UNITED NATIONS UNIVERSITY
UNU-WIDER

WIDER Working Paper 2017/111

Project appraisal

A revival is long overdue

Clive Bell*

May 2017

United Nations University World Institute for Development Economics Research
wider.unu.edu

Abstract: This paper makes the case for the systematic appraisal of public sector projects using shadow prices as the signals of social scarcities. In so doing, it attempts to redress the balance between estimating inputs and outputs, central though that task is, and valuing them correctly. The account of how to estimate shadow prices for this purpose pays particular attention to the social discount rate and how to treat uncertainty at both the project and economy-wide levels. There follows an illustration of the ideas and apparatus in practice using an evaluation of India's large rural roads programme (PMGSY) in connection with a World Bank loan. A discussion of the more general uses of shadow prices in evaluating decisions in the public sector is complemented by a series of proposals for the implementation of systematic procedures for project appraisal in government and international donors.

Keywords: cost-benefit analysis, shadow prices, uncertainty
JEL classification: D61, D81, H43, O22

Acknowledgements: I am indebted to Tony Addison and Lyn Squire for valuable comments on earlier drafts. Responsibility for all surviving errors of analysis, interpretation and opinion is mine alone.

* South-Asia Institute, Heidelberg University, Germany, clive.bell@urz.uni-heidelberg.de

This paper is commissioned for the upcoming UNU-WIDER 30th Anniversary Conference publication on 'Mapping the Future of Development Economics', as part of the UNU-WIDER project on 'Development policy and practice—competing paradigms and approaches'.

Information and requests: publications@wider.unu.edu

ISSN 1798-7237 ISBN 978-92-9256-335-6

Typescript prepared by the Author.

The United Nations University World Institute for Development Economics Research provides economic analysis and policy advice with the aim of promoting sustainable and equitable development. The Institute began operations in 1985 in Helsinki, Finland, as the first research and training centre of the United Nations University. Today it is a unique blend of think tank, research institute, and UN agency—providing a range of services from policy advice to governments as well as freely available original research.

The Institute is funded through income from an endowment fund with additional contributions to its work programme from Denmark, Finland, Sweden, and the United Kingdom.

Katajanokanlaituri 6 B, 00160 Helsinki, Finland

1 Introduction

Half a century ago, development planning was still in its heyday and project appraisal stood rather in its shadow, like a dutiful but less attractive sibling. Yet in order to realise a plan, specific projects must be undertaken; for otherwise, the plan would remain in the abstract. At that time, the emphasis was on ensuring that the various branches of the economy would develop together in a consistent way: in this sense, a plan was to be more than merely a collection of individual projects, though they still had to be specified. By no means, however, did this fully answer the question of what were good projects, as opposed to bad ones. Indeed, consistency alone seems a quite inadequate criterion on which to judge how to allocate investment funds. Should a doubling of fertiliser output, for example, be achieved by using process A or process B – or should domestic demand be met, in part at least, by imports? And loosening the straitjacket of 'consistency' somewhat, would it not be better further to improve the road network instead of constructing all of the planned fertiliser plants? To answer such questions, the notion of what is 'good', as opposed to 'bad', must be made precise, and that is the starting point of any system of project appraisal.

As drawing up, and attempting to execute, five-year plans fell out of fashion as a means of promoting economic growth, so project appraisal began to come into its own. In 1968, the OECD published a *Manual of Industrial Project Analysis*, volume II of which (Little and Mirrlees, 1968) was to exert a profound influence on social cost-benefit analysis more generally. The equally notable UNIDO *Guidelines for Project Evaluation* (Dasgupta, Marglin and Sen, 1972) appeared shortly thereafter. The central concerns of both were how to value consumption, savings, public income, foreign exchange and unskilled labour when some key markets function poorly and taxes and quotas are strongly distortionary. In such settings, market prices will deviate quite sharply from

the marginal social values of resources, so that using the former to evaluate projects may lead to serious errors when deciding which to accept and which to reject. What is required instead is a set of *shadow prices*[1] that take into account such concerns and market distortions, and so send the right signals of scarcities. In deriving these prices, an appeal to some form of social welfare function is unavoidable, and an aversion to inequality has its due place in both manuals.

It follows that a project may be bad, in the sense of reducing welfare, for two quite different reasons. First, it may function poorly in the purely technical sense. To give examples, experience may indicate that steel plants employing a particular process produce output of variable quality and have to be shut down rather frequently. Nor is it especially advisable to make school curricula tougher when teachers are inadequately trained and have trouble mastering the current ones. The second, separate and ultimately decisive reason is that a project may fail to pass muster when evaluated at the right prices. For whereas a project may promise to function perfectly well technically and yield a profit at market prices, the value of its outputs at shadow prices can fall short of the shadow cost of producing them. Undertaking it would therefore reduce welfare.

It was natural that economists should devote themselves to the task of estimating shadow prices, especially when the engineers are largely in agreement about what will work well, as is arguably the case for many industrial and some types of infrastructure projects. For shadow prices hold for all projects, whatever their engineering merits. A resulting stream of theoretical refinements and practical applications found its way into the literature on cost-benefit analysis, admirably surveyed by Drèze and Stern (1987) and Squire (1989). The method even made some inroads into the operations of the

[1]The French term *prix fictifs* is somehow more appealing.

World Bank, underpinned in house by Squire and van der Tak (1975). By the early 1990s, however, the scholarly stream had largely dried up, and the practice of shadow pricing had virtually died out in the World Bank.

One reason for the latter was almost surely the Bank's increasing involvement in health and education, spheres in which estimating benefits is inherently more difficult and troublesome than, say, looking up the rated output of a steel or fertiliser plant. Yet this cannot be the whole story. For estimating the benefits generated by dams and roads is also highly demanding, and the Bank had been involved in that part of the business since its founding. At all events, there occurred a shift towards the social sectors not only in donors' lending, but also in scholars' interests, and the latter saw the rise of so-called impact evaluation and the current fashion for random control trials. Thus, the pendulum has swung from estimating the right prices at which to do the accounting to measuring the effects of interventions in spheres from which well-established, 'engineering' certainties are largely absent. At the same time, little if any attention is now paid to the possibility that using market prices to value inputs and outputs can lead to serious errors in the assessment of social profitability.

Thus motivated by the aim of improving public investment decisions, this paper calls for a revival of project appraisal based on the use of shadow prices. The public sector's activities are notably wide-ranging, be it providing schools and engaging teachers to educate children and youth, conducting vaccination campaigns, or constructing and maintaining roads, dams and other infrastructure. Valuing the benefits within a shadow-pricing framework is an intrinsic part of the task of estimating them. It will be argued, moreover, that shadow prices find important applications outside the evaluation of projects, narrowly construed. How to treat uncertainty, which confronts all decisions in practice, is also given a prominent place. Fluctuations in world prices and

trade, and the vagaries of the weather and other macroeconomic shocks at home, constitute an inherently stochastic environment. It is joined by a project's idiosyncratic risks: at least some of its inputs and outputs will be random variables, perhaps correlated with the economy's systemic ones. With the general aim in view, the approach advocated emphasises simplicity underpinned, as far as possible, by robustness. Taking a leaf out of another profession's book, it has been said that, at the end of the day, an engineer must come up with an answer, if only an approximate one.[2]

Such advocacy will cut little ice if unaccompanied by practical proposals for implementation. There is a lesson to be drawn here from the status of the central bank and the central statistical office in government. It is widely accepted that both should be politically independent, and in many countries over the course of the past three decades, both institutions have attained, and then asserted, their independence. It is surely natural that this movement should extend to any bodies that are constituted and charged with the evaluation of public investment decisions. To preserve experience and pool expertise, moreover, such a body should take the form of a permanent office, rather than a motley collection of ad hoc commissions. Proposals along these lines are developed towards the end of the paper.

The plan of the paper is as follows. It begins, in Section 2, with an intuitive account of the various elements comprising project appraisal and the difficulties that arise in practice. Only then does it proceed to give a formal account of how to estimate shadow prices (hereinafter SPs), which serves both as a foundation for the whole argument and as an introduction for those readers who are unfamiliar with the topic. Section 3 deals with the shadow prices of goods and labour, Section 4 with the social discount rate, and Section 5 with uncertainty. To illustrate the method in action, with short-cuts and

[2]An Introductory Lecture delivered by Sir John Baker to Cambridge freshmen in 1963.

approximations placed in the foreground, Section 6 outlines how the social profitability of India's very large rural roads programme was evaluated in connection with a World Bank loan. Returning to the general theme, shadow prices find diverse uses in the analysis of public sector decisions, and these are taken up in some detail in Section 7. Following all this advocacy, Section 8 addresses the operational question of how to revive the practice of project appraisal based on shadow-pricing through organisational and institutional initiatives in government and international donor agencies. Section 9 rounds off the paper with a concluding discussion.

2 Projects and Profitability

In order to grasp the difficulties confronting decision-makers in the public sector, it is enlightening to begin by asking how rational entrepreneurs should set about evaluating a private project. For their tasks share many elements in common and the underlying logic – though not the ultimate goal – is much the same. First, the entrepreneur needs a design, or 'blueprint', for producing whatever goods and services he or she intends to sell. If properly specified, this will involve a detailed description of what is to be done, and how and when it is to be done. Uncertainty, in the form of possible deviations from the 'blueprint', now intrudes. What other outcomes are possible, and what are their chances of occurrence?

Secondly, the markets for the inputs and outputs must be assessed. What can be charged for the products, and what will be the costs of labour and materials? Since the plant or equipment will normally have a lifetime exceeding one year, these questions concern future as well as current prices. The future being what it is, the former are not known with certainty. Like the chances that the plant or process will function in

a particular way, they are conjectures in the entrepreneur's mind, though they may be based, in part, on past experience and market forecasts.

Thirdly, he or she needs a decision rule that will weed out bad projects, but not good ones. The textbooks tell us that their goal is (or should be) to maximise profits; but if the project has a lifetime exceeding one accounting period, which is inherently so if it involves the use of durable inputs that cannot be sold off earlier, then revenues and costs at different dates must be made commensurable for the notion of profit maximisation to be precise. It is widely accepted that, ignoring risk, the right way to render alternative streams of revenues and costs comparable is to calculate their respective net present values (NPV). To complete this calculation, therefore, the entrepreneur needs yet one more price, namely, his or her particular discount rate. Having settled on it, the correct decision rule is clear: accept all (compatible) projects that yield a positive NPV and reject those that do not.

Drawing up blueprints is likewise the starting point for developing project proposals in the public sector. If they involve roads, dams or other types of infrastructure, the blueprints will have an engineering look about them, though the sheets dealing with the projects' effects will rest, in part at least, on conjectures. Projects in the so-called 'soft' sectors of education and health usually involve less engineering and stronger conjectures. Consider, for example, reforming the training of teachers and school curricula, or a vaccination campaign. Even so, there is no escaping the necessity of drawing up a blueprint, however conjectural it may appear. Expertise, in-house or through consultants, is indispensable in bringing evidence to bear, especially to curb politicians' wilder flights of fancy.

Once the blueprint has taken sufficiently definite form for the proposal to be assessed, the inputs and outputs must be valued. Here the private and public calculi part

company; for market prices may not reflect social scarcities. The entrepreneur books costs and revenues at the prices actually paid and received. Any taxes are entered as a regrettable necessity, but subsidies are readily accepted, leaving the net payment as the Treasury's business. In contrast, the public sector's bookkeeping is, in this respect, comprehensive. If, for example, cement is taxed, the public cost of building roads will be lower than the invoice charges for cement; for the Treasury claws back with one hand part of the funds it has given to the spending ministry with the other. More troublesome still, some of the outputs produced by public projects are not only subsidised, but using market prices – even if the markets exist – as a guide is controversial. How, for example, should one value improved educational attainment or reductions in morbidity? On the input side, there is the central question of whether the market wage rate sends the right signal when there is involuntary unemployment or the labour market is inflexible in certain ways and wage goods are taxed or subsidised. To cap matters, calculating the proposal's projected NPV – if that is indeed the right way to conclude the whole procedure – requires an appropriate, *social* discount rate. How is that to be estimated? All these questions must be addressed, and answered, if the decision of whether to undertake the project is to be well-founded.

In order to convey some flavour of how one might proceed in practice, we shall draw on an appraisal (Bell, 2012a) of a World Bank loan to support India's rural road programme, *Pradhan Mantri Gram Sadak Yojana* (hereinafter PMGSY). Launched in 2000, its ultimate aim is to draw India's villages into the mainstream, especially in three domains. First, with improved, all-weather connections to markets, villagers should face more favourable prices for inputs and outputs, which will raise their incomes and sharpen their incentives to cultivate more intensively, pursue new activities and invest in new methods. Secondly, by reducing the time spent travelling to and from

school (and in the rainy season, by making the trip actually possible), an all-weather road should improve the attendance not only of pupils, but also of their teachers, thus promoting the formation of human capital. Thirdly, by likewise improving villagers' access to timely treatment, especially in the event of accidents and bouts of acute sickness, the connection should lower morbidity and mortality.

By the end of 2010, when the Bank's loan was approved, accumulated expenditures had already amounted to about US $14.6 billion, with an estimated further US $40 billion required to complete the program by 2020 (World Bank, 2010). One could be forgiven for expecting that the decision to embark on one of Asia's largest public sector programmes was supported by some specific economic analysis, with or without SPs; but a diligent search of possible sources, including numerous discussions with officials of various ranks, uncovered none. The appraisal for the Bank had to start, therefore, essentially from scratch.

The next step, however, is to lay the formal foundations without regard to the particular exigencies of a project chosen for the purposes of illustration.

3 Shadow Prices

Like any other concern, the public sector uses inputs to produce outputs. Let z_i denote the sector's net output of good i, where the usual convention that inputs carry a negative sign applies, so that z_i is the algebraic sum of its gross output and gross input of i. The vector $\mathbf{z}$ denotes the public sector's net output of all goods. Taking the current $\mathbf{z}$ as given, a public sector project is simply a change therein, denoted by $\Delta\mathbf{z}$, which covers possible retrenchments as well as new undertakings and extensions. Is the particular proposal $\Delta\mathbf{z}$ socially profitable, in the sense that implementing it will

improve social welfare?

In order to answer this central question, we need to value each and every element of the bundle $\Delta\mathbf{z}$ in such a way that the resulting aggregate value reveals whether undertaking the project will improve, or diminish, welfare. This motivates the following definition: the shadow price of a good or factor service is the improvement in social welfare that would result if an additional (small) unit thereof became available to the public sector. Though precise, this sounds rather abstract, so let us move to money-metric utility and hence to prices we can observe.

A very important and extensive category of goods comprises all those that are traded on world markets. For a small open economy, their prices, denominated in dollars, euros or whatever, will be parametrically given. If an additional unit of such a good becomes available to the government and the economy exports the good in question, the said unit can be sold at the f.o.b. world price and so augment public revenue by that amount; if the good is imported, the additional unit can be sold domestically, thereby increasing public revenue by the c.i.f. world price. This is the intuitive argument underpinning the celebrated 'border price rule': the vector of the SPs of goods that an economy trades freely in world markets at parametrically given prices is proportional to the vector of their corresponding world prices, f.o.b. or c.i.f. as appropriate. The scalar multiple in question is the shadow price of public income; for changes in availability, when so exploited, yield a particular change in public income measured in units of foreign exchange, which then results in a change in social welfare, depending on how the government responds to changes in revenue. In order to rid ourselves of this remaining complication where the border price rule is concerned, we can choose an appropriate numéraire, namely, public income not committed to any particular purpose and measured at border prices (Little and Mirrlees, 1974: 358). With this choice,

the SP of a traded good is equal to its world price.

No such elegant line of argument exists for goods that, in virtue of prohibitive transportation costs or restrictive policies, are either not internationally traded at all or, if actually traded, subject to a binding quota. The argument proceeds instead by deriving the marginal social cost of producing the good in question. Let a_{ij} denote the input of good i needed to produce one additional unit of good j, there being n goods altogether. For simplicity, let there be just one sort of labour: the associated marginal input thereof is denoted by l_j. Then the marginal social cost of good j is $\sum_{i=1}^{n} \pi_i a_{ij} + \pi_l l_j$, where π_i and π_l denote the SPs of good i and labour, respectively. Now partition the set of all goods into the subsets of those that are traded and those are not, and label them $T = \{i : 1, \ldots, m\}$ and $N = \{i : m+1, \ldots, n\}$, respectively.[3] The border price rule, when 'normalised' as above, yields $\pi_i = p_i^* \ \forall i \in T$, where p_i^* denotes the world price of good i. Given that the SP of a non-traded good is equal to its marginal social cost of production, we obtain

$$\pi_j = \sum_{i=1}^{m} p_i^* a_{ij} + \sum_{i=m+1}^{n} \pi_i a_{ij} + \pi_l l_j, \ j \in N. \tag{1}$$

It is seen that there are $n - m$ linear equations in $n - m + 1$ unknowns. If the system is of full rank and the shadow price of labour is known, one can solve for the (unique) vector of the SPs of non-traded goods. In compact form, we have

$$\boldsymbol{\pi}'_N = (\mathbf{p}^{*\prime} \mathbf{A}_{TN} + \pi_l \, \boldsymbol{l}\,')(\mathbf{I} - \mathbf{A}_{NN})^{-1}, \tag{2}$$

where the matrix of marginal input-output coefficients $\mathbf{A} = (a_{ij})$ is partitioned into the sub-matrices corresponding to T and N, and a prime denotes a transpose. Eq. (2) reveals that the SP of a non-traded good is its direct and indirect traded content,

[3]The latter includes those traded goods that are not traded at the margin.

valued at world prices, plus the social cost of labour employed, directly and indirectly, in producing it. The procedure has its origins in Tinbergen's semi-input-output model (1966). In practice, the matrix $\mathbf{A}$ is obtained from an aggregate input-output table, on the assumption that the marginal and average coefficients are equal.[4]

There remains the task of deriving and estimating the SP of labour, commonly called the shadow wage rate. When the government employs a worker, two kinds of effects can follow. First, the worker will be drawn out of his or her existing activity, perhaps of involuntary unemployment. As a result, other individuals may also change what they are doing, in a whole chain of adjustments, perhaps involving a sort of jobs ladder. In a poor, still largely agrarian economy, the ultimate result will be that one or more individuals leave peasant agriculture. Be that as it may, the employment of the additional worker in the public sector will normally cause output to fall elsewhere in the economy. When valued at SPs, this reduction in output, denoted by the bundle $\Delta\mathbf{y}(w)$, is termed the social opportunity cost of labour.

The second effect stems from any associated changes in the level of private incomes and how such changes are allocated between consumption and savings. Let the newly engaged worker receive the wage w, net of any direct taxes. If, after the whole chain of adjustments sketched above, the level of private incomes, including the said w, does not change, that will be the end of the matter, and the shadow wage rate will be equal to the social opportunity cost of labour. It is not uncommon, however, for wages in the public sector to carry a premium, so that w will often exceed the ensuing losses elsewhere. Suppose, for simplicity, that the gain in private incomes, ΔM, is wholly spent on the basket of goods $\mathbf{b}$. This act of consumption is socially costly, to be precise, $\boldsymbol{\pi} \cdot \mathbf{b}$ at SPs. Yet the private enjoyment thereof is also socially valuable. The shadow price

[4]Such tables are produced, at lengthy, varying intervals, by the central statistical offices of many countries.

of public income, λ_f, now forces its way back into the picture. For the normalisation imposed in connection with the border price rule leaves open the possibility that there is a premium on public over private income, and in the absence of lump-sum taxation, both theory and an extensive empirical literature establish that there is such a premium – unless the private beneficiaries have incomes below the poverty line. Let $\theta \equiv 1/\lambda_f$, so that $1-\theta$ is the premium on public income. Then the shadow wage rate is given by

$$\pi_l = \boldsymbol{\pi} \cdot (\mathbf{y}(w) + (1-\theta)\mathbf{b}). \tag{3}$$

Given ΔM, the bundle $\mathbf{b}$ can be found from household expenditure surveys. Let the marginal share of income spent on good i be β_i and its consumer price be q_i. Then $b_i = (\beta_i/q_i)\Delta M$, so that (3) also fully reflects any taxes or subsidies on the basket $\mathbf{b}$. The sum $\sum_{i=1}^{n} \pi_i \cdot (\beta_i/q_i)$ is the cost, at SPs, of the additional consumption generated by a unit increase in private income: it is called the *consumption conversion factor*. Estimating ΔM, θ and $\Delta\mathbf{y}(w)$ may well be quite formidable tasks in practice, but in principle at least, the system is complete. The border price rule, (2) and (3) yield the entire vector of shadow prices.

We are now in a position to answer the question, is the project $\Delta\mathbf{z}$ socially profitable? If its scale is sufficiently small that any linearisation involved in (2) and (3) remains valid, then according to the above definition of SPs, it will be socially profitable if its value at those prices, $\boldsymbol{\pi} \cdot \Delta\mathbf{z}$, is positive, and unprofitable if negative.[5] These SPs serve as signals of profitability to the public sector just as market prices do so to private firms. It should be noted that, by definition, the alternative of not undertaking the project, $\Delta\mathbf{z} = \mathbf{0}$, yields a social profit of zero.

Projects almost invariably have a lifetime stretching over a number of years, and

[5]In an abuse of notation, $\boldsymbol{\pi}$ will include the shadow wage rate π_l as needed.

yet time has not been mentioned in the above account. One can, of course, attach the dates of availability to all goods and their corresponding prices, but this device, though perfectly rigorous, leaves vital questions unanswered. It is therefore attractive, not least as an aid to intuition, to introduce the familiar concept of discounting. Analogously to shadow prices at a point in time, we need an appropriate social discount rate.

4 The Social Discount Rate

The formal definition of any discount rate is, the rate at which the value of the numéraire is falling over time. Given the choice of uncommitted public income measured at border prices, the social discount rate at time t is $1-(\lambda_{ft}/\lambda_{ft-1})$; equivalently, the social discount factor at t is $\lambda_{ft}/\lambda_{ft-1}$. The task of estimating λ_f fairly precisely in a one-period setting can be tackled with some confidence; but estimating the whole sequence $\{\lambda_{ft}\}_{t=1}^{t=T}$ is one of a wholly different order, for it involves a full-blown dynamic optimisation problem over some, possibly infinite, time horizon. Confronted with such a task, it is natural to seek practical alternatives that are unlikely to lead to serious errors.

The border price rule for goods suggests an analogous rule for foreign borrowing. Suppose, suspending belief for the sake of argument, the government has the option of borrowing or lending as much as it pleases in international capital markets at the parametric rate ρ_f. It can then transform public income in one period into public income in the next period at the parametric rate $1+\rho_f$, just as one tradable good, i, can be transformed into another, j, at the rate p_i^*/p_j^* through international trade. The social discount rate at time t is therefore ρ_{ft}. History tells us that the rate ruling in these markets varies over time, so it will still be necessary to forecast the sequence $\{\rho_{ft}\}_{t=1}^{t=T}$,

the difficulty and importance of which need no emphasis. In principle, however, the task is finished.

The drawback to this argument is that poor and middle-come countries are borrowers that are almost invariably rationed in these markets, so that their marginal cost of public funds is greater than ρ_f. The margin in question must therefore be identified, and the corresponding λ_f estimated. Three 'pure' variants come to mind, combinations of which can be employed in the light of what particular adjustment is expected to occur if the public sector undertakes additional investment.

In the first variant, public investment crowds out private investment one-for-one. The associated opportunity cost, and hence the social discount rate, is the rate of return on private investment, calculated – it must be emphasised – at SPs. Estimating it is no mean task. In the second variant, a new public sector investment displaces some reference bundle of other public sector projects, also one-for-one. The social discount rate is then the rate of return on this bundle, again calculated at SPs. In practice, one can assume that the said alternative is a representative collection of past public sector projects, suitable data on which, it is hoped, are available.

In the third variant, public investment comes wholly at the expense of private consumption. Given the choice of public income as numéraire, the premium on public income relative to private income now enters once more; but it turns out that, under certain conditions, the said premium will not come into play. Suppose the economy is growing fairly steadily, with per capita output and consumption both increasing at the long-run annual rate g. Let private preferences over dated consumption be representable by the felicity function $u(c_t) = c_t^{1-\eta}/(1-\eta),\ \eta > 0$. Since, by hypothesis, c_t is growing at the rate g, $u'(c_{t+1})/u'(c_t) = (1+g)^{-\eta} \approx 1 - g \cdot \eta$, in view of the fact that g will be at the very most 7 per cent. Hence, the rate at which the marginal utility of

consumption is falling, namely, the *consumption rate of interest*, is simply $g \cdot \eta$. Now, the hypothesis of steady growth also implies that all other key macroeconomic magnitudes are growing at the per capita rate g. Hence, the premium on public income must be constant, so that λ_{ft} is also falling at the rate $g \cdot \eta$, thus yielding the social discount rate. To give a familiar example in the macroeconomics literature, let $u = \ln c_t$ ($\eta = 1$). Then the social discount rate is simply g.

Marrying the many-commodity structure of Section 3 to these proposals for estimating the social discount rate involves a further difficulty, which must be addressed, namely, the relationship between nominal and real magnitudes. World prices at time t, $\mathbf{p}_t^*$, are denominated in dollars or euros, and hence are nominal. The vital part of the argument involving relative world prices, p_{it}^*/p_{jt}^*, is unaffected, but the numéraire is a nominal magnitude. Recalling (2) and (3), the same holds for consumer prices, the wage rate in the public sector and the gain in private incomes ΔM, all suitably dated. No difficulties arise, therefore, if the social discount rate is the rate at which the country can freely borrow on world capital markets, since the latter is a nominal rate, denominated in dollars, euros or whatever. If, however, such borrowing is expected to bump up against a limit in all periods, one or other of the above variants will be in play. Since that involving the displacing of private consumption has the clear attraction of simplicity and aggregate per capita consumption c_t is a real magnitude, this variant demands detailed attention.

Given world prices and commodity taxes at time t, we obtain the consumer prices of tradables at t. Likewise, given the linear technologies $(\mathbf{A}, \boldsymbol{l})$ and the wage rate, and assuming price equals marginal (equals average) cost, we obtain the producer prices of non-tradables, and hence, given commodity taxes or subsidies, their corresponding

consumer prices. Let $\kappa(\mathbf{q}_t)$ denote the cost-of-living index at t.[6] Then the gain in real private income associated with the nominal gain ΔM_t is $\Delta M_t/\kappa(\mathbf{q}_t)$. If this is wholly consumed, it will correspond exactly to the change in aggregate consumption. It is evidently convenient to keep the nominal magnitudes in Section 3, appropriately dated, together with the SPs so derived. The corresponding social discount rate is then the nominal rate $\kappa(\mathbf{q}_t) \cdot g \cdot \eta$. In what follows, the social discount rate, however it be derived, will be denoted by ρ_t.

It is readily conceded that implementing the above structure involves a formidable amount of work in forecasting a whole variety of variables over a rather long time horizon. In this, the public sector faces the same difficulties as private firms. Mistakes are inevitable, but the task is unavoidable. We now turn to the ensuing uncertainties.

5 Uncertainty

Appraising any project involves, in the final reckoning, valuing its inputs and outputs. This act of bookkeeping must take account of uncertainty from two sources. First, there is the 'technical' performance of the project itself, which may not live up to its designers' and managers' expectations. These risks are idiosyncratic to the project. Secondly, there is the evolution of the economic environment in which the project is undertaken, especially movements in prices but also changes in market conditions. The SPs, which are used to value inputs and outputs, must reflect this systemic uncertainty, which besets any proposed project. Now, a rational, risk-neutral decision maker will calculate its *expected* net present value. Should the public sector do the same, using appropriately estimated SPs? Little and Mirrlees (1991) argue that this criterion is

[6] If the functional form of private preferences over goods is known, or assumed, this will be the corresponding Könus price index.

strongly defensible in practice, and they supply a correction in the event that the project's returns are correlated with movements in national income. Yet for all its welcome simplicity and widespread acceptance, this criterion still warrants discussion, whereby the distinction between the two sources of uncertainty must be maintained.

5.1 Economy-wide shocks

The essential argument for the validity of using expected values is set out in Arrow and Lind (1970) in connection with the social discount rate. Suppose a project's net returns are sufficiently small, statistically independent of national income and distributed in lump-sum form over a large population. Then the project will have a vanishingly small effect on each individual's net income, which is itself random. In the next, vital step, Arrow and Lind prove that when the costs of the individual risks arising from the project are summed up over the whole population, the resulting aggregate cost thereof also goes to zero as the population becomes very large. Little and Mirrlees (1974: 328-331) arrive at the same result by a somewhat different route, but without imposing the condition that the project be very small. Now, national income is itself an aggregate, and in the light of the importance accorded to estimating the SPs of individual goods and labour in the foregoing sections, the question arises as to whether the argument remains valid in settings with more than one good. Here, Little and Mirrlees sound a note of caution: 'It is not so clear that this assumption [of an aggregate good] is innocuous [...]' (p. 329). A closely related question is, how should SPs be estimated when the environment is stochastic?

Let it be granted that it is valid to use expected values, not only as a decision rule but also – and this is entailed – as the basis for estimating SPs. Returning to the framework of Sections 3 and 4, we begin with the world prices of traded goods, which

are now to be regarded as random variables. Under the stated assumption, the vector of shadow prices of traded goods in year t is equal to $\mathbb{E}[\mathbf{p}_t^*]$, where $\mathbb{E}$ is the expectation operator. That is to say, the vector $\mathbf{p}_t^*$ must be forecast for each year over the time horizon up to T. Only the point forecasts matter; any forecasting bands are irrelevant.

The next step is to obtain the SPs of non-traded goods and labour using (2) and (3), which involve the technologies described by ($\mathbf{A}$, $\boldsymbol{l}$) and the various factors determining the shadow wage rate, namely, $\mathbf{b}, \Delta M$, θ and $\Delta\mathbf{y}(w)$. Leaving aside, for the moment, the fact that world prices are stochastic, all these parameters and quantities are arguably random variables; for they are influenced by shocks that stem from events in the domestic economy, such as the weather, epidemics and changes in market conditions and policies. What is more, they are almost certainly correlated, influenced as they are by these common events. Taking expectations of the r.h.s. of (3) and (4), respectively, we have the corresponding SPs:

$$\boldsymbol{\pi}'_N = \mathbb{E}[(\mathbf{p}^{*\prime}\mathbf{A}_{TN} + \pi_l \boldsymbol{l}')(\mathbf{I} - \mathbf{A}_{NN})^{-1}] \tag{4}$$

and

$$\pi_l = \mathbb{E}[\boldsymbol{\pi} \cdot (\mathbf{y}(w) + (1-\theta)\mathbf{b})]. \tag{5}$$

Hence, even under the not implausible assumption that $\mathbf{A}, \boldsymbol{l}, \mathbf{b}, \Delta M$, θ and $\Delta\mathbf{y}(w)$ are statistically independent of the world prices of traded goods, the resulting shadow prices of non-traded goods and labour will not be equal to those yielded by inserting the mean values of the elements of $\mathbf{A}$, $\boldsymbol{l}, \mathbf{b}, \Delta M$, θ and $\Delta\mathbf{y}(w)$, namely,

$$\boldsymbol{\pi}'_N = (\mathbb{E}[\mathbf{p}^{*\prime}] \cdot \mathbb{E}[\mathbf{A}_{TN}] + \pi_l \mathbb{E}[\boldsymbol{l}'])(\mathbf{I} - \mathbb{E}[\mathbf{A}_{NN}])^{-1} \tag{6}$$

and

$$\pi_l = \boldsymbol{\pi} \cdot \mathbb{E}[(\mathbf{y}(w) + (1-\theta)\mathbf{b})]. \tag{7}$$

For (4) involves sums of products of a whole array of random variables, some of which are surely correlated.[7]

This finding is unwelcome news for practitioners. If taking mathematical expectations is valid in order to arrive at the social values of resources, then (4) must be employed, which requires the whole structure of correlations among $\mathbf{p}^{*\prime}$, $\mathbf{A}$, $\boldsymbol{l}$, $\mathbf{b}$, ΔM, θ and $\Delta\mathbf{y}(w)$. In practice, the task of estimating just their mean values will be formidable enough, and the temptation to use (6) instead may be nigh on irresistible, even though it will introduce errors into the appraisal of any project.

There remains the question of whether taking mathematical expectations is a valid rule to estimate SPs; for an argument that holds when there is but one aggregate good, in the form of national income, may not hold when there are two or more goods. For a small open economy with two traded goods, one non-traded good and labour, the said rule is strictly valid only under certain restrictions on preferences over goods and for bearing risks (Bell, 2016), though it turns out that the size of the errors yielded by the rule will be rather small under the assumptions about preferences that are likely to be made in practice. This reassuring finding does not, however, rule out the possibility that the model itself is faulty in some way, in which event the SPs derived from it are likely to deviate from their true values. Misspecification of the model is therefore a source of systemic risk that affects the process of evaluation. To sum up, practitioners will doubtless want to employ (6) and (7), but they should do so with due caution.

[7] In particular, changes in the global climate may well involve greater variability.

5.2 The project's performance

Projects seldom work out exactly as envisaged in the plans drawn up by their designers, engineers and managers. Building a road may run into unforeseen difficulties with drainage; new crop varieties may not respond to fertilisers on farmers' fields as they do on experimental stations; a steel plant may not get along well with local coal; or, on a happier note, school meals may yield unexpectedly large improvements in children's physical and cognitive development. What all of these examples have in common is that they are outcomes that are arguably independent of the common shocks to which the economy is subject. That is to say, although the project defined by $\Delta\mathbf{z}$ involves a set of random variables, the latter are arguably independent of the stochastic factors that enter into the determination of the SPs to be used in evaluating $\Delta\mathbf{z}$. If that be so and the project is sufficiently small, it will be valid to use the mean values of the components comprising $\Delta\mathbf{z}$: the project should be accepted if, and only if, $\boldsymbol{\pi}\cdot\mathbb{E}[\Delta\mathbf{z}] > 0$.

For practitioners, this is the good news that accompanies the weaknesses of (6). The fundamental difficulty, then, lies in estimating SPs for an economy beset by stochastic shocks, and not in the uncertainty surrounding the project's performance. All that is needed for the latter is to extract an estimate of $\mathbb{E}[\Delta\mathbf{z}]$ from the technical experts and managers.

5.3 Monte Carlo simulations

Decision-makers in government may not be persuaded by the reduction of all risks to expected values, for all its grounding in economic argument. Their intuition may well lead them to demand some estimate of the dispersion of outcomes, perhaps with special attention to the left-hand tail. In so doing, they implicitly reject the said argument and the assumptions on which it is based. To address their concerns, one can resort to

Monte Carlo simulations.

The first, elementary step is to accept the SPs, howsoever they have been estimated, and concentrate on the project's idiosyncratic risks. For this purpose, the technical experts and managers must provide not merely the vector $\mathbb{E}[\Delta \mathbf{z}]$, but the joint distribution of the whole set of variates $(\Delta Z_1, \Delta Z_2, \ldots, \Delta Z_n)$. This is a tall order. In practice, they are likely to respond by simply giving the range for each variate, perhaps with the implicit hint that the variates are to be regarded as independently distributed. In the absence of any further hints, one could assume that each variate is distributed uniformly or as a symmetric triangle on its given interval. On this basis, one makes, say, 10,000 random draws, calculating the net present value at SPs of each 'project' so drawn. By the central limit theorem, the resulting distribution of the NPVs will be approximately normally distributed, with an expected value of $\boldsymbol{\pi} \cdot \mathbb{E}[\Delta \mathbf{z}]$. The tails will not, therefore, play much of a role – unless the project takes the form of a nuclear power plant or an imposing high dam.

The second step is to treat the SPs as random variables. Such randomness will arise even when they are estimated on the basis of expected values. For in practice, the expected values of world prices and all the other elements appearing in (2) and (3) will not be the population means, but rather a set of sample estimates thereof, and the latter are random variables. It should be emphasised that most of the associated standard errors are likely to be small, especially in relation to those of the variates themselves, but they are present all the same. To give an example, the expected value of the world price of traded good i, $\mathbb{E}[p_i^*]$, can be thought of as a point forecast with an associated standard error, whereby the latter is surely much smaller than the standard error of the variate P_i^*.

Suppose, therefore, that the said expected values are normally (or uniformly) and

independently distributed. A single, random draw from their joint distribution yields a particular constellation $(\mathbf{p}^{*\prime}, \mathbf{A}, \boldsymbol{l}, \mathbf{b}, \Delta M, \theta, \Delta\mathbf{y}(w))$. The border price rule, (2) and (3) then yield the corresponding vector of SPs. On the basis of this draw of SPs, the first step is employed to obtain a set of values of the project's NPV. A large number of random draws to yield vectors of SPs can be made, each followed by the first step, thus yielding a grand distribution of the NPV that reflects both systemic and idiosyncratic risks, where it should be recalled that the dependence of the SPs of non-traded goods and labour introduce some degree of correlation among them.

There remains the caveat that the underlying model in Section 3 may be faulty and that the use of expected values almost surely entails some errors in the resulting SPs. These awkward possibilities will need closer investigation if the above procedure yields a distribution of the project's NPV whose mean is not strongly positive and whose non-positive tail is rather large.

6 An Example: Rural Roads

Armed with the apparatus and results of Sections 3-5, we now provide an account of how a World Bank loan to finance part of PMGSY was appraised. The task was to be accomplished within about 15 months, by the equivalent of about one full-time person. A resort to strong simplifications and various short-cuts was therefore unavoidable, and these have a prominent place in the sketch that follows.

6.1 The procedure

One of the first hurdles to be encountered was the non-availability of even a rudimentary set of SPs off the shelf. Since their estimation in detail within the deadline was out of the

question, attention was focussed on three parameters, namely, the social discount rate, the premium on public income, and the shadow cost of constructing and maintaining the new roads over an assumed life-cycle of 30 years.

The social discount rate plays a leading role in the appraisal of such long-lived projects. Pressed for time in making the appraisal, using the consumption rate of interest, $g \cdot \eta$, had evident appeal. Private consumption per capita had been growing rapidly for over two decades, and a continuation of the current rate of about 4.5 per cent p.a. could be expected for some time to come. As noted in Section 4, the choice $\eta = 1$ is ubiquitous in the macroeconomics literature; but Little and Mirrlees (1974, 1991) argue in favour of $\eta = 2$. Taking a rather stringent position, therefore, the test rate of discount ρ was set at 10 per cent[8] – though a good case could also be made for a lower value, set alternatively at 5.5 per cent. Since these are real rates and there was no estimate of the vector $\boldsymbol{\pi}$ to draw on, all costs and benefits were estimated at 2010 prices.

Dealing with the direct costs of extending the road network was relatively straightforward. Transport specialists provided estimates of the range of costs per km. for the construction and maintenance phases separately, distinguishing between hilly and flat connections. These furnished an important component of $\mathbb{E}[\Delta \mathbf{z}]$, in keeping with Section 5.2, but the ranges were kept for subsequent sensitivity analysis. As for the associated SP, one senior specialist claimed that in his extensive experience in various parts of the world, he had scarcely ever encountered an accounting ratio outside the range 0.75 to 0.90, and since this accorded with the sort of estimates for construction in the older literature, the mean value $\mathbb{E}[\pi/p\,] = \ 0.825$ was selected. This choice and the discount rate of 10 per cent imply that the present value of the lifetime costs, at

[8] It was learned later that Murty and Goldar (2007) had also arrived at this upper estimate, but by a different route.

SPs, of extending the network by 1 km. is $64,000.

The latter amount has acquired a particular connotation, thus adding spice to the question, is it worth the resulting benefits? Estimating the latter, in contrast to estimating the costs, involves a variety of substantial difficulties. The classic approach is to estimate the savings in the direct costs and time of transporting people and goods, making an allowance for changes in the volume of traffic. Since PMGSY's beneficiaries are overwhelmingly peasant farmers and agricultural labourers, who go about their business in a setting quite removed, in certain ways, from the textbook case, and who are very likely, moreover, to enjoy benefits in the spheres of education and health, the classic approach seemed wanting, even if traffic studies and forecasts were available. The alternative chosen was to estimate the changes in their incomes from production and employment, and then to add an estimate of the monetary value of the attendant improvements in education and health, whereby the latter involves a good measure of speculation in *any* approach to estimating benefits.

Since the state pays for the roads and the benefits accrue to peasants, labourers and a few others, the premium on public income must enter into the reckoning. A strong case can be made that it is zero (i.e., $\theta = 1$); for the roads are built to serve villages located mostly in remoter and rather backward regions, whose inhabitants are also mostly poor. Hence, the money-metric benefits that the villagers obtain from such a connection, once estimated, count the same as the cost of providing it, at $64,000 per km.

Now, a road is a local public good, so the larger the population it serves, the greater are the benefits it generates – all else being equal. There is little chance of congestion on PMGSY's rural feeders, and though the extra traffic might be felt in the local towns, ignoring these external effects should not involve significant errors. The details

of the process for selecting villages to receive a connection need not concern us here. Suffice it to say that more populous ones are favoured and the officials who manage the programme supplied estimates of the numbers of villages to be connected, the distribution of their populations according to the 2001 census, and the distribution of the lengths of the connections to be provided. That leaves the task of estimating the benefits received by a representative family, which, when suitably scaled up, can be compared with the aggregate cost of thus extending the network.

The commercial sphere is the natural starting point, since the basis for making the associated estimates is conceptually clear and empirically least speculative. As with SPs, no specific, directly useful studies of the effects of rural roads could be discovered. India's sheer heterogeneity also posed problems. A resurvey of 30 villages in an upland tract in eastern India was commissioned, 6 of which turned out to have received a connection since the last survey round in 2004. This was supplemented by field trips to areas in the north-east and north-west of the country and the importation of results from an analysis of the effects of providing rural roads in Bangladesh (Khandker et al. 2009). For the substantial sample of Bangladeshi villages surveyed, the authors estimate that per capita household consumption increased by 11 per cent. This estimate was combined with data for the said 30 Indian villages to yield estimates of the 'commercial' benefits in the poor states, with a check involving a comparison of the estimated reductions in transport costs from the resurvey with those obtained by Khandker et al. (2009). The latter, classic approach was employed to obtain estimates for the relatively commercialised hilly tracts in the north-west. Aggregate benefits for each state were obtained by scaling up, using the village lists of planned connections.

The next step was to calculate the net present value of the estimated stream of benefits under the assumption that it would be constant over the 30-year horizon,

and then to compare it with the corresponding net present value of the costs of the planned extensions, both calculated at the alternative discount rates of 5.5 and 10 per cent. Benefits comfortably exceeded costs at the lower rate in those poor states with relatively large villages, but were somewhat lower than costs at the higher rate. Costs exceeded benefits at both rates in the north-western state, where villages are also generally somewhat less populous. Thus, the final reckoning rests on the sizes of the benefits in the spheres of education and health.

Two radically different methods were used to estimate them. In the resurvey villages, both the sampled households and informally assembled focus groups were asked to rank the importance of the benefits in the three spheres. Their responses pointed unambiguously to those in education and health combined being at least on a par with the 'commercial' ones. Taken at face value, this implies simply doubling the latter to get total benefits, a step which yields clear social profitability even with a test discount rate of 10 per cent, except in hilly states with small villages, which pass muster at about 7 per cent. This way of ascertaining willingness to pay has evident weaknesses, so resort was made to a more formal, albeit still speculative approach. An overlapping generations model developed to analyse the long-term economic effects of the AIDS epidemic in Kenya was re-calibrated to the rural setting in the Indian states, using data on school attendance, morbidity and mortality from the resurvey as well as fragments from other sources (Bell, 2012b). According to this model, the ratio of commercial to non-commercial benefits in the first 20-year period is about two to one, falling to four to five in the following period. To sum up, one can conclude that taken as a whole, PMGSY is almost surely socially profitable at the test rate of 5.5 per cent, and very probably so at 10 per cent.

6.2 Discussion

The emphasis lavished on estimating $\boldsymbol{\pi}$ in Section 3 sits ill at ease with its almost complete absence from Section 6.1. The reason for the latter is clear, but does it cast doubt on the conclusion concerning PMGSY's social profitability? Before attempting to answer this question, however, the fact that a set of SPs is most unlikely to be available off the shelf requires a more general comment.

The long-term aim must be to ensure off-the-shelf availability by establishing the right institutional arrangements and building up the technical capacity to estimate SPs. In the meantime, practitioners will be faced with the difficulties described in Section 6.1. The importance of the social discount rate is beyond dispute, and estimating it can, with certain reservations, proceed independently of the steps in Section 3, with some sensitivity analysis to be on the safe side. It is also clear, however, that the apparatus of Section 3 provides a framework for thinking about the problems of appraisal and assessing what sort of short-cuts are defensible, even if actually estimating a whole set of SPs is out of the question. This framework was invaluable in guiding the procedure in Section 6.1.

Returning to the question of whether that procedure is robust, providing a firm answer is impossible, but the principles involved are clear. Suppose the inhabitants of a village get a new road. Facing new prices, they may well change their production plans, and will in any event obtain some additional income, which they will spend on some bundle of goods. All this bookkeeping is done at market prices, and the additional income duly appears in Section 6.1, with $\theta = 1$. These changes in production plans and consumption should, however, be valued at SPs. Given $\theta = 1$, it then follows that the calculations in Section 6.1 will be correct if, and only if, the value of the change in net output at SPs is equal to its value at market prices multiplied by the

consumption conversion factor. Thus, the size and sign of the error, if any, resulting from the procedure in Section 6.1 can be established only if the entire set of SPs is known. Whether the error is a serious one is a matter of surmise. The consumption conversion factor is almost surely less than 1, despite subsidies on essential food items, but irrigation, fertilisers and electricity are likewise subsidised. Similar considerations apply to the evaluation of benefits in the spheres of education and health, the services of which are largely provided by the public sector without fees. To the extent that the villagers make heavier use of these services, any resulting costs of providing them must enter into the reckoning at SPs. On balance, the above conclusion concerning PMGSY's social profitability likely remains tenable; but some doubt has crept in, which cannot be resolved on the basis of present knowledge.

The case study exemplifies the pervasiveness of uncertainty, both about the size and nature of the programme's effects and how to value them in assessing its social profitability. By and large, mean values have been used, supplemented by some sensitivity analysis in respect of some key parameters. For such large schemes, it would be desirable to undertake a Monte Carlo analysis that encompasses the joint distributions of the programme's effects and how they are valued, so that decision makers can gain a clear view of the distribution of its net present value, as the measure of its social profitability. Creating the capacity to meet this need, among others, will be taken up in Section 8.

7 The Diverse Uses of SPs

Given the vector of SPs and the social discount rate, establishing whether any proposed project whose performance is perfectly certain – or assumed to be so – is socially prof-

itable, is matter of simple accounting. In the presence of uncertainty, the calculations are more complicated, but still manageable. It will now be demonstrated that, with the apparatus of Sections 3-5 at hand, what constitutes a 'project' covers a whole variety public sector decisions. For the sake of simplicity, uncertainty will be kept largely in the background.

7.1 Public vs. private: production and goods

The classic *Manuals* on social cost-benefit analysis were written at a time when governments were heavily involved in building steel and textile mills, and all sorts of plants in various branches of industry, a fact that clearly influenced how the texts were written and the authors' choice of examples. Times have changed, and most governments have largely retreated from this sphere of investment activity, even though many of them have yet to privatise – or dismantle – much of what their predecessors built. This retreat raises a question of central importance: what is the right scope for public production, as distinct from public provision? This question motivates Devarajan et al.'s (1997) thoughtful attempt to revive – and reorient – project appraisal, in which they pay a great deal of attention to specifying private production as the counterfactual and downplay the role of SPs. It will now be argued, on the contrary, that SPs force their way into the reckoning.

Consider the familiar, and celebrated, example of the shoe factory.[9] It is readily conceded that there is no *prima facie* case for it to be a public-sector undertaking, as opposed to leaving the investment decision to private investors, who will do their calculations at market prices, with imports as a further alternative if the factory is privately unprofitable. Let the project, if privately undertaken, involve the net output

[9]The authors surely had Rosenstein-Rodan (1943) in mind when choosing it, though they do not cite him.

vector $\Delta\mathbf{y}$. It is privately profitable if, and only if, $\mathbf{p} \cdot \Delta\mathbf{y} > 0$, where $\mathbf{p}$ denotes the vector of market prices. The public undertaking involves the net output vector $\Delta\mathbf{z}$, where there is every reason to suppose that $\Delta\mathbf{z}$ and $\Delta\mathbf{y}$ will differ materially. For whereas goaded by profit-seeking, owners and managers will keep to their tasks, including keeping the workers to theirs, the more relaxed setting of the public sector, often with restrictions on firing at all levels, will tend to encourage padding and perhaps waste. It will be no surprise, therefore, that what is a privately profitable shoe factory may well be a loss-making public enterprise – at market prices $\mathbf{p}$.

This is not, however, the end of the matter. For suppose investment in the shoe factory is privately profitable because there is a stiff tariff on imported shoes or the government has granted generous tax holidays. It is then quite possible that this private action is socially unprofitable: to be precise, $\boldsymbol{\pi} \cdot \Delta\mathbf{y} < 0$. Now, $\Delta\mathbf{y}$ is the firm's private information, the government's knowledge of which will be somewhat incomplete or hazy. Yet any substantial investment requires a licensing procedure, in which much information must be divulged and compliance with regulations is supposed to be ensured. The government should, moreover, be able to form a fairly good idea of what $\Delta\mathbf{z}$ involves, which should also help in getting a fairly reliable estimate of $\Delta\mathbf{y}$. Obtaining an estimated $\Delta\mathbf{y}$ and assessing it at $\boldsymbol{\pi}$ will surely add to the burden of the licensing procedure for all concerned, but if the project involves substantial investment, incurring the associated cost is likely to be socially worthwhile. However this assessment turns out, whether having the factory in the public sector is socially profitable remains an open question. If the scale of padding and waste is large enough, the chances that the difference between shadow and market prices will rescue the project $\Delta\mathbf{z}$ will be slim.

Consider next, towards the other end of the spectrum spanning private and public

goods, vaccination against an infectious disease, an act that benefits not only the individual who is vaccinated, but also all other members of the community.[10] This external benefit plays no part in the individual's decision, unless he or she is both somewhat informed about epidemiology and altruistically disposed. The private costs comprise those of the treatment and trip to the clinic, the opportunity cost of time so spent and, sometimes overlooked, the expected costs of complications in the event of an adverse reaction, which can, in rare cases, lead to lifetime disability or death. The fact that the external benefit does not enter fully, if at all, into the private calculus almost invariably results in a level of immunisation in the population that is socially sub-optimal when vaccination is privately provided, even in fairly affluent societies; but free treatment, though strongly desirable, may not suffice to yield an optimal level in poor ones.

This conclusion is so well known as scarcely to bear repetition, but a little reflection suggests that the chief problem in assessing the social profitability of projects in this sphere lies elsewhere, namely, how to value reductions in morbidity and mortality. Hypertension, diabetes and childhood malnutrition, for example, are serious, but non-infectious conditions, and – abstracting from any pooling through the provision of social insurance – treating them effectively generates purely private benefits. Estimating these benefits is indeed a formidable task, and until this has been tackled successfully, the job of wrestling with the size of any external benefits ought to be put on the back burner. Expressed somewhat provocatively, what is the point of trying to determine, say, the total number of infections averted (it will exceed, but lie rather close to, one) when one individual is vaccinated, when it is to be multiplied by a individual benefit whose size unknown? This is not to dispute the importance of externalities, but rather

[10]No vaccine confers absolute immunity, but the chances of an outbreak fall as the proportion of the population with substantial immunity rises.

to point out a more pressing task of estimation.

After going through the argument in Sections 3 and 4, the reader will naturally ask how, if at all, the private benefits of reduced morbidity and mortality enter into the SPs and the social discount rate derived therein. The brief answer is, there is no natural place for them, and external effects are ruled out by assumption. They must be grafted on somehow; and one approach has been sketched in Section 6. Suffice it to say that the approach is firmly in the framework of willingness-to-pay, i.e., money-metric utility, consistent with Sections 3 and 4. This proposal is in keeping with the engineer's dictum that, at the end of the day, the calculations must enable one to come up with an answer, if only an approximate one.

7.2 Counterfactual and fungibility

There is always an alternative, however unsatisfactory, to some proposed course of action. Leaving things as they are is an obvious one, and it underpins the definition of a public sector project as a change in the existing net supply vector $\mathbf{z}$. Returning to the shoe factory, the alternative to having it in the public sector is cast as leaving the investment decision to the private sector, while noting an important second alternative in the form of imports. In view of Sections 3 and 4, the latter alternative prompts the question, what is the underlying counterfactual in the derivation of $\boldsymbol{\pi}$ and the social discount rate?

The answer is that the border price rule, (2) and (3) are derived from perturbations to a general equilibrium system, in which the world prices of traded goods are given, the producer prices of non-traded goods are equal to their respective marginal costs of production and, rather generally formulated, employing an additional worker in the public sector results in a change in net output elsewhere in the economy and a change

in private income. For most economies and most traded goods, the assumption that world prices are parametrically given is surely unexceptionable. If, for some reason, a government does impose quantitative restrictions that make such a good non-traded at the margin, then it will belong in the latter category. The assumption concerning the producer prices of non-traded goods is rather strong, and those needed to get at $\Delta\mathbf{y}(w)$ and ΔM are stronger still. Imperfect competition in the branches producing non-traded goods will violate the former assumption, and the scope for making errors in estimating the shadow wage rate using (3) requires no further elaboration. The issue, then, is not whether the SPs so derived are saddled with errors, but rather whether the system from which they are derived yields estimates that are sufficiently robust.

In this connection, the shadow price of public income, λ_f, deserves careful scrutiny. If, for example, the public sector obtains an additional unit of a traded good, government revenue will increase by the corresponding world price. What the government does with this extra revenue, for good or ill, determines λ_f: that is to say, SPs depend, in general, on the government's behaviour (Bell and Devarajan, 1983). This involves the matter of fungibility. Once more, it is Rosenstein-Rodan who furnishes a fine, albeit less well-known illustration. In a memorandum written for the World Bank in 1948, he points out the possible unintended consequences of financing an infrastructure project. This goes as such into the Bank's loan portfolio; but the recipient government would have undertaken it anyway, and exploits the resulting budgetary freedom to spend on some other undertakings, which otherwise would have been infeasible. As a consequence, a chain of other changes occur, mostly in the private sector, with the end result that the Bank has financed, *inter alia*, a bordello.

This example reveals that the estimation of λ_f involves certain difficulties that were partly swept under the carpet in Section 3 by the adroit choice of numéraire. How

the government responds to changes in revenue, be it by changing certain tax rates, or adjusting borrowing, or even expenditures themselves, must be specified explicitly if SPs are to be derived. Once this is settled, and it is an important task in practice, the structure in Section 3 is complete – though elaborating it to include the supply of, and demand for, the services of bordellos can be omitted.

7.3 Foreign aid

A government considers some proposed project, to be financed out of domestic resources. Since the project would stretch over many years, let it be denoted by the sequence $\{\Delta \mathbf{z}_t\}_{t=0}^{t=T}$. It will be socially profitable if, and only if, its NPV, $\sum_{t=0}^{t=T}(\boldsymbol{\pi}_t \cdot \Delta \mathbf{z}_t)/(1+\rho)^t$, is positive, where, for simplicity, the social discount rate is assumed to be constant.

An aid donor becomes aware of this project, or perhaps even suggested it to the government as a promising possibility in the first place. Now, donors insist on particular standards, from the tendering process and technical evaluation, to the protection of the environment and social groups who might be displaced. Hence, what the donor proposes, which is denoted by $\{\Delta \mathbf{z}_t^a\}_{t=0}^{t=T}$, need not be at all the same as $\{\Delta \mathbf{z}_t\}_{t=0}^{t=T}$. Associated with the donor's proposal, moreover, is a financial plan, under which the recipient is to make a net payment of d_t in foreign currency in period t, where d_t will be negative in the construction phase. As the recipient government sees it, the social profitability of the donor's proposal is $\sum_{t=0}^{t=T}(\boldsymbol{\pi}_t \cdot \Delta \mathbf{z}_t^a - d_t)/(1+\rho)^t$. The government is therefore faced with three mutually exclusive choices – financing the project domestically, or with foreign aid, or not undertaking the project at all – and it should choose that yielding the highest net present value at SPs.

The donor may rightly bemoan the laxer regulatory standards associated with the

domestic proposal: perhaps its only remedy is to sweeten the financial terms in order to impose more stringent standards. The donor also faces a further complication. This paper's strong advocacy of the use of SPs and the social discount rate applies to donors just as much as recipients, so donors would have to devote resources to estimating them. There can be no guarantee, however, that the donor's estimates will be the same as the recipient's, even if both use the same framework, for example, that laid out in Sections 3 and 4. The near certainty that they will differ in important respects offers donors and recipients the potentially useful opportunity to engage in a constructive dialogue to resolve them.

Closely related to foreign aid is the restructuring of foreign borrowing, broadly construed. Long-term interest rates in world capital markets have rested at historically low levels for quite some years now, which raises the question of whether poor countries should seize any opportunities that come their way of taking on new loans at fixed current rates, in the expectation that rates will resume more normal levels in the not-too-distant future. It can be argued, for example, that now is the time to use such loans, if available, to finance infrastructure projects. This reasoning is intuitively appealing, but it does not escape the problem of fungibility; for a relaxation of the government's budget constraint through inter-temporal arbitrage also makes feasible other public expenditures.

Let an economy have repayment obligations $\{D_t\}_{t=0}^{t=T}$, and suppose these can be converted into $\{D'_t\}_{t=0}^{t=T}$. Given the choice of numéraire and the sequence of social discount rates $\{\rho_t\}_{t=1}^{t=T}$, swapping the latter for the former is socially profitable if, and only if,

$$\sum_{t=0}^{t=T} D'_t \Big/ \left(\prod_{k=1}^{k=t} (1+\rho_k) \right) > \sum_{t=0}^{t=T} D_t \Big/ \left(\prod_{k=1}^{k=t} (1+\rho_k) \right).$$

The same calculus also applies to any proposed restructuring of the economy's foreign

debt.

The fungibility of public funds is reflected, satisfactorily or otherwise, in the procedure for deriving the sequence of social discount rates. If the swap or restructuring is sufficiently large, it may well affect the social discount rate, in which event the sequence $\{\rho_t\}_{t=1}^{t=T}$ will have to be re-estimated.

7.4 Tax reform

A public sector project is an intervention in the space of quantities, with possible effects on market prices. A tax reform is an intervention in the space of market prices, almost surely with effects on quantities. Since SPs are bound up with market prices, various questions arise if a tax reform – as opposed to a change in taxes in response to undertaking a project – is, at some point, on the table.

According to the border price rule, changes in taxes on traded goods have no effects on their shadow prices. If such changes result in changes in the prices of intermediate goods or consumer prices, however, then a reform of the taxation of traded goods will normally affect the SPs of non-traded goods and the shadow wage rate, the latter through the consumption conversion factor as well as, quite possibly, the social opportunity cost of labour. The same applies if there are changes in the taxes on non-traded goods. To give an example involving both categories, suppose an imported good is subject to a binding quota, which the government is subsequently persuaded to abolish in favour of a tariff less than the effective level under the quota. Under the latter, the SP of the good in question is equal to its marginal cost of production at SPs; after the reform, it will be the c.i.f. world price. The difference may be quite large, and there will be spillover effects on the SPs of non-traded goods and the shadow wage rate. When estimating SPs, therefore, it is important to bear in mind what tax

reforms might be implemented over the course of time.

In this connection, it must be emphasised that the argument in Sections 3 and 4 rests on the assumption that the only changes in taxes that occur are those needed to balance the government's (inter-temporal) budget under whatever *fixed* rule is thought to govern its behaviour in this sphere. Suppose there is a tax reform that will improve social welfare. If implemented, it will normally result in a change in SPs. That the reform and the evaluation of specific projects must be kept rigorously separate is established by the following example. Let a project be socially unprofitable at the pre-reform SPs, but the minister in question wants it anyway. A young economist in his ministry tells him that she has discovered a welfare-improving tax reform under which the project would pass muster at market prices. To sell the project to the Treasury and then the cabinet, it can be further stated – correctly if the reform is good enough and the project is not too large – that the *combined* package of project and reform is socially profitable. If, however, the project is socially unprofitable at the post-reform SPs, some damage will result and the clever discovery will have been abused.

In a further twist, it turns out that the hypothetical discovery depended on the use of SPs. In the standard analysis of a tax reform, producer prices are assumed to be fixed, so that adding commodity taxes or subsidies yields consumer prices. If, however, there are distortions in domestic production, as is commonly the case, then producer prices will not be equal to SPs, and whether a reform is welfare-improving can be established only by using SPs. The full argument is set out in Ahmad and Stern (1991), who painstakingly apply it to the reform of indirect taxes in Pakistan. The essential point is that any reform will affect social welfare not only partially, holding the endogenous policy variables fixed, but also by changing the vector of the private sector's net excess demands for goods, and the cost of any changes in the latter must be calculated at

SPs.

8 Implementation

Implementing – and maintaining – systematic appraisal of public expenditure programmes based on SPs requires not only technical and organisational capacity, but also the willingness to employ it, so far as possible without exception. Establishing the former seems the lighter task.

The generation of development economists who grew up with the central questions and general framework of Sections 3 and 4 is now largely in, or nearing, retirement. As attention and fashion changed in favour of other topics, moreover, so project appraisal began to disappear, at least two decades ago, not only from the pages of refereed journals, but also from graduate schools' syllabi. Most young development economists nowadays know all about so-called impact evaluation, but it is unlikely that many could fit it into a wider, rigorous scheme for assessing the social profitability of government interventions in the presence of distortions, and fewer still will be familiar with the above framework. The first step, therefore, is to persuade some of them and those in somewhat maturer years, both inside government and out, to invest in learning the trade as a sub-field of public economics. This investment should be fairly modest, for nothing especially exotic or technically abstruse is demanded. Rather, there is every need for the exercise of judgement, patience and stamina in assembling and analysing data and evidence, as in any serious undertaking in applied economics. One clear professional incentive is the opportunity to influence for the better how public revenues are spent. Yet this will cut little ice unless the work is so organised – and enjoys such standing – within government that its findings exercise a palpable influence over actual

decisions. If that can be accomplished, project appraisal will have fair chances of making a modest comeback in graduate programmes and the applied journals, thus securing technical competence in the longer run.

What, then, is the right way of organising of the work within government? A central plank of the whole argument is that different projects should be evaluated, as far as possible, using common criteria and a common yardstick. This calls for a strong degree of centralisation. Little and Mirrlees (1974: 100-104), for example, set out the case for what they term a 'Central Office of Project Evaluation', or COPE, whose task it would be to estimate a comprehensive set of SPs in the light of a general long-term investment plan for the economy. Individual projects would be assessed according to their estimated NPVs at SPs, whereby the latter would be revised from time to time as the economy develops. One should add that in order to avoid political capture, COPE should be independent of the spending ministries; for it is supposed to weed out bad projects without considering which ministries have proposed them. In particular, the SPs must apply uniformly.

Establishing such a centralised body as COPE does not, however, rule out having some in-house technical capacity in each spending ministry. There are, indeed, good reasons to complement COPE with such 'satellite' offices, each answerable to its own ministry, not to COPE. First, with technical knowledge available in house, prospectively bad projects can be weeded out at an earlier stage, thus freeing up resources the better to prepare the case for more promising ones, and so sharpening competition among ministries under common rules. Secondly, with the SPs before them and understanding how they have been estimated, the in-house analysts will have clearer signals as to which items of costs and benefits need careful attention if the project is to pass muster when scrutinised by COPE. Thirdly, possessing special expertise in

the ministry's remit, especially in estimating and thinking about benefits, they may be able to construct more reliable estimates of the associated social valuations than their colleagues in COPE. Fourthly, some rotation between COPE and the satellites in the ministries would give the central staff a better grasp of how flexible to be when evaluating very heterogeneous proposals using a set of SPs and a yardstick common to all.

COPE's influence over decisions will depend on where it is placed in the structure of government. One could argue that it should be under the wing of the Treasury, so as to strengthen the latter's control over the spending ministries. The drawback is that the spending ministries would then have every incentive to make common cause against COPE's assessments, doubtless citing all sorts of special circumstances and considerations; and in any decisions at cabinet level, the Treasury would very likely find itself isolated. A more promising arrangement would be to establish COPE as a cabinet office, with its own (junior) minister. This would leave the Treasury in the less exposed position of having to argue only for broad levels of spending, while setting the spending ministries at loggerheads over how aggregate spending is to be allocated among them, with COPE's assessments now in play as ammunition. Although the Treasury would be somewhat above the fray, the case for it having its own in-house 'satellite' office applies just as well as to the spending ministries. For this would enable the Treasury to keep an eye on any possible undue influence of those ministries on COPE. The central office and the ring of satellites would have incentives to engage in some mutual monitoring where the estimation and use of SPs is concerned, thus providing an answer, even if not a wholly satisfactory one, to the old question, '*quis custodiet ipsos custodes?*'.[11]

[11] 'Who will guard the guards themselves?'

A more radical alternative would be to place COPE outside the sphere of partisan politics by having it report directly to the whole elected body: the General Accountability Office of the U.S. is arguably the prime example of constitutional agencies of this type. The proposed 'satellites' answerable to their ministries would, however, remain essential. It is unclear whether the polities of many countries could tolerate such an independent agency, with its power to influence spending decisions, within government.

One may well ask whether small countries will possess the know-how and means to implement such a system of appraisal. As it is, their ministries of finance must deal with the task of analysing as well as collecting taxes, in view of which adding an agency to evaluate public projects does not appear to be unduly burdensome. Their ministries of finance are, moreover, frequently the recipients of technical assistance provided by donors, who could equally well do the same in the matter of project appraisal – if they possessed the expertise.

A still more compelling reason for international donors, multilateral and bilateral alike, to establish in-house capacities of their own is to improve their own decisions. The potential problems that arise are well exemplified by looking at the World Bank, wherein a concerted effort to introduce shadow pricing was made in the 1970s, with an accompanying manual of how to do it (Squire and van der Tak, 1975). Although the Bank has an audit agency, the Independent Evaluation Group (IEG), which reports not to the President, but rather the whole Board of Executive Directors, it would be absurd to charge it with the task of estimating a set of SPs for each and every borrowing country. That is a job for the Regional Vice-Presidencies, where the country economic work is done under the direction of their respective Chief Economists. With this prime task necessarily decentralised among the satellites, IEG would need its own capacity to monitor how the latter are performing and to enforce sufficient uniformity of method

and procedure across the regions. IEG's seal of approval, when warranted, would then carry some weight when the loan proposal is presented by the Region to the Bank's Board. By imposing a common standard and acting in keeping with its title, IEG would also temper any tendency towards loan-pushing, to which the Bank, like other institutions of its kind, is by no means immune.

Almost all bilateral donors would rightly consider the overhead costs of establishing an in-house IEG and regional satellites as not to be borne. What is individually insupportable, however, may become justifiable when spread more widely. The Nordic Group already co-operates closely in matters of development assistance, and the EU Commission in Brussels would surely not be averse to adding yet another agency.

If governments and donors were to take these proposals seriously, they would still need to set some priorities when executing them; for expertise and time will remain scarce, and estimating a whole set of SPs carefully is burdensome. Since so much public investment is long-term in nature, getting the social discount rate right is a compelling need. The same holds for the shadow wage rate, even in middle-income countries. Concentrating on these two parameters, with the robust support of the border price rule, is surely the correct way to proceed. How to assess and deal with risk, which is pervasive at the level of both the individual project and the economy as a whole, completes the agenda.

9 Concluding Discussion

No case for a revival of project appraisal based on shadow prices can be fully persuasive without a word of caution about its limitations. The concept of willingness-to-pay appears several times in the above argument and the general approach savours rather

of money-metric utility. It must be added at once that aversion to inequality easily finds a place in the whole scheme, one from which it exerts a potentially important influence on the whole structure of shadow prices. The usual, but not invariable, assumption is that the underlying social welfare function is Bergson-Samuelsonian in form; and once embraced, going a step further by assuming additive separability in individual utilities is nigh on irresistible in practical applications. Such an ethical foundation is soundly rejected by some economists. Libertarians and some on the Marxist left alike attack it as essentially authoritarian, others will have no truck with utilitarianism in any form. Yet even those who are prepared to make welfare judgements based on individual utilities can have serious misgivings about where this might lead when purchasing power over goods is the yardstick, even with a full allowance for aversion to inequalities in purchasing power.

Consider, for example, a multi-purpose dam. Not only would it produce electricity and irrigate a multitude of farms downstream, but also flood a vast area upstream, thus displacing many people, who would have to be resettled elsewhere or at least compensated monetarily. The dissolution of their communities would surely be a hard blow in itself, and if the natural environment is part and parcel of their way of life and culture, its destruction would be gravely damaging to their welfare, broadly conceived. India's Narmada scheme, which was embroiled in just such a controversy, is a well-known example. Some feared that when deprived of their livelihoods and communities, many of these tribal families would wind up in the slums of large cities, with much of the compensation squandered on drink. How is all this to weigh in the final reckoning?

Other social scientists, lawyers and activists of various stripes quite naturally seize on exactly this question. The reply must be modest: social cost-benefit analysis of the kind advocated in this paper does not – indeed cannot – provide a complete answer

when the question is so formulated. What it can claim to do is to provide a rigorous and transparent way of obtaining a cardinal measure of many of the chief consequences of a course of public action. It yields an answer that is admittedly wanting in some respects; but it places those who wish, for perfectly good reasons, to extend the scope of the evaluation to things that are not within its compass, to make their case in the light of that answer. That is to say, it seeks to set the terms within which the debate is conducted, when many of the arguments will be about what is commensurable. In so doing, it will doubtless draw the accusation that the economists are up to their old hegemonic tricks once more.

To close, a word of encouragement to those young economists who are undeterred by the prospect of such strife. Although the theoretical underpinnings are mature and largely secure, there are still plenty of hard problems to be solved, especially where estimating and valuing the benefits in certain spheres is concerned. The measurement and treatment of risks is a similarly promising topic for research. Finally, there is the task of combining all these elements in a unified structure in order to evaluate actual public projects. These are all worthy intellectual challenges.

References

Ahmad, E. and Stern, N.H. (1991), *The Theory and Practice of Tax Reform in Developing Countries*, Cambridge: Cambridge University Press.

Arrow, K.J. and Lind, R.C. (1970), 'Uncertainty and the Evaluation of Public Investment Decisions', *American Economic Review*, 60: 364-378.

Bell, C. (2012a), 'Estimating the Social Profitability of India's Rural Roads Program: A Bumpy Ride', Policy Research Working Paper No. 6168, World Bank, August.

Bell, C. (2012b), 'The Benefits of India's Rural Roads Program in the Spheres of Goods, Education and Health: Joint Estimation and Decomposition', Policy Research Working Paper No. 6169, World Bank, August.

Bell, C. (2016), 'Shadow Prices for a Small Open Economy under Uncertainty. Which Expected Values are Valid?', mimeo.

Bell, C. and Devarajan, S. (1983), 'Shadow Prices for Project Evaluation under Alternative Macroeconomic Specifications', *Quarterly Journal of Economics*, 97: 454-477.

Dasgupta, P., Marglin, S. and Sen, A.K. (1972). *Guidelines for Project Evaluation.* New York: United Nations Industrial Development Organization.

Devarajan, S., Squire, L. and Sethaput Suthiwart-Narueput, S. (1997), 'Beyond Rate of Return: Reorienting Project Appraisal', *The World Bank Research Observer*, 12 (1): 35-46.

Drèze, J. and Stern, N.H. (1987), 'The Theory of Cost-Benefit Analysis', in Auerbach, A.J. and Feldstein, M. (eds.), *Handbook of Public Economics*, vol. 2, Amsterdam: North-Holland.

Graham, D.A. (1981), 'Cost-Benefit Analysis under Uncertainty', *American Economic Review*, 71: 715-725.

Khandker, S.R., Bakht, Z., and Koolwal, G. B. (2009), 'The Poverty Impact of Rural Roads: Evidence from Bangladesh', *Economic Development and Cultural Change*, 57: 685-722.

Little, I.M.D. and Mirrlees, J.A. (1968), *Manual of Industrial Project Analysis*, volume

II, OECD Development Centre, Paris.

Little, I.M.D. and Mirrlees, J.A. (1974), *Project Appraisal and Planning for Developing Countries*, London: Heinemann.

Little, I.M.D. and Mirrlees, J.A. (1991), 'Project Appraisal and Planning Twenty Years On.' In: Proceedings of the World Bank Annual Conference on Development Economics 1990. Washington, D.C: World Bank.

Murty, M.N., and Goldar, B.N. (2007), 'Estimation of National Parameters for Project Appraisal in India', mimeo, Institute of Economic Growth, Delhi, April.

Rosenstein-Rodan, P. (1943), 'Problems of Industrialisation of Eastern and South-Eastern Europe', *Economic Journal*, 53: 202-211.

Squire, L., and van der Tak, H.G. (1975), *Economic Analysis of Projects*, Baltimore: Johns Hopkins University Press.

Squire, L. (1989), 'Project Evaluation in Theory and Practice', in Chenery, H.B. and Srinivasan, T.N. (eds.), *Handbook of Development Economics*, vol. 2, Amsterdam: North-Holland.

Tinbergen, J. (1966), 'Some Refinements of the Semi-Input-Output Method', *The Pakistan Development Review*, 6 (2): 243-247.

World Bank (2010), 'Project Appraisal Document (PAD)', Report No. 57081-IN, Washington DC, July.