Technical Guidance Note on Assessing the Welfare and Distributional Impacts of Private Sector Participation in Infrastructure Interventions

This Note introduces key evaluation approaches and methods that can produce credible evidence about the welfare and distributional impacts of infrastructure interventions. Evaluating both types of impact is crucial because the ultimate objective of public policy is to improve the well-being of the population, and particularly that of the poorest and the most vulnerable in society.

The Note focuses on infrastructure because across the developing world, inadequate infrastructure remains a major constraint to achieving the Sustainable Development Goals (SDGs) and the World Bank Group's (WBG's) goals of ending extreme poverty by 2030 and promoting shared prosperity (income growth of the bottom 40 percent in every country). To bridge the infrastructure gap, the WBG and other development actors are committed to leveraging all sources of finance, expertise, and solutions, including those of the private sector.

Expanding access to and improving the quality of infrastructure in sectors such as transport, energy, water, telecommunications, health, and education can increase economic productivity and also positively impact people's well-being. Hence, it is crucial for WBG staff to assess the impacts of infrastructure interventions across the whole range of those impacted while also taking into account all of the elements involved in financing, implementing, and operating the intervention.

The theory of change and methodologies introduced here are relevant for all projects and programs regardless of the role of the public and private sectors in financing and delivery. The design and implementation of infrastructure interventions must have a sound evidence base that ensures that the interventions address the key development constraints and achieve maximum social benefits.

The intended audience of this short note are project teams in the World Bank, International Finance Corporation (IFC), Multilateral Investment Guarantee Agency (MIGA), other development agencies, and client governments. The objective of this note (which is complemented by a more detailed technical guidance note and case studies illustrating the application of some of the tools) is twofold: to explain the importance of assessing welfare and distributional impacts of infrastructure projects and to provide project teams and other decision-makers with the information to make choices about what, when, and how to assess or evaluate.

What an Impact Evaluation of an Infrastructure Intervention Should Assess

Impact evaluations are designed to measure and value the causal effects of interventions on the outcomes of interest. Policy makers who are seeking to promote equity also wish to know how those effects are distributed across relevant socioeconomic groups. An effective evaluation of the welfare and distributional impacts of an infrastructure intervention must, therefore, provide credible evidence about the effectiveness of the intervention; and who benefits (or suffers losses) from the intervention, how and why they benefit, for how long, and at what cost. These questions are relevant throughout the policy and project development cycle: ex ante (at the design stage), during implementation, and ex post (after implementation). The issue of effectiveness relates to the attribution question: Are the intended development outcomes attributable to the intervention? Answering this question requires finding an independent variation in exposure to the intervention and then linking it to the changes observed in the outcomes of interest.

A policy intervention is a means-ends relationship wherein resources are transformed into activities designed to solve a perceived social problem. Resource scarcity means that efficiency is an important consideration in the evaluation of interventions. In fact, efficiency gains are what typically motivate the involvement of the private sector in infrastructure development. The question of whether private sector participation in infrastructure development leads to better outcomes underpins the cascade algorithm. This algorithm asks WBG staff to consider if there is a sustainable private sector solution that limits public debt and contingent liabilities. If this exists, it should be pursued; if not, policy changes, then risk mitigation, and then public funding should be considered. In making this assessment, WBG staff are advised to ensure that the costs and benefits of private versus public solutions are accurately assessed, and that equity and affordability concerns for consumers are addressed.[1]

To account for the distributional implications of private sector participation in an infrastructure intervention, the evaluation should consider how the causal relationship between the intervention and the intended outcomes might change with different sources of finance and result in different outcomes for various groups of people and locations. When calculating an infrastructure intervention's net benefits, it is necessary to include the total cost for capital, as well as long-term, recurring expenditures, since operations and maintenance costs can be even higher than the initial capital investment for investments in certain sectors (like transport).

A key consideration in assessing the impact of private sector involvement is the interdependence of funding and financing. Funding relates to the question of who ultimately pays for the full cost of infrastructure services. Financing designates who provides the upfront resources to build and start operating the infrastructure. Both funding and financing are driven by the potential for cost recovery, pricing, and other regulatory decisions, as well as the timeline of the cash flow.

Financing could be fully public, fully private, or a combination of the two. This interdependence could lead to a trade-off between financial viability and inclusion. For example, the risk of excluding poor or credit-constrained consumers by increasing user fees could limit the extent to which policy makers can pursue a project's financial viability. Since financing typically comes from the government's general budget, user fees, or taxes, the total cost can affect not only the sustainability of the infrastructure but also people's welfare as the government may have less to spend on social services such as education and healthcare.

A complete evaluation of the impacts of an infrastructure intervention, and how these are distributed across the affected population, requires assessing not only the impact of the intervention on the economy but also how it affects human welfare (income, consumption, and the non-monetary dimensions of human development such as health and education). The potential benefits from infrastructure depend on community as well as household and individual characteristics. Differences in these characteristics determine both the causal effects of infrastructure interventions and how they are distributed across the impacted population.

Developing and Using a Theory of Change for Infrastructure Interventions[2]

Assessing the welfare impacts of an infrastructure intervention and how these are distributed requires understanding the logic behind the intervention. This logic is commonly expressed in the form of a theory of change, which describes the causal relationship between the intervention and the intended outcomes. Indeed, an intervention is a means-ends relationship wherein social resources are transformed into individual and social outcomes through a set of activities that are subject to the influence of contextual factors.

A theory of change has three components: (i) a causal chain, (ii) outside conditions and influences, and (iii) key causal assumptions. The causal chain consists of causal mechanisms that convey causal influence from the intervention on the intended outcomes. Outside conditions and influences are contextual factors that are unrelated to the intervention but that may strengthen or reduce the causal relationship between the intervention and the outcomes of interest. Finally, causal assumptions state what needs to happen for the causal mechanisms to work as expected. Therefore, each such assumption reflects a risk that could prevent realization of the theory of change.

The development of a theory of change starts by identifying the motivation that underlies the decision to undertake the intervention. The core theory of change presented below is based on the following assumptions: (i) the ultimate goal of public policy is to improve the well-being of the beneficiaries, and (ii) individuals derive well-being from the best bundles of market and nonmarket goods and services they can afford given the socioeconomic constraints they face. Changes in these constraints are the core channels through which policy interventions affect well-being at the individual or household level. Policy outcomes and impacts emerge when agents interact within institutions in the roles of consumers and producers of goods and services. The model of causality underpinning these transmission channels is the standard economic model of individual behavior and social interaction.

The figure below depicts a framework for developing theories of change for infrastructure projects, and for other initiatives, based on the economic model of causality discussed above. In general, the underlying causal mechanisms involve technical relationships, institutional arrangements, and individual behavior. The bottom block of the core theory of change depicted in the figure indicates that the welfare impact varies depending on the policy-relevant characteristics of the beneficiaries. These characteristics are the key determinants of how impacts are distributed. The impact of an intervention on a household depends on the household's attributes and on the circumstances it faces.

Core Theory of Change

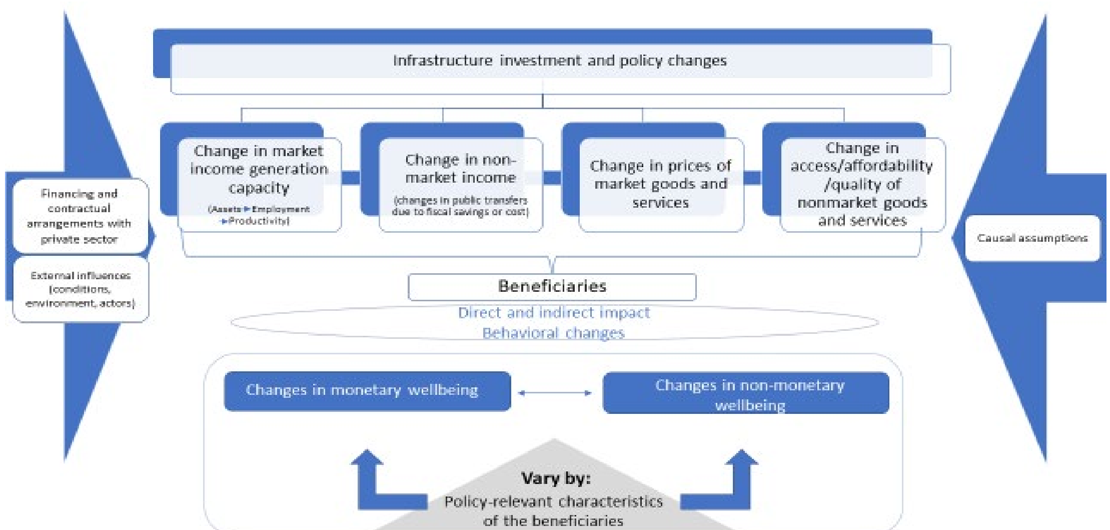

Source: World Bank Group staff

Arguably, the results chain is the simplest and clearest representation of the theory of change used by development programs. Consider the case of an investment in electricity infrastructure. The results chain might include the following:

- **Inputs**: Financing and other resources
- **Activities**: Investment in electricity generation, transmission, and distribution
- **Outputs**: Expanded electricity network
- **Outcomes**: Improved indoor air quality, educational and health outcomes, improved employment possibilities
- **Impacts**: Improved socioeconomic welfare

Examples of causal assumptions underlying this process include: (i) the legal and regulatory framework governing the energy sector is favorable; (ii) the intervention is well designed; (iii) electricity provision is properly targeted, and the associated infrastructure is well maintained; (iv) polluting sources of energy are replaced by electricity (improving indoor air quality); (v) better lighting increases children's study time (improving educational outcomes).

When making decisions about and evaluating the private sector's role in a specific intervention, it is important to view the welfare and distributional implications of this intervention in the context of the impacts of broader policies. Consider the hypothetical scenario where the private sector's role in an infrastructure project produces efficiency gains but has adverse equity effects on user tariffs that are essential for cost recovery. Complementary policies such as providing targeted subsidies for poorer users (for example, in the form of transfers or vouchers) can mitigate or reverse these adverse impacts. This can be achieved without giving up the efficiency gains from private sector participation or weakening the beneficial outcomes of the project. A rich set of World Bank materials for Poverty and Social Impact Analysis (PSIA) provides rigorous and practical approaches for evaluating ex ante the welfare and distributional impacts of policies. These materials also include many examples of how the approaches have been applied to assess the impacts of national and sectoral reforms.

While the four channels of causal influence are presented separately in the figure, it is important to keep in mind that they are interrelated and produce simultaneous effects. Computable general equilibrium (CGE) models are an interpretation of the theory of change underlying the figure above. They provide an analytical framework to handle the complex and simultaneous causal relationships that drive the impacts of policy.

IFC's development impact assessment framework, AIMM (Anticipated Impact Measurement and Monitoring), provides a way to reflect the relevant elements of the project's theory of change and the uncertainty around the realization of potential effects. Sector-specific infrastructure AIMM frameworks include a variety of components that describe different effects of a comprehensive theory of change for sub-sectors (e.g., power and transport infrastructure such as airports, ports, or roads). In addition, a likelihood assessment that addresses implementation and sector-specific, country-specific, and policy-related risks is included in the analysis to reflect the uncertainty around ex ante expectations and ultimately provide risk-adjusted development impact assessments.

The theory of change should incorporate realistic assumptions about the sectors involved. The potential for cost recovery is higher for an electricity project than for one that improves rural roads or sanitation. Institutional factors such as the extent of corruption and the public sector's capacity to enforce regulations are also important.

There are several possible uses for theories of change in policy making, and they are crucial in designing, managing, and evaluating interventions. The theory of change provides a rationale for the intervention by demonstrating how the intervention is the best solution for the specific problem, given the circumstances. During implementation, the theory of change can inform the design of monitoring systems to guide management of the intervention. Each element of the theory (inputs, activities, outputs, outcomes, and causal mechanisms) requires relevant performance indicators.

In the context of an evaluation, the theory of change can help define key evaluation questions and make plausible claims about impact. These questions concern the information that decision-makers and other stakeholders seek about an intervention's performance. They can also clarify uncertain aspects of the intervention theory.

Contribution analysis is a theory-based approach to causal inference that seeks to answer cause-and-effect questions by checking the observed outcomes against the intervention's theory of change. The notion of contribution is based on the idea that an intervention works alongside contextual factors to produce the observed outcomes. The analysis demonstrates a plausible association between the intervention and the observed outcomes, and also identifies and assesses alternative explanations for this association. Association between the intervention and the observed outcomes is considered plausible if the relevant activities have been carried out as specified in the theory of change, and the outputs and outcomes are consistent with the results expected in the theory of change. To turn this association into a causal relationship entails using evidence and logical arguments to rule out alternative explanations for the results.

Ex Ante Impact Evaluation

Once they have identified a key problem to solve, policy makers are interested in discovering the best course of action for doing so. To determine the most socially desirable intervention within a set of feasible alternatives, the following logical steps are commonly used: (i) identify and describe alternative policy options; (ii) estimate the likely consequences of each alternative; (iii) determine the value of each alternative on the basis of its consequences and a specific metric; and (iv) rank the alternatives using the values computed in (iii). The best policy option within the feasible set is the most desirable according to the chosen criterion.

Ex ante estimation of the welfare and distributional impacts of assignable infrastructure interventions is commonly undertaken with a cost-benefit analysis (CBA), which systematically identifies and evaluates the likely outcomes of alternative interventions. CBA can be viewed as comparing outcomes in two states of the world: one with the intervention and one without. This assumes that the project team has carefully considered all the relevant alternatives and then selected the one that is most likely to prevail in the absence of the intervention. For instance, if the policy objective is to achieve greater transport connectivity, alternative interventions could include upgrading the existing rail infrastructure to high-speed rail, building a new highway, increasing airport capacity, or doing nothing. If preliminary analysis of these alternatives shows that upgrading the existing rail infrastructure is the next best alternative to building a new highway, then CBA would focus on these two options.

The costs associated with an intervention are the benefits that must be given up by allocating resources to the chosen intervention. This is consistent with the idea of assessing the return on a resource engaged in a socioeconomic activity based on its opportunity cost, which indicates what the resource would have earned under the next best alternative use. Also, the only benefits that matter in CBA are the marginal (incremental) benefits. These are the benefits that would accrue beyond those that would have accrued in the counterfactual (alternative) state.

Ex ante analysis is necessarily based on models as the analysis must predict the consequences of policy. Conventional CBA tends to rely on partial equilibrium modeling in order to compute the consequences of interventions because CBA uses a simple supply-and-demand framework that represents the primary market. Although the partial equilibrium framework is valid for small-scale infrastructure, it has limitations when assessing the impacts of large-scale infrastructure interventions because the effects of these are likely to spill over beyond their primary markets. In such cases, general equilibrium modeling provides an appropriate framework since it accounts for both the direct and indirect effects of large interventions or policy changes that concern private sector participation.[3] To estimate distributional impacts, the general equilibrium model needs to be linked to a microsimulation, which can project the impact of changes in macro variables on household and individual outcomes. An example of a macro-micro integrated model to assess distributional impact ex ante is the EPIQ (Economy-wide Private Impact Quantification) model developed by IFC. Existing microsimulation tools developed by the World Bank, such as ADePT and WELCOM, can also be combined with computable general equilibrium (CGE) or financial computable general equilibrium (FCGE) models as appropriate.[4]

When appraising infrastructure investments with private sector participation, the welfare and distributional impacts must also be analyzed in relation to how they are financed. An FCGE model is an appropriate framework as it integrates the real and financial sectors of the economy and can be used to trace the flows of financial and real resources among socioeconomic agents.[5] For example, in the case of a highway expansion project, the following financing options could be available: (i) tax revenues, (ii) government bonds, and (iii) private financing. By using an FCGE model, the consequences of these three alternatives could be considered. Regardless of the exact model used, the analysis of distributional impacts must be conducted at the household level or with policy-relevant socioeconomic groups. The benefits of using a general equilibrium model, however, must be balanced against the costs and data requirements of setting up such a model. Generally speaking, a CGE or FCGE model is more appropriate when an infrastructure investment is large enough to have economy-wide impacts through multiple channels, including indirect and induced impacts.

Policy outcomes are the effects on people due to attaining the immediate objective of the intervention. For example, an intervention in the transport sector could lead to change in the following outcomes: income, travel time, fatal accidents, and air pollution—all of which are measured in different units. To assess the desirability of different policy options, their favorable and adverse consequences (benefits and costs) must be made commensurate. This, in turn, requires defining and estimating commensurate measures of value for the policy outcomes. The value of these outcomes is given by the corresponding changes in individuals' well-being.[6] The economic approach to valuation defines individual well-being as satisfying individuals' preferences. However, individuals' preferences are not directly observable. Instead, they can be inferred by the welfare values of policy outcomes—either from the real-world choices that people make in markets (the revealed preference approach) or from people's responses to hypothetical questions about their preferences (the stated preference approach).

When people choose less of one commodity and substitute more of another, they reveal something about the relative worth that those two commodities have for them. If the monetary value of one of the commodities involved in this tradeoff is available, then the monetary value of the other commodity can be inferred from the observed tradeoff ratio. The value of policy outcomes is commonly expressed in terms of either the willingness to pay (WTP) for a desirable outcome or the willingness to accept (WTA) compensation for an undesirable outcome. In other words, the amount of money a person is willing to pay for something (or to accept as compensation for going without it) reveals how much a thing is worth in terms of that person's well-being.

Public policy is made for society as a whole. Thus, the selection of the most desirable policy option entails ranking the alternative policy options based on the social value of their outcomes. Ranking policy options is commonly based on a social evaluation function known as a social welfare function. This function is based on individual valuations of policy outcomes, and it provides a rule for aggregating individual levels of well-being into an indicator of social welfare. This rule reveals the value judgments that govern social evaluation. Therefore, the most desirable policy option is the one with the highest level of social welfare when all of the feasible options have been considered.

Mainstream CBA focuses on demonstrating that the intervention will achieve economic efficiency, and little attention is paid to how the costs and benefits will be distributed across policy-relevant socioeconomic groups. Typically, the decision-making rule in conventional CBA aggregates the costs and benefits from each alternative across both individuals and time periods in order to arrive at a net present value (NPV). This NPV is then compared with the NPV of the next best alternative. This rule implicitly values only efficiency in resource allocation and not how the benefits vary across the population. In fact, such a rule, even though it appears to be neutral, has a built-in bias against low-income winners and losers. This is because in comparison with low-income people, those with higher incomes typically have higher willingness to pay for a favorable outcome and demand higher compensation when the outcome is unfavorable. This influences the economic valuation of policy outcomes, and it implies that the policy options that generate value for high-income groups will be seen as more valuable than those that generate value for low-income groups. Thus, when the benefits and costs are added up, the benefits going to wealthier individuals count more than those going to the poor.

This built-in bias can be addressed through the use of distributional weights or social weights that are assigned to the outcomes for individuals or groups in order to reflect the value that society places on such outcomes. The use of social weights requires a disaggregated analysis of the change in net benefits created by the project for different policy-relevant socioeconomic groups. Meaningful weights can be assigned based on the net benefit going to each group, and the results can then be aggregated to make a weighted sum of net benefits that is consistent with how a society values efficiency and distribution. For example, if a society decides to evaluate policy options just on the basis of how the incomes of poor individuals are improved, the social weights could be such that the net benefit of the project, and its alternatives, would be measured by the reduction in the poverty gap. The latter measures the distance of the average poor person from the poverty line, and it ignores any improvement in the incomes of those above the poverty line.

Ex Post Impact Evaluation

Choosing a policy on the basis of rigorous ex ante analysis does not guarantee that it will effectively address the target problem. Predictions made by using ex ante analysis are uncertain due to unforeseen events, unaccounted-for consequences, and misunderstood causal mechanisms. It is therefore important to evaluate policy interventions ex post as well as ex ante. The difference between ex ante and ex post analysis is not a matter of methodology but rather a matter of timing and purpose. Ex ante analysis is a prospective assessment of how the intervention will address the targeted problem under specific assumptions based on an underlying theory of change, whereas ex post analysis is a retrospective assessment of the performance and results of the intervention. This means that CBA can also be conducted ex post, and when it is, CBA seeks to determine whether the intervention resulted in an efficient allocation of resources.

Ex post evaluations help validate the predictions of ex ante assessments that have been conducted for the same project. They also provide critical information for ex ante assessments (including CBAs) of future (and comparable) projects by validating potential theories of change and assumptions, and generating parameter estimates that can be fed into ex ante models. Although an ex post evaluation can be completed only after the project is complete and has had time to generate impacts, it requires upfront planning at the inception stage of the project to implement an evaluation strategy at different points of the project's life cycle. The strategy might include a series of data collection efforts, with the frequency and timing of data collection depending on the nature of the project and the type of evaluation strategy adopted.

Ex post impact evaluation focuses on the consequences of interventions and therefore must credibly deal with the attribution problem. The strategies available for coping with the attribution problem depend on whether the intervention is assignable or not. An intervention is assignable when it can be allocated to some observational units (such as households, firms, or communities) and not to others. There are three broad categories of infrastructure interventions: (i) small-scale infrastructure interventions (e.g., rural electrification, rural roads improvement, and certain urban transport schemes); (ii) large-scale infrastructure interventions (e.g., port development, power distribution networks, and transnational railways); and (iii) policy interventions (e.g., adoption of a public-private partnership model). Whereas small-scale interventions tend to be assignable, typically large-scale ones are not.

The main challenge in evaluating the consequences of a policy intervention is determining the counterfactual (i.e., what would have happened in the absence of the intervention). In the case of assignable interventions, to infer the causal effect of the intervention, it is possible to create a counterfactual using non-assigned units and perform counterfactual comparisons. Such comparisons are commonly based on two single-difference comparisons: (i) with and without the intervention, and (ii) before and after the intervention. The validity of causal conclusions based on with-and-without comparisons is threatened by selection bias (i.e., when, on average, the treated and untreated units do not have the same characteristics). Selection bias is more likely when exposure to the intervention is by choice (self-selection or administrative selection) rather than when it is random. One solution to this—random assignment, which ensures that the distribution of both observed and unobserved characteristics prior to the intervention is the same for both the treatment and control groups—is difficult to implement for most infrastructure investments, even if they are small scale and assignable.

It may seem that the selection-bias problem can be avoided by measuring the outcome for the same group of treated units before and after the treatment. However, one cannot conclusively claim that the observed change in the outcome revealed by this comparison is attributable only to the intervention. External events unrelated to the intervention may also affect the outcome. Thus before-and-after comparisons are vulnerable to history. In fact, when the full impacts take a long time to emerge, as is the case with some infrastructure interventions, even the with-and-without comparisons for similar groups are vulnerable to history. This is particularly the case if the outcome of interest is the improvement of living standards. A time frame of several years creates opportunities for shocks and spillover effects from other investment programs to affect the comparison and treatment areas. In other words, history threatens the validity of any causal inference that ignores these confounders.

Double-difference comparisons that involve both before-and-after and with-and-without comparisons are the most effective way to counter the validity threats that result from both selection bias and history. This requires finding a comparison group that is not exposed to the treatment, but is similar to the treatment group in all other respects, and then measuring the relevant outcomes for both groups, before and after treatment. If both groups are indeed subject to the same external events, then whatever happens to one group also happens to the other, except for the intervention. Under these circumstances, the comparison group acts as a recorder of history. Hence, any difference in outcomes between the two groups must be due to the intervention.

The methods most commonly used to evaluate the impact of assignable interventions belong to the design-based approach to causal inference. To identify and estimate causal effects, this approach includes both experimental and nonexperimental methods. But even in the case of small-scale infrastructure, the opportunity to undertake random assignment to apply experimental methods is limited. Policy makers tend to assign small-scale infrastructure on the basis of local conditions. This may lead to endogeneity (selection bias), given that such conditions are also likely to affect the potential benefits, which tend to be derived and conditional. Furthermore, the benefits may take a long time to materialize. These issues have made the Difference-in-Differences (DD)[7] method one of the most effective tools for evaluating assigned infrastructure interventions. This method uses panel data that is collected at different points in the intervention cycle, including before and after the intervention. Depending on the type of endogeneity, in order to obtain unbiased results, the DD method may have to be combined with other methods (e.g. Propensity Score Matching[8] or the Instrumental Variable method[9]). A Regression Discontinuity Design method[10] is useful when there is a discontinuity in the assignment of an intervention.

When an intervention is non-assignable, causal relationships cannot be inferred from observed patterns in the data since a counterfactual cannot be created from the non-exposed units. Such situations demand theory-based approaches to evaluation. These are guided by an explicit theory of change about how the intervention causes the intended outcomes. On the basis of a plausible theory of change, the impact of the intervention is assessed in a way that accounts for both the underlying causal mechanisms and the implementation processes. As discussed earlier, such an evaluation can be couched in terms of contribution analysis. The analysis of different components of the causal chain, and of the causal assumptions, necessarily involves the use of mixed methods (qualitative and quantitative). Qualitative information may shed some light on the factors that drive infrastructure placement and selection bias.

Conducting a quality impact evaluation requires analyzing the theory of change underlying the intervention. In addition, the analysis of different components of the causal chain and of the causal assumptions necessarily involves the use of mixed methods. In this context, qualitative information may shed some light on the factors that drive infrastructure placement and selection bias. Alternatively, the theory of change can be expressed with a structural model (e.g., a CGE model—as discussed earlier) of individual behavior and social interaction. A CGE model can simulate the counterfactual for infrastructure interventions that have geographically dispersed effects on a variety of variables—for example, a trunk road, a national railroad, or a change in regulations that leads to economy-wide impacts. This approach will lead to the quantification of both direct and indirect effects. With respect to private sector participation, counterfactuals may be generated by considering different sources of financing and different contractual arrangements for service delivery within an FCGE model. A CGE or FCGE model can be used for both ex ante and ex post analysis. The crucial difference is that ex post analysis will have the benefit of considering observed changes in variables to calibrate the model and validate its results, whereas ex ante analysis can only do so with historical data. Typically, the ex post evaluation is an update of the CGE or FCGE simulation conducted ex ante, based on new data collected after the project is completed.

Most evaluations of the effects of policy and infrastructure interventions focus on estimating mean impacts such as the average treatment effect, or the average treatment effect on those treated by the intervention. While mean impacts may answer the question of whether or not an intervention worked, they do not provide evidence about who benefited from the intervention and why they benefited. Infrastructure interventions (and development interventions in general) are expected to have differential effects on individuals or households, depending on their attributes and the circumstances they face. This heterogeneity is what drives the distributional impact of such interventions.

Analyzing the distributional impacts of an intervention requires considering those impacts across policy-relevant socioeconomic groups, households, or individuals. A regression model that includes interaction between the treatment and household or individual characteristics can provide a framework for analyzing systematic variation in mean impacts across socioeconomic groups. For example, such a model could be used to assess the distributional impact of expanding irrigation infrastructure. Other approaches to assess the distributional impact of interventions include analyzing quantile treatment effects that capture responses to treatment across the entire distribution of welfare, and general equilibrium modeling. As is the case with ex ante assessment, an ex post CGE model can be combined with microsimulation to assess distributional impacts.

Implications for Data

A credible evaluation of the welfare and distributional impacts of an infrastructure intervention requires both a valid strategy for impact identification and measurement, along with relevant and reliable data. As noted earlier, impact evaluation of assignable interventions has relied more often than not on the DD method or a combination of DD with either the PSM or IV method. Effective implementation of these methods determines the nature of the needed data and the frequency of data collection.

Assessing the impacts of assignable infrastructure interventions entails panel data measuring outcomes, including pre-intervention baseline data. Information is needed on both the assigned and non-assigned units in order to create an appropriate counterfactual. In particular, the assessment requires data on observables when combining a DD with a PSM method. Such data include baseline characteristics that may affect selection for the intervention. Additional data are required to control for changing circumstances over time, to explore variation of impacts across policy-relevant socioeconomic groups, and to account for exogenous shocks such as natural disasters or weather-related events. With respect to the frequency of data collection, for impacts to emerge there should be a sufficiently extensive time interval between the baseline and follow-up surveys.

In the case of non-assignable interventions, a Social Accounting Matrix (SAM) or a Financial Social Accounting Matrix (FSAM) should be used with data that are disaggregated according to the sectors and institutions that are relevant for analyzing the policy issue concerned. In conducting distributional impact analysis, the SAM (or FSAM) must be consistent with the available household survey data.

New technologies can significantly expand data availability and reduce data costs. For example, detailed mapping of infrastructure quality and availability across geographic locations can greatly help ex ante and ex post evaluations. Spatially disaggregated data derived from leveraging new technologies—such as highresolution satellite images processed with machine learning algorithms—can fill critical data gaps in infrastructure and reduce the cost of data collection.

Summary

To sum up, assessing the welfare and distributional impacts of infrastructure interventions with private sector participation, whether ex ante or ex post, does not require the invention of new approaches or methods; rather it calls for the judicious application of well-known methods of impact evaluation to measure disaggregated impacts across different socioeconomic groups. In particular, analysts must widen the scope of the evaluation to consider the welfare and distributional implications of the actual or potential mode of financing and the contractual arrangements for service delivery. All these elements, along with the characteristics of the infrastructure itself, influence the casual relationship between the intervention and well-being.

In conclusion, when making decisions about the extent and modalities of private sector participation in infrastructure vis-à-vis the public sector, it is important to consider the tradeoffs between financial viability and equity, or between efficiency and equity. When conducted with a distributional lens, ex ante impact analysis can inform these decisions with evidence about the likely trade-offs associated with different modes of funding and financing. Ex post evaluations have a key role to play as well, by measuring how much of the predicted impacts and trade-offs were realized and enabling actual cost-benefit calculations that can in turn inform the planning and design (and ex ante analysis) of future projects and policies.

Acknowledgement

This technical guidance note was prepared by a World Bank team, consisting of Ambar Narayan, Clive Harris, Esther Illouz, Luciana Silva, Essama Nssah (Consultant), Cynthia Gears (Consultant), Afua Entsuah (Consultant), Ann Bishop (Consultant), with the assistance of Shuai Alexander You (Intern). The work was undertaken under the guidance of Carolina Sanchez-Paramo, Global Director, Poverty and Equity Global Practice; Imad Fakhoury, Director, Infrastructure Finance, PPPs & Guarantees; Fatouma Toure Ibrahima, Practice Manager, PPP Group; and Benu Bidani, Practice Manager, Poverty & Equity Global Unit. Guidance was also provided by Abha Joshi-Ghani, Jordan Schwartz, and Mamta Murthi during the inception and concept stages. Vivien Foster, Chief Economist Infrastructure, provided valuable advice to the team during the early stage of the work and co-chaired the Concept Note Review. The team is grateful for inputs and comments provided by Issa Faye and Camilo Mondragon Velez of IFC and Moritz Nikolaus Nebe, Jessica Charles Wade and Mena Cammett of MIGA, and to peer reviewers – Andrew Dabalen, Arnaud Dornel, Fan Zhang, Mona Hadad, Richard Damania , Richard Spencer, Stefan Apfalter, and Vincent Palmade – for their insightful comments at different stages of the review process. The team would also like to thank Jemima T. Sy, Program Manager of PPIAF that partly financing the project, and Svitlana Orekhova, Program Officer, PPIAF for coordinating the Trust Fund activities.

---

[1] "Maximizing Finance for Development: Leveraging the Private Sector for Growth and Sustainable Development" prepared by the World Bank Group for the October 14, 2017, Development Committee Meeting.

[2] For a comprehensive discussion of the theory change framework, see the full Technical Guidance Note

[3] A general equilibrium model is a logical representation of a socioeconomic system wherein the behavior of all participants is compatible. A socioeconomic system is a set of market and nonmarket institutions that govern social coordination to cope with the scarcity of resources. General equilibrium modeling is the workhorse for the quantification of the consequences of shocks and policy interventions that are likely to affect the whole economy. The standard economic model of individual behavior and social interaction serves as a template for most applied general equilibrium models. General equilibrium is achieved by the configuration of relative prices such that, for each market, demand is equal to supply. Policy analysis requires a computable or applied model that usually takes the form of a system of equations describing the supply and demand sides of the economy, along with budget constraints and

equilibrium conditions. The necessary data for a computable general equilibrium (CGE) model are usually organized in a social accounting matrix (SAM) that reflects the circular flow of economic activity over the chosen time period.

[4] Lokshin et al. (2013); Araar et al (2018); and IFC (2018). IFC and the World Bank are also currently working together to build an integrated macro-micro tool to assess the distributional impacts of private sector investments including in infrastructure.

[5] With an FCGE model, real transactions cover supply and demand interactions across commodity and factor markets, while financial transactions related to the operation of the loanable fund market reflect choices made by agents about the composition of their portfolios. The empirical implementation of an FCGE model requires a data set organized within a financial social accounting matrix (FSAM).

[6] Well-being is the source of value because the ultimate goal of public policy is to improve the well-being of the target population.

[7] The difference-in-differences (DD) method compares the change over time in the outcomes of the treated units with those of the untreated units. In particular, the intervention's impact is calculated by subtracting the change over time in the outcome for the comparison group from the change over time in the outcome for the treated group.

[8] Propensity score matching (PSM) compares treated units with untreated units with the same conditional probability of participation. The validity of this method hinges on the assumption that matching units on their propensity scores makes them also comparable along unobservable dimensions. This may be too strong an assumption, however, given that it is virtually impossible to rule out whether there are unobserved characteristics that differ between the treated and the matched untreated units.

[9] The instrumental variable (IV) method is commonly used when selection bias stems from endogenous placement or endogenous participation. This method can generate independent variation in treatment to the extent that it has a direct effect on the choice of the treatment, but it affects the outcome of interest only through its direct effect on treatment. The second condition implies that there is no direct relationship between the instrument(s) and the outcome of interest, which is known as the exclusion restriction. Basically, IVs: influence the likelihood that an individual will participate in the intervention, are independent of individual characteristics, and are not under the control of the individual.

[10] Regression discontinuity design (RDD) exploits discontinuity in the assignment mechanism with respect to an assignment variable, also known as a forcing or running variable. Usually discontinuities stem from rules that govern eligibility for the intervention. In particular, eligibility is determined on the basis of a threshold or cut-off point in the range of the assignment variation. Units on one side of the cut-off point are eligible, while those located on the other side are excluded from the intervention.