IP&E
INVESTMENT & PENSIONS EUROPE
SPRING 2015

EDHEC-Risk Institute
Research Insights

EDHEC-RISK DAYS SPECIAL

EDHEC-RISK
Institute

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

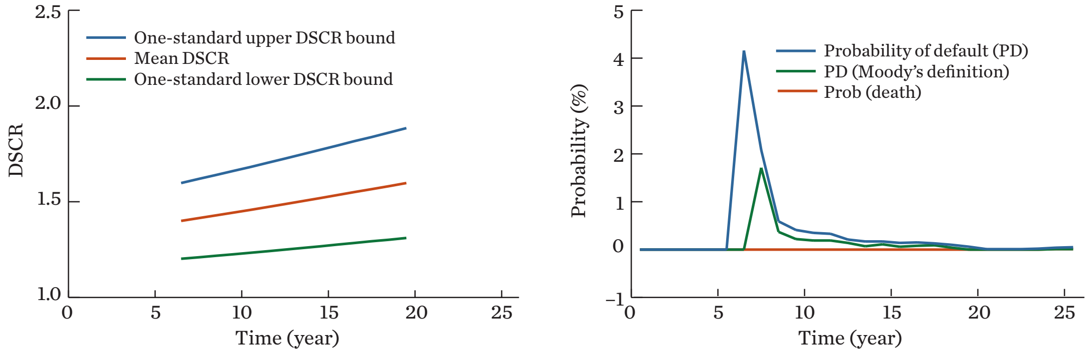

In the right-hand figure, the blue line indicates the probability of technical default (PD), the green line the probability of hard default (Moody's definition) and the red line the probability of bankruptcy.

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

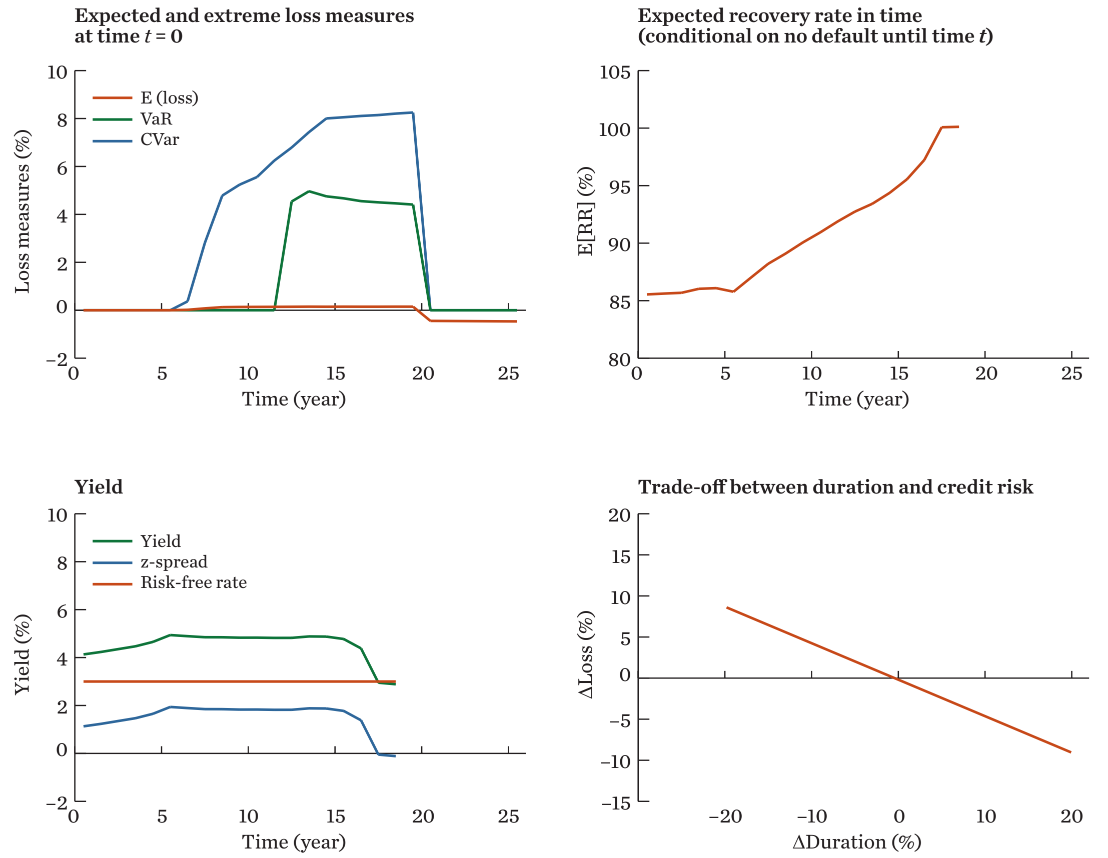

Next steps: data collection and portfolio construction

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 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.

Building a long-term investment benchmark for privately-held infrastructure equity

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

In new research[1] drawn from the work of the Meridiam/Campbell Lutyens research chair at EDHEC-Risk Institute, we propose the first valuation framework dedicated to privately-held infrastructure equity investments.

Following the roadmap to create long-term infrastructure investment benchmarks described in Blanc-Brude (2014), we develop a framework that takes into account the challenges of valuing privately-held and seldom-traded infrastructure equity investments, with the aim of designing a methodology that can be readily applied given the current state of empirical knowledge and, going forward, at a minimum cost in terms of data collection.

Three challenges

The valuation of unlisted infrastructure pro-

►

---

[1] Blanc-Brude, F., and M. Hasan (2015). The Valuation of Privately-Held Infrastructure Equity Investment. Meridiam/Campbell Lutyens research chair at EDHEC-Risk Institute on Infrastructure Equity Benchmarking. Singapore: EDHEC-Risk Institute, March.

◄ ject equity stakes requires three significant challenges to be addressed:

- ***Endemic data paucity***: while primary and secondary market prices can be observed, sufficiently large and periodic samples, representative of different types of infrastructure projects at each point in their multi-decade lifecycle, are unlikely to be available every year in each regional market.
- ***The term structure of expected returns***: the nature of such investments requires the estimation of a term structure of discount factors at different points in their lives that reflects the change in their risk profile. Indeed, in expectation, infrastructure investments can exhibit a dynamic risk profile determined by the sequential resolution of uncertainty, the frequent deleveraging of the project company's balance sheet or the existence of a fixed-term to the investment, which creates a time-varying duration.
- ***The absence of a unique price*** for a given investment in unlisted infrastructure, which springs from the fact that there is no traded equivalent to the payoff of infrastructure project equity. It follows that prices are partly driven by investor preferences and that substantial bid/ask spreads are likely.

The first point is partly a mundane aspect of the difficulties encountered when collecting data on private investments, but also a reflection of the nature of long-term equity investment in infrastructure. Indeed, the type of infrastructure projects that have been financed in the past are not necessarily representative of investment opportunities today. Thus, even if year-23 dividends for projects that were financed 24 years ago can be observed today, they may not be good predictors of dividends in projects financed three years ago, 20 years from now. For example, projects financed in the early 1990s may have been in sectors where fewer projects exist today (eg, telecoms) or rely on contractual structures or technologies that are not relevant to long-term investors in infrastructure today (eg, coal-fired merchant power).

If data paucity is an endemic dimension of the valuation of privately-held infrastructure equity investments, we must start from the premise that we cannot observe enough data simply to derive prices empirically. Instead, we acknowledge a position of relative ignorance and aim to build the possibility of improving our knowledge into our approach as new observations that can be used to update models of dividend distributions become available.

The second point about the term structure of expected returns has long been made in the finance literature: using such constant and deterministic discount rates is defective if projects have multiple phases and project risk changes over time as real options are exercised by asset owners.

It also amounts to assuming that the risk-free rate, asset beta and market risk premium are constant and deterministic, when we know that such variables are time-varying and stochastic. Moreover, the internal rate of return (IRR) of individual investments cannot be easily used to estimate performance at the portfolio level, as the IRR of a portfolio is not the same as the weighted average IRRs of individual investments.

Thus, using methodologies based on discounting at a constant rate, while common in the corporate sector, is inadequate for the purposes of long-term investors who need performance measures that can help them make hedging, risk management, and portfolio management decisions.

The third point (the absence of unique pricing measures) is a reflection of what are usually labelled 'incomplete markets' – ie, the fact that the same asset can be valued differently by two investors – and yet this does not constitute an arbitrage opportunity (and therefore the bid-ask spread does not narrow) because transaction costs are high and because in the absence of complete markets, investors' heterogeneous preferences partly explain prices.

The existence of a range of (or bounds on) values is also impacted by market dynamics: if a new type of investor (eg, less risk averse) enters the private infrastructure equity market, the range of observable valuations for similar assets may change. Likewise, if some investors want to increase their allocations to unlisted assets, given the limited available stock of investable infrastructure projects at a given point in time, their valuations may rise, but not that of others (who may sell).

Hence the important point that the required rate of return or discount rate of individual investors' infrastructure equity is fundamentally unobservable: it cannot be inferred from observable transaction prices since it is both a function of the characteristics of the asset (eg, cash flow volatility) and individual investor preferences.

Existing approaches are inadequate

Because of these challenges, existing approaches developed to value private equity investments are mostly inadequate for the purpose of valuing unlisted infrastructure project equity.

In our review of the literature we identify three groups of valuation techniques: repeat sales, public market equivalents and factor extraction from cash flows. Importantly, these techniques all imply that enough data can be observed to compute a price.

The repeat sales approach assumes that asset betas can be inferred from discrete and unevenly timed transaction observations after correcting for price staleness and sampling bias, while the public market equivalent approach implies that public asset betas can be combined to proxy the return of unlisted assets. Cash flow-driven approaches are less normative and aim to derive the unobservable rate of return of unlisted assets by decomposing their implied returns into traded and untraded components ex post facto – that is, once all cash flows have been observed and can be related to equally observable market factors.

Thus, these approaches cannot be directly applied to privately-held infrastructure investments, the value of which is determined by streams of expected and risky cash flows that mostly occur in the future, and for which few comparable realised investments exist today.

Existing approaches also typically fail to take into account the subjective dimension of asset pricing in the unlisted space and compute asset betas and alphas as if a unique pricing measure existed – ie, as if all investors had similar preferences, and in some papers, as if private equity exposures could always be replicated with a combination of traded assets.

Endogenously determined discount factors

To the extent that infrastructure dividend cash flows can only be partially observed today, their expected values cannot be decomposed into exogenous factors (markets, the economy, etc), the future value of which is not known today and would be very perilous to predict 30 years from now.

Instead, we must derive the relevant discount factors endogenously – ie, using observable information about each private investment in infrastructure equity including, as suggested above, its contractual characteristics, location, financial structure etc, as well as the value of the initial equity investment made, which is also observable.

Hence we argue that a robust valuation framework for equity investments that solely create rights to future (and yet largely unobserved) risky cash flows, as is the case of privately-held infrastructure equity, requires two components:

- A model of expected dividends and conditional dividend volatility, calibrated to the best of our current knowledge;
- A model of endogenously determined discount factors, that is, the combination of expected returns implied by the distribution of future dividends, given observable investment values.

In other words, as for any other stock, the valuation of privately-held equity in infrastructure projects amounts to deriving the appropriate discount rates for a given estimate of future dividends. But while this process is implicit in the pricing mechanism of public stock markets, in the case of privately-held equity with distant payoffs, we have to derive the relevant parameters explicitly, taking into account the characteristics of infrastructure assets.

Dividend distribution model and required data

The dividend stream or cash flow process can be described as state-dependent and we introduce a new metric for infrastructure project dividends: the equity service cover ratio (ESCR), which is computed as the ratio of realised-to-base case dividends.

The base case equity forecast of infrastructure equity investments, while not necessarily accurate, provides a useful and observable quantity, which by definition spans the entire life of each investment. Thus, we propose to describe the behaviour of equity cash flows in infrastructure projects as a function of this initial forecast, in order to create metrics allowing direct comparisons between different equity investments.

In our research, we show that the value of the ESCR at each point in the lifecycle of infrastructure equity investments can be used as a state variable describing the dynamics of the cash flow process. In combination with a given project's base case dividend forecast (which is known at the time of investment), knowledge of the distribution of the ESCR at each point in time is sufficient to express the expected value and conditional volatility of dividends.

The fact that new observations are not redundant today (we can still learn about the dynamics of dividends in infrastructure investment by collecting new data) justifies the need for an ongoing and standardised reporting of these cash flows to keep learning about their true distribution and value the infrastructure investments made today, tomorrow.

Filtering implied market values (and their bounds)

Since the term structure of expected returns of individual investors/deals is unobservable and lies within a range (or bounds) embodying market dynamics at a given point in time, we adapt the classic state-space model mostly used in physical and natural sciences to capture the implied average valuation (or state) of the privately-held infrastructure equity market at one point in time and its change from period to period. Using such a model also allows us to capture the market bounds on value implied by observable investment decisions for a given ►

◀ stream of expected cash flows.

The objective of state-space models is parameter estimation and inference about unobservable variables in dynamic systems – that is, to capture the dynamics of observable data in terms of an unobserved vector, here the term structure of discount factor. Hence, we have an observation equation relating observable data to a state vector of discount factors, and a state equation, which describes the dynamics of this state, from one observation (transaction) to the next. Each transaction corresponds to a new state – ie, a given term structure of discount factors matching the price paid in that transaction (the initial investment) with expected cash flows, which may or may not be the same as the previous transaction's.

Given a stream of risky future dividends, if the price paid in the current transaction is different from that paid in the previous one, it must be because the valuation state has shifted. The valuation state can change due to a change in investor preferences between the two deals, or due to a change in the consensus risk profile of that kind of investment (eg, projects with commercial revenues after a recession), or because of a change in the overall market sentiment (the average) valuation.

Thus, by iterating through transactions, we may derive an implied average valuation state (a term structure of discount factors) and its range, bounded by the highest and lowest bidders in the relevant period.

Later, when dividend payments are realised, period returns can be computed using the discounted sum of remaining cash flows as the end-of-period value (given the implied term structure of discount factors at that point).

In our research, we define the observation equation using a dynamic version of the standard Gordon growth model (discounted dividends) and the state equation using an autoregressive model of the term structure of expected returns which can be derived from the kind of factor models of expected excess returns that are commonly found in the literature. We take the view that expected returns are a function of conditional dividend volatility.

In a simple, linear setting, we show that we can iterate through observable investments, while estimating model parameters on a rolling basis, to capture both the implied expected returns (and discount factors) during a given reporting period and track these values and their range (arbitrage bounds) from period to period.

1.

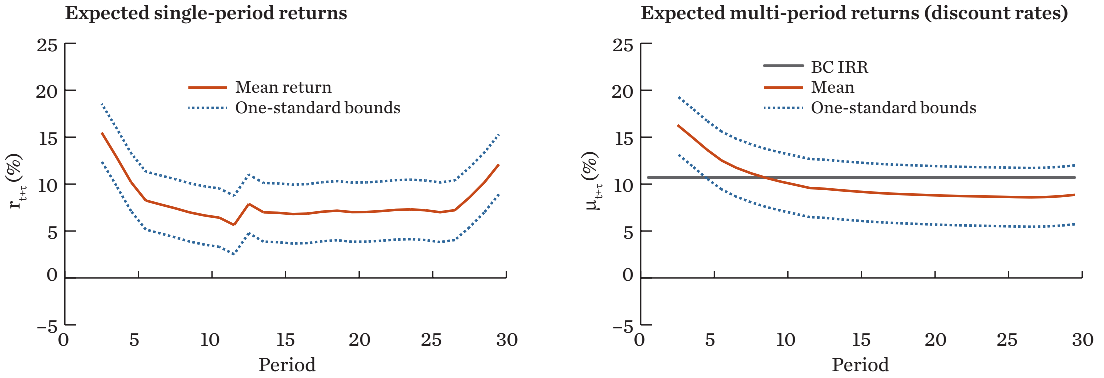

2.

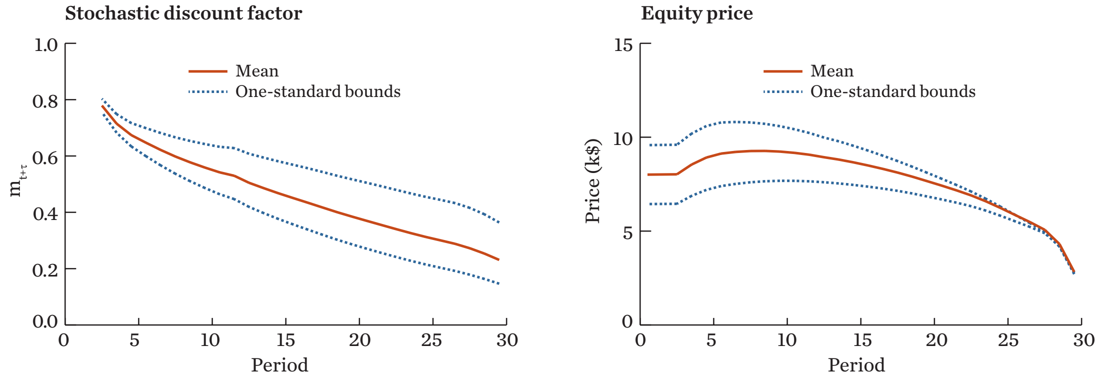

3.

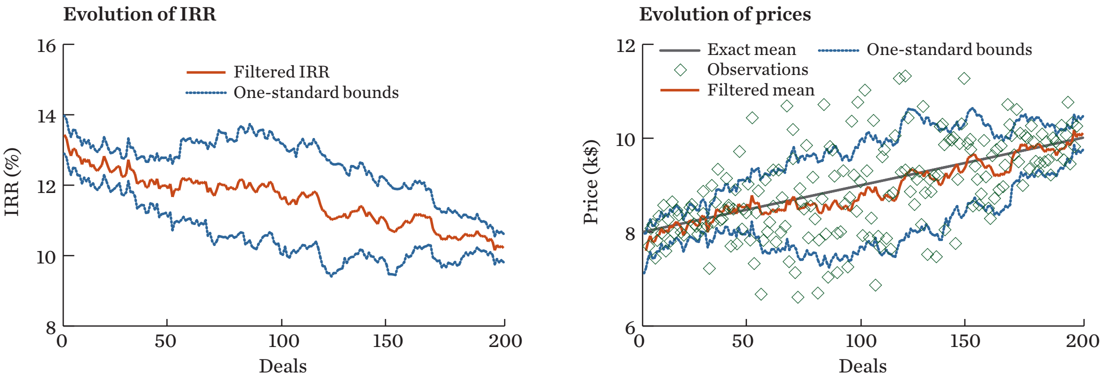

Illustration

As an illustration of our approach, we apply the dividend and pricing models to a generic case of privately-held infrastructure investment, assuming an expected ESCR and ESCR volatility profile (including the probability of receiving no dividends in any given period).

Given a base case dividend scenario inspired by an actual infrastructure project financed in Europe in the last decade, we obtain a full distribution of future dividends and apply our valuation framework to this assumed dividend process for an (equally assumed) range of investment values. Some of the key outputs are shown in the following figures.

Figure 1 shows the resulting filtered term structure of expected period and multi-period (average) expected returns filtered from a range of 20 initial transactions.

Figure 2 shows the resulting values of the dividend discount factor[2] at the time of valuation and the expected average price and its range for this group of transactions.

Finally, figure 3 shows how we can implement this model with rolling parameter estimation to track the implied average expected returns and price of consecutive transactions from period to period.[3]

These results spring from model inputs that are only inspired by existing data and a number of intuitions about privately-held infrastructure equity investments, and can only be considered an illustration. However, they show clearly that, with well-calibrated cash flow models and a transparent valuation framework, the kind of performance measures that have so far been unavailable to long-term investors can readily be derived and monitored over time, as new investments are made.

Future steps

Next steps include the implementation of our data collection template to create a reporting standard for long-term investors and the ongoing collection of the said data. Beyond, in future research, we propose to develop models of return correlations for unlisted infrastructure assets in order to work towards building portfolios of privately-held infrastructure equity investments. These developments will take place with the support of, and in collaboration with, the financial industry and its regulators.

References

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

Blanc-Brude, F., and M. Hasan (2015). The Valuation of Privately-Held Infrastructure Equity Investment. Meridiam/Campbell Lutyens research chair at EDHEC-Risk Institute on Infrastructure Equity Benchmarking. Singapore: EDHEC-Risk Institute, March.

---

[2] Using continuously compounded (log) returns, the discount factor is simply the exponent of minus the total return from the valuation date until the relevant period.

[3] In this example, the average price investors are willing to pay for the same infrastructure asset is assumed to increase continuously (perhaps because investors increasingly value assets that pay predictable dividends in bad states of the world) but the range of prices investors are willing to pay to buy a stake in this (unchanged) dividend process is also assumed to change. Initially it is assumed to widen (say that new investors become active in this market and have different preferences or views on risk); halfway through the 200 observed transactions, the range of valuations is assumed to start shrinking (perhaps there is now a greater consensus amongst investors about risk or more traded assets allowing replication).