

Chief Economist's Response: An Evaluation of World Bank Research, 1998—2005.

The *Evaluation of World Bank Research: 1998-2005* was initiated by the Chief Economist and agreed with the Senior Management of the Bank and the Board. This response is that of the Chief Economist, and takes into account the dual responsibilities of his office, to head DEC and to provide broad oversight over the Bank's overall research program.

We are grateful to the Panel and the Evaluators for their work in producing this evaluation, work which often went beyond the call of duty. It is deep and thoughtful, and draws insights from among the best research expertise available. The report will be extremely valuable as an input to DEC management, and also of utmost interest for managers outside DEC who are responsible for research done by or commissioned by their respective units.

We are encouraged by many points made in the report, including its confirmation that, as suggested by the record of publication, much of the Bank's research is of leading quality, its support for strong in-house research capacity, and its conclusion that "Bank researchers have delivered value to the institution that is much larger than could be reasonably expected from the very small share of the budget that they command." Particularly in the current environment of contracting research budgets this is a valuable endorsement.

The report makes a number of detailed criticisms and draws attention to aspects of Bank research and research management that the panel believes can be improved. At the same time, it demonstrates explicitly and, even more frequently, implicitly that research in the Bank requires trade-offs over a complex set of objectives and constraints. It must be relevant and useful to Bank operations and to policymakers world-wide, as well as being academically sound and stimulating further advance. Managers will find the Panel's and the Evaluators' opinions extremely valuable in navigating these trade-offs and in weighing the many other opinions that they receive from colleagues in operations, practitioners, journal editors, peer researchers, etc.

The purpose of this response is to help us chart the road ahead. Our intention is to set out how we see the issues addressed in the report in the context of the Bank, to indicate where we broadly agree with the report's recommendations and observations on Bank research, and where we may have reservations. It is not to set out the details of how we plan to respond to every aspect of the report. Some require considerable thought and consultation, not only among research managers in DEC but also with staff in other parts of the Bank and with the Bank's top management.

Strengthening the Quality of Research and Research Management

We support the objectives set out in the Panel’s recommendations—to ensure the integrity and credibility of Bank research, to ensure its independence from the Bank’s immediate interests and policy advocacy, and to continue to contribute robustly to development knowledge. We are encouraged that the Evaluators have scored 61% of the studies they read to be of above average or superior quality with respect to contributing to increased knowledge, and a further 28% as average, but we recognize the need to maintain effective checks and balances in order to achieve the objective of maintaining a high-quality Bank research program.

While noting the academic excellence and policy-relevance of much of the reviewed research, the Panel also argues that a significant proportion of Bank research is “undistinguished” and makes a number of suggestions for how to reduce this proportion. We appreciate the Panel’s candor, and also the confirmation it gives to the difficulty of assessing the quality of research objectively. For example, sometimes the opinions of the Panel differ from those of the Evaluators; they also note that some of the research most critiqued by the report has appeared in leading academic journals. We find no objective evidence that the Bank’s research “tail” is longer than that of other leading institutions. Also, work that can appear undistinguished from an academic perspective can have important operational impact. Nonetheless, we share the Panel’s belief that we should continuously try to strengthen the quality of research – and will continue to explore ways to do so.

The report advocates stronger managerial oversight to strengthen the average quality of research. Within DEC, and also considering research funded by the Research Support Budget (RSB), strengthened oversight could include more frequent reviews of large projects while under implementation and of the contribution of programs after completion. We will explore how these approaches might be operationalized. Since a substantial volume of research is non-DEC, strengthening research management also calls for consultation with managers elsewhere in the Bank on how to strengthen the Chief Economist’s advisory role and oversight of Bankwide research more broadly, including that part of research that does not go through the Research Committee and the RSB. Both suggestions have budgetary implications and the latter would require careful delineation of the (sometimes fine) dividing line between research and other analytical work (mainly Economic and Sector Work), which is often more closely directed towards immediate regional and operational concerns and therefore may have different objectives.

Related to this, the report recommends more careful and central control over the quality of consultants commissioned to undertake research. It is hard to respond concretely without more specifics on the basis for this recommendation. In practice, various factors shape the consultant selection process. The ideal consultant is frequently not available on short notice; in some cases, the Bank’s pay scales make it difficult to attract those consultants in very high demand; capacity-building, an objective endorsed by the Panel, also sometimes argues for a trade-off between research excellence/experience and other

factors. Where capacity-building is an important objective, it should be clarified in proposals.

The Panel notes the importance of maintaining good data on research costs and outputs. This is indeed essential, but even with accurate attribution of costs, ascertaining “value for money” is difficult with research, particularly in an operational context where standard tools tell only part of the story. Many projects will yield little payoff but (as noted in the report) a few will have enormous impact several years after their completion and publication, and it is not always easy to predict which these will be. We agree in particular with the importance of ensuring accurate reporting of outputs, and will work to ensure that the bibliography of Bank research products (which has improved substantially as a result of the efforts undertaken to support the evaluation) is maintained and is also made accessible to improve outreach. DECRG has recently revived its Annual Report listing all outputs.

The Panel suggests that the volume of Bank Flagships should be reduced in order to better ensure their quality. We agree with the need to revisit the objectives and scope of Bank flagship reports and, drawing on recently-commissioned evaluations of other parts of DEC, are currently exploring approaches towards reducing the volume of DEC reports where possible. We should clarify, however, that many of the Flagships reviewed by this research evaluation (and in general, assessed as being of good quality) may not be formally designated as research products, and that they are often intended to target specific Bank clients or broader interests in addition to their specific research content.

Some questions are raised in the report on the organization of research and other units within the Bank. These issues were not part of the mandate of the evaluation, in part because the Bank’s organizational structure responds to many demands, and not primarily to the needs of research.

Restructuring the Balance of Research

The Panel stresses the value of case studies and analytical narratives that can draw on Bank experience. We concur. This is entirely in line with our philosophy that good research starts not with a preferred method, but with an important question, and then looks for appropriate and feasible methods for answering it. We respect the Panel’s opinion that cross-country regressions, CGE models and matching methods for impact evaluation may not always be the most appropriate technique, and that randomized evaluations may be appropriate more often than is sometimes recognized. On the other hand, there is no general agreement within the scientific community on such methodological preferences so we would not wish to be constraining in this regard. The important issue is that all empirical work should be done well – for example, as the Panel demands, that claims to identification and causation are plausible and duly supported - whatever the particular empirical approach selected. We are already seeking to greatly increase the coordinated use of impact evaluations in Bank and other development operations, via the DIME program.

The Panel suggests that increasing emphasis should be given to encouraging proposals in areas where research is weak or there is no ongoing work. The recent Bank-wide research consultation carried out under the auspices of the Research Committee was designed partly to identify such lacunae, and has led to the commissioning of several forward-looking reviews intended to encourage further research in some of these understudied areas. We will consult further within DEC and elsewhere in the Bank to see whether additional steps can be helpful.

We agree with the Panel that the economics profession gives too little weight to replication and would certainly encourage it as part of a sound applied and policy research program. Indeed, some replication goes on already among researchers, often leading to publications that might be judged “undistinguished” or to no publication at all. The Networks, which are charged with developing and disseminating “best practice” advice, also often undertake such tasks in their respective areas. Establishing a dedicated replication unit within DEC does not seem optimal: it would be hard to staff with strong researchers, and would not have expertise across the range of topics that DEC covers. Rather, we will respond to the Panel’s advice by encouraging staff to undertake replications in their own areas where these seem to be useful. The DIME initiative, for example, encourages independent yet coordinated evaluations of sets of similar interventions to feed into meta-evaluations.

Strengthening the Bank’s Data Function

The Panel strongly endorses the Bank’s already central role in development data, and makes a number of valuable suggestions for further strengthening and centralizing this role, including through the possible creation of a Chief Statistician position. Timely, accurate and consistent data are at the core of the results agenda, and we support the thrust of the recommendations in this area. Implementing the ambitious agenda suggested by the report would require senior management to endorse the Bank’s central data role, and channel substantial financial and human resources into the area over the long term. It would require some changes in our business model, both internally—e.g., how a Chief Statistician could vet *all* data collection without stifling research—and in terms of how the Bank works with clients, for example, on access to information (such as household surveys) and broader disclosure policy. These options will need to be considered together with those of the recent review of the data group of DEC, which also focuses on data issues, but more from the country operations side than from a research perspective, and they also need to be discussed more widely in the Bank.

In response to the call for the need to develop standard protocols on release and revision of data, a policy for access to research data was recently established for DECRG. It provides for access within six months of the first working paper using the data-set or within two years of collection. We are mindful of the constraints that some countries impose on data jointly owned with them, but we continue to encourage more open access. We are mindful also of the need to develop policies for ethical review of surveys where needed and for better quality control of the shared data, and expect to move forward soon on a pilot case in this area.

Improving Outreach and Relationships with Academia

The Panel makes a number of recommendations to improve outreach and relationships with academia, including improving and revitalizing the visitor program and strengthening the research function of the ABCDE. The evaluation also pointed to the need to improve the accessibility of the DEC website, and to increase efforts to foster collaborative work with developing country researchers.

These are useful recommendations, which in some cases require more study to develop specific plans. We will take them up in further management discussions, and with the Research Committee.

The Panel also urges the Bank to help build research capacity by working more extensively with researchers from developing countries. We agree with this objective and will increase our efforts, but bearing in mind the difficulties sometimes experienced in finding suitable collaborators in countries with low research capacity. We note that the RSB has provided research capacity-building support to a number of institutions for several years, and has also supported capacity-building through the Global Development Network.

Responses to Observations on the Bank's Research

The Evaluators have produced fairly detailed reports on some two hundred research products. We are grateful for their diligence and insight, and while we do not agree with every one of their comments and criticisms we recognize their reports as a very valuable guide for improving ongoing and future research.

The Panel has drawn on the Evaluators' reports in preparing their overall evaluation, but has been frank in admitting that their opinions do not always reflect those reports. It has made a number of specific criticisms in passing, some of which we accept, some of which we do not, and it has also highlighted three major concerns. While many of the issues involved are complex and technical and cannot be entered into here, we very briefly note our reactions to the tenor of these critiques, starting with those of a more methodological nature.

First, the Panel argues that certain techniques have been over-used. As noted above, we should seek techniques after we identify questions, and we also recognize that tastes differ. However, we believe that the Panel's impressions that research on aid effectiveness, finance and growth and trade relied almost exclusively on cross-country regressions is not fully consistent with a broader reading. The first two programs drew extensively on other sources—e.g., project-level data, sectoral data and country case studies. Similarly, of the total output of the DECRG trade team between 2000 and 2006, less than 10 percent employed CGE methodologies, suggesting that concerns about the balance of the trade research portfolio may be overstated. The Panel's perception may reflect the use that has been made of CGE analysis to assess potential impacts of the

Doha Round in flagship reports and speeches by senior management, a process in which researchers' reservations about method were a brake on undue enthusiasm over the benefits of more open trade, as expressed by others.

Of course both cross-country analysis and CGE modeling must be appropriately explained and caveated, but both play an indispensable role in the research that the Bank is called upon to do. For many questions—including many in the areas of macroeconomics and growth—*countries* are the only choice of unit of observation that matches the key independent variable of interest, namely national policies or institutions, and often the lack of temporal variation makes cross-country analysis essential. Indeed, some of the best-known papers of some of the evaluators rely on cross-country regressions, and the conclusions of these studies, too, are not without controversy. Similarly, much of the trade analysis requested by countries ('relevant' research in the Panel's terms) is of future trade policy options. Careful data collection and statistical analysis to derive parameters for CGEs is good practice, but there is no effective alternative to modeling to answer this sort of *ex ante* question. Much of this kind of work is done with local researchers to help local policy decision-making.

Second, the report suggests that in some areas the methodology used in research has been insufficiently tested or that researchers have not sufficiently alerted users to possible limitations of the methodology. The main area here is that of poverty mapping. The main question for poverty mapping is whether it can be used to improve targeting in poor countries unable to devote the statistical resources available to map poverty through more complete census data, as is possible in the USA. Such countries have no need to use poverty mapping because they have full information from census sources already. Bank support for such techniques in developing countries does not represent promotion of "rejected" methods, rather that high-income countries have better alternatives. The arguments on poverty mapping are technical and subject to ongoing debate, including with Panel members. However, we would note procedurally that the work was extensively discussed before publication and carefully refereed by the world's leading econometrics outlet. We agree with the Panel that the important issue is to test the reliability of poverty mapping empirically. A recent paper on Mexico has done this and found the approach satisfactory, and we expect further evidence from other studies to confirm that the approach is useful. We welcome further debate on this important subject through normal scientific channels. A technical response is posted at <http://econ.worldbank.org/programs/poverty>.

Third, the report argues that in some cases, research has been used to underpin Bank advocacy, without taking care to ensure that the empirical results were sufficiently robust. Highlighted here are the two areas of aid effectiveness and pension reform. The report also notes the use of research on the long-run impact of growth on poverty.

On aid effectiveness, as the Panel agrees, the proposition that aid is more productive under conditions of sound economic management is intuitively appealing. In *Assessing Aid* the proposition was supported by project-level data and case studies, as well as by

cross-country regressions (which were given four paragraphs of coverage in a 100 page report). It should also be noted that the findings reported in Assessing Aid did not provide “justification for Bank policy” but in fact were critical of Bank policy of the time. We feel that this research program stimulated a useful debate within academia and development agencies that goes on to this day.

Pension reform has been a contentiously debated area, both inside and outside of the Bank. Many studies – by the Bank and other researchers – in the late 1980s and early 1990s emphasized – perhaps excessively – the promise of funded systems, while not paying sufficient attention to the problems of poverty alleviation for old age. This was for several reasons, including the experience of the transition countries, and the severity of the problems faced by many established defined-benefit programs combined with the political difficulties experienced by many countries in reforming them. It should also be noted that when this work was initiated, there was virtually no research on old age security in developing countries to provide hindsight on the implications of different approaches. Subsequent research, as well as the operational experience of the Bank, has become more balanced, as evidenced by more recent publications.

The critique of "Growth is Good for the Poor" draws attention to the possible impact of measurement error in shaping its result (a possibility also noted in that study) and also that this work has featured more strongly in Bank advocacy than other Bank research that reached different conclusions on the relationship between distribution and growth. Like the Panel we believe that more needs to be done to resolve the growth-poverty issue definitively, but we also believe that the Bank’s research community has not been biased in reporting its various results. These have also drawn attention to the role of growth in poverty alleviation.

We take from these points in the Evaluation the need to ensure balanced research review especially in contentious areas which are important for policy and especially where managers’ own research is involved, and also where research is brought together to provide a more institutional view on an important issue, as in a Policy Research Report. Internal and external debate is key to this process, and could be led by the Chief Economist.

Conclusion

As noted previously, this response is only the start of the process of internalizing the very useful insights gained from the evaluation, a process that, we are sure, will not be confined to discussions within Bank Management but which will involve debates among researchers, both inside and outside the Bank. Such debates, including on methodology and on how research can best be applied to help approach development problems, is an essential part of the process of knowledge creation. We are deeply indebted to the Panel and the Evaluators and thank them again for their contribution to this process.

wb18913
L:\DECRS\Evaluation of Research\Final Version\Chief Economist Response 11-27-06 FINAL.doc
11/27/2006 9:24:00 AM