A rigorous yet implementable framework to measure the performance of private infrastructure debt

Frédéric Blanc-Brude, Research Director, EDHEC Risk Institute–Asia;
Majid Hasan, Research Assistant, EDHEC Risk Institute–Asia

Estimating the performance of infrastructure debt instruments has become a recurring question for both long-term investors and prudential regulators. In a new paper[1], we propose the first robust valuation, risk measurement and data collection framework for private infrastructure project loans.

We focus on those performance measures that are the most relevant to investors at the strategic asset allocation level, and to prudential regulators for the calibration of risk weightings, including expected loss, expected recovery rates, loss given default, value-at-risk (VaR), expected shortfall or CVaR, duration, yield and z-spread. We also determine parsimonious data collection requirements. Hence, we can realistically ►

1 Blanc-Brude, F., M. Hasan and O. R. H. Ismail (2014). Unlisted Infrastructure Debt Performance Measurement. Natixis research chair on Infrastructure Debt Investment Solutions. Singapore: EDHEC-Risk Institute. This paper and the present article are drawn from the Natixis research chair at EDHEC-Risk Institute on the Investment and Governance Characteristics of Infrastructure Debt Instruments.

◄ expect to deliver these performance measures at a minimal data collection cost.

Marking to business models

As for any security, the valuation of infrastructure project finance loans consists of modelling or observing cash flows and deriving their present value. However, cash flow data for large, representative samples of projects and spanning their entire lifecycle are not yet available. It is one of the objectives of this article and of the Natixis research chair on infrastructure debt at EDHEC-Risk Institute to determine what data needs to be collected and to build a global database of infrastructure project debt metrics.

Until such a large dataset has been built, we must proceed in two steps: we first model the cash flows of generic but commonly found infrastructure project financing structures and calibrate these models with existing data, allowing for new calibrations when larger datasets become available. Second, given a generic cash flow model, we build a valuation model to derive the relevant return and risk measures. So we are effectively marking to a business model that reflects the available knowledge of private infrastructure project credit risk today.

This exercise also yields the list of data points that need to be collected to update and better calibrate this model and improve our knowledge of the performance of infrastructure debt.

We show that documenting the dynamics of the debt service cover ratio (DSCR) in infrastructure project finance vehicles, in combination with information about initial leverage, covenant triggers and the size of the loan's 'tail' (ie, the difference between the original loan maturity and the life of the project) is sufficient to derive key credit risk metrics in infrastructure project finance, including default frequencies and distance to default.

In particular, we show that knowledge of the statistical distribution of the DSCR in infrastructure projects is sufficient to predict default and compute distance to default measures, allowing the development and implementation of a powerful structural credit risk model à la Merton (1974).

In terms of data availability, we know that DSCRs are typically monitored and recorded by infrastructure creditors, while the base case debt service and other project characteristics are documented at financial close. Hence, the data that is required is observable.

Families of infrastructure debt financing

We focus on two families of financial structures, which correspond to a large number of real-world infrastructure projects and their associated debt securities: 1) merchant infrastructure and 2) contracted infrastructure.

Merchant infrastructure refers to those projects that generate revenue by selling their output or service in a market, and hence are exposed to commercial risks, while contracted infrastructure projects receive contracted revenue in exchange for providing a pre-agreed output or service, and bear little to no market risk.

Examples of merchant infrastructure projects include a power plant that sells electricity at market prices or a road collecting tolls from users. Examples of contracted projects may include schools and hospitals that receive a fixed payment from a government entity upon the satisfactory delivery and maintenance of an infrastructure, or an energy project financed on the back of take-or-pay purchase agreement.

1. Rising debt service cover ratio and probability of default

Generic 25-year merchant infrastructure project with five-year construction period and seven-year tail

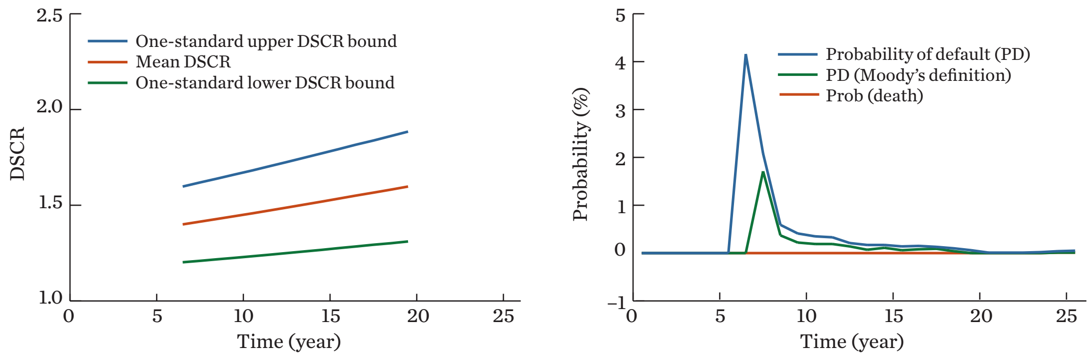

In the right-hand figure, the blue line indicates the probability of technical default (PD), the green line the probability of hard default (Moody's definition) and the red line the probability of bankruptcy.

These two types of project structures correspond to different underlying business risks, and as a consequence, that are structured in different ways. As illustrated in figure 1, merchant infrastructure projects are structured with a rising mean DSCR and a longer tail. A rising DSCR implies that lenders get paid faster than equity owners, and a longer tail increases the value of lenders' security. In other words, lenders demand an increasing mean DSCR and a longer tail to protect themselves against higher and increasing cash flow volatility, which results from higher revenue risk.

In contrast, contracted projects are structured with a flat DSCR and shorter tails, as lenders demand less protection against default due to lower expected underlying revenue risk.

Of course, other generic infrastructure project financing structures exist, even though they tend to be a combination of these two types – eg, shadow toll roads collect a volume-based income paid typically by a government.

The value of the tail

Combined with the base case debt service, the distribution of the DSCR in generic projects can also be used to infer the expected value and volatility of the cash flow available for debt service (CFADS) of a typical infrastructure project.

However, the base case determined at financial close can change following a breach of covenant or a hard default, leading to the restructuring of the debt schedule and its extension in the loan's tail. To take into account these potential changes in the debt schedule and value the tail, we model the debt renegotiation process to determine the outcome of restructuring after either a technical (covenant-driven) or a hard default (of payment). A new debt service is determined by taking into account what each party would lose in the absence of a workout.

Thus, we can determine the cash flows to infrastructure project creditors in every future state of the world. To value these cash flows, we take a so-called structural approach and develop a version of the Merton (1974) model that takes into account the characteristics of infrastructure project finance debt, and combine value of the debt service in different states by adapting the Black and Cox (1976) decomposition to the case of project finance workouts.

Key findings: a low but dynamic risk profile

For parameter estimates corresponding to current industry practices, we find that the debt of both types of generic infrastructure projects discussed, merchant and contracted, exhibit highly dynamic risk profiles. Our results show the importance of 'valuing the loan's tail' to get an accurate picture of the risk profile of infrastructure project debt. Overall, we find that the different aspects of credit risk in infrastructure projects can largely be explained by their DSCR profiles, tail values, and the costs of exit relative to the cost of renegotiation for lenders.

Using Moody's definition of default in project finance by which each loan is only allowed to default once (Moody's [2014]) our model predicts marginal default frequencies in line with reported empirical averages: trending downwards from just under 2% at the beginning of the loan's life to almost zero after 10 years, in the case of merchant infrastructure, and flat at 0.5% for contracted projects – ie, public-private partnerships (see for example Moody's [2014]).

Overall, risk levels are found to be relatively low and recovery relatively high: Expected loss (EL) never rises above 2%; VaR and CVaR, while they increase towards the end of the loan's life as the value of the tail is exhausted, never reach levels higher than 6% and 10% respectively; expected recovery rates are always in the 80% to 100% range.

In the case of merchant infrastructure projects, the probability of both technical and hard defaults (PD), and of hard defaults only (Moody's definition) shown in the top-left quadrant of figure 2, goes down sharply post construction, while expected (EL) and extreme (VaR, CVaR) losses tend to rise throughout the loan's life. Similarly, in the case of contracted infrastructure projects, while PD stays almost constant during the loan's life, the severity of losses increases with time.

The diverging trends in the distribution of defaults and losses are a consequence of debt restructurings upon defaults. Even if defaults are concentrated in a certain period of time, debt restructuring spreads losses over the entire life of the project. Hence, losses tend to increase with time, as the cumulative number of defaults (and hence restructurings) accrue losses near the end of the loan's life. However, part of the losses suffered during the loan's life is recovered in the loan's tail, thus reducing overall expected losses.

The size of losses for both DSCR families is primarily influenced by lenders' exit value net of exit costs. Exit costs determine the aggregate loss of value (debt plus equity) if the debt owners take over the project company upon a hard default and do not renegotiate with the original equity investors.

The higher the exit costs, the lower the value that lenders can obtain by taking over the project company after a hard default, and the lower their bargaining power in negotiations with original equity holders. Hence, lenders may have to suffer losses even in otherwise low risk projects like contracted infrastructure because replacing the equity owners upon a hard default, while it is in their power, can be very costly in some cases.

As a consequence, ongoing monitoring of the special purpose entity (SPE) conducted is required of lenders in the project in order to avoid ever having to contemplate exercising their option to exit, in particular, technical default triggers (eg, a low DSCR or loan-life cover ratio) allow lenders to intervene and maximise their recovery rates long before more expensive options to restructure, sell or liquidate the SPE ever arise.

Finally, our approach also allows us to derive expected returns and yield measures and highlights the fact that the ability to reschedule debt upon technical and hard default creates a trade-off between credit risk and duration risk. That is, to reduce the credit losses upon default, investors have to extend the maturity of their loan further in the tail, and have to bear a higher interest rate risk due to a higher duration. This trade-off is shown in the bottom-right quadrant of figure 2.

2. Credit risk metrics for merchant infrastructure

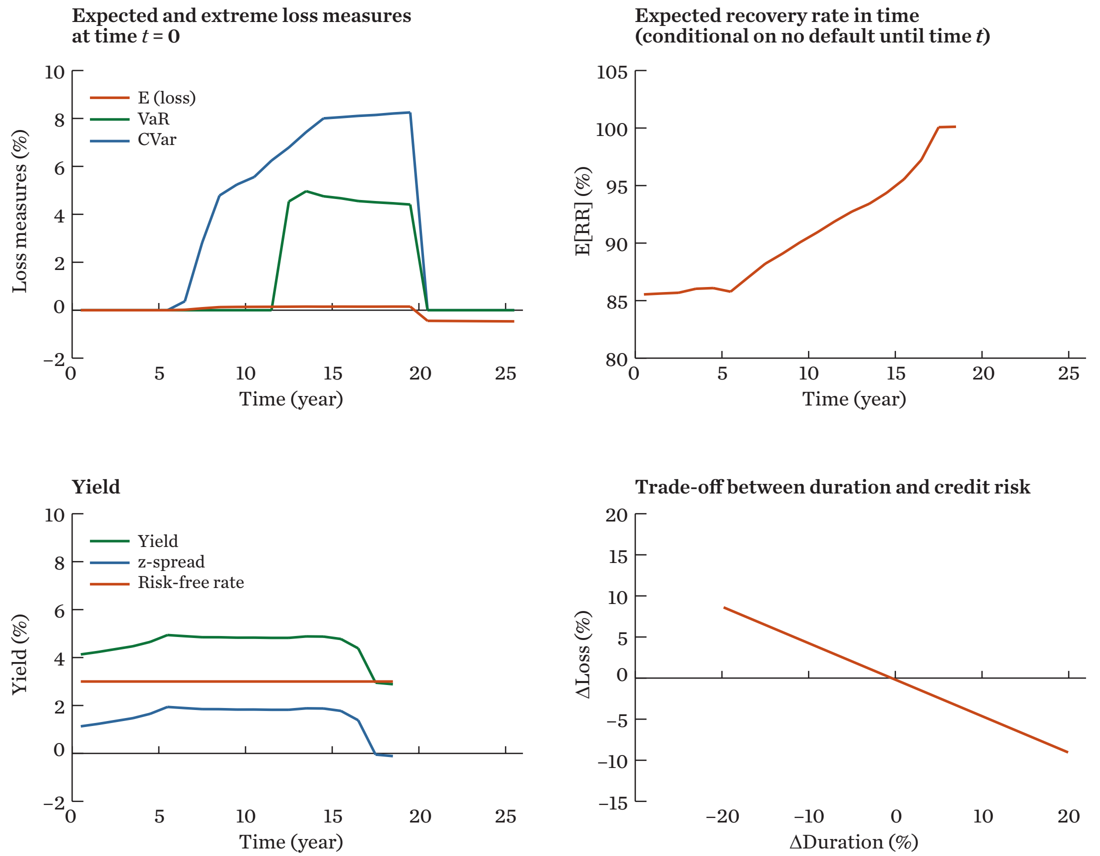

Next steps: data collection and portfolio construction

Thus, with a parsimonious set of inputs that consists of the parameters of the DSCR distribution across different types of generic projects, the base case debt schedule and a number of variables defined in the covenants at financial close, infrastructure project finance loans can be valued at any point in time, and their risk/return profile can be constructed spanning the entire life of the loan.

Our study delivers the first three steps of the roadmap defined in Blanc-Brude (2014) with respect to infrastructure debt investment: defining the most relevant underlying financial instrument, designing a valuation framework that is adapted to its private and illiquid nature, and the determination of a standard for data collection and investment performance reporting in infrastructure debt investment.

Next steps include active data collection to better calibrate our model of dynamics, before moving to the portfolio level of the analysis, towards long-term investment benchmark in infrastructure debt.

The research from which this article was drawn was produced as part of the Natixis Investment and Governance Characteristics of Infrastructure Debt Instruments research chair at EDHEC-Risk Institute.

References

Black, F., and J. C. Cox (1976). Valuing Corporate Securities: Some Effects of Bond Indenture Provisions. *Journal of Finance* 31(2): 351–67.

Blanc-Brude, F. (2014). Benchmarking Long-Term Investment in Infrastructure. EDHEC-Risk Institute position paper, July.

Blanc-Brude, F., M. Hasan and O. R. H. Ismail (2014). Unlisted Infrastructure Debt Performance Measurement. Natixis research chair on the Investment and Governance Characteristics of Infrastructure Debt Instruments. Singapore: EDHEC-Risk Institute.

Merton, R. (1974). On the Pricing of Corporate Debt: The Risk Structure of Interest Rates. *Journal of Finance* 29: 449–70.

Moody's Investors Service(2014). Default and Recovery Rates for Project Finance Bank Loans, 1983–2012.