Comment

Francisco H. G. Ferreira: When major infrastructure investments are large enough to have general equilibrium effects, how can their impacts on poverty or the distribution of incomes be assessed? What about the effects of substantial increases in a country's ability to extract rents from its natural resources? Is it possible that such gains might worsen the welfare of the very poor, even while contributing to increases in aggregate national income? Bussolo, de Hoyos, and Medvedev investigate the distributional effects of the Panama Canal expansion—a large infrastructure project that enhances Panama's ability to extract rents from its geographical position—to shed new light on these questions.

The expansion of the Panama Canal represents a massive infrastructure investment, estimated to cost approximately 40 percent of Panama's 2003 GDP, over a seven-year horizon. Such a large investment could arguably be expected to have multiplier effects on output and employment across the entire country. In addition, both by attracting large inflows of dollar-denominated loans during the construction phase and by raising canal revenues during the subsequent operation phase, it might affect the country's real exchange rate and, hence, relative domestic prices. Both of these channels could mediate the effects of the canal expansion on poverty and inequality, through the general equilibrium of the economy.

The existence of such general equilibrium effects poses serious problems for most approaches normally used to evaluate the distributional impact of infrastructure interventions. One might, a priori, consider estimating the impact of the canal expansion by comparing changes in outcomes for households located near the Canal Zone with those living farther away, using propensity score matching techniques analogous to those used by Escobal and Ponce to evaluate the effects of rehabilitating rural roads in Peru.[1] Alternatively, if a suitable instrument were found, one might adopt

1. Escobal and Ponce (2002).

an instrumental variables (IV) strategy such as that employed by Duflo and Pande to evaluate the impact of irrigation dams on district-level poverty in India.[2] However, if relative prices are changing across the entire economy, and if multiplier effects are affecting labor demand even in far-flung areas of the country, then instruments are unlikely to be valid, and comparison groups are almost certainly contaminated.

The strategy adopted in this paper is to tackle the general equilibrium effects head-on, by means of a macro-micro model that combines a computable general equilibrium (CGE) model with a household-survey-based microsimulation module. Additional expenditures and borrowing associated with ex ante plans for the canal expansion are fed into a CGE built for Panama and are allowed to affect investment, output, labor demand, and so on, relative to a baseline ("business-as-usual") scenario. The model ultimately predicts certain changes in equilibrium prices, wages, and labor allocation, which are fed down to earnings and occupational choice equations estimated on household survey data. Those equations allow the authors to simulate changes in wages and employment patterns at the level of individuals and households, in a manner consistent both with the preexisting conditional distributions observed in the microeconomic data and with aggregate changes predicted by the general equilibrium model.

My main methodological comment is that the failure to contrast model predictions with ex post historical data represents something of a missed opportunity for this paper. As the authors acknowledge, and as appendix D makes abundantly clear, the use of a CGE combined with a microsimulation model necessarily implies reliance on a large number of assumptions. "Identification of impact," even with all the caveats about scenarios, rather than forecasts, is obviously conditional on all of these assumptions. The justification for embarking on such an exercise is, as noted above, that the intervention in question will plausibly have significant general equilibrium effects, so that alternative techniques would lead to confounded estimates. Even so, such an assumption-heavy technique calls for as much validation by real data as possible.

The paper relies on a social accounting matrix (which underpins the CGE) and a household survey, both from 2003. It also uses population and GDP growth forecasts taken at some (not precisely specified) point in the last decade. When simulating the canal expansion scenario, for comparison with the business-as-usual benchmark, it relies on an ex ante plan of expenditures for the 2007–14 period. We are now in 2012, and much of the relevant

2. Duflo and Pande (2007).

macroeconomic data exist for Panama until 2011. Household survey data exist at least for 2008 (in the form of the *Encuesta de Niveles de Vida—ENV 2008*) and possibly for later years. It is difficult to imagine that some of these data could not have been used to shed light on the performance of the model for the business-as-usual scenario between 2003 and 2008 and, for the macroeconomic module, until more recently. If the baseline scenario had been simulated for 2003–11, for example, and compared with real data, one would certainly learn a lot about how reliable the 2003–14 estimates are likely to be. To be clear: the suggestion would not have been to recalibrate the model using more recent time-series data, but to assess the performance of the dynamic computable general equilibrium model, calibrated on a previous period, on a fresh time sample not used for calibration. Similarly, the ENV 2008 household data might have been used to compare actual poverty and inequality statistics for that year with a simulation run for 2003–08. The performance of such an exercise would be enormously informative for how confident one might be of the simulation results presented for 2014 and 2020. In general terms, if the growing macro-micro simulation literature wants to be taken seriously, despite its heavy reliance on all sorts of assumptions, it has to face the data.[3]

Abstracting for a moment from the above caveat, what can policymakers from Panama learn from this paper? First, despite a modest contribution to faster economic growth (mostly in the operation phase, after 2014), there is a real risk that the canal expansion may have a regressive effect on the income distribution. The model-predicted effect on poverty is small, both in the construction phase (2003–14) and in the operational phase (2014–20). In the latter period, a small decline in the poverty headcount coexists with an increase in the depth of poverty, driven by increasing inequality. Welfare losses are concentrated among the rural poor who do not migrate from rural to urban areas or, more accurately, who do not diversify away from agriculture into nonagricultural activities. The losses are driven by rising domestic prices, rather than by falling nominal incomes. Most of the income gains from the canal expansion are concentrated in the top half of the income distribution, among those in the formal nonagricultural sector, where both employment and wages expand. Even though most of the direct increases in labor demand

3. Ferreira and others (2008) is one of the few papers that compares simulation results from a macro-micro model with actual ex post data. In this study of an exchange rate devaluation in Brazil, model performance was found to be uneven (p. 160).

are in the construction sector, canal construction is relatively skill intensive, and the skill premium is simulated to rise as a result of the expansion.

From the viewpoint of poverty reduction, these results suggest little hope that the benefits of the large investments and rising exports associated with the Panama Canal expansion will trickle down to the poorer segments of society. In line with the common characterization of the canal as an enclave, with relatively few employment linkages with the rest of the economy, it seems that whatever multiplier effects the expansion might engender are likely to be confined to the higher-skill segments of the labor force. In addition, because of the effects on exchange rate and price levels, the expansion might well end up raising the cost of living of those poor people who do not benefit from it in any way. Although the margins of error around these scenarios are likely to be large, given all of the assumptions that underpin them, the broad contours that arise from the exercise seem plausible enough. It is then difficult to disagree with the authors that if the Panamanian government is at all concerned with poverty reduction, it ought to seriously consider setting aside some, if not all, of the extra public revenues likely to be generated by the expanded canal for targeted transfers to its poorest citizens. Even if perfect targeting is difficult, it seems clear that some of the most vulnerable, including indigenous people living in rural areas, would not see any benefit from the canal expansion unless it were through this channel of public redistribution.

That is an important lesson from thinking about the general equilibrium effects of this large infrastructure investment and its distributional consequences. If the government takes heed and plans the redistribution in advance, the paper will already have made a contribution. Greater confidence on specific results, however, would require validation of some of the simulation exercises against ex post data.