Fungibility and the Impact of Development Assistance

Evidence from Vietnam's Health Sector

Adam Wagstaff

The World Bank
Development Research Group
Human Development and Public Services Team
December 2008

Abstract

How can the impact of aid be estimated in the presence of fungibility? And how far does fungibility reduce its benefits? These questions are analyzed in a context where a donor wants to target its efforts on a specific sector and specific geographic areas. A traditional differences-in-differences method comparing the change in outcomes between the target and nontarget areas before and after the project risks misestimating the project's benefits. The paper develops an alternative estimation method in which intersectoral fungibility reduces project benefits insofar as government spending has a smaller impact in the sector to which the funds leak than in the target sector, while intrasectoral fungibility reduces benefits

insofar as the donor is able to leverage productivity increases in government spending in the target areas. The methods are applied to two contemporaneous World Bank health projects that set out to target assistance on approximately one-half of Vietnam's provinces. Aid is not apparently fungible between Vietnam's health sector and other sectors, but is fungible across provinces within the health sector. Differences-in-differences yield an insignificant impact on infant mortality, while the use of the new method yields a statistically significant impact of around 4 per 1000 live births. The results, however, are ambiguous on the costs associated with intrasectoral fungibility.

This paper—a product of the Human Development and Public Services Team, Development Research Group—is part of a larger effort in the department on health financing and service delivery arrangements. Policy Research Working Papers are also posted on the Web at <http://econ.worldbank.org>. The author may be contacted at awagstaff@worldbank.org.

The Policy Research Working Paper Series disseminates the findings of work in progress to encourage the exchange of ideas about development issues. An objective of the series is to get the findings out quickly, even if the presentations are less than fully polished. The papers carry the names of the authors and should be cited accordingly. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development/World Bank and its affiliated organizations, or those of the Executive Directors of the World Bank or the governments they represent.

Fungibility and the Impact of Development Assistance: Evidence from Vietnam's Health Sector

by

Adam Wagstaff^a

^a Development Research Group, The World Bank, Washington DC, USA

Corresponding author and contact details: Adam Wagstaff, World Bank, 1818 H Street NW, Washington, D.C. 20433, USA. Tel. (202) 473-0566. Fax (202)-522 1153. Email: awagstaff@worldbank.org.

Keywords: foreign aid, fungibility, impact evaluation, child mortality.

Acknowledgements: My thanks to Kiersten Johnson for answering queries about the DHS, to Joseph Capuno, Martin Rama and Minh Van Nguyen for sharing their provincial expenditure data, and to Misha Lokshin and Zurab Sajaia for answering queries on the use of movestay. The findings, interpretations and conclusions expressed in this paper and any errors are entirely mine; the views expressed do not necessarily represent those of the World Bank, its Executive Directors, or the countries they represent.

I. INTRODUCTION

Much has been written on the issue of aid fungibility—the notion that governments alter their planned expenditure allocations so that foreign aid does not necessarily get spent where it is allocated but rather is diverted elsewhere.¹ However, the bulk of the literature to date has focused on the issue of *whether* aid is fungible, most focusing on intersectoral fungibility (the issue of whether aid intended for, say, the transport sector in fact ends up financing defense projects) rather than intrasectoral fungibility (the issue of whether aid intended for, say, projects in poor rural areas ends up financing projects in affluent urban areas).² By contrast, few studies to date have examined two key questions, namely: How can one estimate the impact of development assistance in the presence of fungibility? And how far does fungibility reduce the benefits of aid? The few that have are aggregate studies based on cross-country regressions; they examine whether the conclusions reached in previous cross-country regressions linking foreign aid to outcomes such as economic growth and child mortality are modified when the (estimated) degree of fungibility of aid is taken into account.³

This paper presents a microeconomic perspective on the two questions mentioned above. The scenario is a common one facing multilateral and bilateral aid agencies, namely a donor wanting to target its efforts on a specific country, and a specific sector

¹ Pack and Pack (1990; 1993) seem to be the first published articles on the subject.

² Studies of intersectoral fungibility are several, and include Pack and Pack (1993), Feyzioglu, Swaroop and Zhu (1998) and Swaroop, Jha and Rajkumar (2000). Studies of intrasectoral fungibility are rarer, and include van de Walle and Mu (2007), who examine fungibility across Vietnamese provinces in assistance for road rehabilitation.

³ Pettersson (2007b) re-estimates the macroeconomic growth regressions of Burnside and Dollar (2000) distinguishing between aid that is fungible and aid that is non fungible; he finds no evidence of fungible aid being any less effective than nonfungible aid. Pettersson (2007a) includes the estimated fungibility of aid in a cross-country regression of child mortality on GDP and aid volumes; he finds the fungibility of aid has no independent effect on child mortality.

and specific geographic areas within the country. Any intersectoral fungibility results in the donor's resources leaking out of the sector, while any intrasectoral fungibility results in its resources leaking out of the target areas into the nontarget areas. Given the possibility of both types of fungibility, the first question arises of whether classic impact evaluation methods will yield an accurate estimate of the impact of development assistance. The paper shows first that the traditional differences-in-differences (diffs-in-diffs) comparing the change in outcomes between the target and nontarget areas before and after the donor's project risks misestimating the benefits of the project.

The paper goes on to propose an alternative approach to estimating the benefits of a project in such circumstances that requires the specification and estimation of a model linking the outcomes of interest to government spending, allowing the donor's project to modify this relationship by increasing the productivity of government spending in the target areas. Intersectoral fungibility reduces the benefit insofar as government spending in the sector where the funds leak to is less productive than the sector where the aid was intended, while intrasectoral fungibility reduces the benefits of aid insofar as the donor is able to leverage productivity increases in government spending in the target areas. The method thus sheds light on the second question posed above, namely how far fungibility reduces the benefits of development assistance. The methods are applied to two contemporaneous World Bank health projects that set out to target assistance on approximately one-half of Vietnam's provinces. The results are compared to the results obtained when the classic diffs-in-diffs method is used: statistically significant impacts of the projects are obtained using the former, but not when using the latter.

II. GOVERNMENT EXPENDITURE ALLOCATIONS AND FUNGIBILITY: THEORY

Consider a specific sector—call it sector H . Suppose that a donor divides the country in question into two areas: area 1 which it targets with its loans or grants; and area 2 which it does not target. Government spending on sector H in areas 1 and 2 (gross of any loan or grant) is measured along the x- and y-axes respectively of Figure 1. The initial distribution of government spending across the two areas is shown by point a , where government spending is X_{10} in area 1 and X_{20} in area 2. Area 2 has a higher level of spending, perhaps because government spending is decentralized and intergovernmental fiscal transfers do not equalize fiscal capacity.

Figure 1: General case of government spending allocations and development assistance

Suppose that in the absence of the donor's project, the government would have raised spending over the period being studied in both areas: to X_{11} in area 1 and X_{21} in

area 2, giving the new distribution shown by point b . This point is known to the government but is not observed by the analyst. Indeed, the donor may never know its location either. The donor provides resources for the government to use in area 1 equal to an amount D , equal to the distances df and bc . It is assumed that the donor can force the government to keep spending at least X_{10} in area 1, and can monitor its project sufficiently tightly so as to ensure that all of D is spent in area 1.

If the government went ahead with its planned spending increase and added the donor's resources to its own resources, its new budget line would be the kinked line passing through points d, f, e and c . The donor might naively assume that the government would select point c on this line, using its own planned extra resources to raise its own spending in areas 1 and 2, taking it from point a to point b , and then using the donor's loan to increase spending in area 1, taking the government from point b to point c .

The reality is likely to be different. First, the government might well decide not to operate on the outermost budget constraint at all. It might decide instead to react to the project by raising its spending on sector H less than planned. The loan crowds out part of the planned increase in government spending, and some of the resources that were to have been spent in sector H now get spent in other sectors. In effect, the government spends part of the donor's money in other sectors. This is the issue of *intersectoral fungibility*. If the government does decide to spend part of the loan or grant outside sector H , the new budget line would lie inside the outermost budget line in Figure 1. The most extreme case would be that the government cancels its spending increase altogether, and its total health spending rises by an amount equal to the loan or grant compared to the

pre-project outcome (i.e. point a). This would produce the dotted kinked budget line passing through points a and g . In this case, the government would likely be at a corner solution, i.e. at point g , government spending in area 1 would increase by the amount of the loan, and spending in area 2 would stay unchanged. In this example, this outcome is unlikely given that the government would have increased its spending on sector H to a higher level even without the loan. The budget line is therefore likely to be at least as far away from the origin as the line passing through point b , albeit with a kink above X_{13} .

In addition to the possibility of some of the loan getting spent outside the sector of interest, there is the possibility that some of the loan in effect gets spent in area 2. This is the case of *intrasectoral fungibility*, the fungibility in this case being between the geographic area(s) targeted by the project and the rest of the country. Suppose for the moment that all the aid stays intended for sector H stays there, so the budget constraint is the kinked line passing through points d, f, e and c . As drawn, Figure 1 shows point e as the government's preferred point. The government moves from point a to point e using (in the case of area 1) all of the resources provided by the donor (equal to $X_{13}-X_{12}$) and some of the extra resources it planned to use anyway. Note that the government increases spending on area 1 by an amount equal to $X_{12}-X_{10}$ which is smaller than the originally planned increase, given by the distance $X_{11}-X_{10}$. In area 1, then, the donor's project crowds out some of the planned increase in government spending. By contrast, in area 2, the project results in government expenditure increasing by *more* than planned, the increase there being $X_{23}-X_{20}$ rather than the originally planned increase of $X_{21}-X_{20}$. Even though the government complies with the donor's requirement that all of D must be spent in area 1, it in effect spends some of the foreign aid in area 2, which becomes a

beneficiary of the project, even though from an accounting perspective none of the resources that finance the extra spending in area 2 come from the donor.

Point e in Figure 1 is the point the government would have selected if its budget had increased by the same amount as the development assistance. (In the case of a pure budget increase, the government could have chosen a point on the dotted section to the northwest of point f .) The government's choice of point e when it receives the aid is the case of full fungibility. If, by contrast, the government had chosen a point to the southeast of point e , aid would be described as not fully fungible; the extreme case is when the government chooses point c , in which case the aid is said to be completely non-fungible. In this latter case, and in the intermediate cases of partial fungibility (points between e and c), there is said to be a flypaper effect—at least some of the aid ‘sticks’ to area 1 where it was intended to be spent, as flies stick to flypaper.

Figure 1 is the case where even though the project crowds out some of the planned increase in spending in area 1, the increase in spending in area 1 (relative to the pre-project level) is larger than the amount of the loan; the government supplements the loan with its own resources. Figure 2 shows the case where this does not happen. Here the government would like to select point e on the dotted section of the outermost budget line. It cannot do so, however, because the donor successfully forces it to spend at least $X_{10}+D$ in area 1. The optimal feasible point for the government is the corner solution at point f . The government moves from point a to point f using (in the case of area 1) none of the extra resources it planned to use ($X_{11}-X_{10}$) and relies exclusively on the resources provided by the donor (equal to $X_{13}-X_{10}$). In area 1, the project crowds out *all* of the

planned increase in government spending. By contrast, as in Figure 1, the project results in government expenditure increasing in area 2 by *more* than planned, the increase there being $X_{23}-X_{20}$ rather than the originally planned increase of $X_{21}-X_{20}$. In this example area 1 benefits by more than area 2, in contrast to Figure 1 where they benefited more or less equally.

Figure 2: Special case involving a corner solution

In both examples, compared to what would have happened otherwise, the donor's project raises the amount of resources going to area 1. But it also raises the amount of resources going to area 2. In neither example is it the case that *compared to what would have happened in the absence of the project* area 1 sees its government spending on sector H rise by the full amount of the loan. Area 2 in effect “captures” part of the total spending by the donor.

III. HOW FUNGIBILITY UNDERMINES TRADITIONAL IMPACT EVALUATION METHODS

Fungibility—between and within sectors—substantially complicates the job of evaluating the impacts of donor projects. When evaluating impacts, we are not interested in how the project changes the resources going to specific sectors and their allocation within the sector—the issue discussed in the previous section. Rather, we are interested in the impact of the project in terms of use of services, their quality, and outcomes. The relevant counterfactual is what these would have been in the absence of the project. The importance of the fungibility discussion is that all areas of the country—even *those not formally covered by the project*—may benefit from the project, since they may see a larger increase in government spending than would have been possible without the project. As a result, the quality of facilities, people’s use of them, and outcomes may all increase *across the entire country* by more than they would have done in the absence of the project.

Donors are presumably not indifferent to improved outcomes in nonproject areas. Indeed, it seems more reasonable to assume that they attach a positive value to improvements in, say, the quality of health facilities that occur in areas they have not targeted for their projects. Whether the value to the donor of changes in nonproject areas is the same as the value to the donor of changes in project areas is likely to depend on why the donor selected the particular areas for support. If it selected the areas more or less at random, or if they selected them because they felt the areas stood a better chance than other areas of achieving positive results with its money, the value attached to quality improvements in non-targeted areas ought to be equal to the value attached to quality

improvements in the targeted areas. If, on the other hand, the donor felt that people living in the targeted area were somehow more deserving (perhaps because the areas have historically been deprived ones), the value attached by the donor to quality improvements in the non-targeted areas might be somewhat lower. But even in this case, the value is most unlikely to be zero.

Given fungibility, and given that donors are likely to value improvements in quality, utilization and outcomes whether they occur in the areas they target or in other areas, the estimation of the benefits of a project that targets specific areas becomes problematic. In this section, we show what the true benefits are and compare these with the benefits estimated by the traditional approach that compares the changes before and after the project between the project and nonproject areas (the so-called ‘double difference’ approach). We see in this section that the latter approach is likely to result in an unreliable estimate of the true benefits of the project.

Let Y_1 be the outcome of interest in area 1 and Y_2 the outcome of interest in area 2. We denote by Y_{13} and Y_{23} the outcomes corresponding to the government expenditure allocation X_{13} and X_{23} , i.e. point e in Figure 1, and analogously for other points in the chart. Suppose that the donor values outcomes in areas 1 and 2 equally, and its indifference curves are straight lines with slopes of -1.⁴ Then the benefit of the project—in the donor’s eyes—is equal to:

$$(1) \quad B = [Y_{13} - Y_{11}] + [Y_{23} - Y_{21}],$$

⁴ A more sophisticated analysis would use a social welfare function for the donor that would allow for convex donor indifference curves, thereby allowing for the possibility that the donor is averse to inequalities in outcomes between the project and nonproject areas.

which is equal to the difference between the observed outcome in area 1 and the counterfactual outcome in area 1, plus the difference between the observed outcome in area 2 and the counterfactual outcome in area 2. In terms of Figure 1, we are interested in the outcomes (in the two areas combined) at point e , which is the observed final expenditure allocation, compared to the outcomes at point b , which is the counterfactual expenditure allocation that would have occurred in the absence of the project.

Compare eqn (1) to the classic double-difference (DD) approach, where a project's benefits are estimated by comparing the change in outcomes in the project area with the change in outcome in the nonproject area. Ignore for the moment the fact that the Y_{11} and Y_{21} are not observed. The 'true' DD (the DD one would like to estimate) is equal to:

$$(2) \quad DD^* = [Y_{13} - Y_{11}] - [Y_{23} - Y_{21}].$$

Combining eqns (1) and (2) we get:

$$(3) \quad DD^* - B = [Y_{13} - Y_{11}] - [Y_{23} - Y_{21}] - [Y_{13} - Y_{11}] - [Y_{23} - Y_{21}] = -2[Y_{23} - Y_{21}],$$

and therefore

$$(4) \quad DD^* = B - 2[Y_{23} - Y_{21}] \leq B,$$

so the true double-difference underestimates the true benefits of the project; the gap is equal to twice the increase in the outcome in the nonproject areas compared to the outcome that would have occurred in the absence of the project. Only in the case where there is no within-sector fungibility does the true double difference, DD^* , estimate the benefit of the project, B : in this case, we have $X_{23}=X_{21}$ and hence $Y_{23}=Y_{21}$, so that $DD^*=B$.

This ignores the fact that Y_{11} and Y_{21} are unobservable. The obvious DD one would estimate is not DD^* but rather the *feasible* double-difference:

$$(5) \quad DD^0 = [Y_{13} - Y_{10}] - [Y_{23} - Y_{20}].$$

Combining eqns (5) and (3) we get:

$$(6) \quad \begin{aligned} DD^0 - B &= [Y_{13} - Y_{10}] - [Y_{23} - Y_{20}] - [Y_{13} - Y_{11}] - [Y_{23} - Y_{21}] \\ &= [Y_{11} - Y_{10}] - [Y_{23} - Y_{20}] - [Y_{23} - Y_{21}]. \end{aligned}$$

In the special case where there is no fungibility, we have $Y_{23}=Y_{21}$, giving

$$(7) \quad DD^0 - B = [Y_{11} - Y_{10}] - [Y_{21} - Y_{20}].$$

A sufficient condition for the right-hand side to be zero is that the government's preferences are homothetic, in which case points a and b in Figure 1 would lie on a straight line through the origin, and expenditures are equally productive in both areas. In general, we would not expect this to be the case, and typically DD^0 will over- or underestimate the benefits of the project. The sign of any bias cannot be determined *a priori*.

IV. ESTIMATING IMPACTS IN THE PRESENCE OF FUNGIBILITY

To get at the true benefits, we need to make some assumptions about the links between spending and outcomes, and how this relationship is modified by the project.

Quantifying impacts

Assume that the country consists not just of two areas, but rather multiple areas (e.g. provinces, counties, districts, etc.), some of which are designated project areas and

some of which are not. Assume we have government spending data at this level. Assume too we have data on outcomes (and non-spending determinants of these outcomes) at the micro level (in the application below the unit is the individual, but the method can be applied to, say, a school or a health facility). We use i to denote the i th individual, and $j=0,1,3$ to indicate whether we referring to the (observed) baseline scenario ($j=0$), the (unobserved) scenario without the project ($j=1$), and the (observed) scenario with the project ($j=3$). Thus:

$$(8) \quad \begin{aligned} Y_{1j} &= \sum_i Y_{i1j} = N_{1j} \bar{Y}_{1j} \\ Y_{2j} &= \sum_i Y_{i2j} = N_{2j} \bar{Y}_{2j} \end{aligned},$$

where N_{1j} and N_{2j} are the number of individuals in areas 1 and 2 respectively in scenario j , and $\bar{Y}_{1j}$ and $\bar{Y}_{2j}$ are the mean values of Y_{i1j} and Y_{i2j} respectively.

Assume the data are available at two dates: before the project is implemented (which corresponds to point a in Figure 1), and after the project has been implemented or at least after it has begun implementation (which corresponds to point e). The data need not be panel data. Of course, we have no data corresponding to the counterfactual scenario, point b , since this is never observed. This absence of data is addressed below. Let $t=0,1$ denote whether the time period is prior to the project's start or afterwards. We assume that for the two observed sets of data, outcomes are linked to government spending X_i (other determinants are omitted from the equations below for simplicity; they are not excluded from the empirical model) as follows:

$$(9) \quad \begin{aligned} Y_{i1t} &= \alpha_1 + \theta_1 t_i + X_{it} \beta_1 + X_{it} \phi + \varepsilon_{i1t} \\ Y_{i2t} &= \alpha_2 + \theta_2 t_i + X_{it} \beta_2 + \varepsilon_{i2t} \end{aligned}$$

Eqn (9) allows the project to affect the productivity of government spending (in the period after the project has begun implementation) through the interaction between X and t in the project-area part of the equation. In addition, the project may result in a larger or smaller shift in the intercept (θ_1 and θ_2 may be different from one another). The specification also allows the initial intercepts to be different in the project and nonproject areas (α_1 and α_2 may be different from one another) and the initial spending productivity to be different (β_1 and β_2 may be different from one another).

Assume for the moment we have estimates of the relevant parameters in the outcome-generating equation (9). The relevant averages and totals corresponding to point e in Figure 1 are equal to:

$$(10) \quad \begin{aligned} \bar{Y}_{13} &= \alpha_1 + \theta_1 + \bar{X}_{13}(\beta_1 + \phi); & Y_{13} &= N_{13}[\alpha_1 + \theta_1 + \bar{X}_{13}(\beta_1 + \phi)] \\ \bar{Y}_{23} &= \alpha_2 + \theta_2 + \bar{X}_{23}\beta_2; & Y_{23} &= N_{23}[\alpha_2 + \theta_2 + \bar{X}_{23}\beta_2] \end{aligned}.$$

Getting the averages and total corresponding to point b is problematic since we have data only for points a and e . Suppose that in the event the project had not gone ahead, the outcomes in areas 1 and 2 at time $t=1$ would have been generated by:

$$(11) \quad \begin{aligned} Y_{i11} &= \alpha_1 + \theta_2 + X_i\beta_1 + \varepsilon_{i11} \\ Y_{i21} &= \alpha_2 + \theta_2 + X_i\beta_2 + \varepsilon_{i21} \end{aligned},$$

This gives:

$$(12) \quad \begin{aligned} \bar{Y}_{11} &= \alpha_1 + \theta_2 + \bar{X}_{11}\beta_1; & Y_{11} &= N_{11}[\alpha_1 + \theta_2 + \bar{X}_{11}\beta_1] \\ \bar{Y}_{21} &= \alpha_2 + \theta_2 + \bar{X}_{21}\beta_2; & Y_{21} &= N_{21}[\alpha_2 + \theta_2 + \bar{X}_{21}\beta_2] \end{aligned}.$$

Suppose that the numbers in the two groups remain constant across the three scenarios:

$$(13) \quad \begin{aligned} N_{10} &= N_{11} = N_{13} = N_1 \\ N_{20} &= N_{21} = N_{23} = N_2 \end{aligned} .$$

Then we can write:

$$(14) \quad \begin{aligned} Y_{13} - Y_{11} &= N_1 [\alpha_1 + \theta_1 + \bar{X}_{13}(\beta_1 + \phi)] - N_1 [\alpha_1 + \theta_2 + \bar{X}_{11}\beta_1] \\ &= N_1 [\theta_1 - \theta_2 + (\bar{X}_{13} - \bar{X}_{11})\beta_1 + \phi\bar{X}_{13}] \end{aligned}$$

and

$$(15) \quad \begin{aligned} Y_{23} - Y_{21} &= N_2 [\alpha_2 + \theta_2 + \bar{X}_{23}\beta_2] - N_2 [\alpha_2 + \theta_2 + \bar{X}_{21}\beta_2] \\ &= N_2 \beta_2 (\bar{X}_{23} - \bar{X}_{21}) \end{aligned} .$$

Substituting these into eqn (1), and assuming $\beta_1 = \beta_2 = \beta$, gives:

$$(16) \quad \begin{aligned} B &= [Y_{13} - Y_{11}] + [Y_{23} - Y_{21}] \\ &= N_1 [\theta_1 - \theta_2 + (\bar{X}_{13} - \bar{X}_{11})\beta + \phi\bar{X}_{13}] + N_2 \beta (\bar{X}_{23} - \bar{X}_{21}) \\ &= N_1 [(\theta_1 - \theta_2) + \phi\bar{X}_{13}] + \beta [N_1 \bar{X}_{13} + N_2 \bar{X}_{23} - (N_1 \bar{X}_{11} + N_2 \bar{X}_{21})] \\ &= N_1 [(\theta_1 - \theta_2) + \phi\bar{X}_{13}] + \beta (X_3 - X_1) \\ &= N_1 [(\theta_1 - \theta_2) + \phi\bar{X}_{13}] + \beta (X_1 + D - X_1) \\ &= N_1 [(\theta_1 - \theta_2) + \phi\bar{X}_{13}] + \beta D \end{aligned} .$$

where D is the donor's total outlay, $X_1 = N_1 \bar{X}_{11} + N_2 \bar{X}_{21}$ is government spending in scenario 1, and we have assumed that all the donor's money sticks in sector H , i.e.

$X_3 = N_1 \bar{X}_{13} + N_2 \bar{X}_{23} = X_1 + D$. To capture the possibility of fungibility *between* the sector H and other sectors, eqn (16) could be replaced by

$$(17) \quad B = N_1(\theta_1 - \theta_2) + \phi N_1 \bar{X}_{13} + \beta \lambda D.$$

where λ is the fraction of an extra dollar in aid to sector H that ends up being spent in sector H . If aid intended for sector H is completely non-fungible between sectors, λ is equal to 1.00. If aid is fully fungible between sectors, λ will be equal to the government's marginal propensity to spend (MPS) on sector H from its own revenues. So, for example, if a one-dollar rise in government revenues increases its spending on sector H by 15 cents, λ would be equal to 0.15. If aid is incompletely fungible, i.e. if there is a 'flypaper effect', λ will be somewhere between the government's MPS (0.15 in the example) and 1.00.

The interpretation of eqn (17) is straightforward. The donor's project adds to *national* government spending on sector H an amount of money equal to D (or more generally λD). This improves outcomes *in the country as a whole* by an amount equal to βD (or more generally $\lambda \beta D$). Fungibility between sector H and other sectors reduces the size of this effect. The donor may have additional impacts on outcomes *in the project areas*. The project may bring about increases in the productivity of government spending in project areas, reflected in a nonzero value of ϕ . This raises outcomes *in the project areas* by an amount $\phi N_1 \bar{X}_{13}$. Finally, the donor's efforts may result—in the project areas—in a larger increase in the intercept in the relationship linking outcomes to spending than would have occurred in the absence of the project, reflected in θ_1 differing from θ_2 , and giving an additional benefit equal to $N_1(\theta_1 - \theta_2)$.

How impacts are affected by fungibility

Eqn (17) also sheds light on the question of how inter- and intrasectoral fungibility affects the size of a project's benefits. Intersectoral fungibility is reflected in a value of λ that is less than 1, and resources get diverted from sector H to other sectors. The use of eqn (17) to compute benefits assumes that all resources lost from sector H through fungibility have no impact on the outcome of interest: for each dollar so lost, the outcome of interest falls by an amount β . This is clearly a strong assumption, and may be unwarranted. Spending in sectors other than sector H may have a beneficial effect on the outcome of interest. For example, if sector H is the health sector and the outcome of interest is child health, leakage of funds from the health sector to water and sanitation projects may not matter much. But government spending in other sectors may be less beneficial to child health, and spending in some sectors may be positively detrimental (e.g. spending on energy projects might result in pollution that might cause respiratory and other health problems). If λ turns out to be less than 1, one ought to estimate β for the sectors where the funds leak to and compute a net loss from fungibility.

Intrasectoral fungibility is reflected in $\bar{X}_{13}$ being lower than would have otherwise been the case (and $\bar{X}_{23}$ being higher): with zero intrasectoral fungibility, $\bar{X}_{13} = \bar{X}_{11} + \bar{D}$ and $\bar{X}_{23} = \bar{X}_{21}$ (cf. Figure 1). If the donor leverages productivity increases in its area (but not in the area not targeted by the donor), i.e. ϕ is nonzero and the same sign as β , intrasectoral fungibility will result in a lower benefit than would have otherwise been the case. The cost of intrasectoral fungibility is the forgone additional benefits brought about by the donor's productivity-enhancing measures. This cost could

be calculated. One could estimate the government's marginal propensity spend from years prior to the start of the project, thereby obtaining the counterfactual amounts $\bar{X}_{11}$ and $\bar{X}_{21}$. The cost of intrasectoral fungibility could then be computed as ϕ_1 times N_1 times the difference between the estimated counterfactual value of $\bar{X}_{13}$ and its actual value. Note that intrasectoral fungibility has no cost if $\phi_1=0$, i.e. if the donor does not leverage any productivity increases in the target areas.

V. THE CASE OF TWO WORLD BANK HEALTH PROJECTS IN VIETNAM

The World Bank's Population and Family Health Project and National Health Support Project in Vietnam⁵ are just two examples of donor projects that fit the scenario outlined above.⁶

Around half of the budget of each project was directed to 37 of Vietnam's 64 provinces, with the overall objective of improving primary care—especially maternal and child care—in commune and district health centers in the selected provinces.^{7,8} The

⁵ Details of the projects are to be found in World Bank (1995; 2004; 2006; 2007), which can be downloaded from <http://web.worldbank.org/WBSITE/EXTERNAL/PROJECTS/0,,menuPK:115635~pagePK:64020917~piPK:64021009~theSitePK:40941,00.html>.

⁶ In actual fact, the projects were cofinanced by the Vietnamese government, the World Bank (more precisely the International Development Association), and other donors, the Dutch government and the Swedish Development Agency (SIDA) in the case of the health support project, and the Asian Development Bank and the German Kreditanstalt für Wiederaufbau (KfW) in the case of the population and family health project.

⁷ Both projects aimed to disburse around \$ 130 million. The percentage allocated to the province-specific interventions was 47% in the case of the health support project and 49% in the case of the population and family health project. The rest of the two projects' funds were not targeted on specific provinces, and instead addressed national issues and programs. The rest of the health support project budget was directed at two areas: national priority programs concerning malaria, TB and ARI (40%); and health planning and management capacity in the health ministry (4%). The rest of the population and family health project budget was directed at family health and family planning issues nationwide, and included: a component on information, education and communication (10%); a component on contraceptive supplies (19%); a component on family planning (4%); and a component on service delivery model innovation (6%).

⁸ The population and family health project targeted 15 provinces (this grew to 18, and subsequently 20), and the health support project targeted a further 15 (this also grew to 18). The number of provinces grew during implementation, in part because some provinces split into two, and in part because new provinces were added to both projects. By the end,

measures implemented to achieve this goal, which were similar in the two projects, included: the upgrading of facilities; the purchase of equipment, vehicles and furniture; the provision of essential drugs; in-service training; and outreach services. Criteria for selecting provinces included: low GDP per capita; low health expenditures per capita; below-average health indicators; implementation capacity in the provincial government's department of health; the availability of other donor assistance; and a reasonably even geographic coverage. The two projects were intended to run in parallel: both were approved by the Bank's Board of Directors in January 1996, became 'effective' (i.e. active) in mid-1996, began disbursing in late-1996 and early 1997, and were originally scheduled to close in 2003.⁹

The problem of fungibility and the consequent inappropriateness of comparisons of project and nonproject provinces were recognized in the so-called 'implementation completion report' of the population and family health project (World Bank 2004) but not in that of the health support project (World Bank 2007).¹⁰ The latter cites evidence of matched (single-difference) comparisons between facilities in project provinces and facilities in nonproject provinces as support of the project's impact.¹¹ By contrast, the population and family health project report argued that conventional evaluation methods were flawed in part because there was an explicit understanding between the government and the Bank that the government would devote extra resources to the provinces not

37 of Vietnam's 64 provinces were covered by one or other of the projects; one province—Dac Lac—appears on both project lists.

⁹ In the event, the health support project was extended three times and eventually closed in 2006. The population and family health project was extended only once and only by three months; it closed in September 2003.

¹⁰ The problem does not seem to have been appreciated early on in the population and family health project, however. The Bank's Independent Evaluation Group's project performance assessment report (World Bank 2006) was critical of the evaluation efforts of the project in part on the grounds that "the approach of using use of nonproject provinces as the counterfactual was fundamentally flawed" (World Bank 2006 p.27).

¹¹ This evidence referred to is apparently that subsequently published in Fritzen (2007).

covered by the project and would aim to provide a similar set of interventions in nonproject provinces to those being provided to project provinces through the two projects.¹² This, it was argued, made it impossible to conduct a formal impact evaluation of the project.¹³ The same point was made by the authors of a report on the population and family health project by the Bank's Independent Evaluation Group (World Bank 2006), who were forced to resort to examining national trends in key indicators without trying to determine what could be attributed to the project.

The behavior of the Vietnamese government appears to have been precisely that illustrated in Figure 1. It apparently responded to the World Bank projects by scaling back its planned increase in spending in the project provinces (area 1), relying largely on the Bank to finance extra spending (compared to the initial level) in these areas. The Bank projects thus allowed the Vietnamese government to achieve higher levels of spending (gross of Bank-provided funds) in both project *and* nonproject provinces. In such a situation, the use of differences-in-differences would indeed fail to provide a reliable estimate of the projects' overall benefits.

The approach outlined above provides a better way of approaching the estimation of the combined benefits of the two Bank projects. The benefits are twofold. First, the two projects combined allowed government spending nationally to increase by more than it would have otherwise done. If the donor assistance had stuck with the health sector, the two projects combined—including the spending targeted at specific provinces and other spending—would have added around US\$ 200 million over the course of the seven-year

¹² See pp. 2 and 7 of the population and family health project report (World Bank 2004).

¹³ See p.2 of the population and family health project report (World Bank 2004).

implementation period, equivalent to around US\$ 0.35 per capita per annum. Total government spending on health—from all sources, including development assistance—averaged around \$US 5.50 per capita per annum over the period 1997-2003. So, the two projects combined raised government health spending by around 7% per annum. The first element of the benefits of the projects is the extra health generated by this extra spending: this is the βD term in eqn (16).

There is a second potential element to the projects' benefit, namely that the projects may have leveraged increases in the productivity of government health spending in the project provinces. In terms of eqn (16), this would be reflected in θ_1 and θ_2 being different (the intercept in the outcome-spending equation shifted up by a different amount during the project in project provinces than nonproject provinces), and/or a nonzero value of ϕ (after the project begins, government health spending has a bigger or smaller effect on outcomes at the margin in project provinces than in nonproject provinces). If indeed ϕ is nonzero, the addition to the projects' benefit is equal to ϕ times the total amount of government spending in the project provinces, the reason being that the projects leverage productivity increases not only on the marginal expenditures funded by the foreign assistance but also on the expenditures funded domestically.

How might a project leverage productivity increases in government health spending in project provinces? There are three obvious possible channels. The first is that the donor might impose strict procurement procedures—e.g. requiring competitive bidding—through which input prices are driven down, at least on inputs purchased with donor finance. The World Bank does, in fact, have strict rules on procurement, and while

both projects' implementation completion reports comment on how challenging these rules were to the government unit responsible for managing the two Bank projects, their use seems to have led to cost savings. The completion report of the population and family health project, for example, notes that through "the assiduous use of bidding procedures" the actual cost of the civil works element of the primary care component of the project was almost half that expected (\$US 14 million instead of \$US 25 million). This enabled additional upgrading and equipment purchases beyond that planned. The second channel is that the project may have altered the input mixes in the health facilities of project provinces compared to what they would have otherwise been, by specifying not only *how* inputs were to be procured but also *what* was to be procured. Again, the project did, in fact, have strict rules about the drugs and equipment to be purchased: the report of the population and family health project, for example, comments that "the project was designed as a *well coordinated package* of hardware and software linking civil works, equipment, drugs, training, ..., etc." (World Bank 2004 p14, emphasis in original). Finally, the projects may have leveraged increases in productivity of government spending in project provinces by putting in place stricter staff supervision, better training programs, new protocols, etc. that resulted in more health improvements being obtained from a given bundle of inputs. Again, this is also plausible. Training in family planning and reproductive health, for example, was more advanced in project than nonproject provinces.

While each of these channels is plausible, a couple of points are worth noting. First, nonproject provinces may also have benefited from the same measures, though it is unclear whether these benefits began accruing before 2003. Apparently, the bidding and

contracting procurement procedures were adopted nationwide by the health ministry but only after the project's end. Whether the government followed the project's rules and approaches vis-à-vis the other two channels is unclear from the project documents. Nonetheless, it seems at least plausible that the project may have had spillover effects that enhanced the productivity of government spending in nonproject provinces. Second, while the project may indeed have increased the productivity of government spending at least in project provinces, the limited scale of the project needs to be kept in mind. The project will have affected the prices and composition only of those inputs financed by the project, not of all inputs financed out of domestic government and donor funds. Insofar as the inputs are those most relevant for the outcomes in question, the leveraging could have been quite high; the more marginal and the less comprehensive the inputs, the smaller the likely leveraging of productivity increases.

VI. MODEL AND DATA

This section explains the methods and data used to estimate the benefits of the two World Bank health projects using the approach outlined in section IV. The outcome examined is infant mortality—deaths among children in the first year of life. Both Bank projects explicitly aimed to reduce infant mortality, and used it as a monitoring indicator. It has the considerable merit of being an outcome indicator, rather than an indicator of utilization or process.

Model

In the case of infant mortality, Y_{1j} and N_{1j} are to be interpreted respectively as the number of infant deaths and live births in area 1 in scenario j , $\bar{Y}_{1j}$ is the fraction of infants dying before their first birthday in area 1 in scenario j , and Y_{i1j} as a dummy variable indicating whether the i th child in scenario j survived through to his or her first birthday.

The parameters required to compute the project's benefits from eqn (17) are estimated using a model along the lines of eqn (9):

$$(9') \quad \begin{aligned} Y_{i1t} &= \alpha_1 + \theta_1 t_i + X_{it}^1 \beta_1 + X_{it}^1 t_i \phi_1 + X_{it}^2 \delta_1 + \varepsilon_{i1t} \\ Y_{i2t} &= \alpha_2 + \theta_2 t_i + X_{it}^1 \beta_2 + X_{it}^1 t_i \phi_2 + X_{it}^2 \delta_2 + \varepsilon_{i2t} \end{aligned},$$

where we have allowed factors other than government spending (now labeled X^1) to influence outcomes (the other determinants being labeled X^2), and for the possibility that after implementation of the Bank project government spending may become more productive in the nonproject provinces (ϕ_2 is not necessarily zero). In addition, we require a selection equation explaining where the project is located:

$$(18) \quad \begin{aligned} Y_{i1} &\text{ observed if } Z_i \gamma + u_i > 0 \\ Y_{i2} &\text{ observed if } Z_i \gamma + u_i \leq 0 \end{aligned},$$

where Z_i is a vector of determinants of project placement, which includes X^1 and X^2 .

Eqns (9') and (18) together comprise an endogenous switching model. Various special cases are of interest, the most obvious being where we impose the two restrictions $\beta_1 = \beta_2$ and $\phi_2 = 0$, both of which are imposed in the derivation of eqn (16). Other special cases of interest are where restrictions are imposed on the covariance structure of the

three error terms. Denote by ρ_1 and ρ_2 the covariances between u and ε_1 and u and ε_2 , and by σ_1^2 and σ_2^2 the variances of ε_1 and ε_2 . The general case where $\rho_1 \neq \rho_2$ and $\sigma_1^2 \neq \sigma_2^2$ is consistent with what Heckman et al. (2006) call ‘essential heterogeneity’, i.e. there are unobserved individual-specific returns to the project and these affect the placement of the project. A special case of interest is where $\rho_1 = \rho_2$ (but $\sigma_1^2 \neq \sigma_2^2$): this is the case where there are unobserved individual-specific returns to the project but these do not affect the placement of the project (cf. e.g Basu *et al.* 2007). A further special case of interest is where $\sigma_1^2 = \sigma_2^2$ (and $\rho_1 = \rho_2$): this is the case where there are no unobserved individual-specific returns to the project, which is the case handled by the traditional Heckman (1978) endogenous dummy variable model.

Data and variables

The two years we use are 1997 and 2003, both years when the Demographic Health Survey (DHS) was fielded in Vietnam.¹⁴ The fieldwork for the 1997 DHS was undertaken in July-October, after the two Bank projects had started (Committee for Population Family and Children of Vietnam and ORC Macro 2003). However, both projects were slow to start disbursing, and little had happened by the time the 1997 DHS went into the field; for example, as of October 1997, the health support project had disbursed only 3% of the total amount disbursed by the end of the project (World Bank 2007). The two DHS’s were fielded in the same 205 communes in both years¹⁵; these communes fell in 41 of Vietnam’s 64 provinces, 19 being population and family health

¹⁴ This is not by accident: the DHS was funded by the population and family health project with the intention that the two surveys would be used to monitor and evaluate the project. In the event, the only public evaluation of the project (Fritzen 2007) was undertaken using an altogether different survey.

¹⁵ Different commune identifiers were used in the two surveys, however.

project provinces and 12 being health support project provinces. The DHS thus covers 30 project provinces and 11 nonproject provinces.¹⁶

The outcome we explore is, as already mentioned, infant mortality. We have data at the level of the individual child indicating when the child was born, if the child is still alive, and if not when it died.¹⁷ We classify a child as an infant death if it died within its first year. We focus on children who were born at least a year before the interview (this ensures we have a full year of exposure to the risk of death), and who were born not more than five years before the interview (this ensures that the births in the 2002 survey we use occurred after the start of the two projects). In the two years combined, we have 4568 children born between 1 and 5 years prior to the survey.

We use government health spending data at the provincial level from the years 1997 and 2003, and convert them to a per capita basis and to 1997 prices using the GDP deflator for Vietnam.¹⁸ In addition to government health spending, X^1 , we include as our X^2 vector a vector of wealth quintile dummies, the expectation being (and, as it turns out, broadly consistent with the data) that the survival prospects of newborns will tend to be better in more wealthy households.¹⁹ The DHS does not include income or consumption, and like others (cf. e.g. Filmer and Pritchett 2001) we proxy household wealth through a

¹⁶ One province—Dac Lac—appears on both project lists. See footnote 8.

¹⁷ One limitation of the DHS (World Bank 2006) is that it collects birth histories only from (ever) married women.

¹⁸ The provincial government health expenditure data cover spending in facilities operated by the provincial government or lower levels of government, however financed. In addition to including spending financed from the provincial government's own budget and transfers from central government (whether general or specifically for so-called 'national' programs), it also includes spending financed by local 'people's committees', foreign and domestic donors, payments by Vietnam's social health insurance scheme, and user fees levied at government facilities. It excludes spending in central government facilities in a province, such as national teaching hospitals. In 2000, spending in central government facilities accounted for only 17% of total government health spending.

¹⁹ The exclusion of other household-level variables—such as mother's education—means that the wealth quintile coefficients likely pick up multiple determinants of child survival that are correlated with the household wealth index. The focus of the estimation is on the coefficients on government spending, so not including these other factors explicitly ought not to matter.

linear combination of a long list of variables capturing a household's ownership of household durables, the materials of its house, etc. We use Principal Components Analysis (PCA) to derive the weights: using the first principal component, we obtain a z-score with zero mean and unit variance that is increasing in household wealth, and from this we form quintiles of wealth. To allow for increasing wealth between 1997 and 2002, we conducted the PCA on the pooled 1997 and 2002 data, and the weights from this exercise were then used to construct the wealth score in each year. The resultant wealth index has a zero mean in the pooled sample, but a higher mean in 2002 than in 1997 (the 2002 mean was positive and the 1997 mean negative).

Among the Z -variables in the selection equation, we include the province's GDP per capita, and the fraction of a province's government expenditures (on all sectors, not just health) covered through its own revenues. Given that the projects targeted poorer provinces, we would expect GDP per capita to enter in the program placement equation with a negative sign, and given provinces with institutional capacity were favored, we might expect the fraction of government expenditures financed locally to enter with a positive sign. Both coefficients, in the event, have the expected sign. It might be argued that the first of these variables could belong in the outcome equation; after all, studies have shown that, controlling for government health spending, GDP continues to exert an influence on infant mortality (cf. e.g. Filmer and Pritchett 1999; Bokhari *et al.* 2007). These are, however, aggregate-level studies; here, by contrast, we are exploring the impact of government health spending on child-level mortality data, holding constant the wealth of the child's household. One would have to believe that GDP per capita in a

child's province has an effect on his or her survival prospects independently of its impact on government health spending, conditional on the wealth of the child's household.²⁰

VII. ESTIMATES OF PROJECT BENEFITS

This section uses the approach outlined above to estimate the benefits of the two Bank health projects in Vietnam. We first report some results of regressions aimed at shedding light on the degree of fungibility between the health sector and other sectors, and within the health sector between different provinces. We also present some naïve *diffs-in-diffs* estimates of the impact of the Bank projects on infant mortality. These results help motivate the analysis that follows.

Fungibility and naïve *diffs-in-diffs* results

The regressions in columns (1) and (2) of Table 1 shed light on the issue of intersectoral fungibility. These equations, like those used in other studies of fungibility (cf. e.g. Devarajan *et al.* 1999), regress (provincial) government health expenditures (gross of development assistance (ODA) to the health sector) on GDP per capita, and net provincial revenues or expenditure (for *all* sectors, net of ODA).²¹ The data are cross-section and for the year 2000. Regressions are reported for both provincial total revenues and provincial total expenditures. Arguably the latter is a better measure of the resources

²⁰ We explored empirically the possibility that these two variables might belong in the outcome equation, keeping one as the identifying instrument and including the other in the outcome equation. In the general model, convergence was achieved with the revenue share variable included in the outcome equation but not when GDP per capita was included. When restrictions were imposed on the models (see below), convergence was typically achieved for GDP per capita and was always achieved for the revenue share variable. The coefficient on GDP per capita was not significant in the outcome equation, nor was the coefficient on the revenue share variable. We also explored other possible Z-variables, such as population density, communes per 1000 population, etc., but none had a significant effect in the selection equation.

²¹ In computing net (total) revenues and expenditures, only health sector ODA was subtracted, due to lack of data on ODA to other sectors.

that could have been devoted to the health sector given Vietnam's system of revenue- and expenditure-sharing between provincial and central government; some of the revenues collected at provincial level are passed to the center and distributed across provinces as transfers.²² In both columns (1) and (2), the coefficient on ODA is significantly larger than that on the provincial government's total resources, and is not significantly different from one. This suggests that fungibility of domestic and external resources between the health sector and other sectors may not be a major issue in Vietnam, and that a value of one for the parameter λ in eqn (17) may be a reasonable approximation pending more in-depth econometric work on intersectoral fungibility.

²² On fiscal decentralization in Vietnam, see Government of Vietnam-Donor Working Group on Public Expenditure (2000).

Table 1: Estimates of fungibility and naïve diffs-in-diffs estimates of project impact

	(1)		(2)		(3)		(4)		(5)	
			Provincial government health expenditures						Infant mortality	
	coef	t	coef	t	coef	t	coef	t	coef	t
GDP	-0.003**	-2.158	0.000	0.613	-0.004**	-2.555	-0.000	-0.331		
ODA	1.275***	5.939	1.060***	6.671						
Net prov. revenues	0.006**	2.414								
Net prov. expenditures			0.046***	7.970						
Prov. revenues					0.009***	2.711				
Prov. revenues x Bank project					0.001	0.255				
Prov. expenditures							0.057***	7.412		
Prov. exp. X Bank project							-0.002	-0.511		
2002 dummy									-0.002	-0.333
Bank project									0.012*	1.825
Bank project x 2002 dummy									-0.008	-0.884
Constant	36.284***	9.877	9.537***	3.727	43.598***	9.455	8.632***	2.623	0.018***	4.460
Additional t-tests:										
ODA diff. from own resources		5.90		6.32						
ODA diff. from 1		1.28		0.38						

Notes: .01 - ***, .05 - **, .1 - *. t-statistics in column (5) reflect standard errors adjusted for clustering at the province level. ODA is overseas development assistance.

Columns (3) and (4) shed light on the issue of intrasectoral fungibility, and specifically between the provinces targeted by the two Bank projects and Vietnam's remaining provinces. If domestic and external resources are fungible between the project and nonproject provinces, an increase in the province's total resources ought to have the same impact on government health spending in project provinces as in nonproject provinces. The equations in columns (3) and (4) regress the same dependent variable—provincial government health expenditures (gross of overseas development assistance (ODA) to the health sector—on total provincial resources gross of ODA (either revenues or expenditures, as before) and this variable interacted with a dummy indicating whether the province is a project province. In neither column is the coefficient on the interaction term anywhere near significant; this suggests that domestic and external resources in Vietnam are fungible within the health sector across provinces. This is consistent with the opinions expressed in the Bank's population and family health project documents and with the government behavior depicted in Figure 1.

The results in columns (3) and (4), together with the conclusions of section III, suggest that a simple *diffs-in-diffs* comparing the change in infant mortality between the project and nonproject provinces is likely to misestimate the benefits of the Bank projects. Column (5) reports the results of this *diffs-in-diffs* calculation, where the infant mortality dummy (at the child level) is regressed on a 2002 dummy, a project province dummy, and the interaction between the two; the coefficient on the latter indicates the *diffs-in-diffs* estimate of the project's impact on the infant mortality rate (cf. e.g. Ravallion 2008). The coefficient is negative and implies a reduction in the infant

mortality rate of 8 per 1000 live births; however, it is not significantly different from zero. This provides a motivation for employing the methods outlined in section IV to estimate the benefits of the project.

Estimates of project impacts

Table 2 reports the results of specification tests of the endogenous switching model.²³ Model (1) is an unrestricted model. Model (2) imposes the restrictions $\beta_1=\beta_2$ (government spending has the same effect in the project and nonproject provinces) and $\phi_2=0$ (government spending in the nonproject provinces has the same effect after the start of the project as before), which are imposed in the derivation of eqn (16). This pair of restrictions is not rejected: the prob value of the relevant likelihood ratio (LR) test is 0.66. Model (3) imposes—in addition to the pair of restrictions just mentioned—several further restrictions, namely: $\delta_1=\delta_2$ (wealth has the same effect in the project and nonproject provinces), $\alpha_1=\alpha_2$ (the intercepts are the same in the project and nonproject provinces), and $\theta_1=\theta_2$ (the upward shift in the outcome-expenditure relationship between 1997 and 2002 is the same in the project and nonproject provinces). The prob value of the relevant LR test statistic is 0.18, so these restrictions are also consistent with the data at conventional levels. Model (4) imposes the further restriction $\phi_1=0$, i.e. the project does not leverage increases in the productivity of government spending in project provinces in terms of infant mortality. This restriction is also consistent with the data (the prob value of the relevant LR test is 0.22). Model (5) imposes the further restriction $\rho_1=\rho_2$, i.e. the error terms of the two outcome equations are equally highly correlated with the error of the selection equation (any heterogeneity is—to use Heckman’s

²³ We estimate the model using the Stata routine movestay (Lokshin and Sajaia 2004).

terminology—‘inessential’). This restriction is also consistent with the data at conventional significance levels. Finally, model (6) imposes the restriction $\sigma_1 = \sigma_2$, i.e. the errors of the two outcome equations have the same variance. This restriction is decisively rejected by the data.

Table 2: Model specification tests and estimates of implied averted infant deaths

	(1)	(2)	(3)	(4)	(5)	(6)
	No restrictions	$\beta_1 = \beta_2$ $\phi_2 = 0$	$\delta_1 = \delta_2$ $\alpha_1 = \alpha_2$ $\theta_1 = \theta_2$	$\phi_1 = 0$	$\rho_1 = \rho_2$	$\sigma_1 = \sigma_2$
Log-Likelihood	556.25	555.84	551.39	550.64	550.47	509.80
No. parameters to be estimated	29	27	21	20	19	18
LR test statistic vs. previous model		0.826	8.908	1.489	0.341	81.351
Prob value of above LR test		0.662	0.179	0.222	0.559	0.000
LR test statistic vs. model (1)		0.826	9.734	11.223	11.564	92.914
Prob value of above LR test		0.662	0.350	0.997	1.000	0.000
Implied change in infant deaths	-6117	-1294	-11288	-5669	-5133	-5439
t-statistic for change	-0.43	-0.11	-1.95	-1.87	-1.81	-1.85
% change in infant deaths	-13%	-3%	-24%	-12%	-11%	-11%

Notes: .01 - ***; .05 - **, .1 - *. t-statistics in column (5) reflect standard errors adjusted for clustering at the province level.

The most parsimonious model consistent with the data is thus model (5). This specification implies that any benefits accruing from the Bank projects occur not through leveraging increases in the productivity of government spending or producing a faster downward shift in the mortality-expenditure relationship but rather by allowing government spending to increase nationwide. The conclusion about the channels by which the Bank projects affected infant mortality is, however, less clear-cut than the LR tests imply. Table 3 presents parameter estimates for the models numbered (1), (2), (3) and (5) in Table 2. In model (1), provincial government health spending has a significant negative impact on infant mortality but only in project provinces and only in 2002 after

the projects had been (largely) implemented. This is true also of model (2) where the restrictions $\beta_1=\beta_2$ and $\phi_2=0$ are imposed. When the further restrictions in model (3) are also imposed— $\delta_1=\delta_2$, $\alpha_1=\alpha_2$ and $\theta_1=\theta_2$ —the coefficient on government spending interacted with the 2002 dummy drops in absolute size and is no longer significant; by contrast, the coefficient on health spending without any year interaction (i.e. the coefficients β_1 and β_2 , which are constrained to be equal to one another) increases in absolute size and becomes significant. Imposing the further restrictions in moving from model (3) to model (5)—namely $\phi_1=0$ and $\rho_1=\rho_2$ —reduces somewhat the absolute value of β_1 ($=\beta_2$) but it stays significant. The restrictions under model (5), which imply the projects did not leverage any productivity increases, are consistent with the data. Yet in model (2), the coefficient on the interaction between spending and the 2002 dummy is significant, implying the opposite.

Table 3: Parameter estimates for selected models

	(1)		(2)		(3)		(5)	
	coef	t	coef	t	coef	t	coef	t
<i>Nonproject areas</i>								
Wealth quintile no. 2	-0.003	-0.227	-0.002	-0.197	-0.014	-1.547	-0.013	-1.503
Wealth quintile no. 3	-0.023***	-2.812	-0.023***	-2.822	-0.020**	-2.506	-0.020**	-2.446
Wealth quintile no. 4	-0.026***	-3.128	-0.025***	-3.070	-0.024***	-3.523	-0.023***	-3.512
Wealth quintile no. 5	-0.011	-0.634	-0.011	-0.674	-0.019**	-2.090	-0.019**	-2.008
2002 dummy	0.010	0.706	0.001	0.117	0.005	0.676	-0.000	-0.085
Govt. health expend. (GHE)	0.035	0.151	-0.046	-0.296	-0.150**	-2.070	-0.120*	-1.807
GHE x 2002 dummy	-0.147	-0.952						
Constant	0.026	1.575	0.030**	2.375	0.045***	5.192	0.042***	6.435
<i>Project areas</i>								
Wealth quintile no. 2	-0.021*	-1.789	-0.021*	-1.792	-0.014	-1.547	-0.013	-1.503
Wealth quintile no. 3	-0.018	-1.578	-0.018	-1.591	-0.020**	-2.506	-0.020**	-2.446
Wealth quintile no. 4	-0.022**	-2.301	-0.022**	-2.302	-0.024***	-3.523	-0.023***	-3.512
Wealth quintile no. 5	-0.030***	-2.961	-0.030***	-2.949	-0.019**	-2.090	-0.019**	-2.008
2002 dummy	0.029*	1.793	0.023	1.600	0.005	0.676	-0.000	-0.085
Govt. health expend. (GHE)	0.137	0.385	-0.046	-0.296	-0.150**	-2.070	-0.120*	-1.807
GHE x 2002 dummy	-0.716*	-1.861	-0.535**	-2.071	-0.161	-1.229		
Constant	0.040***	2.903	0.045***	3.997	0.045***	5.192	0.042***	6.435
<i>Selection equation</i>								
GDP	-1.564***	-3.387	-1.564***	-3.387	-1.563***	-3.382	-1.564***	-3.383
Revenues as % expenditures	1.759*	1.836	1.759*	1.836	1.760*	1.839	1.759*	1.836
Constant	5.116***	4.474	5.116***	4.474	5.112***	4.468	5.115***	4.468
$\ln \sigma_{\text{nonproject}}$	-2.038***	-23.934	-2.037***	-23.950	-2.035***	-23.800	-2.036***	-23.845
$\ln \sigma_{\text{project}}$	-1.835***	-23.108	-1.835***	-23.102	-1.834***	-23.067	-1.834***	-23.039
$\rho_{\text{nonproject}}$	-0.023	-0.336	-0.021	-0.362	0.050	1.525	0.019	0.810
ρ_{project}	-0.013	-0.198	-0.012	-0.182	0.008	0.107	0.019	0.810

Notes: .01 - ***; .05 - **; .1 - *. t-statistics reflect standard errors adjusted for clustering at the province level. Selection equation also includes variables in outcome equation. Model numbers correspond to those in Table 2.

Thus the evidence is ambiguous on the questions of whether the projects leveraged productivity increases in government spending in respect of infant mortality and hence whether the intrasectoral fungibility resulted in any forgone benefits. What can be done is to produce estimates of the overall effect of the projects on infant survival, without trying to establish which mechanisms were responsible. Table 2 shows the estimated number of infant deaths averted for each model specification. These have been computed using eqn (17). Some clarifications are in order. Given the evidence above concerning intersectoral fungibility, the parameter λ has been set at 1. N_1 is the number of live births in the project provinces.²⁴ $\bar{X}_{13}$ is average government spending on health in the project provinces and has been computed using the data used to estimate the models in the above tables. D is the total amount of donor outlays per annum. However, in deriving eqns (16) and (17), it was implicitly assumed that N_1 and N_2 together comprised the total number of beneficiaries of the donor support. In this exercise, however, N_1 and N_2 refer to a subpopulation, namely those born in the relevant year. This means that D needs to be scaled down accordingly, since other age groups will also have benefited from the Bank credits. Live births in 1999 totaled 1.35 million, equivalent to 1.7% of Vietnam's population of 78 million. Using 1.7% as the scaling factor would, however, seem inappropriate as both projects prioritized the health of (expectant) mothers and children. A figure of 10% was used instead: hence in computing averted infant deaths, we have replaced D in eqn (17) by $0.1D$.

²⁴ This number is for 1999, and has been computed from provincial data on the infant mortality rate and the number of children aged 0. The data are from Vietnam Central Census Steering Committee (2000).

The estimated number of averted deaths range from 1,294 in model (2) to 11,288 in model (3). To put this in context, the number of infant deaths in 1999 in Vietnam was around 48,000, so the estimates range from a 3% reduction to a 24% reduction. The reason for the much higher estimate from model (3) is clear from Table 3, namely that government spending has a sizeable (equal) effect in both project and nonproject provinces (the coefficient β) that is reinforced in the case of the project provinces by a sizeable additional negative effect through the interaction term ϕ . At the same time, the model constrains the intercept to shift up by the same amount in the project and nonproject provinces, i.e. $\theta_1 = \theta_2$. Four of the six models produce an estimated percentage reduction of 11-12%; three of these estimates are statistically significant at the 10% level, one of which is the estimate from model (5), which is our data-consistent parsimonious model. These estimates translate into a reduction in the infant mortality rate of around 4 per 1000 live births. This number is half that obtained from the naïve *diffs-in-diffs* estimate, though that estimate—unlike this—was not significantly different from zero.

VIII. SUMMARY AND CONCLUSIONS

The concern of the paper is with two questions that have been largely unaddressed in the literature on fungibility. How can one estimate the impact of development assistance in the presence of fungibility? And how far does fungibility reduce the benefits of aid? Both are approached from a microeconomic perspective, the setting being a donor wanting to target its efforts on a specific country, and on a specific sector and specific geographic areas within the country. Two types of fungibility are possible: *intersectoral* fungibility, which results in the donor's resources leaking out of the sector, and

intrasectoral fungibility, which results in its resources leaking out of the target areas into the nontarget areas.

The paper begins by showing that given the possibility of both types of fungibility, a traditional differences-in-differences comparing the change in outcomes between the target and nontarget areas before and after the donor's project risks misestimating the benefits of the project. The paper goes on to develop an alternative approach to estimating the benefits of a project in such circumstances that also sheds light on the issue of how far fungibility reduces the benefits of development assistance. The method requires the specification and estimation of a model linking the outcomes of interest to government spending, allowing the donor's project to modify this relationship by increasing the productivity of government spending in the target areas. Intersectoral fungibility reduces the benefit insofar as government spending in the sector where the funds leak to is less productive than the sector where the aid was intended for, while intrasectoral fungibility reduces the benefits of aid insofar as the donor is able to leverage productivity increases in government spending in the target areas.

The paper applies the methods to two contemporaneous World Bank health projects that set out to target assistance on approximately one-half of Vietnam's provinces. There was a clear understanding in this case between the client government and the donor that the government would focus its expenditures and efforts on the provinces the bank was not targeting. For the most part, the Bank reviewers of these projects were aware that such behavior would render the nonproject provinces invalid as a counterfactual; they therefore shied away from a formal estimation of the project's

benefits. The results in section VII of the paper suggest that aid is *not* fungible between Vietnam's health sector and other sectors in the economy, but *is* fungible across provinces within the health sector. The use of diff-in-diffs with micro-data yields an insignificant effect of the projects on infant mortality. By contrast, the use of the methods proposed in the paper yield a statistically significant impact equivalent to a reduction in the infant mortality rate of around 4 per 1000 live births. While the econometric results suggest that there are no costs in this case from intersectoral fungibility (because there is none), they are less clear on the costs of intrasectoral fungibility. The most general model specification finds a statistically significant parameter capturing the degree to which the Bank was successful in leveraging increases in the productivity of government spending in its project provinces. However, the most parsimonious data-consistent model, which includes a variety of parameter restrictions, includes a zero restriction on this parameter. The data are unable, therefore, to determine whether intrasectoral fungibility in this particular pair of projects was associated with a lower-than-feasible reduction in infant mortality.

References

- Basu, A., J. J. Heckman, S. Navarro-Lozano and S. Urzua (2007). "Use of Instrumental Variables in the Presence of Heterogeneity and Self-Selection: An Application to Treatments of Breast Cancer Patients." Health Econ **16**(11): 1133-57.
- Bokhari, F. A., Y. Gai and P. Gottret (2007). "Government health expenditures and health outcomes." Health Econ **16**(3): 257-73.
- Burnside, C. and D. Dollar (2000). "Aid, Policies, and Growth." American Economic Review **90** 4: 847-68.
- Committee for Population Family and Children of Vietnam and ORC Macro (2003). Vietnam Demographic and Health Survey 2002. Calverton, MD, Committee for Population, Family and Children and ORC Macro.
- Devarajan, S., A. S. Rajkumar and V. Swaroop (1999). What does aid to Africa finance? Washington, DC, World Bank. Policy Research Working Paper #2092.
- Feyzioglu, T., V. Swaroop and M. Zhu (1998). "A panel data analysis of the fungibility of foreign aid." World Bank Economic Review **12**(1): 29-58.
- Filmer, D. and L. Pritchett (1999). "The impact of public spending on health: does money matter?" Social Science and Medicine **49**: 1309-1323.
- Filmer, D. and L. Pritchett (2001). "Estimating wealth effects without expenditure data or tears: An application to educational enrollments in states of India." Demography **38**(1): 115-132.
- Fritzen, S. A. (2007). "Legacies of primary health care in an age of health sector reform: Vietnam's commune clinics in transition." Social Science & Medicine **64**(8): 1611-1623.
- Government of Vietnam-Donor Working Group on Public Expenditure (2000). Vietnam - Managing Public Resources Better. Public Expenditure Review. Hanoi, Joint Review of the Government of Vietnam-Donor Working Group on Public Expenditure Review,.
- Heckman, J. J. (1978). "Dummy Endogenous Variables in a Simultaneous Equation System." Econometrica **46** 4: 931-59.
- Heckman, J. J., S. Urzua and E. Vytlacil (2006). "Understanding Instrumental Variables in Models with Essential Heterogeneity." The Review of Economics and Statistics **88**(3): 389-432.
- Lokshin, M. and Z. Sajaia (2004). "Maximum likelihood estimation of endogenous switching regression models." Stata Journal **4**: 282-289.
- Pack, H. and J. R. Pack (1990). "Is Foreign-Aid Fungible - the Case of Indonesia." Economic Journal **100**(399): 188-194.
- Pack, H. and J. R. Pack (1993). "Foreign Aid and the Question of Fungibility." Review of Economics and Statistics **75** 2: 258-65.
- Pettersson, J. (2007a). "Child mortality: Is aid fungibility in pro-poor expenditure sectors decisive?" Review of World Economics **143**(4): 673-693.
- Pettersson, J. (2007b). "Foreign Sectoral Aid Fungibility, Growth and Poverty Reduction." Journal of International Development **19** 8: 1074-98.
- Ravallion, M. (2008). Evaluating Anti-Poverty Programs. Handbook of Development Economics vol. 4. R. Evenson and T. Schultz. Amsterdam, North Holland: 3787-3846.

- Swaroop, V., S. Jha and A. S. Rajkumar (2000). "Fiscal Effects of Foreign Aid in a Federal System of Governance: The Case of India." Journal of Public Economics **77 3**: 307-30.
- van de Walle, D. and R. Mu (2007). "Fungibility and the Flypaper Effect of Project Aid: Micro-evidence for Vietnam." Journal of Development Economics **84 2**: 667-85.
- Vietnam Central Census Steering Committee (2000). 1999 Population and Housing Census: Sample Results. Hanoi, The Gioi Publishers.
- World Bank (1995). Staff Appraisal Report. Vietnam National Health Support Project. Report No. 14971 Washington DC, World Bank.
- World Bank (2004). Implementation Completion Report. Vietnam Population and Family Health Project. Report No. 28400. Washington DC, World Bank.
- World Bank (2006). Project Performance Assessment Report. Vietnam: Population and Family Health Project. Washington DC, World Bank.
- World Bank (2007). Implementation Completion and Results Report. Vietnam National Health Support Project. Report No. ICR527. Washington DC, World Bank.