

Does Ignoring Heterogeneity in Impacts Distort Project Appraisals? An Experiment for Irrigation in Vietnam

Dominique van de Walle and Dileni Gunewardena

Could the simplifying assumptions made in project appraisal be so far from the truth that the expected benefits of public investments are not realized? Using data for Vietnam, commonly used estimates of the benefits from irrigation investments based on means are compared with impacts assessed through an econometric modeling of marginal returns that allows for household and area heterogeneity using integrated household-level survey data. The simpler method performs well in estimating average benefits nationally but can be misleading for some regions, and, by ignoring heterogeneity, it overestimates gains to the poor and underestimates gains to the rich. At moderate to high cost levels, ignoring heterogeneity in impacts results in enough mistakes to eliminate the net benefits from public investment. When irrigating as little as 3 percent of Vietnam's nonirrigated land, the savings from the more data-intensive method are sufficient to cover the full cost of the extra data required, ignoring other benefits from that data.

The methods used in practice for project appraisal simplify reality in certain respects. Appraisals implicitly assume that any economic losses arising from the use of these methods are of second-order importance. Frequently, appraisals resort to rapid assessments of project benefits, conducting the analysis for a representative project and/or beneficiary. This article focuses on just one of the common simplifying assumptions in project appraisal.¹ We question whether ignoring heterogeneity this way—by looking only at aggregates—biases the estimates of both aggregate benefits and their distribution. Would collecting better data or using better methods make an appreciable difference in the social welfare outcomes of public investments? Would it be worth investing the extra resources needed to make a more thorough assessment? Indeed, could the defi-

Dominique van de Walle is with the Development Research Group at the World Bank and Université des Sciences Sociales, Toulouse; Dileni Gunewardena is with the Department of Economics, University of Peradeniya, Sri Lanka. The authors' e-mail addresses are dvandewalle@worldbank.org and demelgun@slnet.lk. For comments and suggestions, the authors thank Lionel Demery, Gaurav Datt, Shanta Devarajan, Gunnar Eskeland, Jeff Hammer, David Hughart, Jean-Jacques Laffont, Peter Lanjouw, Jennie Litvack, Howard Pack, Vijaya Ramachandran, Martin Ravallion, three anonymous referees, and seminar participants at the International Food Policy Research Institute, the World Bank, and Université des Sciences Sociales, Toulouse.

1. For an overview of project evaluation in theory, see Dreze and Stern (1990).

© 2001 The International Bank for Reconstruction and Development / THE WORLD BANK

ciencies of certain methods be so great as to eliminate the entire social gains of the investments?

We address these questions for a method that is often used for assessing rural infrastructure investments. This method assesses the impact of a project on a representative farm household in a particular geographic area and it is a stylized version of crude but common methods found in practice. We call this the quick-and-dirty (QD) method, reflecting a relatively rapid evaluation procedure based on the calculation of means. We use the QD method to calculate the implied average benefit of irrigation in Vietnam. We compare this value with the marginal benefit estimated by an alternative, more sophisticated econometric method for estimating project impacts on farm profits at the farm household level: the slow-and-clean (SC) method. The SC method comes closer to a well-defined theoretical ideal and is about as sophisticated a method as one would expect to find in a small research project set up for project evaluation.

Comparison of stylized QD and SC methods allows us to estimate the potential gains from using more theoretically sound but costly methods, where those gains are assessed by the same criteria used to assess the projects themselves. In trying to identify the benefits of better data and methods to support project appraisal, we deliberately chose to characterize our QD method by something close to the worst of what is done by practitioners. Comparison of the methods indicates the maximum potential benefits of using better methods. If the maximum gains turn out to be small, it is clear that devoting resources to better evaluation would be wasted. Conversely, large gains would suggest that increased evaluation efforts could have high returns.

Here we are concerned with both project impact on average incomes and impact on the distribution of income. In practice QD methods often appeal to both efficiency and equity criteria for project selection. For example, appraisals of rural infrastructure projects in developing countries often argue that because the project is to be located in a poor area, it will help reduce poverty. However, this may be deceptive. Several factors that are typically hidden by QD methods influence the benefits of physical infrastructure investments. There may be complementarities between physical and human infrastructure such that the returns to an individual household depend in part on its level of human capital (van de Walle 2000). If wealthier households have higher human capital, they may also have higher gains. Returns to irrigation on the family farm may further depend on household size and composition in settings with underdeveloped labor markets. The size of landholdings may also matter, again with obvious potential skewness of benefits. In our setting—rural Vietnam in 1992–93—many such household characteristics are likely to be exogenous to the project or not influenced by it over an appreciable period of time. Hence, project benefits will vary with these characteristics.² We aim to see whether a project analysis that ignores

2. This would not be the case in settings with more flexible markets.

these sources of heterogeneity might seriously misinform policy conclusions about the impact of public investments on poverty.

Section I outlines the theoretical ideal and the principles underlying the SC and QD methods. The section briefly discusses the setting and the data that we use to implement the methods. Section II compares and contrasts results obtained by the alternative approaches including implications for distributional assessments, for project selection, and for the net social gains from public investments. Section III concludes.

I. METHODS

This section presents the theoretical ideal and describes two approximations—the SC and QD methods. The section describes the setting and data for our analysis and explores whether we can predict how the choice of methods will affect the results of project appraisal.

The Ideal

An important input to the appraisal of irrigation projects is assumed to be an estimate of the gain in farm profit from irrigating given amounts of previously unirrigated land.³ The ideal method would start with a general specification of the profit function for a farm household. We measure farm household profit from crop production by total revenue minus total production costs, which we call net crop income. It is assumed to be a function of output and input prices (p), the amount of annual nonirrigated (L^N) and irrigated (L^I) cropland, and other relevant variables (z). The generic profit function is

$$(1) \quad \pi_j = \pi(p_j, L_j^N, L_j^I, z_j).$$

Equation 1 gives the maximum profit received by the j th household for $j = 1, \dots, n$. The vector z includes other fixed factors and parameters of the production function used by the j th household.

In the case where a complete set of perfect markets exists for all farm outputs and inputs, variables influencing consumption decisions, such as the prices of consumer goods and the size and demographic composition of the household, would not alter the maximum profit from farming. However, when markets are incomplete—so that the conditions required for separability of production and consumption decisions do not hold—such variables will spill over into production decisions (Strauss 1986). For example, in Vietnam, rural labor markets are thin or nonexistent, reflecting the dual effects of the past socialist organization of rural production and reliance on subsistence farming as well as possibly high supervision costs and limited mobility in the early stages of transition. Variables such as family size and composition influence the amount of labor available for

3. We can abstract from whether the investments are public or private because, either way, an estimate of the gains in farm yields is needed.

farming and, hence, maximum profits. Therefore, z may include factors besides the parameters of the farm household's production function. The specification in equation 1 can thus be made general enough to encompass market effects of credit or labor market failure.

Now consider a project that involves irrigating amounts ΔL_j^I of previously unirrigated land for each of n households (possibly zero land for some). The benefit to the j th farm household is given by the increment to its profits from farming:

$$(2) \quad B_j = \pi(p_j, L_j^N - \Delta L_j^I, L_j^I + \Delta L_j^I, z_j) - \pi(p_j, L_j^N, L_j^I, z_j).$$

One would then calculate the average benefit ($\Sigma B_j / n$) or some distribution-weighted benefit. In the special case in which one unit of land is irrigated, it is useful to define the *marginal benefit function* as

$$(3) \quad MB_j = \pi(p_j, L_j^N - 1, L_j^I + 1, z_j) - \pi(p_j, L_j^N, L_j^I, z_j).$$

If we knew the profit function, the task of calculating project benefits would be complete. This section describes two approximations to this ideal, one of which—the sc method—is undoubtedly more accurate than the other but is still an approximation. But first we need to describe some key features of our data.

Setting and Data

We test irrigation project appraisal methods using data from the Vietnam Living Standards Survey (vnlss) of 1992–93. This is a nationally representative, high-quality, household consumption survey covering a sample of 4,800 households.⁴ The data include detailed coverage of agricultural production and incomes that allows us to construct a comprehensive measure of annual crop incomes net of all production costs. The survey also collects detailed information on land assets, including quality of plots, and other inputs to crop production, including family labor inputs. For the welfare measure, we use total household per-capita expenditures (including the imputed value of consumption from own production), appropriately deflated to allow for spatial cost of living differences.⁵

Vietnam is a largely agricultural economy. In 1992–93, 84 percent of the rural labor force aged six years or older claimed agriculture as their primary occupation. A majority of households are engaged in small-scale subsistence farming, relying almost exclusively on household labor and traditional inputs. According to official sources, corroborated by the 1992–93 vnlss data, about half of the country's arable cropland is under irrigation (Vu and Taillard 1993). Irrigated land is defined by the vnlss to include land benefiting from any kind of water

4. A detailed description of the data set is given in Glewwe (1994).

5. We use a Laspeyres spatial price index covering the rural and urban areas of Vietnam's seven regions. It is constructed based on prices collected by the vnlss in the 120 surveyed rural communes. We also use urban price data from official sources. Expenditures are likewise deflated to January 1993 prices using the government-constructed monthly consumer price index. Further details are available in Glewwe (1994).

management system—such as pumps—that prevents flooding or drought. It is generally agreed that there is great potential for an expansion of the area served by new irrigation infrastructure as well as by the rehabilitation of long non-functioning irrigation networks (Barker 1994).⁶ Such investments have not been undertaken due to the combination of historical factors, such as war, highly constrained public budgets, and lack of access to credit.

The current distribution of access to cropland and irrigation varies across regions, but much less so within regions due to past land reform. In general, land endowments are distributed relatively equally in the north. They are distributed less equally in the south where, on average, the poor have access to less than half the amount of land compared with the nonpoor (van de Walle 1996). The existing distribution of irrigation is somewhat more equal than that of land. Given the current distribution, it cannot be argued that investment in irrigation would necessarily benefit the poor more than the rich.

Although Vietnam has been undergoing reform since 1986, markets were still relatively underdeveloped in 1992–93, when land remained under state ownership and land markets were illegal. Fieldwork suggests that labor markets, though generally thin everywhere, did not exist at all in some parts of the country. Mobility was severely restricted because access to social services and transactions to do with land, housing, and credit were officially linked to an individual's residency permit and new permits were not easily acquired (UNDP 1998). Using the same data as this article, van de Walle (2000) found evidence that household demographics and human capital exert considerable influence on farm household crop incomes. This reflects an environment where, for the most part, households do not have the option of hiring in or out workers and/or skills.

A Slow and Clean Approach Incorporating Heterogeneity

The sc method works by assuming a functional form for the profit function that is estimated by regression methods using suitable microdata—in this case the 1992–93 VNLSS. The chosen specification allows a number of variables—including land, demographic and education variables, and regional dummy variables for Vietnam's seven regions—to have direct effects on the marginal returns from irrigated and nonirrigated cropland. For the sc method used here, the profit function is assumed to have the following parameterization:⁷

$$(4) \quad \pi_j = \alpha + \beta_j^N L_j^N + \beta_j^I L_j^I + \gamma z_j + \delta d_j + \varepsilon_j$$

where

6. Note that “new irrigation infrastructure” could include both large-scale projects and smaller investments, such as bore holes or diesel pumps. We do not preclude the latter.

7. Due to labor valuation difficulties, profit is not net of household labor inputs. Our profit function is thus literally more an income function. This difference does not affect the calculation of household consumption gains. On the sensitivity of the estimates of farm profits, see van de Walle (1996).

$$(5) \quad \beta_j^N = b_0^N + b_1^N d_j + b_2^N z_j + b_3^N L_j^N$$

and

$$(6) \quad \beta_j^I = b_0^I + b_1^I d_j + b_2^I z_j + b_3^I L_j^I$$

and where d is a vector of regional dummy variables that are assumed to fully capture the variation in prices faced in each region.⁸ The error term ϵ_j is assumed to be independently and identically normally distributed.

The country's regions are made up of provinces, districts, and, at the lowest level, communes. Dummy variables for 119 out of the 120 sampled communes are included in the intercept of the profit function (d in equation 4) to capture variations in prices and any other spatial, cross-commune variations in omitted or fixed factors, such as land and soil quality. Thus, prices of outputs and variable inputs are assumed to vary between but not within communes. The commune dummies pick up the effects of geographical, social, and physical infrastructure variations at the commune level. In addition, we collapse the commune dummies into seven regional dummy variables (d in equations 5 and 6) and interact them with irrigated, nonirrigated, and other types of land, thus permitting regional effects on the marginal returns to land.

By allowing nonlinearity in land and interaction effects with other variables, the above specification is a reasonably flexible functional form for the present purposes. However, our SC method could possibly be improved. For example, as is common in the literature and indeed is true of most rapid assessment methods of project appraisal, our SC method ignores general equilibrium and dynamic welfare effects. There is clearly a continuum of QD as well as SC methods for estimating benefits. As we have explained, our purpose here is to see whether policy conclusions and choices are significantly altered when comparing a "worst case" with a "best case" common evaluation practice.

The vector z includes other land in agricultural production; land tenure variables; education, health, and demographic variables; and location-specific, agro-ecological variables. As discussed, a range of variables is included in z to capture characteristics specific to a transition economy in which markets are still underdeveloped. In other settings, there would be concerns about possible endogeneity of household characteristics. However, thin or missing markets allow us to assume that these characteristics are predetermined and that we are looking at benefits over a period of time that do not allow households to change their characteristics. There will be long lags before such behavioral responses, if they occur, bring higher returns. In the present setting, it is reasonable to assume that the irrigation project does not change such household characteristics as education and household demographics for the period over which we measure gains.

We use ordinary least squares on a sample of the 3,049 farm households in the data set (including some urban farm households). Table 1 defines the vari-

8. Because the article focuses on irrigated and nonirrigated annual crop land, we show the marginal returns only for these land types. However, the specification allows for the same variables to interact with other land types that are contained in the z vector.

ables and gives sample statistics. Table 2 presents the regression results.⁹ The equation's explanatory power, as reflected in an adjusted R^2 of 0.58, is unusually high for a regression on cross-sectional household-level data. Irrigated and nonirrigated cropland are both found to have high but diminishing impacts on crop income. Household size has a positive effect, as does its interaction with many land variables. One notable exception is size interacted with irrigated land, which has a pronounced negative effect on crop income. This coefficient is probably picking up a tendency for larger irrigated farms to not be constrained by family size; the more irrigated land a household has, the less it is dependent on family labor for farm household production.

Education has strong effects. The primary education of the household head is convex in its impact, suggesting increasing returns to schooling. Interaction effects between primary education variables and irrigated land tend to be large and positive. There are also significant commune fixed effects and spatial differences in the effects of irrigated and nonirrigated land and other land types.¹⁰

From the regression model, we derive a marginal benefit function for irrigation that allows for heterogeneity across households. The marginal benefit from irrigating one unit of previously unirrigated land is given by the difference in the derivatives of the regression function with respect to irrigated land and non-irrigated land:

$$(7) \quad MB_j = (b_0^I - b_0^N) + (b_1^I - b_1^N) d_j + (b_2^I - b_2^N) z_j + 2(b_3^I L_j^I - b_3^N L_j^N).$$

Table 3 provides the estimates of equation 7. As expected, marginal benefits from irrigation decrease as irrigated area increases, and they increase with the amount of unirrigated land. The results show the strong influence of the household's demographic and education endowments on the gains from moving one unit of nonirrigated land into irrigation. We evaluate this function at the sample variable means to calculate the mean marginal benefit.

A Quick and Dirty Approach

The sc method is demanding in a number of respects. Special microdata are required to capture heterogeneity—namely, an integrated household survey that contains information for all the relevant variables for the sampled households.

9. We tried a number of functional form specifications, including linear, semi-log, and double log forms, with and without quadratics in land and education (for further discussion, see van de Walle 1996). The presented linear model with quadratics in land and education variables was found to perform best. Full regression details are available from the authors.

10. Biases due to endogenous explanatory or omitted variables that are correlated with included variables is a potential issue here. However, as a result of past land reform and distribution processes, land and irrigation inputs can reasonably be treated as exogenous at the household level. Possible omitted variable bias is more worrying. The regressions control for omitted between-commune variance through the commune dummies. But there may also be latent heterogeneity in, say, land or soil quality within communes. Still, including land quality in the regression did not reveal any sign of such bias (van de Walle 1996). This could be more of an issue in the south—salinity and acidity are common problems in the Mekong Delta—although they were not observed in the VNLSS data.

TABLE 1. Variable Definitions and Summary Data

Variable	Definition	Mean	Standard deviation
<i>cropinc</i>	Net household crop income, 1993 dong	2,282,069	2,391,173
<i>sick</i>	Dummy for household member being sick in last year	0.933	0.250
<i>sexhbb</i>	Gender of household head	0.809	0.393
<i>hbsize</i>	Size of the household	5.033	1.992
<i>prop06</i>	Proportion of household members 6 years old and younger	0.156	0.176
<i>prop716</i>	Proportion of household members 7–16 years old	0.213	0.204
<i>pfadlt</i>	Proportion of female adults (over age 16)	0.327	0.169
<i>pmadlt</i>	Proportion of male adults (over age 16)	0.282	0.160
<i>hed1</i>	Years of primary education of household head	4.379	1.114
<i>hed2</i>	Years of post-primary education of household head	2.513	2.842
<i>oed1</i>	Years of primary education of other adult household members (over age 16)	6.872	5.372
<i>oed2</i>	Years of post-primary education of other adult household members (over age 16)	4.111	5.287
<i>irrigated</i>	Irrigated annual crop land area (m ²)	2,267.58	3,997.50
<i>nonirrigated</i>	Nonirrigated annual crop land area (m ²)	2,605.92	5,632.34
<i>perennial</i>	Perennial land area (m ²)	678.43	2,169.50
<i>forest</i>	Forest land area (m ²)	279.22	1,970.98
<i>waterland</i>	Water surface land area (m ²)	122.89	1,203.53
<i>otherland</i>	Other land area (m ²)	217.50	2,106.11
<i>proplt</i>	Proportion of annual land that is long-term	0.20	0.380
<i>propauct</i>	Proportion of annual land that is auctioned	0.023	0.092
<i>proppriv</i>	Proportion of annual land that is private	0.227	0.341
<i>propshare</i>	Proportion of annual land sharecropped/rented	0.043	0.165
<i>propall</i>	Proportion of annual land that is allocated	0.507	0.431
<i>urban</i>	Dummy variable for urban residence	0.057	0.231
<i>nu</i>	Dummy variable for the Northern Uplands region	0.183	0.387
<i>rr</i>	Dummy variable for the Red River Delta region	0.275	0.447
<i>nc</i>	Dummy variable for the North Coast region	0.178	0.383
<i>cc</i>	Dummy variable for the Central Coast region	0.090	0.286
<i>cb</i>	Dummy variable for the Central Highlands region	0.020	0.139
<i>mk</i>	Dummy variable for the Mekong River Delta region	0.20	0.40

Source: Glewwe, 1992–93.

TABLE 2. Regression Results for Crop Incomes

Variable	Coefficient	<i>t</i> -ratio
<i>urban</i>	1,093,640	1.08
<i>sick</i>	-318,465.4	2.59
<i>hbsize</i>	81,451.7	2.12
<i>prop06</i>	-586,514.9	1.18
<i>hed1</i>	-468,646.3	2.69
<i>hed1*hed1</i>	67,862.6	2.66
<i>oed1*oed1</i>	-1,932.7	2.48
<i>oed2</i>	21,023.8	1.24
<i>irrigated</i>	352.40	4.27
<i>irr*irr</i>	-0.0030	4.42
<i>nonirrigated</i>	238.40	3.81
<i>nonirrig*nonirrig</i>	-0.0036	9.60
<i>perennial</i>	-277.04	1.73
<i>perennial*perennial</i>	-0.0097	6.43
<i>forest</i>	-372.88	1.35
<i>forest*forest</i>	-0.0026	1.38
<i>waterland*waterland</i>	-0.0401	3.80
<i>otherland</i>	-611.27	1.22
<i>otherland*otherland</i>	-0.0024	1.32
<i>propauct</i>	1,116,555	2.54
<i>proppriv</i>	325,505.5	1.49
<i>propall</i>	470,198.4	2.10
<i>hed1*irrigated</i>	47.87	6.06
<i>hed1*otherland</i>	-113.39	2.58
<i>hed2*irrigated</i>	-6.46	1.53
<i>hed2*perennial</i>	21.90	2.53
<i>hed2*forest</i>	23.01	1.62
<i>hed2*waterland</i>	72.61	1.58
<i>hed2*otherland</i>	33.45	1.51
<i>oed1*irrigated</i>	20.74	8.03
<i>oed1*nonirrigated</i>	7.27	3.38
<i>oed1*perennial</i>	5.42	1.21
<i>oed1*forest</i>	-21.37	1.72
<i>oed1*otherland</i>	-49.20	4.66
<i>oed2*irrigated</i>	-4.179	2.18
<i>oed2*nonirrigated</i>	1.741	1.04
<i>oed2*perennial</i>	-10.914	2.65
<i>oed2*otherland</i>	33.814	4.00
<i>hbsize*irrigated</i>	-35.991	6.94
<i>hbsize*nonirrigated</i>	4.639	1.13
<i>hbsize*perennial</i>	52.933	4.62
<i>hbsize*forest</i>	37.473	1.68
<i>hbsize*otherland</i>	79.081	2.62
<i>pfadlt*irrigated</i>	-176.63	2.16
<i>pfadlt*nonirrigated</i>	-137.02	2.07
<i>pfadlt*perennial</i>	610.10	3.33
<i>pfadlt*otherland</i>	1941.44	4.18
<i>pmadlt*irrigated</i>	-162.40	1.71
<i>pmadlt*perennial</i>	289.39	1.92
<i>prop716*irrigated</i>	155.85	2.03

(continued)

TABLE 2. (continued)

Variable	Coefficient	<i>t</i> -ratio
<i>rr*irrigated</i>	271.75	4.06
<i>rr*forest</i>	135.35	1.03
<i>mk*irrigated</i>	-67.71	1.94
<i>mk*perennial</i>	-158.92	2.94
<i>nu*irrigated</i>	255.81	3.47
<i>nu*perennial</i>	-199.58	2.27
<i>nu*otherland</i>	434.86	1.20
<i>nc*perennial</i>	-218.53	3.02
<i>nc*otherland</i>	528.01	1.39
<i>cc*irrigated</i>	-203.38	3.63
<i>cc*nonirrigated</i>	-152.25	2.62
<i>cc*perennial</i>	-228.68	1.26
<i>ch*irrigated</i>	-973.79	1.66
<i>ch*nonirrigated</i>	-134.37	2.65
<i>ch*perennial</i>	310.57	4.37
<i>ch*waterland</i>	5,195.78	1.31

Number of observations: 3,049

$F(233, 2815)$: 19.06

Prob > F : 0.0000

R^2 : 0.6120

Adjusted R^2 : 0.5799

Root MSE: 1.5e+06

Notes: See table 1 for variable definitions. For ease of presentation we have not reported results for the 119 commune dummies, many of which are significant. The table also omits regressors with *t*-ratios less than 1. The model contained the following variables: demographic composition variables, *prop716*, *pfadlt*, *pmadlt*, and interactions with land variables; education variables *hed2*, *hed2*², *oed1*, *oed2*², and interactions with land; land variable *waterland* and interactions between types of land and regions; and *proplt* and *propshare*.

Source: Authors' calculations.

The sc method requires econometric modeling. Without such data and methods (or the resources to obtain them), the project appraiser has little option but to do a rapid assessment using less than ideal data to evaluate the impact of expansion of irrigation infrastructure. The essence of the QD method is to estimate the marginal benefit function using simple averages that can be readily calculated in the field or using simple preexisting (nonintegrated) surveys.

To guide our characterization of the QD approach, it is useful to look at what is done in practice. In general the aim is to assess the income gain over preproject rain-fed crop incomes for a representative farmer and a given farm size. The literature on the economic evaluation of irrigation projects emphasizes the importance of estimating quantity changes based on the budgets of farms with and without the project for different sizes and types of farms (for example, Bergmann

TABLE 3. Marginal Benefit Function for Irrigation

Variable	Coefficient	<i>t</i> -ratio
Intercept	114.0	1.24
Irrigated land	-0.006	4.42
Nonirrigated land	0.0072	9.60
Household head's primary education	47.87	5.56
Household head's other education	-5.131	1.07
Other adult's primary education	13.472	4.69
Other adult's other education	-5.92	2.57
Household size	-40.63	7.04
Proportion of female adults	-39.61	0.44
Proportion of male adults	-100.14	1.00
Proportion of children 7-16 years old	148.1	1.94
Red River	315.81	3.05
Mekong Delta	-72.08	3.44
Northern Uplands	274.64	1.63
North Coast	98.67	1.14
Central Coast	-51.13	0.75
Central Highlands	-839.42	1.41

Source: Authors' calculations.

and Boussard 1976, Brown 1979, Gittinger 1982, Londero 1987, OECD 1985). Yet, actual practice in project appraisals exhibits considerable variation and is often far from ideal. Project appraisal staff typically go to a target area, observe the amount of land that can be irrigated in the catchment area, and estimate the incomes of nonirrigated farms. Predictions of farm output gains from the irrigation project are made partly on the basis of such field observations and often draw on an assumed model of one or more seemingly representative farm households. Sometimes province-level statistics on cropping patterns, intensities, and yields are employed; sometimes estimates are based solely on the field visit.

Such methods appear to be the norm in the practice of international development banks, including the World Bank. However, there is heterogeneity in the quality of the inputs. In reviewing the World Bank's recent irrigation project appraisals, we found that it is often hard to figure out exactly what has been done based on the available documentation.¹¹ Some of the project analyses reviewed appear to have used finely disaggregated assessments of output gains by geographical area, by type of crop, or by allowing for heterogeneity in farm size through using representative farm models. But in the majority of cases, we found methods that do not appear to have allowed for heterogeneity and that tend to assume that farmers will benefit equally. We did not find any evaluations that accounted for farm household-level characteristics likely to influence marginal benefits, such as human capital or household size.

11. We reviewed each of the 19 irrigation project appraisal reports prepared by the World Bank since 1992.

Here we propose a QD method that is probably close to a “worst-case” characterization of the methods found in practice. Appraisals often do something better than our QD method. We deliberately chose a benchmark that requires minimal information to implement and is close to one extreme of the range of common practice. The gains from using the SC relative to our QD aim to measure the maximum potential benefits and provide a clear vantage point for judging the results.

In implementing the QD method, we take advantage of the availability of farm-level data on crop incomes from the VNLSS consumption survey. Although we use household survey data to carry out this exercise, that is solely a matter of convenience—there is nothing inherent in our procedure that requires such data. We do not use the integrated nature of the survey (whereby a wide range of different types of data are obtained for the same sample). Rather we use the survey to estimate simple means, such as those from special purpose surveys or field trips. The same approach could be enacted through a rapid assessment survey of a project area or by drawing on data collected by a small agricultural survey.

There is an advantage to using the same survey for calibrating both methods—it allows us to control for differences due to sampling. We may otherwise find differences in the results that are due to nothing more than sampling errors. By the same token, it could be argued that basing the QD estimates on a statistically sound household survey renders them less “dirty” than the typical rapid appraisal estimates using less rigorous sampling methods.

Our aim is to calculate the difference in the value of crop incomes net of costs per area of irrigated land versus the same area of nonirrigated land. This difference is a measure of the average benefit from irrigation allowing for any difference in production costs associated with a change in irrigation.

We use the survey data to approximate what the project appraisal would do in the field. The appraiser is unlikely to pay attention to nonfarmers or farmers producing crops that do not require irrigation. This leads us to estimate the mean over a restricted sample of the survey households, including only households that are primarily engaged in the production of rice, other food crops, or annual industrial food crops—typically the major users of irrigation. We restrict the sample to households whose income from these sources comprises 90 percent or more of their total crop income. The excluded households have a greater dependence on income from perennial industrial crops, fruit, and forest tree crops.

It is unlikely that a rapid field appraisal would be able to identify exactly households that have only irrigated (or nonirrigated) land. We therefore allow for some probable margin of error and further limit the sample to households that have 90 percent or more irrigated or nonirrigated land (as opposed to 100 percent) and calculate mean net crop incomes for these groups. The difference between these amounts expressed per unit of irrigated and nonirrigated land gives us our measure of the average benefit. We calculate average benefits for the national level as well as for regional subsamples. The latter is done for six of Vietnam’s seven re-

gions—excluding the Central Highlands, where there is very little irrigation and no households with 90 percent or more of their land under irrigation.¹²

In trying to emulate the approach that a rapid appraisal might adopt, we have made a number of choices and assumptions in calculating the average benefits from irrigation as described above. To check the sensitivity of the estimates to our choices, we also calculate the means under alternative assumptions. We experiment by increasing (decreasing) the sample to include households with a smaller (larger) share of crop income from rice, other food crops, and annual industrial food crops and by using more or less precision in defining the sample of only irrigated and nonirrigated land. The results indicate reasonably similar magnitudes.

The QD method just described has a number of obvious limitations. These include the neglect of heterogeneity across households and of behavioral effects. For example, the household's level of human capital and its size and demographic composition may influence the returns to irrigation infrastructure. Impacts will then vary according to how such characteristics vary across households. Furthermore, these characteristics will not be changed by the irrigation project, so they should be controlled for when assessing project benefits.

As discussed above, there is diversity in the amount of effort put into estimating quantity changes in practice. Our QD method characterizes worst practice. Some project appraisals do better, such as by allowing for some heterogeneity across regions and farms. But many are unlikely to do any better at a similar level of aggregation. The comparison of our SC and QD methods allows us to judge how much extra effort in evaluation is warranted.

Can We Predict What Difference the Choice of Method Makes?

Before we compare empirical results from the two methods, we consider whether something can be said a priori about the comparison. Can we expect the QD method to over- or underestimate the marginal benefits? Consideration of two stylized cases is sufficient to establish that there can be no theoretical presumption as to which method will indicate higher benefits.

Consider the relationship between profits and irrigated and nonirrigated land holding under the following assumptions:

- *Case 1.* The function relating profits to the amount of irrigated land is strictly concave (marginal profit declines as land increases throughout), whereas that of the function relating profit to nonirrigated land is linear.
- *Case 2.* The reverse holds; profit is linear in irrigated land but concave in nonirrigated land.

Let us also assume that profits are zero when there is no land. In case 1, both the QD and SC methods give exactly the same estimate of the lost profit from having

12. Vietnam's seven regions include the Northern Uplands, Red River, North Coast, and Central Coast in the North and the South East, Mekong Delta, and excluded Central Highlands in the South.

one unit less nonirrigated land. But the QD method overestimates the marginal gain from an extra unit of irrigated land (because the concavity implies that the average profit per unit of land will be higher than the marginal profit). Thus, in case 1, the QD method overestimates the gain as assessed by the SC method. The ordering reverses in case 2 because the QD method overestimates the loss from one unit less of nonirrigated land, and the two methods agree on the gain from extra irrigated land.

It is possible to construct a number of other examples that relax these assumptions and exhibit similar ambiguity. The above argument is sufficient, however, to illustrate that there can be no general supposition as to which of the two methods will give a higher estimate of the gains from irrigation. That is an empirical question to which we now turn.

II. QUICK AND DIRTY VERSUS SLOW AND CLEAN

This section turns to the empirical implementation of the two methods described above and attempts to assess the gains from the SC method relative to the QD method and how those gains are distributed.

Comparing Mean and Marginal Benefits

For the national sample, the QD method gives a mean benefit of 316 dongs per year per square meter of irrigated land. The mean marginal benefit calculated by the SC method is not far off at 304 dongs.¹³ Given the same database, the means are unlikely to be very different. However, on disaggregating to the regional level, we find both larger variation between the estimates using the two methods for a particular region and a striking variation in the mean regional gains whatever the method employed. Table 4 presents the regional benefit estimates.¹⁴

13. In 1992–93, 10,000 dongs were roughly equal to US\$1. To check for sensitivity of the QD measure to our assumptions, we increased (decreased) the sample to include households with a smaller (larger) share of crop incomes from rice, other food crops, and annual industrial food crops and used more or less precision in defining the sample of only irrigated and nonirrigated land. We kept the assumptions about crop income shares but made four adjustments. First, we found a national mean of 306 dongs per square meter by assuming 100 percent (instead of 90 percent) precision in defining the sample of households with only irrigated or nonirrigated land. Second, assuming less precision (75 percent) in defining households with only irrigated or nonirrigated land gave an average benefit of 320 dongs per square meter. Keeping the assumption of 90 percent precision in defining households with only irrigated or nonirrigated land but changing the sample to include those with a larger or smaller share of income from rice, other foods, and annual industrial crops also yielded results of similar magnitude. Third, for households where 100 percent of crop income is from these crops, the mean benefit was 345 dongs per square meter. Fourth, for households where 75 percent or more of crop income is from these crops, the mean benefit was 310 dongs per square meter. Thus, under all assumptions, the national QD estimates slightly overestimate the gains from irrigation.

14. The analysis does not hinge on whether the SC and QD means are significantly different from one another. Even if one could not reject the null hypothesis that the means are the same, it would still be possible to have radically different policy conclusions and large losses from project appraisals based on QD point estimates. Nonetheless, using a conventional *t*-test, at the 5 percent level of confidence one can reject the null that the means are the same for all but the South East region.

There is clearly a diverse regional dimension. As one would expect, it is more accurate to use the QD regional estimates than the QD national mean. Figure 1 plots the SC marginal benefits (on the vertical axis) against the QD mean benefits at the household level. The figure provides an overall summary of our results. (Recall that the SC estimates are household-specific and the QD estimates are the same for all households within a specific region. The figure also identifies the regional means.) The 45-degree line indicates the line of equality between the estimates. The horizontal distance between the line of equality and the SC regional means indicates the amount by which the QD mean over- or underestimates the SC mean. This distance is greater for regions in the North (especially in the Northern Uplands and in the Red River region) than in the South (South East and Mekong Delta). In each region, the QD method both over- and underestimates the SC method. The scatter plot of SC estimates indicates a wide distribution of household-level marginal benefits within each region that is not captured by the QD method.

Distributional Implications

We have seen that there is much variation in mean and marginal benefits across regions. Our next question concerns how much the SC estimates vary with per-capita expenditures and what that means for the QD approximation. In figure 2, we plot the line of best fit estimated using nonparametric methods for the QD measures (which reflect only regional variation) across per-capita expenditures.¹⁵ We overlay this line with the line of best fit of the household-level measures estimated by the SC method against per-capita expenditures. The figure displays a positive association of benefits with per-capita expenditure. In general, there is a tendency for the QD method to overestimate (underestimate) benefits to households with low (high) per-capita expenditure. This emphasizes the main distributional implication: the regional QD method tends to overestimate benefits for the poor and underestimate them for the better-off.

The distributional results partly reflect the influence of household-specific characteristics as shown by the marginal benefit function (table 3). In addition to an important regional dimension, the marginal benefits from irrigated land are significantly affected by the household's size and its level of primary education. Household size has a negative effect on the marginal benefit. Primary education has a positive influence on marginal returns. The survey data also show that larger and less-well-educated households tend to be poorer in terms of per-capita expenditures (World Bank 1995). Thus, in the aggregate, both variables tend to lower marginal benefits for less-well-off households and to raise them for households at higher per-capita expenditure levels. Table 3 shows that, even allowing for regional differences, there are household-specific characteristics that influence the potential benefits from irrigation. These factors are missed by the QD method.

15. The estimates are obtained using Cleveland's (1979) locally weighted, smoothed scatter plots—a method of nonparametric estimation that is of the nearest-neighborhood type.

TABLE 4. Benefits by Region (dongs/year/m²)

Project appraisal method	Northern Uplands	Red River	North Coast	Central Coast	South East	Mekong Delta
(1) Slow and clean	449	488	272	102	175	100
(2) Quick and dirty	541	392	342	48	184	129
(2)-(1)	92	-96	70	-54	9	29

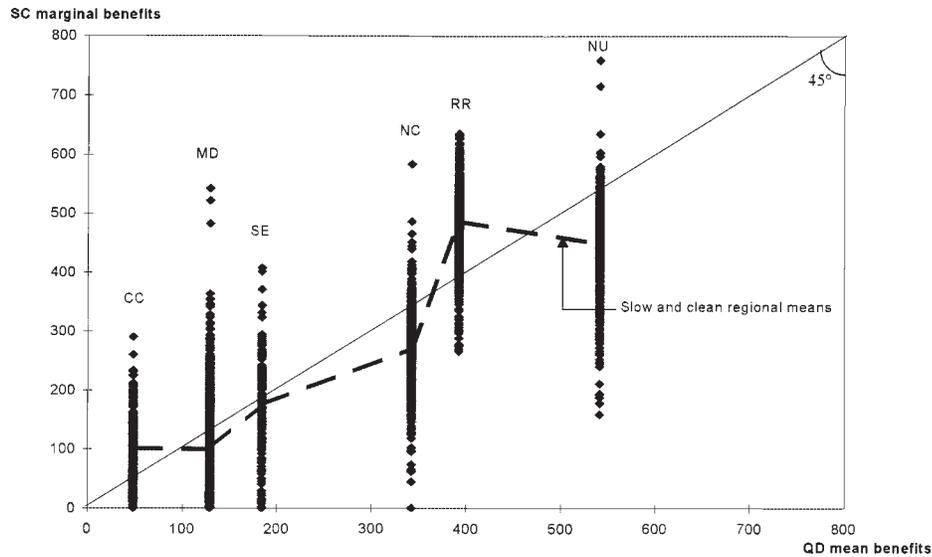
Source: Authors' calculations.

Distributional Impacts of an Irrigation Expansion

So far we have explored the marginal and mean benefits associated with a small change. We have shown that these benefits vary regionally and with household expenditures. Together these factors have distributional implications. However, in determining the distributional impact of a policy of irrigation infrastructure investment, the existing distribution of land and access to irrigation will influence the outcomes. Farms with larger amounts of nonirrigated land are those best positioned to gain from such a policy, but we have seen that the level of their gains will depend on a number of factors not connected to their land assets.

To examine the distributional implications of irrigation infrastructure investments, we simulate a policy of a 10 percent expansion in the area of currently

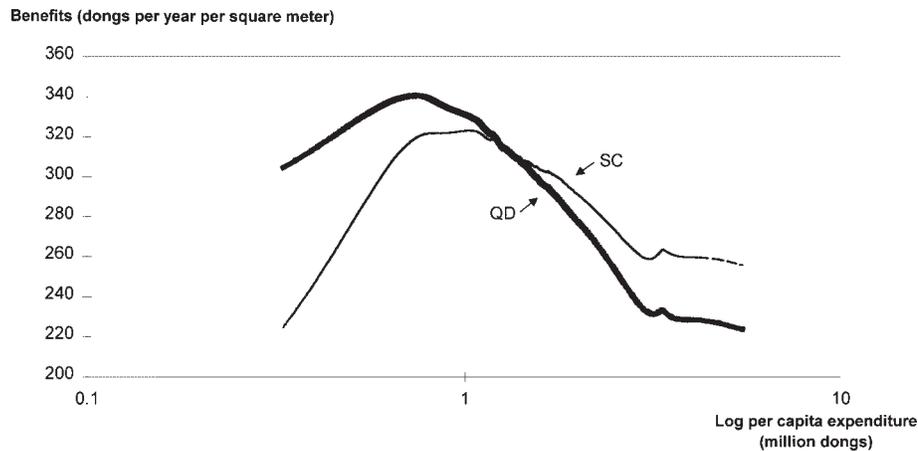
FIGURE 1. A Comparison of the Slow-and-Clean and Quick-and-Dirty Methods of Project Appraisal



Notes: SC marginal benefits and QD mean benefits are measured in dongs per year per square meter of irrigated land. Values are given for the following regions: Central Coast (CC), Mekong Delta (MD), South East (SE), North Coast (NC), Red River (RR), and Northern Uplands (NU). In 1992-93, 10,000 dongs were roughly equal to US\$1.

Source: Authors' calculations.

FIGURE 2. Distribution of Benefits Estimated Using the Slow-and-Clean and Quick-and-Dirty Methods



Notes: QD refers to the quick-and-dirty method and SC to the slow-and-clean method of project appraisal. In 1992–93, 10,000 dongs were roughly equal to US\$1.

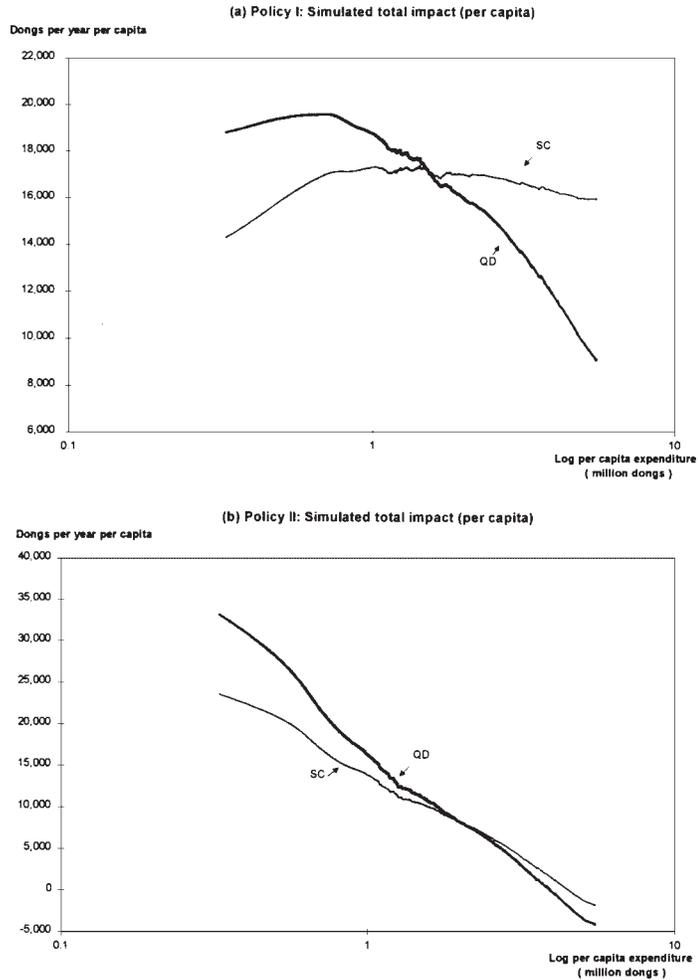
Source: Authors' calculations.

nonirrigated land into irrigation. The distribution across households of that expansion is enacted under two scenarios. Policy I simply increases the amount of land covered by irrigation similarly for each household subject only to feasibility as determined by current land and irrigation holdings and the expansion being considered. Policy II is explicitly pro-poor in that it distributes the 10 percent irrigation expansion only to households with low per-capita land holdings—namely, less than 610 square meters per person.¹⁶

Figure 3 presents the results of these simulations. The impact of irrigation expansion is simply the (household-level) benefit per square meter times the (household-level) increase in irrigated land. We present the results in terms of the lines of best fit for total impacts per capita and household impacts as a percentage of household expenditure. Although the distribution of total impacts under policy I has a slight inverted U distribution (figure 3a), the more explicitly poverty-targeting policy II displays a negative association of impacts with per-capita expenditure (figure 3b). However, both policies are inequality reducing, in that the impact as a percentage of household expenditure declines as per-capita expenditure increases (figures 3c and 3d). How well does the QD-based distribution approximate the SC-based distribution of total impacts? As with the estimates of benefit per square meter (figure 2), the QD measure overestimates the total impacts at the lower end of the welfare distribution and underestimates them at the upper end of the distribution.

16. Given the existing distribution of land, this policy extends irrigation to the nonirrigated land of all households with landholdings less than 610 square meters per person.

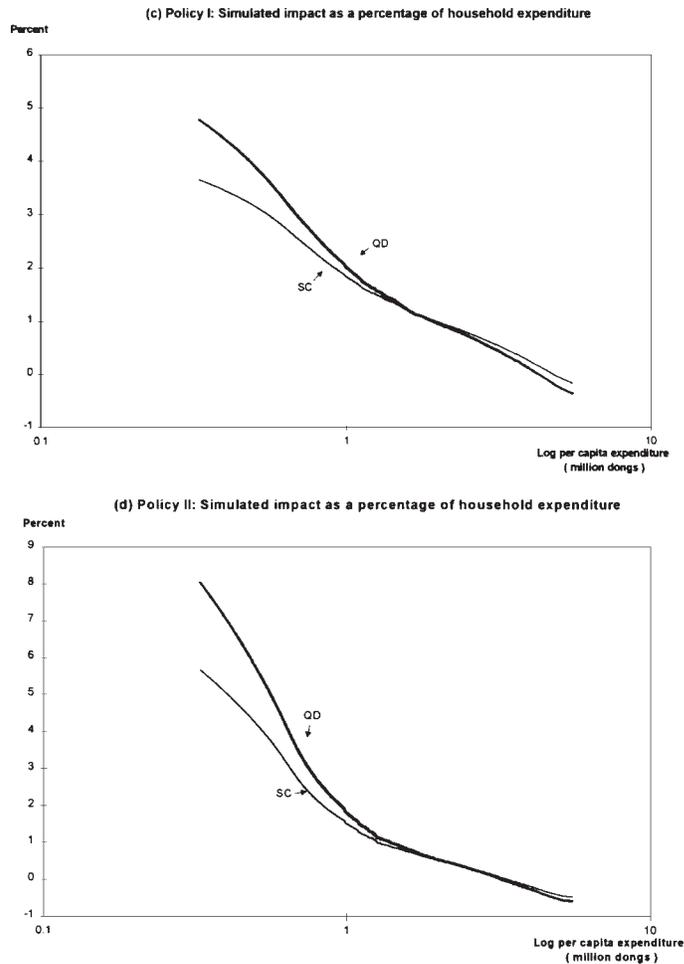
FIGURE 3. Measures of the Impact of an Expansion in Irrigation



Impact on Project Selection

Irrigation investment projects should ideally be approved where benefits exceed costs. It is interesting to ask how the two methods we have reviewed would differ in determining where and which projects to undertake. Table 5 conducts such an exercise using hypothetical (linear) costs. Sampled farm households are used to define a distribution of potential projects where the data for each household can be thought of as determining the characteristics of a potential project area. Each area may contain many more than one farm household. At each unit cost level, we ask how many projects would be accepted by the two methods of calculating household-level benefits. The overall results indicate that the SC and QD estimates tend to be in agreement at both ends of the cost spectrum, but there is a greater margin of error when the cost is around 400 dongs per year per square meter. There is also regional variation.

FIGURE 3. (continued)



Notes: QD refers to the quick-and-dirty method and SC to the slow-and-clean method of project appraisal. In 1992–93, 10,000 dong were roughly equal to US\$1.

Source: Authors' calculations.

Table 6 gives an overall assessment of the outcomes under the QD approach. The table presents summary data on acceptance rates for projects under both the QD and SC acceptance rules for various cost levels. The QD and SC rules agree on the selection of projects most of the time for project costs up to 400 dong. However, there is greater disagreement at 400 dong per square meter and higher.

We also present the net benefits that would be realized were project selection to be decided according to each method. Benefits for both sets of projects are evaluated using the SC estimates to work out actual benefit levels. These are then expressed per square meter. The difference in net benefits realized under each rule gives us some idea of the potential losses due to use of the QD method. In

TABLE 5. Irrigation Projects Approved under QD and SC Benefits

Cost (dong/m ²)	Benefit rule	Projects accepted by region (%)						Total projects accepted (%)
		Northern Uplands	Red River	North Coast	Central Coast	South East	Mekong Delta	
200	SC	99	100	87	6	45	13	68
	QD	100	100	100	0	0	0	65
250	SC	98	100	67	1	17	7	61
	QD	100	100	100	0	0	0	65
300	SC	95	99	32	0	5	3	52
	QD	100	100	100	0	0	0	65
350	SC	89	97	9	0	2	1	46
	QD	100	100	0	0	0	0	47
400	SC	79	92	3	0	1	1	41
	QD	100	0	0	0	0	0	19
425	SC	69	88	1	0	0	1	38
	QD	100	0	0	0	0	0	19
450	SC	53	81	1	0	0	0	33
	QD	100	0	0	0	0	0	19
475	SC	38	65	0	0	0	0	26
	QD	100	0	0	0	0	0	19
500	SC	26	45	0	0	0	0	17
	QD	100	0	0	0	0	0	19

Source: Authors' calculations.

particular, the last column in table 6 gives the percentage of the benefits realized under the SC method that is lost under the QD rule for accepting projects.¹⁷ The loss is small for low-cost projects, which is not surprising because one is unlikely to make serious errors in this case, as most potential projects will be accepted. However, the loss from the QD method rises rapidly for project costs of about 375 dong per square meter; at a unit cost of about 450 dong, the entire project benefit is lost when project selection is based on the QD method.

Figure 4a plots the loss under the QD method against the unit cost. Figure 4b plots the loss under the QD method against the proportion of projects accepted under the SC rule that is a decreasing function of the unit cost. The sharp increase in the loss from the QD method starts to set in at a point when less than half of all projects should be accepted under the SC rule. The realized net benefits under the QD method fall to zero when about one third of the projects should be accepted or, alternatively, at a point where about one fifth of all projects would be accepted under the QD decision rule.

We have said much about the potential benefits of adopting the SC method but little about its costs. Integrated household surveys, such as the VNLSS, can seem expensive propositions. For example, depending on how one accounts for

17. Note that columns 3 and 5 in table 6 give average net benefits, averaged over all accepted projects. To obtain total benefits it is therefore necessary to multiply columns 3 and 5 by columns 2 and 4, respectively. The percentage of benefits lost from using the QD method given in column 6 reflects accepted projects that should have been rejected and rejected projects that should have been accepted.

TABLE 6. Overall Ex Post Assessment of the Quick and Dirty Method

Cost (dongs per m ²)	Projects on which the rules agree (percentage of all projects) (1)	QD rule		sc rule		Loss from using QD instead of sc (%) (6)
		Projects accepted (%) (2)	Average net benefit realized for projects accepted under QD rule (dongs per m ²) (3)	Projects accepted (%) (4)	Average net benefit per m ² (dongs) (5)	
200	91.80	64.97	216.08	68.15	212.32	2.97
250	91.17	64.97	166.08	61.09	183.88	3.94
300	85.71	64.97	116.08	52.36	160.34	10.16
350	95.35	46.77	122.30	45.94	129.48	3.82
400	69.52	18.70	49.43	41.15	91.52	75.46
425	69.12	18.70	24.43	38.00	73.05	83.55
450	68.38	18.70	-0.57	32.85	57.46	100.57
475	69.89	18.70	-25.57	25.66	44.95	141.47
500	73.47	18.70	-50.57	17.40	35.21	254.41

Notes: For all projects accepted, benefits are evaluated using the sc estimates and expressed on a per m² basis. Column 6 = 1-[column 2×column 3]/[column 4×column 5].

Source: Authors' calculations.

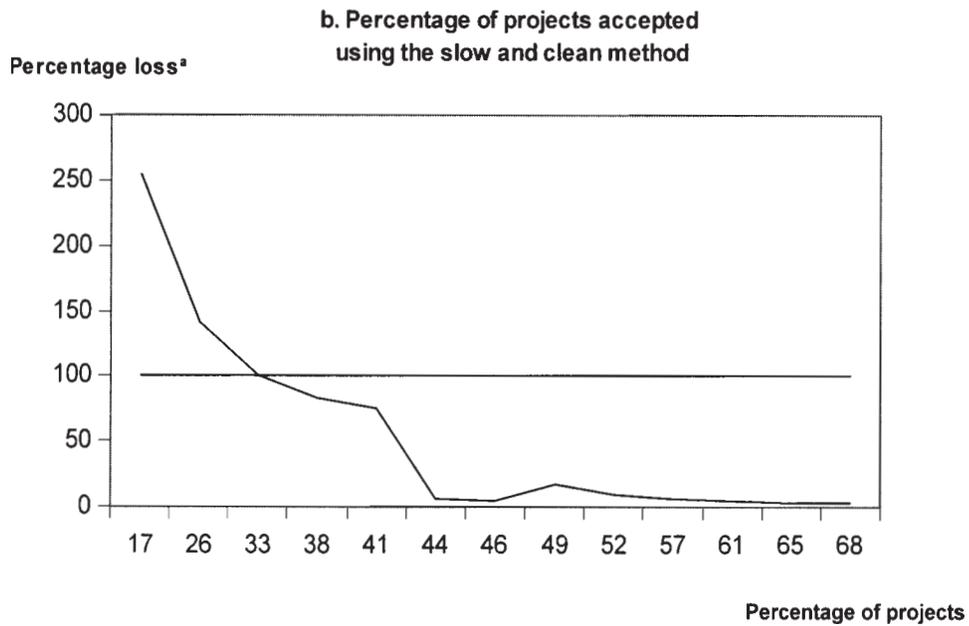
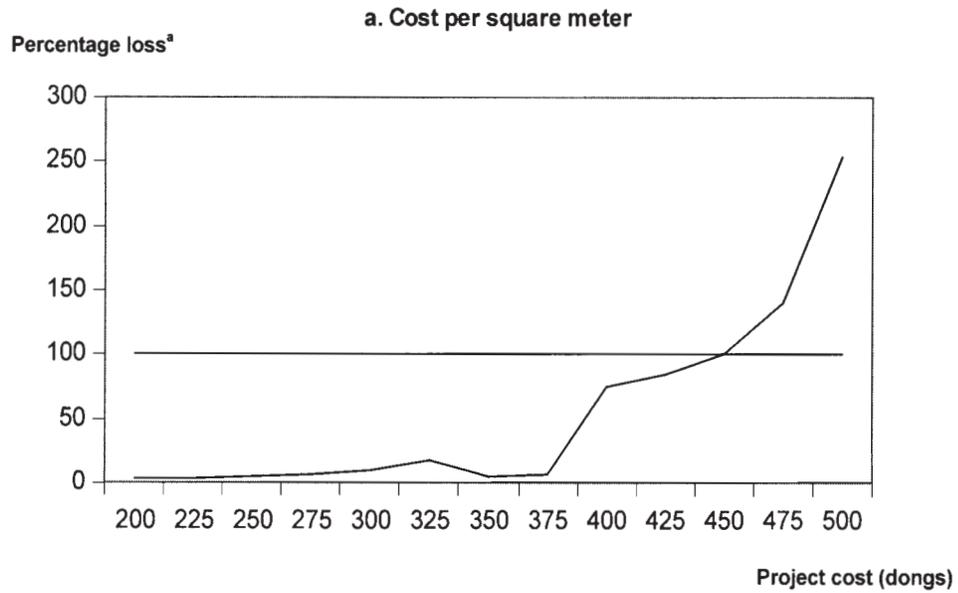
various inputs (such as jeeps and computers that will continue to be used after data collection), the 1992–93 VNLSS cost between US\$500,000 and US\$1 million. However, our results suggest that this outlay is dwarfed by the potential benefits from using the sc method for project selection.¹⁸ Let us assume that the VNLSS cost \$1 million, the upper bound of estimates. Let us also assume a low project unit cost of 200 dongs per square meter. Under these conservative assumptions, we calculate that the realized benefit from using better data to select projects would be sufficient to cover the cost of data collection when irrigating as much as only 3 percent of currently nonirrigated arable land in Vietnam. Of course the data would also have other benefits.

III. CONCLUSIONS

Ex post evaluations often find that realized benefits from irrigation projects are less than had been predicted at the appraisal stage (Carruthers and Clark 1983). Problems in project implementation are usually blamed. But errors in appraisal and project selection may also explain the gap between ex ante and ex post evaluation results. Extensions of the type of analysis in this article to other settings would be needed to fully test that hypothesis. Although it would in many respects be more difficult than our stylized case study, it would also be desirable to apply similar methods to the evaluation of appraisals for specific projects subsequently imple-

18. The total benefit from using the sc method is simply the loss per square meter from using the QD method. The total benefit is given in table 6 by (column 4 × column 5) – (column 3 × column 2) for a given project unit cost times the project size (that is, multiplied by the size of the project in square meters).

FIGURE 4. The Percentage Loss from Using the Quick-and-Dirty Method Instead of the Slow-and-Clean Method for Project Appraisal



^aFor details on calculating the percentage loss, see table 6.
 Note: In 1992–93, 10,000 dongs were roughly equal to US\$1.
 Source: Authors' calculations.

mented. More ex post analyses are also needed, explicitly aimed at testing findings like ours and other common hypotheses to explain worse than expected outcomes.¹⁹

We have provided an experiment in assessing how much a common “quick-and-dirty” method that ignores heterogeneity in project impacts might distort the appraisal. To do this, we have compared the method to a more arduous “slow-and-clean” approach and assessed what benefits can be expected from the extra effort needed to implement the latter approach. As in all empirical research, our results are to some extent data-specific. However, there has been so little research into the benefits of doing more thorough ex ante project evaluations that even one case study can be instructive. There do appear to be some potential lessons for other settings.

Estimation of the marginal benefit function on national data can help guide information gathering for irrigation project appraisals as well as for improving the design of the projects in a specific country or region. For example, the findings for Vietnam suggest that simultaneously investing in the education of irrigation project households found to be lacking in human capital would tend to increase the impact of irrigation investments on average yields, as well as to make them more pro-poor.

We find that at the national level in Vietnam, QD mean benefits are similar to SC mean benefits. At the regional level, there is greater variation. One cannot determine a priori whether the QD method would over- or underestimate SC mean benefits, and our empirical results show that there is no tendency for QD to only over- or underestimate the SC benefits.

Our results indicate that the QD method overestimates the progressivity of a policy of irrigation expansion in Vietnam. The QD results suggest a more pro-poor distribution of gains than the SC method. This difference occurs when we define gains as the benefit per square meter, as well as when we take into account the existing distribution of land and access to irrigation. The QD method ignores household and regional heterogeneity, and the poor have characteristics that reduce the gains relative to the mean.

We find that the QD method rejects some projects that should have been accepted, and accepts some that should have been rejected. We used the SC method to evaluate the realized benefits when project selection is based on the QD method, and compared the results to the benefits realized when the SC method is used for project selection. The comparison suggests that the QD method entails only small losses when project unit cost is low, and hence most potential projects are easily accepted. However, the loss from the QD method rises sharply at medium to high project cost levels. Indeed, at unit costs at which about one third of the potential projects should be accepted (and would yield large net benefits), the QD method incorrectly chooses enough projects to wipe out the realized benefits.

Although a potentially profitable set of projects exists in our experiment, selection based on a method that ignores impact heterogeneity realizes a net loss. The savings from the more data-intensive method are sufficient to cover the costs of its

19. For example, another common hypothesis put forward to explain overestimation of benefits by irrigation appraisals is that technical coefficients used to determine yields assume full adoption of technical packages by farmers. This does not appear to happen frequently, although it is often unclear why. The VNLS survey is unable to throw light on this issue.

extra data requirements (even ignoring other benefits from the data) for projects irrigating more than 3 percent of unirrigated land. This case study suggests that unless project size and cost are small or distributional impacts are irrelevant, there can be large social benefits to investing in more thorough evaluations.

REFERENCES

- Barker, Randolph. 1994. "Agricultural Policy Analysis for Transition to a Market-Oriented Economy in Viet Nam: Selected Issues." FAO Economic and Social Development Paper 123, Food and Agriculture Organization, Rome.
- Bergmann, Hellmuth, and Jean-Marc Boussard. 1976. *Guide to the Economic Evaluation of Irrigation Projects*, rev. ed. Paris: OECD.
- Brown, Maxwell. 1979. "Farm Budgets: From Farm Income Analysis to Agricultural Project Analysis." World Bank Staff Occasional Paper 29, Johns Hopkins University Press, Baltimore, Md.
- Carruthers, Ian, and Colin Clark. 1983. *The Economics of Irrigation*. Liverpool: Liverpool University Press.
- Cleveland, W.S. 1979. "Robust Locally Weighted Regression and Smoothing Scatter Plots." *Journal of the American Statistical Association* 74(368):829–36.
- Dreze, Jean, and Nicholas Stern. 1990. "Policy Reform, Shadow Prices, and Market Prices." *Journal of Public Economics* 42(1):1–45.
- Gittinger, Price. 1982. *Economic Analysis of Agricultural Projects*. Baltimore, Md.: Johns Hopkins University Press.
- Glewwe, Paul. 1994. "Vietnam Living Standards Survey (VNLSS), 1992–93: Basic Information." World Bank, Poverty and Human Resources Division, Policy Research Department, Washington, D.C.
- Londero, Elio. 1987. *Benefits and Beneficiaries: An Introduction to Estimating Distributional Effects in Cost-Benefit Analysis*. Washington, D.C.: Inter-American Development Bank.
- OECD (Organisation for Economic Cooperation and Development). 1985. *Management of Water Projects: Decision-making and Investment Appraisal*. Paris: OECD.
- Strauss, John. 1986. "The Theory and Comparative Statics of Agricultural Household Models: A General Approach (Appendix)." In Inderjit Singh, Lyn Squire, and John Strauss, eds., *Agricultural Household Models: Extensions, Applications and Policy*. Baltimore, Md.: Johns Hopkins University Press.
- UNDP (United Nations Development Program). 1998. "The Dynamics of Internal Migration in Vietnam." UNDP Discussion Paper 1, Hanoi, Vietnam.
- van de Walle, Dominique. 1996. "Infrastructure and Poverty in Vietnam." Living Standards Measurement Study Working Paper 121, World Bank, Policy Research Department, Washington, D.C.
- . 2000. "Are Returns to Investment Lower for the Poor? Evidence for Irrigation in Rural Vietnam." Policy Research Working Paper 2425, World Bank, Policy Research Department, Washington, D.C.
- Vu, Tu Lap, and Christian Taillard. 1993. *Atlas du Viet Nam*. Montpellier and Paris: Reclus, La Documentation Française.
- World Bank. 1995. "Vietnam: Poverty Assessment and Strategy." East Asia and Pacific Region, Washington, D.C.