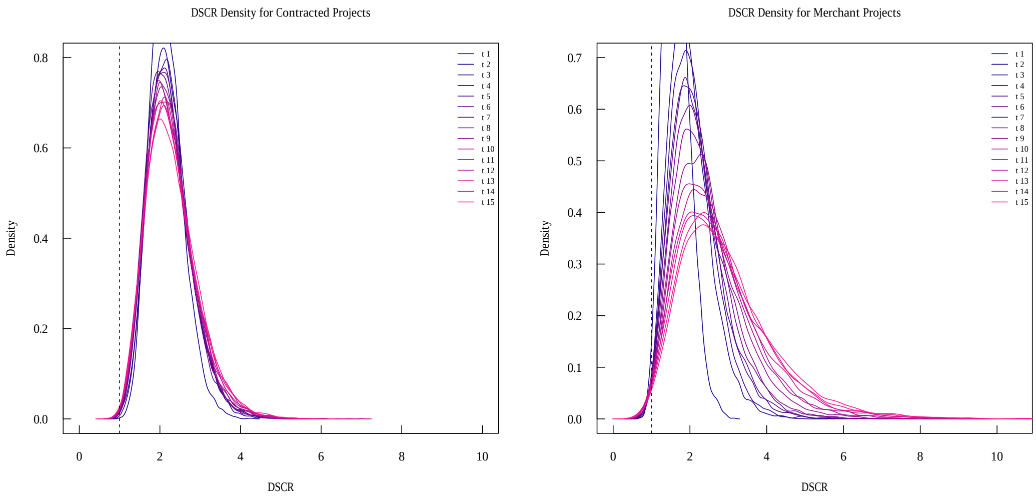

from about 25 to 45%. Thus, the merchant exhibits a much more pronounced increases in its mean and volatility of DSCR over time. Table 11 shows this clearly by documenting the growth rates of the DSCR mean and volatility for the two families.

Figure 25 shows the resulting estimates of DSCR densities for both project families at time *t*. Densities for merchant projects are more spread out, and have heavier right tails in almost all operation years, compared to DSCR densities for contracted projects, consistent with the panel regression results reported in Section 2.3.

Moreover, densities for the merchant projects drift further to the right compared to their contracted counterparts as investment time unfolds, consistent with the more pronounced change in mean and standard deviation shown in Table 10.

4. Stochastic DSCR Model

Figure 26: DSCR trajectories in the state (m, σ) plane, for both families.

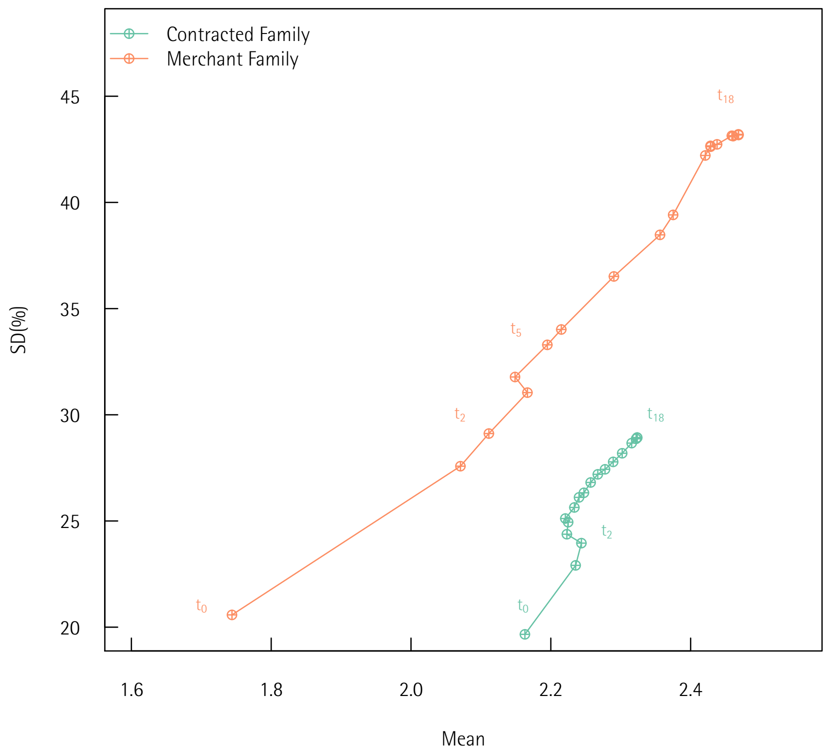

Figure 26 shows the DSCR trajectories of the two families in the DSCR mean-volatility plane.

4.2 Credit Risk Implications

An immediate result from the filtered estimates of DSCR densities in infrastructure projects is the probability that realised DSCR values fall below a default threshold.

Figure 27 shows the probability that $DSCR_t$ falls below the hard default threshold at each point in time. This is consistent with Moody's definition of default (missing one payment) and we note that this result is very similar to the trends in marginal default frequencies reported by Moody's (2015).

Figure 27 also shows the probability that $DSCR_t$ falls below a "technical" default threshold of 1.05. Having derived the entire distribution of DSCRs at each point time, we can now predict technical as well as hard defaults, an important improvement on the reduced form approach employed by rating agencies.

Indeed, technical defaults are the most common in project finance and they trigger valuable options to step-in for creditors, which largely explains the high reported level of recovery in infrastructure project debt.

Overall, these results suggest that merchant family shows a more dynamic credit risk profile compared to contracted family.

4. Stochastic DSCR Model

Figure 27: Probabilities of hard and soft defaults for contracted and merchant families, computed as the probabilities of $DSCR_t$ falling below 1.0 and 1.05, respectively.

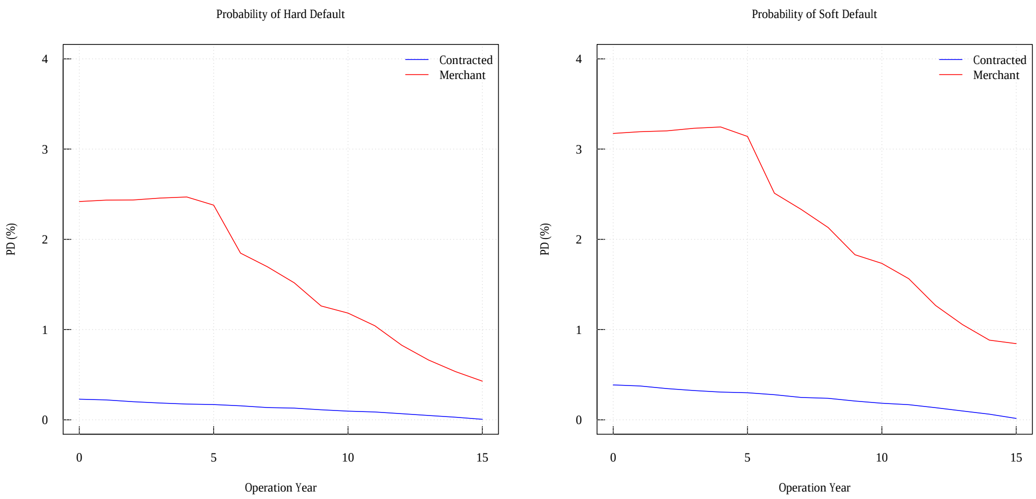

Table 11: Growth rate of the DSCR mean and SD for the two families.

	Contracted	Merchant
Mean	0.45%	2.44%
SD	1.09%	4.38%

As discussed above, once an infrastructure project is in default, this lognormal DSCR process stops for a least one period. Likewise, DSCRs may take such high values that they may not be characterised as following a lognormal distribution anymore. In the next chapter, we present our estimates of the DSCR state transition model outlined in Chapter 3.

5. DSCR State Transition Model

5. DSCR State Transition Model

In this chapter, we describe our results for time-varying state transition probabilities of the DSCR between the default, risky and safe states discussed in Chapter 3.

We first estimate the transition probabilities between either the risky or the safe state and default state in Section 5.1 and then the probability of transition between the risky and safe states conditional on the project not being in default in Section 5.2. Section 5.3 combines these estimates and summarises the probability of being in any given state at each point in time.

5.1 Transitions To and From the Default State

We estimate transition probabilities using Bayesian updating of prior transition probabilities once a number of transition counts have been observed.

Defaults are rare in project finance and our sample remains limited in size and therefore allows us to observe only a handful of defaults (5). This makes calibration using observed default counts less reliable as our choice of prior will dominate the results. Instead, the predicted count of defaults is obtained using our estimates of the probability of default for the two families reported in Chapter 4.

Say that $N = n_d + n_r + n_s$, where n_r, n_s, and n_d are number of DSCR observations in the risky, safe, and default states, respectively, and N is the total number of DSCR observations at time t. Then if p_d is the probability of default at that time, we can write:

$$\begin{aligned} N &= p_d N + n_r + n_s \\ N(1 - p_d) &= n_r + n_s \end{aligned} \tag{5.1}$$

and

$$\begin{aligned} n_d + n_r + n_s &= \frac{n_r + n_s}{(1 - p_d)} \\ (n_r + n_s)p_d &= n_d(1 - p_d) \\ n_d &= \frac{(n_r + n_s)p_d}{1 - p_d} \\ &= \frac{n_{rs} p_d}{1 - p_d} \end{aligned} \tag{5.2}$$

Hence, we can obtain a default count from our previous estimate of p_d and counts of $n_{rs} = n_r + n_s$ in each period.

To estimate the number of transitions from the default state, we start with the time to emergence from default of 2.3 years, reported in Moody's (2015), and provide the following calibration: we assume that each project has a 10% probability of emergence from default within the first year, a 50% probability of emergence in the second year conditional on not having emerged in the first year, and a 40% probability of emergence from default in the third year conditional on not having emerged in the first two years. This leads to an average time of emergence of 2.3 years. Figure 29 shows probability of emergence from default.

This allows us to estimate the number of projects that transition from the default to default states, n_{dd}, and from the default to one of the two no-default states (risky or safe), $n_{drs} = n_{dr} + n_{ds}$, in each period. Once the number of transitions have been

5. DSCR State Transition Model

Table 12: Bayesian estimates of transition probabilities from the default state for both families. Initial priors $\alpha^- = 15$ and $\beta^- = 1$, the posterior values for each period are used as priors in the next period.

	Contracted									Merchant								
	p_d	n_{rs}	n_d	n_{dd}	n_{drs}	α	β	p_{dd}	p_{drs}	p_d	n_{rs}	n_d	n_{dd}	n_{drs}	α	β	p_{dd}	p_{drs}
1	0.37	105	0.44	0.44	0.00	15	1	0.94	0.06	0.37	33	2.86	2.86	0.00	18	3	0.84	0.16
2	0.37	119	0.83	0.44	0.39	16	2	0.90	0.10	0.37	47	2.75	0.17	2.57	20	6	0.77	0.23
3	0.37	111	0.80	0.41	0.40	17	3	0.87	0.13	0.37	40	0.30	0.15	0.16	21	6	0.76	0.24
4	0.36	107	0.93	0.39	0.37	18	4	0.84	0.16	0.36	38	1.43	0.14	0.13	21	8	0.73	0.27
5	0.36	93	0.86	0.33	0.35	19	4	0.81	0.19	0.36	27	0.29	0.10	0.12	22	8	0.73	0.27
6	0.36	79	0.75	0.28	0.30	19	5	0.79	0.21	0.36	26	0.24	0.09	0.09	22	8	0.72	0.28
7	0.35	69	0.66	0.25	0.25	20	6	0.77	0.23	0.35	20	0.21	0.07	0.08	22	9	0.72	0.28
8	0.36	50	0.54	0.18	0.22	20	6	0.76	0.24	0.36	16	0.16	0.06	0.06	22	9	0.72	0.28
9	0.35	34	0.40	0.12	0.16	20	7	0.75	0.25	0.35	11	0.13	0.04	0.05	22	9	0.72	0.28
10	0.34	24	0.29	0.08	0.11	21	7	0.75	0.25	0.34	9	0.09	0.03	0.03	22	9	0.71	0.29
11	0.34	15	0.20	0.05	0.07	21	7	0.74	0.26	0.34	6	0.07	0.02	0.03	22	9	0.71	0.29
12	0.33	11	0.13	0.04	0.05	21	7	0.74	0.26	0.33	5	0.05	0.02	0.02	22	9	0.71	0.29
13	0.33	10	0.10	0.03	0.03	21	7	0.74	0.26	0.33	4	0.04	0.01	0.01	22	9	0.71	0.29
14	0.32	8	0.08	0.03	0.03	21	8	0.74	0.26	0.32	3	0.03	0.01	0.01	22	9	0.71	0.29
15	0.29	6	0.06	0.02	0.02	21	8	0.74	0.26	0.29	4	0.03	0.01	0.01	23	9	0.71	0.29

Figure 28: Transition probabilities from the default state for contracted and merchant projects in operation time.

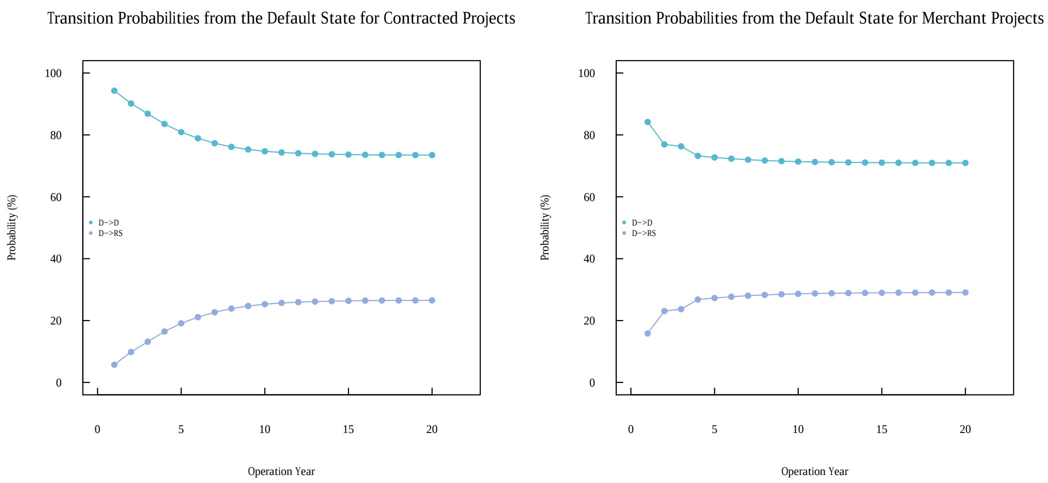

computed, the transition probabilities from the default state are estimated using the procedure described in Chapter 3. That is, the probability of transitioning from default to default state is

$$p_{dd} = \beta(\alpha_{dd}, \beta_{dd}) \tag{5.3}$$

$$p_{drs} = 1 - p_{dd} = 1 - \beta(\alpha_{dd}, \beta_{dd}) \tag{5.4}$$

We note that the p_{drs} is different from the probability of emergence of a single project, as this is also affected by the number of projects that are currently in the default state and due to emerge but may have defaulted at different points in time. p_{drs} is the probability that all the projects currently in default will emerge from default in the next period. Thus, it is the probability of emergence from default of an average project in a diversified portfolio of projects within a given project family.

Table 12 shows the estimated number of projects as well number of transitions from the default to default states and default to no-default states. Once the number of transitions from the default state have been obtained, we can compute the implied probabilities of transition from the default state. These transition probabilities are also reported in Table 12, and shown in Figure 28.

5. DSCR State Transition Model

Figure 29: Probability of emergence from default.

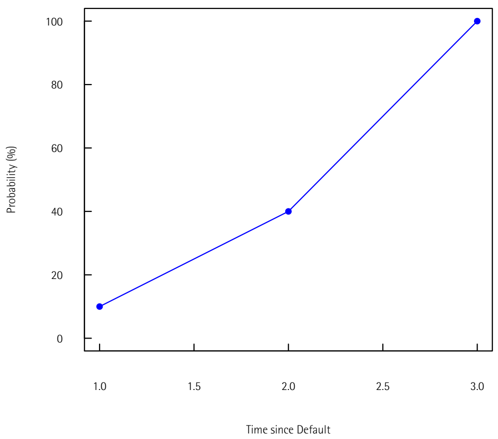

For both families, we find that the probability of transitioning out of the default state, conditional on having been in default at the previous period increases as a project matures. That is, in a diversified portfolio of loans within a single family, more loans are likely to emerge from default in later years of the project life.

5.2 Transitions Conditional on No Default

Next, we move on to computing the probabilities of transition between safe and risky states, given that the project is not in default.

Tables 13 and 14 report the Bayesian estimates of the posterior parameters for the probabilities of staying within the same state from period to period. We start with a uniform prior for the probability of staying within the same state. That is, p_{rr}, p_{ss}, and p_{dd} 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 the procedure described in chapter 3:

$$\alpha_{ii} = \alpha_{ii} + n_{ii} \tag{5.5}$$

$$\beta_{ii} = \beta_{ii} + N_i - n_{ii} \tag{5.6}$$

where $i \in \{r, s, d\}$, 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:

$$p_{ii} = \frac{\alpha_{ii}}{\alpha_{ii} + \beta_{ii}} \tag{5.7}$$

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 tend to become more stable.

Figure 30: Transition probabilities from the safe state for contracted and merchant projects in operation time.

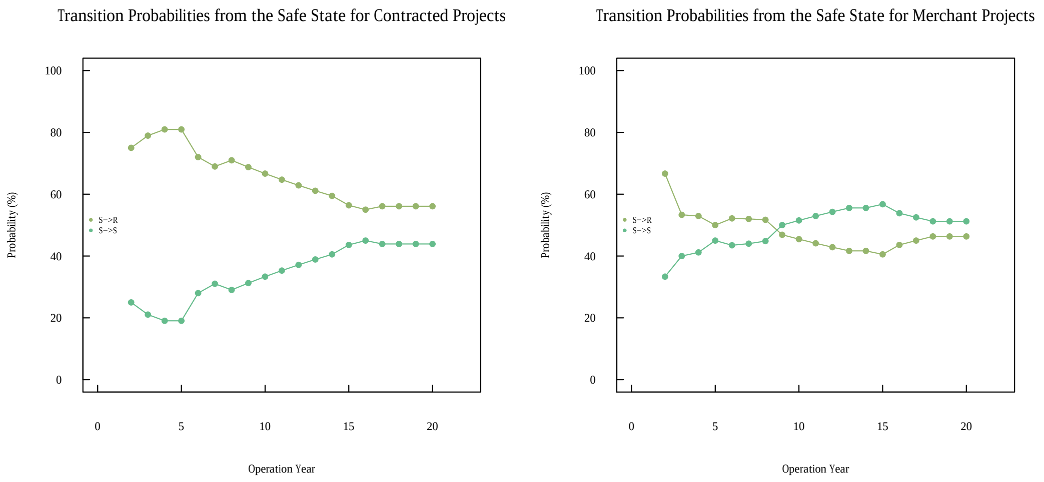

Figure 31: Transition probabilities from the risky state for contracted and merchant projects in operation time.

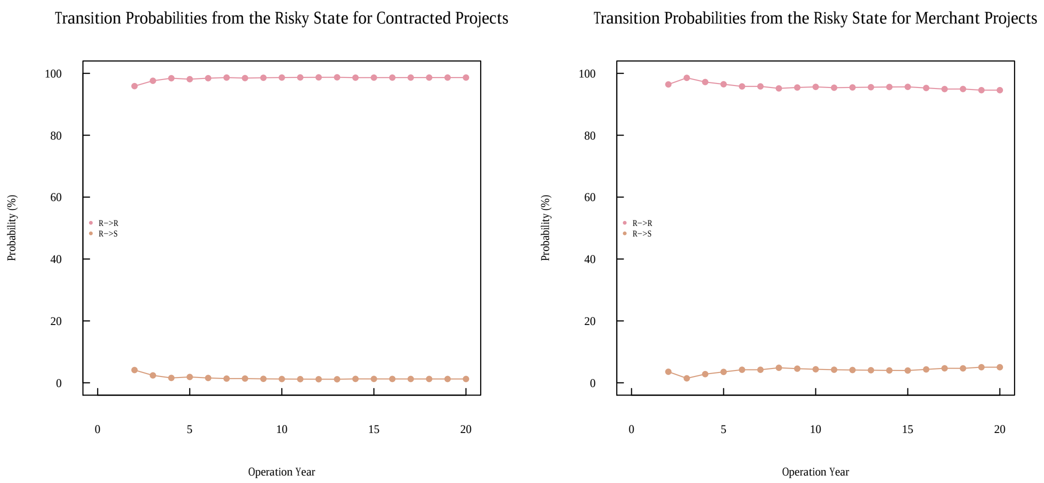

Table 13: Bayesian estimates of state transition probabilities for contracted projects. Initial prior is assumed to be $\alpha^{-} = \beta^{-} = 1$ (equivalent to a Uniform distribution), and the posterior values for each period are used as prior for the next period.

	risky-to-risky					safe-to-safe							
	N_r	n_{rr}	α^{+}_{rr}	β^{+}_{rr}	p_{rr}	N_s	n_{ss}	α^{+}_{ss}	β^{+}_{ss}	p_{ss}	p_r	p_s	p_d
1	95	92	93	4	0.96	10	2	3	9	0.25	0.87	0.10	0.00
2	112	111	204	5	0.98	7	1	4	15	0.21	0.91	0.07	0.00
3	109	109	313	5	0.98	2	0	4	17	0.19	0.93	0.06	0.00
4	107	104	417	8	0.98	0	0	4	17	0.19	0.95	0.04	0.00
5	89	89	506	8	0.98	4	3	7	18	0.28	0.95	0.04	0.00
6	75	75	581	8	0.99	4	2	9	20	0.31	0.95	0.04	0.00
7	67	65	646	10	0.98	2	0	9	22	0.29	0.95	0.04	0.00
8	49	49	695	10	0.99	1	1	10	22	0.31	0.95	0.04	0.00
9	33	33	728	10	0.99	1	1	11	22	0.33	0.95	0.04	0.00
10	23	23	751	10	0.99	1	1	12	22	0.35	0.95	0.04	0.00
11	14	14	765	10	0.99	1	1	13	22	0.37	0.95	0.04	0.00
12	10	10	775	10	0.99	1	1	14	22	0.39	0.95	0.04	0.00
13	9	8	783	11	0.99	1	1	15	22	0.41	0.95	0.04	0.00
14	6	6	789	11	0.99	2	2	17	22	0.44	0.95	0.04	0.00
15	5	5	794	11	0.99	1	1	18	22	0.45	0.95	0.04	0.00

Table 14: *Bayesian estimates of state transition probabilities for merchant projects. Initial prior is assumed to be $\alpha^- = \beta^- = 1$ (equivalent to a Uniform distribution), and the posterior values for each period are used as prior for the next period.*

	risky-to-risky					safe-to-safe							
	N_r	n_{rr}	α^+_{rr}	β^+_{rr}	p_{rr}	N_s	n_{ss}	α^+_{ss}	β^+_{ss}	p_{ss}	p_r	p_s	p_d
1	26	26	27	1	0.96	7	2	3	6	0.33	0.63	0.08	0.02
2	41	41	68	1	0.99	6	3	6	9	0.40	0.72	0.06	0.02
3	38	36	104	3	0.97	2	1	7	10	0.41	0.76	0.04	0.01
4	35	33	137	5	0.96	3	2	9	11	0.45	0.78	0.03	0.02
5	24	22	159	7	0.96	3	1	10	13	0.43	0.79	0.03	0.01
6	24	23	182	8	0.96	2	1	11	14	0.44	0.80	0.04	0.01
7	16	14	196	10	0.95	4	2	13	16	0.45	0.80	0.04	0.00
8	13	13	209	10	0.95	3	3	16	16	0.50	0.80	0.04	0.00
9	10	10	219	10	0.96	1	1	17	16	0.52	0.80	0.04	0.00
10	8	7	226	11	0.95	1	1	18	16	0.53	0.80	0.04	0.00
11	5	5	231	11	0.95	1	1	19	16	0.54	0.80	0.04	0.00
12	4	4	235	11	0.96	1	1	20	16	0.56	0.80	0.04	0.00
13	4	4	239	11	0.96	0	0	20	16	0.56	0.81	0.04	0.00
14	2	2	241	11	0.96	1	1	21	16	0.57	0.80	0.05	0.00
15	2	1	242	12	0.95	2	0	21	18	0.54	0.80	0.05	0.00

Table 15: *Average state transition probabilities between default and no-default states for contracted and merchant projects.*

	Contracted		Merchant	
	D	*RS*	*D*	*RS*
D	0.78	0.22	0.64	0.36
RS	0.00	1.00	0.00	1.00

Table 16: *Average state transition probabilities (conditional on no default) between risky and safe states for contracted and merchant projects.*

	Contracted		Merchant	
	R	*S*	*R*	*S*
R	0.99	0.01	0.94	0.06
S	0.58	0.42	0.48	0.49

Figure 32: *Probability in being one of the three DSCR states for contracted and merchant projects.*

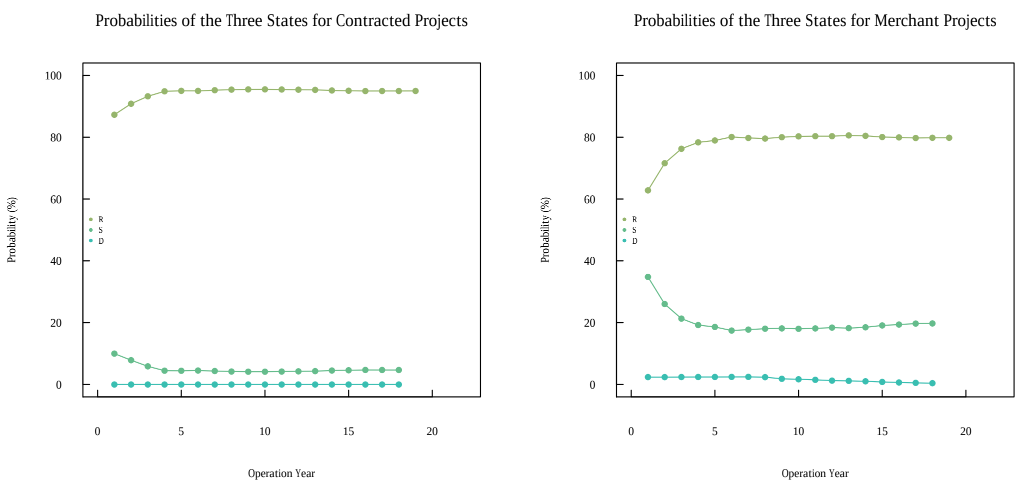

5. DSCR State Transition Model

Figure 30 shows transition probabilities in each period between safe and risky states conditional on not being in the default state. In this figure, we see that the transition probabilities from the risky to the safe state start at a relatively high level (40%) and tend to increase in time for merchant projects but never reach levels higher than 50%. For contracted projects, the probability of transitioning from the risky to the safe state increases faster but from a lower level of around 20%, also to reach the 50% region by year 15. Hence, contracted projects take longer than merchant projects to create a potential for significant upside, signalled by a high DSCR (i.e. high free cash flow).

Figure 31 shows the same per period transition probabilities from the risky state. Generally speaking, the risky DSCR state is highly persistent i.e. the probability of staying in this state once the process is in it, is high for both project families, indicating that the projects are also likely to revert to this state. The probability of staying in the risky state is never 100% because the likelihood of a positive or negative shock pushing the DSCR process into the default or safe states is never zero.

Still, this result is significant for the purpose of asset pricing and conditional risk measurement because we have documented the risky state to be well-behaved, following a Gaussian process. This is significant to help improve the robustness of return and correlation estimates derived from this type of cash flow model.

5.3 State Transition Probabilities

Finally, figure 32 shows the probabilities of being in one of three states for contracted and merchant projects.

For contracted projects the probability of being in the risky state is much higher compared to the probability of being in the other two states. Contracted projects are more likely to stay in the "normal" risky state.

For merchant projects, the probability of being in the risky state is lower, while the probabilities of being in the default and safe states are higher compared to the corresponding probabilities for contracted projects. Thus, merchant projects are found to have more diverse DSCR trajectories in state space, and each state is less persistent (stable).

Table 16, which shows the matrix of transition probabilities between the 3 states, confirms this. It shows that the probability of staying within the state is much higher than switching to the other state, indicating a persistent state-dependence in both states. Moreover, persistence is higher for the risky state. The safe state is less stable and the DSCR process tends to (mean-)revert to the risky state.

As we argued earlier, path dependency can be an important dimension of infrastructure investment insofar as assets are more or less heterogenous and it can be difficult to fully diversify a portfolio of very large and bulky assets. Our results above suggest that contracted infrastructure is

5. DSCR State Transition Model

more homogenous than merchant projects, which are more likely to follow paths that diverge strongly from the mean of the population.

In the next chapter, we examine the role of path dependency in individual projects and how our approach allows us to track the dynamics of individual investments in the context of what what we have learned about other projects through previous observations.

6. DSCR Path-Dependency

6. DSCR Path-Dependency

In preceding chapters, we have focused on a population of investment infrastructure projects. While these projects can be divided into distinct families of DSCR processes, our results remained focused on broad trends within each family of projects.

However, we know that building a large, well-diversified portfolio of infrastructure projects is not necessarily available even to large, long-term investors (see Blanc-Brude, 2013, for a discussion) and that project specific risk is likely to remain a concern. Moreover, certain creditors such as commercial banks are required by prudential rules (a.k.a. Basel III) to report expected losses "line by line" i.e. without taking into account any diversification of the loan portfolio risk arising from their lending to multiple projects of different types and in different countries.

Hence, in this chapter, we present illustrative examples of how our DSCR state transition and filtering models can be used to track the DSCR behaviour of individual projects.

In what follows, we consider two examples: a negative shock to the DSCR path of a merchant project and a project which transits between risky, safe and default states during its life.

We show that applying Bayesian inference to these cases allows expectations to be revised in terms of future cash flows and cash flow volatility, thereby enabling us to revisit expected default frequencies and therefore expected loss and return measures, taking into account the path dependency of individual projects.

6.1 Tracking a Project with a Jump in the DSCR

We first consider a project that follows an oft-observed trajectory: while it remains in the risky state throughout its life, it starts off with a relatively high DSCR, implying a merchant-type structure with relatively high DSCR volatility, but later on undergoes a large downward shift in its realised DSCR level, e.g. as the result of a negative demand shock, while its DSCR realised volatility from that point onwards also decreases markedly.

A concrete case of such a trajectory could be a toll road experiencing significant loss of traffic after an economic recession, but for which the residual "baseload" traffic is much less volatile than before the shock, and still sufficiently high to keep the DSCR out of the default state.

Such a project would not be adequately captured by the mean DSCR process of its original family, even though this was the best available starting point to anticipate its behaviour at t_0.

In this illustration, we know the "true" underlying DSCR process that is otherwise unobservable, as discussed in chapter 3 and how it is impacted by the negative demand shock. The point of the example is to show that as we observe realised information, our estimates of the true process can quickly converge to the true value and

6. DSCR Path-Dependency

then track it as it evolves during the life of the investment.

Here, the project is assumed to have a life of 20 years, starts with a DSCR of 2.0 and follows the following *true* dynamics for the first 10 years

$$log(DSCR)_{t+1} = log(DSCR)_t + N(0, 0.08) \tag{6.1}$$

In year 11, the project suffers a shock, causing its *true* mean DSCR to go down to 1.3 and it then replicates the following dynamics for the next 10 years:

$$log(DSCR)_{t+1} = log(DSCR)_t + N(0, 0.04) \tag{6.2}$$

Hence, at t_{11} the expected DSCR has dropped by 35% below its original trajectory, but DSCR volatility has halved.

Next, we draw a single realised DSCR value from this underlying process for each of the 20 periods, thus simulating 20 DSCR annual observations, and use the filtering approach described in chapter 4 to estimate the underlying parameters.

Figure 33 shows the filtered DSCR mean and standard deviation along with the realised DSCR values and the true standard deviation of the project.

The figure shows that filtered estimates of mean and standard deviation follow the underlying trajectory i.e. the estimated parameters of m and p are reasonably accurate and track its true values very closely.

As soon as the DSCR diverges from its original trajectory the filter takes this new information into account, and if the divergence persists, future estimates of the expected value of $DSCR_t$ are updated accordingly.

Likewise, initial estimates of the volatility of $DSCR_t$ on the right panel of figure 33 are corrected as information about the lower realised volatility becomes integrated into each posterior value.

The ability to revise the DSCR dynamics of individual projects directly leads to the revision of their risk metrics.

For example, figure 34 shows the implied probabilities of the project's DSCR falling below the level of 1.15, 1.10, and 1.0 as the DSCR dynamics are revised in time i.e. the probabilities of dividend lockup, soft default, and hard default, respectively.

The figure suggests that the negative jump in the DSCR, combined with the lower realised volatility of DSCR, has no noticeable effect on the project's probability of hard default, a negligible impact on probability of soft default, but a noticeable impact on the probability of a dividend lockup.

We note that without this combined assessment of the DSCR mean and volatility, a drop of 35% of the project free cash flow could qualify it for a (possibly unwarranted) credit downgrade.

6. DSCR Path-Dependency

Figure 33: Filtered DSCR quantiles and standard deviation for a single project.

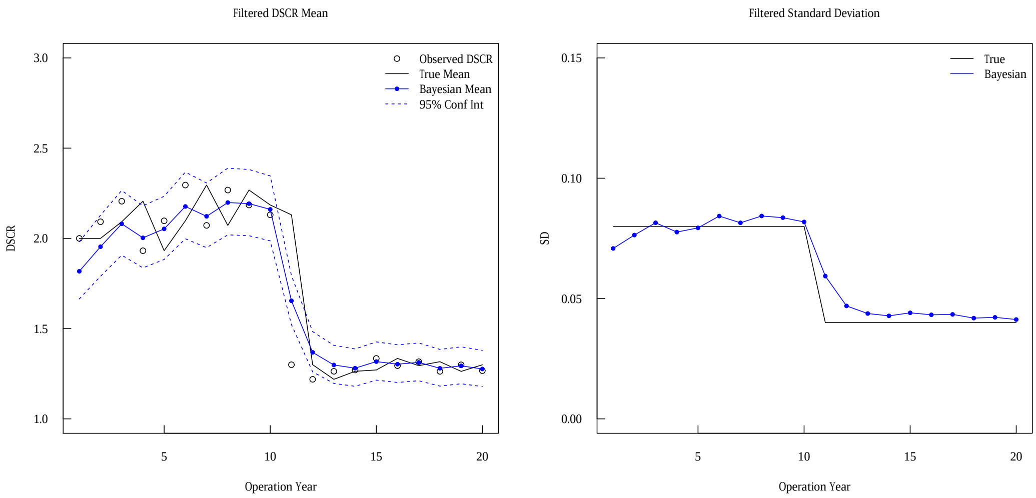

Figure 34: Project's distance to lockup, soft, and hard defaults and the corresponding probabilities of lockup, soft, and hard defaults, estimated using filtered DSCR mean and volatility.

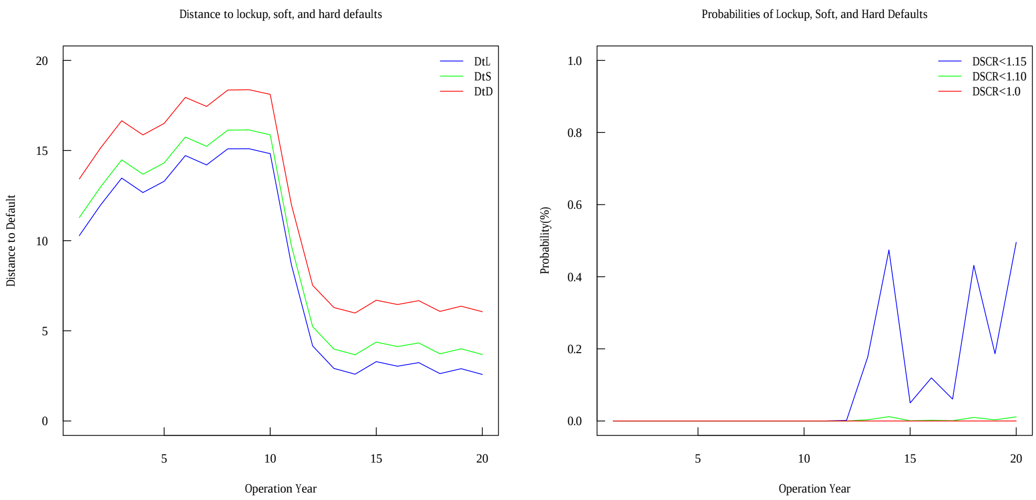

6.2 Tracking a Project Changing State

In this example, we assume the existence of a project that first starts in the risky state, stays in it for three years, then transitions into the safe state after year three, but goes into default in the subsequent year. It eventually emerges from default after two periods and stays in the risky state until maturity.

As in the previous example, we assume that the project follows autoregressive dynamics, so that the mean DSCR in the next period is simply the realised DSCR in the previous period. The project is assumed to have a constant variance of 5%.

In this example, we need to predict DSCR dynamics post-emergence from the default state. Hence, the filtered DSCR distribution for the whole family can be combined with

Figure 35: Weights for forecasts based on project and family priors, and the corresponding DSCR forecasts.

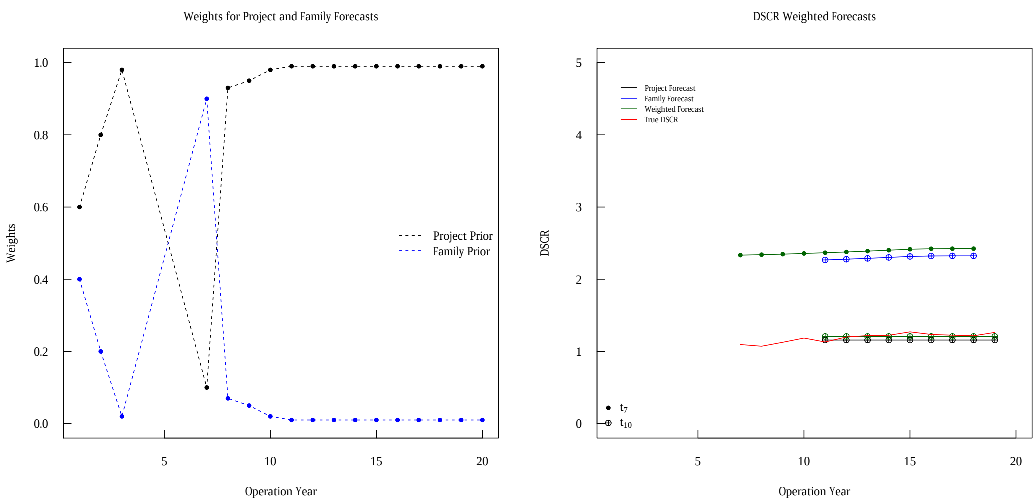

Figure 36: Filtered DSCR quantiles and standard deviation for a project with state transitions.

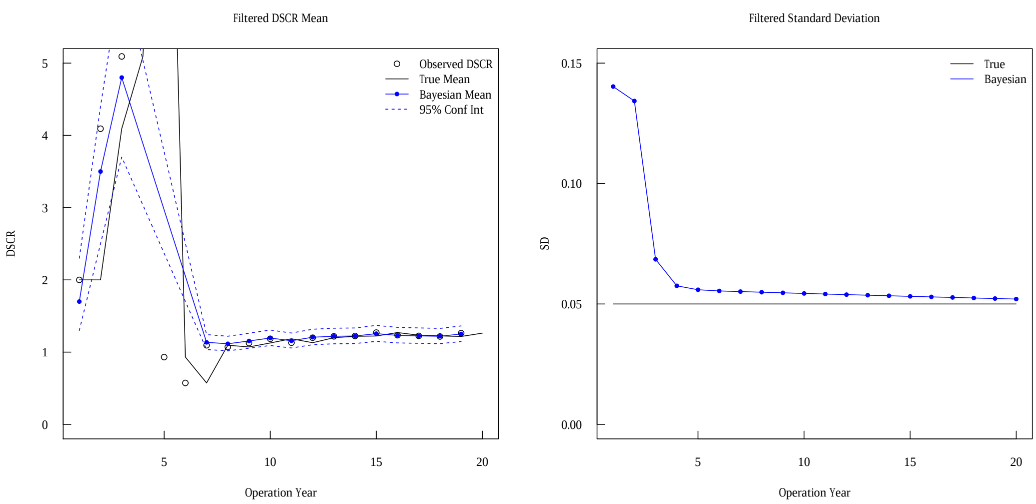

the project specific DSCR distribution to form a more reliable DSCR forecast.

Once the project goes into default, we use a weighted average of the family-specific DSCR distribution and the project-specific DSCR to forecast the new DSCR distribution when it comes out of the default state.

To combine the family-specific DSCR distribution with project-specific DSCR forecasts, we compare realised DSCR in every period after the project's emergence from default with the DSCR forecasts implied by both the family- and project-specific priors, and compute the sum squared deviation for both forecasts.

Then we compute the updated (posterior) DSCR distribution in every period using project- and family-specific DSCR prior, weigh the two implied distributions by their

6. DSCR Path-Dependency

sum squared deviation. Thus the posterior distribution of the project can be written as

$$\mu^{+} = \frac{SS_p}{SS_p + SS_f}\mu_f^{+} + \frac{SS_p}{SS_p + SS_f}\mu_p^{+} \tag{6.3}$$

$$p^{+} = \frac{SS_p}{SS_p + SS_f}p_f^{+} + \frac{SS_p}{SS_p + SS_f}p_p^{+} \tag{6.4}$$

μ_f^{+} and p_f^{+} are the posterior mean and precision of DSCR distribution obtained by updating the family-specific prior using the realised project DSCR in the current period, and μ_p^{+} and p_p^{+} are the posterior mean and precision of DSCR distribution obtained by updating the project-specific prior using the realised project DSCR in the current period.

This weighting scheme allows us to better capture the project's DSCR distribution immediately after its emerges from default i.e. when we have little information about its new dynamics. If we used only the project-specific prior, which was last updated before the project transitioned into the safe state, the forecasted DSCRs will be much higher. As new information about the project's new behaviour transpires in periods following emergence, the weights on project-specific forecasts increase, and the DSCRs are effectively computed using only the project-specific forecasts.

This is shown on figure 35. The project forecasts carry more weight in the first three periods, but drop in year 7 when the project emerges from default and new "learning" has to occur. In this case, the new DSCR is much lower than it was before project went into default. The weights for project-specific forecasts then increase again as realised DSCRs stabilise at a level considerably lower than the level forecasted by the family-specific prior.

Black, blue, and green curves indicate project-specific, family-specific, and weighted forecasts, respectively, for the 7th and 10th years. We see that the weighted forecasts are always closer to the true DSCR level (the red curve). Figure 36 shows the filtered DSCR mean and volatility using weighted forecasts, and it tracks the true values fairly well, apart from the first few years where filtered DSCR volatility exceeds true DSCR due to prior uncertainty.

Finally, Figure 37 shows the DSCR trajectories for the two families and the two projects described in this chapter, in the DSCR mean-variance plane. It can be clearly seen that the individual projects can exhibit fairly different trajectories compared to the DSCR families. To the extent that investors cannot fully diversify or have to report line-by-line credit risk measures, the ability to model the path dependency of the $DSCR_t$ process allows project specific risks to be taken into account, while also integrating available information about each project's reference process or family.

6. DSCR Path-Dependency

Figure 37: Individual DSCR state (m_t,p_t) trajectories in time

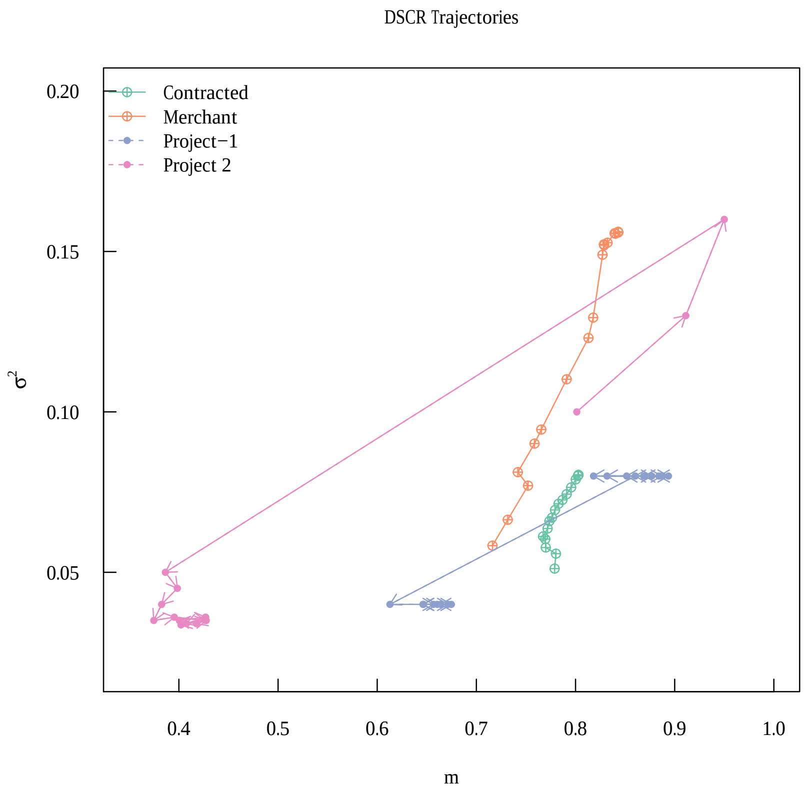

7. Conclusions

7. Conclusions

In this paper, we have conducted the first empirical study of DSCR dynamics in infrastructure project finance. We hand collected 15 years of realised DSCR data for more than 200 projects in Europe and the United States, covering two broad categories of projects: those receiving a contracted income and those exposed to merchant or commercial risks.

We find that these two business models can be considered to correspond to two distinctive DSCR processes, with statistically different mean and variance parameters and following different project time dynamics. We also find that contracted infrastructure DSCRs in the cross section are much less affected by macro-variables or the business cycle than merchant projects.

Initial findings using linear regression and panel models suggest that the DSCR profile of individual projects and families of projects is highly non-linear, auto-regressive and heteroskedastic (variance is not constant). However, we also show that realised DSCRs can be fitted to a lognormal process up to their 90^{th} and 85^{th} quartiles for contracted and merchant projects, respectively, which allows the development of an easily tractable model of parameter inference.

Hence, we propose a two-dimensional modelling strategy combining, first a three-state transition model between which the DSCR process is assumed to transit: a risky state in which it is indeed an autoregressive lognormal process; a default state defined by a threshold corresponding to DSCR = 1 in which the DSCR process stops until it emerges from default; and a safe state, corresponding to high realised values above the "good-lognormal-fit" quantile, in which case, as long as the DSCR stays in that state, the project debt is considered risk-free.

Second, we propose to use a straightforward implementation of so-called particle filtering models to infer the parameter values of the DSCR's lognormal process in the risky state i.e. when documenting and tracking the volatility of the DSCR matters, because it is a direct measure of credit risk.

We show that such a framework allows us to derive the dynamics of DSCRs in well defined groups of projects as well as individual projects, including tracking the individual DSCR "path" followed by investments that do not necessarily correspond to the median infrastructure project.

For instance, we show that the ability to infer both the expected value and the volatility of the DSCR process allows us to take a much more informed view on the credit risk of projects that substantially deviate from their base case e.g. a negative DSCR shock accompanied by a significant reduction of DSCR volatility does not necessarily lead to an increase of expected default frequencies.

Such analyses will be further developed as new data is collected and standardised to improve our ability to track the DSCR path of individual and groups of infrastructure projects.

8. Appendix: Calculations

8. Appendix: Calculations

For the most part, annual accounts are represented faithfully from the statements provided. However, statements may not always include a cash flow statement, which is required to understand the underlying cash flows of infrastructure assets. As a result, cash flow statements are estimated from the income statement and balance sheet.

The starting point for all calculations is the operating profit after tax. This value is the accrual profit and does not reflect the actual cash flows the firm experienced over the year. As a result, adjustments are made to take into account the non-cash impacts of the accrual accounts.

8.1 Operating Cash Flows

To obtain operating cash flows, the following adjustments are made to:

1. Operating profit after tax
2. add Depreciation and amortisation - these are non-cash charges so are added back
3. add Changes in Finance Receivable - This is the main investment for project financed companies with off-take arrangements. This is because they have to be accounted for as a finance lease. Any change in this amount represents cash received on top of the interest recognised in the income statement.
4. add Interest Expense - this is added back as it is not an operating cash flow, instead this is a financing item
5. add (Increase)/Decrease in receivables and prepayments - any increase (decrease) in receivables and payments means that revenue is recognised (not recognised), but the cash has not be received (cash has been received). As a result the operating profit needs to be adjusted.
6. add Increase/ (Decrease) in payables - any increase in payables means that expenses have been recognised but not paid (a decrease means a liability has been paid, but no expense recognised)
7. add Movement in Taxes - for the same reason as accounts receivables and payables, if this goes up, the tax expense has been recognised but not paid, if it is negative, then the tax has been paid, but not recognised.
8. add Increase/ (Decrease) in Non Current Provisions - any increase in provisions means that expenses have been recognised but not paid (a decrease means a liability has been paid, but no expense recognised).
9. add Increase/ (Decrease) in Unearned Income - any increase in unearned income means that cash has been received but not recognised as revenue, instead it is an obligation to provide services in the future (a decrease means the services have been provided and revenue recognised but no cash received).

8.2 Investing Cash Flows

The investing cash flows are made up of two line items. These are:

8. Appendix: Calculations

1. Change in Investments - takes into account any changes in the short-term investments
2. Investment in Project - measures any change in PP&E and increases in the finance debtor. But subtracts any increase in the Asset Revaluation Reserve.

8.3 Financing Cash Flows

1. Repayment of Senior Debt - any decrease in the non-current and current senior debt is assumed to be a repayment.
2. Repayment of Mezzanine Debt - any decrease in the non-current and current mezzanine debt is assumed to be a repayment.
3. Repayment of Equity Bridge - any decrease in the non-current and current equity bridge is assumed to be a repayment.
4. Repayment of Shareholder Loans - any decrease in the non-current and current shareholder loans is assumed to be a repayment.
5. Repayment of Bonds - any decrease in the non-current and current bonds is assumed to be a repayment.

The working assumption for all interest is that any that is recognised as an expense during the period in the income statement, is paid in the same period. This is not an issue for capitalised interest during construction, but for projects where interest is capitalised as a result of an inability to pay, it may create some issues.

1. Interest Expense-Senior Bank Loans - any interest recognised as an expense during the period is treated as a cash outflow during the period.
2. Interest Expense-Mezzanine Bank Loans - any interest recognised as an expense during the period is treated as a cash outflow during the period.
3. Interest Expense-Bonds - any interest recognised as an expense during the period is treated as a cash outflow during the period.
4. Interest Expense-Other Interest Bearing - any interest recognised as an expense during the period is treated as a cash outflow during the period.
5. Interest Expense-Shareholder Loans - any interest recognised as an expense during the period is treated as a cash outflow during the period.

Debt drawdowns are all calculated the same way. If the difference between the sum of current and non-current debt of this year is greater than the sum of the current and non-current debt of last year, it is assumed that the debt has increased. This assumption is poor when dealing with index-linked debt securities like RPI linked bonds. However without the necessary detail supplied in the accounts, it is the best alternative at the moment.

1. Drawdown of Senior Debt
2. Drawdown of Mezzanine Debt
3. Drawdown of Shareholder Loans
4. Drawdown of Bonds
5. Drawdown of Equity Bridge
6. Initial Equity Investment - this is calculated as the difference between the prior

8. Appendix: Calculations

year's paid up capital and the current year's paid up capital. Any increase is assumed to be an initial equity investment.

7. Dividends Paid - is calculated by summing of the prior year's retained earnings and the current year's profit and subtracting the current year's retained earnings. If there is a difference, it is assumed to be the dividend paid for the year. Dividends can be declared and not paid so in addition to this, any change in the dividend declared account from the prior year to this year is added.

8.4 Prior DSCR beliefs and meta-parameters

For a given prior belief about the arithmetic mean DSCR (E($DSCR_t$)) and arithmetic variance of DSCR at time t ($\sigma^2_{DSCR_t}$), we can compute the prior parameters m (location) and *sigma* (scale) of the log data thus:

$$\sigma^2 = ln(1 + \frac{\sigma^2_{DSCR_t}}{E(DSCR_t)^2})$$

$$m = ln(E(DSCR_t)) - \frac{1}{2} ln(1 + \frac{\sigma^2_{DSCR_t}}{E(DSCR_t)^2})$$

$$= ln(E(DSCR_t)) - \frac{1}{2}\sigma^2$$

which follows from the definition of the lognormal density with mean $exp(m + \frac{\sigma^2}{2})$ and variance $(exp(\sigma^2) - 1)exp(2\mu + \sigma^2)$.

The value of parameter m follows a Gaussian (normal) distribution with meta-parameters μ (mean) and δ (precision). The prior value of μ is simply the prior value of m and the prior value of precision δ is set to a small number (implying a large variance), e.g. 0.01.

The value of precision parameter p follows a Gamma distribution for which we need to derive the shape (a) and rate (b) meta-parameters. We first give p a prior expected value and variance. The prior mean of p (call it μ_p) is simply the inverse of σ^2, the initial prior for the scale of the log data. The prior variance of p (call if var_p) is set to a large number relative to the prior expected precision, e.g. ten times p.

The initial prior values of a and b are then computed as

$$a = \frac{\mu_p^2}{var_p}$$

$$b = \frac{\mu_p}{var_p}$$

which follows form the definition of the Gamma density function with mean $\mu_p = ab$ and variance $var_p = ab^2$.

9. Appendix: Data Tables

9. Appendix: Data Tables

9. Appendix: Data Tables

Table 17: Realised DSCR quantiles post-construction for merchant projects, 1999-2014

Year	N	10th-Q	20th-Q	50th-Q	80th-Q	90th-Q	Mean	SD
1	45	0.000	1.014	2.130	8.266	18.261	3.036	2.108
2	54	1.214	1.399	1.836	4.392	13.514	2.236	0.897
3	54	1.073	1.274	1.828	3.118	6.232	1.937	0.491
4	48	1.109	1.229	1.749	3.803	7.864	1.978	0.660
5	42	1.027	1.243	1.685	3.062	6.156	1.828	0.517
6	35	1.049	1.243	1.994	4.496	11.575	2.330	0.985
7	30	1.142	1.399	1.732	5.444	10.576	2.151	0.921
8	22	1.135	1.531	2.483	8.795	259.456	3.368	2.215
9	18	1.259	1.387	2.139	17.609	44.053	3.035	2.075
10	14	1.099	1.223	1.855	4.094	54.122	2.065	0.850
11	11	1.099	1.245	1.481	4.855	1,212.018	1.960	1.112
12	9	1.101	1.302	1.830	13.453	109.168	2.312	1.435
13	5	2.022	2.054	2.347	181.986	540.371	2.404	0.365
14	9	1.051	1.634	2.223	6.209	14.705	2.412	0.605
15	4	2.693	2.936	16.550	30.205	30.477	16.550	18.796

Table 18: Realised DSCR statistics post-construction by sector group, contracted infrastructure, 1999-2014

Year	N	10th-Q	20th-Q	50th-Q	80th-Q	90th-Q	Mean	SD
Commercial and Industrial	22	2.085	2.284	2.638	3.617	4.301	2.701	0.279
Environmental Services	6	0.969	1.255	1.461	1.503	1.637	1.461	0.010
Oil and Gas	7	2.039	2.511	3.527	3.629	3.772	3.456	0.217
Telecom	9	2.176	2.346	2.655	4.707	6.096	2.986	0.861

Table 19: Realised DSCR statistics post-construction by sector group, merchant infrastructure, 1999-2014

Year	N	10th-Q	20th-Q	50th-Q	80th-Q	90th-Q	Mean	SD
Energy	258	0.888	1.203	1.644	3.705	8.486	1.860	0.579
Environmental Services	8	1.004	1.196	2.597	3.837	5.140	2.494	0.854
Government Services	2	1.293	1.318	1.393	1.467	1.492	–	–
Oil and Gas	57	1.636	1.819	3.738	35.604	419.386	6.548	7.168
Telecom	9	2.352	2.413	2.968	5.646	6.923	3.522	1.182

Table 20: Realised DSCR statistics post-construction for contracted infrastructure, by initial leverage, project size

Year	N	10th-Q	20th-Q	50th-Q	80th-Q	90th-Q	Mean	SD
Leverage Below Median	509	1.472	1.714	2.284	3.828	5.876	2.428	0.521
Leverage Above Median	438	1.456	1.755	2.278	3.008	5.200	2.271	0.308
Size Below Median	491	1.435	1.681	2.231	3.688	5.587	2.331	0.475

Table 21: Realised DSCR statistics post-construction for merchant infrastructure, by initial leverage, project size

Year	N	10th-Q	20th-Q	50th-Q	80th-Q	90th-Q	Mean	SD
Leverage Below Median	227	1.148	1.425	2.814	18.942	232.373	4.105	3.136
Leverage Above Median	37	1.225	1.527	2.382	4.060	5.848	2.350	0.624
Size Below Median	127	1.243	1.548	2.727	7.661	23.784	3.175	1.280

9. Appendix: Data Tables

Table 22: Number of reporting firms in project time, by revenue risk family and by sector

	Total	Contracted	Merchant	C&I	Energy	Env. Ser.	Gov. Ser.	O&G	Telecom	Transp.
1	280	156	112	6	77	13	119	13	5	47
2	254	145	97	5	64	12	114	13	5	41
3	253	145	96	5	64	12	114	12	5	41
4	249	145	92	5	61	12	114	12	5	40
5	243	144	87	5	56	12	113	12	5	40
6	237	142	83	5	53	12	112	12	5	38
7	229	138	79	5	51	12	108	11	4	38
8	216	130	75	5	49	11	99	11	4	37
9	197	115	72	5	47	10	85	10	4	36
10	172	95	67	3	45	10	68	9	3	34
11	148	74	64	3	45	10	51	7	3	29
12	127	58	59	3	40	10	36	7	3	28
13	109	45	54	2	35	9	26	7	3	27
14	99	38	51	2	33	9	20	7	2	26
15	92	34	48	2	31	9	17	7	1	25
16	83	31	43	2	27	9	15	5	0	25
17	69	27	34	2	22	8	11	5	0	21
18	57	21	28	1	16	8	8	5	0	19
19	48	18	22	1	11	8	7	5	0	16
20	42	16	18	1	9	8	6	5	0	13
21	33	13	13	1	4	8	5	4	0	11
22	30	12	12	1	4	7	5	3	0	10
23	26	10	10	1	3	7	5	2	0	8
24	19	4	10	1	3	6	1	1	0	7

Table 23: Number of reporting firms and DSCR statistics post-construction in calendar time.

	Contracted			Merchant		
Year	*N*	*Median*	*SD*	*N*	*Median*	*SD*
1997	1.000	2.332	–	3.000	7.860	–
1998	2.000	1.665	–	4.000	16.490	20.481
1999	4.000	2.384	0.334	9.000	2.475	0.873
2000	8.000	3.565	1.002	10.000	3.659	95.928
2001	9.000	2.878	5.284	13.000	2.835	6.168
2002	14.000	3.102	0.845	12.000	1.759	0.561
2003	21.000	2.583	0.714	15.000	2.549	2.317
2004	32.000	2.430	0.388	21.000	1.879	1.771
2005	43.000	2.502	0.595	19.000	2.123	0.822
2006	62.000	2.275	0.541	24.000	2.258	1.070
2007	80.000	2.151	0.399	24.000	1.665	0.898
2008	86.000	2.334	0.324	24.000	2.229	0.893
2009	105.000	2.272	0.274	33.000	2.685	1.856
2010	114.000	2.246	0.358	36.000	2.092	0.546
2011	121.000	2.127	0.324	41.000	1.880	0.474
2012	127.000	2.097	0.328	44.000	1.963	1.043
2013	124.000	2.161	0.409	47.000	1.704	0.694
2014	47.000	2.437	0.535	37.000	1.565	0.517

9. Appendix: Data Tables

Table 24: OLS regressions for contracted and merchant families, and key sector within each family.

	Dependent variable: log(DSCR)						
		Contracted				Merchant	
	All	Gov. Services	Transport	All	Energy	Oil & Gas	Transport
Op Ph 1	0.451	0.824	0.048	0.196	2.454**	4.179***	−0.000
	(0.446)	(0.585)	(0.799)	(1.681)	(1.010)	(1.358)	(1.135)
Op Ph 2	0.522	0.870	0.378	0.679	2.887***	5.668***	0.365
	(0.450)	(0.588)	(0.833)	(1.698)	(1.053)	(1.570)	(1.190)
Op Ph 3	0.918**	1.258**	0.788	0.749	3.279***	5.631***	0.218
	(0.460)	(0.598)	(0.834)	(1.726)	(1.154)	(1.658)	(1.222)
Op Ph 4	0.489	0.833	0.274	0.044		6.355***	0.375
	(0.495)	(0.720)	(0.871)	(1.841)		(1.878)	(1.437)
Com & Ind	0.066	–	–	–	–	–	–
	(0.129)	–	–	–	–	–	–
Energy	−0.875***	–	–	–	–	–	–
	(0.200)	–	–	–	–	–	–
Env Ser	–	–	–	−0.311	–	–	–
	–	–	–	(0.732)	–	–	–
Oil & Gas	–	–	–	0.830**	–	–	–
	–	–	–	(0.402)	–	–	–
Telecom	−0.843***	–	–	−0.052	–	–	–
	(0.210)	–	–	(0.648)	–	–	–
Transport	0.193**	–	–	−0.515	–	–	–
	(0.069)	–	–	(0.318)	–	–	–
1999	2.184***	0.262	3.077***	1.216	–	−2.888	1.978
	(0.498)	(0.740)	(0.819)	(1.988)	–	(2.610)	(1.376)
2000	0.455	−0.087	0.322	1.456	−1.844	−0.844	2.605**
	(0.457)	(0.740)	(0.757)	(1.966)	(1.588)	(2.615)	(1.288)
2001	0.771*	0.075	1.254	0.774	−1.418	−2.121	1.198
	(0.438)	(0.585)	(0.770)	(1.968)	(1.588)	(2.612)	(1.308)
2002	0.747*	0.077	1.031	0.072	−2.246	−0.750	1.076
	(0.420)	(0.566)	(0.734)	(1.979)	(1.523)	(3.919)	(1.310)
2003	0.664	0.003	1.170	1.533	−0.143	0.218	1.344
	(0.413)	(0.552)	(0.741)	(1.960)	(1.526)	(2.998)	(1.273)
2004	0.399	−0.153	0.765	0.933	−0.354	−4.204	2.029
	(0.404)	(0.540)	(0.749)	(1.956)	(1.474)	(3.328)	(1.305)
2005	0.609	0.097	0.835	0.421	−1.503	−4.303	0.995
	(0.398)	(0.533)	(0.759)	(1.954)	(1.443)	(4.369)	(1.265)
2006	0.618	0.150	0.599	0.302	−1.955	−3.526	2.139
	(0.396)	(0.531)	(0.742)	(1.954)	(1.393)	(3.286)	(1.341)
2007	0.555	0.045	0.616	0.521	−1.555	−4.384	1.796
	(0.393)	(0.529)	(0.722)	(1.949)	(1.358)	(4.031)	(1.314)
2008	0.550	0.043	0.579	0.717	−1.156	−3.868	2.011
	(0.392)	(0.528)	(0.719)	(1.955)	(1.389)	(3.573)	(1.411)
2009	0.525	0.043	0.408	1.116	−0.812	−3.219	2.803**
	(0.392)	(0.527)	(0.727)	(1.946)	(1.365)	(3.324)	(1.343)
2010	0.511	0.037	0.548	0.487	−1.553	−4.206	2.048
	(0.391)	(0.527)	(0.719)	(1.949)	(1.364)	(3.013)	(1.327)
2011	0.468	−0.021	0.390	0.576	−1.847	−2.549	2.086
	(0.391)	(0.527)	(0.722)	(1.942)	(1.352)	(2.728)	(1.326)
2012	0.430	−0.042	0.406	0.812	−1.500	−3.370	2.412*
	(0.391)	(0.527)	(0.719)	(1.952)	(1.387)	(3.859)	(1.321)
2013	0.428	−0.055	0.582	0.507	−1.610	0.865	1.344
	(0.392)	(0.527)	(0.726)	(1.942)	(1.325)	(3.891)	(1.350)
2014	0.499	0.020	0.682	0.451	−1.970	−3.979	2.680*
	(0.398)	(0.531)	(0.969)	(1.958)	(1.354)	(4.409)	(1.487)
Leverage	−2.242***	−2.153***	−6.831***	−0.693*	−0.279	0.497	−1.426**
	(0.166)	(0.167)	(1.045)	(0.417)	(0.587)	(6.690)	(0.589)
Log assets	0.136***	0.122***	0.088	0.151*	0.340**	−0.233	0.037
	(0.022)	(0.021)	(0.172)	(0.085)	(0.152)	(0.716)	(0.101)
Constant	0.753	1.337**	5.142**	−0.703	−1.004	6.344	−0.287
	(0.461)	(0.582)	(2.414)	(1.968)	(1.956)	(6.895)	(1.364)
Observations	842	709	98	275	138	39	82
R^2	0.288	0.226	0.540	0.219	0.123	0.717	0.441
Adj. R^2	0.279	0.219	0.419	0.192	0.105	0.221	0.285
F Statistic	13.192***	9.566***	4.254***	2.115***	0.871	1.169	1.494
	(df = 25; 816)	(df = 21; 687)	(df = 21; 76)	(df = 32; 242)	(df = 19; 118)	(df = 26; 12)	(df = 28; 53)

Note: $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

9. Appendix: Data Tables

Table 25: Explanatory power of investment start years in explaining realised DSCRs in PLS regression model. Both contracted and merchant regressions control for sectors, but their coefficients are omitted from this table to save space.

	Dependent variable: log(DSCR)	
	Contracted	Merchant
Leverage	-2.377^{***}	-1.416^{***}
	(0.200)	(0.382)
log of Total Assets	0.141^{***}	0.253^{***}
	(0.023)	(0.094)
1986	–	0.701
	–	(0.846)
1993	–	2.896^{***}
	–	(0.724)
1995	–	0.296
	–	(0.756)
1996	1.357^{***}	-0.770
	(0.189)	(0.636)
1997	0.804^{***}	0.122
	(0.184)	(0.585)
1998	0.115	-0.837
	(0.176)	(0.612)
1999	0.300^{*}	-0.252
	(0.173)	(0.670)
2000	0.271	5.297^{***}
	(0.188)	(0.752)
2001	0.038	-0.603
	(0.159)	(0.653)
2002	-0.061	-1.401^{**}
	(0.159)	(0.666)
2003	0.053	-0.809
	(0.154)	(0.687)
2004	0.029	-1.032^{*}
	(0.155)	(0.620)
2005	0.258	-1.337^{**}
	(0.158)	(0.652)
2006	0.248	-1.326^{**}
	(0.158)	(0.626)
2007	0.077	-1.291^{*}
	(0.167)	(0.700)
2008	0.081	-1.743^{**}
	(0.201)	(0.800)
2009	0.047	-1.257^{*}
	(0.277)	(0.660)
2010	0.379	-1.325^{*}
	(0.358)	(0.676)
2011	0.005	-1.587^{*}
	(0.426)	(0.835)
2012	–	-1.360
	–	(1.104)
Observations	890	350
T	1-15	1-15
n	135	67
R^2	0.226	0.499
Adj. R^2	0.219	0.458
F Statistic	9.346^{***} (df = 27; 865)	11.109^{***} (df = 30; 334)
Note:	$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$	

9. Appendix: Data Tables

Table 26: Fixed year effects in panel regression model controlling for investment start year and sector (previous table).

	Dependent variable: log(DSCR)	
	Contracted	Merchant
Year 1	0.970***	2.264***
	(0.211)	(0.684)
Year 2	0.829***	1.988***
	(0.210)	(0.682)
Year 3	0.736***	2.050***
	(0.210)	(0.690)
Year 4	0.803***	1.915***
	(0.210)	(0.691)
Year 5	0.766***	1.921***
	(0.210)	(0.687)
Year 6	0.808***	1.709**
	(0.211)	(0.709)
Year 7	0.799***	2.012***
	(0.213)	(0.729)
Year 8	0.797***	1.720**
	(0.215)	(0.733)
Year 9	0.817***	1.677**
	(0.222)	(0.728)
Year 10	0.816***	1.691**
	(0.229)	(0.774)
Year 11	0.844***	1.059
	(0.251)	(0.807)
Year 12	0.822***	1.305
	(0.251)	(0.896)
Year 13	0.933***	1.494*
	(0.265)	(0.827)
Year 14	1.078***	2.473***
	(0.279)	(0.866)
Year 15	0.733**	2.443***
	(0.300)	(0.866)
Observations	890	350
T	1-15	1-15
n	135	67
R^2	0.241	0.526
Adj. R^2	0.230	0.445
F Statistic	6.275*** (df = 43; 849)	6.097*** (df = 56; 308)
Note:	*p<0.1; **p<0.05; ***p<0.01	

9. Appendix: Data Tables

Table 27: Panel regression of DSCRs controlling for financial ratios and project time (next table)

	Dependent variable: log(DSCR)	
	Contracted	Merchant
Profit Margin	0.886^{***}	-0.513^{*}
	(0.295)	(0.283)
Asset Turnover	0.851^{***}	0.360
	(0.261)	(0.277)
Cash Return on Assets	0.055	0.184
	(0.413)	(0.794)
Capex/Revenue Ratio	0.0001	0.00000
	(0.0002)	(0.00000)
Capex Coverage	−0.00004	−0.00000
	(0.0002)	(0.00000)
Commercial and Industrial	0.540	–
	(0.376)	–
Energy	0.072	1.201^{***}
	(0.370)	(0.259)
Government Services	0.933^{***}	–
	(0.117)	–
Telecom	0.797	–
	(0.666)	–
Oil and Gas	–	3.178^{***}
	–	(0.368)
Transport	0.987^{***}	1.414^{***}
	(0.183)	(0.266)
Observations	239	198
T	1-15	1-15
n	79	63
R^2	0.101	0.187
Adj. R^2	0.097	0.179
F Statistic	2.579^{***} (df = 10; 229)	5.451^{***} (df = 8; 190)
Note:	$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$	

9. Appendix: Data Tables

Table 28: Fixed year effects in panel regression model controlling for financial ratio and sector (previous table)

	Dependent variable: log(DSCR)	
	Contracted	Merchant
Year 1	0.570	1.539***
	(0.403)	(0.417)
Year 2	0.634	1.221***
	(0.408)	(0.405)
Year 3	0.388	1.249**
	(0.410)	(0.505)
Year 4	0.650	1.022**
	(0.430)	(0.516)
Year 5	0.658	0.701
	(0.453)	(0.535)
Year 6	0.620	0.715
	(0.467)	(0.589)
Year 7	0.507	1.328**
	(0.480)	(0.609)
Year 8	0.478	1.028
	(0.522)	(0.680)
Year 9	0.685	0.778
	(0.514)	(0.717)
Year 10	1.051*	1.376*
	(0.589)	(0.808)
Year 11	0.966*	−0.069
	(0.579)	(0.923)
Year 12	1.074*	0.233
	(0.587)	(1.110)
Year 13	1.245**	0.525
	(0.586)	(0.776)
Year 14	1.185**	1.444
	(0.587)	(1.008)
Year 15	0.598	1.302
	(0.631)	(1.080)
Observations	239	198
R^2	0.175	0.309
Adj. R^2	0.156	0.262
F Statistic	1.741** (df = 26; 213)	2.506*** (df = 30; 168)
Note:	*p<0.1; **p<0.05; ***p<0.01	

References

References

- Black, F. and J. C. Cox (1976). Valuing corporate securities: Some effects of bond indenture provisions. *The Journal of Finance 31*(2), 351–367.
- Blanc-Brude, F. (2013). Towards efficient benchmarks for infrastructure equity investments. *EDHEC-Risk Institute Publications*, 88.
- Blanc-Brude, F. (2014, June). Benchmarking Long-Term Investment in Infrastructure. *EDHEC-Risk Institute Position Paper*.
- Blanc-Brude, F., M. Hasan, and O. R. H. Ismail (2014). Unlisted Infrastructure Debt Valuation & Performance Measurement. *EDHEC-Risk Institute Publications July*.
- Blanc-Brude, F. and R. Strange (2007). How Banks Price Loans to Public-Private Partnerships: Evidence from the European Markets. *Journal of Applied Corporate Finance 19*(4), 94–106.
- Chemmanur, T. J. and K. John (1996). Optimal Incorporation, Structure of Debt Contracts, and Limited-Recourse Project Financing. *Journal of Financial Intermediation 5*(4), 372–408.
- Ciochetti, B. A., Y. Deng, G. Lee, J. D. Shilling, and R. Yao (2003). A proportional hazards model of commercial mortgage default with originator bias. *The Journal of Real Estate Finance and Economics 27*(1), 5–23.
- Esty, B. C. (2001, September). Structuring Loan Syndicates: A Case Study Of The Hong Kong Disneyland Project Loan. *Journal of Applied Corporate Finance 14*(3), 80–95.
- Esty, B. C. (2002). Returns on project financed investment. *Journal of Applied Corporate Finance 15*(1), 71–86.
- Esty, B. C. (2003). The economic motivations for using project finance. *Harvard Business School*, 28.
- Esty, B. C. and W. L. Megginson (2003). Creditor Rights, Enforcement, and Debt Ownership Structure: Evidence from the Global Syndicated Loan Market. *The Journal of Financial and Quantitative Analysis 38*(1), 37–59.
- Fink, D. (1997, December). A compendium of conjugate priors. *Montana The Magazine Of Western History*, 1–47.
- Flannery, M. J., J. F. Houston, and S. Venkataraman (1993). Financing multiple investment projects. *Financial Management*, 161–172.
- Gatti, S. (2013). *Project finance in theory and practice*. Academic Press.

References

- Harris, M. and A. Raviv (1990). Capital structure and the informational role of debt. *Journal of Finance*, 321–349.
- Jensen, M. C. and W. H. Meckling (1976). Theory of the firm: Managerial behavior, agency costs and ownership structure. *Journal of Financial Economics 3*(4), 305–360.
- John, T. A. and K. John (1991). Optimality of project financing: Theory and empirical implications in finance and accounting. *Review of Quantitative Finance and Accounting 1*(1), 51–74.
- Kealhofer, S. (2003). Quantifying credit risk I: default prediction. *Financial Analysts Journal*, 30–44.
- Modigliani, F. and M. H. Miller (1958). The cost of capital, corporation finance and the theory of investment. *The American Economic Review 48*(3), 261–297.
- Moody's (2012, January). Default and recovery rates for project finance bank loans 1983-2010. Technical report, Moody's Investors Service, London, UK.
- Moody's (2014). Global P3 Landscape: A region-by-region round-up. Technical report, Moody's Analytics, New York.
- Moody's (2015). Default and recovery rates for project finance bank loans, 1983-2014. Technical report, Moody's Investors Service.
- Myers, S. C. (1977). Determinants of Corporate Borrowing. *Journal of Financial Economics 5*(2), 147–175.
- Shah, S. and A. Thakor (1987). Optimal capital structure. *Journal of Economic Theory 42*, 209–243.
- Sorge, M. and B. Gadanecz (2008). The term structure of credit spreads in project finance. *International Journal of Finance and Economics 13*, 68–81.
- Standard and Poor's (2013). Project Finance Default & Recovery Study. Technical report, Standard & Poor's, London, UK.
- Yescombe, E. R. (2002). *Principles of Project Finance*. Academic Press.

About Natixis

About Natixis

Natixis is the international corporate, investment, insurance and financial services arm of Groupe BPCE, the 2nd-largest banking group in France with 36 million clients spread over two retail banking networks, Banque Populaire and Caisse d'Epargne.

With more than 16,000 employees, Natixis has a number of areas of expertise that are organised into three main business lines: Corporate & Investment Banking, Investment Solutions & Insurance, and Specialised Financial Services.

A global player, Natixis has its own client base of companies, financial institutions and institutional investors as well as the client base of individuals, professionals and small and medium-size businesses of Groupe BPCE's banking networks.

Listed on the Paris stock exchange, it has a solid financial base with a CET1 capital under Basel 3(1) of Eur12.6 billion, a Basel 3 CET1 Ratio [11] of 11% and quality long-term ratings (Standard & Poor's: A / Moody's: A2 / Fitch Ratings: A).

11 - Based on CRR-CRD4 rules published on June 26, 2013, including the Danish compromise - no phase-in except for DTAs on loss carry-forwards. Figures as at June 30, 2015

Natixis is a recognised player in the infrastructure space and has notably obtained the following rankings in 2014:

- #1 Arranger in France for PPP, Concessions or DSP by Le Magazine des Affaires
- #10 Global MLA for Project Finance by Thomson Reuters
- #10 Global Bookrunner for Project Finance by Thomson Reuters
- #9 Americas Advisory mandates won for Project Finance by Thomson Reuters

More information on Natixis infrastructure expertise available at: http://cib.natixis.com/infrastructure.

About the EDHEC Infrastructure Institute-Singapore

About the EDHEC Infrastructure Institute-Singapore

EDHEC*infra* addresses the profound knowledge gap faced by infrastructure investors by collecting and standardising private investment and cash flow data and running state-of-the-art asset pricing and risk models to create the performance benchmarks that are needed for asset allocation, prudential regulation and the design of new infrastructure investment solutions.

A Profound Knowledge Gap

Institutional investors have set their sights on private investment in infrastructure equity and debt as a potential avenue towards better diversification, improved liability-hedging and reduced drawdown risk.

Capturing these benefits, however, requires answering a number of difficult questions:

1. **Risk-adjusted performance measures** are needed to inform strategic asset allocation decisions and monitoring performance;
2. **Duration and inflation hedging properties** are required to understand the liability-friendliness of infrastructure assets;
3. **Extreme risk measures** are in demand from prudential regulators amongst others.

Today none of these metrics is documented in a robust manner, if at all, for investors in privately-held infrastructure equity or debt. This has left investors frustrated by an apparent lack of adequate investment solutions in infrastructure. At the same time, policy-makers have begun calling for a widespread effort to channel long-term savings into capital projects that could support long-term growth.

To fill this knowledge gap, EDHEC has launched a new research platform, EDHEC*infra*, to collect, standardise and produce investment performance data for infrastructure equity and debt investors.

Mission Statement

Our objective is the creation a global repository of financial knowledge and investment benchmarks about infrastructure equity and debt investment, with a focus on delivering useful applied research in finance for investors in infrastructure.

We aim to deliver the best available estimates of financial performance and risks of reference portfolios of privately-held infrastructure investments, and to provide investors with important insights about their strategic asset allocation choices to infrastructure, as well as support the adequate calibration of the relevant prudential frameworks.

We are developing unparalleled access to the financial data of infrastructure projects and firms, especially private data that is either unavailable to market participants or cumbersome and difficult to collect and aggregate.

We also bring advanced asset pricing and risk measurement technology designed to answer investors' information needs about long-term investment in privately-held infrastructure, from asset allocation to prudential regulation and performance attribution and monitoring.

What We Do

The EDHECinfra team is focused on three key tasks:

1. **Data collection and analysis**: we collect, clean and analyse the private infrastructure investment data of the project's data contributors as well as from other sources, and input it into EDHECinfra's unique database of infrastructure equity and debt investments and cash flows. We also develop data collection and reporting standards that can be used to make data collection more efficient and reporting more transparent.

About the EDHEC Infrastructure Institute-Singapore

This database already covers 15 years of data and hundreds of investments and, as such, is already the largest dedicated database of infrastructure investment information available.

2. **Cash flow and discount rate models**: Using this extensive and growing database, we implement and continue to develop the technology developed at EDHEC-Risk Institute to model the cash flow and discount rate dynamics of private infrastructure equity and debt investments and derive a series of risk and performance measures that can actually help answer the questions that matter for investors.

3. **Building reference portfolios of infrastructure investments**: Using the performance results from our asset pricing and risk models, we can report the portfolio-level performance of groups of infrastructure equity or debt investments using categorisations (e.g. greenfield vs brownfield) that are most relevant for investors' investment decisions.

Partners of EDHEC*infra*

Monetary Authority of Singapore

In October 2015, the Deputy Prime Minister of Singapore, Tharman Shanmugaratnam, announced officially at the World Bank Infrastructure Summit that EDHEC would work in Singapore to create "usable benchmarks for infrastructure investors."

The Monetary Authority of Singapore is supporting the work of the EDHEC Singapore Infrastructure Investment Institute (EDHEC infra) with a five-year research development grant.

Sponsored Research Chairs

Since 2012, private sector sponsors have been supporting research on infrastructure investment at EDHEC with several research Chairs that are now under the EDHEC Infrastructure Investment Institute:

1. The EDHEC/NATIXIS Research Chair on the Investment and Governance Characteristics of Infrastructure Debt Instruments, 2012-2015
2. The EDHEC/Meridiam/Campbell Lutyens Research Chair on Infrastructure Equity Investment Management and Benchmarking, 2013-2016
3. The EDHEC/NATIXIS Research Chair on Infrastructure Debt Benchmarking, 2015-2018
4. The EDHEC/Long-Term Infrastructure Investor Association Research Chair on Infrastructure Equity Benchmarking, 2016-2019
5. The EDHEC/Global Infrastructure Hub Survey of Infrastructure Investors' Perceptions and Expectations, 2016

Partner Organisations

As well as our Research Chair Sponsors, numerous organisation have already recognised the value of this project and have joined or are committed to join the data collection effort. They include:

- The European Investment Bank;
- The World Bank Group;
- The European Bank for Reconstruction and Development;
- The members of the Long-Term Infrastructure Investor Association;
- Over 20 other North American, European and Australasian investors and infrastructure managers.

About the EDHEC Infrastructure Institute-Singapore

EDHEC*infra* is also :

- A member of the Advisory Council of the World Bank's Global Infrastructure Facility
- An honorary member of the Long-term Infrastructure Investor Association

Origins and Recent Achievements

In 2012, EDHEC-Risk Institute created a thematic research program on infrastructure investment and established two Research Chairs dedicated to long-term investment in infrastructure equity and debt, respectively, with the active support of the private sector.

Since then, infrastructure investment research at EDHEC has led to more than 20 academic publications and as many trade press articles, a book on infrastructure asset valuation, more than 30 industry and academic presentations, more than 200 mentions in the press and the creation of an executive course on infrastructure investment and benchmarking.

Testament to the quality of its contributions to this debate, EDHEC infra's research team has been regularly invited to contribute to high-level fora on the subject, including G20 meetings.

Likewise, active contributions were made to the regulatory debate, in particular directly supporting the adaptation of the Solvency-2 framework to long-term investments in infrastructure.

This work has contributed to growing the limited stock of investment knowledge in the infrastructure space.

Significant **empirical findings** already include:

- The first empirical estimates of construction risk for equity and debt investors in infrastructure project finance;

- The only empirical tests of the statistical determinants of credit spreads in infrastructure debt since 2008, allowing controlling for the impact of market liquidity and isolating underlying risk factors;

- The first empirical evidence of the diversification benefits of investing in greenfield and brownfield assets, driven by the dynamic risk and correlation profile of infrastructure investments over their lifecycle;

- The first empirical documentation of the relationship between debt service cover ratios, distance to default and expected default frequencies;

- The first measures of the impact of embedded options in senior infrastructure debt on expected recovery, extreme risk and duration measures;

- The first empirically documented study of cash flow volatility and correlations in underlying infrastructure investment using a large sample of collected data covering the past fifteen years.

Key **methodological advances** include:

- A series of Bayesian approaches to modelling cash flows in long-term investment projects including predicting the trajectory of key cash flow ratios in a mean/variance plane;

- The first fully-fledged structural credit risk model of infrastructure project

About the EDHEC Infrastructure Institute-Singapore

finance debt;

- A robust framework to extract the term structure of expected returns (discount rates) in private infrastructure investments using conditional volatility and initial investment values to filter implied required returns and their range at one point in time across heterogenous investors.

Recent **contributions to the regulatory debate** include:

- A parsimonious data collection template to develop a global database of infrastructure project cash flows;
- Empirical contributions to adapt prudential regulation for long-term investors.

Infrastructure Research Publications at EDHEC (2012-16)

Infrastructure Research Publications at EDHEC (2012-16)

EDHEC Publications

- Blanc-Brude, F., T. Whittaker and M. Hasan. Cash Flow Dynamics of Private Infrastructure Debt (March 2016).
- Blanc-Brude, F., T. Whittaker and M. Hasan. Revenues and Dividend Payouts in Privately-Held Infrastructure Investments (March 2016).
- Blanc-Brude, F., and M. Hasan. The Valuation of Privately-Held Infrastructure Equity Investments (January 2015).
- Blanc-Brude, F., M. Hasan and O.R.H. Ismail. Performance and Valuation of Private Infrastructure Debt (July 2014).
- Blanc-Brude, F., Benchmarking Long-Term Investment in Infrastructure (June 2014).
- Blanc-Brude, F., and D. Makovsek. How Much Construction Risk do Sponsors take in Project Finance. (August 2014).
- Blanc-Brude, F. and O.R.H. Ismail. Who is afraid of construction risk? (March 2013)
- Blanc-Brude, F. Towards efficient benchmarks for infrastructure equity investments (January 2013).
- Blanc-Brude, F. Pension fund investment in social infrastructure (February 2012).

Books

- Blanc-Brude, F. and M. Hasan, Valuation and Financial Performance of Privately-Held Infrastructure Investments. London: PEI Media, Mar. 2015.

Peer-Reviewed Publications

- F. Blanc-Brude, S. Wilde, and T. Witthaker, "Looking for an infrastructure asset class Definition and mean-variance spanning of listed infrastructure equity proxies", 2016 (*forthcoming*)
- Blanc-Brude, F., M. Hasan, and T. Witthaker, "Benchmarking Infrastructure Project Finance - Objectives, Roadmap and Recent Progress", Journal of Alternative Investments, 2016 (*forthcoming*)
- R. Bianchi, M. Drew, E. Roca and T. Whittaker, "Risk factors in Australian bond returns", Accounting & Finance, 2015

Infrastructure Research Publications at EDHEC (2012-16)

- Blanc-Brude, F. "Long-term investment in infrastructure and the demand for benchmarks," JASSA The Finsia Journal of Applied Finance, vol. 3, pp. 57–65, 2014.
- Blanc-Brude, F. "Risk transfer, self-selection and ex post efficiency in public procurement: an example from UK primary and secondary school construction contracts," Revue d'Economie Industrielle, vol. 141, no. 1st Quarter, pp. 149–180, 2013.
- Blanc-Brude, F. , H. Goldsmith, and T. Valila, "A comparison of construction contract prices for tradition- ally procured roads and public–private partnerships," Review of Industrial Organization, vol. 35, no. 1-2, pp. 19–40, 2009, ISSN: 0889-938X. DOI: 10.1007/s11151-009-9224-1.
- Blanc-Brude, F. , H. Goldsmith, and T. Valila, "Public-private partnerships in europe: an update," EIB Economic & Financial Reports, p. 24, 2007.
- Blanc-Brude, F. and R. Strange, "How banks price loans to public-private partnerships: evidence from the european markets," Journal of Applied Corporate Finance, vol. 19, no. 4, pp. 94–106, 2007.
- Blanc-Brude, F. , H. Goldsmith, and T. Valila, "Ex ante construction costs in the european road sector: a comparison of public-private partnerships and traditional public procurement," EIB Economic & Financial Reports, vol. 2006/1, 2006.
- O. Jensen and F. Blanc-Brude, "The handshake: why do governments and firms sign private sector participation deals? evidence from the water and sanitation sector in developing countries," World Bank Working Papers, Wold Bank Working Paper Series, no. October 2005, p. 25, 2006.

For more information, please contact:
Karen Sequeira on +65 6438 0030
or e-mail: karen.sequeira@edhec.edu

EDHEC Infrastructure Institute-Singapore
EDHEC Business School Asia-Pacific
One George Street - #07-02
Singapore 049145
Tel.: +65 6438 0030

edhec.infrastructure.institute